Chapter 5
Institutional Cybersecurity in a Clinical Research Setting

Michal Kouril and John Zimmerly

Abstract The principal challenge facing IT groups that support research on a daily basis lies in striking a fine balance: On one hand researchers must share data and use cutting edge analytic tools on a variety of computing platforms to enhance their creativity and productivity. On the other hand, much of the data that supports translational research contains personal health information derived from patients' medical records. Hospitals are justifiably concerned about the highly sensitive nature of the data and the need to comply with a myriad of federal, state, and local laws, and contractual regulatory requirements that demand high levels of security and access control. A number of frameworks exist to help with the process. In this chapter we discuss these challenges and the approaches taken at a research intensive children's hospital to put a policy framework in place that enacts standards for security, risk evaluation and mitigation, monitoring, testing, and design of the IT infrastructure. These protect the institution, while enabling collaboration and innovation among researchers. We stress the organizational need for a close and collaborative relationship between IT groups that support research and those charged with support of the medical center's clinical and business operations. It is also important to recognize that technology alone cannot assure security. Institutional policies and user security awareness education also play key roles in assuring that confidential information is in fact protected.

Keywords Access control • Authentication • Cybersecurity • Data-centric • Data sharing • Firewall • Network • Operating systems • Protected health information • Security • Security awareness

M. Kouril, Ph.D. (✉)
Departments of Pediatrics and Biomedical Informatics, Division of Biomedical Informatics, Cincinnati Children's Hospital Medical Center, University of Cincinnati College of Medicine, 3333 Burnet Avenue, ML-7024, Cincinnati, OH 45229-3039, USA
e-mail: michal.kouril@cchmc.org

J. Zimmerly
Division of Information Services, Cincinnati Children's Hospital Medical Center, 3333 Burnet Avenue, ML-9009, Cincinnati, OH 45229-3039, USA
e-mail: john.zimmerly@cchmc.org

J.J. Hutton (ed.), *Pediatric Biomedical Informatics*, Translational Bioinformatics 10,
DOI 10.1007/978-981-10-1104-7_5

5.1 Introduction

Translational research typically requires access to data that reside in geographically distributed data warehouses that are called upon by a team of collaborators using a variety of software applications. Information technology support of this type of translational research requires networks to permit data sharing, while maintaining high levels of security. Without networking, investigators cannot access the applications, servers, and other resources in a distributed environment. Given this criticality to providing services, networking can be the major infrastructure that provides an environment where access to data is restricted to authorized users and the overall system is protected from malicious attacks by unauthorized individuals (data breaches). According to the Verizon 2011 Data Breach Report (Verizon 2011), which outlines the incidents investigated by Verizon security response teams and publicly accessible data, the top 2 of 10 threats to systems involved direct hacking against servers. Given this threat, network security and sound network security practices can provide a large layer of protection against current threats to the security of the environment.

5.2 Secure Network Design

In this section, secure network design practices along with models for architecting secure networks will be covered including data-centric network, firewalls, intrusion prevention and detection systems (IPS, IDS), and secure remote access including virtual private networking (VPN).

5.2.1 Data-Centric Networking

The traditional model for building enterprise networks takes into account the different trust zones within which applications must be accessed and published. These typically include: (1) the internal zone for trusted employees and machines on the Local Area Network (LAN), and those authorized to remotely access the network; (2) the Demilitarized zone (DMZ), a subnet that provides limited connectivity to the internal network; and (3) the external untrusted Internet zone. This model has worked well for most organizations, providing a barrier between internal and external environments, as well as segregating different populations of users. However, access requirements for most applications are evolving, as technology changes so this model may not work for all situations. Medical centers are finding that users are increasingly using mobile devices, requesting access from all locations, and requiring more help with troubleshooting, when accessing applications from different locations. Given this, the barriers between the internal and external environment are

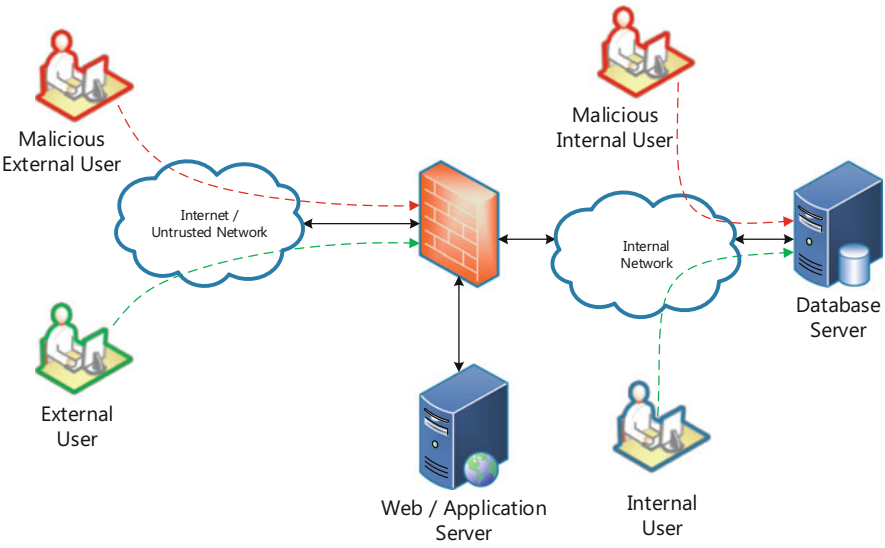

Fig. 5.1 Typical architecture of a network with different zones of trust

becoming less clear and less stringent than they were originally intended. To address this, the concept of data-centric networks has evolved. The aim is to re-architect networks around protecting data, unify the internal/external experience, and build core protections for what must be protected; the data.

