**Group Project**

Isaiah Bristol, Josh Seaford, and Ian Hagmann

School of Cybersecurity, Old Dominion University

IT 417 Management of Info Security

Professor Kalburgi

Dec. 1, 2025

Roles:

Isaiah Bristol: a, i, ii, iii

Ian Hagmann: iv, v, vi, vii

Josh Seaford: viii, ix, x, xi

# Table of Contents

# Introduction (a)

The goal of this project is to design a secure network for the Strome College of Business. Our mission is to ensure that both academic and administrative data remain confidential and available, all while maintaining its integrity. The Strome College of Business is housed in Constant Hall. Our network design will be able to support the needs of multiple classrooms, labs, and offices. The network must be able to accommodate up to six hundred computers across three subnets. Additionally, two domain controllers, network-attached storage (NAS), a web server, and an internal network using private IP addresses will be supported. To achieve these needs, we will use both Cisco and Windows products. The network will consist of a Cisco Catalyst 8500 series edge router, a dual firewall setup (Cisco 3100 and 4200 Series), and a multi-layered switch setup. Cisco Catalyst 9400 series core switches will be paired with Cisco Catalyst 9300 series access switches. Windows 2025 will provide the necessary capabilities to run the two domain controllers, NAS, and internal web server.

**Router:** Cisco Catalyst 8500 Series Edge Router

This is a high-performance router capable of segmentation and secure WAN functionality. This will serve as the primary router, which will connect the Strome College of Business to the Internet. It will also provide routing for all internal traffic across subnets.

**Justification:**

This router has high throughput, which makes it capable of supporting high volumes of traffic across encrypted tunnels. It provides scalability, should the number of hosts increase. This router is reliable and flexible, which makes it an excellent choice for this network.

**External Firewall:** Cisco Secure 4200 Series Firewall

The 4200 firewall is the first line of defense against external traffic and threats. It protects all internal networks and systems by filtering inbound internet traffic and enforcing security policies. This model comes with integrated Firewall Threat Defense (FTD) software that is capable of packet inspection, malware detection, and advanced logging. Using this

software, the firewall can be set to either an IPS or an IDS configuration.

**Justification:**

The Cisco Secure 4200 Series Firewall has high throughput capabilities and can handle high volumes of encrypted traffic. The built-in IPS/IDS security features make this ideal for any situation where strong security is necessary. The firewall's scalability helps it fit perfectly in a university environment.

**Internal Firewall:** Cisco Secure 3100 Series Firewall

The Cisco Secure 3100 Series firewall will act as an internal segmentor between the lab, classroom, and faulty subnets. This will provide even more security across the network by protecting the VLANs from each other, should one of them become compromised. For example, if a student were to plug a USB device into a lab computer, which resulted in harm to the lab subnet, the firewall would prevent further spread to the student and faculty subnets. Having this internal router will reduce the risk of malware and internal threats while allowing the other components of the network to function at their highest capability. This firewall model also comes with IDS/IPS capabilities, making it an even greater security asset.

**Justification:**

An internal firewall ensures that internal traffic is carefully monitored. Because Old Dominion's academic environment often includes students plugging in new devices, this is essential for maximum security. The combination of both firewalls provides a greater range of security capabilities. Both the Cisco Secure 4200 series firewall and 3100 series have IPS/IDS integrations. However, each one can only be set to one mode at a time. Setting the external firewall to IDS allows it to be an alarm should any harmful traffic try to make its way into the network from the internet. Setting the internal firewall to IPS would enable it to prevent incoming traffic from reaching any critical systems and prevent internal threats from spreading across the network. Using this firewall as a second layer of defense is necessary for the university environment where it will exist.

**Core Switches:** Cisco Catalyst 9400 Series Switch

The Cisco Catalyst 9400 Series switches will serve as core aggregation switches across the network. This will create

redundancy on the network for internal routing. This centralizes switching for subnets that house the lab, classroom, and faculty devices.

**Justification:**

These switches have a high bandwidth suitable for an academic environment. With hundreds of devices already spread across the internal subnets, these switches provide scalability should it become a need in the future. It has strong performance capabilities and security features that will pair well with and take the load off the access layer switches.

**Access Switches:** Cisco Catalyst 9300 Series Switch

These switches will provide connectivity to classrooms, labs, and faculty offices. They will connect endpoints securely, which will allow the internal firewall and core switches to focus on their role.

**Justification:**

They provide even more security with their seamless integration with the other Cisco products. The advanced security features and integration make them excellent choices to pair with the rest of the network components.

**Windows Server 2025:** Domain Controllers, NAS, and Web Server

The network will use Windows Server 2025 to deploy the domain controllers, NAS file server, and internal web server.

Deploying two Windows Server domain controllers provides centralization for both authentication and authorization services. Active Directory, DNS, and DHCP are also included. Two controllers allow for redundancy, should there be a hardware or software failure.

The NAS file server provides storage for a wide range of project files, research, and miscellaneous data. Secure file sharing and permissions for access are encompassed within the Windows file server.

An internal web server allows the Strome College of Business to keep its hold on internal applications, resources, and other services. The server can also integrate with the NAS file server and domain controllers to further secure student and faculty access across the network.

**Justification:**

Windows sets the standard for managing information and data. Windows Server 2025 uses advanced encryption, Windows Defender features, and Credential Guard. This protects the domain controllers, NAS server, and internal web server from malware and unauthorized access. Windows Server 2025 provides backup services and keeps services available should one part of the system fail. If one domain controller fails, the other maintains its capabilities. Should a hardware or software outage occur, NAS will be protected against data loss. Centralized management ensures that security standards and policies will be carried out across the network. It is also scalable, working alongside the Strome College of Business as it grows.

