



4TH EDITION

Windows Server 2025 Administration Fundamentals

A beginner's guide to managing and administering
Windows Server environments

Table of Contents

Preface

Part 1: Introducing Windows Server and Installing Windows Server 2025

1

Network Fundamentals and Introduction to Windows Server 2025

Technical requirements

Understanding hosts, nodes, and client/server architecture

What is a computer network?

Types of computer networks

Understanding computer network components

Understanding Computer Network Architectures

Overview of IP addressing and subnetting

IPv4 network addresses

IPv6 network addresses

[IPv4 subnetting](#)

[Getting to know the server](#)

[Understanding server hardware and software](#)

[Understanding server sizes, form factors, and shapes](#)

[Understanding Network Operating System](#)

[Windows Server overview](#)

[Linux Server overview](#)

[macOS Server overview](#)

[Overview and editions of Windows Server 2025](#)

[Windows Server eras overview](#)

[Windows Server 2025 overview](#)

[Windows Server 2025 editions](#)

[Key differences between Windows Server 2025 and Windows Server 2022](#)

[Minimum and recommended system requirements for Windows Server 2025](#)

[Chapter exercise 1.1 – downloading Windows Server 2025](#)

[Downloading Windows Server 2025](#)

Chapter exercise 1.2 – downloading Windows Admin Center

Downloading Windows Admin Center

Summary

Questions

Further reading

2

Installing Windows Server 2025

Technical requirements

Understanding disk partitioning and storage options

Understanding partition schemes

Overview of storage options

Accessing the advanced startup options

Exploring boot configurations and startup options

Understanding boot options in UEFI

Getting to know the startup process in BIOS

A different firmware program for booting modern computers

Understanding TPM

A crucial test for server hardware

GPT and the boot programs

A database for booting Windows OS

What is the bootloader?

What is the boot sector?

How to use the boot menu?

How does Safe Mode operate?

Windows setup and disk configuration errors

Installation options for Windows Server 2025

Understanding the role of your server

Pre-installation checks – resource compatibility checks

Which installation option for Windows Server 2025 should I choose?

Comparing Nano Server and Server Core

Using logs to diagnose installation failures

Network connectivity and domain joining

Activation and licensing issues

Various methods for deploying Windows Server 2025

Clean install

Deploying with the MDT

In-place upgrade

Migration

Deploying in Azure

Summary

Questions

Further reading

3

What to Do After Installing Windows Server 2025

Technical requirements

Understanding and managing devices and drivers, including Plug and Play, IRQs, and driver signing

Understanding computer devices and device drivers

Managing devices and device drivers

Customizing the Start menu for efficient navigation

[Working with devices and Device Manager](#)

[Understanding PnP, IRQ, DMA, and driver signing](#)

[Managing and optimizing registry entries and service accounts](#)

[Windows Server registry](#)

[Services Control Manager and Windows Server services](#)

[Accessing and managing the Windows registry and services](#)

[Performing initial server setup for better performance and security](#)

[Initial settings for Windows Server](#)

[Managing configuration drift with PowerShell Desired State Configuration](#)

[Validating hardware stability with memory testers and burn-in applications](#)

[Chapter exercise – performing an initial Windows Server configuration](#)

[Using Server Manager to configure the initial settings for Windows Server](#)

[How to use Server Configuration for Windows Server initial setup](#)

[Summary](#)

[Questions](#)

[Further reading](#)

[Part 2: Setting Up Windows Server 2025](#)

[4](#)

[Directory Services in Windows Server 2025](#)

[Technical requirements](#)

[Understanding the AD infrastructure in Windows Server 2025](#)

[Addressing the importance of AD](#)

[Core protocols and services supporting AD](#)

[Tools and roles for administering AD](#)

[Adding and configuring the AD DS role](#)

[Understanding DCs](#)

[Understanding domains](#)

[Understanding the Domain Tree](#)

[Understanding the Forest](#)

[Understanding the Child Domain](#)

[Understanding Operations Master Roles](#)

[Understanding the difference between domains and workgroups](#)

[Understanding trust relationships](#)

[Understanding functional levels](#)

[Exploring the concept of namespaces](#)

[Sites explained](#)

[Exploring replication](#)

[Understanding the schema](#)

[Microsoft Passport explained](#)

[Exploring DNS fundamentals and configurations in Windows Server 2025](#)

[Understanding how DNS works](#)

[Installing the DNS role](#)

[Understanding the role of hosts and lmhosts files](#)

[Understanding hostnames](#)

[Understanding DNS zones](#)

[Getting to know WINS](#)

[The UNC explained](#)

[Managing OUs and default containers](#)

[Understanding OUs](#)

[Default containers explained](#)

[Understanding hidden default containers](#)

[The purpose of default container types](#)

[Delegating authority within an OU](#)

[User and group management within AD](#)

[Domain accounts explained](#)

[Understanding the Local Accounts](#)

[The User Profiles Explained](#)

[Understanding Computer Accounts](#)

[Understanding Group Types](#)

[Getting to know default groups](#)

[Understanding group scopes](#)

[Group nesting explained](#)

Chapter exercise – installing the AD DS and DNS roles and promoting the server to a DC

Summary

Questions

Further reading

5

Adding Roles to Windows Server 2025

Technical requirements

Understanding server roles and features in Windows Server 2025

Roles and features overview

Role services explained

Understanding server features

An overview of Server Manager

Exploring application server roles and their implementations

Understanding the email server in Windows Server 2025

Understanding the database server

[Understanding the collaboration server](#)

[Understanding the monitoring server](#)

[Understanding the Data Protection Server](#)

[Configuring web services and their roles in Windows Server 2025](#)

[IIS Explained](#)

[WWW overview](#)

[Understanding an FTP](#)

[Worker processes and how to access them?](#)

[Installing more features for IIS](#)

[Sites overview](#)

[Ports overview](#)

[What is SSL?](#)

[How do certificates work?](#)

[Setting up remote access roles and their functionalities](#)

[How to use Remote Assistance](#)

[How does RSAT work?](#)

[Explaining RDS](#)

[How to manage RDS CALs](#)

[Setting up RDG](#)

[What is a VPN?](#)

[Explaining App-V](#)

[Understanding multiple ports](#)

[Deploying file and print services for network environments](#)

[File Services overview](#)

[PDS Role overview](#)

[What is a local printer?](#)

[Network printer explained](#)

[Understanding printer pooling](#)

[Internet printing overview](#)

[Understanding Web Printing Management](#)

[Understanding Printer Driver Deployment](#)

[Understanding User Rights and Permissions Management](#)

[NTFS permissions explained](#)

[Understanding Share Permissions](#)

[Configuring User Rights](#)

[Monitoring file server activities](#)

[Chapter exercise – installing webserver \(IIS\) and PDS roles](#)

[Setting up a Web Server \(IIS\) role](#)

[Installing a PDS role](#)

[Summary](#)

[Questions](#)

[Further reading](#)

[Part 3: Configuring Windows Server 2025](#)

[6](#)

[Group Policy in Windows Server 2025](#)

[Technical requirements](#)

[Understanding GP fundamentals in Windows Server 2025](#)

[GPO's default location](#)

[Managing Group Policy Objects \(GPOs\)](#)

[Managing administrative templates](#)

[Best practices for Group Policy Management](#)

