

Noroff

School of Technology and Digital Media

LNXE Corp – Exam Project 04

Secure Virtualized Infrastructure Design and Implementation for LNXE Corp

Tommy André Slettvoll
[Project04-LNXECorp / Project Exam 1]
Submission Date: 4 March 2026

Table of Contents

1. Introduction	3
2. Main	5
2.1 Network Architecture & Topology Design	5
2.2 IP Addressing Scheme	7
2.3 Virtualisation Architecture (VMware Design)	8
2.4 Server Operating System Selection & Justification	9
2.5 Web Server Implementation (Nginx)	10
2.6 Backup Server Implementation (rsync strategy)	11
2.7 Authentication & Access Control Design	12
2.8 Firewall & Network Security Configuration	13
2.8.1 Stateful Firewall Concept	14
2.8.2 Port Exposure Analysis	15
2.8.3 Internal vs External Filtering	15
2.8.4 Logging and Monitoring	16
2.9 Wireless Network Security Implementation	16
2.10 OS Hardening Measures	17
2.10.1 Minimal Installation and Service Reduction	17
2.10.2 SSH Hardening Configuration	18
2.10.3 User Privilege Management	18
2.10.4 File Permissions and Ownership	18
2.10.5 System Updates and Patch Management	19
2.10.6 Firewall Integration as Hardening	19
2.10.7 Log Monitoring and Security Awareness	19
2.10.8 Hardening Philosophy and Industry Alignment	19
2.11 Administrative Access Controls	20
2.12 Security Best Practices Alignment	20
2.13 Risk Analysis & Threat Considerations	20
2.13.1 External Threats	21
2.13.2 Internal Threats and Lateral Movement	21
2.13.3 Data Loss and System Failure	22
2.13.4 Service Misconfiguration and Vulnerability Exposure	22
2.13.5 Wireless Network Risks	23
2.13.6 Residual Risk and Limitations	23
2.14 Design Decisions & Trade-Off Analysis	24
2.14.1 Choice of Type 2 Hypervisor	24
2.14.2 Internal-Only Network Design	25
2.14.3 Backup Strategy Selection	25
2.14.3.1 Backup Server Architecture	26
2.14.3.2 Secure Backup Communication	27
2.14.3.3 Backup Execution	28
2.14.3.4 Restore Validation	29
2.14.4 Separation of Web and Backup Roles	29
2.15 Defence-in-Depth Model Applied to LNXE Corp	29
2.15.1 Physical / Host Layer	30
2.15.2 Network Layer Security	30
2.15.3 Host / Operating System Layer	31
2.15.4 Application Layer Security	31
2.15.5 Administrative & Human Layer	31
2.15.6 Layer Interaction & Redundancy	32
2.16 Business Continuity & Disaster Recovery Considerations	32
2.16.1 Recovery Time Objective (RTO)	33
2.16.2 Recovery Point Objective (RPO)	33
2.16.3 Backup vs Redundancy	33
2.16.4 Ransomware & Integrity Protection	34
2.16.5 Testing & Validation of Recovery	34
2.16.6 Scalability & Future Continuity Planning	35
2.16.7 Risk-Aware Resilience Design	35
3. Practical (Testing & Results)	36
3.1 Network Connectivity Validation	37
3.2 SSH Authentication Security Testing	37
3.3 Backup & Restore Verification	39
3.4 Firewall Validation	41
3.5 Wireless Security Testing	42
3.6 Security Validation Summary	43
4. Summary & Conclusion	45
5. References	46
6. AI Usage Disclosure	48
Tools Used	48
Purpose of AI Usage	48
Example of AI Usage	48

1. Introduction

This project, developed for LNXE Corp (a fictional start-up), demonstrates the design and implementation of a secure server-client architecture. It aims to showcase practical expertise in virtualisation, networking, operating systems, and security. The project features a realistic enterprise IT environment built in accordance with current best practices, with an emphasis on security by design to ensure that protection mechanisms are integrated seamlessly from the outset.

The company LNXE Corp is in its growth phase, with about 10 employees. The employees work primarily in software development, and thus require a reliable, secure and scalable IT environment. The development team would depend heavily on having access to internal services, secure remote administration and even protected data storage. As LNXE Corp is expected to grow, the infrastructure must be able to adapt and be designed to support future expansion without compromising any features, such as security or performance.

The Chief Executive Officer (CEO) of LNXE Corp has requested that this infrastructure be implemented using a server-client model, with a strong focus on security, OS hardening, and adherence to recognised best IT Security Practices.

To meet these requirements, the solution for this project will be deployed using VMware Workstation Pro (25H2), which will host multiple virtual machines simulating a realistic enterprise environment. Virtualisation is used in enterprise environments because it efficiently utilises hardware resources, improves scalability, and enhances security through system isolation (VMware, 2023). By using a Type 2 Hypervisor, this project will simulate a realistic enterprise deployment, while still remaining suitable for an academic environment. VMware Workstation Pro (25H2) provides the flexibility to deploy, configure, and test multiple operating systems simultaneously.

The objective is to demonstrate competence in virtualisation, networking, operating system administration, and security across multiple technical domains.

This project focuses on maintaining secure server and client configurations, ensuring controlled network access, and safeguarding sensitive data through strong authentication mechanisms and regular backups.

The infrastructure will consist of:

- ❖ One Ubuntu Server, Acting as the Web Server
- ❖ One Ubuntu Server, Acting as the Backup Server
- ❖ Linux Client and Windows 11 Client (for validation and testing)

- ❖ A segmented internal network with both wired and wireless access

Ubuntu was chosen as a server for its long-term support and high stability. It also has a strong security track record, especially in enterprise environments (Canonical, 2024).

Using both Linux and Windows clients will enable cross-platform testing, as mixed operating systems are very common in modern organisations. (Odom 2020)

As mentioned earlier, the core design principle will focus on security rather than treating it as an afterthought. The project will also use the principles of least privilege, secure authentication, hardened operating systems, firewall rules, and controlled network access. These measures ensure users and services receive only the access they require to perform their tasks (Saltzer and Schroeder, 1975).

Very secure authentication mechanisms are implemented, including SSH key-based authentication instead of traditional password-based logins, which significantly reduces the risk of brute-force attacks that lead to credential compromise (NIST, 2020). Direct root access to the servers is disabled, and administrative tasks are performed through controlled privilege escalation. While we now have access control, the project will add operating system hardening, firewall configurations, and controlled network communication to further reduce the attack surface on deployed systems. Preventing direct exposure of servers to external networks, Using Internal IP addressing, and network segmentation to separate different types of devices. Many of these security measures are a widely accepted industry standard and follow accepted security best practices for secure infrastructure designs (CIS, 2023)

To meet the assignment requirements, the project demonstrates:

- ❖ A secure network implementation with an internal IP address protection.
- ❖ A documented network topology using Cisco Packet Tracer.
- ❖ Secure server configuration and operating system hardening.
- ❖ Encrypted remote access using SSH keys.
- ❖ A functional backup solution with verified restore capabilities.
- ❖ A proof of concept demonstration that is supported by screenshots and testing results

This project covers knowledge and skills from many modules in the Network and Security Administration program at Noroff. These include Linux (LNX), Networking Infrastructure (NET), Operating and File Systems (OFS), Microsoft Server Technologies (MST), and Introduction to Information Security (IIS). By combining these, the project shows how these skills integrate in a real-world IT environment.

2. Main

2.1 Network Architecture & Topology Design

The network infrastructure for this LNXE Corp project will balance security, usability, and scalability. It will support both server and client systems. The design's core is to protect internal resources from unauthorised access and ensure reliable connections for employees and administrators.

To achieve this, the network is implemented as a private internal network with clear roles, access paths, and security boundaries.

A layered network architecture will be used for all servers and client systems. This ensures that internal resources are not exposed to the external network by separating servers, client devices, and wireless access into logical segments. This approach is a well-established enterprise networking practice. (Network Segmentation: Why it's Important, How to Implement It and How it Improves Cyber Security, 2024) Segmentation is used to reduce the attack surface, limit lateral movement, and simplify access control and monitoring (Cisco, 2023). All internal systems will operate using private IP addressing.

The network consists of:

- ❖ A router manages traffic between network segments and acts as the default gateway for internal devices.
- ❖ A switch provides wired connections for servers and desktop clients.
- ❖ A wireless access point (WAP) provides a secure wireless connection for both laptops and mobile devices.
- ❖ A dedicated server system runs both the Web Server and the Backup Server.
- ❖ A client system represents employee workstations and laptops.

The router is the central control point for network traffic. It enforces routing policies and, in later project stages, will restrict management access to authorised administrative systems. In a real Enterprise/IT environment, the router handles Network Address Translation (NAT, which allows multiple devices on a private network to access external networks using a single public IP address), firewalling (blocking or allowing traffic based on security rules), and logging (recording network activity for monitoring and troubleshooting). Although this project is simulated, its logical design reflects real enterprise network architecture (Odom, 2020).

The switch is the primary distribution layer. It connects all the wired devices in the internal network. This means that both the server system and the client systems are directly connected to the switch via Ethernet. This ensures stable, high-bandwidth communication.

