Form Approved
OMB No. 0990-0379
Exp. Date 09/30/2020

Volume 2: Cybersecurity Best Practices for Medium and Large Healthcare Organizations

Table of Contents

Introduction

Cybersecurity Best Practices for Medium Healthcare Organizations

Medium healthcare organizations perform critical functions for the Healthcare and Public Health (HPH) Sector. These organizations include critical access hospitals in rural areas, management practice organizations that support physician practices, revenue cycle or billing organizations, mid-sized device manufacturers, and group practices. Generally speaking, a medium healthcare organization employs hundreds of personnel, maintains between hundreds and a few thousand Information Technology (IT) assets, and may be primary partners with and interfaces between small and large healthcare organizations. It's typical for a medium organization to have several critical systems that are interconnected to enable work activities in support of the organization's mission.

These organizations tend to have a fairly diverse inventory of assets that support multiple revenue streams. They also tend to have narrow profit margins, limited resources, and limited flexibility to implement robust cybersecurity practices. For example, unless connected to a larger system, it is rare for a medium organization to have a 24x7 security operations center (SOC). Yet, these organizations are essential to care for patients in more rural settings, for example, a rural critical access hospital or an operational support provider for health care organizations located in rural areas.

Medium organizations tend to focus on preventing cybersecurity events, implementing a more rigid security policy with few exceptions permitted. This is often due to a lack of resources required to support a more open and flexible cybersecurity model that larger organizations can afford. Medium organizations usually struggle to obtain cybersecurity funding that is separate and distinct from a standard IT budget. The top security professional in an organization of this size might often feel overwhelmed by compliance and cybersecurity duties, wear multiple hats, and be constrained around execution plans.

Medium organizations operate in complex legal and regulatory environments that include but not limited to:

- ONC Certified Electronic Health Information Technology interoperability regulations
- Medicare Access and CHIP Reauthorization Act of 2015 (MACRA)/Meaningful Use
- Multiple enforcement obligations under the Food and Drug Administration (FDA)
- The Joint Commission accreditation processes
- HIPAA/HITECH requirements
- The Payment Card Industry Data Security Standard (PCI-DSS)
- Substance Abuse and Mental Health Services Administration (SAMHSA) requirements
- The Gramm-Leach-Bliley Act for financial processing
- The Stark Law as it relates to being able to provide services to affiliated organizations
- The Family Educational Rights and Privacy Act (FERPA) for those institutions participating within Higher Education
- The Genetic Information Nondiscrimination Act (GINA)
- The new General Data Protection Regulation (GDPR) in the European Union

IT Assets Utilized by Medium Organizations

Medium organizations may have up to a few thousand IT assets, with a mix of dozens to a hundred information systems. All assets are capable of having cybersecurity vulnerabilities and are susceptible to cyber threats. The important factor in securing these assets is understanding their relationship within the organization's IT ecosystem, how they are leveraged and used by the workforce, and the data that is generated, stored, and processed within. Not all assets are equally important: some are mission critical and must be fully operational at all times while others might be able to be offline for days or weeks without harming the organization's mission. Some assets have large repositories of sensitive data that represent a significant risk, but are not as critical to the enterprise business drivers. In all cases, IT assets are used by the organization for a business reason and should be protected with proper cyber hygiene controls.

Examples of assets that can be found in medium organizations include but are not limited to:

- Static devices used by the workforce, such as shared workstations and clinical workstations used strictly for patient care with select mobile devices, such as laptops and smartphones. There may not be a large number of mobile devices due to their increased cost.
- Large numbers of "Internet of Things" (IoT) devices, for example, medical devices, smart televisions, printers and copiers, and security cameras.
- Data that includes sensitive health information stored and processed on devices, servers, applications, and the cloud. These data include names, medical record numbers, birth dates, social security numbers, diagnostic conditions, prescriptions, and potentially highly sensitive mental health, substance abuse, or sexually transmitted disease information. These sensitive data are generally referred to as Protected Health Information (PHI) or Personal Identifiable Information (PII).
- Assets related to the IT infrastructure, for example, firewalls, network switches and routers, Wi-Fi (both corporate and guest), servers supporting IT management systems, and file storage systems (cloud or on premise).
- Applications or Information Systems that support the business processes. This includes Human Resource (HR) or Enterprise Resource Planning (ERP) systems, pathology lab systems, blood bank systems, medical imaging systems, pharmacy systems, revenue cycle systems, supply chain/materials management systems, specialized oncology therapy systems, radiation oncology treatment systems, and data warehouses (e.g., clinical, financial).

Personal devices, often referred to as BYOD (Bring Your Own Device), generally are not permitted due to the organization's inability to dedicate the security controls required to secure their access.

Cybersecurity Best Practices

Medium organizations should consider the *Baseline Practices* discussed in each best practice presented in this volume. Nothing in this volume is intended to exclude medium organizations from adopting advanced best practices. Advanced best practices that are determined to be relevant should be considered for adoption by any organization.

Cybersecurity Best Practices for Large Organizations

Large healthcare organizations perform a range of different functions. These organizations may be integrated with other healthcare delivery organizations, academic medical centers, insurers that provide healthcare coverage, clearinghouses, pharmaceuticals, or medical device manufacturers. In most cases, a large organization employs thousands of employees, maintains tens of thousands to hundreds of thousands of IT assets, and has intricate and complex digital ecosystems. While smaller organizations operate with a few critical systems, large organizations can have hundreds if not thousands of interconnected systems with complex functionality.

Large organizations have the broadest mission scope and a large volume of assets to fulfill this mission. Even so, they often struggle to obtain funding to maintain security programs as well as control of their assets (shadow IT, rogue devices, unmanaged/unpatched devices), how sensitive data flows in and out of the organization's environment, and understanding system boundaries and segmentation that determine where one entity's responsibilities end and another's starts.

The missions of large organizations are diverse and varied. They include providing standard general practice care, providing specialty or subspecialty care for complicated medical cases, conducting innovative medical research, providing insurance coverage to large populations of patients, supporting the healthcare delivery ecosystem, and supplying and researching new therapeutic treatments (such as drugs or medical devices). Overall, these organizations are responsible for complicated healthcare issues and tend to operate in complex environments.

Large organizations operate in a legal and regulatory environment that is as complicated as their digital ecosystems. It includes but not limited to:

- ONC Certified Electronic Health Information Technology interoperability standards
- CMS Medicare Access and CHIP Reauthorization Act of 2015 (MACRA)/Meaningful Use
- Multiple obligations under the Food and Drug Administration (FDA)
- The Joint Commission accreditation processes
- HIPAA/HITECH requirements
- Minimum Acceptable Risk Standards (MARS) for payers
- State privacy and security rules
- Federal Information Security Modernization Act (FISMA) requirements for federal contracts and research grants through agencies such as the National Institutes of Health (NIH)
- The Payment Card Industry Data Security Standard (PCI-DSS)
- Requirements under the Substance Abuse and Mental Health Services Administration (SAMHSA)
- The Gramm-Leach-Bliley Act for financial processing
- The Stark Law as it relates to being able to provide services to affiliated organizations
- The Family Educational Rights and Privacy Act (FERPA) for those institutions that participate in Higher Education

- The Genetic Information Nondiscrimination Act (GINA)
- The new General Data Protection Regulation (GDPR) in the European Union.

IT Assets Utilized by Large Organizations

As mentioned, large organization operations are supported by a complicated ecosystem of IT assets. All assets are capable of having cybersecurity vulnerabilities and are susceptible to cyber threats. The important factor in securing these assets is understanding their relationship within the organization's IT ecosystem, how they are leveraged and used by the workforce, and the data that is generated, stored, and processed within. Not all assets are equally important: some are mission critical and must be fully operational at all times while others might be offline for days or weeks without harming the organization's mission. Some assets have large repositories of sensitive data that represent a significant risk, but are not as critical to the enterprise business drivers. In all cases, IT assets are used by the organization for a business reason and should be protected with proper cyber hygiene controls.

Examples of assets that can be found in large organizations include but are not limited to:

- Devices used by the workforce such as mobile phones, tablets, voice recorders and laptop computers for dictation (all with Internet connectivity).
- Personal devices that are often referred to as BYOD.
- Extremely large deployments of IoT assets that include medical devices, smart televisions, printers and copiers, security cameras, refrigeration sensors, blood bank monitoring systems, building management sensors and more.
- Data that includes sensitive health information stored and processed on devices, servers, applications and the cloud. These data include names, medical record numbers, birth dates, social security numbers, diagnostic conditions, prescriptions, and potentially highly sensitive mental health, substance abuse, or sexually transmitted disease information. The sensitive data captured here are generally referred to as PHI or PII.
- Assets related to the IT infrastructure, for example, firewalls, network switches and routers, Wi-Fi (corporate and guest), servers supporting IT management systems, and file storage systems (cloud or on premise).
- Applications or Information Systems that support the business processes. This includes ERPs, pathology lab systems, blood bank systems, medical imaging systems, pharmacy systems (retail and specialized), revenue cycle systems, supply chain/materials management systems, specialized oncology therapy systems, radiation oncology treatment systems, data warehouses (clinical, financial, research), vendor management systems, and so much more.

Document Guide – Cybersecurity Best Practices

This volume is intended to provide medium and large organizations with best practices to reduce the impact of these five currently prevailing threats:

- Email Phishing Attacks
- Ransomware Attacks
- Loss of Theft of Equipment or Data
- Internal, Accidental or Intentional Data Loss
- Attacks Against Medical Devices that can Affect Patient Safety

The intent of this section is to help you navigate the document this document. Each best practice is broken up into three core segments: Baseline Practices, Advanced Practices and the Key Risks that are Mitigated by the Practice. Additionally, each Best Practice contains a series of suggested metrics to measure the effectiveness of those practices.

"Baseline Practices" apply to both Medium and Large organizations. "Advanced Practices" apply to Large organizations, and any other organization that is interested in their adoption.

A summary of each of the Best Practices is outlined in the following ten tables.

Best Practice 1: Email Protection Systems	
Data that may be affected	Passwords and PII
Baseline Practices	A. Basic Email Protection Controls B. MFA for Remote Access C. Email Encryption D. Workforce Education
Advanced Practices	A. Advanced and Next Generation Tooling B. Digital Signatures C. Analytics Driven Education
Key Mitigated Risks	• Email Phishing Attacks • Ransomware Attacks • Accidental or Intentional Data Loss

Best Practice 2: Endpoint Protection Systems	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	A. Basic Endpoint Protection Controls
Advanced Practices	A. Automate the Provisioning of Endpoints B. Mobile Device Management C. Host Based Intrusion Detection/Prevention Systems D. Endpoint Detection Response E. Application Whitelisting F. Micro-segmentation/virtualization strategies
Key Mitigated Risks	• Ransomware Attacks • Theft or Loss of Equipment or Data

218

Best Practice 3: Identity and Access Management	
Data that may be affected	Passwords
Baseline Practices	A. Identity B. Provisioning, Transfers, and Deprovisioning Procedures C. Authentication D. Multi-Factor Authentication for Remote Access
Advanced Practices	A. Federated Identity Management B. Authorization C. Access Governance D. Single-Sign On (SSO)
Key Mitigated Risks	• Ransomware Attacks • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

219

Best Practice 4: Data Protection and Loss Prevention	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	A. Classification of Data B. Data Use Procedures C. Data Security D. Backup Strategies E. Data Loss Prevention (DLP)
Advanced Practices	A. Advanced Data Loss Prevention B. Mapping of Data Flows
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss

220

Best Practice 5: IT Asset Management	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	A. Inventory of Endpoints and Servers B. Procurement C. Secure Storage for Inactive Devices D. Decommissioning Assets
Advanced Practices	A. Asset Pre-Configuration B. Automated Discovery and Maintenance C. Integration with Network Access Control
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

221

Best Practice 6: Network Management	
Data that may be affected	-
Baseline Practices	A. Network Profiles and Firewalls B. Network Segmentation C. Intrusion Prevention Systems D. Web Proxy Protection E. Physical Security of Network Devices
Advanced Practices	A. Additional Network Segmentation B. Command and Control Monitoring of Perimeter C. Anomalous Network Monitoring and Analytics D. Network Based Sandboxing/Malware Execution E. Network Access Control (NAC)
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss • Medical Devices and Patient Safety

222

Best Practice 7: Vulnerability Management	
Data that may be affected	-
Baseline Practices	A. Host/Server Based Scanning B. Web Application Scanning C. System Placement and Data Classification D. Patch Management, Configuration Management, Change Management
Advanced Practices	A. Remediation Planning

Key Mitigated Risks	• Ransomware Attacks • Insider, Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices

223

Best Practice 8: Security Operations Center and Incident Response	
Data that may be affected	-
Baseline Practices	A. Security Operations Center B. Incident Response C. Information Sharing and ISACs/ISAOs
Advanced Practices	A. Advanced Security Operations Center B. Advanced Information Sharing C. Incident Response Orchestration D. Baseline Network Traffic E. User Behavior Analytics F. Deception Technologies
Key Mitigated Risks	• Phishing Attacks • Ransomware Attacks • Loss or Theft of Equipment • Insider, Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices

224

Best Practice 9: Medical Device Security	
Data that may be affected	ePHI
Baseline Practices	A. Medical Device Management B. Endpoint Protections C. Identity and Access Management D. Asset Management E. Network Management
Advanced Practices	A. Vulnerability Management B. Security Operations and Incident Response C. Procurement and Security Evaluations D. Contacting the FDA
Key Mitigated Risks	• Attacks Against Connected Medical Devices and Patient Safety

225

Best Practice 10: Cybersecurity Policies	
Data that may be affected	-
Baseline Practices	A. Policies

Advanced Practices	-
Key Mitigated Risks	• Email Phishing Attacks • Ransomware Attacks • Loss or Theft of Equipment or Data with Sensitive Information • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

226

227

Cybersecurity Best Practice #1: Email Protection Systems

Best Practice 1: Email Protection Systems	
Data that may be affected	Passwords and PII
Baseline Practices	A. Basic Email Protection Controls B. MFA for Remote Access C. Email Encryption D. Workforce Education
Advanced Practices	A. Advanced and Next Generation Tooling B. Digital Signatures C. Analytics Driven Education
Key Mitigated Risks	• Email Phishing Attacks • Ransomware Attacks • Accidental or Intentional Data Loss

According to the 2017 Verizon Data Breach Report, "weak or stolen passwords were responsible for 80% of the hacking related breaches" (Zaw 2017). This report further identifies the phishing attack (which is a hacking attack) as the most common first point of unauthorized entry into an organization. After monitoring 1,400 customers and 40 million simulated phishing campaigns, the PhishMe 2017 Enterprise Resiliency and Defense Report concluded that the average susceptibility of an organization to a phishing attack is 10.8% (How susceptible are you to enterprise phishing? 2017). Though other areas of significant threat exist, including in the web application space, the effectiveness of phishing attacks allows attackers to bypass most perimeter detections by "piggy backing" on the legitimate workforce user. Think about it, if an attacker obtains an employee's password and that employee has remote access to the organization's IT assets, the attacker has made significant progress toward penetrating the organization.

The two most common methods deployed by phishing attacks include a credential theft attack and a malware dropper attack. An organization's cybersecurity practices must address these two attack vectors – leveraging email to conduct a credential harvesting attack on the organization and delivery of malware through an email that can compromise the organization's endpoints. For both vectors, email is the vector leverage and the focus for additional security controls.

Baseline Practice

A. Basic Email Protection Controls (NIST Framework: PROTECT)

Basic protections that should be implemented in any email system are standard anti-spam and anti-virus filtering controls, both of which are implemented directly on the email platform. These controls assess inbound and outbound emails from known malicious senders or patterns of malicious content. Table 3 provides a list of suggested security implementations for email protection controls.

Control	Description
Real-time Blackhole List (RBL)	Community-based lists of IP addresses and host names of known or potential spam originators. Consider Spamhaus, Spamcop, DNSRBL or lists provided by your email technology.

Distributed Checksum Clearinghouse (DCC)	The DCC is a distributed database that contains a checksum of messages. Email messages are run through a checksum algorithm and then checked against the database. Depending upon the threshold of checksum matches, these can be determined to be spam or malicious messages.
Removal of Open Relays	Open Relays are Simple Mail Transfer Protocol (SMTP) servers that enable the relay of third party messages. SMTP is critical for the delivery of messages but must be configured to allow messages only from trusted sources. Failure to do this may permit a spammer or hacker to leverage the "trust" of your mail server to transmit malicious content.
Spam/Virus check on outbound messages	Spam/Virus checks of outbound emails can detect malicious content, indicating a compromised account and potential security incident in the organization. Email spam/virus rules should be reviewed as part of *Best Practice #8: Security Operations Center and Incident Response*.
Antivirus check	All email content should be scanned against an antivirus (AV) engine with up-to-date signatures. If possible, this control should unpack compressed files (such as zip files) to check for embedded malware.
Restrict the "Send As" permission for Distribution Lists	Limit distribution lists to essential members. Distribution lists can be popular mechanisms for a compromised account to disseminate malicious content and should not be accessible to large numbers of users.
Implement Sender Policy Framework (SPF) Records	A Domain Name System (DNS) record identifies which mail servers are permitted to send email on behalf of your domain. This enables the receiving mail server to verify the authenticity of the sending mail server.
Implement Domain Key Identified Mail (DKIM)	A method of email authentication that leverages cryptography to ensure email messages are sent from authorized email servers. A public key is stored within the organization's DNS as a text record (TXT). All messages sent from that domain are digitally signed with a DKIM signature that can be validated through the DNS public key TXT record.
Implement Domain-based Message Authentication Reporting and Conformance (DMARC)	An authentication technology that leverages both SPF and DKIM to validate that an email's "From:" address (i.e., the sender) is valid. DMARC enables the receiving mail system to check SPF and DKIM records which ensures conformance to the sending host as well as the From: address. It instills trust that the sending party's email address has not been spoofed, a common attack type to trick users into opening malicious emails.

Table 3. Email protection controls

In most cases, email protection controls do not operate alone. When combined to evaluate an organization's emails, they contribute information that provides a more complete assessment of each

message. In modern systems, this is accomplished by scoring email content on each pass through the protection controls.

It is highly recommended to leverage this scoring technique and to set at least three thresholds: OK for Delivery, Quarantine, or Block/Drop. Each email should be scored to determine which of the three thresholds are met. Based on that threshold, automated actions should be executed. Emails cleared for delivery pass through automatically for additional processing. Block/Drop emails are discarded and never seen by the user. Quarantine actions allow the user to evaluate the message in a secured environment, not the user's regular email box, for final verification. In most cases, these quarantine messages are delivered on a daily basis to the user in a single email digest for verification.

Adding X-Headers to the delivery of email messages is a good way to flag potential spam or malicious email before sending it to the user. There are two common methods for deploying this:

Spam X-Header: If a message is scored to a threshold where it cannot be definitely classified as spam/malicious, it can be tagged with an X-Header. The Subject or top of the Body of the message is modified to include a [POSSIBLE SPAM] tab. This advises the user to verify the legitimacy of the message before opening it.

External Sender X-Header: Another common practice is to add an [EXTERNAL] tag or message to inbound messages from external senders. This tag can be configured to be highly visible, such as "**WARNING: Stop. Think. Read. This is an external email.**" This method is effective at catching messages that might be spoofed or pretend to come from within the organization. It also informs the email recipient to be cautious when clicking links or opening attachments from these sources.

DMARC: If you leverage DMARC, you might consider exempting the External Sender X-header tag for messages that pass the DMARC authentication. Generally speaking, this will help the email users understand the trust environment setup and when it is necessary to be extra vigilant.

In addition to tagging messages that fail DMARC authentication, messages can be tagged, or digitally signed, when they originate from approved hosting or cloud based services with a legitimate need to spoof an internal address. This is common for communications platforms, such as marketing systems, emergency management communications systems or alert management systems.

B. Multifactor Authentication (MFA) for Remote Email Access (NIST Framework: PROTECT)

It is a common and expected practice for sensitive information to be shared and submitted through email systems. Email is the primary mechanism used by most organizations to communicate electronically. It is also a common practice to access email remotely as the workforce has become increasingly mobile.

Given the prevalence of credential harvesting attacks, if remote email systems are made available, the only control that prohibits a malicious user from accessing sensitive information within transmitted emails is a password. The susceptibility of organizations to phishing attacks makes this a critical exposure.

As discussed in *Cybersecurity Best Practice #3: Identity and Access Management*, Two Factor Authentication, or Multi Factor Authentication, is the process of verifying a user's identity using more than one credential, such as a password. The most common method is to leverage a soft token. The soft token is a second credential that can be delivered through a mobile phone or tablet, two devices that most of have close at hand. The soft token could be the delivery of a short message service (SMS)

text message with a code or an application installed on the phone that provides the code and/or asks for independent verification after a successful password entry.

Implementation of Two Factor Authentication on your remote access email platform mitigates the risk of a compromised credential, such as a user password. With Two Factor Authentication, a hacker requires both the phone and the user's password, which significantly reduces the likelihood of a successful attack. This is one of the most effective controls to protect your organization's data.

C. Email Encryption (NIST Framework: PROTECT)

Email is the most common form of delivering content, including sensitive information, to other members in an organization. Though this might not be the preferred method for most organizations, one must assume users will leverage this common and easy-to-use communication channel.