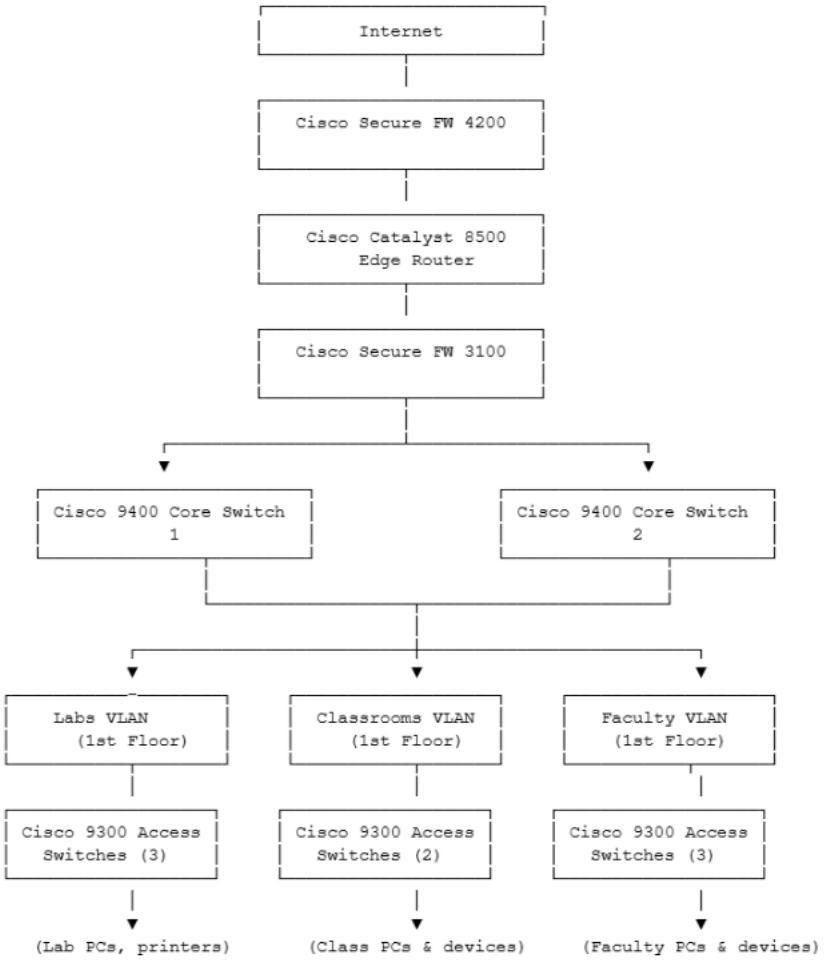

# Possible threats and attacks (i)

The Strome College of Business will face a wide variety of threats that could compromise the data and infrastructure it houses. Because university environments are publicly accessible and multiple users might use the same device, they are subject to unique threats. It is important to identify all possible threats and vulnerabilities to prepare for worst-case scenarios. This section aims to do just that. We will dive into multiple attack types and discuss the unique threats they pose to the Strome College of Business.

**Malware:**

Malware attacks pose a high risk to the resources and services of the Strome College of Business. Malware could be introduced by multiple agents, even those not meaning to. Personal student devices like laptops and USB thumb drives run the risk of spreading things like viruses and worms if they are infected. Lab and faculty devices can become the source of malware if the user visits an untrusted website or accesses a link that exposes them.

**Phishing Attacks:**

Phishing attacks are an ever-present threat, especially in a university context. Fraudulent emails are frequent among university students. These attacks could cause harm if they succeed in their goal. These goals range from stealing login credentials to installing malware on the network. The loss of login credentials can lead to a threat actor gaining unauthorized access to Strome and subsequently, Old Dominion's broad range of resources. A phishing attack on a personal device can also introduce malware into the network, as mentioned in the previous section.

**Internal Threats:**

Unauthorized access by users on the network could compromise the confidentiality, integrity, and authenticity of data and information within the network. Although students should have access to the college's resources, they should not be able to access faculty and administrative information. Likewise, not all faculty members should be able to access every piece of information. Only that which is essential to complete their jobs. Such allowances could compromise the security of the network and connected systems.

**Physical Security:**

Damage to the physical components of the network could pose a significant threat to the availability of its resources. This could come from unauthorized access to Constant Hall wiring closets and server rooms. Damage to network hardware could impact the connectivity of the system. Even if no physical damage is done, tampering with settings on the switches and routers could result in performance issues and increased exposure to other vulnerabilities. Again, the introduction of malware could pose a serious threat to the Strome College of Business. Weak physical security could potentially lead to the introduction of malware directly to key network devices. There is also a threat to academic devices in general. Lab and classroom PCs could be the target of theft. Weak physical security could result in the loss of these devices.

**Human Factors:**

The weakest parts of security in any network are humans. As previously mentioned, unsuspecting users could expose themselves and the network to harm if they carelessly click on a link and download malware, or connect a USB device that was already carrying some. Humans could even affect the network before it even goes online. IT personnel could incorrectly configure network segments and allow access points for potential threats.

**Steps for Prevention:**

It will take a robust, multilayered effort to mitigate the potential for threats to impact the network.

As mentioned in the introduction, the network will be separated across three subnets (Labs, classrooms, and offices) to help stifle movement within the network. The implementation of role-based access controls will prevent unauthorized access to unnecessary resources. Students will not be able to access faculty resources, nor will they be able to access and edit network attributes.

The university has already implemented multi-factor authentication via the Duo Mobile app. This will help in the event that login credentials are compromised. Setting mandatory password expiration windows will also help combat credential theft.

Regular updates and network maintenance to patch vulnerabilities will be crucial in securing the network.

Scheduled maintenance and downtimes will help keep a consistent eye on possible weak points.

All of the Cisco and Windows products come with security functions. The multilayered firewall setup creates a strong safeguard against suspicious activity and malware. The built-in intrusion detection and prevention sensors will help contain a threat should one try to make its way into the network.

As an IT team, we will ensure that only essential ports are open between VLANs and subnets. This will decrease the number of vulnerable access points.

Physical security will also be bolstered to prevent damage, theft, and network manipulation. Wiring and network closets will be locked. Only IT personnel will be able to unlock it with their IDs. Cameras in and around Constant Hall will help deter someone from entering sensitive areas and tampering with hardware.

Logging and audit trails that are under regular review will help bolster security by giving IT staff the necessary information they need to improve security.

# Planning, Organization, Risk Analysis, and Policies (ii)

The first place to start when planning and conducting risk analysis is to identify key assets. Including hardware, the key assets in this network are the domain controllers, the NAS file server, the internal web server, computers across the classroom, lab, and faculty levels, and the networking infrastructure. Each asset has a different level of vulnerability that must be addressed in order to maintain confidentiality, integrity, and availability.