Figure 5.1 outlines the setup of a typical network with the different trust zones, using an application server that has a database backend. External users and any malicious users on the Internet will be restricted from accessing the database server directly and have limited access to the Web/Application server through firewall filtering. However, any user who is on the internal network (malicious or not) will have direct access to the database server, where the application data may be stored. Given vulnerabilities or misconfigurations that may be present within the server or database application, users may be able to bypass the application protections and pull data directly from the database, whether they have been authorized or not. Given the scale and size of most enterprise networks, the malicious user in Fig. 5.1 could be accessing via remote offices, affiliates, or other locations that may not have the same level of physical security as a main campus or office may have.

To address the risks posed by the direct access to data as shown in existing network design, data-centric networking looks at building rings of protection around data. Figure 5.2 provides a high-level overview of this design philosophy. This architecture follow similar security designs used within CPU privilege rings (so-called Ring 0 access) (Cruse 2016). Starting from the center of the circle and moving outward, the rings and access become progressively less trusted. No user or process can jump a ring without going through a ring above it. This forces all access

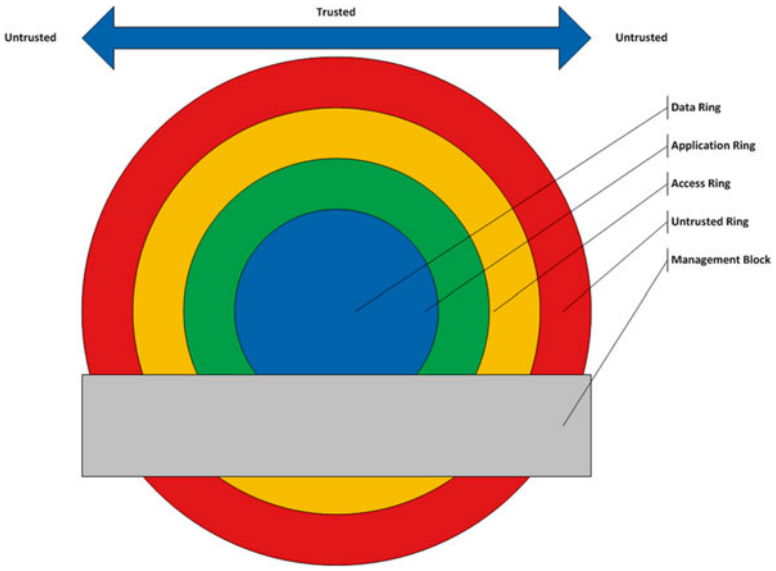

Fig. 5.2 Building rings of protection around data

to data through known applications and paths that can be hardened, monitored, and protected using internal controls. This layering approach can also help in the principal of containment of malicious activity. Rings include:

- Data Ring: This would include all structured and unstructured data sources such as database servers, file servers, etc.
- Application Ring: This would include all application servers serving content or publishing data that is accessed, used, or manipulated by users.
- Access Ring: This is the main ring that presents the initial login, authentication, and authorization to the users. This would typically include VPN, proxy servers, etc. that are used for access to the environment.
- Untrusted Ring: This is anything outside of the environment. This can include both the Internet and internal network segments.
- Management Block: This would include management systems used by system administrators to maintain the environment, such as patching, authentication, backups, etc.

Figure 5.3 shows a network with the same type of users as the network in Fig. 5.1, but has been re-architected in accord with the data-centric network model. As shown in the Figure, there is no direct access to the database server or application except from the management layers. Users can access an application only through the access server, which will present the data back to them. All access to the application is through the external firewall, to the access server, and then to the application. At these initial rings, filtering for users, locations, etc. can be put in place to further limit the scope of access by users. None of the traffic is ever sent to the application

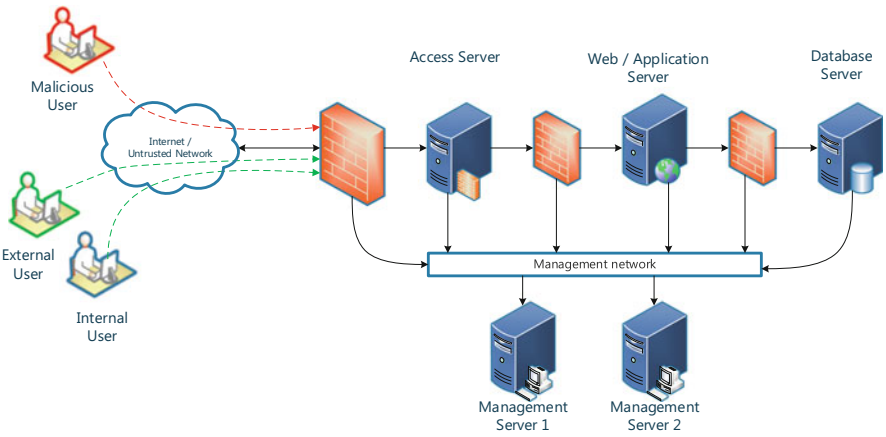

Fig. 5.3 The network shown in Figure 5.1 has been re-architected in accord with principles of the data-centric network model

or database server until it has been properly authorized to be legitimate traffic within the environment. The data-centric network model reduces the access footprint on the database server from the typical 50–200 ports that may be accessible in the traditional model down to zero ports because only the application server can access the database server. This reduces the risk of unauthorized internal users accessing the database server. It also decreases the possibility that malicious users can gain access to the internal network or access the server to launch attacks. Some of the other benefits of this model include:

- Define Firewall Access: Since all application and data access must go through a firewall interface, specific inventories of port usage are gathered during the deployment of applications and can be used for auditing and compliance, when reviewing servers and applications in the environment.
- Unified User Experience: A single method of access is used by all authorized users regardless of location, i.e. the process of access is the same, whether they are at an internal (e.g. at work) or external (e.g. at home, travelling).
- Enforced Standards/Monitoring: Since all access to the applications from within the internal and external zones is funneled through a limited number of access methods in the access ring, more monitoring of access and enforcement of standards can be enabled to reduce the potential of successful attacks on the environment.
- Outbound Access Protections: Since all the application and database servers are on internal limited access segments, outbound access from the servers can be denied. This will reduce the possibility of an attacker gaining access to the system and installing software to copy data to outside sources. A user who gains access to the internal network through the firewall can only access services/resources that have been pre-authorized. This reduces the possibility of data escaping the environment.