Real-life applications of Group Policy

Exploring GP processing mechanisms and order of precedence

Configuring GPO settings

GPO application

GP editors overview

Local Group Policy Editor

Applying local GPOs

Exploring GPO settings categories

Chapter exercise – examples of GPOs for system administrators

Renaming the administrator account

Renaming the guest account

Blocking the Microsoft accounts

Prohibiting access to the Control Panel and PC settings

Denying access to all removable storage classes

Summary

Questions

Further reading

7

Virtualization with Windows Server 2025

Technical requirements

Understanding virtualization fundamentals in Windows Server 2025

Emphasizing the connection between Hyper-V and cloud computing

Virtualization modes

Performance considerations in virtualization

Adding and configuring the Hyper-V role on Windows Server 2025

Hyper-V architecture

Hyper-V installation requirements

Nested virtualization

Exploring Hyper-V Manager for VM administration

Understanding key Hyper-V Manager functions

Configuration settings in Hyper-V

[How to make and adjust VHDs](#)

[Adjusting the RAM of a VM](#)

[Virtual networks in Hyper-V Manager](#)

[Understanding checkpoints](#)

[VHD and VHDX formats](#)

[Migrating from VMware to Hyper-V](#)

[Adjusting VM settings](#)

[Working with VMs](#)

[Best practices for VM startup and recovery settings](#)

[Real-world applications of Hyper-V for modern IT environments](#)

[Chapter exercise – installing the Hyper-V role on Windows Server 2025](#)

[Summary](#)

[Questions](#)

[Further reading](#)

8

[Storing Data in Windows Server 2025](#)

Technical requirements

Understanding storage technologies and their evolution in Windows Server 2025

Exploring different storage types

Understanding ATA and SCSI interfaces

PCI and PCIe overview

Explaining local storage

Exploring storage architectures and their implications for network environments

Understanding block-level and file-level storage

Understanding how adapters and controllers operate

Data transmission in storage devices

Overview of storage protocols and their roles in data transmission and access

Communication protocols in storage devices

File-sharing protocols

HBA and FC switches

iSCSI hardware

Explaining S2D

[An introduction to dedup](#)

[Storage tiering](#)

[Overview of Network Automated Tiered Storage Control in Windows Server 2025](#)

[Real-world applications of storage technologies in managed security service providers](#)

[Managing server storage using Server Manager and Windows PowerShell](#)

[Managing storage using Server Manager](#)

[Managing storage with Windows PowerShell](#)

[Understanding RAID principles and configurations](#)

[RAID variants](#)

[RAID implementation methods](#)

[SDS overview](#)

[Fault tolerance with S2D](#)

[High Availability](#)

[Understanding primary storage concepts and optimizing storage solutions in Windows Server 2025](#)

[Understanding HDDs](#)

Understanding Solid-State Drives (SSDs)

Optical disk drives (ODDs) and optical disks (ODs)

Understanding basic disks

Understanding dynamic disks

Changing a basic disk to a dynamic disk

Optimizing disk performance

Getting to know the mount points

Filesystem overview

Mounting a VHD

DFS explained

Chapter exercise – enabling Dedup on Windows Server 2025

Summary

Questions

Further reading

Part 4: New and Enhanced Features in
Windows Server 2025

Active Directory Domain Services (AD DS) Enhancements

Technical requirements

Overview of AD DS enhancements in Windows Server 2025

Key enhancements in Active Directory Domain Services for Windows Server 2025

Significant security improvements and enhanced authentication mechanisms in Windows Server 2025

Integration with cloud services and hybrid environments

Implementing the 32k database page size for scalability

Optimized replication for large-scale environments

Enhanced backup and recovery mechanisms

Understanding schema updates and extending AD schema capabilities

Managing schema conflicts and versioning

Best practices for schema design and maintenance

Utilizing AD object repair for enhanced object management

Identifying and diagnosing AD Object issues

Repairing and restoring AD objects

[Best practices for AD object management and recovery](#)

[Practical applications of diagnostic tools in Active Directory](#)

[Key diagnostic tools for Active Directory management](#)

[Chapter exercise – implementing 32k database page size in Windows Server 2025](#)

[Summary](#)

[Questions](#)

[Further reading](#)

10

[Configuring SMB over QUIC in Windows Server 2025](#)

[Technical requirements](#)

[Introduction to SMB over QUIC in Windows Server 2025](#)

[Overview of SMB and QUIC](#)

[Historical context and evolution](#)

[Key benefits of implementing SMB over QUIC](#)

[Understanding security considerations and encryption protocols](#)

[Public key infrastructure \(PKI\)](#)

[How PKI secures network communications](#)

[Implementing PKI in Windows Server 2025](#)

[Best practices for PKI management](#)

[Encryption mechanisms in SMB over QUIC](#)

[Configuring security settings](#)

[Best practices for security](#)

[Optimizing SMB over QUIC performance](#)

[Network configuration for optimal performance](#)

[Hardware considerations](#)

[Performance tuning and monitoring](#)

[Troubleshooting SMB over QUIC implementations](#)

[Identifying common issues](#)

[Solutions and fixes](#)

[Preventative measures](#)

[Understanding the relevance of SMB over QUIC in different network environments](#)

Chapter exercise – configuring and enabling SMB over QUIC in Windows Server 2025

Summary

Questions

Further reading

11

Implementing New Security Enhancements in Windows Server 2025

Technical requirements

Overview of new security enhancements in Windows Server 2025

Understanding improved access controls

Overview of advanced threat detection

Microsoft Defender Antivirus for Windows Server 2025

What are automated response systems?

Enhancing authentication and authorization mechanisms

Biometric authentication explained

Implementing biometric authentication

[Overview of Conditional Access policies](#)

[Configuring Conditional Access policies](#)

[Understanding OAuth 2.0 integration in Windows Server 2025](#)

[Implementing OAuth 2.0](#)

[Securing communication channels with TLS and other protocols](#)

[TLS explained](#)

[Overview of other secure protocols \(HTTPS, IPsec, and SSH\)](#)

[How to monitor secure channels](#)

[Leveraging security features in Windows Server 2025 and beyond with Azure integration](#)

[Licensing and cost considerations for deploying advanced security features](#)

[Microsoft Defender for Servers Plan 1 and Plan 2](#)

[Implementing security best practices in Windows Server 2025](#)

[Overview of patch management](#)

[Understanding audit logging](#)

[Why regular security assessments?](#)

[Integration and requirements for Microsoft Defender for Endpoint on Windows Server 2025](#)

[Note on security baselines and monitoring drift](#)

[Chapter exercise – configuring firewall rules, enabling TLS encryption, and setting up audit logs](#)

[Exercise 11.1 – Configuring firewall rules](#)

[Exercise 11.2 – Enabling TLS encryption](#)

[Exercise 11.3 – Setting up audit logs](#)

[Summary](#)

[Questions](#)

[Further reading](#)

12

[Managing Updates with Hotpatching, Azure Arc, and More in Windows Server 2025](#)

[Technical requirements](#)

[Introduction to server hotpatching with Azure Arc in Windows Server 2025](#)

Overview of hotpatching

Benefits of using Azure Arc

Windows Server 2025 compatibility

Applying hotfixes and updates using hotpatching

Preparing for hotpatching

Executing Hotpatches

Post-patching validation

Managing the server lifecycle and updates efficiently

Lifecycle management strategies

Automating update management

Monitoring and reporting

Troubleshooting hotpatching implementations

Common issues and solutions

Diagnostic tools and techniques

Best practices for troubleshooting

Chapter exercise – setting up Azure Arc in Windows Server 2025

Summary

Questions

Further reading

Part 5: Managing and Maintaining Windows Server 2025

13

Tuning and Maintaining Windows Server 2025

Technical requirements

Understanding server hardware components and their roles in Windows Server 2025

Processor overview

Understanding memory

Understanding the disc

Understanding the network interface

Understanding 32- and 64-bit architectures

Understanding external drives

Impact of external USB drives on server performance

Understanding graphics cards

[Cooling essentials](#)

[Power supply basics](#)