The Strome College of Business will implement policies that are intended to be a guide for exercising best practices. There are acceptable use policies to outline proper internet and device usage. Such policies play a crucial part in protecting key assets.

These policies will use information from a Threat assessment compiled into a Threat-Vulnerability-asset worksheet. This worksheet will help organise things like assets and vulnerabilities, evaluate risk, and plan how to manage it all.

**Table 1**

*Threat-vulnerability-asset analysis*

| Asset | Threat(s) | Vulnerability | Controls |
|---|---|---|---|
| Domain Controllers | Credential theft, System Compromise, Privilege escalation. | Weak passwords, Outdated Windows Server, Misconfigured AD and group policy | Credential Guard, segmenting and isolation, dual domain controllers |
| NAS File Server | Ransomware, unauthorised access, Human misuse | Weak access controls, unencrypted traffic | IPS and IDS on firewalls, access control, and backups |
| Internal Web Server | Unauthorized access, Malware | Default Configurations, Weak credentials | TLS, traffic filtering, regular updates |
| Classroom/Lab/Faculty Computers and devices | Phishing, Malware, Student devices, Physical tampering | Outdated software, Unsecured devices, USB connections | Segmentation, subnetting, internal firewall, inactive device ports, secured devices |
| Network Hardware/Infrastructure | Malware, DoS attacks, and Physical tampering | Default device credentials, Open ports, Weak physical security | Locked wiring closets, segmentation, subnetting, IDS/IPS, port authentication |

# Measures to ensure confidentiality and authenticity (iii)

The network focuses heavily on security to protect the confidentiality and authenticity of all data and information. The infrastructure consists of Cisco Secure 4200 and 3100 firewalls, Cisco Catalyst 9400 and 9300 switches, and Windows Server 2025. These play a major role in protecting data and information. Each device provides its own level of security; however, they are most effective when put together. The layered defenses ensure that sensitive information will be protected from unauthorized access while protecting its validity. The firewalls and router play the greatest roles in encryption. They bolster the security and support VPN tunneling, which further secures communication between the network and the internet.

The perimeter 4200 series firewall is the network’s barrier to external threats and intrusion. It prevents unauthorized external access to information. This helps protect the system's confidentiality by inspecting inbound and outbound traffic. It uses strong packet filtering and FTD intrusion prevention. This not only prevents an external source from coming in and accessing data, but also protects users against sending that data out of the network. Only encrypted data and authenticated information can pass through the firewall. The firewall prevents connections to illegitimate sources that might attempt to spoof their addresses or change header information. This ensures that all inbound traffic is from authentic web clients. Additionally, the firewall offers VPN and encryption capabilities to ensure that any network communications are secure. Strict encryption standards maintain the confidentiality of all network traffic. Credential requirements help maintain authenticity between internal and external connections. Only trusted clients will be able to communicate with the network.

The 8500 series edge router controls the flow of traffic between the internal network and the Internet. Its support of VPN technology ensures that connections within the network are secure and reach where they are supposed to. The router also authenticates the source and destination of shared files to protect the system. It discards anything that does not match the routing policies to prevent illegitimate sessions across the network. The router also manages encrypted tunnels. Again, these tunnels contribute to protecting confidentiality. The router’s ability to authenticate VPN IDs ensures that only authentic communication flows can occur.

The second firewall (Cisco Secure 3100 Series Firewall) protects the Windows domain controllers, NAS, and internal web server by ensuring that all shared information that originated within the network is carefully monitored and filtered. It also protects against internally spoofed sessions by authenticating device communications. It works in conjunction with the router and external firewall to contain and prevent internal threats from reaching other parts of the system. It also adds another layer of VPN encryption for traffic. In the event that an internal device is compromised, this helps prevent it from intercepting data and directory information. It verifies the authenticity of internal hosts to prevent a threat actor with inside access from communicating with essential services. This prevents lateral movement within the network by illegitimate devices. This makes communication more reliable within the Windows environment.

## Access control policies and implementation (iv)

Access controls are a technical security measure that define who can access what resources and the way that they can. In a network that uses access control, a user must first authenticate with a server running access control software by expressing who they are and proving it through things like passwords, fingerprint scans, or physical ID cards. Then the server will, in our case, send the user a ticket-granting ticket (TGT), which allows the user to request access to specific resources. Our network running a Microsoft Windows environment will be using Kerberos for access control/authentication on the two Active Directory servers. Another important facet of access control software is that it audits security-related events like failed authentication attempts, denied authorisation requests, and changes to the access control configuration. The necessary control of resources on the network also means that any computer device on the network has to authenticate itself with an access control server. This prevents rogue devices from being plugged into the Ethernet cables or open outlets in the rooms and accessing the network.

There are several ways to go about defining who can access what in an organization that range in ease. The two overarching types are discretionary access controls (DACs) and nondiscretionary access controls (NDACs), the former of which leaves it up to users to grant access and the latter, which is controlled by the organization. We are definitely going to go with non-discretionary access controls for our network, but it is further subdivided into different strategies. One strategy is called role-based access controls (RBACs), and it bases access on the user's role or position within the organization. The next is mandatory access controls (MACs), which classify data based on its sensitivity and assign these levels of access to users. Then there are attribute-based access controls (ABACs) that rely on characteristics to regulate access. I see role-based access controls to be a good fit because we can define set roles for students, professors, and any other specific group without needing granular control. Users can also be outfitted with multiple roles, adding their privileges together as if a professor is working on a research project that needs to have its access controlled. The roles we make should follow the principle of least privilege, meaning that users are only given the minimum permissions they need.

The student role would be the simplest because they won't need access to many of the resources. Of course, the student role will be assigned to all student accounts, and any special

exceptions will require new roles. The student role will only be able to log in from the lab computers, as the faculty offices should clearly not have students logging in, and the regular classrooms only have a computer for the instructor. Students will not be permitted to use printers in the building, as there is already a particular system for students to use the printers elsewhere. They will have access to a student folder in the NAS, in which they have to create their own subfolder for use. Students will not be given special access to the web server, only ports 443 and 80 for regular web traffic.