- Reduced Horizontal Attack Surface: Since most environments are only as secure as their weakest server or application, limiting access between rings and systems within the rings reduces the possibility of a lower-security application server being compromised and then used as a pivot point to launch attacks against more secure systems.

There are a number of challenges that must be overcome in using the data-centric network model. These include:

- Performance: In a typical application/database scenario, data access is relatively simple and is over a switched network or routed across a core router (Fig. 5.1). Typically, performance is excellent. In a data-centric network, all application database access must be sent through a firewall with limited access (Fig. 5.3). This can impact database access times and cause application performance issues based on the bandwidth the firewall can handle. Large infrastructure firewalls or firewall service modules in core routers/switches can address this and allow for setup of different network and trust zones.
- Network Re-Architecture/Change: Implementing the data-centric model requires a re-design of existing environments and for example, moving servers, retraining staff, rerouting connections, and implementing additional firewalls. This is expensive,can be quite stressful to the IT staff, and can impact operations
- User Access: When database servers are on internal networks, users may become complacent/accustomed to accessing servers directly when they are in their office/workspace. In the new model, users at the office, as well as at home, will be required to VPN or use other remote access methods to access databases and other unstructured data within the environment. Users may require substantial training to learn new processes.
- Firewall Changes: When there are no firewalls or other protections in place, server owners and administrators may not have accurate inventories of what ports are required for an application to run. When moving to the data-centric model, issues can arise because non-standard ports are being used and may not have been added to new firewall rules, thus causing issues with application access/operation. Proper sniffing and network monitoring can help identify the ports so that firewall policies can be established before applications are migrated to the new environment.
- Application Development: Developers writing applications for the data-centric environment need to make sure they are writing their applications to use standard database ports or confined port ranges to ensure they can work through different layers of firewalls. When ephemeral ports and other dynamic ranges are used, this can cause issues in negotiating firewall policies.
- Outbound Application Access: Many applications today are built to have auto-update or other web service integration to pull in and process data used within the application. In an environment with limited to no outbound-access, application calls to external resources must be inventoried and allowed to use the infrastructure or other proxies to access. Open source products such as Squid Proxy

(Squid 2016) or built-in proxies in firewalls can be used to allow access to a restricted number of external resources.
- Desktop applications: Many application are design to run on user desktops with connectivity to a remote database or filesystem. In the data-centric network model these application are usually moved to the VDI (Virtual Desktop Infrastructure) setup residing within the protected network. Alternatively a VPN connection is granted to the backend data servers for those users. Generally we try to avoid granting a direct access through the firewalls to the backend resources – if it becomes necessary very narrow firewall rules are highly recommended.

When implementing the data-centric network model, it is important to assess the environment and use of applications within the environment. Given the right architecture and use of the model, the footprint of threats to applications and data within the environment can be drastically reduced with improvement in security.

5.2.2 Intrusion Prevention/Detection

Once the secure infrastructure and firewalls are properly implemented, unwanted traffic is stopped. Given the frequency of attacks and the number of potential vulnerabilities in all systems, best practices dictate that monitoring be in place to alert when possible attacks are taking place within the environment. Continuous monitoring will help respond to attacks in a timely manner, reduce the potential impact of the attack, and reduce the possibility of the attack occurring again within the environment. To achieve this, Intrusion Prevention Systems (IPS) or Intrusion Detection Systems (IDS) can be configured to monitor key ingress/egress points and critical segments of the network to alert and/or block attacks as they occur:

- IPS – actively interfering with unwanted traffic
- IDS – merely monitoring the traffic and alerting or reporting on observed anomalies or known threats

There are different types of IPS/IDS as well as different models for deployment to be considered, when planning to implement monitoring of the environment. The main types of IPS/IDS include (Scarfone and Mell 2007):

- Network-Based: monitors network traffic across particular network segments or devices and analyzes the network and application protocol activity to identify suspicious activity. The IPS/IDS simply applies predetermined (mostly vendor provided) rule set to detect anomalous traffic.
- Network Behavior Analysis (NBA): examines network traffic to identify unusual traffic flows that may represent threats, such as distributed denial of service (DDoS attacks), certain forms of malware, and policy violations. The IPS/IDS first learns typical network behavior and then reports deviations from the baseline.

- Host-Based: monitors the characteristics of a single host and the events occurring within that host for suspicious activity.

Network-based IPS/IDS is generally the most prevalent type of monitoring, but host-based monitoring is a rapidly improving technology and provides detection of a broader range of internal and external attacks. The host-based IPS/IDS are typically bundled within a suite of tools in addition to e.g. antivirus. When planning to deploy an IPS/IDS system, the deployment model plays a key role in determining what the system can provide and how it affects the entire environment. In a traditional IDS role, the system is placed in an out-of-band mode (Fig. 5.4), whereby a monitor port or tap is copying all traffic to the monitoring system. IDS analyzes the traffic to detect an intrusion/attack. In this scenario, a possible attack generates an alert to a console or response team that responds to mitigate the attack and address the vulnerability that lead to the attack. This model is used in environments with concerns over blocking legitimate traffic, and the possible impact an inline device may have on flow of legitimate traffic.

In the Intrusion Prevention Mode all traffic is routed directly through the IPS system, inspected, then sent to the destination host/network (Fig. 5.5). In this scenario, traffic designated as an attack is dropped directly inline and is not sent to the destination. This model of implementation can potentially affect performance of legitimate traffic leading to careful evaluation of the placement prior to deployment.