Stable communication among critical infrastructure components, especially between the web server, the backup server, and client systems used for testing and validation, is essential. Centralising this wired connection through a switch supports future implementation of more security controls, such as VLAN segmentation and port-based access restrictions.

The complete network topology using Cisco Packet Tracer is presented in Figure 1 below. The design reflects logical segmentation between the server infrastructure, routing layer, switching layer, and wireless access components.




Figure 1 – Logical Network Topology for LNXE Corp (Cisco Packet Tracer)

As shown in Figure 1, the switch (LNXE-SW01) acts as the central distribution point for wired devices. The router (LNXE-RTR01) manages routing policies and, in a real enterprise environment, would enforce NAT and perimeter security controls. The wireless access point (LNXE-WIFI01) provides encrypted wireless access while remaining logically connected to the internal LAN.

In this simulated environment, routing is represented conceptually within Packet Tracer, while VMware host-only networking provides the practical implementation of internal network segmentation.

2.2 IP Addressing Scheme

All internal systems operate using private IP addressing in accordance with RFC 1918 (IETF, 1996). The use of private addressing prevents direct exposure of internal systems to external networks and reduces the overall attack surface.

Internal network segmentation was implemented to logically separate servers and client systems while maintaining controlled routing via a central gateway.

The router functions as the default gateway for the internal network and manages routing between segments.

The following addressing scheme was implemented:

<i>Device</i>	<i>Role</i>	<i>IP Address</i>	<i>Network</i>
LNXE-WEB01	Web Server	<i>192.168.10.10</i>	LAN
LNXE-BAK01	Backup Server	<i>192.168.10.20</i>	LAN
LNXE-CLI01	Linux Client	<i>192.168.10.30</i>	LAN
WIN11CLI01	Windows Client	<i>192.168.10.40</i>	LAN

This structured addressing scheme improves manageability, simplifies troubleshooting, and supports future scalability.

2.3 Virtualisation Architecture (VMware Design)

The LNXE Corp environment is hosted using VMware Workstation Pro (Type 2 hypervisor) to provide isolated, repeatable infrastructure suitable for testing and security validation (VMware, 2023).

A Type 2 hypervisor was selected for its practicality in an academic environment. It allows rapid deployment of multiple virtual machines on a single physical host while maintaining logical separation between server and client roles.





Figure 2 - VMware Workstation Pro showing the deployed virtual machines for the LNXE Corp environment.

To support segmentation and controlled outbound connectivity, two logical network types were implemented:

- ❖ Internal LAN (Host-only):
 - ❖ A private, isolated network used for all server-to-server communication and internal management traffic. This ensures that internal services are not directly exposed to external networks.
- ❖ NAT (Controlled outbound access):
 - ❖ Used only when necessary to allow systems to perform updates and access package repositories. This adapter is treated strictly as an egress path and is not used for internal service communication.



Name	Type	External Connection	Host Connection	DHCP	Subnet Address
Exam Project LAN	Host-only	-	Connected	-	192.168.10.0
Exam Project NAT	NAT	NAT	Connected	Enabled	192.168.40.0

Figure 3 - VMware virtual networks used for segmentation (Host-only LAN) and controlled outbound updates (NAT)

LNXE-WEB01 and LNXE-BAK01 were configured with dual network interfaces to maintain separation between internal communication and outbound update traffic.




Figure 4 - LNXE-WEB01 dual-NIC configuration with internal LAN access and NAT egress for updates.

Bridged networking was intentionally avoided because it would place systems directly on the physical network, increasing exposure and weakening the security justification of an internal corporate infrastructure. (Best Practices for Mitigating Risks in Virtualized Environments, 2015)

This design reflects enterprise-style segmentation principles while remaining appropriate for a controlled lab environment.

2.4 Server Operating System Selection & Justification

Ubuntu Server LTS was selected due to its long-term support lifecycle, stability, and strong security maintenance model (Canonical, 2024).

Ubuntu LTS provides:

- ❖ Extended support cycles
- ❖ Stable kernel releases
- ❖ Consistent security patching
- ❖ Wide enterprise adoption

Using both Linux and Windows client systems enables cross-platform validation and reflects the mixed operating-system environments commonly found in modern organisations.

2.5 Web Server Implementation (Nginx)

The Web Server (LNXE-WEB01) hosts internal web services to demonstrate server-client communication.

Nginx was selected as the web server software due to its lightweight architecture, high performance, and strong security (Nginx, 2024).

The web server was installed using the Advanced Package Tool:

sudo apt update

sudo apt install nginx

The installation completed successfully without errors.

The web service operates exclusively within the internal network and is not directly exposed to external networks.

```
slettvoildadmin@lnxe-web01:~$ sudo apt install nginx
[sudo] password for slettvoildadmin:
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
nginx is already the newest version (1.24.0-2ubuntu7.6).
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
slettvoildadmin@lnxe-web01:~$ sudo systemctl status nginx
● nginx.service - A high performance web server and a reverse proxy server
   Loaded: loaded (/usr/lib/systemd/system/nginx.service; enabled; preset: enabled)
   Active: active (running) since Tue 2026-02-24 15:20:24 UTC; 15min ago
     Docs: man:nginx(8)
   Process: 14051 ExecStartPre=/usr/sbin/nginx -t -q -g daemon on; master_process on; (code=exited, status=0/SUCCESS)
   Process: 14053 ExecStart=/usr/sbin/nginx -g daemon on; master_process on; (code=exited, status=0/SUCCESS)
   Main PID: 14082 (nginx)
    Tasks: 3 (limit: 4548)
   Memory: 2.4M (peak: 5.3M)
      CPU: 21ms
   CGroup: /system.slice/nginx.service
           └─14082 "nginx: master process /usr/sbin/nginx -g daemon on; master_process on;"
             └─14085 "nginx: worker process"
               └─14086 "nginx: worker process"

Feb 24 15:20:24 lnxe-web01 systemd[1]: Starting nginx.service - A high performance web server and a reverse proxy server...
Feb 24 15:20:24 lnxe-web01 systemd[1]: Started nginx.service - A high performance web server and a reverse proxy server.
slettvoildadmin@lnxe-web01:~$
```

Figure 5 - Nginx Install

Because the firewall was configured with a default-deny inbound policy, HTTP (TCP 80) and HTTPS (TCP 443) were explicitly permitted through UFW. All other inbound ports remain blocked.

This configuration supports the principle of least privilege and minimises unnecessary service exposure.

To validate functionality, the web service was accessed from an internal client:

http://192.168.10.10

(Validation screenshots are provided in Section 3)

The default Nginx welcome page was successfully displayed, while confirming:

- ❖ Network Connectivity
- ❖ Firewall rule functionality
- ❖ Web server service availability
- ❖ Correct client-server communication

2.6 Backup Server Implementation (rsync strategy)

To ensure operational continuity and data resilience, a dedicated backup server (LNXE-BAK01) was deployed within the internal network.

The backup solution is based on rsync over SSH, providing encrypted data transfer and efficient incremental synchronisation (Samba Project, 2024).

Rsync was selected due to:

- ❖ Reliability
- ❖ Low system overhead
- ❖ Incremental synchronisation capability
- ❖ Suitability for small-to-medium enterprise environments

(Rsync: Streamlining Data Transfer and Synchronization, 2025)

Backups are stored under:

/srv/backups/web01/

Access to the backup server is restricted using SSH key-based authentication combined with host-based firewall rules. Only **LNXE-WEB01** (192.168.10.10) is permitted to initiate SSH connections.

Detailed architectural design and validation of restoration are documented in Section 2.14.3.

2.7 Authentication & Access Control Design

Secure administrative access is a critical component of infrastructure protection.

Remote management is performed exclusively via SSH, using public-key authentication.

Password-based authentication was disabled to mitigate brute-force attacks and credential compromise. Cryptographic authentication mechanisms significantly reduce the risk of unauthorised access compared to password-based methods (NIST, 2020).

The SSH daemon configuration was modified in:

/etc/ssh/sshd_config

Key parameters:

PasswordAuthentication no

PermitRootLogin no

PubkeyAuthentication yes

AuthenticationMethods publickey

Administrative actions are performed using sudo with controlled privilege escalation, ensuring accountability and auditability.



```
# Example of overriding settings on a per-user basis
#Match User anoncvs
#    X11Forwarding no
#    AllowTcpForwarding no
#    PermitTTY no
#    ForceCommand cvs server
PasswordAuthentication no
KbdInteractiveAuthentication no
UsePAM no
PermitRootLogin no
PubkeyAuthentication yes
AuthenticationMethods publickey
```

Figure 6 – SSH configuration showing password authentication disabled and root login restricted.

To validate the configuration, both positive and negative testing scenarios were performed:

- ❖ Attempting login without a valid key (denied)
- ❖ Attempting password-only login (denied)
- ❖ Successful login using a valid SSH key (granted)

These tests confirm that interactive password authentication is disabled and that access is restricted to authorised key holders only. (*Validation screenshots are provided in Section 3*)

2.8 Firewall & Network Security Configuration

The infrastructure incorporates host-based firewall controls to regulate inbound and outbound network traffic.