Professors, on the other hand, may require different levels of access depending on what they are working on. Special roles can be added for collaboration between professors, which would require separation from other staff. The professor's role would be able to log in to only the classroom and faculty offices subnets, as we don’t want to allow students to try and log in as a professor in the labs. This also means that the instructor computers in the labs will be on the classroom subnet. Professors will be given access to the printers on the necessary ports and only the ones in shared spaces or the room they are currently in to avoid confusion. As for the NAS, professors will be given access to their own folder and several groupings of shared ones. They will also have the ability to read the student folder and its subfolders, in case that comes in handy, also meaning that students will have no expectation of privacy in those folders. The professor's role will also be permitted to access the web server on ports 443 and 80, and if they need to add webpages to it, they will need another special role.

## Firewalls, policies, and implementation (v)

Firewalls are the first line of defense when it comes to computer network security. They include a set of rules that determine appropriate action based on the information that is passed through them. These rules are compared to incoming packets in order, usually from top to bottom. Firewalls can be categorized in a number of ways based on factors such as what they deal with or how. What I am going to be talking about are packet-filtering firewalls that operate primarily in the network layer on packets. This type of firewall's rules include information from both the network and transport layers of network packets, like IP addresses involved, source and destination port numbers, protocol, and more. A further distinction for packet-filtering firewalls is whether they are by default allowing traffic through, called blocklisting, or by default blocking all traffic, called allowlisting. This naming scheme seems backward at first glance, but these names refer to the nature of the other rules set in the firewall. Allowlisting presents more of a burden on security staff when they need to make changes in exchange for stricter security. Blocklisting would be more likely to let malicious packets through and would rely more on other security measures on hosts to protect them initially.

The firewalls implemented in our network are going to require specific policies for their operation and changes. We are going to be running two firewalls, one on the outside to protect the router, and another inside the router to create a DMZ and protect the subnets. The external firewall will have rules in a blocklist configuration to ensure blocking of known attacks and protocols we don't want to let in, like Telnet. The internal firewall will then run an allowlist configuration for its rules to provide greater security and control. Rules being added to the firewalls or changed will undergo testing outside of business hours and will be overseen by at least two security personnel. Firewall rules that will affect more packets than others will be placed sooner in the logical flow of the firewall. Access to configuration interfaces for the firewalls will be restricted to direct wired connections on a usually empty port. This also means that physical access to the firewall will need to be controlled. Modern-day firewalls often come with other features, such as DHCP server capabilities; we will keep unnecessary services disabled. The rules in tables 2 and 3 aren't exact and will need adjustments when the full context of our network comes in, like with blocking certain IP groups, such as private IPs coming in externally.

**Table 2**

*Pseudo rule set for external firewall*

| Rule # | Source Address | Source Port | Destination Address | Destination Port | Action |
|---|---|---|---|---|---|
| 1 | External | Any | Internal | 67 | Block |
| 2 | External | Any | Internal | 68 | Block |
| 3 | External | Any | Internal | 21 | Block |
| 4 | External | Any | Internal | 23 | Block |
| 5 | External | Any | Internal | 161 | Block |
| 6 | External | Any | Internal | 7 | Block |
| 7 | Not [dedicated port] | Any | [Firewall IP] | 443 | Block |
| 8 | Any | Any | Any | Any | Allow |

**Table 3**

*Pseudo rule set for internal firewall*

| Rule # | Source Address | Source Port | Destination Address | Destination Port | Action |
|---|---|---|---|---|---|
| 1 | Internal | Any | External | 88 | Allow |
| 2 | Internal | Any | External | 443 | Allow |
| 3 | Internal | Any | External | 80 | Allow |
| 4 | Internal | Any | External | 53 | Allow |
| 5 | Internal | Any | External | 123 | Allow |
| 6 | [Classroom Subnet] | Any | [Labs Subnet] | [printer port #] | Allow |
| 7 | [dedicated port] | Any | [Firewall IP] | 443 | Allow |
| 8 | Any | Any | Any | Any | Block |

# Intrusion detection systems policies and implementation (vi)

Intrusions are when an attacker tries to gain access to an information system or otherwise disrupt the system. An Intrusion Detection System (IDS) will attempt to identify possible intrusions and generate alerts for security staff to investigate. There are three main types of IDS: those that run on hosts called host-based IDS (HIDS), those that run on network equipment called network-based IDS (NIDS), and those that incorporate data from both hosts and network devices called a hybrid IDS. A step beyond just detecting an intrusion would be actively preventing it, and that is just what an intrusion detection and prevention system (IDPS) does. Both IDS and IDPS suffer from issues when they are inaccurate, either misassigning benign traffic as malicious or letting through malicious traffic without detecting it. Both are also not able to effectively catch unknown attacks or zero-day attacks. There are also three main approaches to detecting intrusions: signature-based detection, anomaly-based detection, and stateful protocol analysis. Signature-based detection compares traffic to known attack patterns, or signatures, because many attacks have clear distinctions from normal traffic. Anomaly-based detection instead compares current traffic to a baseline, known-good traffic pattern, which is good for catching abnormally high traffic flow. Lastly, stateful protocol analysis uses the firewall's state table of protocol connections to compare them with known-good, benign protocol profiles.

The external firewall will be running an IDS, while the internal firewall will run an IDPS. The IDS running on the external firewall will give us good metrics on how well our other defenses perform collectively. The IDS and IDPS are similar to firewalls in needing to run continuously to be effective. Because of the computing overhead imposed by HIDS, as a cost-saving measure, they can be optionally installed on workstations that have plenty of computing power. Information from the IDS and IDPS can be used to improve firewall rules. If a large number of attempted intrusions are coming from a particular port that we are not making use of. Also similar to firewalls, changes should be tested outside of business hours, and a team should be ready to respond to reports of false positives on benign traffic. Audit logs will be maintained in the NAS in a set maximum size that allows for at least 6 months of records, where the oldest records are discarded to make room for new ones. Because the traffic pattern can vary by semester--

and spike when classes are doing something, signature-based detection would be preferred.