Deploying encryption techniques within an email message is an important protection control for users to leverage. Multiple techniques exist, though the most common invoke a third party application to conduct encryption by tagging the outbound message in some form. This can be done by putting a trigger in the subject line (e.g., #encrypt, #confidential), or by invoking it through the email client itself. It all depends upon the technology solution deployed.

When organizations have established partnerships with third parties, fully encrypted transparent email delivery can be provisioned between the two entities' email systems. In this model, leveraging transport layer security (TLS) enables both systems to be configured to require TLS encryption when sending or receiving messages from one another. This ensures the messages are delivered over the Internet in a manner that cannot be intercepted.

Whichever encryption technique is implemented, the organization's workforce must be trained to leverage the technique when transmitting sensitive information. This best practice may be integrated into the Data Protection best practices discussed in *Cybersecurity Best Practice #4: Data Protection and Loss Prevention*. Messages that are not encrypted by the user as required can be automatically encrypted or simply blocked.

D. Workforce Education (NIST Framework: PROTECT)

A study released in 2017 determined the average measured susceptibility of an organization for phishing attacks is 10.8%. Maintaining a workforce to be vigilant and aware of cyberattacks is incredibility important. Organizations should implement security awareness programs that provide context around email-based attacks. The challenge presented to security departments is how to deliver a distilled message or spot a technical attack when the workforce's knowledge level does not match the hacker's level of sophistication. For example, it's easy to make a phishing email appear to originate from the company itself, incorporating logos, department names, and management names. It's difficult to train your entire workforce to detect that fake message.

When implementing information security and cybersecurity programs, a few key principles should be considered. Leverage the following techniques from the 2015 HBR Article from Keith Ferrazzi, (Ferrazzi 2015):

- *Ignite each managers' passion to coach their employees* – Engage and train your management team. Leverage them to communicate security practices and information to staff in all areas of the organization.

- *Deal with the short-shelf life of learning and development needs* – Security information changes continuously. Implement continuous and ongoing campaigns to maintain awareness of current trends, issues and events.
- *Teach employees to own their career development* - Customize cybersecurity training to the needs of employees in different positions or units in the organization. Develop training that is clearly relevant to the user's job.
- *Provide flexible learning options* – Provide options, including on-demand and mobile training solutions that allow the workforce to schedule and complete training independently.
- *Serve the learning needs of virtual teams* – Recognize that many employees work remotely and virtually. Training solutions should fit within the work environment of virtual employees.
- *Build trust in organizational leadership* – Leaders must be open and transparent and lead by example. Managers must demonstrate to the workforce that they are fully engaged in security strategy and committed to successful execution of security controls and techniques.
- *Match different learning options to different learning styles* – Effective training accommodates the different learning styles and requirements of employees who function in diverse work environments within a single organization. Consider multiple options to conduct a single training course to maximize training effectiveness and efficiency.

Organizations should implement a multifaceted training campaign that engages users to catch phishing through multiple channels. Points to include in your training campaign are:

Sender Verification: Users should look very carefully at the sender of the email message. It is common to spoof the organization's name by changing a simple character, for example, "google.c0m" rather than "google.com." Be on the lookout for emails where the organization's name is given with a separate email domain, such as "ACME.google.com" rather than "acme.com."

Follow the Links: Every link in an email message is suspect. Organizations should limit the use of links in corporate messages to those that are required. Users should hover the cursor over each link to check the corresponding URL and determine if it is credible. Specifically, URLs that are mismatched – the name of the link in the email does not match the corresponding URL – are highly suspect.

Beware of Attachments: Though it can be difficult to determine if an attachment is malicious based on the content of an email message, often there are clues. Be wary of messages that require immediate action, for example, such as "you must read this right away." Be cautious when receiving attachments from senders you do not regularly correspond with. It's important to detect malicious attachments which may contain malware or exploit scripts that permanently compromise your computer.

Suspect Content: In most cases, hackers entice you to follow a link or open an attachment. They will use messages to play with your curiosity and emotions. These messages vary widely from urgent messages such as "your account will be deactivated unless you re-register" to scary messages such as "the IRS is suing you and you must fill out the attached form." Hackers also prey on hopes and desires. Examples of these messages include "You have won a $100 Amazon Gift card!" to the old fashion Nigerian Prince messages which everyone knows about.

As you establish your awareness campaigns, keep this simple goal in mind: you want your workforce to be human sensors detecting malicious activity and reporting these incidents to your cybersecurity department. As they say in the New York subway systems, "If you see something, say something." The earlier that security personnel become aware of a phishing attack, the faster they can execute *Cybersecurity Best Practice #8: Security Operations Center and Incident Response*.

The following are a list of recommended channels to leverage for these awareness campaigns:

Ongoing and Targeted Training – Although these not the most effective means of awareness, phishing content should be added to your organization's ongoing privacy and security training.

Monthly Phishing Campaigns – The most effective means of training your workforce to detect a phishing attack is to conduct simulated phishing campaigns. Your authorized security personnel or third party provider crafts and sends phishing emails to your employees. These emails are embedded with tracking components (like link clicks). Employees who detect the email as a phishing attack as well as those who don't detect the attack and open the email or emailed links are identified with the appropriate training and feedback provided as soon as possible after the event. These methods provide a cause-and-effect training opportunity and are incredibly effective. Consider conducting phishing simulations on at least a monthly basis for the entire workforce. Specialized simulations can be developed for the higher risk areas within your organizations. These could be based on the type of department (such as finance and human resources) or on data received that identifies your highest risk users.

Departmental Meetings – Leverage departmental meetings to disseminate information on information security and cybersecurity events and trends. Brief presentations or informal conversations provide face-to-face context and build relationships between security personnel and the organization's workforce. These relationships encourage a continuous dialogue that elevates the visibility of cybersecurity across the organization.

Email Campaigns – Leverage the email platform to deliver a pointed message or alert about specific attacks. Leverage Secure Multipurpose Internet Mail Extensions (S/MIME) or other digital certificates as evidence that these messages are authentic. Remember that the attackers will attempt to do the same thing!

Newsletters – Work independently or with your marketing departments, develop and distribute your own cybersecurity newsletter. Write articles that explain how to catch a phishing attack – better yet, provide an example of an actual phishing attack, highlighting the warning signs that might have prevented the attack.

Advanced Practices

A. Advanced and Next Generation Tooling

Many sophisticated solutions have been developed to help combat the phishing and malware problem. These solutions can be referred to as Advanced Threat Protection services. They use threat analytics and real-time response capabilities to provide protection against phishing attacks and malware.

The following list describes some of these tools:

URL Click Protection via Analytics – In a modern phishing attack, the hacker will create a web page on the Internet for the purpose of harvesting credentials or dropping malware. The hacker will conduct an email campaign, sending emails with a link to a web page that does not have malicious content. The organization's traditional spam and anti-virus protections clear the email and it is delivered to the user. As soon as the emails are sent and delivered, the hacker changes the link's web page to the newly

created malicious web page. This allows the hacker to bypass many traditional email protections and leaves the organization to rely on the user's vigilance and awareness.

Protection technologies that rely on analytics leverage the ability to re-write links that are embedded in an email message. The re-written URL instead points to a secure portal that applies analytics to determine the maliciousness of the request at the time of the click. The message is protected no matter where or when it is linked. The technologies tend to leverage the cloud and numerous sensors throughout the install base to check these sites in real-time. They can also push down blocks of these discovered malicious sites ahead of time to inoculate the organization.

Attachment Sandboxing – Another common attack technique is to leverage attachments with embedded malware, malicious scripts, or other local execution capabilities that compromise vulnerabilities on the endpoint where the attachment is launched. These attachments bypass traditional signature-based blocks by leveraging multiple obfuscation techniques that alter the attachment content to provide a completely different signature.

Sandboxing technologies leverage virtual environments to open these attachments proactively and determine what behaviors occur after the attachment is opened. The protection system determines whether a file is malicious based on these behaviors, such as system calls, registry entry creation, file downloading and a slew of other checks.

Automatic Response – Another useful technique is to leverage mechanisms that automatically rescind or remove email messages that are categorized as malicious after delivery to a user's mailbox. Leveraging analytics described earlier in this section, cybersecurity response teams will remove these messages from the user's mailbox. This manual process requires assessing characteristics of the malicious email message, searching the mailbox environments in the organization and deleting messages that match the assessed characteristics. This time consuming process is difficult to run in a 24x7 operation and can be dangerous.

These technologies are capable of identifying the signature of a delivered email. When advanced threat tools determine that a previously clean message has become malicious, it can automatically delete that email message from all user mailboxes in the organization. This reduces the labor involved in manual processes and provides the consistency of automation.

B. Digital Signatures

Digital signatures allow a sender to leverage public/public key cryptography to cryptographically sign an email message. This does not encrypt the message itself. It validates that a received message is from a verified sender and hasn't been altered in transit.

As long as trusted root certificates are used to create the S/MIME certificate used in digital signatures, most modern email clients will check and provide verification automatically through an icon on the message itself. This is useful when training your workforce to determine the validity of an email.

A word of caution: many email protection technologies change the content of an email messages (e.g., tagging subject lines, re-writing URLs). Digital signature technology that maintains the integrity of an email will fail when these other protection techniques are deployed. Currently, there is no method to resolve this problem.

C. Analytics Driven Education

Cybersecurity departments leverage data and analytics from regular email protection platforms and advanced threat protection systems to identify the most frequently targeted users in your organization.

These users might not be the ones you think are highly susceptible, such as the CEO or Finance workforce. With the systems discussed in this section, SOCs can identify these targets and implement increased protections (e.g., lower thresholds for spam/malware checking, delayed processing time for attachments) as well as provide on-the-spot and targeted education. Informing these individuals of their high-risk profile instills a heightened sense of awareness and increased vigilance.

Threats Mitigated

1. Email Phishing Attacks
2. Ransomware Attacks
3. Internal, Accidental or Intentional Data Loss

Suggested Metrics

- Number of malicious phishing attacks prevented on a weekly basis. The goal is to ensure that systems are working. A reduction in attacks prevented indicates system misconfiguration.
- Number of malicious URLs / Attachments delivered via email discovered and prevented on a weekly basis. The goal is to measure the effectiveness of advanced tools, like click protection or attachment protection.
- Number of account resets on a weekly basis. This is based on users who accessed a malicious website. It assumes that a registered click indicates compromised credentials, so be sure to change the credential before it can be further compromised). The goal is to leverage education to keep this number as low as possible.
- Number of malicious websites visited on a weekly basis. The goal is to establish a baseline understanding, then strive for improved awareness through education activities that train employees to avoid malicious websites.
- Percentage of users in the organization who are susceptible to phishing attacks based on results of internal phishing campaigns. This provides a benchmark to measure improvements to the workforce's level of awareness. The goal is to reduce this percentage as much as possible, realizing that it's nearly impossible to stop all users from opening phishing emails. Secondary Goal: Potentially correlate to percentage of susceptible users to the number of malicious websites visited or number of malicious URLs that are opened.
- List of the top 10 targeted users each week with corresponding activity. For example, how many phishing emails are targeted to the top three users compared to the rest of the workforce? What positions do these users hold in the organization? Is there are correlation between the user, the user's position and the number of phishing emails received? What inferences can be made? What conclusions can be reached? The goal is to conduct targeted awareness training to these individuals, advising them that they are targeted more often than other users and increasing their vigilance as well as their ability to detect and report phishing attacks.
- Average Time to Detect and Time to Respond statistics for phishing attacks on a weekly basis. Time to Detect measures how long the phishing attack was in progress before the cybersecurity department was aware of it. Response times measure of how quickly the

cybersecurity department neutralizes the messages to end the attacks. The goal is that both of these metrics should be as low as possible – establish a baseline to understand current state and set goals to improve performance.

References

- (Configure your spam filter policies 2017)
- (Using RBL and DCC for spam protection 2007)
- (Use DMARC to validate email in Office 365 2017)
- (Ferrazzi 2015)

Cybersecurity Best Practice #2: Endpoint Protection Systems

Best Practice 2: Endpoint Protection Systems	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	A. Basic Endpoint Protection Controls
Advanced Practices	A. Automate the Provisioning of Endpoints B. Mobile Device Management C. Host Based Intrusion Detection/Prevention Systems D. Endpoint Detection Response E. Application Whitelisting F. Micro-segmentation/virtualization strategies
Key Mitigated Risks	• Ransomware Attacks • Theft or Loss of Equipment or Data

Endpoints are the assets used by the workforce to interface with an organization's digital ecosystem. Generally speaking, endpoints include desktops, laptops, workstations, and mobile devices. Current cyber attacks target endpoints as frequently as networks. Implementing baseline security measures on these assets provides a critical layer of threat management. As the modern workforce becomes increasingly mobile, it's essential for these assets to interface and function securely.

The endpoints that form the majority of our computing environment are no longer static devices that exist in the healthcare organization's main network. Organizations commonly leverage virtual teams, mobility and other remote access methods to complete work. In some cases, endpoints rarely make it to the corporate network. It's important to build cybersecurity hygiene principles with these characteristics in mind.

Baseline Practice

A. Basic Endpoint Protection Controls (NIST Framework: PROTECT)

Table 4 describes basic endpoint controls with practices to implement and maintain them.

Control	Description	Implementation Specification
Antivirus (AV)	Technology that is capable of detecting known malicious malware through the use of signatures, heuristics and other techniques.	• Push AV packages out using endpoint management systems that interface with Windows and Apple operating systems (OS). • Develop metrics to monitor the status of AV engines, signature updates and health. • Dispatch Field Services/Desktop Support for malware that is detected and not automatically mitigated. • Leverage Network Access Control (NAC) to conduct a validation check prior to enabling network access.

Full Disk Encryption	Technology that is capable of encrypting an entire disk to make it unreadable for unauthorized individuals.	• Ensure encryption is enabled on new endpoints that are acquired by the organization. • Connect encryption management to endpoint management systems that interface with both Windows and Apple OS. • Develop metrics to monitor the status of encryption. • Dispatch Field Services/Desktop Support teams to resolve encryption errors. • Use anti-theft cable locks to lock down any device that cannot support environment. • Leverage NAC to conduct a validation check prior to enabling network access
Hardened Baseline Images	Configure the endpoint operating system in the most secure manner possible.	• Limit usage of local administrator accounts. Enable only those local administrative rights required by the user. Use a separate account dedicated to this purpose. • Enable local firewalls and limit inbound access to the endpoint to only those ports that are required. • Disable weak authentication hashes (e.g., LANMAN, NTML Version 1.0). • Prevent software from auto-running/starting, especially when using thumb drives. • Disable unnecessary services and programs. • Permit usage only of known hardware encrypted thumb drives for writing data. • Review and consider the implementation of Security Technical Implementation Guides (STIG).
Patching	The process to ensure that endpoint operating systems and third party applications are patched regularly.	• Establish an endpoint management system and distribute operating system patches during regular maintenance windows. • Automatically update and distribute patches to third party applications that are known to be vulnerable, such as Internet browsers, Adobe Flash, Acrobat Reader, and Java. • Develop metrics to monitor patch status. Review on a weekly basis. • Dispatch Field Services/Desktop Support for endpoints that fail to patch.

Local Administrative Rights	The provisioning of privileged access to users for the purpose of installing or updating application and operating system software.	• Limit local administrative rights deployed to endpoints. Leverage endpoint management systems to install new programs and patch systems. • For users that require administrative rights, deploy a local account with administrative privileges that is separate from the general user account. Never allow a general user account to operate with administrative privileges as this increases vulnerability to malware and client-side attacks.

Table 4. Basic Endpoint Controls Mitigate Risk at Endpoints

Organizations should reference *Cybersecurity Best Practice #5: IT Asset Management*, to determine if their endpoints meet IT Asset Management (ITAM) requirements. Examples include maintaining a proper inventory of endpoints, re-imaging endpoints as they are redeployed, and securely removing endpoints from circulation when decommissioned.

Advanced Practice

A. Automate the Provisioning of Endpoints (NIST Framework: PROTECT)

It is a challenge to manage thousands of endpoints consistently, especially when it is comprised of manually executed processes. Most organizations do not have the necessary resources to run such an operation without leveraging additional capabilities.

Value-Added Resellers (VARs) that sell endpoints through your supply chain can provide a mechanism that preconfigures the endpoint before delivering it to your enterprise. It requires the organization to build a "gold image" with a series of checklists and configuration procedures, which is provided to the VAR. This approach helps to ensure a consistent and resilient deployment.

In some cases, vendors provide the ability for an organization to provision devices centrally. For example, Apple provides this service for its devices. Leveraging the Device Enrollment Program (DEP) enables an organization to simplify enrollment and security management of the endpoint. This is accomplished by entering the serial number or order number of the new device within the DEP. That initiates a series of device configuration tasks that are specific to your organization's requirements. Further information can be found within Apple's Device Enrollment Program Guide. (DEP Guide 2015)

B. Mobile Device Management (NIST Framework: PROTECT)

Mobile devices, such as smartphones and tablets, present their own management challenges. Multiple security configuration options exists for these devices which should be configured consistently to comply with organizational security policies.

Mobile Device Management (MDM) technologies manage configurations of devices that are connected to these systems. In addition to configuration management, they may offer application management and containerization. All three options are important to consider, especially for organizations that leverage personal devices in business operations.

As most mobile devices travel on and off the organization's network, it's important to consider leveraging cloud-based MDM systems to enable consistent check-in. If cloud-based systems cannot be leveraged, then the on premise MDM systems must be accessible from the Internet through VPN connectivity or in the organization's demilitarized zone (DMZ).

Configuration Management – At minimum, ensure that passcodes are leveraged and encryption is enabled. Ensure the device locks automatically after a predefined period of time (perhaps 1 minute). Implement device wipe functions after a series of unsuccessful logins (consider 10 unsuccessful logins). Limit the time that an email can reside on the mobile device (consider 30 days maximum).

Consider leveraging "Always on VPN" to protect the device when connecting to unsecured wireless networks. Consider prohibiting the installation of unsigned applications.

Application Management – Malicious applications reside in app stores and appear to be legitimate, such as PDF Readers or Netflix. In reality, these applications contain malicious code that provides access to data elsewhere on the mobile device. MDM solutions use whitelisting or blacklisting techniques to limit the installation of these malicious applications. Both should be considered, especially for devices that run on the Android platform – which is an open platform that accepts a wide range of applications.

Containerization – Organizations that leverage BYOD policies should consider containerization technologies. These technologies allow business data on a mobile device to be segmented and processed separately from personal data. In these models, business applications exist only in the hardened container on the mobile devices. Examples of these applications include email, calendaring and data repositories. This allows the organization to wipe the container and clear business data from the device when the workforce member leaves or changes position in the organization. It limits the ability for personally downloaded malicious applications to access business data.

C. Host Based Intrusion Detection/Prevention Systems (HIDS/HIPS) (NIST Framework: DETECT)

HIDS and HIPS technologies leverage an intrusion protection method similar to that used by Network-based Intrusion Detection and Prevention systems. These technologies can be deployed on the endpoint and designed specifically to detect attack patterns launched against the endpoint. These attacks can originate at the endpoint's network, or through client-side attacks that occur when using email or browsing the web.

HIDS and HIPS technologies are usually managed through central management systems. They are deployed through central endpoint management systems used to manage endpoint software and patching. These types of systems should be configured to auto-update against their command server. The command servers should be configured to pull down fresh signatures of attack indicators regularly.

D. Endpoint Detection and Response (EDR) (NIST Framework: DETECT)

EDR technologies bridge the gap between execution and processing that occurs in an organization's fleet of endpoints. These agent-based technologies allow cybersecurity departments to query large fleets of endpoints for suspicious running processes, file actions, and other irregular activities.

EDR enables a large-scale response to malware outbreaks. If malware is installed in the organization's environment, cybersecurity professionals can 'reach in' and 'remove' the malware from thousands of devices using a single action. Finally, EDR technologies provide cybersecurity departments with forensic capabilities that supplement incident response (IR) processes.

E. Application Whitelisting (NIST Framework: PROTECT)

Application Whitelisting technologies permit only applications that are known and authorized to run, rather than identifying applications that should not be permitted to run. They are based on the assumption that it is impossible to identify and blacklist, or block, every malicious application.

Organizations should maintain a current inventory of all software resident on endpoints to facilitate complete and consistent maintenance and patching to protect against client-side attacks. (CIS Control 2: Inventory of Authorized and Unauthorized Software 2018).

Configuration of application whitelisting is complex and outside of the scope of this guide. Interested organizations should read NIST Special Publication 800-167: Guide to Application Whitelisting. (Guide to Application Whitelisting 2015)

F. Micro-segmentation/Virtualization Strategies (NIST Framework: PROTECT)

Technologies called "micro-virtualization" or "micro-segmentation" assume that the endpoint will function in a hostile environment. These technologies work by preventing malicious code from operating outside of its own operating environment. The concept is that every task executed on an endpoint (e.g., click on a URL, open a file) can be run in its own sandboxed environment and is prohibited from interoperating between multiple sandboxed environments.