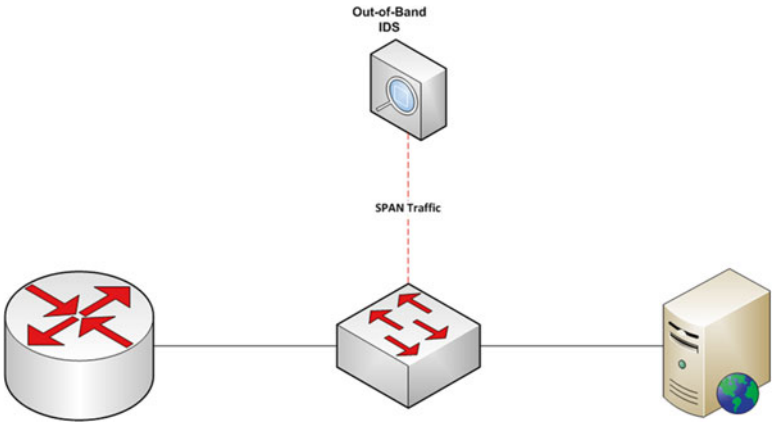

Fig. 5.4 Out-of-band intrusion detection system

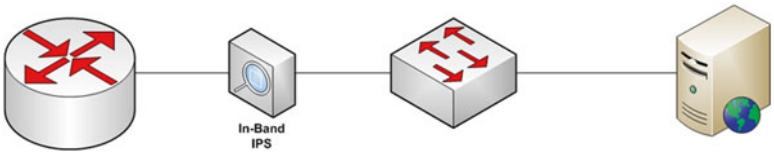

Fig. 5.5 In-band intrusion prevention system

However, in terms of attack mitigation, IPS provides the quickest response and most effective way of detecting and responding to attacks.

When planning to deploy an Intrusion Prevention System, options are available that can help address the typical concerns. The key options include:

- Performance: When purchasing/building an IPS, the network must have enough bandwidth to handle the amount of traffic that will be forwarded through the device. A number of vendors have technical specifications published for appliances, but purchasers must test the inline inspection using traffic simulators to ensure the system can handle the volume of traffic it will be inspecting in their institution. In any case, based on the IPS model, number of deployed rules, level of inspection and traffic patterns the impact on the traffic flow can be significant with possible business impact.
- Network Taps: When placing the system inline, bypass network taps must be purchased or built so that systems can be taken offline without affecting traffic flow in the environment. In maintenance windows, bypass taps can be put in bypass mode allowing reboots, updates, etc. without affecting the traffic in the environment.
- Placement: An IPS deployment for every network segment within an environment would be very costly. When planning to implement IPS, the institution should plan to place it in critical segments or ingress/egress points to get as much coverage as possible, without being required to deploy multiple devices to get complete coverage of the network. Systems placed within access infrastructure and key uplinks to routers/switches can give coverage across the environment, while not requiring multiple expensive devices to cover all segments of the network.
- Tuning/Configuration: Each vendor provides a set of signatures (configuration rules to identify bad behavior). These signatures are updated regularly to reflect the emerging threats. A number of implementations fail because signatures have been disabled, tuned out of use, or removed. Most vendors have so called 'Low-Fidelity' or 'High-Accuracy' signatures, which accurately distinguish between legitimate and illegitimate traffic, thereby minimizing false-positives that may block legitimate traffic. At a minimum, high accuracy signatures must be enabled for true blocking, whereas less accurate signatures are enabled for detection only. The number of enabled signatures might affect the performance of the IPS due to additional processing needed to match the signatures to the network traffic.
- Spend Time Tuning: When deploying the system, it is critically important to spend time developing a baseline to learn what is normal traffic in the environment, i.e. tuning the system. Ad-hoc configurations and deployments without adequate tuning can lead to a situation where signatures that have not been properly tested must be disabled and are not monitoring the environment. Adequate time spent tuning and building the signatures will lead to more effective monitoring and blocking.

Once an IPS is in place and properly tuned, it can provide a wealth of knowledge and visibility into the environment. In networks with high traffic flow, inevitably

both true and false alerts will be generated. Alert fatigue may occur. True alerts of attacks may be ignored or missed in the myriad of alerts that are generated by the IPS. To address this, a plan for monitoring and responding to alerts must be developed before the IPS is implemented. Unless there is a 24/7 presence of staff within the environment, organizations should evaluate the possibility of using a Managed Security Services Provider (MSSP) to monitor the system and alert staff when an attack is detected. An MSSP augments current staff and will help them gain knowledge of current trends. Such knowledge can be helpful in identifying actual attack traffic and reduce the amount of time spent monitoring and responding to attacks.

IPS/IDS systems complement firewalls. They allow monitoring of traffic that has been authorized for possible malicious content and facilitate responding to attacks in a timely manner. Properly deployed and configured within the environment, IPS/IDS provide additional layers of protection for the institution, which handles sensitive information and is potentially liable for breaches of security. Many modern firewalls have IPS/IDS built in and presented as part of the Unified Threat Management (UTM).

5.2.3 Remote Access/VPN

Translational research is inherently collaborative with both researchers and data frequently located at multiple geographically separate sites. Researchers need remote access to multiple networks. An infrastructure to support remote access must be implemented to allow users to access the private environment securely and reliably. There are many methods of remote access, ranging from traditional Virtual Private Networking (VPN) to other methods such as Remote Desktop/Terminal Services. Advantages and disadvantages of each must be considered when planning the infrastructure to permit remote access. It is important to take into account all the different methods that are going to be used/required/supported for remote access to the environment. Given the data-centric model for the application infrastructure, all access to applications can be considered ‘remote’ since there is no internal network. Looking at remote access in this way is valuable, when deciding what is going to be supported in the environment. To begin, planners must inventory:

1. the applications within the environment that are going to be accessed;
2. how users will access the applications; and
3. how administrators will access the applications.

Most recently developed applications utilize World Wide Web (www or web services), or have some form of web service client that accesses the application. Given this, some form of web remote access will be required and is generally provided through web reverse proxy servers or some other form of remote access or gateways. Administrators of applications may require access to terminal services or to some other type of client/server application. Table 5.1 can guide an inventory of the type of remote access that is going to be required for the environment.