## Host hardening, including update policies and implementation (vii)

When managing many computers in an organization setting, it is important that they are kept up to date, running the same configurations, and hardened against attacks. The most important part of system hardening starts with a good base, the operating system. In our case, we are running Microsoft Windows and will need to make a fair few changes to its default configuration. The operating system also needs to be kept up to date, so we will need to use Active Directory to automatically update all the computers at the same time. Windows updates should, in general, be stable, but testing updates on one machine first would be required for stable operation. Another part of stable operation is not letting the runtime get too high, so computers should be restarted at a minimum of once a week. Several default services on the Windows computers will need to be turned off, like Remote Desktop and drive and printer sharing services. Microsoft Defender Antivirus should also be sufficient in keeping end hosts secure from known malware. The end hosts will also follow the group policy from the AD domain controllers. If we cannot afford to have a full-on penetration testing team, which I’m guessing we cannot, then at least the most impactful and widespread attacks similar to EternalBlue should be tested.

# Security for software/applications (viii)

Any software and applications used in Constant Hall services, whether managing residents, network services, intranet portals, or other administrative and user-facing activities, constitute a significant attack surface when not properly secured. Application security is the protection of the confidentiality, integrity, and availability of data and services in an application. The measures below provide a broad program encompassing secure development, configuration standards, and continuity lifecycle of the application-security program for a system's lifecycle.

**Core Software/Application Security Tactics**

1. Implement secure development practices (Secure SDLC) as the end-to-end practices.
    a. Now, all new applications or major enhancements should have a process with stages that are part of an industry-standard development project lifecycle that encompasses the phases of requirements, design, implementation, testing, deployment, and maintenance. Security needs to be built into every phase, not attached at the end. Such a "security-by-design" strategy prevents vulnerabilities that could be built into the system.
    b. Developers themselves need to be trained in secure coding techniques (input validation and error handling, secure session and credential handling, data protection) to decrease the risk of such injection, authentication, or unsafe errors that could leak sensitive information.
2. Adhere to the Standards You Know How to Use (like OWASP/ASVS):
    a. Ensure the underlying security requirements on applications are satisfied by using the OWASP (OWASP/ASVS) framework. The ASVS framework is designed to systematically evaluate authentication methods, session management and data protection security protocols, input validation, and other security controls.
    b. Use the core secure-coding checklist that covers proper input validation and encoding output and safe session and credential handling, as well as secure file and database access, cryptographic practices, and safe error handling and logging.
3. Secure Configuration and Hardening:

   a. Application servers, web servers, and database servers must be securely configured, turn down unnecessary services, change default credentials, and apply security patches promptly.
   b. Apply the "least privilege" approach: applications, services, and user accounts should run only the required minimum permissions while using applications, services, and user accounts. This will help in reducing the potential damage when an application or account is compromised.
4. Input Validation, Output Encoding, and Data Protection:
   a. The input needs to be very strict (whitelist expected values), and output encoding is mandatory, to protect against injection attacks (SQL injection, script injection and/or cross-site scripting).
   b. Applications have sensitive data (resident data, passwords, data in the cloud), and for that information to be available and protected, it should be encrypted in the storage and communication environment (via HTTPS or through secure protocols) to avoid being disclosed or intercepted.
5. Continual Testing and Vulnerability Review:
   a. Carry out security testing, including static analysis (code review/SAST), dynamic testing (penetration testing/DAST), dependency scanning, and configuration auditing before deployment and before major changes.
   b. Regularly schedule assessments of applications in production: regular vulnerability scanning, patching cycles, and routine code audits to identify new risks and resolve vulnerabilities promptly.
6. Session Management and Access Control:
   a. Use strong authentication and authorization mechanisms: ensure strong password or credential policies, secure session handling, role-based access restrictions on sensitive functionality, and do not implement hard-coded credentials to maintain application code or configuration.
   b. Sessions and cookies should be protected (secure and HTTP-only flags, timeouts, inactivity logout) to mitigate hijacking or session-based attacks.
7. Logging, Monitoring, and Incident Response Integration:
   a. Enforce logging in the applications: log accesses, failed logins, configuration changes, data access or modification, errors, and suspicious behaviour. These logs are passed through into wider monitoring and incident-response processes for Constant Hall, so that

a security incident can be swiftly discovered and investigated if it exists.

   b. Keep logs secure: logs should themselves be secured from unauthorized access or tampering, and sensitive data should be redacted or encrypted as necessary.

8. Secure Third-Party Components and Dependencies:
   a. If you plan a project on third-party libraries, frameworks, or services, keep it in the correct state of practice and address/mitigate known vulnerabilities. Do not add unapproved, deprecated, or insecure components.
   b. Cloud or SaaS/PaaS solutions ensure secure credential/key management, encryption, least privilege, and minimum privilege access are used, and only services that comply with institutional security requirements.
9. Plan for Implementation for Constant Hall Applications:
   a. Ensure all in-house developed or maintained applications employ a Secure SDLC process: document requirements, threat modeling early, provide for secure coding and design reviews, enforce structured peer code reviews, and mandate security testing before deployment.
   b. Use OWASP ASVS (Level 2) as the baseline for application security verification for any application dealing with resident data, administration, or network services.
   c. Harden servers and databases: lock down configurations, disable default accounts, enforce strong credential management, and patch regularly.
   d. Use encryption when it comes to sensitive data: enforce encryption at rest and in transit (databases, storage, and network using TLS).
   e. Implement a schedule of application-security testing: static and dynamic testing before release, periodic scanning of production software, audits of dependencies, and periodic patching.
   f. Develop rigorous access controls in apps: role-based permissions, session management, secure authentication, and authorization.
   g. Implement consistent logging processes for application events and tie the logs into a broader incident-response and monitoring system in Constant Hall.
   h. Audit third-party content that is hosted in your environment and make sure services used within the

environment follow appropriate and secure procurement or vetting procedures in line with your policies and plans.

**Why is this Important for Constant Hall?**

Residence halls, like Constant Hall, hold a lot of very sensitive staff (student identity, contact information, possibly some personal finance or academic information) information and may operate web and/or network apps in support of services/account management or for the management (facility) of a campus. Any possible vulnerability in that application, an injection problem, an unauthorized authentication, insecure data storage, or an improper configuration, can cause a data loss that could then be exploited for unauthorized access, system loss, or worse. With all of this involved and in high stakes (personal privacy, regulatory compliance, resident trust, service availability), using a strong application security program is imperative.