Since most malware is installed by launching incremental processes after gaining an initial foothold, this strategy can be effective by eliminating that second launch. Additionally, once the malicious task has been completed, the microenvironment is torn down and reset. Further configuration advice for these technologies is specific to the technology deployed.

Threats Mitigated

1. Ransomware Attacks
2. Loss or Theft of Equipment or Data

Suggested Metrics

- Percentage of endpoints encrypted based on a full fleet of known assets and measured weekly. The first goal is to achieve a high percentage of encryption, somewhere around 99%. Achieving 100% encryption is nearly impossible as defects always exist. The percentage of endpoints encrypted will vary as new assets are discovered, which is why it should be measured weekly.
- Percentage of endpoints that meet all patch requirements each month. The first goal is to achieve a high percentage of success. A second goal is to ensure there are practices to patch endpoints for third party application vulnerabilities, as well as OS level applications, and have the ability to determine the effectiveness of those patches. Without the metric there might not be checks and balances in place to ensure satisfactory compliance with expectations)
- Percentage of endpoints with active threats each week. The goal is to ensure practices are in place to respond to AV alerts that aren't automatically quarantined/protected. This indicates there could be active malicious action on an endpoint. An endpoint with an active threat should be reimaged using general IT practices and managed using a ticketing system.
- Percentage of endpoints that run non-hardened images each month. For this metric, the goal is to check assets for compliance with the full management IT practices, identifying those assets that do not comply. To do this, a key or token is placed on the asset indicating that it is managed through a corporate image. Separate practices need

to be implemented for those assets that are not managed this way to ensure they are properly hardened.

- Percentage of local user accounts that are configured with administrative access each week. The goal would be to keep this number as low as possible, granting an exception to only those local user accounts that require such access.

References

- (Security Technical Implementation Guides n.d.)
- (CIS Control 3: Secure Configurations for Hardware and Software on Mobile Devices, Laptops, Workstations and Servers 2018)
- (CIS Benchmarks 2018)
- (Guide to Application Whitelisting 2015)

Cybersecurity Best Practice #3: Identity and Access Management

Best Practice 3: Identity and Access Management	
Data that may be affected	Passwords
Baseline Practices	A. Identity B. Provisioning, Transfers, and Deprovisioning Procedures C. Authentication D. Multi-Factor Authentication for Remote Access
Advanced Practices	A. Federated Identity Management B. Authorization C. Access Governance D. Single-Sign On (SSO)
Key Mitigated Risks	• Ransomware Attacks • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

Identity and Access Management (IAM) is a program that encompasses the processes, people, technologies and practices relating to granting, revoking and managing access for users. Given the complexities associated with healthcare environments, access management models are critical for limiting the security vulnerabilities that can expose organizations. A common phrase used to describe these programs is "*enabling the right individuals to access the right resources at the right time.*" Most access authentication methods rely on usernames and passwords which is proven to be a weak model based on the success of phishing and hacking attacks.

Establishing IAM controls requires a distinct and dedicated program to accommodate its high level of complexity and numerous points of integration. A toolkit for establishing an IAM program can be found on the Educause. (Toolkit for Developing and Identity and Access Management Program, 2013)

This section will focus on the critical elements of an IAM program required to manage threats relevant to the Healthcare and Public Health Sector.

Baseline Practice

A. Identity *(NIST Framework: PROTECT)*

As defined by NIST 800-63-3, "*…digital identity is the unique representation of a subject engaged in an online transaction*" (Digital Identity Guidelines). A common principle to follow is "One person, One identity, Multiple contexts." In healthcare, a person can have the context of a patient, payor, or even employee of the health system. For clinical staff, one person can have one identity, but that person's context and ability to practice specialties will depend on context by country, practice area, or hospitals where the person has a business or employee relationship.

Within the United States (U.S.), a person is are provided with a unique Social Security Number (SSN). Similarly, a person who joins an organization should be given a unique identifier. By the way, that unique identifier should not be used as a secret authenticator, the way a person's SSN is often used. The unique identifier is not the authenticator.

A person's identity should be established through onboarding systems of record. The most common of these systems is the ERP, or Human Resources (HR) System. When onboarding new employees in the organization, HR business processes identify and establish the new employee in the organization. Many processes are executed during onboarding, such as background checks, employment verification, and

preparation for payroll. They provide solid identity proofing activities that can be used to verify the employee's identity going forward. They are the trigger to generate a person's newly formed digital identity. These identity-proofing practices refer back to the need to understand a person's relationship and context within the organization(s).

It is imperative that IAM programs and functions align with HR business practices and business processes in general.

Identities maintain a series of attributes that describe common elements for the user. The series of attributes come from the system of record, whether that is an HR system, Contingent Workforce System, Medical Staffing Office, or other system in the organization's ecosystem. Examples of common elements include a person's name, location, telephone number, email address, job title/job code, and specialization/practice data.

Attributes from the system of record are transmitted to the IAM system, enriching the identity data and facilitating the flow of information to systems for sign on, access management, and other cyber security and business related functions. These attributes can be consumed by the organization and leveraged for future authorization components. In addition to common descriptive attributes, there may be other system-defined attributes. Further examples of this could be role or affiliations (general terms used to describe populations of individuals, such as clinicians, staff, staff nurse, visitor, student, etc.), attributes used to authorize eligibility for various systems (using techniques called Attribute Based Access Control), or other technical constructs.

Users defined under proper identity management processes will include more than your employees. These processes should account for and consider volunteers, locums, contractors, students, visiting scholars, visiting nurses, physician groups staff, billing vendors, visiting residents, special statuses (such as emeritus professors) and third party vendors that require access to provide services to your organization. Each of these identity types must have an approved channel to serve as the System of Record, where the identity proofing activities will occur.

All of this identity information can be stored in a single repository to be consumed for other purposes. IAM systems, by principle, are an aggregate of system of record data. IAM systems, by nature, should be a system of last resort when none exists (e.g., Contingent Workforce if HR/business does not have a solution). In the most basic context, the following principles should be followed:

- Enumerate all authorized sources of identity within the organization. This is often referred to as the Systems of Record. Examples of these sources include HR systems, vendor management systems, contingent workforce systems, medical staffing office/practice office, and student information systems.
- Ensure all users are provided with a unique identity and identifier. Smaller organizations without multiple constituents may consider using the Employee ID number from the System of Record. Larger organizations with many constituents should establish a unique identifier for each user and reconcile. Do not use a SSN as the unique identifier! An individual with multiple contexts should have one person record/one unique identifier that ties to all contexts.
- The integrity and uniqueness of digital identities should be maintained. Never reuse identities for different people. People come and go throughout the life of an organization. Their records should be maintained perpetually.

- Proper identity management enables the automation of functions such as system access and authentication. Enumerate and establish attributes that are critical to provide context to the identity and required for access and authentication controls. For complex and large organizations, this is an important principle to ensure consistent application of attributes and, subsequently, the ability to automate authentication and authorization needs, such as automated provisioning and de-provisioning.
- Store identity information in a database or directory capable of registering identity information and associated attributes. These are an aggregate of Systems of Record data. Consider specialized tools for organizations with multiple constituent types.
- Leverage a single namespace to establish user accounts within the organization. Tie these user accounts back to the identity so the individual can always be traced to their digital identity.

B. Provisioning, Transfers and De-Provisioning Procedures (NIST Framework: PROTECT)

Once digital identities and user accounts are established, users must be provisioned access to Information Systems prior to using them. It is important to ensure that provisioning processes follow organizational policies and principles, especially in the Healthcare environment.

Under HIPAA, a key principle relates to the Minimum Necessary rule. That is, organizations should take reasonable steps to limit uses, disclosures or requests of PHI to the minimum required to accomplish the intended purpose. These same principles reduce the attack surface of potentially compromised user accounts. By limiting access, you can limit the scope of a ransomware outbreak or data attack.

The following principles should be followed for provisioning:

- Identify common systems that all users will need to access and the most basic access rights required for each of those systems. This is generally considered to be "birthright entitlements."
- Define these entitlements in organization policies, procedures or standards. There should be documentation that describes the access rights that all users are entitled to receive.
- Establish procedures and workflows that ensure consistent provisioning to birthright entitlements. Consider leveraging specialized tools to automate this process for accuracy and reliability. A best practice is not to automate bad or unknown workflows.
- Establish procedures and workflow that enable provisioning required in addition to birthright entitlements, such as auxiliary or ancillary systems. Pay special attention to cloud-based systems. Consider leveraging Federated Access Management tools that automatically provision access in the cloud.
- Consider a two-part process that allows a user to request access with a second individual required to approve the request prior to granting access. A common approach is to designate an employee's supervisor as the approving party.

Leverage IT ticketing systems and build the workflow process into the ticketing system. This establishes consistency in the approval processes, automates the ability to gain approvals, and documents the access that was granted. It is important to ensure that access provisioning processes are auditable.

In addition to establishing a robust process to grant access to users, it is equally important to have a similar process to remove access at the right time. Failure to remove access promptly after an employee's relationship with the organization is terminated may result in unauthorized or malicious access to systems.

The following principles should be followed for de-provisioning:

- Establish procedures to terminate access to user accounts. Ensure these procedures are executed promptly at the time of termination. Consider leveraging tools that automate this process after being informed by the SoR. The SoR is usually the HR systems although other systems may trigger the process to terminate access.
- Ensure your termination process, whether manual or automated, includes session termination steps to prevent active sessions (e.g. email on phone) from remaining active after the employee leaves the organization.
- Establish an "urgent termination" process, which flows outside of the normal termination procedures. These should be used when a sensitive termination process is being executed, such as an involuntary termination.
- Ensure that termination procedures include critical business systems as well as ancillary or auxiliary systems. Pay special attention to cloud-based systems that are accessible outside of the organization's standard network. These assets will remain accessible to the user if the de-provisioning process isn't completed. Consider leveraging Federated Access Management tools to de-provision access to cloud-based systems automatically.
- Build automatic timeouts for non-use in critical systems. These timeouts can catch edge cases where de-provisioning procedures are not executed, ultimately reducing the exposure to unauthorized access.

Removal of access should occur when a user terminates its relationship with the organization and when users transfer to new functions in the organization. For example, if a patient care service (PCS) manager transfers to the nursing department, access granted when the user was a PCS manager should be removed prior to granting access required by the user as a nurse. This helps to eliminate the accumulation of unnecessary access rights by a user.

C. Authentication (NIST Framework: PROTECT)

User accounts must leverage authentication techniques to properly assert the user's identity in the digital ecosystem. The most common and, unfortunately, the weakest method for authentication relies on password credentials. That said, password-based authentication systems will be used for the foreseeable future and organizations should develop solid practices.

Centralize Authentication – Central authentication systems, such as Lightweight Directory Access Protocol (LDAP) directors or Active Directory, should be leveraged to the greatest extent possible. Tying authentication mechanisms back to these central systems enables enterprise management of credentials. The access rights of your user base can be managed from a single location. This is incredibly important when access needs to be de-provisioned in a timely and automated manner.

Passwords are the most common credential used to authenticate users. The strength and management of passwords are paramount. Password strength perspectives have been constructed to combat brute force or password guessing attacks. Assuming that you can limit the exposure to brute force and guessing attacks, NIST recommends the following techniques as part of SP 800-63 (SP 800-63B 2017):

- Limit the rate of speed at which authentication attempts can occur. Spacing out each password attempt by a second or two severely limits the ability of automated systems to brute force the password.
- Ensure the use of cryptographically strong hashing and salting for password storage.
- Use pass phrases in place of passwords. Require a minimum of 8 characters and permit up to 64 characters.
- Implement the use of dictionary-based password checking and compromised password black-lists. Prohibit users from establishing risky passwords.

Privileged Account Management – Centralized authentication should be leveraged for general user access and privileged administrative accounts. Privileged administrative accounts must be separate from the general user accounts. For example, an IT administrator should be provisioned a minimum of two accounts: one account for use completing day-to-day activities and a separate administrative account with access limited to those systems required by the IT administrator. This second step is critically important – the use of privileged accounts during the course of normal day-to-day business may expose these accounts to malware attacks giving an attacker elevated access to the organization's environment. This exposure should be limited as much as possible. Consider the following controls for managing your privileged accounts:

- Rotate passwords regularly.
- Escrow privileged systems credentials, making them unique for each system or device.
- Link privileged access to problem, change or service tickets in the organization's ticketing system.
- Require the use of a jump server when elevating privileges. Ensure full recording and auditing of the jump server.
- Require brokered access to a privileged account that registers the exact user using the privileged account and records all actions taken.
- Require multi-factor authentication for all privileged accounts used interactively.
- Conduct regular reviews of privileged access.
- Limit actions that privileged accounts can take through the use of access control lists. Check for the use of sensitive commands and alert accordingly.

For further details on how to securely configure privileged access, consider the following resources:

- Windows: (Implementing Least-Privilege Administrative Models 2017)
- Linux: (Controlling Root Access n.d.)

Local Application Authentication – There may be cases where applications do not support a centralized authentication model. Although there are fewer on premise systems that cannot bind to centralized authentication systems, these systems are still prevalent in the healthcare environment. As organizations migrate applications to the cloud, it's easy to accidentally instantiate a cloud-based service that lacks robust authentication and focuses on local user accounts.

Whenever a local authentication system is used, solid access control procedures must be maintained to manage user accounts. This requires designating a responsible IT owner who will manage and regularly

review these accounts. Failure to do this allows users access to systems for longer than necessary and is especially risky when an employee leaves the organization and continues to have access to these systems. Consider the following extra controls:

- Designate an IT owner for each legacy/cloud-based system.
- Establish a distribution list in your organization which includes your IT owners as members. Submit terminations out to these IT owners as they occur.
- Ensure IT owners comply with standard operating procedures for the onboarding, review and, most importantly, termination of users.
- Regularly audit the practice of these manual processes. Ensure compliance with regular account review and termination procedures.

Monitor Authentication Attempts - Regular user accounts and privileged user accounts should be monitored for security and compliance purposes. These details are discussed further in *Cybersecurity Best Practice #8: Security Operations Center and Incident Response.* Details on methods to monitor these types of authentication logs can be found on the SANS white paper: (Keys to the Kingdom: Monitoring Privileged User Actions for Security and Compliance 2010).

D. Multi-Factor Authentication (MFA) for Remote Access (NIST Framework: PROTECT)

Multi-Factor Authentication (MFA) systems require the use of several authentication methods to verify a user's identity. MFA systems are commonly described as leveraging at least two of the following: something you know, something you have, and something you are. In these models, at least two of those three categories must be correctly addressed before the user's identity will be verified for access.

The most common MFA techniques are the use of passwords and one-time codes that are delivered to the user out-of-band from the authentication technique. For example, customer authentication to access sensitive banking information. Most banks have MFA capabilities, which require the customer to enter a password (something you know) followed by a verification code that is texted to the customer's smart phone (something you have). This is the most basic functionality of these systems.

MFA should be implemented on remote access technologies in the organization to limit the exposure of password credentials that could be compromised through phishing or malware attacks. This is the single most effective tool at limiting an attacker's ability to compromise the organization's environment. Consider implementing MFA on the following types of technologies:

Virtual Private Networks (VPN) – These are systems that allow remote network access to your environment. VPNs should be configured to limit access of the user based on RBAC or ABAC rules and to enable MFA.

Virtual Desktop Environments – These are environments where virtual terminal sessions can be exposed to remote access that allows your employees to work remotely. Although highly useful for workforce flexibility, these systems can be compromised easily without MFA authentication.

Remote Email Access – If your organization is going to permit remote email access, MFA should be enabled to limit the risk of compromised credential access in the email system. It is common for healthcare environments to store PHI with these systems and this exposure could result in a breach of sensitive information.

Advanced Practice

A. Federated Identity Management (NIST Framework: PROTECT)

Federated Identity Management is the mechanism that enables identity information to be shared between organizations in a trusted manner so that identities can be leveraged from home institutions inside of a greater ecosystem. In healthcare organizations, it is commonplace for providers, payors, and other affiliates to work together in an integrated manner. In complex, large environments, multiple organizations operate in a joint model, with different HR practices inside each of the systems.

Rather than creating identities in each organization involved in this joint operation, the use of federated identity management tools and processes allows identity assertions of the home institutions to be used throughout the federation.

Consider the following example: a clinician is part of a practice group that is credentialed within a regional hospital. From the hospital's perspective, this clinician is not an employee but must be credentialed and cleared through the medical staff office with access to the electronic medical record. From the practice group's perspective, this clinician has been on-boarded through standard HR background checks and processes. If the practice group and the hospital were operating within a federation, the clinician's "home" identity could be established from the practice group and asserted to the hospital as part of the clearance processes. If the clinician's relationship with the practice group changes, this identity information would be revoked within the hospital based upon asserts from the federation. These processes would be completely automated.

The same model can be leveraged when working with third party vendors who have a large workforce that provides support or staff-augmentation capabilities. In this model, the third party is the "home institution" that requires access to resources in the organization. To monitor the activities of each of those workforce members would be a highly complicated and largely manual process that is likely to be ineffective. A federation can solve this problem.

In complex environments such as large integrated delivery networks, federations are almost a requirement to properly manage the temporal aspects of the identities within.

B. Authorization (NIST Framework: PROTECT)

After authentication has occurred, the mechanism to obtain specific access in an information resource, such as an application system, is referred to as authorization. Authorization processes check the level of access that has been granted to a particular credential, and assures that the credential can access only those areas that are pre-authorized. Consider the analogy of traveling at the airport. When you pass through the security lines, your identity is authenticated and then you are authorized to access the terminals based on a ticket for a particular flight. You are not permitted to access any other flight than the one authorized on your ticket.

The limitation of authorization is a critical component, required by the HIPAA Privacy rule under Minimum Necessary. In addition to HIPAA compliance, Minimum Necessary is a leading practice to limit malicious use of credentials. In most cases, when hackers break into systems, they are trying to access the "keys to the kingdom" – privileged access credentials that permit access to the most sensitive resources. Don't risk unauthorized access by granting more access than necessary to your users!

Consider the following techniques to limit authorization to only those components required by the user:

Role Based Access Control (RBAC) – Conduct a high-level role mining exercise to map out the role types that exist within your organization and the access they require. For example, identify access

requirements for clinicians, support staff, unit secretaries, switchboard operators, case managers, and others. The clinician may need access to the medical record (though not necessarily the entire medical record), and the support staff cannot need access at all. By defining the unique requirements for these two roles, you have started on the path towards differentiating access models.

It can be difficult to provision granular-based authorization models based on a user's role. In Healthcare, two individuals might have the same job title and role, yet complete different tasks within the organization. Relying solely on a person's role to grant access will limit the ability for the other authorized responsibilities.

Attribute Based Access Control (ABAC) – Attribute based authentication models considers the attributes associated to a user's identity, the attributes of the information system being accessed, and the context associated with the access request. In this model, a user may be granted an attribute through a standard workflow process to enable the user to be granted access to a specialized function in an information system, but only during business hours or only while on premise. When the user requests this access, ABAC systems check the specific context against these factors to determine if access should be granted. (Attribute Based Access Control 2018).

This highly effective model limits access based on user-specific rules established in the ABAC systems that define access parameters. As an example, a request is made to grant all nurses with access to all patients on a specific floor in a specific hospital to support flexible care requirements. You might be concerned that this access is excessive since the nurse might not be part of a care team for a particular patient. To minimize this impact, you can leverage ABAC, limiting the access to the time when the nurse is present at a specific hospital and only after the nurse has been authenticated using both a password and multifactor authentication. These particular access credentials cannot be used to grant remote access to the same patients or anywhere else within the healthcare system. In this model, a hacker with access to the password and the multi-factor authentication mechanism would be unable to access the patient using those credentials.

C. Access Governance (NIST Framework: IDENTIFY)

On-boarding processes and system access request processing are the most active times during the whole establishment of access process. Once access is established for a user, it can be a challenge to determine if that access continues to be required at a given time in the future. Consider the employee who has worked for an organization for a long time, serving in multiple capacities, placed on special projects and working throughout the organization. Over time, this employee might accumulate more access than was ever intended.

Conducting a manual review of each employee's access to each of an organization's critical application systems would be nearly impossible. Fortunately, specialized tools are available that automate these processes, providing this capability in a self-service capacity to an organization's leadership. These processes are generally referred to as Access Governance. Below are the relevant components:

Tooling - Specialized tools bind to identity management systems and connect to critical business systems to understand the access in place for all users in these systems. These tools need the ability to connect and parse through specific aspects of the applications in question, such as electronic medical records systems, revenue cycle systems, imaging systems, lab systems, and more.

Access rules – Within the tools, specialized rules can be defined based on roles or attributes of users in the organization. For example, it would be important in a financial system to define the difference in roles between accounts payable and accounts receivable. No employee should be capable of access

with both roles. Otherwise, a fraudulent purchase order could be generated and the invoice paid by the same person, resulting in a fraud loss to the organization. Understanding the characteristics and requirements of these critical roles enables you to create automated alerts that control user access.

In addition to the standard segregation of duties checks, some specialized tools compare access profiles of certain users in a role to determine if outliers exist. For example, these systems can look at the accesses that nurses generally have across multiple systems, set a baseline of access based on the normative manner of these accesses, and compare each user against that baseline. A specific nurse's profile that is determined to have excessive access can be reviewed for appropriate adjustments.