Table 5.1 Methods of remote access by various applications

Application	Remote user access	Remote admin access
Public site	http/https	DB Server Studio, https
eCommerce site	https	https
eMail	https, eMail Client (MAPI)	VDI, https

Once the applications and their methods of remote access have been established, a proper set of remote access tools can be selected and implemented. When looking at remote access tools, remember that more than one method will generally be required. Not all applications, users and security models for remote access are going to work for all applications and users in the environment. It is important to select a few options that will work for most of the environment, then work to integrate the few that do not work into what is supported. Table 5.2 lists several options/methods for remote access along with their typical use, advantages and disadvantages.

Once methods for remote access have been selected for a particular environment, administrators can focus on the tools necessary to support all users and can configure them properly. Proper configurations include:

- Limited Access: Users should be restricted to those applications and services for which they have been pre-authorized. The default 'any' or 'allow all' methods of access should not allowed in the environment with an exception of administrator access well protected through a VPN. Different institutional policies are generally required for different classes of users.
- Logging/Monitoring: Logs of access by users regardless of method must be in place. Syslog and other Security Information and Event Management (SIEM) tools can be used to collect and analyze the remote access logs so that they can be generated when needed.
- Central Authentication: Ensure the remote access infrastructure uses a common identity format so that users are not required to remember multiple usernames/passwords and/or tokens to access the environment. Inconvenient access can encourage people to share credentials or methods of remote access, which can lead to elevated access privileges and violation of institutional policies.
- Multi-factor Authentication: Passwords are susceptible to multiple attack vectors including email phishing. Organizations are encouraged to add a second factor to the authentication process for remote access and highly sensitive servers. The second factor is typically in the form of a hardware token, mobile app token, text or call confirmation.
- Encryption: While encryption: While encryption of data in flight within a local network might be optional based on the network inspection needs and level of other mitigating controls – traffic through external networks is typically encrypted by default.

A well-conceived strategy for remote access with methods supported by the enterprise can provide access to internal systems and applications in a secure manner. Systems that are convenient for users permit them to perform their jobs in a secure environment, while ensuring that the institution complies with federal and state laws, regulations, and policies that govern access to sensitive data.

Table 5.2 Targeted use, benefits, and issues with various methods of remote access

Method	Target use	Benefits	Issues
VPN (IPSEC)	Full network access to the environment	Is mature and been around a long time. Is built into most security and firewall devices.	Can be problematic with slow/unreliable connections Gives user full network access (IP access) to the environment Network ports may not be opened from some secure networks
VPN (SSL)	Full network access to the environment	Uses standard https protocols Reliable across slow/unreliable network connections Open from most environments	Gives user full network access (IP access) to the environment
HTTP	Web-based applications over public and private networks	Browser-based Familiar to most users Has handlers/methods for building security onto protocol	Can be insecure in default deployments Very visible attack footprint
VDI (Virtual desktop infrastructure)	Access to a selected application or individual desktop	Thin client – small foot print, low bandwidth requirements unless video intensive applications. Security controls.	Complex setup, compatibility, performance for higher latency networks.
Terminal services/remote desktop	Full console access to remote servers and devices	Traditional desktop experience Full access to remote system Built-into Windows operating systems	Can be insecure in default deployments Older tools may not support authentication/encryption requirements Can give elevated access to systems Is not built for open/public networks using default configuration
Secure shell (SSH)	Remote user login to shell (linux) Remote user tunnel	Secure transport protocol for open/public networks Low cost	Takes more configuration from client side Requires configuration changes on some client machines

5.3 Operating Systems and Cybersecurity

Operating systems of applications are probably second to networks in terms of implications for security of the environment. Operating systems of servers are complex, and. have vulnerabilities that can lead to compromise of applications. It is important to choose operating systems that meet institutional needs, secure them by default, and continuously monitor and improve the security of the operating systems within an institution's particular environment. This section reviews challenges of securing operating systems with specific examples of how these can be addressed.

5.3.1 Configuration

Configuration standards, policies, and processes document the setup steps and procedures for configuring applications, network devices, servers, etc. within the environment. Configuration standards should be documented for each application or device within the environment. This ensures that applications and devices are consistently setup, securely built based on best practices, and are accessible in the event that primary support personnel are not available.

Initial configuration standards must be built for a particular environment. They should include:

- Best Practices: Do not re-invent the wheel. There are a number of sources of standards that have been well tested and can serve as a starting point for an institution (Scarfone and Mell 2007; Quinn et al. 2015; CIS 2016; NIST-NVD 2016).
- Do Not Copy Standards: Many of the available standards are purpose-built for specific regulatory and/or security requirements and may not be fully applicable to your particular organization and its environment. Institutions should use the standards of others as a guide, but standards adopted by an institution must fit its specific needs.
- Conduct Regularly Scheduled Reviews: Once a standard is written, it needs to change over time as technology, regulations, and the environment change. Configuration standards should be reviewed at least annually to make sure they are still relevant and accurate.
- Auditable: Compliance with standards must be auditable. This can be done by manual, scripted, or automated reviews. The Security Content Automation Protocol (SCAP) provides specifications and practical guidance to enable automated vulnerability management and audits of compliance (Quinn et al. 2010; NIST-SCAP 2016). Furthermore modern vulnerability scanners can validate setups (such as OS, database, applications) against published standards such as Center for Internet Security – CIS, Defense Information Systems Agency – Security Technical Implantation Guides (DISA STIG), etc.