Firewall rules were implemented using Uncomplicated Firewall (UFW) to:

- ❖ Allow only required services.
- ❖ Restrict administrative access
- ❖ Block unnecessary ports

Network segmentation was implemented using VMware virtual networking and host-based firewall enforcement.

The following configuration was applied on **LNXE-WEB01**:

- ❖ ***sudo ufw default deny incoming***
- ❖ ***sudo ufw default allow outgoing***
- ❖ ***sudo ufw allow 80/tcp***
- ❖ ***sudo ufw allow 443/tcp***
- ❖ ***sudo ufw allow from 192.168.10.0/24 to any port 22 proto tcp***

This configuration ensures:

- ❖ Administrative access (SSH) is restricted to the internal subnet.
- ❖ Web services (HTTP/HTTPS) are explicitly permitted.
- ❖ All other inbound traffic is denied.

The default-deny policy reflects the principle of fail-safe defaults, meaning access is denied unless explicitly permitted (Saltzer and Schroeder, 1975).

```

slettvoldadmin@lnxe-web01:~$ sudo ufw status verbose
Status: active
Logging: on (low)
Default: deny (incoming), allow (outgoing), disabled (routed)
New profiles: skip

To Action From
--
80/tcp ALLOW IN Anywhere
443/tcp ALLOW IN Anywhere
22/tcp ALLOW IN 192.168.10.0/24
80/tcp (v6) ALLOW IN Anywhere (v6)
443/tcp (v6) ALLOW IN Anywhere (v6)

slettvoldadmin@lnxe-web01:~$ _

```

Figure 7 – UFW firewall status showing restricted SSH access to the internal subnet only.

This layered approach strengthens security by reducing unnecessary exposure while maintaining required service functionality.

2.8.1 Stateful Firewall Concept

Modern firewalls operate in a stateful manner, meaning they track active connection states and automatically permit return traffic for established sessions.

UFW operates on top of iptables, which provides state-aware filtering. This ensures that outbound-initiated traffic receives legitimate return responses without requiring manual rule duplication.

Stateful inspection reduces configuration complexity and lowers the risk of misconfiguration.

2.8.2 Port Exposure Analysis

Only the following ports were permitted:

TCP 22 (SSH)

TCP 80 (HTTP)

TCP 443 (HTTPS)

Limiting open ports reduces potential attack vectors. Each open port represents a possible entry point for exploitation. Restricting open ports and enforcing least privilege network access significantly reduces potential attack vectors in modern network environments (Stallings, 2018).

Unused ports were verified as closed through controlled port scanning during the testing phase.

Excluding legacy services such as FTP and Telnet further reduces risk exposure.

2.8.3 Internal vs External Filtering

In enterprise environments, traffic filtering typically occurs at multiple layers:

- ❖ Perimeter firewall
- ❖ Internal segmentation firewall
- ❖ Host-based firewall

Although this project is implemented in a simulated environment, the logical design reflects this layered filtering model.

Host-based firewall rules ensure that even if network-level segmentation were bypassed, individual servers would still enforce traffic restrictions.

This approach aligns with defence-in-depth principles.

2.8.4 Logging and Monitoring

Firewall logging was enabled to provide visibility into blocked traffic and authentication attempts.

Monitoring log data increases situational awareness and enables early detection of suspicious activity.

In production environments, logs would typically be centralised and monitored through automated alerting systems. While centralised logging was not implemented in this lab, the architectural consideration aligns with enterprise security planning.

2.9 Wireless Network Security Implementation

Wireless connectivity introduces additional risk due to signal propagation beyond physical boundaries.

To mitigate wireless threats such as unauthorised access and eavesdropping, the wireless network was configured using WPA2/WPA3 encryption standards in accordance with NIST recommendations (NIST, 2018).

Security configuration included:

- ❖ Clearly Defined SSID
- ❖ Strong Encryption using WPA2/WPA3 standards
- ❖ Authentication that prevents unauthorised devices from connecting to the wireless network

These measures reduce the likelihood of unauthorised device connection and protect wireless traffic from passive interception.

2.10 OS Hardening Measures

Operating system hardening was implemented to reduce the attack surface and strengthen host-level controls. Hardening practices such as minimising installed services, restricting administrative access, and enforcing secure authentication mechanisms are widely recommended in Linux system administration best practices (Nemeth et al., 2017).

Hardening measures included:

- ❖ Minimal installation
- ❖ SSH configuration security
- ❖ User privilege management
- ❖ File permission control
- ❖ Patch management
- ❖ Firewall enforcement

2.10.1 Minimal Installation and Service Reduction

A minimal Ubuntu Server installation was selected to reduce unnecessary packages and background services.

Active services were reviewed using:

systemctl list-units --type=service

Reducing active services decreases potential attack vectors and improves system stability. (NSA's Top Ten Cybersecurity Mitigation Strategies, 2025)

2.10.2 SSH Hardening Configuration

SSH key-based authentication was enforced. Password-based authentication and direct root login were disabled.

Configuration file:

/etc/ssh/sshd_config

The following parameters were implemented:

- ❖ PasswordAuthentication no
- ❖ PermitRootLogin no
- ❖ PubkeyAuthentication yes
- ❖ AuthenticationMethods publickey

These measures significantly reduce exposure to brute-force attacks (NIST, 2020).

(Validation screenshots are provided in Section 3)

2.10.3 User Privilege Management

Administrative privileges are granted via sudo rather than direct root access.

usermod -aG sudo username

This approach supports the principle of least privilege (Saltzer and Schroeder, 1975).

2.10.4 File Permissions and Ownership

Backup directories were assigned restricted ownership:

chown -R slettvoldadmin:slettvoldadmin /srv/backups

Proper file permissions prevent unauthorised modification or access. (CVE-2023-28079: Insecure File Permissions Vulnerability, 2023)

2.10.5 System Updates and Patch Management

Systems were updated using:

`sudo apt update`

`sudo apt upgrade`

Maintaining updated systems reduces exposure to known vulnerabilities and aligns with CIS recommendations (CIS, 2023).

```
slettvoildadmin@lnxe-web01:~$ sudo apt update
Hit:1 http://archive.ubuntu.com/ubuntu noble InRelease
Hit:2 http://archive.ubuntu.com/ubuntu noble-updates InRelease
Hit:3 http://archive.ubuntu.com/ubuntu noble-backports InRelease
Hit:4 http://security.ubuntu.com/ubuntu noble-security InRelease
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
All packages are up to date.
slettvoildadmin@lnxe-web01:~$ sudo apt upgrade -y
Reading package lists... Done
Building dependency tree... Done
Reading state information... Done
Calculating upgrade... Done
0 upgraded, 0 newly installed, 0 to remove and 0 not upgraded.
slettvoildadmin@lnxe-web01:~$
```

Figure 8 - Successful system update output

2.10.6 Firewall Integration as Hardening

Firewall configuration forms part of OS hardening.

UFW was configured with a default-deny inbound policy, allowing only necessary services.

This ensures layered host-level enforcement independent of network segmentation.

2.10.7 Log Monitoring and Security Awareness

Authentication logs and system logs located in `/var/log/` were reviewed to monitor activity.

Log review enhances detection capability and accountability. (Kent & Souppaya (2006) SP 800-92)

2.10.8 Hardening Philosophy and Industry Alignment

Hardening measures implemented align with recognised industry guidance (CIS, 2023; NIST, 2020).

Security was integrated during system deployment rather than applied retroactively.

2.11 Administrative Access Controls

Administrative access paths are logically separated from standard user traffic.

Direct root login is disabled. All administrative actions require controlled privilege escalation.

This reduces exposure of sensitive management interfaces. (CISA), 2023)

2.12 Security Best Practices Alignment

The infrastructure aligns with:

- ❖ Defence in Depth
- ❖ Least Privilege
- ❖ Secure by Design
- ❖ Separation of Duties

Layered controls ensure that compromise of one mechanism does not result in total system compromise.

Patch management and access control align with CIS benchmark guidance.

2.13 Risk Analysis & Threat Considerations

The design of LNXE Corp's infrastructure was informed not only by functional requirements but also by a structured consideration of potential threat scenarios. Even in a simulated enterprise environment, it is essential to evaluate how systems could be compromised, misconfigured, or disrupted. By identifying risks during the design phase, appropriate security controls can be integrated proactively rather than applied reactively.

2.13.1 External Threats

One of the primary risks in any server-client infrastructure is unauthorised remote access. Attackers commonly target exposed services such as SSH and HTTP using brute-force attacks, credential stuffing, or exploitation of known vulnerabilities.

If password-based authentication is enabled, automated attack tools can attempt thousands of logins per minute, significantly increasing the likelihood of credential compromise.

To mitigate this threat:

- ❖ SSH key-based authentication was enforced
- ❖ Password authentication was disabled.
- ❖ Direct root login was disabled.

Cryptographic authentication mechanisms significantly reduce the likelihood of unauthorised access compared to password-based methods (NIST, 2020).

In addition, internal systems use private IP addresses (RFC 1918), meaning they are not directly routable from external networks (IETF, 1996). This architectural decision reduces exposure to reconnaissance scanning and direct exploitation attempts.