Access Review – Through workflows established with these advanced tools, supervisors within the organization can review the access that their employees currently have in critical environments. This can be done on a regular cadence established by policy. In the case where an employee has retained access that is no longer necessary, the manager can use self-service portals to identify these access violations and flag them for removal. In some systems, once the manager flags an access for removal, it will be automatically stripped.

At the end of an access review, the manager can certify and sign off that the review is accurate. This documentation is useful for audit practices and to demonstrate effective reviews.

D. Single-Sign On (NIST Framework: PROTECT)

Federated Single-Sign On (SSO) is an effective method to authenticate users against centralized credential repositories. SSO techniques abstract authentication principles away from the general Microsoft- or Linux-based methods into a generalized standard that can be implemented cross-platform. These standards conduct authentication processes securely and pass additional identity attribute information for authorization processes in the resources being accessed. Lastly, it has the benefit of requiring only one login during a particular time frame while an active SSO session is enabled. Get rid of those pesky password prompts!

Several federated SSO standards exist, including OpenID, Security Assertion Markup Language (SAML), OAuth and Active Directory Federation Services (ADFS). When implementing cloud-based systems, such as Software-as-a-Service (SaaS) systems, the use of SSO should be a security requirement.

Another SSO system model leverages a second authentication factor at a clinical workstation for easy access within a healthcare provider space. These systems can be configured to require a user to authenticate once per day/shift at a clinical workstation after accessing a clinical setting through use of a password and key card. Subsequent authentications are conducted by tapping the key card to provide secure, easy access to the clinical workstations. These systems provide multi-factor access to the clinical environment while easing the password authentication processes.

Threats Mitigated

1. Ransomware Attacks
2. Internal, Accidental or Intentional Data Loss
3. Attacks Against Connected Medical Devices that can Affect Patient Safety

Suggested Metrics

- Number of alerts generated for excessive access to common systems. For example, "allow any" permissions to core applications SharePoint, file systems, etc.

- Number of users with privileged access, trended over time. The primary goal is to establish a baseline of the normal number of privileged accounts and monitor variances from the baseline.
- Number of automated terminations, trended over time. The goal is to establish a baseline for normal terminations and monitor variances from that baseline. A decrease in the amount of terminations can indicate that the automated systems are not terminating access properly.
- Number of elevated privileged access requests, trended over time. The goal is to establish a baseline to determine how much privileged access is generated over a one-week period and monitor variances from that baseline.

References

- (Attribute Based Access Control 2018)
- (Controlling Root Access n.d.)
- (Implementing Least-Privilege Administrative Models 2017)
- (Keys to the Kingdom: Monitoring Privileged User Actions for Security and Compliance 2010)
- (Digital Identity Guidelines n.d.)
- (SP 800-63B 2017)

Cybersecurity Best Practice #4: Data Protection and Loss Prevention

Best Practice 4: Data Protection and Loss Prevention	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	A. Classification of Data B. Data Use Procedures C. Data Security D. Backup Strategies E. Data Loss Prevention (DLP)
Advanced Practices	A. Advanced Data Loss Prevention B. Mapping of Data Flows
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss

All organizations within the Healthcare and Public Health Sector access, process and transmit sensitive information, such as health information or personally identifiable information. The fundamental data used in operations is highly sensitive, representing a unique challenge to the Healthcare and Public Health Sector. The majority of the workforce in Healthcare must leverage these data to carry out their respective missions.

With that backdrop, Healthcare faces a growing challenge of understanding where data assets exist, how they are used and how they are transmitted. On any given day, PHI is being discussed, processed and transmitted between information systems. Protecting these data requires robust policies, processes and technologies.

As your organization starts shoring up its data protection and prevention controls, it's best to begin by understanding the type of data that exist in the organization, setting a classification schema for these data, then determining how the data are processed. Establish a set of policies and procedures for normal methods of data use and then build in 'guardrail' systems to guide your user base towards these business processes.

Baseline Practice

A. Classification of Data (NIST Framework: IDENTIFY)

There is a vast proliferation of data in Healthcare environments. They can range from the obvious records, including treatment information, billing information, and personally identifiable information (e.g., social security numbers, insurance numbers, driver's license numbers) to research information. It includes the non-obvious, but still important, information such as business strategies and development plans, business finances, employee records, and corporate board materials.

Before establishing policy describing how these varied data types should be used and disclosed, it's best to classify them into high-level categories that provide a consistent framework when developing policies and procedures. Table 5 provides an example of a classification schema with examples of the types of documents that comprise the classification.

Classification	Description	Examples
Highly Sensitive Data	Data that could easily be used for financial fraud, or	Social security number, credit card number, mental health information,

	could cause significant reputational damage.	substance abuse information, sexually transmitted infections/disease.
Sensitive Data	Regulated data, or data that could cause embarrassment to patients or organizations.	Health information, clinical research data, insurance information, human/employee data, board materials.
Internal Data	Data that is not considered sensitive, but should not be exposed publicly.	Policies and procedures, contracts, business plans, corporate strategy and business development plans, internal business communications.
Public Data	All data that has been sanitized and approved for distribution to the public with no restrictions on use.	Materials published on websites, presentations, and research publications.

Table 5. Example of a Data Classification Schema

B. Data Use Procedures (NIST Framework: IDENTIFY)

After data have been classified, procedures can be written that describe how to use these data in the organization based on their classification. These types of procedures include the process of setting usage expectations and labeling the information properly.

Usage and Disclosure – Based on the classification type, data use should be limited appropriately and disclosed in specific methods. Consider the procedures in Table 6:

Classification	Use	Disclosure
Highly Sensitive	1. Must be restricted to only those individuals who have a need to know. 2. Must use extreme caution when handling data.	Only share information internally and only when expressly permitted and when directed by the data owner.
Sensitive	3. Must be restricted to only those individuals who have a need to know.	Only share information internally and only when expressly permitted.
Internal Use	4. Data can be generally used, but care should be considered in its consumption.	Only shared information internally within the organization.
Public	5. No restrictions.	Shared freely shared with no restrictions.

Table 6. Suggested Procedures for Data Disclosure

Be careful when sending information through email. Ensure that sending PHI via email is consistent with ONC guidance. Do not send unencrypted PHI through regular email or text, unless patients have expressly requested receiving their PHI via this means.

Labeling – It is important to label information properly to facilitate implementation of restrictions related to its usage and disclosure. This helps keep the data secure through a few methods. First of all, users will understand how to handle information that is properly labeled. Second, specialized security tools, such as Data Loss Prevention systems, can be configured to discover and control information when it is properly labeled.

At minimum, the labeling process should be done so that it is readily apparent when viewing the information or documenting its classification. Leverage techniques such as placing the classification in the footer of the document. Collaborate with your marketing and communication departments to create document templates based upon these sensitivity types. Leveraging organization-wide document templates enables specialized tokens or signatures to be embedded in the documents and tracked by Data Loss Prevention systems.

C. Data Security (NIST Framework: PROTECT)

Once policies and procedures have been defined, you can establish additional data security methods. Consider the security methods described in Table 7.

Security Method	Description	Considerations
Encryption Data At Rest	Ensure data are encrypted when resident on file systems.	• When using Cloud, enable native encryption capabilities to prevent exposures if cloud provider is hacked. • Ensure full disk encryption is enabled on all workstations and laptops.
Encryption In Transit	Ensure secure transport methods are used for both internal and external movement.	• Ensure websites accessed that contain sensitive data use encrypted transport methods, such as hypertext transfer protocol secure (HTTPS). • Enable internal encryption methods when moving data in the organization. • Never send unencrypted sensitive data outside of the organization.
Data Retention and Destruction	Ensure retention policies are set. Contractually bind third parties to destroy data when terminating contracts.	• Use standard destruction forms and require vendors to attest data has been destroyed pursuant to those forms. • Set retention policies and quote limits on email systems to reduce the amount of data

		that can be exposed. Ensure that legal retention requirements are met. • Establish a purge strategy that includes purge mechanisms.
Scrub production data from Test and Development environments	Ensure that identifiable information is removed when replicating production environments for testing.	• Leverage specialized tools to de-identify data elements within large systems (such as EMRs). • Regularly audit data elements within test and production environments to ensure they are clean.
Mask Sensitive Data within Applications	Restrict users from accessing highly sensitive information, such as SSNs, by masking it unless authorized.	• Permit SSN access only to those members who require it (e.g., registration desks, admitting desks, payor processing).
Limit ability to print / save / export data based on function	Restrict the workforce's ability to export data out of systems that contain sensitive data without proper authorization.	• Encourage users to work within applications. Minimize data exporting by providing the required capabilities to manipulate data within the application. • Implement restrictions on data exports, especially in reporting or database systems that can query and return large datasets. • Potentially remove the ability to print and copy/paste EMR applications or web mail accessed from home.

Table 7. Security Methods to Protect Data

D. Backup Strategies (NIST Framework: PROTECT)

A robust backup strategy for enterprise assets is critically important to daily IT operations. It is equally important to have such a backup strategy in the event of cybersecurity incidents. There will be events that cause an asset, or multiple assets, to be thoroughly compromised. During these events, these routine backups can be the only way to ensure proper execution of the Recovery phase of your IR process. A full decommission and restore to a time before the compromise occurred is the best method to neutralize the compromise.

At minimum, each mission critical asset in your environment should have a backup plan. Backups can be executed using a variety of methods, the most common being disk-to-tape, disk-to-disk, or disk-to-cloud

backups. The integrity of these backups is paramount: these copies are your last line of defense and you want to make sure they are complete and accurate when you need them.

Disk-to-Tape – This method makes backups by accessing designated systems and files and writing all the content to a tape drive, or a tape library. Specialized software, hardware and inventory controls are required. To conduct backups efficiently, you will need the appropriate tape robots and a tape library that is based on the number and size of systems being backed up. These backups can be very large. You'll want to configure a Write Once and Read Many (WORM) option with the tapes. It is of utmost criticality that encryption is enabled in the writing to these tapes. If a tape is lost or stolen, you do not want to have a breach of sensitive data.

There are great advantages to maintaining offline backups. You can rely on these copies to be available when you need them, and attacks against the backup medium itself are as they are offline.

Disk-to-Disk - This method involves taking backup copies from a disk and replicating them to a separate disk storage array that is dedicated to maintaining backup copies. This option is generally lower cost than the other tape strategies discussed in this section. Disk-to-disk backups generally executes more quickly than disk-to-tape. It is important to execute encryption strategies on the backup file, in case the file is copied outside of the organization.

It's important to consider access control of the disk storage system as part of a disk-to-disk approach. With cyberattacks like ransomware, attackers intent to disrupt production and backup files. Attackers that launch ransomware attacks are aware that an organization's first response will be to contain the ransomware and then restore the uncorrupted files from a backup source. If they can compromise the backup and production files, there is a much higher likelihood that the organization will pay a ransom to get their files back. Access control mechanisms should limit access from the system being backed up to the disk array to only those access channels. Do not permit other access to the array from other accounts, including administrative accounts. Remember, everyone can be a potential target of a ransomware attack, especially administrators.

Disk-to-Cloud – This method is very similar to the disk-to-disk backup. Cloud backup offers multiple added values, however. With a disk-to-cloud backup, you automatically get the resiliency and flexibility of the cloud environment as well as the benefits from investments made by the cloud providers to maintain 100% data availability. Rather than a single-point-of-failure model in backup plans like a disk-to-disk or even disk-to-tape, cloud providers replicate data backups leveraging a cloud infrastructure with multi-fault tolerant capabilities.

As with the disk-to-disk model, it's important to limit access to the cloud backup storage to only the systems and disks that are backed up and the data repository. Never implement a drive that maps to the backup repository. That mapped drive could be the vehicle that delivers the ransomware encryption. Always encrypt backup files to protect your organization if the cloud provider is breached.

Lastly, whatever method of backup is used, it's important to test the recovery of these backups on a periodic basis to ensure data availability. As mentioned previously, your backup process is the last line of defense and must be demonstrated to be trustworthy in a time of need.

E. Data Loss Prevention (DLP) (NIST Framework: PROTECT)

Once standard data policies and procedures are established and the workforce is trained to use them, Data Loss Prevention (DLP) systems should be implemented to ensure that sensitive data are used in compliance with these policies.

Multiple DLP solutions exist and can be applicable depending on the types of data access channels that need to be monitored. Traditionally, DLP systems monitor channels that include email, file storage, endpoint usage, web usage, and network transmission. All these channels should be considered.

A challenge with DLP systems is to determine which methods will be used to positively identify sensitive information. Within a healthcare environment, that can be tricky. Generally, there are two approaches and both have limitations:

- Identify sensitive data based on dictionary words that may trigger the inclusion of sensitive data. These dictionaries include robust language repositories that identify health information. The challenge with this technique is related to the terminology. Medical terms are used in the regular course of business without being attributed to any particular sensitive information. This can lead to a high rate of false positives, forcing the workforce to apply prevention practices that are not necessary.
- Identify sensitive data based up identifiers that are known to be sensitive. This may include leveraging tokens that are embedded in documents classified as sensitive and the actual identifiers used for patients. When leveraging document identifiers, the number of false positive rates will be drastically reduced. However, the workforce must be trained on the proper classification. When leveraging actual identifiers used for patients, the false positive rates will be lower as it is positive confirmation. This requires extracting information from the EMR in a regular cadence to load these identifiers into the system. Extra precautions must be taken so that these large data sets are not exposed.

Once your identification methodology is established, DLP systems can be configured to monitor data access channels of interest and make policy decisions based on the data types and the access channel. It's best to provide direct feedback to the user when the data policy has been violated to avoid recurrence. This real-time feedback will help users to adjust their data usage behaviors. Data channels are presented in Table 8 for your consideration.

Data Channel	Implementation Specification	Considerations
Email	Implement inline through SMTP routing for email messages being delivered outside of the organization.	• Define thresholds of risky behavior. Implement a DLP block for these thresholds (e.g., over 100 records of PHI in the email). • Define thresholds of risky behavior. Implement a DLP encrypt action for these thresholds, forcing the message to be encrypted before delivered.
Endpoint	Install DLP agents on managed endpoints that can apply data policies.	• Standardize and deploy encrypted thumb drives to users that require mobile storage options. • Prevent the copying of data to unencrypted thumb drives, or force encryption when copying data.

		• Control the use of non-controlled peripherals and/or storage devices (e.g., backup of iPhones on devices). Permit only when specifically authorized. • Conduct data discovery scans of data resident on endpoints – exposing data on the endpoint so the user can make data destruction decisions.
Network	Implement through Switched Port Analyzer (SPAN) ports from egress network points or through Internet Content Application Protocol (ICAP) on web proxies.	• If online, prevent the leakage of identified unencrypted sensitive data based upon thresholds defined (e.g., files that contain 100 records of PHI). • If out of band, activate IR procedures to contain data leakages that occur through the network.

Table 8. Data Channels Enforce Data Policies

Advanced Practice

A. Advanced Data Loss Prevention (NIST Framework: PROTECT)

After implementing basic DLP controls, you should consider expanding your DLP capabilities to monitor other common data access channels. Table 9 recommends methods for your consideration.

Data Channel	Implementation Specification	Considerations
Cloud Storage	Leverage cloud access security broker (CASB) systems to monitor data flows into cloud systems.	• Label data identified as sensitive. Implement digital rights and encryption to limit access to sensitive data. • Ensure that cloud storage file systems and shares do not expose sensitive data in an "open sharing" construct without authentication (i.e., do not permit the use of sharing data through a simple URL link).
On Premise File Storage	Point discovery scanning systems at known file servers or other large data repositories.	• Conduct regular DLP scans against the file systems to scan and identify sensitive data. • Leverage tools to query security access permissions for each file that contains sensitive data. Define thresholds for excessive access and set alerts if these are crossed. Forward alerts to the SOC for response as described in *Cybersecurity Best Practice #8: Security Operations Center and Incident Response.*

		• Determine staleness of records with sensitive data. Consider executing data destruction practices for records that have not been opened or viewed for an extended period of time. • Determine data ownership of sensitive files identified in file storage systems, leveraging automated tools. Establish workflow options that allow the data owner to provide input into access permission reviews of their sensitive files.
Web Based Scanning	Configure DLP systems to crawl known public websites for sensitive information.	• Conduct a "spearing" exercise, similar to methods deployed by search engines. Compare files and results posted on websites against DLP matching policies and respond quickly to any sensitive data that is exposed. • Conduct manual searching activities on a periodic basis over exposed websites. Look for files that may contain large amounts of sensitive data (e.g., XLS(X), CSV, TXT and PDF).

Table 9. Expanded DLP for Other Data Channels

B. Mapping of Data Flows (NIST Framework: IDENTIFY)

After data business practices are defined, it's advisable to describe these processes in a data map. Data maps should include the following components: applications that house sensitive data, standard direction movement of data, users of applications and data, and methods used to store and transmit data.

Conducting this type of mapping and potentially adding to a larger Enterprise Architecture (EA) reference architecture enables an organization to identify data protection and monitoring requirements.

Threats Mitigated

1. Ransomware Attacks
2. Loss or Theft of Equipment or Data
3. Internal, Accidental or Intentional Data Loss

Suggested Metrics

- Number of encrypted email messages, trended by week. The goal is to establish a baseline of encrypted messages sent. Be on the lookout for spikes of encryption (which could indicate data exfiltration) and no encryption (which could indicate that encryption is not working properly).
- Number of blocked email messages, trended by week. The goal is to detect large numbers of blocked messages which could indicate potential malicious data exfiltration or user education training.

- Number of files with excessive access on the file systems, trended by week. The goal is to enact actions that limit access on the file storage systems to sensitive data, create tickets and deliver to access management
- Number of unencrypted devices with access attempts, trended by week. The goal is to use this information to educate the workforce on the risks of removable media.

References

- (SP 800-122: Guide to Protecting the Confidentiality of Personally Identifiable Information (PII) 2010)

Cybersecurity Best Practice #5: IT Asset Management

Best Practice 5: IT Asset Management	
Data that may be affected	Passwords, ePII, ePHI
Baseline Practices	E. Inventory of Endpoints and Servers A. Procurement B. Secure Storage for Inactive Devices C. Decommissioning Assets
Advanced Practices	D. Asset Pre-Configuration E. Automated Discovery and Maintenance A. Integration with Network Access Control
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

The process by which organizations manage IT assets is generally referred to as IT Asset Management (ITAM). ITAM is critically important to understand and ensure that proper cyber hygiene controls are in place across all assets in the organization. These processes increase the visibility of cybersecurity professionals in the organization and reduce the unknowns.

ITAM processes should be implemented for endpoints, servers and networking equipment. The best practices in this section assist and support every other best practice identified in this publication. ITAM best practices can be difficult to implement and sustain. They should be baked into every life cycle stage of IT operations to establish and sustain data accuracy and integrity. For each asset, this includes procurement, deployment, maintenance, and decommissioning. Though each type of asset deploys different methods during its life cycle, the life cycle itself is consistent.

The Financial Sector, as part of its public-private partnership with NIST National Cybersecurity Center of Excellence (NCCOE), has written a detailed ITAM best practice guide. Though specific to the Financial Sector, the methodologies discussed in the guide are easily applied to the Healthcare and Public Health Sector. *NIST Special Publication 1800-5b: IT Asset Management 2015*

Baseline Practice

A. Inventory of Endpoints and Servers (NIST Framework: IDENTIFY)

The first ITAM component that should be implemented is a build out of the inventory repository. This critical technology component provides a normalized, consistent approach that organizations can use to store inventory data.

Important data elements should be captured for each asset in the ITAM, including the following:

- AssetID (primary key)
- Hostname
- Purchase Order
- Operating System
- MAC Address
- IP Address
- Deployed to (User)

- Last Logged on User
- Purchase Date
- Cost
- Physical Location

The build out of a robust ITAM becomes your single source of truth for all IT assets in your organization. This repository will be maintained and trusted to be highly accurate and actionable.

Special consideration should be given to the differences between ITAM systems and device management systems. Device management systems, which connect to IT devices such as endpoints and servers, can automate the management and maintenance of these assets. They are highly effective at executing tasks such as software discovery, patch management and monitoring performance. Device management systems cannot account for the addition and removal of IT assets or answer the inevitable question, "where did that laptop go?" They manage an organization's devices at a single point in time and are not workflow driven.

IT Service Management tools (e.g., ticketing systems) can be integrated with IT General Controls to ensure accuracy and precision of assets through standard performance management activities. (CIS Control 1: Inventory of Authorized and Unauthorized Devices 2018)

C. Procurement (NIST Framework: IDENTIFY)

Once the ITAM system is implemented and configured, it is important to tie normal supply chain processes with the ITAM processes. The goal is to leverage supply chain processes to proactively register each technology asset, endpoint, server or networking equipment into the ITAM system as it is acquired.

To achieve this, IT organizations must work with Supply Chain departments to streamline technology acquisition channels. When technology acquisitions are specifically categorized, a trigger can be established to capture details of each technology purchase. At a minimum, this can be the generation of a ticket in the IT ticketing systems that prompts a designated IT professional to manually capture details for the new asset when it is acquired. This can be accomplished physically at a shipping dock, or virtually for virtual technology purchases.