[Exploring physical ports](#)

[Performance monitoring tools and methodologies in Windows Server 2025](#)

[Implementing a systematic approach to performance monitoring](#)

[Applying performance monitoring procedures](#)

[Establishing server baselines](#)

[Utilizing Performance Monitor, Resource Monitor, and Task Manager for performance tuning](#)

[What Does Performance Monitor do?](#)

[What does Resource Monitor do?](#)

[What does Task Manager do?](#)

[Monitoring with Azure Monitor for Arc-enabled servers](#)

[Interpreting performance counters for optimizing server performance](#)

[Creating Data Collector Sets](#)

[The purpose of performance logs and alerts](#)

[Chapter exercise – the performance logs and alerts service](#)

[Activating the service for performance logs and alerts](#)

[Navigating to the PerfLogs folder](#)

[Generating performance data logs](#)

[Creating performance counter alerts](#)

[Summary](#)

[Questions](#)

[Further reading](#)

[14](#)

[Updating and Troubleshooting Windows Server 2025](#)

[Technical requirements](#)

[Managing updates for the OS, drivers, and applications in Windows Server 2025](#)

[Keeping Windows Server up to date](#)

[Installing updates on Windows Server 2025](#)

[Updating Microsoft programs](#)

[Updating third-party programs](#)

Updating applications with Winget

Configuring Windows Update to check for device drivers

WSUS overview

Troubleshooting methodologies and best practices

Best practices, guidelines, and procedures

How do you troubleshoot effectively?

Comparing systematic and specific approaches in troubleshooting

Understanding troubleshooting procedures

Understanding ITIL

Implementing business continuity strategies in Windows Server 2025 environments

What is a DRP?

Business continuity and disaster recovery differences

How does data redundancy work?

What is clustering?

Maintaining BC through backup, restore, and DR planning

Backup and restore basics

[Restoring AD](#)

[Volume Shadow Copy Service explained](#)

[Understanding folder redirection](#)

[Power redundancy explained](#)

[Implementing a DRP for Windows Server 2025](#)

[Utilizing Event Viewer to monitor system logs and perform troubleshooting](#)

[Overview of Event Viewer](#)

[Chapter exercise - using Event Viewer to monitor and manage logs](#)

[Configuring centralized monitoring in Windows Server 2025](#)

[How do you apply filters to Event Viewer logs?](#)

[How do you modify the default location for logs?](#)

[Summary](#)

[Questions](#)

[Further reading](#)

[Part 6: Studying and Preparing for the AZ-800 Certification Exam](#)

15

Understanding Microsoft Certifications and Preparing for the AZ-800 Exam

The value of Microsoft certifications

The impact of Microsoft role-based certifications

Identifying the target audience for Microsoft certifications

Skills measured in Microsoft certification exams

AZ-800 Exam Study Guide

Deploy and manage AD DS in on-premises and cloud environments (30–35%)

Manage Windows Servers and workloads in a hybrid environment (10–15%)

Manage virtual machines and containers (15–20%)

Implement and manage an on-premises and hybrid networking infrastructure (15–20%)

Manage storage and file services (15–20%)

Success strategies and preparation tips for Microsoft certifications

Essential resources for Microsoft certification preparation

[Navigating the Microsoft certification exam registration process](#)

[Exam day guidelines for Microsoft certification](#)

[New validity and renewal requirements for Microsoft certifications](#)

[Summary](#)

[Further reading](#)

[Appendix](#)

[Index](#)

[Other Books You May Enjoy](#)

Preface

Windows Server 2025 is Microsoft's latest server operating system, continuing the legacy of the Windows NT family of operating systems built upon the Windows 11 platform. This version enhances the server operating system with robust performance, increased security features, and seamless cloud integration, empowering organizations to harness the full potential of modern IT solutions. The objective is clear: to democratize cloud capabilities and provide users with the tools necessary for success in an increasingly competitive digital landscape.

This book begins by introducing network fundamentals and Windows Server 2025, establishing a solid foundation in networking essentials before diving deep into the functionalities of Windows Server 2025. The journey unfolds systematically, moving from the installation process in [Chapter 2](#) to post-installation tasks in [Chapter 3](#). You will learn how to efficiently configure directory services and add essential roles in *Chapters 4 and 5*.

Later chapters will guide you through the intricacies of managing user and computer settings and introduce you to the world of server virtualization. You will then explore data storage solutions, followed by insights into **Active Directory Domain Services (AD DS)** enhancements.

Next, you will discover the configuration of SMB over QUIC and learn about the implementation of new security enhancements in Windows Server 2025. The book then delves into hotpatching with Azure Arc, followed by performance tuning and maintenance strategies.

Finally, the book equips you with the skills needed to manage updates and resolve issues effectively. The book concludes with a chapter that will guide you through the essentials of Microsoft certifications and provides valuable tips for exam success.

Through hands-on exercises, practical scenarios, and expert insights, this book aims to provide you with a comprehensive understanding of Windows Server 2025, empowering you to manage complex tasks efficiently. Each chapter includes a concept summary and questionnaire to reinforce your learning, ensuring that you have the knowledge and skills necessary to excel as a system administrator.

By the end of this book, you will be well prepared to administer and manage Windows Server 2025 confidently, with the added opportunity to pursue Microsoft certifications, such as preparing for the AZ-800 exam, should you choose to challenge yourself further.

Who this book is for

This book is designed for IT professionals, such as system administrators, network engineers, and IT managers, who are starting their journey in Windows Server 2025 administration. It is also suitable for those who wish to update their knowledge with the latest tools and features in Windows Server 2025. Whether you are responsible for managing server infrastructure, ensuring network security, or optimizing server performance, this book will provide you with the essential skills and knowledge to excel in your role. If you want to learn from this updated version of Windows Server, then this book is for you.

What this book covers

[*Chapter 1*](#), *Network Fundamentals and Introduction to Windows Server 2025*, introduces essential network concepts, including hosts, nodes, IP addressing, and subnetting. It covers **Network Operating Systems (NOS)** and provides an overview of Windows Server 2025, its editions, system requirements, and key improvements over Windows Server 2022. Practical exercises on downloading the installation media and Windows Admin Center are included.

[*Chapter 2*](#), *Installing Windows Server 2025*, guides you through various installation methods for Windows Server 2025, such as clean installations, network installations using WDS, and unattended installations with Windows ADK and MDT. It also covers in-place upgrades, network service migration, and testing in Azure, with practical exercises on setting up WDS and deploying Windows Server 2025.

[*Chapter 3*](#), *What to Do After Installing Windows Server 2025*, covers critical post-installation tasks, including managing device drivers, working with the Windows Server registry, and performing initial server configurations. Topics include setting up the computer name, network, firewall, and system updates, with practical exercises to reinforce your understanding.

[*Chapter 4*](#), *Directory Services in Windows Server 2025*, explores AD DS and **Domain Name System (DNS)**, covering domains, forests, domain controllers, and DNS functionalities. It includes managing user and computer accounts using **Organizational Units (OUs)** and groups, with a hands-on exercise on installing AD DS and DNS roles and promoting a server to a domain controller.

[*Chapter 5*](#), *Adding Roles to Windows Server 2025*, focuses on configuring server roles and features, including application server roles, web services, remote access, and file and print services. It covers best practices for selecting server hardware and monitoring performance, concluding with a practical exercise on installing web server (IIS) and print server roles.

[*Chapter 6*](#), *Group Policy in Windows Server 2025*, introduces **Group Policy (GP)** essentials, covering the configuration and management of **Group Policy Objects (GPOs)** on local servers and domain controllers. It includes GP processing, order of precedence, and using the **Group Policy Management Console (GPMC)**, with practical exercises on implementing GPOs.