Constant Hall can greatly reduce its software-related risk with usability and function by infusing security throughout the application development lifecycle, implementing hardened deployments, implementing strong access controls, adding an encryption layer, and conducting continuous testing and monitoring. This defense-in-depth strategy adheres to good practices in the industry and ensures that software and application security become a fundamental part of a security architecture rather than an afterthought.

# Data protection measures (ix)

Protecting data sensitive to Constant Hall, including personal information, resident information, configuration files, logs, and institutional records, is a cornerstone of the security program. The data protection measures presented in this paper will use the best practices from the industry and standards, which will also be applied at the network, storage, access, and backup layers to ensure confidentiality, integrity, and availability of the data.

**Important Data Protection Strategies**

- Performing Data Classification and Least Privilege Access. Data should be classified based on sensitivity before applying technical controls (e.g., public, internal, sensitive/resident personal data, critical infrastructure data). This classification dictates which protections are necessary.
- When dealing with sensitive or personal data, access would follow the principle of least privilege: only those participants whose roles require the data could have access, and only the minimum permissions we need to do so (read vs. write, say).
- Encryption at rest and in transit. All sensitive data that is stored on servers, shared hard drives, and backups must be encrypted at rest. Similarly, moving data across the network (e.g., endpoints from a computer to servers or between subnets) should use encrypted protocols (e.g., TLS/HTTPS, secure VPN, SSH, SFTP) that cannot be intercepted.
    - For campus computers: those that store or access sensitive data, disk or full device encryption should be enabled.
    - Encryption keys must be safe: they need to be kept secure, rotated periodically as needed, and granted only to authorized users.

**Strict data handling policies and secure storage**

- The infrastructure of data storage should adhere to secure-storage standards, as described in storage security advice, for example, ensuring stored data remains safe from unauthorized access, tampering, or theft.
- Policy on data handling should also be defined and enforced: data should be securely deleted when it is no

longer required; temporary files or backups should follow retention and disposal rules to avoid leaking this data.

- Data minimization is critical as well, saving and collecting data only for use by the institution, avoiding unwanted information about individuals or sensitive data.
- Routine backup and storage from an area outside the office.
- For data availability and recovery, in the event of hardware failure, ransomware, accidental deletion, or disasters (and more), backups should be scheduled with regularity.
- Strong strategies such as 3-2-1 (3 total copies on at least 2 different media types plus 1 copy off-site) are highly recommended.
- Backup copies should themselves be encrypted, and if off-site or in cloud storage, encrypted robustly and authenticated securely.
- A regular testing run of backups (restore trials) guarantees stable backup recovery in instances of data loss.

**Access control and monitoring, auditing and logging**

- All access to sensitive data, success or failure, should be recorded. Access log files, modification records, modification logs, and backup/restore events should be stored for a specified period (with periodic review to identify any irregularities, violations, or unauthorized access and re-access).
- Alerts can be set up to raise red flags (such as unusually large data flows, repeated failed logins, or access outside of normal hours) and to allow for early warning and remediation.

**Secure data disposal and crypto-erases as and when needed:**

- Secure deletion can be used when data is no longer required (when old records, logs, backups, or storage devices are being retired)
- The trusted method is “crypto-shredding,” which involves securely deleting the encryption key and making data unreadable and irrecoverable for encrypted data.
- Physical media (old hard disk drives, backup tapes) should be thoroughly sanitized or destroyed in accordance with disposal policy, especially if they contain sensitive personal or institutional information.

**Regular security review and compliance check-up:**

- Regularly scheduled safety and security audits, vulnerability assessments, and examination of data protection policies mean that the way controls have been designed or put in place remains effective as well as dynamic in response to new threats, evolving regulations, and institutional needs, an extremely valid criterion for any institution's environment.
- This is particularly important in a higher-education residence hall environment, as seen here.

**Integration with Existing Controls:**

- These controls are complemented by the access control and firewall policies outlined in the security measures.
- Things that often seem to matter to Constant Hall:
    - Encryption schemes, for instance, work hand in hand with access controls when it comes to unauthorized reads; backups serve as a hard belt against firewalls or access controls being bypassed, and in case of failures, logs produced by access control systems or network devices can be used to feed into monitoring and incident response procedures.
- When you take the combined efforts of these systems, they create a holistic approach towards the defense and maintenance of an organization's data lifecycle, one that reduces chances of data breaches, corruption, or loss while at the same time providing data that must be available and ready for its restoration whenever needed.

**Why They Matter to Constant Hall**

- Given that Constant Hall likely handles sensitive resident information (student identities, contact information, possibly financial or academic data), configuration data for network infrastructure, system logs, and backups, implementing strong data protection is not optional; it is essential.
- The massive amounts of personal data at universities and residence halls make these sites an increasingly high-profile target: in reality, they can suffer from high turnover.
- Moreover, data protection goes beyond protection against external threats such as hackers, malware, and data thieves; it also eliminates risk from insider misuse, accidental deletion, and hardware failure.
- The institution protects the privacy of residents and maintains compliance with and supports the service with the

right backups, encryption, and other data access controls in place to provide continuity of operation.

Adopting the extensive data protection program, such as data classification, encryption at rest and in transit, secure data storage, backup and offsite copies, access auditing and logging, secure disposal of data, and security monitoring by Constant Hall will ensure well-preserved sensitive data throughout its whole life. These measures, along with access control, firewall, incident response, and disaster recovery components, complete an integrated multi-layered security architecture that builds upon a comprehensive layered security architecture to ensure the confidentiality, integrity, availability, and resilience of resident and institutional data.