In more advanced organizations, the procurement process may be automated to capture salient details for an asset. This reduces the manual labor required and exposure to human error collecting the data.

As the asset is acquired, it is critical to tag it with an asset tag. These tags can be physical or logical. The most important aspect of the tagging process is ensure that the asset has a unique ID that will be used to identify it in the ITAM system. Using other fields (e.g., hostname, IP address, MAC address) as the unique ID may cause complications as these fields may change, potentially creating a duplicate record.

D. Secure Storage for Inactive Devices (NIST Framework: IDENTIFY)

Assets that are not in circulation should be returned to the appropriate IT department for secure storage. Storage areas (e.g., lockers, cages, rooms) should be secured with physical access controls. Access should be limited to those who require it. Physical access controls may include badge readers, video camera surveillance, and door alarms.

If an asset is identified for redeployment, it should be securely imaged to deploy a "fresh" computer system for the new user. This ensures that old sensitive data is removes and the asset has a clean bill of health.

When an asset is sent to storage for redeployment or processing, the ITAM system should be updated to reflect a change of ownership and new physical location (i.e., storage) for the asset. If the asset is redeployed or decommissioned, the ITAM system should be updated again to reflect its new status.

E. Decommissioning Assets (NIST Framework: PROTECT)

It is critically important to properly dispose of retired assets. These assets may contain sensitive information. When executing destruction and certification procedures, update the ITAM in indicate that the device has been decommissioned. This establishes a permanent record in your asset management source of truth, the ITAM. The following procedures should be completed when decommissioning an IT asset.

Central Collection – IT assets should be collected and stored in centralized, physically locked areas prior to decommissioning. Your workforce must be trained to turn in any asset that they no longer use.

Central Destruction/Wipe – Assets that are collected for decommission must complete a secure process to destroy or electronically wipe the storage media. This ensures that devices are properly sanitized before leaving the organization's possession for destruction. Permanent removal of storage media may be completed by your IT organization or an external service provider. It is a good practice to obtain and archive a certificate of destruction for audit purposes.

Record Keeping – Once the IT asset has been cleared for removal from the organization, the ITAM record of the asset information should be registered for destruction or decommissioning. The certificate of destruction can be stored in the ITAM record for easy access. It's highly advisable not to delete the asset record. Instead, update the asset's status in the ITAM system to reflect that it is decommissioned and no longer owned by the organization. You may need to refer to the asset record in the future.

Advanced Practice

A. Asset Pre-Configuration (NIST Framework: PROTECT)

Most large Value Added Resellers (VARs) can preconfigure an asset being purchased by the organization before shipping it to the organization. This is most commonly used for user workstations and laptops, but may be completed with server-based software.

In this method, IT organizations build a "gold-image" for the asset. This image is a fully configured template that includes IT and cybersecurity baselines. A series of tasks are identified that are normally completed by the IT department, such as establishment of encryption, installation of antivirus, and registration of the asset in the ITAM.

Providing these standard tasks to the VAR enables an organization to outsource this common work to third party providers who are contractually bound to complete it accurately and consistently. This provides a higher level of assurance that assets acquired by the organization are properly configured and secured.

B. Automated Discovery and Maintenance (NIST Framework: IDENTIFY)

Once your ITAM system is in place and your procurement processes are registered, the challenge is to maintain these records. A fictitious example describes a common IT asset in a large organization with the following characteristics:

- Number of endpoints: 10,000
- Number of servers: 1,000
- Number of data elements managed per asset: 11
- Total number of data elements required to maintain accurate details: 121,000

It is very difficult to manually maintain 121,000 data elements. After an asset is acquired by an organization, it is often deployed throughout its lifecycle in unforeseen ways. For example, a new laptop may be issued to a user. That user may leave the organization, turning in the laptop to a supervisor. The supervisor may assign the laptop to the new employee who fills the open position. Unless IT is informed and ITAM is updated, the asset record for the laptop, now assigned to a different user, will be wrong.

Another classic example relates to an upgrade or hardware change to an existing asset. This asset might change operating system or patch levels. Maintaining that information manually in the ITAM is nearly impossible.

Automated discovery systems can maintain these records and account for both scenarios described above. In the case where an asset changes hands to a new user, discovery tools can register login occurrences for the "assigned user" and for the "actual logged in user." If a threshold is triggered indicating that the assigned user no longer logs in and a different user continually logs in, that may prompt a change in ownership workflow. This workflow may be automated without requiring intervention or manually completed by generating a ticket to validate the change of ownership. In the case of OS and patching levels, automated discovery systems can provide snapshot views of current patching levels for assets. When these snapshots are compared by cybersecurity vulnerability management systems, vulnerabilities due to obsolete software versions will be identified across the fleet.

C. Integration with Network Access Control (NIST Framework: PROTECT)

The practices outlined so far assume normal processes for acquisitions. There are times, however, when IT assets are integrated in the organization by means other than standard supply chain channels. Examples include personal devices (referred to as Bring-Your-Own-Device, or BYOD) and assets that are donated or provided free-of-charge as part of a third party contract.

Without an oversight function, it is difficult to detect and track these assets. These outliers can be controlled by integrating your Network Access Control and ITAM systems. Further details can be found in *Cybersecurity Best Practice #6: Network Management*.

Threats Mitigated

1. Ransomware Attacks
2. Loss or Theft of Equipment or Data
3. Internal, Accidental or Intentional Data Loss
4. Attacks Against Connected Medical Devices that May Affect Patient Safety

Suggested Metrics

- Percentage of devices added to asset management system through procurement channels, trended over time. The goal is to establish a baseline and achieve a higher percentage over time.
- Number of devices added to the ITAM as a result of NAC, trended over time. The goal is to analyze spikes that occur after initial deployment and may indicate a problem capturing or maintaining asset records.
- Number of devices properly removed from asset management system using proper decommissioning channels, trended over time. The goal is to ensure devices are properly decommissioned. Lack of execution of these processes over a period of time may indicate a compliance issue.

References

- (NIST SPECIAL PUBLICATION 1800-5b: IT ASSET MANAGEMENT 2015)
- (CIS Control 1: Inventory of Authorized and Unauthorized Devices 2018)
- (FIPS 199: Standards for Security Categorization of Federal Information and Information Systems 2004)

Cybersecurity Best Practice #6: Network Management

Organizations leverage IT networks as a core infrastructure to conduct business operations. Without networks, there would be no interoperability capability. It is equally critical to deploy networks securely to limit exposure to and potential impacts of cyber attacks.

Best Practice 6: Network Management	
Data that may be affected	-
Baseline Practices	A. Network Profiles and Firewalls B. Network Segmentation C. Intrusion Prevention Systems D. Web Proxy Protection E. Physical Security of Network Devices
Advanced Practices	A. Additional Network Segmentation B. Command and Control Monitoring of Perimeter C. Anomalous Network Monitoring and Analytics D. Network Based Sandboxing/Malware Execution E. Network Access Control (NAC)
Key Mitigated Risks	• Ransomware Attacks • Loss of Theft of Equipment or Data • Accidental or Intentional Data Loss • Medical Devices and Patient Safety

Baseline Practice

A. Network Profiles and Firewalls (NIST Framework: PROTECT)

An effective network management strategy includes the deployment of firewalls to enable proper access inside and outside of the organization. Firewall technology is far more advanced than standard router based access lists, and is a critical component of modern network management. Organizations should deploy firewall capabilities throughout the following areas: on WAN pipes to the Internet and perimeter, across data centers, in building distribution switches, in front of partner WAN/VPN connections and over wireless networks.

There should be clear boundaries that determine how traffic is permitted to move throughout the organization, leveraging a default-deny rule set whenever possible. At the perimeter, inbound and outbound rules must be configured with a default-deny rule set to limit accidental network exposures. This often complicated process can be achieved by establishing security zones through network segmentation.

Consider limiting the outbound connections permitted by assets in your organization. This can be a challenge to implement across the board. However, for particular zones of high sensitivity, egress limiting can prevent malicious callbacks or data exfiltration. The SOC should monitor these egress logs.

Firewall rules may change when technology is added or removed. A robust change management process should include reviewing every firewall to identify necessary changes. These change requests should comply with standard IT Operations Change Management processes and be approved by cybersecurity departments before any firewall is modified.

As part of standard rule management for firewalls, it's important to periodically review firewalls to ensure they are properly structured as required by cybersecurity teams. Consider a monthly or quarterly review of the highes risk rule sets.

B. Network Segmentation (NIST Framework: PROTECT)

A fundamental method to limit cyberattacks is the practice of partitioning networks into security zones. These zones can be based on sensitivity of assets within the network (e.g., clinical workstations, general user access, guest networks, medical device networks, building management systems, IOT networks) or standard perimeter segmentations (e.g., DMZ, middleware, application servers, database servers, vendor systems). Examples of standard network zones are:

Perimeter Defenses – Most organizations host services that are accessed through the Internet. A robust defense strategy should be deployed to monitor these "front doors" (CIS Control 12: Boundary Defense 2018).

Best practices for perimeter defenses include:

- Implement highly restrictive rules on inbound and outbound ports and protocols. Leverage default-deny rules in firewalls and enable access only when clearly understood.
- Restrict DMZ from Middleware, Application and Database servers. DMZ controls are critical as these servers are exposed to the Internet and have a large threat footprint.
- Restrict the ability for DMZ servers to log in directly to servers on the inside network, specifically using remote desktop protocol (RDP), server message block (SMB), secure shell (SSH) or other remote access ports (tcp/3389, tcp/445, tcp/139, tcp/22).
- Ensure that local administrator passwords are unique to each DMZ server and do not use these passwords for any other server in the organization.
- Ensure that DMZ servers cannot connect directly to the Internet. Instead, these servers should access the Internet through outbound proxy services. Outbound proxy rules should limit the sites, URLs, IPs and ports that a DMZ server can access to only those whitelisted sites required for updates or application functionality. Be cautious of whitelisting hosting organizations like AWS: they may be used by malicious actors to download malware to a compromised server.
- Consider this type of posture for partner WAN links or site-to-site VPN connections. Do not permit access to systems/applications that are not required by the user.

Data Center Networks – Servers in the data center should be segmented into appropriate zones. Several different layers of segmentation may occur within the Data Center networks, including

- Database servers;
- Application Servers; and
- Middleware.

*Critical IoT Asset*s – It is important to restrict access to assets that have a potentially high impact to the business or patients if compromised. A common factor of these devices is that management and patching of security vulnerabilities may be limited. Examples include medical devices, security cameras, badge readers, temperature sensors, and building management systems. Generally speaking, these

assets exist outside of the data centers. Without proper segmentation, they may infiltrate general access networks. To achieve segmentation in the physical buildings, leverage multiprotocol label switching (MPLS) to build out virtual networks and place these restrictions behind core firewalls.

Vendor Access – Vendor access should be limited based on need. It should be temporary and provided only to access required information. Some assets are managed exclusively or accessed by third party vendors. These vendors may need continual access to the organization's network. It's important to segment this vendor access from other networks and limit the vendor's ability to access other parts of your corporate network. Whether these networks exist inside or outside of the data center, the principles are the same. In 2015, Target was the victim of a cyberattack leveraging these exact channels (Anatomy of the Target data breach: Missed opportunities and lessons learned 2015). Common examples include building management systems, security systems, physical access controls, and persistent tunnels required to enable cloud functionality.

*General Access Network*s – The vast majority of your workforce will operate on general access networks. These are "edge" networks that provide connectivity back to the services offered in data centers, the Internet, or other assets. These networks require a sense of openness when communicating with services that are hosted by the organization. However, restrictions should be implemented that prohibit the assets in one general access network from communicating with the assets in another general access network. This critical control that can help stop the outbreak and spread of malware and ransomware attacks.

Guest Networks – It is common for most organizations to provide guest access to the Internet. This is especially important in provider organizations visited by patients and their friends and families. Access to the Internet is a core value of provider organizations. However, it must be restricted and controlled appropriately. These restrictions should exist on wireless networks, where it is most common, as well as wired networks often located in public spaces or conference rooms. Explicitly prohibit access to the internal network: guest users should access the organization through the same front door as the rest of the Internet. Lastly, as much as possible, limit the ability of your workforce to access guest networks.

C. Intrusion Prevention Systems (NIST Framework: PROTECT)

An Intrusion Prevention System (IPS) is important for your network perimeter, data center and partner connections. These systems are capable of reading network traffic to detect and potentially prevent known attacks.

Today, the effectiveness of these signature-based systems is not as prevalent as they used to be. However, they do serve as vital input to an organization's SOC that provide context to the types of attacks that occur. Though they might not identify every single attack, they provide information you're your IR team to conduct forensic capabilities.

D. Web Proxy Protection (NIST Framework: PROTECT)

Web proxy systems provide incredibly important protections against modern phishing and malware attacks. These systems are implemented at the perimeter of the network or in the Cloud to provide protections for your mobile workforce. As most of these attacks are web-based, web proxy systems provide users with friendly error pages explaining that the user has been restricted from accessing a known malicious website. This is great feedback for users. When configured properly, web proxy systems leverage the following methods to limit client-side attacks:

Reputation Blocking – Many blackhole lists are available publicly or through ISACs/ISAOs that prevent users from accessing malicious websites. These are usually fed into proxy systems through automated list and feeds.

Organizational Block Lists – As part of an organization's IR, malicious websites and other sites can be identified based on actual attacks against the organization. Web proxy systems are critical shut-off points to limit access to websites quickly.

Category Blocking – Most modern and commercial web proxy technologies will pre-categorize websites on behalf of the organization. Considering the millions of websites that exist, this is an incredibly useful service. Consider blocking categories that contain malicious, suspicious, or illegal websites.

E. Physical Security of Network Devices (NIST Framework: PROTECT)

Network devices are deployed throughout an organization's facilities. Inside the general user space, physical data closets that contain network devices must be secured. Additionally, it is useful to limit network ports on switches. Consider the following controls:

- Data and network closets should be locked always. Consider using badge readers instead of traditional key locks to monitor access.
- Disable network ports that are not in use. Ensure that procedures are in place to maintain ports in shutdown mode until an activation request is submitted and approved.
- Establish guest networks in conference rooms that are configured to access only these networks.

Advanced Practice

This section includes methods to detect and potentially prevent cyberattacks against an organization's network. These methods should be engineered into network management practices. Once established, cybersecurity departments can follow *Cybersecurity Best Practice #8: Security Operations Center and Incident Response* to monitor and respond to attacks on the network.

A. Additional Network Segmentation (NIST Framework: PROTECT)

As your network expands, other strategies can be deployed to maintain secure segmentation. Consider the following:

Required VPN Access for Data Center – Consider leveraging a VPN, or bastion hosts, that must be enabled before access is granted to privileged servers in the data center. These VPN or bastion hosts should have Multi-Factor Authentication. Only authorized IT administrators should be granted access. Logs should be routed to the SOC for monitoring.

B. Command and Control Monitoring of Perimeter (NIST Framework: DETECT)

Layered command and control (C2) traffic is a common mechanism used by hackers to maintain access to compromised computers. These are beacons, typically outbound from the computer, that check back in to a central server. This control can detect where an attacker has maintained persistence. There are many methods to look for C2 traffic. These include:

Direct to compromised server via Internet Protocol (IP) or Internet Control Message Protocol (ICMP) – In this method, traffic runs over the network using outbound ports or protocols that are generally open (e.g., HTTP, HTTPS or ICMP protocols). C2 traffic can be encrypted or cleartext, depending on the attacker's level of sophistication. The attacker must have compromised a series of servers or stood up

virtual servers that are accessible. This method tends to be less effective than others as it is easy to shut down offending systems once the compromise has been detected. When shut down occurs, the attacker loses persistence control.

DNS queries – In this method, the attacker establishes control using a DNS query embedded in malware that is downloaded to a computer. As long as the DNS record is maintained, the servers that maintain C2 communications can switch out and flex as they are discovered. Generally speaking, this method is also fairly easy to detect and resolve. When the DNS name has been identified, the organization can implement a DNS Sinkhole. This Sinkhole can be an entry on the local cache in the organization's DNS resolvers to remove a non-existent IP address, such as 127.0.0.1. Once the Fully Qualified Domain Name is identified, these DNS registrations can be taken down through Abuse reporting to DNS hosting services.

Fast flux DNS queries – In this method, the hacker leverages DNS to maintain persistence, knowing that the DNS registrations will likely be taken down at some point. When this occurs, malware downloaded to the local client and C2 services runs an algorithm that checks the first several bytes of well-known sites (e.g., cnn.com, nbc.com) to create and register fake DNS names on the organization's DNS resolvers. These domains tend to live for 24 hours or less. Using the same algorithm, the clients switch to the next domain until command is re-established. These methods are fairly successful. Defense mechanisms for Fast Flux DNS queries require analytics that relate to local DNS lookups and discover "gibberish" domain names. "Oiewr921ai.evil.com" is an example of a gibberish domain name.

C. Anomalous Network Monitoring and Analytics (NIST Framework: DETECT)

A variation on C2 monitoring is to analyze network traffic, rather than focus on a particular vector or attack style. This requires specialized technologies that can profile inbound and outbound network traffic. Some versions of these tools provide "Deep Inspection," which allows the full contents of a packet to be analyzed, categorized and built into massive databases of network-based metadata.

Once metadata is gathered on the network traffic profile, analytics can be conducted to look for outliers, anomalous traffic, and other highly sophisticated methods of discovery. These tools are not preventative in nature. They are intended to widely increase the SOC's visibility, facilitating detection, confirmation or validation of suspicious actions. These tools are highly useful in replaying events that occurred as part of an attack to support network forensic activities.

D. Network Based Sandboxing / Malware Execution (NIST Framework: DETECT)

By monitoring common protocols that allow downloading of binaries and files, organizations can check a download prior to permitting it to run on the organization's devices. These binaries, executables, or even data files (e.g., docx, xlsx) are run in a virtual environment that looks for malicious activities when the file executes.

Common methods include:

- Watching what registry keys are queries, amended, added, or deleted;
- Monitoring for outbound network connections;
- Launching processes in memory; and
- Conducting anomalous system calls.

Tools that facilitate automated sandboxing look for suspicious outputs or actions rather than attempting to base actions on a particular signature of a particular configuration.

To be effective, these technologies monitor network flows. This can occur passively or actively. Passive systems monitor network traffic at the stream level, not residing in line with the communication flows. Active systems insert themselves inline to the communication flows and conduct checks on the fly, denying access to downloaded files until they are cleared.

These systems provide protection against malicious files. However, they do not provide protection against active attacks inside your network.

E. Network Access Control (NAC) (NIST Framework: PROTECT)

NAC systems are engineered to automatically profile new IT assets that connect to network resources, such as wireless networks, wired networks or VPN. They help ensure that the controls discussed in *Cybersecurity Best Practice #2: Endpoint Protection* are in place on each asset. NAC systems execute these controls on a real-time basis when the asset connects to the network.

NAC systems are highly effective at discovering personal devices leveraged on the network (BYOD). They can be configured to permit authorized BYOD devices to access the network or prohibit them entirely.

When basic NAC controls are implemented and you can monitor the security of endpoints that connect to your network, there are other interesting and advanced techniques that can be leveraged to provide checks and balances for general IT controls.

One example is to integrate your NAC solution with your ITAM repository. As discussed in *Cybersecurity Best Practice #5: IT Asset Management,* ITAM repositories should be populated using your organization's standard procurement processes. That said, not all processes run perfectly and there are other ways that assets are integrated into an organization's environment. This often occurs due to human error or sidebar procurement channels that are not leveraged consistently.

Configuring your NAC solution to check against your ITAM enables assets to be profiled spontaneously, providing self-directed work streams to users. This can be achieved by doing the following:

- Set up application programming interfaces (API) between the NAC solution and the ITAM solution that enable read and write options.
- Query the ITAM database when an asset connects to the network. If the asset does not exist, present the user with a splash page.
- Determine if the asset is organizationally owned (purchased with organizational funds) or personally owned (purchased by the user).
- Register the selection, conduct the NAC security scan, and publish the results in the ITAM.
- Execute IT general controls that reconcile assets that are out of compliance with standard Asset Management procedures. This can include:
    - Ensuring that appropriate monitoring controls are in place;
    - Registering the asset with the right identifiers (Asset IDs); and,
    - Updating asset ownership based on actual human interaction.

These highly effective mechanisms provide visibility into the actual devices being used on the network, increasing ITAM accuracy and consistency.

Threats Mitigated

1. Ransomware Attacks
2. Loss or Theft of Equipment or Data
3. Internal, Accidental or Intentional Loss of Sensitive Data
4. Attacks Against Connected Medical Devices that May Affect Patient Safety

Suggested Metrics

- Number of assets on the network that have not been categorized, trended over time. The goal is to establish a process to register and understand all assets on the network. After the baseline is complete, minimize the number of uncategorized assets.
- Number of organizationally owned assets discovered using NAC that were not previously categorized through Asset Management procedures, trended by month. The goal is to monitor this lagging metric that measures effectiveness of the supply chain and IT operations processes. Increases in the number of organizationally owned assets that were not previously categorized indicates that standard processes are not being executed properly. Implement continuous improvement processes for IT operations.
- Percentage of assets that comply with security policies, trended by week. The goal is to establish a baseline, then set stepwise goals to improve compliance over time. Ultimately, compliance percentage should range from 95% to 99%.
- Number of malicious files captured/secured with advanced networking tools (sandboxing), trended by week. The goal is to capture all malicious files. An extended trend of no detected malicious files may indicate that sandboxing solutions are not working.
- Number of malicious C2 connections discovered and removed, trended by week. The goal is a weekly report that shows all detected C2 connections are mitigated successfully.
- Number of approved servers/hosts in the DMZ compared to hosts in the DMZ, trended by week. The goal is for zero servers/hosts in the DMZ that are not understood. IT Operations practices should be reviewed if servers are added that were not previously authorized.