2.13.2 Internal Threats and Lateral Movement

Security threats are not limited to external actors. A compromised internal workstation could attempt lateral movement toward critical infrastructure if segmentation and access controls are insufficient.

To reduce this risk:

- ❖ Network segmentation was implemented.
- ❖ Administrative access paths were logically separated.
- ❖ SSH access was restricted to specific internal IP ranges
- ❖ Least privilege principles were enforced.

This layered control model aligns with defence-in-depth principles (CIS, 2023). Even if an internal system were compromised, firewall rules and privilege restrictions limit the potential for escalation.

Limiting administrative privileges through sudo-based escalation further reduces impact in the event of credential theft (Saltzer and Schroeder, 1975).

2.13.3 Data Loss and System Failure

Data loss due to hardware failure, accidental deletion, or ransomware represents a significant operational risk.

Without a tested backup strategy, business continuity could be severely disrupted.

To address this risk:

- ❖ A dedicated backup server was deployed.
- ❖ Backups were performed using rsync over SSH.
- ❖ Restore procedures were tested and validated.

NIST contingency planning guidance emphasises that backup procedures must be tested to ensure reliability during actual incidents (NIST, 2010).

The restore validation process confirms that backup data is functional, not merely present.

2.13.4 Service Misconfiguration and Vulnerability Exposure

Misconfiguration is one of the most common causes of security incidents. Exposing unnecessary ports, running outdated services, or enabling unused components increases the attack surface.

To reduce misconfiguration risk:

- ❖ A default-deny inbound firewall policy was implemented.
- ❖ Only required ports (22, 80, 443) were allowed.
- ❖ Minimal server installations were used.
- ❖ Systems were patched regularly.

This approach aligns with CIS secure configuration benchmarks (CIS, 2023).

Reducing the attack surface is often more effective than relying solely on reactive security mechanisms.

2.13.5 Wireless Network Risks

Wireless networks introduce additional risks due to signal propagation beyond physical boundaries.

Threats include:

- ❖ Unauthorised device connections
- ❖ Credential brute-forcing
- ❖ Passive traffic interception

To mitigate these risks:

- ❖ WPA2/WPA3 encryption was configured
- ❖ Strong authentication credentials were applied.
- ❖ Access control was enforced at the wireless layer.

Encrypted wireless communication reduces susceptibility to eavesdropping and credential compromise (NIST, 2018).

2.13.6 Residual Risk and Limitations

Despite the implementation of layered controls, no infrastructure can be considered entirely risk-free.

Residual risks include:

- ❖ Zero-day vulnerabilities
- ❖ Insider threats
- ❖ Human misconfiguration
- ❖ Advanced persistent threats

In a production environment, additional controls such as multi-factor authentication (MFA), intrusion detection systems (IDS), and centralised logging would further strengthen the security posture.

The objective of this project is not to eliminate all risk but to demonstrate structured, risk-aware infrastructure design.

By identifying credible threat scenarios and implementing layered mitigation strategies, the LNXE Corp infrastructure achieves a balanced and practical level of security aligned with industry best practices.

2.14 Design Decisions & Trade-Off Analysis

The infrastructure design for LNXE Corp required several architectural decisions balancing realism, security, resource constraints, and academic feasibility.

Each decision involved trade-offs between performance, exposure, complexity, and operational control.

2.14.1 Choice of Type 2 Hypervisor

VMware Workstation Pro (Type 2 hypervisor) was selected instead of a Type 1 hypervisor such as VMware ESXi.

Advantages of Type 1 (ESXi):

- ❖ Stronger isolation from the host OS
- ❖ Lower virtualisation overhead
- ❖ More realistic enterprise deployment

Reasons for selecting Type 2:

- ❖ Academic hardware constraints
- ❖ Faster deployment
- ❖ Easier testing and rollback
- ❖ Sufficient isolation for lab purposes

The trade-off involves slightly reduced realism and marginally higher overhead compared to bare-metal virtualisation. However, VMware Workstation still provides robust isolation and supports secure network segmentation, making it suitable for demonstrating enterprise concepts (VMware, 2023).

The decision reflects a practical balance between realism and academic constraints.

2.14.2 Internal-Only Network Design

The infrastructure was intentionally designed as an internal-only network using host-only virtualisation networking.

Advantages:

- ❖ Eliminates real-world external attack exposure
- ❖ Safer for the student lab environment
- ❖ Simplifies security validation

Trade-off:

- ❖ Does not fully simulate perimeter firewall or DMZ deployment
- ❖ Limited demonstration of external routing controls

Although a DMZ architecture would increase realism, it would also increase configuration complexity and potential exposure on a non-enterprise physical host.

The selected design prioritises safety while still demonstrating segmentation, access control, and layered security principles.

2.14.3 Backup Strategy Selection

A dedicated backup server (**LNXE-BAK01**) was deployed to provide resilience and operational continuity.

The backup strategy uses rsync over SSH for secure, incremental file synchronisation.

This design reflects the principle of separation of duties, ensuring backup storage remains isolated from production workloads.

2.14.3.1 Backup Server Architecture

The backup server was deployed as a separate Ubuntu Server virtual machine with a dual-network interface configuration:

- ❖ Internal LAN Interface (192.168.10.20/24) – Dedicated exclusively to communication with the web server.
- ❖ NAT Interface (192.168.40.0/24) – Used for controlled outbound connectivity, including system updates and package repository access.

Only the NAT interface maintains a default route, preventing unintended routing of internal traffic.

Backup data is stored under:

/srv/backups/web01

```
slettvoidadmin@lnxe-bak01:~$ sudo mkdir -p /srv/backups/web01
[sudo] password for slettvoidadmin:
slettvoidadmin@lnxe-bak01:~$ sudo mkdir -p /srv/backups/web01
slettvoidadmin@lnxe-bak01:~$ sudo chown -R slettvoidadmin:slettvoidadmin /srv/backups
slettvoidadmin@lnxe-bak01:~$ ls -la /srv/backups
total 12
drwxr-xr-x 3 slettvoidadmin slettvoidadmin 4096 Feb 24 20:44 .
drwxr-xr-x 3 root          root          4096 Feb 24 20:44 ..
drwxr-xr-x 2 slettvoidadmin slettvoidadmin 4096 Feb 24 20:44 web01
slettvoidadmin@lnxe-bak01:~$ _
```

Figure 9 – Backup directory structure on LNxE-BAK01.

Trade-off Analysis:

A single-server combined Web + Backup role would reduce resource usage but create a single point of failure.

Separating roles increases resource consumption but significantly improves resilience and security isolation.

This decision prioritises security over hardware efficiency.

```
slettvoidadmin@lnxe-bak01:~$ ip a
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host noprefixroute
        valid_lft forever preferred_lft forever
2: ens33: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 00:0c:29:02:d3:ea brd ff:ff:ff:ff:ff:ff
    altname enp2s1
    inet 192.168.10.20/24 brd 192.168.10.255 scope global ens33
        valid_lft forever preferred_lft forever
    inet6 fe80::20c:29ff:fe02:d3ea/64 scope link
        valid_lft forever preferred_lft forever
3: ens34: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc fq_codel state UP group default qlen 1000
    link/ether 00:0c:29:02:d3:f4 brd ff:ff:ff:ff:ff:ff
    altname enp2s2
    inet 192.168.40.128/24 metric 100 brd 192.168.40.255 scope global dynamic ens34
        valid_lft 1796sec preferred_lft 1796sec
    inet6 fe80::20c:29ff:fe02:d3f4/64 scope link
        valid_lft forever preferred_lft forever
slettvoidadmin@lnxe-bak01:~$ ping 8.8.8.8 -C 2
PING 2 (8.8.8.8) 56(124) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=128 time=27.0 ms (DIFFERENT ADDRESS!)
64 bytes from 8.8.8.8: icmp_seq=2 ttl=128 time=26.4 ms (DIFFERENT ADDRESS!)
64 bytes from 8.8.8.8: icmp_seq=3 ttl=128 time=26.2 ms (DIFFERENT ADDRESS!)
```

Figure 10 – Network configuration of LNxE-BAK01 showing dual-NIC separation and default route via NAT

2.14.3.2 Secure Backup Communication

Backup traffic is encrypted using SSH key-based authentication.

Password authentication is disabled.

Firewall rules restrict SSH access to LNXE-WEB01 (192.168.10.10) only.

This ensures that:

- ❖ Only the web server can initiate backup sessions.
- ❖ The backup server is not accessible from client systems.
- ❖ External systems cannot reach the backup service.
- ❖ Backup storage is protected from lateral movement within the network.

```
slettvoildadmin@lnxe-bak01:~$ sudo ufw default deny incoming
Default incoming policy changed to 'deny'
(be sure to update your rules accordingly)
slettvoildadmin@lnxe-bak01:~$ sudo ufw default allow outgoing
Default outgoing policy changed to 'allow'
(be sure to update your rules accordingly)
slettvoildadmin@lnxe-bak01:~$ sudo ufw allow from 192.168.10.10 to any port 22 proto tcp
Rules updated
slettvoildadmin@lnxe-bak01:~$ ufw enable
ERROR: You need to be root to run this script
slettvoildadmin@lnxe-bak01:~$ sudo ufw status verbose
Status: inactive
slettvoildadmin@lnxe-bak01:~$ sudo ufw enable
Firewall is active and enabled on system startup
slettvoildadmin@lnxe-bak01:~$ sudo ufw status verbose
Status: active
Logging: on (low)
Default: deny (incoming), allow (outgoing), disabled (routed)
New profiles: skip

To          Action      From
--          -
22/tcp      ALLOW IN    192.168.10.10
```

Figure 11 – UFW configuration restricting SSH access to WEB01 only.