# Risk assessment with updated controls and cost-benefit analysis (x)

**Risk Assessment**

- **Unauthorized external intrusion**
  - Before controls: Medium to High likelihood, High impact (data breach, network compromise) - High risk.
  - After controls: Firewall and network segmentation prevent unauthorized access. Logging and monitoring aid in detection, and access controls limit credential misuse.
  - Residual risk: Low to Medium.
- **Internal misuse or privilege abuse (insider threat/students/staff misuse)**
  - Before controls: Medium likelihood, Medium-High impact (data leakage, resource misuse) - Medium risk.
  - After controls: RBAC, strict permissions, regular audits, and immediate revocation of access once an account change/departure occurs.
  - Residual risk: Low to Medium.
- **Malware or ransomware spreading across the network**
  - Before controls: Medium likelihood, High impact (system infection, data loss) - High risk.
  - After controls: Network segmentation and VLAN isolation prevent lateral spread, firewall, NGFW restricts numerous types of malicious traffic, and host permissions restrict installations to limit execution rights.
  - Residual risk: Medium (there is retained risk, but impact and likelihood are reduced).
- **Misconfiguration or overly broad firewall/network rulesets**
  - Before controls: Medium likelihood (especially over time as rules change), High impact (accidental exposure, unauthorized traffic) - Medium-High risk.
  - After controls: Documented firewall ruleset, rule justification, regular review and audits, restrictions to admin access to firewall management.
  - Residual risk: Low to Medium (depends upon review diligence).
- **Data exfiltration from internal hosts to outside (sensitive data leak)**
  - Before controls: Medium likelihood, High impact (data breach, compliance/legal issues) - High risk.

- After controls: Monitoring of outgoing traffic, logging and alerting, egress filtering via firewall, and strict access controls will restrict access to sensitive data.
- Residual risk: Medium (less likely and more detectable).

**Cost-Benefit Analysis**

- **Costs:**
  - Purchase and deployment of a firewall or next-generation firewall (NGFW).
  - Configuration of VLANs/subnets, network segmentation and firewall rules.
  - Administrative overhead: network admin time for access control management, auditing, log review, and updates.
  - Possible user inconvenience or friction: stricter authentication, restricted access, and limited permissions may limit the use of services.
  - Training requirements for staff/admins to perform control operations correctly and for users to understand policy.

- **Benefits**
  - Reduced risk of both external/internal intrusions, malware propagation, data theft, and network compromise. Segmentation and access control are used to contain any breach of one host to prevent widespread damage.
  - Real-time monitoring and logging to identify and respond faster.
  - Safeguarding sensitive data (personal, administrative or institutional), minimizing potential legal/compliance challenges, reputational damage, and recovery costs.
  - More structured, manageable authorization model, easy to audit, manage, and enforce security policies.
  - Long-term return on investment: By avoiding even a single significant breach, you could save far more than the costs of controls in downtime, remediation, and reputational damage.

The proposed access control and firewall for Constant Hall enhances the security posture significantly and provides a significant cost benefit. By significantly decreasing the chance

and potential consequences of external intrusions and intentional abuse, such controls reduce the organization's susceptibility to high-stakes data breaches, which, recent reports have shown, cost organizations, on average, millions of dollars per incident. Additionally, firewalls and network segmentation inhibit the spread of malware or ransomware by preventing one compromised host from compromising the entire network. The costs on an ongoing basis, hardware or NGFW purchase, network configuration, administrator time, training of users, are small compared to the potential damage that could come from a breach, financial losses, lost time, damage to reputation, plus regulatory penalties. All this reinforces why the security investment makes sense: The controls offer "insurance" against relatively rare but extremely costly events and bring long-term benefit through reduced risk, simplified auditing, and greater compliance. That said, residual risk does exist, primarily from advanced threats (zero-day malware, social engineering, endpoint compromise), configuration errors, and an inability to maintain and monitor the system over time. This being the case, the protection value of these controls requires a consistent stance, frequent rule checks, log monitoring and user training in order that the process be followed. On the whole, this proposed security architecture can be considered a strong, inexpensive-level security that supports stable protection of Constant Hall, but for it to work, it will need to be closely monitored to be effective.

# Incident response plans, disaster recovery plans (xi)

Even with robust preventive controls (access control and firewall), organizations should know that some incidents, attacks on the system from malicious actors, malware, insider misuse, hardware failures, and disasters, can remain. An official Incident Response Plan and Disaster Recovery Plan must be in place for swift response to react to and restore service, and minimize the severity.

**Components of an Incident Response Plan:**

1. Preparation: create rules on procedures, gather a Computer Security Incident Response Team and assign roles (incident coordinator, communications lead, technical responders, legal/HR if required), train staff, monitor/logging architecture.
2. Detection & Analysis: Ongoing monitoring (firewall logs, network intrusion detection systems, host-based detection), alerting and incident type classification according to severity and impact.
3. Containment, Eradication & Recovery: isolation of all impacted systems/subnets with revoked credentials, removing malicious software or malware, restoring system access to clean backups, ensuring integrity, applying patches and remedies.
4. Post-Incident Actions: Maintain a written and recorded timeline of the incident, root cause identification, response plan, downtime/impact assessment, results, security policies and controls revision, firewall or access rule adjustment, if necessary, and staff retraining.
5. Communication Plan: Define who is notified in the case of an incident (IT staff, building management, impacted users, institutional leadership) and if needed, external notifications (law enforcement, compliance reporting).
6. Test and Maintenance: regularly run tabletop exercises and live simulations, alter the plan based on changes in infrastructure, threat landscape or post-mortem events.

**Disaster Recovery Plan Elements Key Components:**

1. Risk & threat analysis in advance: Determine potential disaster scenarios (hardware failures, network outages, fire/flood, ransomware attack, data corruption).
2. Define Recovery Objectives: Set Recovery Time Objectives (RTOs) - the time taken to restore systems, and Recovery Point Objectives (RPOs) - acceptable data loss time.
3. Data Backup & Alternate Resources: Regular backup of key data (user data, system configurations), preferably including offline/offsite storage or a separate network segment.
4. Prepared with infrastructure for backup or alternate systems if the main sources fail.
5. Restoration Procedures: Note the process of recovery, including restoration of servers, establishment of network settings, firewall rules, user access, network segmentation, etc.
6. Testing & Maintenance: Routinely test restoration processes (quarterly or semi-annual disaster recovery), verify backups, and update DRP with infrastructural changes.
7. Communication & Escalation: Set policy on who to inform of a disaster event (building ops, IT management, users, campus administration), roles in decision-making, external contacts if needed (utilities, service providers).