[*Chapter 7*](#), *Virtualization with Windows Server 2025*, explores virtualization technology, focusing on Microsoft Hyper-V. It guides you through installing the Hyper-V role on Windows Server 2025, managing virtual environments with Hyper-V Manager, and creating and configuring **virtual machines (VMs)**, with a hands-on exercise on the installation process.

[*Chapter 8*](#), *Storing Data in Windows Server 2025*, covers storage technologies and their roles in server operations, including physical interfaces, disk controllers, and data storage methods. It explores network-based storage systems and protocols such as S2D, SDS, and iSCSI, with a practical exercise on enabling data deduplication in Windows Server 2025.

[*Chapter 9*](#), *Active Directory Domain Services (AD DS) Enhancements*, highlights advancements in AD DS for Windows Server 2025, focusing on scalability, security, and cloud integration. It covers robust authentication mechanisms, enhanced security protocols, and the 32k database page size. Practical instructions on implementing these updates and a hands-on exercise on the 32k page size are included.

[*Chapter 10, Configuring SMB over QUIC in Windows Server 2025*](#), provides an understanding of the SMB over QUIC protocol, its evolution, and its benefits. It covers security considerations, performance optimization, and troubleshooting techniques, with a step-by-step exercise on configuring and enabling SMB over QUIC in Windows Server 2025.

[*Chapter 11, Implementing New Security Enhancements in Windows Server 2025*](#), delves into advanced security and authentication features, including improved access controls, threat detection, and automated response systems. It covers biometric authentication, conditional access policies, and securing communication channels with TLS, HTTPS, IPSec, and SSH. A hands-on exercise on configuring firewall rules and enabling TLS encryption is included.

[*Chapter 12, Managing Updates with Hotpatching, Azure Arc, and More in Windows Server 2025*](#), guides you through implementing hotpatching techniques to ensure seamless server updates with minimal disruption. It covers server hotpatching, using Azure Arc to manage on-premises servers, and highlights benefits such as reduced downtime and enhanced availability. You'll learn to prepare servers for hotpatching, apply updates in real time without reboots, and validate patches. This chapter also discusses automating update management tasks, monitoring server health, and troubleshooting techniques. A hands-on exercise on configuring Azure Arc on an on-premises Windows Server 2025 instance is included.

[*Chapter 13, Tuning and Maintaining Windows Server 2025*](#), provides essential knowledge about server hardware selection and performance evaluation. It covers techniques for monitoring server performance, establishing performance baselines, and generating comprehensive reports. You will learn to identify and address potential performance issues proactively. The chapter concludes with a practical exercise on analyzing performance logs and setting up alerts.

[*Chapter 14, Updating and Troubleshooting Windows Server 2025*](#), delves into updating and troubleshooting within the Windows Server 2025 environment. It discusses the importance of a well-defined plan for managing updates, monitoring, and maintaining servers to ensure business continuity. You'll learn how to implement a backup and restore disaster recovery plan and perform updates to the operating system, server hardware, and third-party software. This chapter introduces the Event Viewer for examining logs and troubleshooting, concluding with a hands-on exercise on using the Event Viewer to monitor and manage logs.

[*Chapter 15, Understanding Microsoft Certifications and Preparing for the AZ-800 Exam*](#), serves as a comprehensive guide for candidates aspiring to attain Microsoft certifications related to Windows Server 2025. It covers the significance of Microsoft certifications, detailing the skills assessed in exams and the importance of role-specific certifications. Practical advice for passing certification exams, including study strategies and the exam registration process, is provided. The chapter also explores the AZ-800 certification exam objectives and includes valuable resources such as study materials and practice exams.

Appendix, Assessments, provides answers to the chapter questions. Each chapter includes questions to help reinforce the concepts and definitions. With this appendix, you can check your answers to those questions.

To get the most out of this book

You must have solid experience working with the Windows 10/11 operating system and have a solid knowledge of computer networks and NOSs.

Make sure you have a computer with a processor that supports virtualization technology and has a minimum of 8 GB or a recommended 16 GB of RAM.

Conventions used

There are several text conventions used throughout this book.

Code in text: Indicates code words in the text, database table names, folder names, filenames, file extensions, pathnames, dummy URLs, user input, and Twitter handles. Here is an example: “You also learned how to use the **regedit** tool.”

A block of code is set as follows:

```
# Example to check the status of Windows Defender Antivirus
$baselineStatus = $false
$currentStatus = (Get-MpPreference).DisableRealtimeMonitoring
if ($currentStatus -eq $baselineStatus) {
    Write-Host "Security baseline is intact."
} else {
    Write-Host "Security baseline drift detected. Review settings."
}
```

Bold: Indicates a new term, an important word, or words you see onscreen. For instance, words in menus or dialog boxes appear in **bold**. Here is an example: “In the **Enter Product Key** window, type the product key and click **OK**.”

Any command-line input or output is written as follows:

```
Get-VMNetworkAdapter -VMName <YourVMName> | Set-VMNetworkAdapter -MacAddressSpoofing On
```

Tips or important notes

Appears like this.

Get in touch

Feedback from our readers is always welcome.

General feedback: If you have questions about any aspect of this book, email us at customercare@packtpub.com and mention the book title in the subject of your message.

Errata: Although we have taken every care to ensure our content’s accuracy, mistakes happen. If you have found an error in this book, we would be grateful if you would report this to us. Please visit www.packtpub.com/support/errata and fill in the form.

Piracy: If you come across any illegal copies of our works on the internet, we would be grateful if you would provide us with the location address or website name. Please contact us at copyright@packt.com with a link to the material.

If you are interested in becoming an author: If there is a topic that you have expertise in and are interested in either writing or contributing to a book, please visit authors.packtpub.com.

Share Your Thoughts

Once you've read *Windows Server 2025 Administration Fundamentals*, we'd love to hear your thoughts! Please [click here to go straight to the Amazon review page](#) for this book and share your feedback.

Your review is important to us and the tech community and will help us make sure we're delivering excellent quality content.

Download a free PDF copy of this book

Thanks for purchasing this book!

Do you like to read on the go but are unable to carry your print books everywhere?

Is your eBook purchase not compatible with the device of your choice?

Don't worry, now with every Packt book you get a DRM-free PDF version of that book at no cost.

Read anywhere, any place, on any device. Search, copy, and paste code from your favorite technical books directly into your application.

The perks don't stop there, you can get exclusive access to discounts, newsletters, and great free content in your inbox daily

Follow these simple steps to get the benefits:

1. Scan the QR code or visit the link below

<https://packt.link/free-ebook/978-1-83620-501-2>

2. Submit your proof of purchase
3. That's it! We'll send your free PDF and other benefits to your email directly

Part 1: Introducing Windows Server and Installing Windows Server 2025

This part provides a comprehensive overview of Windows Server 2025, highlighting its distinctive features and installation methods. You will build a foundational understanding of core concepts and acquire practical skills to perform installations, upgrades, migrations, network-based deployments, and unattended installations.

This part contains the following chapters:

- [*Chapter 1, Network Fundamentals and Introduction to Windows Server 2025*](#)
- [*Chapter 2, Installing Windows Server 2025*](#)
- [*Chapter 3, What to Do After Installing Windows Server 2025*](#)

1

Network Fundamentals and Introduction to Windows Server 2025

In this chapter, we establish the foundation for your journey into **Windows Server 2025**.

Understanding the fundamentals is crucial, so let us get started!

We begin by exploring the essentials of computer networks. Imagine a bustling city of interconnected roads—devices (hosts) communicate with each other via these pathways. We will delve into concepts such as **hosts**, **nodes**, **client-server architecture**, and other network components. Additionally, we will demystify IP addressing and subnetting, which function like the postal codes of the digital world.