This configuration reflects the principle of least privilege and strengthens defence in depth by combining authentication controls with network-level filtering.

```
LNXE-WEB01 x LNXE-BAK01 x WDN11-CL01 x LNXE-CL01 x

slettvoidadmin@lnxe-web01:~$ ssh -o PreferredAuthentications=password -o PubkeyAuthentication=no slettvoidadmin@192.168.10.20
slettvoidadmin@192.168.10.20: Permission denied (publickey).
slettvoidadmin@lnxe-web01:~$ ssh slettvoidadmin@192.168.10.20
Welcome to Ubuntu 24.04.4 LTS (GNU/Linux 6.8.0-71-generic x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/pro

System information as of Tue Feb 24 09:05:05 PM UTC 2026

System load:  0.0          Processes:      217
Usage of /:   36.2% of 18.53GB Users logged in: 1
Memory usage: 12%         IPv4 address for ens34: 192.168.40.128
Swap usage:  0%

 * Strictly confined Kubernetes makes edge and IoT secure. Learn how MicroK8s
   Just raised the bar for easy, resilient and secure K8s cluster deployment.
   https://ubuntu.com/engage/secure-kubernetes-at-the-edge

Expanded Security Maintenance for Applications is not enabled.

0 updates can be applied immediately.

Enable ESM Apps to receive additional future security updates.
See https://ubuntu.com/esm or run: sudo pro status

*** System restart required ***
Last login: Tue Feb 24 20:59:15 2026 from 192.168.10.10
slettvoidadmin@lnxe-bak01:~$ _
```

Figure 12 – Password-based SSH authentication attempt rejected on LNXE-BAK01 and Successful SSH login using public key authentication.

The validation confirms that interactive password authentication is disabled and that access to the backup server is restricted exclusively to public key authentication from authorised internal systems. This significantly reduces exposure to brute-force and credential-based attacks.

2.14.3.3 Backup Execution

Backups are performed using:

```
sudo                                rsync                                -avz                                /var/www/html/  
slettvoidadmin@192.168.10.20:/srv/backups/web01/html/
```

Rsync was selected over alternatives such as full disk imaging or enterprise backup suites due to:

The options used provide:

- ❖ -a – Archive mode (preserves permissions and metadata)
- ❖ -v – Verbose output for verification
- ❖ -z – Compression during transfer

Using rsync over SSH ensures that data is encrypted in transit while maintaining file integrity and minimising bandwidth usage.

This implementation balances simplicity and effectiveness, making it suitable for a small but growing organisation.

2.14.3.4 Restore Validation

To verify operational recoverability, a controlled restore test was conducted. A test file (restore-test.txt) was created on the production web server and included in the backup set. After confirming successful synchronisation to the backup server, the file was deliberately deleted from the production environment to simulate data loss.

The restoration was performed using rsync over SSH from LNXE-BAK01 back to LNXE-WEB01:

```
rsync -avz slettvoledadmin@192.168.10.20:/srv/backups/web01/html/ /var/www/html/
```

Following the restore operation, the file contents were verified:

```
cat /var/www/html/restore-test.txt
```

The successful recovery of the file confirms that:

- ❖ Backup data is accessible.
- ❖ Restoration procedures are functional.
- ❖ Data integrity is preserved.
- ❖ The system can recover from simulated loss scenarios.

This validation demonstrates that the implemented backup strategy supports business continuity objectives and aligns with disaster recovery best practices.

2.14.4 Separation of Web and Backup Roles

Deploying separate servers for Web and Backup functions increases resource usage but improves security and redundancy. Combining both services on a single server would reduce overhead but increase the risk in the event of a compromise.

The decision reflects prioritisation of security over resource efficiency.

2.15 Defence-in-Depth Model Applied to LNXE Corp

The security architecture implemented in this project is not based on a single control or isolated mechanism, but rather on a layered security strategy commonly referred to as Defence in Depth. This principle is widely recognised in information security

frameworks and standards such as NIST and CIS (CIS, 2023; NIST, 2020). The fundamental idea behind defence in depth is that no single security control should be relied upon as the sole protection mechanism. Instead, multiple independent layers of security are implemented so that if one control fails, others continue to provide protection.

In the LNXE Corp infrastructure, defence in depth is reflected across several logical layers:

2.15.1 Physical / Host Layer

Even though this project is implemented in a virtualised lab environment, the concept of physical security still applies. In a real enterprise environment, servers would be physically secured in controlled data centre facilities with restricted access. Unauthorised physical access can bypass many logical controls.

Within this simulated environment, the “physical layer” is represented by the host machine running VMware Workstation Pro. The virtual machines are isolated from each other and from the external environment through virtualisation boundaries. This isolation ensures that compromise of one VM does not automatically result in compromise of the host or other virtual machines.

2.15.2 Network Layer Security

At the network layer, several protective measures were implemented:

- ❖ Private IP addressing (RFC 1918)
- ❖ Logical segmentation of server and client systems
- ❖ Controlled routing via a central router
- ❖ Host-based firewall rules (UFW)
- ❖ Restricted port exposure

Private addressing ensures that internal systems are not directly routable from external networks (IETF, 1996). Even if a misconfiguration occurred at the application layer, network-level restrictions would still limit exposure.

The default-deny inbound firewall policy further strengthens this layer. By explicitly allowing only required ports (22, 80, 443), unnecessary attack vectors are eliminated. This aligns with the principle of fail-safe defaults (Saltzer and Schroeder, 1975).

2.15.3 Host / Operating System Layer

Operating system hardening provides another independent security layer. Measures such as:

- ❖ Disabling root login
- ❖ Enforcing SSH key-based authentication
- ❖ Reducing installed services
- ❖ Applying security updates
- ❖ Restricting file permissions

Reduce each server's internal attack surface. Even if an attacker manages to reach the server over the network, the hardened OS configuration limits the ability to escalate privileges or execute arbitrary services.

This separation between network controls and host controls is important. If firewall rules were accidentally modified or bypassed, SSH configuration and user privilege management would still provide protection.

2.15.4 Application Layer Security

At the application layer, the web server was configured to only provide necessary services. Unused modules or services were not enabled. HTTPS encryption (if implemented) ensures secure communication between clients and the server.

Limiting exposed services reduces opportunities to exploit application-level vulnerabilities. (OWASP 2021) Furthermore, separating Web Server and Backup Server roles prevents compromise of one application from automatically affecting backup integrity.

2.15.5 Administrative & Human Layer

Human error is one of the most common causes of security incidents. Administrative access controls were implemented to reduce this risk:

- ❖ No direct root login
- ❖ Sudo-based privilege escalation
- ❖ Restricted SSH users
- ❖ Logged administrative actions

By enforcing accountability and limiting permanent high-level privileges, the infrastructure reduces the impact of accidental misconfiguration or compromised credentials.

2.15.6 Layer Interaction & Redundancy

The strength of defence in depth lies in the interaction between layers. For example:

- ❖ Even if wireless encryption were broken, firewall rules still restrict access.
- ❖ Even if firewall rules were misconfigured, SSH key enforcement still protects administrative access.
- ❖ Even if user credentials were compromised, limited privileges restrict damage.
- ❖ Even if the Web Server is compromised, the Backup Server remains logically separated.

This layered security approach does not eliminate all risks, but it significantly increases the effort required for a successful attack. Attackers would need to bypass multiple independent controls rather than exploiting a single weakness. (Layered Security: Multiple layers of Cybersecurity Working Together to Defend Your Business, 2023)

The implementation in this project demonstrates structured security thinking rather than isolated technical configuration. By integrating security at multiple architectural levels, the infrastructure achieves a more resilient posture aligned with enterprise best practices (CIS, 2023; NIST, 2020).

2.16 Business Continuity & Disaster Recovery Considerations

While security controls aim to prevent unauthorised access and compromise, no infrastructure can be considered completely immune to failure. Hardware malfunction, software corruption, accidental deletion, ransomware or human error may still result in service disruption or data loss. Because of this, security architecture must also incorporate Business Continuity (BC) and Disaster Recovery (DR) considerations.

Business continuity refers to the ability of an organisation to continue operating during and after a disruptive event. Disaster recovery specifically focuses on restoring systems and data following an incident (NIST, 2010).

In the LNXE Corp infrastructure, resilience is primarily supported by a dedicated backup server using rsync-based synchronisation.

2.16.1 Recovery Time Objective (RTO)

Recovery Time Objective defines how quickly a system must be restored after failure. In this small start-up scenario, an acceptable RTO may be measured in hours rather than minutes.

Using rsync allows for relatively fast file restoration from the backup server to the web server. While this approach may not provide real-time failover capabilities, it offers a practical and sufficient solution for a ten-person organisation.