References

- (CIS Control 12: Boundary Defense 2018)

Cybersecurity Best Practice #7: Vulnerability Management

Best Practice 7: Vulnerability Management	
Data that may be affected	-
Baseline Practices	A. Host/Server Based Scanning B. Web Application Scanning C. System Placement and Data Classification D. Patch Management, Configuration Management, Change Management
Advanced Practices	A. Remediation Planning
Key Mitigated Risks	• Ransomware Attacks • Insider, Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices

Vulnerability management is used by organizations to proactively conduct vulnerability discovery processes. These processes enable the organization to classify, evaluate, prioritize, remediate, and mitigate the technical vulnerability footprint from the perspective of an attacker. The ability to mitigate vulnerabilities before a hacker discovers them gives a competitive edge to the organization along with time to address these vulnerabilities in a prioritized fashion.

Vulnerability scanning comes in multiple flavors. Generally speaking, the most well-known methods are scans against servers (or hosts) and against web applications. Both scan types focus on different considerations.

Baseline Practice

A. Host/Server Based Scanning (NIST Framework: DETECT)

In this model, vulnerability scanners are leveraged to identify weaknesses in an operating systems or third party application that resides on a server. There are two types of scan options: unauthenticated and authenticated.

In the unauthenticated model, the scanner has no extra sets of server privileges and queries the server based on ports that are active and present for network connectivity. Depending on the level of sophistication of the software scanner, each server is queried and checked for vulnerabilities. Scan results provide the perspective of an attacker who lacks server access. Vulnerabilities that rate high in this space should be mitigated first, as they are the most likely entry points to the server for a hacker.

Authenticated scans are conducted by letting the vulnerability scanner log in to the server and query all running software with all running versions. The vulnerability lists are usually compared against a database (maintained by the scanner's vendor) and vulnerabilities are enumerated based on known disclosed issues from the known software version. While this type of scanning provides a much higher degree of accuracy of the enumeration of the vulnerability, it doesn't necessarily provide context that describes how the vulnerability may be exploited. The other advantage of this type of scan is that it will identify client-side vulnerabilities that may exist on the server and otherwise be difficult to discover, such as vulnerable versions of Java.

Most scanning systems can categorize vulnerabilities against the MITRE Common Vulnerability Scoring System (CVSS). The CVSS system helps organizations to prioritize identified vulnerabilities which enables development of a prioritized response. Version 3 of the system considers the following three factors: Base Score, Temporal Score, and Environmental Score. These factors along with sub-factors are used to

calculate vulnerabilities on a scale ranging from one to ten. For more information, refer to (Common Vulnerability Scoring System SIG n.d.).

B. Web Application Scanning (NIST Framework: DETECT)

In this model, specialized vulnerability scanners are leveraged that interrogate a running web application to check for vulnerabilities in the application design. Most web applications run dynamic code, run atop of a web server, interact with middleware, and connect to databases. If the web application is not coded securely, this architecture may enable access to data or systems in ways that are unanticipated.

Common and popular attack types of web applications include structure query language (SQL) Injection, Cross-Site Scripting, and Security Misconfigurations. In these cases, attackers can:

- Bypass web application security controls and pull data directly from the database;
- Steal an already authenticated cookie on a vulnerable website to get access; or
- Leverage misconfigurations that can permit properly formatted commands or scripts to execute privileged content on the webserver itself.

More information can be found on the Open Web Application Security Project (OWASP) Top 10 Website (OWASP Top 10 - 2017: Ten Most Critical Web Application Security Risks 2017).

In all cases, vulnerabilities to web applications with sensitive information represent a high risk to the organization. It is important to understand these vulnerabilities to conduct appropriate and prioritized remediation.

C. System Placement and Data Classification (NIST Framework: IDENTIFY)

Organizations should leverage *Cybersecurity Best Practice #4: Data Protection and Loss Prevention* and *Cybersecurity Best Practice #5: IT Asset Management* to understand IT assets and asset classifications. These best practices answer the question, “How bad would it be if this asset was breached?”

It is important to understand the exposure or each system in your environment. Organizations should leverage *Cybersecurity Best Practice #6: Network Management* to determine the likelihood that each this system can be compromised.

The level of risk related to vulnerabilities in your systems is directly related to the exposure of these systems and the types of data they contain.

D. Patch Management, Configuration Management, and Change Management (NIST Framework: PROTECT)

All organizations should have a routine activity to patch security flaws in their servers, applications (including web applications), and third party software. Although the patching process may vary, large organizations should leverage centralized systems to interrogate servers and determine which software updates should be implemented.

At least monthly, organizations should implement patches that are produced by the vendor community. IT Operations should collect these patches, conduct appropriate regression tests to ensure each patch does not negatively impact the business, and schedule patch implementation during routine change windows. This process should be executed and measured using standard IT Operations activities.

Not all vulnerabilities are created equal. Some are easier to exploit than others. The National Vulnerability Database (NVD) has produced the CVSS. The CVSS is a standard measurement across all industries that normalizes and ranks the severity of a particular vulnerability. (NVD NIST 2018)

Generally speaking, the more a particular vulnerability is exposed, the higher priority an organization will assign to mitigate it. Exposure may be more critical than the actual impact to an asset considering that hackers attempt to gain a foothold on organizational assets before conducting additional internal attacks. Another factor to consider is the level of active exploitability in the wild. A lower criticality vulnerability may have an active threat against it. In these cases, organizations might want to consider proactively executing IR processes, organizing the response team and quickly patching systems. The WannaCry exploit of 2017 was a classic example of organizations identifying an active threat and quickly implementing patches that were previously neglected.

If your systems are running end of life OS or software, associated vulnerabilities should be identified and steps taken to bring these systems back to a supported state. This may include decommissioning application systems that run on unsupported OS, which may require additional investments by the organization. Once systems are unsupported, it is generally impossible to apply security patches. This may increase the organizations risk posture.

Table 10 provides general guidelines to plan remediation efforts based on criticalities.

Vulnerability Criticality	Days to Mitigate in DMZ	Days to Mitigate in Data Center
Critical	< 14 days	< 30 days
High	< 30 days	< 90 days
Medium	< 90 day	< 180 days
Low	< 180 days	At your discretion

Table 10. Vulnerability Criticalities Drive Mitigation Timeframes

The vulnerability scanning process is a quality check on the effectiveness of an organization's patch management practice, otherwise referred to as a "lagging metric." Organizations with a robust patch management practice are better positioned to mitigate residual vulnerabilities.

In addition to conducting routine patch management activities, organizations should ensure that proper security configuration management activities are in place. Common vulnerabilities can be introduced in systems with insecure configurations. Examples of insecure configuration includes permitting an FTP server to allow anonymous login, making that login accessible to the Internet, or failing to change default account passwords on applications. Organizations that follow *Cybersecurity Best Practice #2: Endpoint Protection Systems* and expand those practices to their servers will be positioned to minimize these issues.

Lastly, consideration needs to be given to changes made to systems, servers, and networks along with the vulnerabilities that may be exposed as a result. A testing plan should be part of the change management process. It should include a vulnerability scan of new network connectivity (such as a firewall change) or a new system function or service. This scan should be conducted during the test

phase of the change process, before the change is implemented into the production environment. This enables the identification and mitigation of security exposures.

Advanced Practice

A. Remediation Planning (NIST Framework: PROTECT)

It is important to better classify and prioritize vulnerabilities that remain after completion of standard patch management practices. These are issues that generally cannot be mitigated with a patch. They may require system configuration changes, code updates, or perhaps even a full-blown version upgrade. The process of resolving these vulnerabilities tends to be more time-consuming and complex.

Similar to risk management activities, remediation efforts should be prioritized to resolve identified vulnerabilities. The most common practice is to first patch identified vulnerabilities and rescan the system to validate the vulnerability is closed. Most vulnerability scanning systems have the ability to track the opening, closure, and reopening of vulnerabilities over time. It is highly encouraged to leverage these metrics for institutional tracking.

Mitigation of some vulnerabilities requires far more effort than a simple patch. In these cases, it is best to develop structured remediation plans that include the following elements:

Remediation Owner – The single individual accountable for ensuring the vulnerability, or vulnerabilities, are addressed. It is important to assign a single owner on remediation plans otherwise they are likely to stall due to a lack of leadership.

Plan – A full description of the remediation plan to be completed. This plan should be developed by the remediation plan owner and security office. Once the plan is approved, execution tasks can be started.

Stakeholders – The individual stakeholders who are responsible for completing tasks in the remediation plan, or organizing other individuals who will complete these tasks. This may include individuals who need to be informed of remediation activities as well as individuals who actually complete the work.

Dates – Major milestone dates and remediation plan due dates must be captured on the remediation plan. This is an incredibly important commitment that must be made by the Remediation Owner.

Status – Periodically, the plan should be updated to remain current. This generally occurs between once a week to once a month. The Remediation Owner may be accountable for providing status updates.

After a remediation plan is completed, a monitoring process should be implemented by the organization's information security office. This monitoring process may include all remediation plans in progress and current activities. The security office may provide support to those activities that are behind schedule. Consider operating such a monitoring process once a week to keep good traction.

Threats Mitigated

1. Ransomware Attacks
2. Internal, Accidental or Intentional Data Loss
3. Attacks Against Connected Medical Devices that May Affect Patient Safety

Suggested Metrics

- Stacked aggregate of vulnerabilities in DMZ trended by month, with vulnerabilities categorized using CVSS categories (Critical, High, Medium, Low, None) and plotted as a simple stacked bar. The goal is to mitigate the most severe vulnerabilities first through

patching and configuration management. Of the remaining vulnerabilities, the most critical should be mitigated within 30 days. The total number of vulnerabilities should be reduced over time.

- Stacked aggregate of vulnerabilities in data center trended by month, with vulnerabilities categorized using CVSS scores and plotted as a simple stacked bar. The goal is to mitigate the most severe vulnerabilities first through patching and configuration management. The total number of vulnerabilities should be reduced over time.
- Number of unmitigated new vulnerabilities introduced into the environment trended by week. The goal is to keep the number of new vulnerabilities as low as possible, defined by your organization's level of risk tolerance.

References

- (CIS Control 4: Continous Vulnerability Assessment and Remediation 2018)
- (OWASP Top 10 - 2017: Ten Most Critical Web Application Security Risks 2017)
- (NVD NIST 2018)
- (Common Vulnerability Scoring System SIG n.d.)

Cybersecurity Best Practice #8: Security Operations Center and Incident Response

Best Practice 8: Security Operations Center and Incident Response	
Data that may be affected	-
Baseline Practices	A. Security Operations Center B. Incident Response C. Information Sharing and ISACs/ISAOs
Advanced Practices	A. Advanced Security Operations Center B. Advanced Information Sharing C. Incident Response Orchestration D. Baseline Network Traffic E. User Behavior Analytics F. Deception Technologies
Key Mitigated Risks	• Phishing Attacks • Ransomware Attacks • Loss or Theft of Equipment • Insider, Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices

Most cybersecurity programs begin by implementing controls designed to prevent cyber attacks against an organization's IT infrastructure and data. This is a good place to start and there is a lot of value in basic cyber hygiene, leveraging the best practices that are discussed in this volume. However, in the modern age of cyber threats, not all attacks can be prevented with these basic controls. It is equally important to invest in and develop capabilities to detect successful attacks and respond quickly to mitigate the effects of these attacks.

A good example is the threat of phishing attacks. Even if organizations followed every best practice discussed in *Cybersecurity Practice #1: Email Protection,* they will still be susceptible to phishing attacks. It is incredibly important to detect in near real-time a phishing attack that successfully infiltrates your environment and to neutralize its effects before widespread theft of credentials or malware installation occurs. This is a classic example of shoring up your detection capabilities (detecting the phishing attack that gets past your basic controls) and response capabilities (neutralizing the effects before serious damage to the organization occurs).

Maintaining these capabilities requires the establishment of an IR program and a SOC to manage the IR along with security engineering that enhances an organization's ability to detect and response to cyberattacks.

Baseline Practice

A. Security Operations Center (NIST Framework: DETECT)

A SOC is an organizational structure that leverages cybersecurity frameworks, people, tools, and processes to provide dedicated cybersecurity operations for an organization. Generally speaking, SOCs are the areas within an organization that dedicate 100% of their time to cybersecurity prevention, detection or response capabilities, providing the execution arm of cybersecurity IR.

A SOC is generally segmented into four main functions, depending on the organization's level of maturity. These functions are:

Engineering – Security Engineering is the process of building new cybersecurity capabilities into the existing toolsets in an environment. Examples include building new alerts within a Security Incident and Event Management (SIEM), establishing new log sources for log management systems, establishing new

analytics patterns for detection, or simply implementing new cybersecurity systems to add capabilities into the environment.

Operations – Security Operations is the process of managing and maintaining the cybersecurity tools within the SOC. This is sometimes referred to as "Keeping the Lights On." This generally means monitoring critical cybersecurity systems to ensure they operate at agreed upon performance levels.

Threat Intelligence – A specific type of function that focuses entirely on how to discover cybersecurity threats that may be relevant to the organization along with the means and methods these threats may use to infiltrate the organization. This function focuses on the threat actors themselves, the tools they leverage, and the digital signatures they leave in the process of conducting their activities. Once these digital footprints are established, sometimes called an Indicator of Compromise (IOC), engineering teams can implement these patterns into the systems and establish IR plays to execute when activated.

Incident Response – IR is the process of conducting a structured and consistent response to any IR plays that have been created. The goal of this function is to:

- Validate an IR play that has been triggered;
- Contain any successful cybersecurity attacks to the organization;
- Eliminate the threat from the environment;
- Recover systems or data that might have been affected by the attack; and,
- Ensure that any attack vectors that were exploited are well understood and fed back to the Security Engineering teams for future prevention or enhanced detection capabilities to further minimize the impacts of those vectors.

It is critically important to instill a continuous feedback loop between your IR and Engineering teams so the organization continues to learn and grow based on the actual success of threats and threat actors.

As SOCs are implemented, a core concept is to ensure that IR teams and handlers apply consistent methods to execute response practices. SOCs and IR teams should establish a Playbook, also known as a Runbook that describes existing detection mechanisms and the procedures to be followed if the detection mechanism is triggered. This may be referred to as a "Play," similar to Plays that football teams maintain in their Playbooks.

Examples of Plays that might be found in a Playbook are provided in Table 11.

Play Category	Play	Description	Source Data
Reconnaissance	Vulnerability scanning sweep of DMZ.	Large number of vulnerabilities being scanned across the DMZ spectrum. May be a single server being scanned over multiple ports or multiple servers being scanned on a single port.	• Server list in DMZ. • Intrusion Detection System (IDS)/Intrusion Prevention System (IPS) logs configured to detect vulnerability scanning. • Firewall logs. • Netflow data.

Reconnaissance	Vulnerability scan from known malicious IPs.	Vulnerability scans of the DMZ or other servers/endpoints exposed to the Internet over channels that are shared and known to be malicious (Indicator of Compromise).	• IOC list from threat sharing sources (e.g., ISACs). • IDS/IPS Logs. • Firewall logs. • Netflow data.
Reconnaissance	Successful access from known malicious Ips.	Successful authentications from known malicious IP addresses. Authentications through standard remote access channels, such as VPN, virtual terminals, jump boxes, or other mechanisms.	• Authentication Logs. • Firewall Logs. • IOC list from threat sharing sources (e.g., ISACs).
Reconnaissance	Internal attacks from third party VPNs.	Detection of attacks coming through partnering third party VPN connections, such as organizations that provide building automation services.	• Firewall Logs (from segmented networks). • IDS/IPS Logs. • Authentication Log.
Persistence	Creation of local user accounts on static systems.	Detection of a local user account being created on an asset, such as a Windows *nix server, where local user account creations normally do not occur. This may indicate malicious activity.	• Logs from local servers.
Persistence	After exploit persistence hold.	Detection of malicious users attempting to maintain permanent access. Look for launch or changing of scheduled tasks, script downloads, and new process creation.	• Critical server lists. • Known process baselines. • Logs from server task or scheduled job management. • URL Filtering logs by server.

Privilege Escalation	Privileged Account brute force success.	Large numbers of invalid login attempts followed by a successful login to a known privileged account.	• Privileged Account List. • Authentication Logs (e.g. Active Directory, Servers).
Privilege Escalation	Default Account password guessing.	Large number of invalid login attempts followed by a successful login to a known default user account.	• Default account list. • Authentication Logs (e.g. Active Directory, Servers).
Privilege Escalation	Interactive login to Service Accounts.	Detection of a service account being used as an interactive login (a user logging in to a terminal session). Service accounts should only be used for applications or services.	• Service account list. • Authentication Logs (Active Directory, servers).
Data Exfiltration	Data Transfer.	Detection of data transfers occurring outside of the organization from servers that normally do not conduct such activities. Must normalize/baseline server network behavior and detect anomalous activities off baseline.	• Netflow data, or Firewall traffic profile data. • List of permitted remote storage sites (e.g. Box).

Table 11. IR Plays Describe Responses to Cyberattacks

In each of these cases, the source data provided will include events or log information that is critical to detect the play being constructed. Specialized security systems can ingest these logs and apply pattern matching, rules matching, and analytics matching capabilities to specific events in the logs to call out potential incidents of interest. These specialized systems are referred to as SIEM systems.

This section provides details for the high level play, what it seeks to accomplish, and the types of source data that must be collected to successfully detect it. The list below will not discuss specific technical log event data that is required. Information on how to configure this information can be found in multiple publications, such as (Swift 2010) and (The 6 Categories of Critical Log Information 2013).

B. Incident Response (NIST Framework: RESPOND, RECOVER)

One of the most basic, and most important, functions in a cybersecurity organization is the IR process. This process provides the organization with standardized procedures to respond to cyberattacks. The

attack may be as simple as an attempted phishing attack against users or a highly sophisticated extortion attack that shuts down digital operations. In both cases, from minimal to significant impact, the organized manner of an IR is critical to manage these threats.

The following procedure is a summary of (SP 800-61r2: Computer Security Incident Handling Guide 2012). Generally, a structured IR process is segmented into the following steps:

Preparation – Before you respond to a cybersecurity incident, it's important to have policies, processes, and procedures in place. This includes the following components:

- *IR Policy* – a policy that defines the categorization and severity of incidents, the stakeholders involved in IR, the roles and responsibilities of each person, the entry criteria when a security incident occurs and the person who is in charge of IR plays. Stakeholders may range from the standard blocking and tackling personnel in IT Operations to legal, marketing and public affairs personnel for high impact incidents. A template IR policy is provided as *Appendix I in the main document*.
- *Cybersecurity Incident Response Team (CIRT)* – The CIRT is a pre-formed and "on the ready" group that knows how to navigate issues when Critical or High severity security incidents arise. This team develops and manages your organizational response. Most commonly, CIRTs formed in the HPH Sector when potential data breaches occur and the organization must manage the potential breach in compliance with HITECH. It is important to identify the Incident Commander, the most senior official who will be in charge of managing cybersecurity incidents. The Incident Commander is usually the CISO or equivalent in an organization. It's important to note that the Incident Commander should not dive into the technical weeds of the incident, but keep the various teams organized and focus on their objectives. For example, Table 12 describes that teams that may be involved in resolving a critical security incident and potential breach.

Team	Description
Executive/Senior Leadership	This is an organization's C-Suite or most senior executives. They provide overall direction and approvals required to resolve significant cybersecurity breaches. These individuals should be kept informed throughout the life cycle of a significant cybersecurity incident.
Cybersecurity Teams	These teams are comprised of people with cybersecurity expertise who understand attacks, vulnerabilities and the methods by which threat vectors are exploited. They provide technical depth and detail to technical teams and execute procedures in the Playbook.
Technical Teams	These teams are comprised of subject matter experts (SME) for the technologies that have been compromised and are engaged in developing and implementing the response. These SMEs may be system owners, system administrators, or other individuals with specialized IT expertise. They take instruction from the cybersecurity teams as part of the Playbook execution.

Legal Teams	These teams are comprised of attorneys in your general counsel (internal or external) that help manage the incident under privilege as well as consult on regulatory expectations.
Public Affairs/Marketing and Communications	These people manage external communications to deliver a consistent voice and message in the event of a high visibility cybersecurity incident. It is critically important to manage the reputation of the organization.
HIPAA Privacy Teams	These teams are responsible for understanding the full extent of a cybersecurity incident that involves PHI. This includes conducting a breach analysis process in compliance with HITECH and providing consult for any patient-facing communications that should occur.