There are a number of good sources for configuration standards and checklists. Examples include:

- Center for Internet Security (CIS)
  - http://benchmarks.cisecurity.org/en-us/?route=downloads.benchmarks
- NIST National Vulnerability Database Configuration Checklist Project
  - http://web.nvd.nist.gov/view/ncp/repository

In addition to these vendor neutral databases, most vendors provide guides to setting up the operating system with proper levels of security. Important items to consider, if a vendor guide is not available or is incomplete include:

- Services/Daemons: Disable all unnecessary services, programs, applications, and daemons that are not explicitly needed by the operating system. This reduces the footprint of the server and reduces the possibility of unneeded software affecting the operating system.
- Limited Access/Authorizations: Only allow local access to those administrators and power users who must directly access the system. This can be achieved through local authentication/authorization requirements along with network level controls for limiting access to administrative ports and services. This might also include disabling or renaming local accounts, changing default passwords, etc.
- Secure Management: When allowing remote management, insecure management protocols such as telnet, rsh, VNC, etc should be disabled and only secure management protocols enabled. This will reduce the possibility of individuals intercepting credentials of administrators and using them for malicious purposes.
- Auditing/Logging Events: System should be configured to log access and changes to the system. This will provide an audit trail of who made what changes and when. This can be useful in investigations of possible breaches of security and in troubleshooting.
- Patching/Updating: Each organization should build requirements into the standards to require periodic patching and updating of systems to ensure they have the latest software and any vulnerabilities that have been identified are patched as soon as possible.

When writing configuration standards, they should not be limited to operating systems and other network based devices, but should include any applications within the secure environment. This helps ensure consistent setups, knowledge transfers, and continued secure builds of applications running within the environment.

5.3.2 *Patching*

Given the size and complexity of applications and operating systems, it is inevitable that vulnerabilities will be found and require patching. This is highlighted in the latest Symantec Internet Threat Report (Symantec 2015), which noted a 30% increase in the number of vulnerabilities in software releases in 2010 compared to 2009. Patches are released to address vulnerabilities, after a period of delay while methods of correction are developed. If left unpatched, the vulnerabilities can lead to compromise of the system. To ensure these patches are applied and the systems are secure from these vulnerabilities, a patching program that includes regular scanning and updating needs to be put in place.

The first step in implementing a patching process is to regularly scan the environment for potential required patches. The scans should not be limited to server and applications, but should also include the network equipment (switches, routers, firewalls, etc) within the environment. Vendors generally release patches and updates for their equipment. Systems such as Windows and Linux have automatic update capabilities that can download and install the available patches for the system. There are products that automate scanning to identify vulnerabilities and available patches. Examples include Windows Server Update Services (WSUS) (Microsoft-WSUS 2016), Freeware OpenVAS (2016), Microsoft Baseline Security Analyzer (MBSA) (Microsoft-MBSA 2016), and Flexera Personal Security Inspector (PSI) (Flexera 2016). Once the list of required patches is developed and the available patches are downloaded, they must be installed on all the affected systems. As verification that patches have been installed properly within the environment, verification scans should be run against all the systems. This can be part of the follow-on monthly scans to ensure continuity and continued updating of systems within the environment.

The real challenge to implementing a patch process is setting a standard and regular process for installing patches. Most vendors have moved to a once a month release of patches (outside of critical updates). Organizations can follow the same schedule and regularly install patches on all systems, applications, and network devices within the environment. This will ensure the continued update and review of all the systems within the environment.

The best practices of patch management suggest a careful selection of patches to apply and apply them in succession to test environment before production for validation and avoid interference with production application. This is not always feasible given the staffing levels and amount of patches being issued by various vendors.

The inevitable argument that is raised is: ‘This will cause issues with my [system, application, device]’. If a patch or update has been shown to have compatibility issues with a particular application or device, there must be a process to document the exception, approve it, and regularly review it. The documentation should include the source of the exception (from support, through testing, etc), who generated the exception (owner), and the expected timeframe for resolving the problem. During

the regular scanning and installing of patches, all exceptions should be reviewed to ensure they are still applicable and required within the environment. Failure to resolve exceptions can lead to accrual of systems and applications within the environment that are out of standard and create potential vulnerabilities that could adversely affect the rest of the environment.

5.3.3 Vendor Management and Evaluation

In the fast paced world of information technology and cyber security many new vendors emerge every day with solutions that are often appealing from the execution, cost, feature and other perspectives. Organizations usually develop institutional standards for vendor and vendor products evaluation to assess conformance with industry standards, attention to security during the development and implementation, long term business viability, customer service, etc.

At minimum the organization should ensure that:

- vendor supported systems have a set schedule and that is agreed to upon contract/support agreement signing
- vendor is held accountable to appropriate standards

5.4 Protecting Sensitive Information

Sensitive information such as Personal Health Information is commonly used in translational research and must be protected. Typically, it is captured and stored in three types of systems: Research Patient Data Warehouse, tools for Electronic Data Capture that may include data storage, and Research Data Storage systems.

5.4.1 Protected Health Information

Protected Health Information, also known as Personal Health Information, (PHI) does not include publicly available information that is lawfully made available to the general public from federal, state, or local government records. Protected health information that cannot be exposed to individuals without permission of the patient is defined as information that when used alone or in combination with other information can identify an individual (whether living or deceased). Table 5.3 lists these identifiers, as defined in the United States by the Health Insurance Portability and Accountably Act (HIPAA) Privacy Rule.