In a larger enterprise, more advanced solutions such as replication clusters or high-availability systems would further reduce RTO.

2.16.2 Recovery Point Objective (RPO)

The Recovery Point Objective defines the acceptable amount of data loss, typically measured in terms of time. For example, if backups are performed daily, up to 24 hours of data may be lost.

In this project, backup frequency can be adjusted depending on business needs. More frequent backups reduce RPO but increase storage usage and administrative overhead.

Understanding RTO and RPO demonstrates that a backup strategy is not only a technical decision but also a business decision. (RTO and RPO explained – MSP Backup, 2022)

2.16.3 Backup vs Redundancy

It is important to distinguish between backup and redundancy. Redundancy (such as RAID or mirrored servers) ensures availability if hardware fails, but it does not protect against accidental deletion or ransomware. If a file is deleted, redundancy simply mirrors the deletion.

Backup, on the other hand, allows restoration of previous data states. The implemented rsync-based backup ensures that deleted or corrupted files can be restored.

This separation of concepts strengthens overall resilience.

2.16.4 Ransomware & Integrity Protection

Modern threat landscapes increasingly include ransomware attacks, where data is encrypted and held for ransom. If backups are directly accessible from compromised systems, attackers may also encrypt backup data.

In a real production environment, additional measures would be recommended, such as:

- ❖ Offline backups
- ❖ Immutable backup storage
- ❖ Backup access restrictions
- ❖ Multi-factor authentication for backup systems

Although not fully implemented in this lab, the logical separation of the Web Server and the Backup Server reduces the risk that compromise of the web service automatically results in total data loss.

2.16.5 Testing & Validation of Recovery

One of the most critical aspects of disaster recovery planning is validation. Creating backups is insufficient if restore procedures are never tested. During the practical phase of this project, restore procedures were executed to verify data integrity and usability.

This validation confirms that backups are not corrupted and that recovery steps are documented and reproducible.

Regular testing of recovery processes aligns with NIST contingency planning recommendations (NIST, 2010).

2.16.6 Scalability & Future Continuity Planning

If LNXE Corp expands beyond 10 employees, the current backup architecture may need to be adjusted. Increased data volume and service complexity may require:

- ❖ Centralised backup management
- ❖ Incremental snapshot solutions
- ❖ Cloud-based offsite replication
- ❖ Automated backup monitoring

However, the current design establishes a foundational resilience strategy suitable for a small but growing organisation.

2.16.7 Risk-Aware Resilience Design

By incorporating backup validation, separation of duties, and structured restore procedures, this infrastructure demonstrates awareness that security is not only about preventing attacks, but also about recovering from them.

The combination of preventive controls (firewall, SSH hardening) and recovery mechanisms (rsync backups, restore testing) ensures a balanced, realistic security posture.

This integrated approach strengthens both operational stability and long-term sustainability of the IT environment.

3. Practical (Testing & Results)

After the infrastructure and security controls were implemented, a series of practical tests were conducted to verify that the system functions correctly and that the security mechanisms operate as intended. These tests were designed to validate connectivity, service availability, access control, backup reliability, and firewall protection.

The validation process focused on confirming that:

- ❖ Internal network communication between hosts operates correctly.
- ❖ The web service is accessible to clients.
- ❖ SSH access is secured using key-based authentication
- ❖ Backup operations successfully transfer data to the backup server.
- ❖ Data can be restored in the event of a failure.
- ❖ Firewall rules restrict unnecessary network access.

Each test is documented with command output and screenshots to provide verifiable evidence of system functionality.


3.1 Network Connectivity Validation

Before testing services, internal network connectivity between hosts was verified. The client machine LNXXE-CLI01 performed an ICMP ping test to the web server LNXXE-WEB01 (192.168.10.10).

Successful responses confirm that:

- ❖ The LAN network configuration is functioning correctly.
- ❖ The web server is reachable from internal clients.
- ❖ Basic IP communication between hosts is operational.

This test verifies that the network infrastructure is correctly configured and that there are no routing or addressing errors.

A terminal window titled 'slettvoid-dev@LNXXE-CLI01: ~' showing network configuration and a successful ping test. The configuration includes an ethernet interface 'ens33' with IP '192.168.10.255' and a link-local IPv6 address. The ping test is performed from 'slettvoid-dev@LNXXE-CLI01' to '192.168.10.10', showing 12 successful packets with 0% loss and an average round-trip time of 0.569ms.

```
altname enp2s1
inet 192.168.10.30/24 brd 192.168.10.255 scope global noprefixroute ens33
    valid_lft forever preferred_lft forever
inet6 fe80::fa7e:d6d5:a457:2f92/64 scope link noprefixroute
    valid_lft forever preferred_lft forever
slettvoid-dev@LNXXE-CLI01:~$ ping 192.168.10.10
PING 192.168.10.10 (192.168.10.10) 56(84) bytes of data.
 64 bytes from 192.168.10.10: icmp_seq=1 ttl=64 time=1.04 ms
 64 bytes from 192.168.10.10: icmp_seq=2 ttl=64 time=0.544 ms
 64 bytes from 192.168.10.10: icmp_seq=3 ttl=64 time=0.487 ms
 64 bytes from 192.168.10.10: icmp_seq=4 ttl=64 time=0.516 ms
 64 bytes from 192.168.10.10: icmp_seq=5 ttl=64 time=0.505 ms
 64 bytes from 192.168.10.10: icmp_seq=6 ttl=64 time=0.799 ms
 64 bytes from 192.168.10.10: icmp_seq=7 ttl=64 time=0.490 ms
 64 bytes from 192.168.10.10: icmp_seq=8 ttl=64 time=0.471 ms
 64 bytes from 192.168.10.10: icmp_seq=9 ttl=64 time=0.511 ms
 64 bytes from 192.168.10.10: icmp_seq=10 ttl=64 time=0.506 ms
 64 bytes from 192.168.10.10: icmp_seq=11 ttl=64 time=0.498 ms
 64 bytes from 192.168.10.10: icmp_seq=12 ttl=64 time=0.461 ms
^C
--- 192.168.10.10 ping statistics ---
12 packets transmitted, 12 received, 0% packet loss, time 11167ms
rtt min/avg/max/mdev = 0.461/0.569/1.041/0.165 ms
slettvoid-dev@LNXXE-CLI01:~$
```

Figure 13 – Successful ping test from LNXXE-CLI01 to WEB01 confirming LAN connectivity.

3.2 SSH Authentication Security Testing

Secure remote access to the servers was validated by testing the SSH authentication configuration. Password-based authentication had been disabled in the SSH configuration file to reduce the risk of brute-force attacks.

Two authentication attempts were performed:

❖ A login attempt using password authentication




Figure 14 – SSH password authentication attempt rejected due to disabled password login.

❖ A login attempt using SSH public key authentication




Figure 15 – Successful SSH login using public key authentication.

The password authentication attempt was rejected, confirming that password-based login is disabled. The subsequent login using the authorised SSH key succeeded, verifying that only key-based authentication is permitted.

This security configuration ensures that remote administrative access is restricted to authorised users possessing the correct private key.

3.3 Backup & Restore Verification


To verify the backup system's reliability, a controlled test was conducted to ensure that data could be successfully backed up and restored if required.

The following procedure was performed:

- ❖ A test file was created on the web server.
- ❖ The rsync backup process was executed.
- ❖ The file was verified on the backup server.
- ❖ The original file was deleted from the web server.
- ❖ The file was restored from the backup server.
- ❖ The file content was verified after restoration.

The rsync command successfully transferred the test file to the backup server. After deleting the original file, the restore process retrieved the file from the backup repository and restored it to the web server. The final verification confirmed that the restored file contained the expected content.


This test demonstrates that the backup solution is fully functional and capable of restoring critical data in the event of system failure.



```
LNXE-WEB01 x LNXE-BAK01 x WBN11-CLI01 x LNXE-CLI01 x
slettvoidadmin@lnxe-web01:~$ rsync -avz /var/www/html/ slettvoidadmin@192.168.10.20:/srv/backups/web01/html
sending incremental file list
created directory /srv/backups/web01/html
./
index.nginx-debian.html

sent 534 bytes  received 84 bytes  412.00 bytes/sec
total size is 615  speedup is 1.00
slettvoidadmin@lnxe-web01:~$ _
```

Figure 16 – Successful rsync synchronisation output



```
LNXE-WEB01 x LNXE-BAK01 x WBN11-CLI01 x LNXE-CLI01 x
slettvoidadmin@lnxe-bak01:~$ ls -la /srv/backups/web01/html
total 12
drwxr-xr-x 2 slettvoidadmin slettvoidadmin 4096 Feb 24 15:20 .
drwxr-xr-x 3 slettvoidadmin slettvoidadmin 4096 Feb 24 21:13 ..
-rw-r--r-- 1 slettvoidadmin slettvoidadmin 615 Feb 24 15:20 index.nginx-debian.html
slettvoidadmin@lnxe-bak01:~$
```