Table 12. Roles and Responsibilities to Respond to a Critical Cybersecurity Incident

- *Playbook and/or Runbook* – An organization's Playbook contains sets of standard operating procedures to respond to different types of cyberattacks. Procedures to respond to a phishing attack are different from those required to respond to a system intrusion or a ransomware attack. Each of these three types of attack is a distinct "Play" in an organization's cybersecurity "Playbook." For each Play, it's important to describe the steps that will be followed to mitigate the attack so that your response is not "made up on the fly." Though each particular attack has its own unique characteristics and nuances, your procedural processes should follow the steps provided in your Playbook for that type of attack. A template playbook is provided in *Appendix I in the main document*.

- *Tools and Technologies* – Once you establish your policies, CIRTS and Playbook, the next level of improvement is to configure your tools and technologies to streamline the execution of your plays. This connects your IR processes back to your Security Engineering processes to create a continual feedback loop that is essential to becoming a resilient organization.

Identification – The first response to any cyberattack is to understand the scope and extent of the attack. The identification phase of an attack sets in motion the process of categorizing and classifying components of the attack based on your policies and procedures. Critical and sophisticated attacks warrant a well-organized and effective response.

For example, a general phishing attack against a small user set that is easily identified as malicious may be assigned a lower level of concern than a targeted phishing attack against a select user base leveraging the nomenclature of your organization. These highly specialized attacks are known to be very successful and can easily compromise a user's credentials or introduce remote access malware into your environment.

An example of the identification exercise in a phishing process may be as simple as the following:

- Receive notification from your user base or through your own detection systems of a phishing attack or campaign.

- Profile and understand the extent and scope of the phishing attack. Determining its level of sophistication and intent.
- Conduct a basic investigation to determine if links were clicked or malware was delivered.

Containment – After the extent and scope of the attack is understood, the next step is to contain the attack before it penetrates further into your organization. This phase is critical and must not be overlooked – less mature organizations may start fixing the vulnerability that was exploited before they stop the attack. Your Playbook should include containment procedures for each of the plays that are addressed. In some cases, this may require shutting down information systems to prevent them from being compromised if they are vulnerable to the attack.

An example of a containment exercise in the phishing process may be as simple as the following:

- Shun any remote access C2 traffic that might be established as part of the attack.
- Change credentials proactively for users who clicked to open a credential theft phishing campaign.

Eradication – This phase of your response focuses your IR effort on eliminating all traces of the attack including the attack foothold. This includes:

- Identification of all emails that were delivered to your user base.
- Removal of these emails from mailboxes of the same user base.
- Reimaging of endpoints where malicious binaries or malware were downloaded to ensure no foothold exists.

Recovery – After the threat is neutralized and all malicious activity is removed from the organization's systems, you must determine whether or not to reactivate the compromised technology. In most cases, the answer to this question will be "of course," since these technologies fulfill a larger purpose in your organization. In cases where legacy technologies were compromised, however, it might not be worth the effort and investment to bring them back online.

In either case, the process to restore the technical capability in the organization is equally important as the process to remove the threats and malicious activities in your systems. As you restore functionality, shut down the vectors that made the attack successful. This may be done by patching an exploited vulnerability or rebuilding an entire system to leverage hardening processes such as those identified in *Cybersecurity Best Practice #2: Endpoint Protection Systems*.

Lessons Learned/After Action Report – Arguably the most important stage of your IR process is a full debrief with your IR teams after the attack is mitigated and systems are returned to full functionality. This debrief should profile the successful attack vectors and identify short term adjustments to introduce enhanced prevention, detection or response capabilities as well as long term strategic elements that require more detailed planning.

For example, if your organization falls prey to a sophisticated phishing attack that results in the theft of multiple credentials followed by the installation of remote access tools and elimination of an advanced persistent threat in your systems, a multifaceted set of mechanisms may be considered for short term and long term improvement. Examples may include:

- Refine a particular play within the Playbook that didn't execute as efficiently as possible. Timeliness is one of the most critical aspects of any response – taking too long to ramp up your IR Playbook increases your exposure to a successful attack.
- Refine and expand logging capabilities to detect threats more quickly. Implementing these capabilities into your SIEMs. Delve into the specific patterns of the attack as much as possible for lessons learned.
- Share attack details and information with participating ISACs and ISAOs. This helps other organizations to prevent validated and vetted threats. It provides greater credence to the intelligence and increases resiliency of the sector as a whole.
- Leverage advanced analytics-based phishing protection tools such as "click protection" or "attachment sandboxing." This usually requires investment and budget allocation by the organization.
- Refocus and prioritize resources to build out greater capabilities to identify and respond to phishing attacks. From a strategic perspective, it's important to refocus your resources in response to a threat that is ramping up against your organization.

A feedback loop from your IR processes back into engineering and operations is paramount to become a resilient organization. This type of feedback loop enhances an organization's cybersecurity capabilities overtime and organically while increasing flexibility and agility in IR response processes.

To read an example case of a "mock attack," consider "A Practical Example of Incident Response to a Network Based Attack" from the SANS Reading Room (Fraser 2017)

Further details associated with an IR Playbook can be found in the SANS Reading Room, "Incident Handler's Handbook" (Kral 2011).

C. Information Sharing and ISACs/ISAOs (NIST Framework: DETECT)

Security engineering and operations activities tend to be centered towards preventing cyberattacks and building out systems that enable streamlined execution of IR functions. That said, not all attacks are equal. They range from simple "script kiddies" that attempt to gain entry to any target to advanced persistent threats backed by substantial resources and a strong desire to gain entry to your organization. The means to differentiate these types of attacks falls under the discipline of Threat Intelligence.

The next level of sophistication is realized through involvement with and participation in Information Sharing and Analysis Centers (ISAC) and Information Sharing and Analysis Organizations (ISAO). There are many ISACs and ISAOs and they tend to focus on a specific vertical (such as NH-ISAC within Healthcare) or community (such as the Population Health ISAO). In all cases, these are associations whose sole function is to bring together like communities for the purposes of sharing cyber intelligence to protect their constituents. The means to share this intelligence vary in sophistication, although most mature ISACs leverage common standards and formats, such as Structure Threat Information eXpression (STIX) and Trusted Automated eXchange of Indicator Information (TAXII), as well as flash reports that profile current attacks. Participation in an ISAC or ISAO offers incredible value to an organization. It connects your cybersecurity professionals with the greater cybersecurity community. At a minimum, your organization should be a member of an ISAC or ISAO.

As with all disciplines, there are multiple levels of maturity within the Threat Intelligence discipline. The most basic sharing of threat intelligence involves consuming lists of "vetted bad IP addresses" or "feeds" from commodity sources. These sources have been well curated to identify where the loudest and most

obvious attack space resides. Organizations can use multiple means to consume these feeds, but it usually originates by subscribing to a daily download of IOCs.

Advanced Practice

A. Advanced Security Operations Centers (NIST Framework: DETECT)

In addition to the practices discussed for SOCs, an organization's move to more advanced realms of security management should include expanding its SOC to a 24x7x365 model. In this model, the SOC is staffed and monitored 24 hours a day, 7 days a week, 365 days a year.

There are multiple methods to achieve this model, all of which include benefits and constraints. Some of these methods are described below:

Fully Outsourced – In the fully outsourced model, all SOC and threat actions are sourced to a third party provide who has the required infrastructure, staff, and capabilities. These models generally tend to leverage sensors provided by the third party that are installed on your networks and used to collect necessary log information that enriches detection and response activities. In these models, SOC analysts actively look for threats and provide your internal IR personnel with specific actions to take once the threats is identified.

This model has the advantage of scale and capability. It is difficult to hire and retain qualified security analysts to provide this dedicated function. Additionally, organizations benefit from the shared intelligence discovered by other clients of the service provider. The main disadvantage is that these analysts tend not to fully complete response actions, requiring engagement from your internal teams. Additionally, investments made by your organization in cybersecurity tools might not be fully leveraged as the service providers are likely to use their own tools.

Fully Insourced – In the fully insourced model, all SOC and threat actions are handled with internal staff and infrastructure. This model requires the build out of a dedicated physical space with the IT infrastructure and tools necessary to support your IR personnel. It requires a combination of skills from security engineers, incident handlers, and threat hunters.

This model has the advantage of situational awareness and an in-depth understanding of the organization's business requirements and nuances. Internal staff are accustomed to the specific needs of the organization. Additionally, internal staff understand the context of an organization's various systems far more intrinsically than an outside service provider. The main disadvantages of this model relate to cost, workforce retention, and threat intelligence. Building out an internal SOC can be a costly proposition if the organization doesn't have existing facilities to support it. Moving to a 24x7 operation requires hiring new employees and supervisors to ensure effective management and coverage during holidays and time off. Lastly, in this model, the organization does not necessarily get current information about threat actions occurring in other organizations.

Hybrid – In the hybrid model, the SOC and incident handling functionalities attempt to leverage strengths of the Fully Outsourced and the Fully Insourced models while minimizing the disadvantages. In this model, organizations contract with a service provider to provides 24x7x365 monitoring and response by remotely accessing the organization's existing security technologies (e.g., SIEMs, IPS, firewalls). The service provider provides facilities and staff for monitoring and response actions and the organization provides the tools and escalation processes.

This model tends to offer flexibility and scaling of existing investments made in cybersecurity technologies, processes, and people. However, it requires a specific and scripted procedural playbooks

to be effective. The organization is required to drive these procedural playbooks and ensure the service provider complies with them. Lastly, in this model, organizations lose some of the situational awareness normally provided by internal handlers. Generally speaking, precise roles and responsibilities must be established to achieve the desired outcome.

B. Advanced Information Sharing (NIST Framework: DETECT)

Leveraging threat intelligence can be challenging. The organization must establish a threat model, ingest data according to the model, and automate the collection and response. Generally speaking, this requires dedicated human and technology resources to be successful.

MITRE has developed a model to manage these threats. "Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK™) is a curated knowledge base and model for cyber adversary behavior. It addresses the phases of an adversary's lifecycle and the platforms that are targeted. ATT&CK is useful for understanding security risks from known adversary behavior, planning security improvements, and verifying defenses work as expected (Adversarial Tactics, Techniques & Common Knowledge n.d.). It is recommended that organizations consider this model in addition to STIX and TAXII automation methods to build out a robust threat intelligence program.

Beyond ISACs and ISAOs, there are individual intelligence gathering organizations or departments within organizations that have a vested interest in getting "deep intelligence" directly from the attacker community. This capability requires substantial investments and specialized talent (think intelligence officers). Generally speaking, this level of maturity is not achievable in most large organizations. However, with proper vetting, the fruits of their labor can assist the HPH Sector immensely.

C. Incident Response Orchestration (NIST Framework: RESPOND)

It can be incredibly difficult to manage a response and leverage intelligence from the many specialized tools that may exist to provide cybersecurity in an organization. Examples of these tools include SIEMs, User Behavior analytics, deception technologies, email protection platforms, and Endpoint Detection and Response technologies. Though tools like SIEMs are designed to ingest information from multiple sources and provide context, this capability is dependent on the extensibility of log data generated by these systems as well as the workflow and process capabilities of the SIEM technology. Generally speaking, SIEMs are fantastic at developing alerts and notifying security resources of emergent issues, but, generally not as robust in the process of executing IR playbooks.

This is where IR Orchestration tools come in handy. Once playbooks have been scripted and vetted, these tools ensure that playbook execution is consistent throughout the process. Without these types of tools, IR consistency must be managed by cybersecurity personnel. These tools enable cybersecurity personnel to focus on the incident rather than the consistent execution and documentation of an incident.

In addition to the workflow components of IR Orchestration, these tools can pull data from system security stacks and present it to the incident responder in a centralized dashboard. Examples of data that may be pulled into this dashboard include: SIEM, Log Data, Dynamic Host Configuration Protocol (DHCP) logs, asset inventories, anti-malware consoles, vulnerability management data, threat intelligence information, identity management systems, and endpoint security technologies. Each of these systems provides a perspective on the particular type of threat that your organization is experiencing.

D. Baseline Network Traffic (NIST Framework: DETECT)

A useful technique to deploy is to baseline your network traffic and implement capabilities to alert upon anomalous changes to the baseline. This can be accomplished by leveraging netflow data and systems that can ingest netflow data. Each system that operates in the ecosystem will have a standard 'digital footprint' for its network communications, and generally will operate within those parameters. By conducting a baseline operation on each of the major and core systems you can compare what is 'expected' to what is actually occurring. This process can be done manually or you can invest in technologies that can automate this process for you.

E. User Behavior Analytics (NIST Framework: DETECT)

User Behavior Analytics (UBA) is a technique that may be considered as the "SIEM for Users." In most modern threats, the threat actors attempt to leverage access that already exists in the user space. Although attackers can generate new accounts for access attempts, they are well aware that most organizations monitor systems for new accounts, especially those with privileged access. The exploitation of existing accounts, however, might go unnoticed.

UBA systems provide analytics context from a user perspective. Similar to conducting a baseline activity over network access, UBAs baseline user activity and actions throughout the organization's digital ecosystem. The tool ingests the most relevant user activity logs from these systems as well as existing authentication and authorization systems. Deviations are discovered after the user has been profiled, enabling IR actions to be executed according to the proper playbooks.

One note of importance: UBA protects against external as well as internal threat actors.

F. Deception Technologies (NIST Framework: DETECT)

Deception technologies expand on the honeypot and honeynet techniques of old, delivering it at scale to the enterprise. These techniques place "fake systems" or "fake breadcrumbs" throughout the digital ecosystem and wait for them to be "tripped." They work on the principle that communications should not occur in a system that serves no purpose in the organization. If such a communication occurs, it should be brought to the attention of the IR teams for further investigation.

Deception technologies discover attackers who have placed a foothold on your organization's network and are attempting to pivot to find targets of interest. These targets may be simple (e.g., file storage systems, email systems) or they may be complicated (e.g., EMR or Imaging systems). In all cases, the goal of the attacker is to leverage access already obtained to pilfer data or conduct an extortion attack (i.e., ransomware). The attacker's approach is to generate hundreds or thousands of these "fake systems" so that it is difficult to differentiate them from real assets.

Your IR teams can profile the threat actor by watching their behavior on these fake systems. For example, technologies exist to create a complete fake file system that interacts and responds like a real file storage system, even generating files that appear legitimate. By watching the threat actor enumerate the file system, your IR teams can develop a high level of certainty of the malicious intent and identify the foothold held by the threat actor on the organization's network.

Threats Mitigated

1. Phishing Attacks
2. Ransomware Attacks
3. Loss or Theft of Equipment

4. Internal, Accidental or Intentional Data Loss
5. Attacks against Connected Medical Devices that May Affect Patient Safety

Suggested Metrics

- Time to detect and respond in aggregate, trended by week. The goal is that an IR response should kick off within X hours after detection of an incident and the incident should be mitigated within X hours after response. Lag time between occurrence and detection of a security incident should be fewer than X days.
- Number of true positive incidents executed by incident category on a weekly basis. Though there is no goal set for this metric, it's important to monitor trends in incidents that occur in your organization. This will inform the larger security strategy over time based on actual threats in your organization.
- Number of backup failures by week. The goal is to minimize the number of backup jobs that fail and to provide continual assurance that backup jobs are executing as intended.
- Number of notable (or critical/high rated) security incidents per week, providing a profiled enumeration of each incident. Each notable security incident should be executed consistently and thoroughly. Each incident should have an After Action Report. The goal is to demonstrate that After Action Reports and incident reports are written for each notable security incident. This will help with the development and implementation of continual improvement processes.

References

- (Fraser 2017)
- (Kral 2011)
- (Swift 2010)
- (The 6 Categories of Critical Log Information 2013)
- (SP 800-61r2: Computer Security Incident Handling Guide 2012)
- (Adversarial Tactics, Techniques & Common Knowledge n.d.)

Cybersecurity Best Practice #9: Medical Device Security

Best Practice 9: Medical Device Security	
Data that may be affected	ePHI
Baseline Practices	A. Medical Device Management B. Endpoint Protections C. Identity and Access Management D. Asset Management E. Network Management
Advanced Practices	A. Vulnerability Management B. Security Operations and Incident Response C. Procurement and Security Evaluations D. Contacting the FDA
Key Mitigated Risks	• Attacks Against Connected Medical Devices and Patient Safety

Healthcare systems leverage many diagnostic and therapeutic methods for patient treatments. These may be technological systems that render and provide detailed images of CT scans, or they may be devices that connect directly to the patient for a diagnostic or therapeutic purpose. These types of devices may have straightforward implementations, such as bedside monitors that monitor the vital signs of a patient during an inpatient stay, or it may have a complicated implementation such as an infusion pump that delivers specialized therapies, such as chemotherapy to a cancer patient, and requires continual updates to its drug libraries. These complex and interconnected devices impact patient safety and well-being, and should be robustly designed and properly secured.

This section focuses entirely upon the methods that Health Delivery Organizations (HDO) can employ to protect connected medical devices. Specifically, it will focus on the actions that HDOs are permitted to take and how to best work with device manufacturers and the U.S. Food and Drug Administration (FDA).

Any device that connects directly to the patient for the purpose of diagnostic or therapies must undergo extensive quality control processes to ensure they are safe for use. Rigorous stipulations are in place for the development and release of such systems. These stipulations are managed by the FDA. The organizations that produce these devices must comply with these regulations, and generally are referred to as device manufacturers. Organizations that purchase these devices and use them for the treatment of patients are the clinical providers. In the context of this relationship, they are referred to as HDOs.

Given the highly regulated nature of these devices and the specialized skills required to make modifications, it is ill-advised for HDOs to make configuration changes without the support of the device manufacturer. Doing so may put the HDO at risk for voiding warranties, result in legal liabilities, and, at worst, harm the patient. Traditional methods used by security programs to secure these assets cannot necessarily be deployed. For example, one cannot simply apply a patch to a vulnerable component of the OS that runs a medical device.

For a practical example of full life cycle management, risk analysis, management, best practices, and detailed configuration specifications to secure wireless pumps (one type of medical devices), consider the information released by NIST and the National Cybersecurity Center of Excellence (NCCOE), *NIST SP 1800-8: Securing Wireless Infusion Pumps In Healthcare Delivery Organizations 2017.*

Baseline Practice

A. Medical Device Management (NIST Framework: IDENTIFY)

Medical devices are a specialized type of "Internet of Things" (IOT), leveraged for clinical diagnostic or treatment purposes within HDOs.

That said, cybersecurity for medical devices follows many of the best practices already discussed in this document, including:

- Cybersecurity Best Practice #2: Endpoint Protections
- Cybersecurity Best Practice #3: Identity and Access Management
- Cybersecurity Best Practice #5: Asset Management
- Cybersecurity Best Practice #6: Network Management
- Cybersecurity Best Practice #7: Vulnerability Management
- Cybersecurity Best Practice #8: Security Operations and Incident Response

Rather than recreating these best practices, HDOs are encouraged to extend the relevant best practice from each section, implementing it appropriately for medical device management.

B. Endpoint Protections (NIST Framework: PROTECT)

As much as feasible, medical devices should have following controls enabled:

Antivirus software – Usually the medical device manufacturer must directly support this software or it must be cleared for operation by the manufacturer. Ensure that a compliant AV technology is enabled. If AV cannot be implemented, compensating controls should enforce an AV scan whenever the device is serviced prior to reconnecting to the device network.

Local firewalls enabled – Medical devices should be configured to communicate only with required systems. Unused services and ports should be disabled as long as they are supported by the manufacturer.

Encryption – If supported by the manufacturer, medical devices should have local encryption enabled in the case the device is stolen.

Default Password Changes – If supported by the manufacturer, default passwords, especially those enabling privileged access, should be changed to a credential used only for the medical device. Do not tie this unique credential to any general systems management credential as you do not want any compromise of those credentials to impact the medical device.

Routine Patching – As part of preventative maintenance cycles, medical devices should be patched with supported cybersecurity patches that are released by the device manufacturers. Given the special sensitivity to the configuration and management of these devices, it is emphasized that patching must not take place on these devices unless cleared by the manufacturer.

C. Identity and Access Management (NIST Framework: PROTECT)

As much as feasible, medical devices should have following controls enabled on them:

Authentication – If supported by the manufacturer, the device should bind its authentication capabilities with systems enterprise authentication domains. This automates termination of access to the device upon termination of employment for the user.

Vendor Support Passwords – Passwords should not be shared among the vendor team. A unique logon credential should be established for each vendor employee. Ensure the manufacturer does not use the same account and password to manage medical devices in your organization and others.

Remote Access – If remote access is required to manage medical devices, MFA capabilities should be deployed. Depending on the deployment scenario, the device manufacturer may be required to support these capabilities. Otherwise, these capabilities should be deployed on a separate component of your existing MFA system to limit exposure in the event that the MFA system is compromised.

D. Asset Management (NIST Framework: IDENTIFY)

As much as feasible, medical devices should have following controls enabled on them:

Inventory – All medical devices should be added to an inventory that is capable of understanding the core components of the medical device itself. This may be your general IT Asset Management inventory as described in *Cybersecurity Best Practice #5: IT Asset Management.* You may have to leverage specialized tools designed specifically for tracking the lifecycle of medical devices. Such systems can be useful for maintaining preventative maintenance schedules.