Next, we shift our focus to the **Network Operating System (NOS)**. Think of the NOS as the conductor of our network orchestra. We will discuss the hardware and software requirements for running a server. Furthermore, we will take an in-depth look at Windows Server 2025 itself—its editions, improvements over **Windows Server 2022**, and exciting new features. Imagine enhanced security protocols, seamless hybrid cloud integration, and cutting-edge technical support—all part of the Windows Server 2025 symphony.

Ready to roll up your sleeves? We will cover the detailed system requirements for Windows Server 2025. Additionally, this chapter will guide you through downloading and preparing the installation media. Consider it your backstage pass to server greatness.

Finally, we will equip you with practical skills. You will download and install Windows Server 2025's trusty sidekick, **Windows Admin Center**. This tool will be your companion as we venture into more advanced topics in the upcoming chapters. By the end of this chapter, you will have a solid grasp of networking principles, NOS essentials, and the specifics of Windows Server 2025. Buckle up—it's a transformative journey that will lead to you mastering server management in today's dynamic IT landscape!

In this chapter, we're going to cover the following main topics:

- Understanding hosts, nodes, and client/server architecture
- Overview of IP addressing and subnetting
- Getting to know the server
- Understanding Network Operating System (NOS)
- Overview and editions of Windows Server 2025
- Key differences between Windows Server 2025 and Windows Server 2022
- Minimum and recommended system requirements for Windows Server 2025

- Downloading Windows Server 2025 and Windows Admin Center

Technical requirements

To complete the exercises in this chapter, ensure you have a PC running **Windows 11 Pro**, equipped with a minimum of **8 GB of RAM**, **500 MB** of available disk space, and an **active internet** connection.

Understanding hosts, nodes, and client/server architecture

As you embark on this section, you might initially question the relevance of learning about computer networking when your primary interest lies in Windows Server. This concern is valid at first glance. However, as you progress deeper into the realm of Windows Server, you will increasingly recognize the importance of a solid understanding of computer networks.

This chapter covers foundational concepts—such as IP addressing, subnetting, and network components—that are integral to effectively managing server environments. By understanding these basics, you'll gain the context needed to approach **Windows Server** tasks with confidence, making it easier to handle real-world configurations and troubleshooting down the line. While the book will focus on Windows Server, these networking principles serve as the backbone of every server environment, giving you the tools to navigate server management with precision.

To grasp the significance of computer networks, let us revisit their origins. The necessity for resource sharing sparked the initial development of networking technologies many years ago, during the 1960s and 1970s. As demand grew, so did the advancement of these technologies, leading to the creation of comprehensive terms and concepts essential for describing computer networks. Thus, terms such as **network types**, **topologies**, **architectures**, and **components** emerged, marking computer networks as one of humanity's monumental communication innovations. The internet exemplifies the profound societal benefits of computer networks, connecting countless computers and bridging geographical distances in communication.

With this background, let us delve into the basics of computer networks.

What is a computer network?

According to the Merriam-Webster dictionary, a network is defined as *a group of people or organizations that are closely linked and work with each other*. Additionally, networking is described as *the exchange of information or services among individuals, groups, or institutions*. These definitions provide a simple, concrete basis for understanding computer networks.

In essence, a computer network is a group of computers connected through networking devices and media to share resources. These resources typically include data, network services, and peripheral devices. For instance, sharing files, applications, printers, and other peripherals is straightforward in a networked environment. It is essential to distinguish between what a computer network is and what it does. The former explains the structure and components, while the latter highlights the benefits and functionalities. *Figure 1.1* illustrates that a computer network comprises interconnected computers sharing resources.

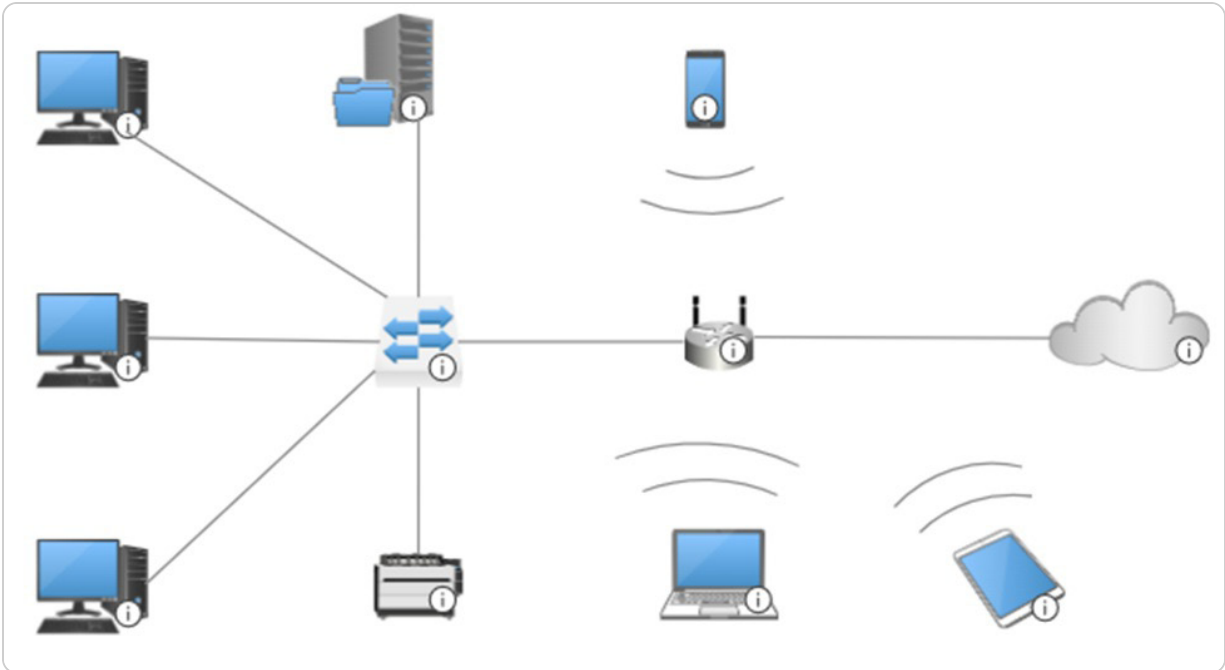


Figure 1.1 – A typical computer network

Computer networks come in various types, each serving different purposes and covering other areas. Let us explore these types individually.

Types of computer networks

Designing and building a computer network is a fascinating process closely tied to its definition. At its core, a computer network requires at least two computers. The number of computers and how they access shared resources determines the categorization of network types, which will be detailed in the following sections. Generally, computer networks are categorized based on the area they cover and their intended purpose. We will discuss some of the most common types of computer networks here.

Personal area network

A **Personal Area Network (PAN)**, depicted in *Figure 1.2*, connects and transmits data among devices within a private area, typically belonging to an individual. For example, in your home office, your laptop, smartphone, printer, and headphones might all be connected via Bluetooth or Wi-Fi. Often referred to as a **Home Area Network (HAN)**, a PAN uses technologies such as Bluetooth and Wi-Fi to interconnect devices.