By combining a formal IRP and DRP, which comes with an integrated firewall and access control policy system, Constant Hall will stand ready to carry out defense-in-depth measures: prevent, detect, respond, and recover.

**Next Steps / Implementation Suggestions:**

1. Finalize the Access Control Policy and the Firewall Policy.
2. Map roles and permissions, network zones.
3. Implement firewall (or NGFW) and VLAN/subnet setup, adopt RBAC and strong authentication setup, configure logging/monitoring infrastructure, and assign security administrator(s).
4. Create the Incident Response Plan and Disaster Recovery Plan documentation, describe the incident response team, work responsibilities, and means of communication.
5. Regular audits, testing (firewall rule review, IRP tabletop exercises, DRP drills), and access reviews should be scheduled.
6. Educate residents and staff of the building and relevant stakeholders on the policies, responsibilities, and response in the event of incidents.

A comprehensive IRP and DRP really increases Constant Hall's resilience because when something does go wrong, whether that's a cyberattack, hardware failure, natural disaster or human error, the operations, data and services of the building can be quickly restored with the least potential for damage. An IRP provides defined methods to detect, contain, and eliminate threats, which decrease the duration and scope of security events and reduce costs and damage reputation. A DRP coupled with other DRP supports key systems and data that are backed up and can be recovered, reducing downtime, data loss and business impact. As a pair, these plans enable Constant Hall to respond to and clean up incidents, and to return to normalized service in an orderly, efficient manner, protecting residents' data, preserving trust and keeping things going even when unexpected events happen.

# Bibliography

Cisco. (2024a, June 4). *Cisco Catalyst 9400 Series Switch Data Sheet*. Cisco. https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9400-series-switches/nb-06-cat9400-ser-data-sheet-cte-en.html

Cisco. (2024b, August 9). *Cisco Catalyst 8500 Series Edge Platforms Data Sheet*. Cisco. https://www.cisco.com/c/en/us/products/collateral/routers/catalyst-8500-series-edge-platforms/datasheet-c78-744089.html?oid=dstetr023042

Cisco. (2024c, September 4). *Cisco Secure Firewall 3100 Series Data Sheet*. Cisco. https://www.cisco.com/c/en/us/products/collateral/security/firewalls/secure-firewall-3100-series-ds.html

Cisco. (2024d, October 9). *Cisco Catalyst 9300 Series Switches*. Cisco. https://www.cisco.com/c/en/us/products/collateral/switches/catalyst-9300-series-switches/nb-06-cat9300-ser-data-sheet-cte-en.html

Cisco. (2024e, October 9). *Cisco Secure Firewall 4200 Series Data Sheet*. Cisco.

https://www.cisco.com/c/en/us/products/collateral/security/firewalls/secure-firewall-4200-ds.html

Microsoft. (2021, July 29). *Kerberos Authentication Overview*. Learn.microsoft.com. https://learn.microsoft.com/en-us/windows-server/security/kerberos/kerberos-authentication-overview

Stallings, W., & Brown, L. (2023). *Computer Security: Principles and Practice* (5th ed.). Pearson Higher Ed.

Whitman, M. E., & Mattord, H. J. (2022). *Principles of Information Security* (7th ed.). Cengage Learning.

# Software/Appliance Specifications

**Cisco Catalyst 9400 Series Switch**

- Modular chassis supporting supervisor redundancy
- Up to 9 Tbps system bandwidth
- Supports 1G/10G/25G/40G/100G uplinks
- Advanced Layer 3 routing (OSPF, EIGRP, BGP)
- MACsec (802.1AE) hardware encryption
- StackWise Virtual for multi-chassis redundancy
- Nonstop Forwarding and Stateful Switchover
- Up to 384 access ports per chassis
- Hot-swappable power supplies and fans
- TrustSec and SGT network segmentation
- Flexible NetFlow and telemetry support

**Cisco Catalyst 9300 Series Switch**

- Up to 1 Tbps stacking bandwidth
- Access layer switching with 1G/2.5G/5G/10G Multigigabit ports
- MACsec hardware encryption
- Layer 3 routing with OSPF, EIGRP, BGP
- Stackable up to 8 switches
- Hot-swappable dual power supplies
- Deep buffer options for high-density user connectivity

- TrustSec segmentation
- Full NetFlow, telemetry, and QoS features

**Cisco Catalyst 8500 Series Edge Router**

- High-speed WAN routing up to 20–100 Gbps
- Hardware crypto acceleration for:
    - IPsec VPN
    - TLS/SSL encryption
- SD-WAN support with segmentation and traffic steering
- High-density 40G/100G interfaces
- Advanced routing protocols (BGP, OSPF, EIGRP, IS-IS)
- MPLS support
- AppQoS and application visibility
- Redundant power and modular line-card options

**Cisco Secure Firewall 4200 Series**

- Up to 1.2 Tbps firewall throughput
- Up to 100 Gbps threat inspection
- Full Firepower Threat Defense (FTD) feature set
- Advanced Intrusion Prevention System (IPS)
- Encrypted Traffic Analytics (ETA)
- IPsec/SSL VPN termination
- Application and identity-based access control

- Clustering and high-availability failover
- Hardware acceleration for encrypted traffic inspection
- Multicontext support for segmentation

**Cisco Secure Firewall 3100 Series**

- Up to 45 Gbps firewall throughput
- FTD intrusion prevention capabilities
- Hardware-accelerated IPsec/TLS decryption and inspection
- Application visibility and control
- Identity-based access policies
- Trustworthy computing module (TPM) for secure key storage
- Active/standby failover support
- Advanced malware protection (AMP)
- VLAN and network segmentation features
- Full VPN support for internal encrypted traffic

**Windows Server 2025**

- Active Directory Domain Services with strengthened security baselines
- Kerberos authentication using AES256 encryption
- Kerberos FAST (Flexible Authentication Secure Tunneling) support
- SMB 3.1.1 with encryption and mandatory signing
- TLS 1.3 enforcement for IIS and system services

- Virtualization-based security, secure boot, and TPM integration
- File and Storage Services with Storage Spaces redundancy
- DNS role with DNSSEC capability
- Group Policy with advanced security templates
- Windows Defender with next-gen malware protection