Wiping – When a medical device is slated for decommissioning, it is critical to ensure that all data resident on the device is wiped. These devices typically are returned to the vendor and potentially resold or delivered to other organization for destruction. You do not want your organization's data to be accessed by these other parties.

E. Network Management (NIST Framework: PROTECT)

As much as feasible, medical devices should have following controls enabled on them:

Segmentation – Given the critical nature of these devices and the general inability to configure them to reduce vulnerabilities, it is critically important to segment these devices separately from general access or data center networks. The ability to restrict access to the device is essential to its safe operation.

Dedicated networks should be set up that are highly restricted. The only traffic allowed on these networks should be profiled based on required operation of the devices connected to that network. Access to the device management systems should be heavily restricted to limit exposure of the management system to being compromised. Lastly, it's important to ensure that these networks are segmented such that any vulnerability scanning systems are <u>not</u> permitted access in a clinical setting. Given the delicate nature of medical devices, execution of a rogue vulnerability scan could disrupt the devices.

As part of the segmentation strategy, review data flows and interfaces between the medical devices and their connected systems. Be sure not to limit the essential functionality of the medical device including the ability for it to be patched remotely if that is required.

Device manufacturers may require installation of their own physical networks in the organization. In these cases, access to the manufacturer's physical network should be limited with the same restrictions as if the HDO were implementing its own segmentation strategy.

Advanced Practice

A. Vulnerability Management (NIST Framework: DETECT)

As much as feasible, medical devices should have following controls enabled on them:

Vulnerability and Risk Categorization – In 2016, the FDA issued the Postmarket Management of Cybersecurity in Medical Devices guidance. (Postmarket Management of Cybersecurity in Medical Devices 2016) This guidance document presents components for the proper management of medical devices once they have been deployed in an HDO. Focusing on the risk to patient safety, this guidance stipulates that manufacturers should implement vulnerability and risk management practices to categorize risks according to the device effectiveness and the potential to cause harm to the patient.

HDOs should work with device manufacturers to enable a common understanding of the framework for the risk categorizations. Upon disclosure of a high risk, HDOs should take an escalated action to secure the device.

Vulnerability Disclosure Programs – Each device manufacturer should have a program that informs HDOs of vulnerabilities that exist in their devices. These programs should have a communication channel to report information and inform parties. HDOs should work with the manufacturers so all parties understand the respective points of contacts in the manufacturer and the HDO.

In addition to these communication channels, other channels exist for the disclosure of medical device vulnerabilities. These include the Industrial Control Systems – Computer Emergency Response Team (ICS-CERT), which the manufacturers can included as part of their vulnerability releases, or the respective ISACs or ISAOs with which the manufacturers participate.

The HDO must have a program in place to accept inbound vulnerability disclosures, evaluate the HDO's exposure to these vulnerabilities, and stand up response actions with the manufacturers to remediate or mitigate each vulnerability according to its level of risk.

Table 13 provides a general rule of thumb for the response (including interim compensating controls) timeframes of medical device vulnerabilities that is in line with expectations in the Postmarket Management of Cybersecurity for Medical Devices guidance.

Vulnerability Criticality	Days
Uncontrolled Risk	
Communicate to Customer	< 30 days
Remediate Risk	< 60 days
Controlled Risk	As defined by routine patching and preventative maintenance

Table 13. Response Timeframes to Resolve Medical Device Vulnerabilities

Software Bill of Materials (SBOM) and Vulnerability Lookups – By leveraging the SBOM registered in the organization's ITAM, the HDO can compare data from the National Vulnerability Database (NVD) against data in the organization's software libraries. This comparison provides the HDO with information on current potential vulnerability postures in the medical device space.

A simple search of the NVD can be conducted by using the web interface located at https://nvd.nist.gov/vuln/search. This search allows HDOs to look up vulnerabilities based upon products that they currently have. It does not require SBOM material to be pre-registered.

Vulnerability Scanning

WARNING: THIS ACTION MUST BE TAKEN WITH EXTREME CAUTION DUE TO THE POTENTIAL IMPACTS ON MEDICAL DEVICES WITHIN THE PRODUCTION ENVIRONMENT. HDOS SHOULD NOT ATTEMPT TO CONDUCT THIS ACTION UNLESS THEY ARE ABSOLUTELY CERTAIN THE MEDICAL DEVICE IS NOT IN PRODUCTION, IS NOT CURRENTLY IMPLEMENTED IN A CLINICAL SETTING, AND IS NOT CONNECTED TO PATIENT.

The final action that an HDO can take to understand its vulnerability posture is to conduct vulnerability scans against the medical devices. There are two opportune times to conduct vulnerability scans against medical devices:

- When the device is first procured and tested before deployment in the production environment; and,
- When a device is taken offline for preventative maintenance and routine patching.

In both scenarios, it is important for the device to be in a highly controlled setting while not connected to a patient. A vulnerability scan can be configured to profile the device and determine if potential vulnerabilities exist, or confirm that vulnerabilities have been mitigated as part of a remediation or patching plan.

To conduct such an exercise, it is best for the cybersecurity team to work with the clinical engineering teams and establish a profiled scan template in the vulnerability management software. This template should allow the scan to be executed only against a specific non-production network and only by specific individuals. To provide further assurance that the vulnerability scan cannot cause harm to the medical devices, IP addresses of the scanners should be blocked as part of the segmentation strategy noted above.

When these preparations are complete, the clinical engineering teams can be granted access to the scanning software in a restricted manner that allows the scan to be run only against the network used for preventative maintenance. Vulnerabilities discovered can be shared with the information security office to determine the relative risks. Upon classification of these risks, the teams should contact the device manufacturer and work together to develop and implement a remediation plan.

B. Security Operations and Incident Response (NIST Framework: PROTECT, RESPOND, RECOVER)

Expanding on the SOC and IR processes found in *Cybersecurity Best Practice #8: Security Operation Center and Incident Response,* HDOs can provide better monitoring, detection and response activities around their medical device ecosystem. To provide visibility into the daily operations of the medical device systems, the following sources should be configured to send logs to the HDO's log management systems, SIEM, or both:

- Firewalls providing segmentation to the medical device network segment;
- Information systems that control the operation of the medical devices;
- Netflow data from the medical device network segment;

- Intrusion prevention systems in front of the medical device network segment; and,
- Logs from any deception technology deployed in the medical device network segment.

With the source of these logs, plays can be enumerated and added into the standard Incident Response playbooks, as described in Table 14.

Play Category	Play	Description	Source Data
Reconnaissance	Vulnerability scanning sweep of Medical Device Segment.	Large number of vulnerabilities being scanned across the medical device network. May be caused by a single server being scanned over multiple ports, or multiple servers being scanned over a single port.	• Medical device management system. • IDS/IPS logs in front of the medical device network, configured to detect vulnerability scanning. • Firewall logs in front of the medical device network. • Netflow data from within the medical device network.
Lateral Movement	Detection of unknown source clients accessing medical device remote access ports.	Detect attacks coming from sources outside of known management sources attempting to gain access to remote access ports (e.g., ftp, ssh).	• Firewall logs in front of the medical device network. • Network data from within medical device network.
Lateral Movement	Triggered decoy within medical device network.	Respond to triggers of decoys being communicated from within or across the medical device network segment. These communications should not occur and indicate malicious or broken processes.	• Deception technology logs from within the medical device network. • Firewall logs in front of the medical device network. • Network data from within the medical device network.

Table 14. IR Plays to Combat Cyberattacks Against Medical Devices

C. Procurement and Security Evaluations (NIST Framework: IDENTIFY)

First and foremost, HDOs must establish a set of cybersecurity requirements during the acquisition of new medical devices. These requirements should be embedded in your contracting processes and

shared with your supply chain and procurement offices. Ideally, these cybersecurity requirements are included in a Request for Information (RFI) or Request for Proposal (RFP) process.

Secondly, technology acquisitions including medical devices require a security evaluation as part of the supply chain process. Though important for the acquisition of medical devices, this is true for any integration of technology into the HDO. Allowing cybersecurity evaluation processes to occur as part of the supply chain process provides an opportunity for cyber risks to be understood, evaluated, and mitigated prior to deployment.

The first process that should be kicked off as part of the medical device acquisition process is a security evaluation of the devices. This evaluation should uncover any risks or flaws in the current design of the medical device and allow transparent communications between the supply chain process, clinical engineering and the manufacturers. To shepherd the process along, the HDO should insist on receiving a Manufacturer Disclosure Statement for Medical Device Security (MDS2). The MDS2 is a standardized form leveraged by most manufacturers and developed by HIMSS and the American College of Clinical Engineering (ACCE). It provides a list of comprehensive cybersecurity questions for medical devices with responses from the manufacturer of the device in question. Questions in the MDS2 include:

- "Can this device display, transmit, or maintain private data (including electronic Protected Health Information)?"
- "Can the medical device create an audit trail?"
- "Can users be assigned different privileged levels within an application based on 'roles' (e.g. guests, regular users, power users, administrators)"; and,
- "Can the device owner/operator reconfigure product security capabilities?"

A copy of the latest MDS2 can be found on the Association of Electrical Equipment and Medical Imaging Manufacturers (NEMA) website (The Association of Electrical Equipment and Medical Imaging Manufacturers 2013). Answers to these questions assist the HDO to complete a meaningful evaluation of the medical device.

After the security evaluation is complete, the cybersecurity program should review and provide input and review into the contract with the manufacturer. This stage should occur in tandem with the supply chain and legal negotiations and leverage a template of security terms of interest from the HDO. These terms should reference the FDA Postmarket Management of Cybersecurity for Medical Devices guidance document, referencing components of the guidance that are critical for the safe operation of the devices.

Armed with the results of the cybersecurity evaluation, scenarios to resolve any unmitigated risks should be included in the contracting process to limit the HDO's liability, especially with constraints around the HDO's ability to alter the medical devices.

*Software Bill of Material*s – The HDO should request a SBOM as part of the procurement process. The SBOM is a list of software components that comprise the medical device. Think of it as the software library, or the ingredients of the recipe, that make up the device. Understanding the software libraries that make up the device facilitates the HDO's ability to understand the impact of vulnerabilities announced by the NVD.

D. Contacting the FDA (NIST Framework: DETECT)

If HDO is stuck managing high risk cybersecurity vulnerability and cannot get support from the medical device manufacturer to mitigate this risk, the HDO's final recourse is to contact the FDA directly to express concern about the vulnerability. Contacts to the FDA should be limited to critical or high risk scenarios, especially those with the potential to cause harm to patients.

The CDRH emergency contact information is provided below:

- Email: AskCyberMed@fda.hhs.gov
- Phone: 301-796-7436

Threats Mitigated

1. Attacks Against Connected Medical Devices and Patient Safety

Suggested Metrics

- Number of medical devices not currently segmented on wireless or wired networks, trended over time. The goal is to limit medical devices on the general access network, data center network, or other location that does not meet the requirements of specific network segmentation strategies.
- Number of unmitigated, high risk vulnerabilities on connected medical devices, trended over time. The goal is to reduce the number of unmitigated risks as close to zero as possible. Each high risk vulnerability should have a remediation action plan, as defined in *Cybersecurity Best Practice #7: Vulnerability Management*.
- Number of medical devices procured that did not receive security evaluation, trended over time. The goal is to reduce the number of procurement actions without security evaluation to as near as zero as possible. Leverage this metric to work with your supply chain and clinical engineering departments to ensure the process is executing as intended.
- Number of medical devices that do not conform to basic endpoint protection best practices, trended over time. The goal is to reduce the number of medical devices that do not meet the basic hygiene management practices or to implement practices for these devices. It is not always possible to reduce this number to zero. Mitigating factors should be leveraged to keep it as low as possible.
- Number of devices that have unknown risks due to lack of manufacturer disclosure information, trended over time. The goal is to ensure device manufacturers have vulnerability disclosure programs and that your organization is tied into these processes.

References

- (Postmarket Management of Cybersecurity in Medical Devices 2016)
- (NIST SP 1800-8: Securing Wireless Infusion Pumps In Healthcare Delivery Organizations 2017)

Cybersecurity Best Practice #10: Cybersecurity Policies

Best Practice 10: Cybersecurity Policies	
Data that may be affected	-
Baseline Practices	A. Policies
Advanced Practices	-
Key Mitigated Risks	• Email Phishing Attacks • Ransomware Attacks • Loss or Theft of Equipment or Data with Sensitive Information • Accidental or Intentional Data Loss • Attacks Against Connected Medical Devices and Patient Safety

Cybersecurity policies must be established for the workforce to understand their expected behaviors within the organization as they relate to cybersecurity. These policies should be written for the various user audiences that exist in the organization. There are differences between the general workforce user, IT user, and high profile/risk users (e.g., Finance, HR, or Health Information Management).

New cybersecurity hygiene controls should be supported by institutional policy to set the proper expectations. Without these policies, it may be unclear to the workforce what level of adherence is required and what activities put the organization at risk for the threat types discussed in this document.

A number of policy templates have been provided in the Appendices.

A. Policies (NIST Framework: IDENTIFY)

There is only one general safeguard for this section and that is a listing of policies that organizations can consider, described in Table 15.

Policy Name	Description	User Base
Roles and Responsibilities	Define all cybersecurity roles and responsibilities throughout the organization. This includes who will establish policy as well as implement and conduct security practices.	All users.
Education and Awareness	Define the mechanisms that will be used to train the workforce on cybersecurity practices, threats and mitigations.	All users. Cybersecurity Department.
Acceptable Use / Email Use	Describe actions that users are permitted and not permitted to take. Explicitly define how email is to be used and leveraged.	All users.
Data Classification	Define how data are to be classified with usage parameters around those classifications.	All users.

Personal Devices	Define the organization's position on the usage of personal devices (i.e., BYOD). If these are permitted, establish expectations for how the devices will be managed.	All users.
Laptop, Portable Device, and Remote Use	Define policies for the security of mobile devices and how they are to be used in a remote setting.	All users. IT Department.
Incident Reporting and Checklist	Define user requirements to report suspicious activities within the organization. Define responsibilities of the cybersecurity department for managing incidents.	All Users. Cybersecurity Department.
Disaster Recovery Plan	Define the practices deployed to recover IT assets in the case of a disaster, including backup plans.	IT Department.
IT Controls Policies	Describe the requirements for IT security controls in a series of policies or a single long policy. Examples include Access Control, Identity Management, Configuration Management, Vulnerability Management, and Data Center management.	IT Department.
IT Acquisition Policy	Define the criteria that must be implemented to ensure proper identification and protection of all IT assets purchased by the organization.	Supply Chain / Procurement Users. IT Department.

Table 15. Cybersecurity Policies to Be Considered

Threats Mitigated

1. Email Phishing Attacks
2. Ransomware Attacks
3. Loss or Theft of Equipment or Data
4. Internal, Accidental or Intentional Data Loss
5. Attacks Against Connected Medical Devices that can Affect Patient Safety

Suggested Metrics

- Number of policies reviewed over a specified timeframe. The goal is to establish a standard practice to review policies and monitor compliance with this standard.
- Number of workforce members review and sign off after reading policies over a specified timeframe. The goal is to establish a standard practice for workforce members

to review applicable policies and attest that this review has been completed, and for the organization to monitor compliance with this standard.

2519 # Appendix A: Acronyms and Abbreviations

2520

Acronym/Abbreviation	Definition
ABAC	Attribute Based Access Control
ACCE	American College of Clinical Engineering
ADFS	Active Directory Federation Services
AHIP	America's Health Insurance Plans
API	Application Programming Interface
ASL	Assistant Secretary for Legislation
ASPR	Assistant Secretary for Preparedness and Response
ATT&CK	Adversarial Tactics, Techniques & Common Knowledge
AV	Anti-Virus
BYOD	Bring Your Own Device
C2	Command and Control
CASB	Cloud Access Security Broker
CEO	Chief Executive Officer
CHIO	Chief Health Information Officer
CHIP	Children's Health Insurance Program
CIO	Chief Information Officer
CIRT	Cybersecurity Incident Response Team
CISO	Chief Information Security Officer
CISSP	Certified Information Security Systems Professional
CMS	Centers for Medicare and Medicaid
CNSSI	Committee on National Security Systems Instruction

COO	Chief Operations Officer
CSA	Cybersecurity Act
CVSS	Common Vulnerability Scoring System
DCC	Distributed Checksum Clearinghouse
DEP	Device Enrollment Program
DHCP	Dynamic Host Configuration Protocol
DHS	Department of Homeland Security
DKIM	Domain Key Identified Mail
DLP	Data Loss Prevention
DMARC	Domain-based Message Authentication Reporting and Conformance
DMZ	Demilitarized Zone
DNS	Domain Name System
DNSRBL	Domain Name System Real-time Blackhole List
DoD	Department of Defense
DOS	Denial of Service
DRP	Disaster Recovery Plan
DSM	Direct Secure Messaging
EA	Enterprise Architecture
EDR	Endpoint Detection and Response
EHR	Electronic Health Record
EMR	Electronic Medical Record
EPHI	Electronic Private Health Information
ERP	Enterprise Resource Planning
FDA	Food and Drug Administration

FERPA	Family Educational Rights and Privacy Act
FIPS	Federal Information Processing Standards
FISMA	Federal Information Security Modernization Act
GDPR	General Data Protection Regulation
GINA	Genetic Information Nondiscrimination Act
HCIC	Health Care Industry Cybersecurity
HDO	Health Delivery Organization
HHS	Department of Health and Human Services
HIDS	Host Based Intrusion Detection Systems
HIMSS	Health Information Management and Systems Society
HIPAA	Health Insurance Portability and Accountability Act
HIPS	Host Based Prevention Systems
HIT	Health Information Technology
HITECH	Health Information Technology Economic and Clinical Health Act
HMO	Health Maintenance Organization
HPH	Healthcare and Public Health
HR	Human Resources
HRSA	Health Resources and Services Administration
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IA	Information Assurance
IAM	Identity and Access Management
IBM	International Business Machines
ICAP	Internet Content Adaptation Protocol

ICMP	Internet Control Message Protocol
ICS-CERT	Industrial Control Systems – Computer Emergency Response Teams
ICU	Intensive Care Unit
IDS	Intrusion Detection System
INFOSEC	Information Security
IOC	Indicator of Compromise
IoT	Internet of Things
IP	Intellectual Property or Internet Protocol
IPS	Intrusion Prevention System or Internet Partner Services
ISAC	Information Sharing and Analysis Center
ISAO	Information Sharing and Analysis Organization
IT	Information Technology
ITAM	Information Technology Asset Management
LAN	Local Area Network
LANMAN	Local Area Network Manager
LDAP	Lightweight Directory Access Protocol
LLC	Limited Liability Corporation
MAC	Media Access Control
MACRA	Medicare access and the Children's Health Insurance Program Reauthorization Act
MARS	Minimum Acceptable Risk Standards
MDM	Mobile Device Management
MDS2	Manufacturer Disclosure Statement for Medical Device Security

MFA	Multi-Factor Authentication
MITRE	The MITRE Corporation
MPLS	Multiprotocol Label Switching
NAC	Network Access Control
NCCIC	National Cybersecurity and Communications Integration Center
NCI	National Cancer Institute
NEMA	Association of Electrical Equipment and Medical Imaging Manufacturers
NH-ISAC	National Healthcare – Information Sharing and Analysis Centers
NIH	National Institutes of Health
NIST	National Institute of Standards and Technology
NNCOE	NIST National Cybersecurity Center of Excellence
NVD	National Vulnerability Database
OCIO	Office of the Chief Information Officer
OCR	Office for Civil Rights
ONC	Office of the National Coordinator (for Healthcare Technology)
OS	Operating System
OWASP	Open Web Application Security Project
PACS	Pictures Archiving and Communication Systems
PCI-DSS	Payment Card Industry Data Security Standard
PCS	Patient Care Service
PHI	Personal Health Information
PII	Personal Identifiable Information

RBAC	Rule Based Access Control
RBL	Real-time Blackhole List
RDP	Remote Desktop Protocol
RFI	Request for Information
RFP	Request for Proposal
ROM	Read Only Memory
SaaS	Software as a Service
SAMHSA	Substance Abuse and Mental Health Services Administration
SAML	Security Assertion Markup Language
SBOM	Software Bill of Materials
SIEM	Security Incident and Event Management
SMB	Server Message Block
SME	Subject Matter Expert
S/MIME	Secure Multi-Purpose Internet Mail Extensions
SMS	Short Message Service
SMTP	Simple Mail Transfer Protocol
SOC/IR	Security Operations Center / Incident Response
SPAN	Switched Port Analyzer
SPF	Sender Policy Framework
SQL	Structured Query Language
SSH	Secure Shell
SSN	Social Security Number
SSO	Single Sign On
STIG	Security Technical Implementation Guide

STIX	Structure Threat Information eXpression
SVP	Senior Vice President
TAXII	Trusted Automated eXchange of Indicator Information
TLS	Transport Layer Security
TXT	Text
UBA	User Behavior Analytics
URL	Uniform Resource Locator
US-CERT	United States Computer Emergency Readiness Team
USB	Universal Serial Bus
VAR	Value Added Reseller
VP	Vice President
VPN	Virtual Private Network
WAN	Wide Area Network
WORM	Write Once Read Many

2521
2522
2523
2524
2525