Figure 17 – Backup files stored on LNXE-BAK01 under /srv/backups/web01/html.

```
LNXE-WEB01 X LNXE-BAK01 X WIN11-CLI01 X LNXE-CLI01 X

slettvoidadmin@lnxe-web01:~$ sudo bash -c 'echo "BACKUP-TEST" > /var/www/html/restore-test.txt'
slettvoidadmin@lnxe-web01:~$ ls -la /var/www/html | grep restore-test
-rw-r--r-- 1 root root 12 Feb 24 21:48 restore-test.txt
slettvoidadmin@lnxe-web01:~$ rsync -avz /var/www/html/ slettvoidadmin@192.168.10.20:/srv/backups/web01/html/
sending incremental file list
./
restore-test.txt

sent 224 bytes received 38 bytes 524.00 bytes/sec
total size is 627 speedup is 2.39
slettvoidadmin@lnxe-web01:~$ _
```

Figure 18 - File exists on backup

```
LNXE-WEB01 X LNXE-BAK01 X WIN11-CLI01 X LNXE-CLI01 X

slettvoidadmin@lnxe-web01:~$ sudo rm /var/www/html/restore-test.txt
slettvoidadmin@lnxe-web01:~$ ls -la /var/www/html | grep restore-test
slettvoidadmin@lnxe-web01:~$
```

Figure 19 – Delete file

```
LNXE-WEB01 X LNXE-BAK01 X WIN11-CLI01 X LNXE-CLI01 X

slettvoidadmin@lnxe-web01:~$ rsync -avz slettvoidadmin@192.168.10.20:/srv/backups/web01/html/ /var/www/html/
receiving incremental file list
./
restore-test.txt

sent 46 bytes received 212 bytes 516.00 bytes/sec
total size is 627 speedup is 2.43
slettvoidadmin@lnxe-web01:~$ cat /var/www/html/restore-test.txt
BACKUP-TEST
slettvoidadmin@lnxe-web01:~$ _
```

Figure 20 – Restore command

3.4 Firewall Validation


Firewall rules were tested to ensure that only the required network services are accessible. The firewall configuration allows only essential ports while blocking all other incoming connections.

A port scan was performed against the web server to verify which services were accessible. The results confirmed that only the expected ports were open:

- ❖ Port 22 (SSH) for secure administration
- ❖ Port 80 (HTTP) for web traffic
- ❖ Port 443 (HTTPS) for encrypted web communication

All other ports were successfully filtered by the firewall, demonstrating that unnecessary services are not exposed to the network.

This validation confirms that the firewall configuration follows the principle of least privilege, minimising the system's attack surface.



```
slettvoild-dev@LNXE-CLI01: ~  
slettvoild-dev@LNXE-CLI01:~$ nmap 192.168.10.10  
Starting Nmap 7.94SVN ( https://nmap.org ) at 2026-03-04 10:25 CET  
Nmap scan report for 192.168.10.10  
Host is up (0.0012s latency).  
Not shown: 997 filtered tcp ports (no-response)  
PORT      STATE SERVICE  
22/tcp    open  ssh  
80/tcp    open  http  
443/tcp   closed https  
  
Nmap done: 1 IP address (1 host up) scanned in 4.91 seconds  
slettvoild-dev@LNXE-CLI01:~$
```

Figure 21 – Nmap scan from CLI01 verifying that only required ports are accessible on WEB01

3.5 Wireless Security Testing

Wireless network security was tested to verify that unauthorised devices cannot connect to the network without valid authentication credentials.

A connection attempt was made from a wireless client device (DEV-LAP01) to the configured WLAN LNXE-DEV-WLAN using incorrect credentials. The wireless client detected the network but was unable to establish a connection due to the WPA2-PSK authentication requirement.

This figure shows the wireless client detecting the available WLAN while remaining unable to authenticate successfully. This demonstrates that the network requires valid credentials before granting access.

The following test was performed from a wireless client device:




Figure 22 – Wireless client detecting the secured WLAN but failing to authenticate due to incorrect credentials.

The access point configuration was also verified to ensure that strong encryption and authentication mechanisms were enabled. As shown in this figure, the wireless network is configured to use WPA2-PSK with AES encryption, requiring a pre-shared key for authentication.




Figure 23 – Wireless access point configured with WPA2-PSK security.

These results confirm that the wireless network enforces authentication controls and prevents unauthorised devices from connecting without valid credentials.

3.6 Security Validation Summary

The testing procedures performed throughout this section demonstrate that the implemented infrastructure operates securely and as designed. Each critical component of the system was validated through practical testing scenarios, confirming both functionality and enforcement of security controls.

Network connectivity testing confirmed that internal hosts can communicate correctly within the private LAN environment while remaining isolated from external networks. This verifies that the IP addressing scheme and network segmentation were implemented correctly.

Web server validation confirmed that the Nginx service is accessible to internal clients and operates through the expected network ports. The successful retrieval of the default Nginx webpage demonstrates that the service is functioning correctly and that firewall rules permit legitimate traffic.

SSH authentication testing verified that password-based login attempts are rejected, while public-key authentication is accepted. This confirms that the SSH hardening configuration effectively prevents password-based brute-force attacks and restricts administrative access to authorised key holders only.

Backup testing demonstrated that the rsync synchronisation process successfully transfers data from the web server to the backup server. The restore validation confirmed that

deleted files can be recovered from the backup repository and that restored data retains its integrity.

Firewall validation confirmed that only the required network ports are exposed while all other incoming connections are blocked. This reduces the system's attack surface and prevents unnecessary services from being accessible.

Wireless security testing demonstrated that unauthorised connection attempts are rejected when incorrect credentials are provided, confirming that authentication mechanisms protect the wireless network.

Together, these validation tests confirm that the implemented infrastructure meets the project's intended security objectives. The system demonstrates controlled network access, hardened administrative authentication, functional backup and recovery procedures, and restricted service exposure. These measures collectively contribute to a secure, resilient infrastructure that meets LNXE Corp.'s operational needs.

4. Summary & Conclusion

This project demonstrated the design, implementation, and validation of a secure virtualised infrastructure for LNXE Corp. The environment was built using VMware Workstation Pro and consisted of multiple virtual machines configured to simulate a small enterprise network environment (VMware, 2024).

The infrastructure included a web server (LNXE-WEB01), a backup server (LNXE-BAK01), and client systems used for testing and administration. Network segmentation was implemented using separate LAN and NAT networks to isolate internal communication while still allowing controlled access to external resources such as software repositories and updates. Network segmentation is widely recommended as a security practice because it limits lateral movement within a network and reduces potential attack surfaces (NIST, 2020).

Security was a primary focus throughout the implementation. Several hardening measures were applied to reduce the system's attack surface and improve administrative security. These measures included disabling password-based SSH authentication, enforcing public-key authentication, implementing host-based firewall rules with UFW, and restricting SSH access to authorised systems only. Hardening remote access and limiting authentication methods are recommended practices for protecting administrative interfaces from brute-force attacks and unauthorised access (CIS, 2023).

A centralised backup solution was implemented using the rsync utility over SSH. This solution allows automated synchronisation of web server data to the dedicated backup server. Backup integrity was validated through restore testing, confirming that files can be successfully recovered in the event of system failure or data loss. Reliable backup strategies are a fundamental component of system resilience and disaster recovery planning (NIST, 2020).

Additional testing procedures were performed to verify the functionality and security of the infrastructure. Network connectivity tests confirmed correct routing and communication between hosts. Web server validation confirmed that the Nginx service was accessible to internal clients. Firewall testing demonstrated that only required ports were exposed, while all other ports remained closed. Wireless security testing confirmed that unauthorised devices cannot access the network without valid authentication credentials.

Overall, the implemented solution demonstrates a practical approach to secure infrastructure deployment. By combining virtualisation, network segmentation, system hardening, and backup strategies, the environment provides a resilient, secure platform that supports enterprise services while maintaining strong security controls.