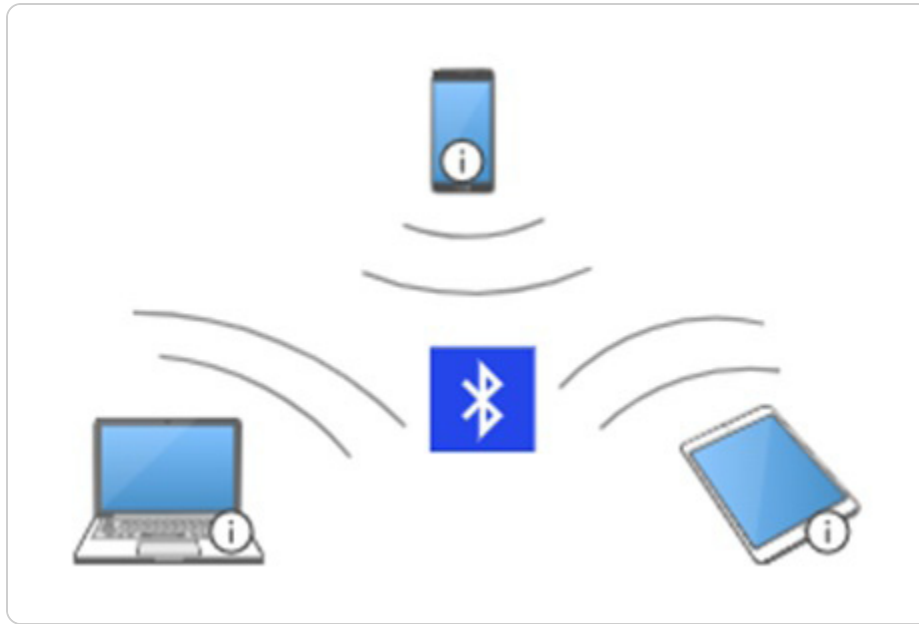


Figure 1.2 – A PAN

Another type of network is the **Local Area Network (LAN)**, which has a much more extensive coverage area compared to a PAN. Let us explore this in the next section.

Local area network

A LAN connects two or more computers within a local area, such as a single room, a floor, several floors, a building, or multiple adjacent buildings. LANs typically use a central device and networking media such as twisted-pair, coaxial, or fiber optic cables to interconnect computers. *Figure 1.3* illustrates an extended LAN that utilizes two switches to connect multiple devices. This configuration enhances network capacity and allows for greater scalability by enabling communication across different segments of the network. The use of multiple switches in the extended LAN ensures efficient data transfer between devices, improving overall network performance and reliability. This setup is commonly used in larger environments to accommodate growing network demands while maintaining optimal performance and minimal latency.

Furthermore, comparing PANs and LANs, a PAN is primarily dominated by portable devices such as smartphones, while a LAN mainly consists of fixed devices. Both cover local areas, but a LAN has a

broader range, potentially spanning an entire building or multiple buildings. A PAN is organized around an individual, whereas a LAN is organized around a specific site.

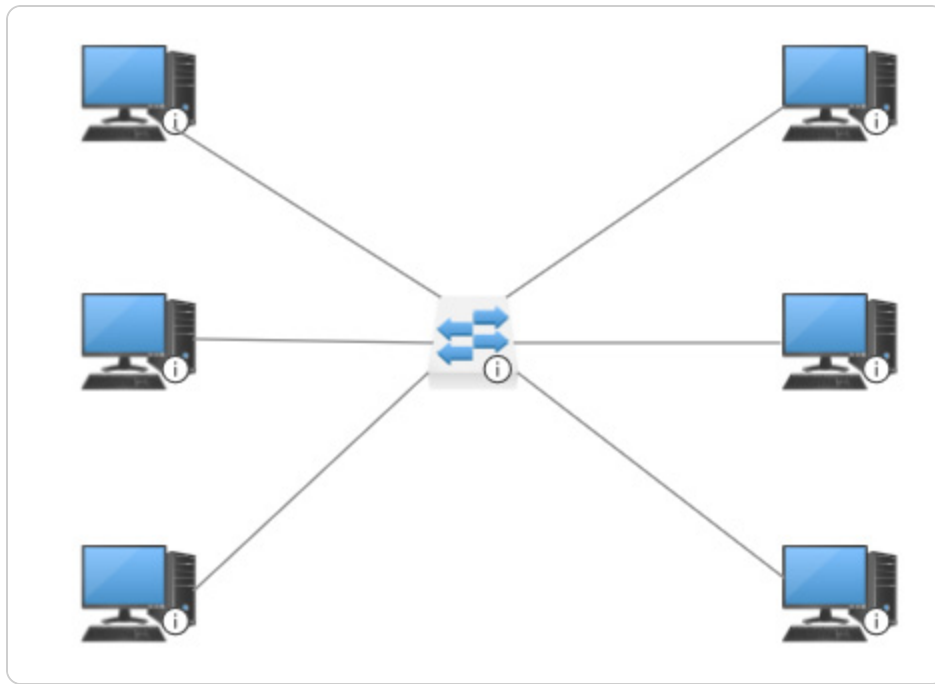


Figure 1.3 – A LAN

The next type of network we will examine is the **Metropolitan Area Network (MAN)**, which has even more excellent coverage than a LAN.

Metropolitan area network

A MAN, illustrated in *Figure 1.4*, connects multiple LANs within a town or city. MANs exist to facilitate resource sharing and access within a metropolitan area. They offer more excellent coverage than LANs but less than **Wide Area Networks (WANs)**. MANs are faster than both LANs and WANs, often using fiber optics and gigabit layer 3 switches for high-speed interconnection.

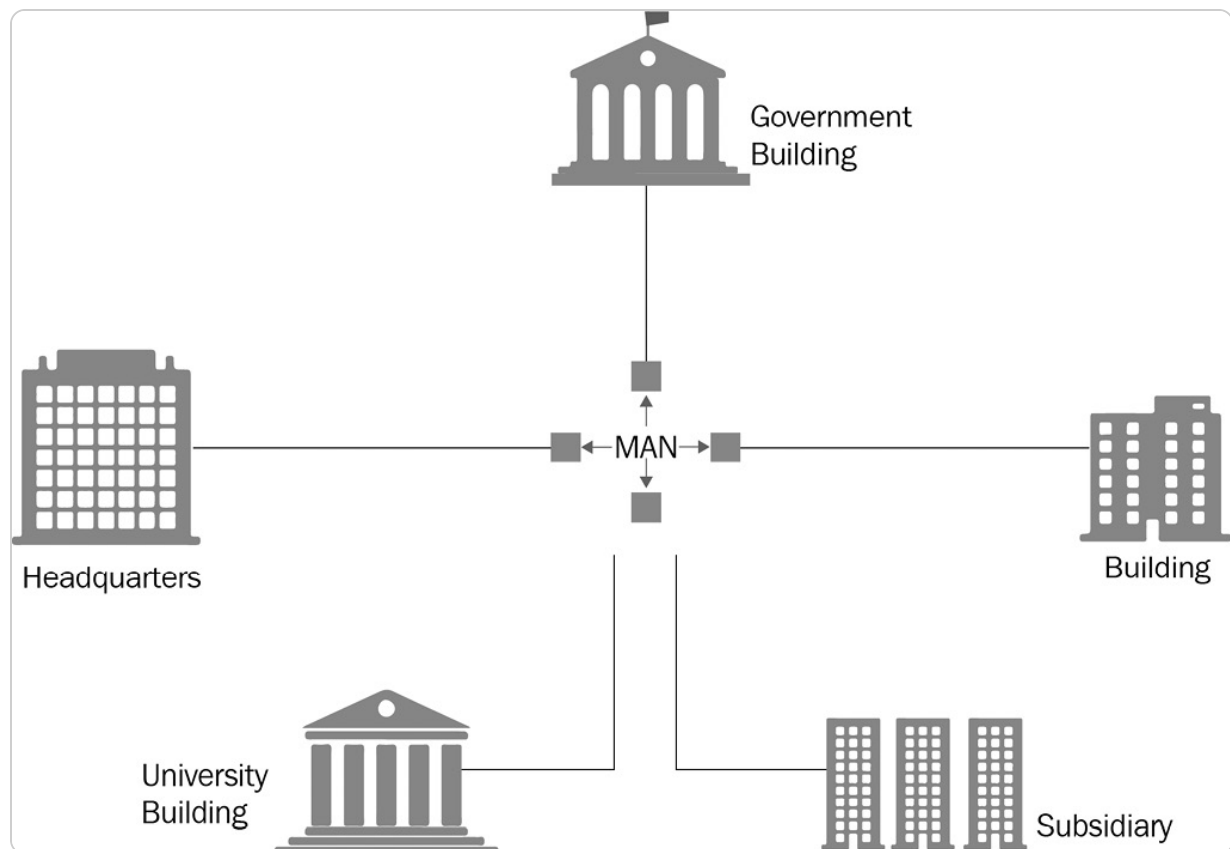


Figure 1.4 – A MAN

Finally, we will understand WANs, which have the most significant coverage.