Table 5.3 List of 18 HIPAA PHI identifiers

	Identifier
1.	Names
2.	All geographical subdivisions smaller than a State, including street address, city, county, precinct, zip code, and their equivalent geocodes, except for the initial three digits of a zip code, if according to the current publicly available data from the Bureau of the Census: (1) The geographic unit formed by combining all zip codes with the same three initial digits contains more than 20,000 people; and (2) The initial three digits of a zip code for all such geographic units containing 20,000 or fewer people is changed to 000.
3.	All elements of dates (except year) for dates directly related to an individual, including birth date, admission date, discharge date, and date of death and all ages over 89 and all elements of dates (including year) indicative of such age (except that such ages and elements may be aggregated into a single category of age 90 or older)
4.	Telephone numbers
5.	Fax numbers
6.	Electronic mail addresses
7.	Social security numbers
8.	Medical record numbers
9.	Health plan beneficiary numbers
10.	Account numbers
11.	Certificate/license numbers
12.	Vehicle identifiers and serial numbers, including license plate numbers
13.	Device identifiers and serial numbers
14.	Web Universal Resource Locators (URLs)
15.	Internet Protocol (IP) address numbers
16.	Biometric identifiers, including finger and voice prints
17.	Full face photographic images and any comparable images
18.	Any other unique identifying number, characteristic, or code (excluding a random identifier code for the subject that is not related to or derived from any existing identifier)

5.4.2 Research Patient Data Warehouse

The Research Patient Data Warehouse (RPDW) is one of the most highly protected components of the IT infrastructure in a research-intensive medical center. Typically it contains information from the electronic medical records of large numbers of patients. See Chap. 6, Research Patient Data Warehousing, for an extended discussion of data warehouses. Protection of PHI within a RPDW requires special attention to:

- Access control
- Auditing of control and data changes
- Strict change control of any alteration of the system

A RPDW typically contains huge amounts of data, e.g. patient's demographics, diagnostic and procedure codes, medical history, results of laboratory tests, images, and genomic data. Accruing, storing, and searching the data in a timely manner

requires high performance systems. Security systems must be carefully designed to avoid overly compromising performance of the warehouse. It requires a balance between ability to process and deliver data in a timely manner and ability to implement security controls that adequately protect the data. Performance can be improved by moving some controls to auxiliary systems surrounding the RPDW:

- Increased physical security of the data center as a trade off for on-disk data encryption
- A dedicated network segment surrounded by a firewall as a trade off for an in-band firewall
- Accessing data through an application server with a tight control rather than a direct database connection
- Application server to database connection might not be encrypted for performance reasons putting an additional burden on security of the application server
- The audit trail might be stored externally rather than stored on the system holding the RPDW

Many of the trade-offs are specific to an institution, its environment, and infrastructure. Institutional specificity might be related to:

- Processes for change control
- Types of queries or in general how the system is utilized
- The process of loading data (nightly full or incremental refresh, continuous/online versus batch load)
- The system of data backups such nightly, online DR, virtual vs. physical tape, etc.
- Method of user access (accessing a single RPDW or accessing a separate datamart that contains an extract of the primary data in a separate database)
- Reporting utilities and the security of service accounts accessing the data
- Utilizing all database built-in security controls – row level security, encryption

5.4.3 Tools for Electronic Data Capture

Tools for electronic data capture are commonly provided to investigators as part of the design and implementation of translational research projects. These allow researchers to collect data to be used in studies, clinical trials, etc. From the point of view of those responsible for security, it is important to note that tools for electronic data capture are typically user facing and hence require strong identity and access management. One such system might host a few or 100 s of studies making it a high level target for unauthorized users. In order to minimize the possibility of data compromise:

- Encryption between the client and the server is a must. This is typically done by SSL (https), but additional layers such as VPN or SSL VPN might be used

- Regularly scheduled changes in passwords must be enforced, typically at 90 day intervals
- Role management – distinguish between study coordinators and data entry personnel. Grant permissions based on need and role.
- Entitlement reviews – ensure that only the people who are entitled to access to particular data sets or studies receive access, and regularly review these entitlements.
- Comply with best practices such as regular server patches, centralized review and storage of the logs, change and configuration management.

5.4.4 Research Data Stores

Biomedical research institutions typically house many data generating groups/cores. In order to share data efficiently, network attached storage (NAS) or online (web) based storage systems are generally used. Much of the data in a medical research institution will contain Protected Health Information, as previously discussed for Research Patient Data Warehouses. There are a several basic considerations, when planning to protect information. In a NAS environment it is important to realize that;

- The data in flight might not be encrypted. The Common Internet File System (CIFS) does not have encryption built-in, whereas the Network File System (NFS) does in version 4 (IAPS 2007). Direct Internet access to file servers should be provided only through VPN.
- Although users can manage permissions to some degree themselves in CIFS and NFS environments, the built-in permissions management might not be ideal. Add on products and careful monitoring of permissions is encouraged in order to keep the permissions well-defined and manageable long term.
- Identity within the NAS is only as strong as the institution's identity management system. Inadvertent additions of users to groups, confusion of users with the same or similar names, unclear definition of groups and the data to which members have access are some of the issues that can plague an identity management system and must be clearly resolved. It is important to develop standardized workflows for addition and deletion of users and groups.
- Audit-trails are essential not only for compliance with laws and regulations, but as a best practice. Many users fail to realize the side effects of, for example, moving or removing data directories. Having the ability through audit trails to quickly identify who did what, when and where helps dispel many misconceptions and improve trust in the system.
- Data classification
- Data storage guidance (what goes where)

Web based data stores typically handle internet access better that NAS, yet can be mounted as a network drive e.g. using WebDAV (Web Distributed Authoring and

Versioning) protocol. With regard to security, some characteristics of web based data stores are:

- Encryption using SSL (https)
- Authorization/permissions are not limited to pre-defined protocols, but are dependent on the hosting application.
- Audit trails are easy to add
- During the upload process, information can be automatically parsed from the data and hence allow more flexible searches
- Typically stores "structured" data

5.5 Role of Users in Protecting Information

The focus of this chapter on institutional cybersecurity has been on technology. In reality, breaches of security with loss of protected or confidential information usually occur because users fail to comply with institutional policies and standards rather than because of technical failures of network security and intrusion prevention systems. For example, theft, loss, and/or misuse of portable devices are one of the most common breaches of security across industries (Verizon 2011). It is critically important that institutions develop and enforce policies defining acceptable and unacceptable uses of information resources. Institutional policies and standards should address such issues as:

- Behaviors of users must be consistent with the mission, vision, and core values of the institution.
- Users must comply with applicable laws, regulations, and policies.
- Users acknowledge that sharing of Confidential Information with unauthorized individuals may result in the waiver of legal protections and may place the user and employer at risk of criminal or civil liability or may damage their financial standing, employability, privacy, or reputation.
- Users are individually responsible for maintaining the confidentiality, security and integrity of their chosen passwords.
- Users must promptly and properly dispose of, sanitize, and/or destroy confidential information in any form (e.g., paper, electronic, on Portable Devices) when no longer useful
- Users consent to institutional auditing and monitoring for excessive or inappropriate personal use, criminal activity, regulatory violations.
- Users understand that violations of any policy may subject them to disciplinary action up to and including termination of employment.
- Users must ensure that: Confidential Information is not visible in an unattended work area or on an unattended workstation; they log off of systems after access; they do not post passwords or user IDs in visible areas.
- Users are responsible for ensuring optimum security of Portable Devices (regardless of device ownership) that access, transmit, or store the institution's

Confidential Information. Users must immediately report any loss or theft of computing devices.
- Users are prohibited from interfering with the intended functioning of the information system, e.g. disabling security devices, inserting viruses, inserting network monitoring devices, installing unapproved software, destroying files.
- Security awareness training – having system users understand the types of threats that they could fall victim too, such as social engineering and phishing campaigns and the risks associated with them.

To summarize, maintenance of high levels of cybersecurity to protect confidential and protected health information requires a combination of best practices in technology and thoughtful institutional policies that are enforced. One without the other precludes success.

A number of frameworks have been developed to aid organizations with risk assessment and mitigation. Our organization is subject to FISMA (Federal Information Security Management Act) for a number of federal contracts. FISMA uses NIST 800-53v4 (JOINT-TASK-FORCE 2013) "Security and Privacy Controls for Federal Information Systems and Organizations" and we follow NIST standards for the annual Information System Security Plan (ISSP).

5.6 Joint Management – Hospital and Research

As discussed in Sect. 4.8, to support research in a research oriented medical center, it is important to define clearly who is responsible for providing IT support to the different missions – patient care, business operations, and research. In the authors' institution, clinical and business operations are supported by hospital information services (IS), while research is supported by biomedical informatics (BMI). The two groups are separately staffed and budgeted, but must collaborate very closely to provide coherent support to the overall research and clinical enterprise. With regard to operation of networks and implementation of cybersecurity, IS and BMI have developed a responsibility matrix (Table 5.4). This has served the organization well and prevented misunderstandings. IS is responsible for installing and maintaining the physical networks, including firewalls. BMI works with IS to design networks

Table 5.4 Distribution of responsibilities for support of research between Hospital Information Services (IS) and Biomedical Informatics (BMI)

Issue	IS	BMI
Physical networks	X/J	J
Security	X/J	J
Authentication	X/J	J

X primary responsibility for operational support, *P* primary responsibility to formulate policies, *J* participates in formulation of policy

that support research programs and to develop standard operating procedures for identity and access management. The goal is to create an environment in research that is secure and meets regulatory requirements with regard to protection of sensitive information. With this in place, the hospital can transfer clinical information from electronic medical records to the research data center with assurance that the information will be properly protected.

References

CIS. Center for internet security. 2016. Retrieved March 21, 2016, from http://www.cisecurity.org.

Cruse A. Processor privilege levels. 2016. Retrieved March 21, 2016, from http://cs.usfca.edu/~cruse/cs630f06/lesson07.ppt.

Flexera. Flexera personal security inspector. 2016. Retrieved March 21, 2016, from http://www.flexerasoftware.com/enterprise/products/software-vulnerability-management/personal-software-inspector/.

IAPS. NFSv4: overview of new features. 2007. Retrieved March 21, 2016, from http://www.iaps.com/NFSv4-new-features.html.

JOINT-TASK-FORCE. Security and privacy controls for federal information systems and organizations, NIST special publication 800–53 Revision 4. Gaithersburg: U.S. Department of Commerce, National Institute of Standards and Technology; 2013.

Microsoft-MBSA. Microsoft baseline security analyzer. 2016. Retrieved March 21, 2016, from http://technet.microsoft.com/en-us/security/cc184923.

Microsoft-WSUS. Windows server update services. 2016. Retrieved March 21, 2016, from http://technet.microsoft.com/en-us/updatemanagement/default.aspx.

NIST-NVD. National vulnerability database checklist program. 2016. Retrieved March 21, 2016, from http://web.nvd.nist.gov/view/ncp/repository?tier=&product=&category=&authority=&keyword=.

NIST-SCAP. Security content automation protocol (SCAP). 2016. Retrieved March 21, 2016, from https://scap.nist.gov/.

OpenVAS. OpenVAS: open source vulnerability scanner and manager. 2016. Retrieved March 21, 2016, from http://www.openvas.org/.

Quinn S, Scarfone K, Barrett M, Johnson C. Guide to adopting and using the security content automation protocol (SCAP) version 1.0, NIST special publication 800–117. Gaithersburg: U.S. Department of Commerce, National Institute of Standards and Technology; 2010.

Quinn S, Souppaya M, Cook M, Scarfone K. National checklist program for IT products – guidelines for checklist users and developers, NIST special publication 800–70 Revision 3. Gaithersburg: U.S. Department of Commerce, National Institute of Standards and Technology; 2015.

Scarfone K, Mell P. Guide to intrusion detection and prevention systems (IDPS), NIST special publication 800–94. Gaithersburg: U.S. Department of Commerce, Technology Administration, National Institute of Standards and Technology; 2007.

Squid. Squid cache proxy project. 2016. Retrieved March 21, 2016, from http://www.squid-cache.org/.

Symantec. Symantec threat report. 2015. Retrieved March 21, 2016, from http://www.symantec.com/threatreport/print.jsp?id=threat_activity,vulnerabilities,malicious_code,fraud_activity.

Verizon. Verizon data breach report. 2011. Retrieved March 21, 2016, from http://www.verizonbusiness.com/resources/reports/rp_data-breach-investigations-report-2011_en_xg.pdf.