5. References

- ❖ Canonical (2024) Ubuntu Server Documentation. Available at: <https://ubuntu.com/server/docs> (Accessed: 4 March 2026).
- ❖ Center for Internet Security (CIS) (2023) CIS Benchmarks. Available at: <https://www.cisecurity.org/cis-benchmarks> (Accessed: 4 March 2026).
- ❖ Cisco (2026) What Is Network Segmentation? Available at: <https://www.cisco.com/site/us/en/learn/topics/security/what-is-network-segmentation.html> (Accessed: 4 March 2026).
- ❖ Cloud Security Alliance (2015) Best Practices for Mitigating Risks in Virtualized Environments. Available at: <https://cloudsecurityalliance.org/artifacts/best-practices-for-mitigating-risks-in-virtualized-environments> (Accessed: 4 March 2026).
- ❖ CISA (2023) Binding Operational Directive 23-02: Mitigating the Risk from Internet-Exposed Management Interfaces. Available at: <https://www.cisa.gov/news-events/directives/binding-operational-directive-23-02> (Accessed: 4 March 2026).
- ❖ IETF (1996) RFC 1918: Address Allocation for Private Internets. Available at: <https://datatracker.ietf.org/doc/html/rfc1918> (Accessed: 4 March 2026).
- ❖ Kent, K. and Souppaya, M. (2006) Guide to Computer Security Log Management (SP 800-92). Gaithersburg, MD: National Institute of Standards and Technology. Available at: <https://csrc.nist.gov/pubs/sp/800/92/final> (Accessed: 4 March 2026).
- ❖ MSP Backup (2025) RTO and RPO explained. Available at: <https://mspbackup.net/rto-and-rpo-explained/> (Accessed: 4 March 2026).
- ❖ National Institute of Standards and Technology (NIST) (2010) SP 800-34 Rev. 1: Contingency Planning Guide for Federal Information Systems. Gaithersburg, MD: National Institute of Standards and Technology. Available at: <https://csrc.nist.gov/publications/detail/sp/800-34/rev-1/final> (Accessed: 4 March 2026).
- ❖ National Institute of Standards and Technology (NIST) (2018) SP 800-153: Guidelines for Securing Wireless Local Area Networks (WLANs). Gaithersburg, MD: National Institute of Standards and Technology. Available at: <https://csrc.nist.gov/publications/detail/sp/800-153/final> (Accessed: 4 March 2026).
- ❖ National Institute of Standards and Technology (NIST) (2020) SP 800-63B: Digital Identity Guidelines – Authentication and Lifecycle Management. Gaithersburg, MD: National Institute of Standards and Technology. Available at: <https://pages.nist.gov/800-63-3/sp800-63b.html> (Accessed: 4 March 2026).
- ❖ National Security Agency (NSA) (2018) NSA's Top Ten Cybersecurity Mitigation Strategies. Available at: <https://www.nsa.gov/portals/75/documents/what-we-do/cybersecurity/professional-resources/csi-nsas-top10-cybersecurity-mitigation-strategies.pdf> (Accessed: 4 March 2026).
- ❖ Nemeth, E., Snyder, G., Hein, T.R. and Whaley, B. (2017) UNIX and Linux System Administration Handbook. 5th edn. Boston: Pearson.
- ❖ Nginx (2024) Nginx Documentation. Available at: <https://nginx.org/en/docs/> (Accessed: 4 March 2026).
- ❖ NIST National Vulnerability Database (NVD) (2023) CVE-2023-28079 Detail. Available at: <https://nvd.nist.gov/vuln/detail/CVE-2023-28079> (Accessed: 4 March 2026).
- ❖ Odom, W. (2020) CCNA 200-301 Official Cert Guide, Volume I. Indianapolis, IN: Cisco Press.

- ❖ OWASP (2021) OWASP Top 10:2021 – A05 Security Misconfiguration. Available at: https://owasp.org/Top10/2021/A05_2021-Security_Misconfiguration/ (Accessed: 4 March 2026).
- ❖ Saltzer, J.H. and Schroeder, M.D. (1975) 'The Protection of Information in Computer Systems', *Proceedings of the IEEE*, 63(9), pp. 1278–1308.
- ❖ Samba Project (2024) rsync Documentation. Available at: <https://rsync.samba.org/documentation.html> (Accessed: 4 March 2026).
- ❖ Stallings, W. (2018) *Network Security Essentials: Applications and Standards*. 6th edn. Harlow: Pearson.
- ❖ VMware (2026) VMware Workstation Pro 25H2 Documentation. Available at: <https://techdocs.broadcom.com/us/en/vmware-cis/desktop-hypervisors/workstation-pro/25H2.html> (Accessed: 4 March 2026).
- ❖ CloudArc (2025) Rsync: Streamlining Data Transfer and Synchronization. Available at: <https://blog.cloudarc.io/p/rsync-streamlining-data-transfer> (Accessed: 4 March 2026).
- ❖ Enterprotect 360 (n.d.) Layered Security: Multiple layers of Cybersecurity Working Together to Defend Your Business. Available at: <https://www.enterprotect.com/360/layered-security> (Accessed: 4 March 2026).

6. AI Usage Disclosure

This project was completed in accordance with Noroff's academic integrity guidelines. Artificial Intelligence (AI) tools were used solely as language assistance tools, not for content generation.

Tools Used

The following AI-assisted tools were used during the writing process:

- ❖ Grammarly (grammar and spelling correction)
- ❖ ChatGPT (used only for grammar review, sentence clarity improvement, and structural feedback)

Purpose of AI Usage

AI tools were used exclusively for:

- ❖ Grammar correction
- ❖ Spelling correction
- ❖ Sentence clarity improvement
- ❖ Structural suggestions
- ❖ Proofreading


No AI tools were used to:

- ❖ Generate original project content.
- ❖ Produce technical configurations
- ❖ Write implementation steps without manual author input.
- ❖ Fabricate test results or screenshots.

All technical configurations, design decisions, testing procedures, and screenshots were produced manually by the author.

Example of AI Usage

Below is an example of how AI was used to refine grammar and structure:



ChatGPT used to identify the issues in the text:





Figure 30A – Example of AI Grammar Review Using ChatGPT



Example 2

Original:

Virtualization is used in enterprise environments because of its efficient and good use of hardware resources, it has improved scalability, and even enhanced security through isolation between the systems (VMware, 2023).

Improved:

Virtualization is used in enterprise environments due to its efficient use of hardware resources, improved scalability, and enhanced security through system isolation (VMware, 2023).

What was improved:

- Fixed comma splice (two sentences improperly joined)
- Removed informal phrasing ("good use")
- Corrected spelling of scalability
- Improved flow and conciseness

Example 3

Original:

This layered approach does not eliminate all risks, but it significantly increases the effort required for a successful attack. Attackers would need to bypass multiple independent controls rather than exploiting a single weakness.

Improved (combined for academic flow):

Although this layered approach does not eliminate all risks, it significantly increases the effort required for a successful attack, as attackers must bypass multiple independent controls rather than exploit a single weakness.

What was improved:

- Combined sentences for smoother academic flow
- Slightly more formal tone

Figure 30B – Example of AI Grammar Review Using ChatGPT

Grammarly used to make sentences clearer and improve them:

The screenshot shows the Grammarly Proofreader interface. On the left, there is a document text with several highlighted areas. On the right, there is a sidebar with the title 'Grammarly Proofreader' and a 'Writing quality' score of 59. Below this, there are tabs for 'All', 'Grammar', 'Clarity', and 'Delivery'. The 'Clarity' tab is selected, showing a list of suggestions. Each suggestion has an 'Improve your text' button and a preview of the suggested change. The suggestions are:

- Improve your text**: The purpose of this project is to demonstrate the practical competence in virtualization, networking, operating systems and security. The demonstration in this project done by designing and implementing an realistic, enterprise, IT environment that follows very modern best practices. This infrastructure is built with a strong focus on security by design, to ensure that protection mechanisms are integrated from day one instead of using it as an afterthought.
- Improve your text**: The company LNXE Corp is described to be in its growth phase, where it has about ten employees. The employees work primarily in software development, and thus require a reliable, secure and scalable IT environment. The development team would depend heavily on having access to internal services, secure remote administration and even protected datastorage. As LNXE Corp is expected to grow, the infrastructure must be able to change and designed in a way that it allows future expansion without compromising any of the features like security or performance.
- Improve your text**: The Chief Executive Officer (CEO) of LNXE Corp have requested that this infrastructure to be implemented using a server-client model, while having a strong focus on security, OS hardening and following the best and recognized IT Security Practices.
- Improve your text**: To be able to meet these requirements, the solution to this project will be deployed using VMware Workstation Pro (25H2), hosting multiple virtual machines that will simulate realistic enterprise environment. Virtualization is used in enterprise environments because of its efficient and good use of hardware resources, it has improved scalability, and even enhanced security through isolation between the systems (VMware, 2023). By using a Type 2 Hypervisor, this project will simulate a realistic enterprise deployment, while still remaining suitable for an academic environment. VMware Workstation Pro (25H2) provides flexibility to deploy, to configure and to test multiple operating systems at the same time.
- Improve your text**: The core objective of this project is to demonstrate practical competence across multiple technical domains like and including virtualization, networking, operating system administration and security.
- Improve your text**: This project focuses on secure configuration of servers and clients, a controlled network access, andthrough strong authentication mechanisms. Ask AI sensitive data protected.

Figure 31A – Grammarly Sentence Improvement Example


Changes and fixes were double-checked to ensure they did not change any technical data or meaning:

The screenshot shows the Grammarly Proofreader interface. On the left, there is a document text with several highlighted areas. On the right, there is a sidebar with the title 'Grammarly Proofreader' and a 'Writing quality' score of 59. Below this, there are tabs for 'All', 'Grammar', 'Clarity', and 'Delivery'. The 'Clarity' tab is selected, showing a list of suggestions. Each suggestion has an 'Improve your text' button and a preview of the suggested change. The suggestions are:

- Improve your text**: The network infrastructure in this project for LNXE Corp will be designed to balance security, usability and scalability, but also support both server and client systems. The core in the design of this network is to protect internal resources from unauthorized access, but also ensuring a reliable connection for both the employees and administrators.
- Improve your text**: To be able to achieve this, the network is implemented as a private internal network with clear roles, access paths and security boundaries.

Figure 31B – Grammarly Sentence Improvement Example

Plagiarism was also checked; the text was found to be irrelevant.



The AI provided suggestions to improve sentence clarity and fix grammatical mistakes. The technical content remained authored and verified by the student.