Wide area network

A WAN, depicted in *Figure 1.5*, covers extensive geographic areas beyond the reach of LANs and MANs. WANs use dedicated telecommunication lines, such as telephone lines, leased lines, or satellites, making them accessible from geographic limitations. The internet is a quintessential example of a WAN.

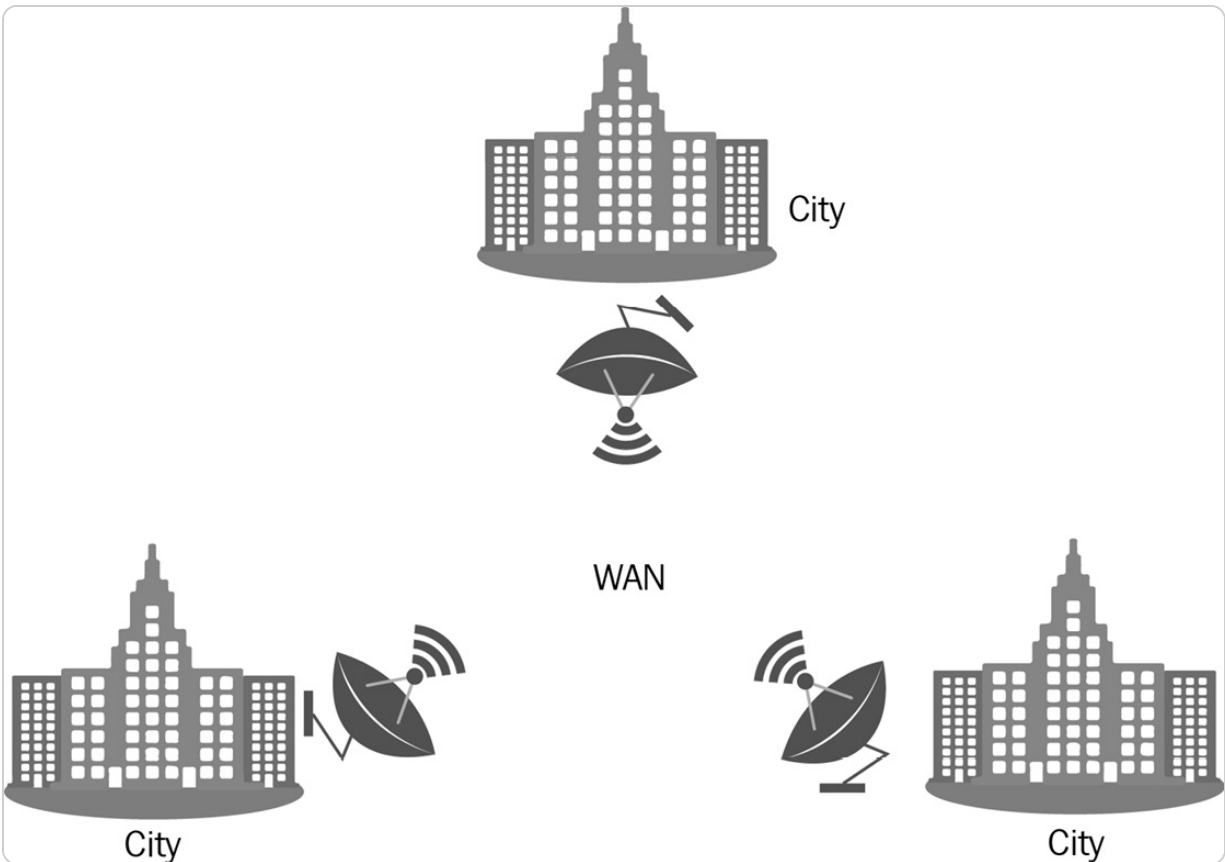


Figure 1.5 – A WAN

NOTE

You can learn more about the types of computer networks at <https://www.lifewire.com/lans-wans-and-other-area-networks-817376>.

After exploring the various types of computer networks, we will examine their underlying components.

Understanding computer network components

Just as **Personal Computers (PCs)** have their components, computer networks also consist of essential elements. While PCs and peripheral devices are familiar to most people, IT professionals focus on components such as networking devices, networking media, and NOSs.

First, let us clarify the roles of clients and servers within a computer network.

Understanding clients and servers

In the context of a computer network, **clients** and **servers** revolve around accessing and providing network resources. Clients typically initiate requests for resources, while servers are responsible for

delivering and managing access to these resources. Both play vital roles in network operations. For instance, as depicted in *Figure 1.6*, a server connected directly to a printer offers printing services to PCs acting as print requesters.

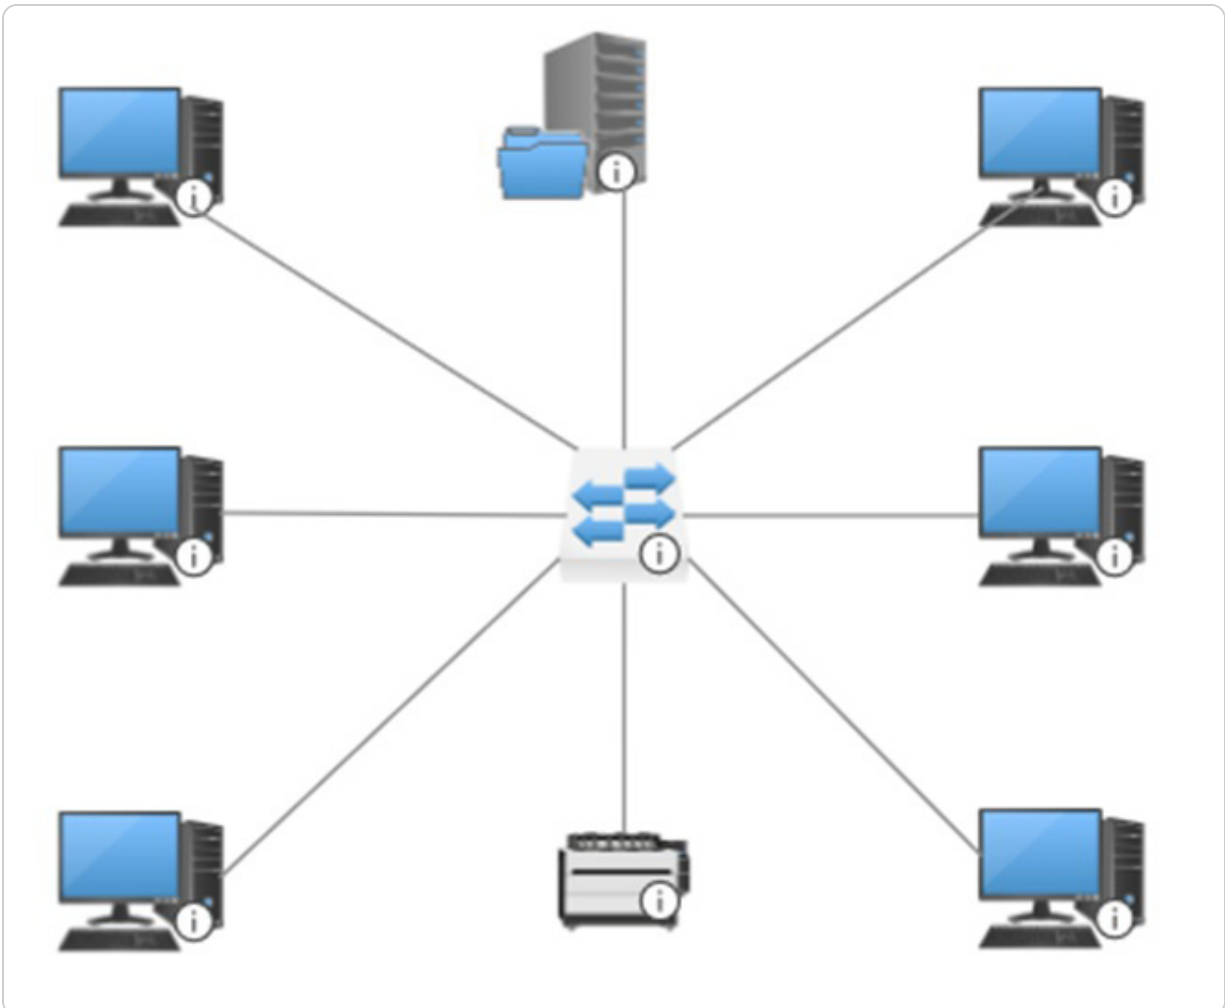


Figure 1.6 – Client and server in a computer network

NOTE

The term server originates from serve, indicating its role in providing beneficial services, as per the Merriam-Webster dictionary. In a computer network, servers fulfill this role by serving clients.

Although clients and servers are fundamental network components, their roles are defined differently in network terminology. Let us explore how they fit into the broader network structure.

Understanding hosts and nodes

Have you ever come across terms such as **hosts** and **nodes** and wondered about their distinctions? While they may initially seem similar, hosts and nodes serve distinct purposes in network communication. All hosts can be considered nodes, but not every node functions as a host. A host refers to any device with an assigned IP address on its network interface actively requesting or providing networking services. Typically, clients, servers, and routers operate as hosts.

NOTE

An **Internet Protocol (IP)** address is a logical sequence of decimal numbers separated by dots that uniquely identifies a host within a computer network.

On the other hand, a node is any device capable of receiving and transmitting network services but lacks an IP address assignment on its interface. Nodes typically have network interfaces used for management purposes. For example, in *Figure 1.7*, PCs and the file server act as hosts, while switches function as nodes.

In a Windows Server environment, understanding the distinction between hosts and nodes is essential, as Windows Server frequently operates in a client-server model. In this model, the server provides resources, and client devices connect to access these services. This setup supports effective resource distribution across networks and is fundamental to IT administration.

As you may know, switches operate by forwarding frames within a local network, ensuring data is efficiently transmitted between devices on the same network segment. In contrast, routers are responsible for forwarding packets across different networks, enabling communication between separate network segments or even across the internet. Both switches and routers play critical roles in network infrastructure, with switches focusing on local traffic management and routers handling more complex inter-network routing to ensure seamless connectivity. Understanding the distinction between these two devices is fundamental for effective network design and troubleshooting.

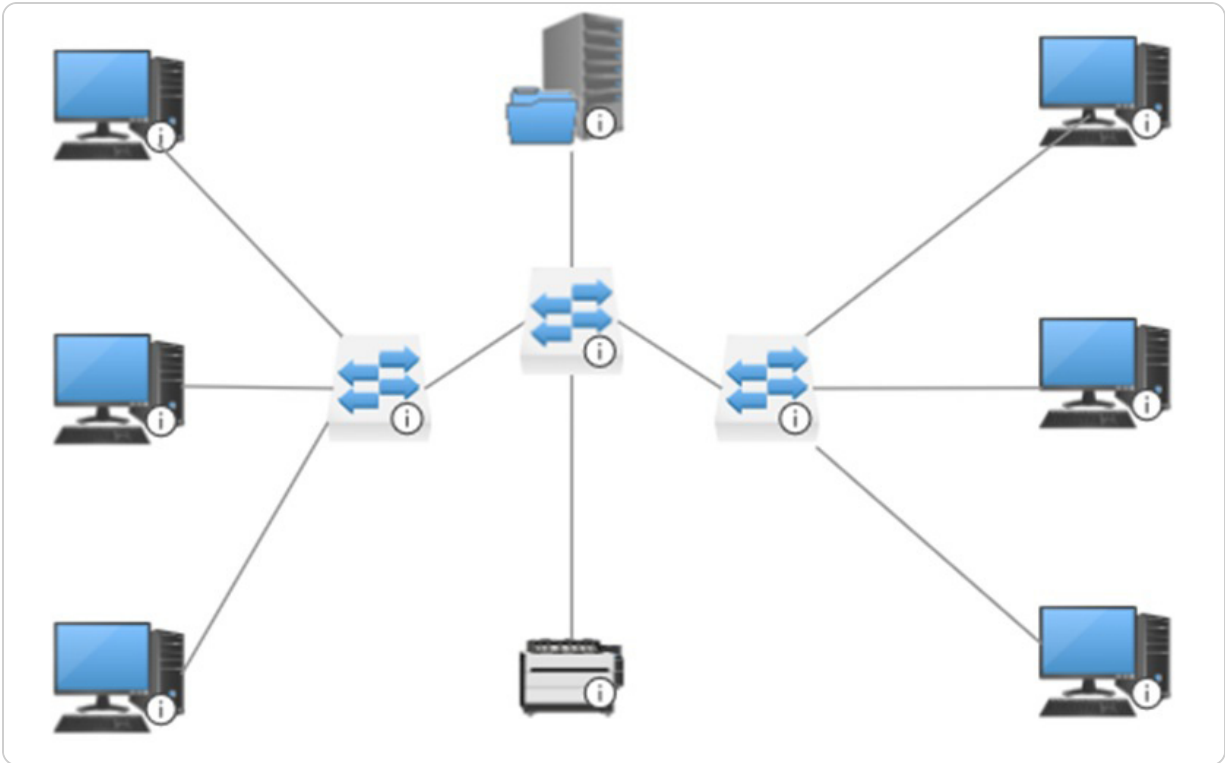


Figure 1.7 – Hosts and nodes in a computer network

Now that we understand a network and its essential components, we can explore network architectures.

Understanding Computer Network Architectures

Discussions about computer networks often involve exploring fundamental and overarching concepts, such as the components that comprise them. That includes considerations of network types based on coverage areas and physical and logical topologies governing their physical layout and structural organization. Computer network architecture encompasses a comprehensive framework that integrates elements such as physical and logical topologies, network components, communication protocols, and operational principles.

Moreover, computer network architecture serves as a design framework enabling computers to communicate using a request and response paradigm. The most prevalent network architectures include **Peer-to-Peer (P2P)** and client/server models.

Let us begin by exploring the P2P network architecture.

Peer-to-Peer (P2P) network architecture

In a **P2P network**, illustrated in *Figure 1.8*, hosts operate without predefined roles. Instead, they dynamically switch roles between the client and server based on their current network activities. For instance, if PC1 requests services from PC2, PC1 acts as the client while PC2 serves as the server. Conversely, if PC2 initiates a request to PC1, PC2 becomes the client and PC1 the server. PANs often exemplify P2P network setups.

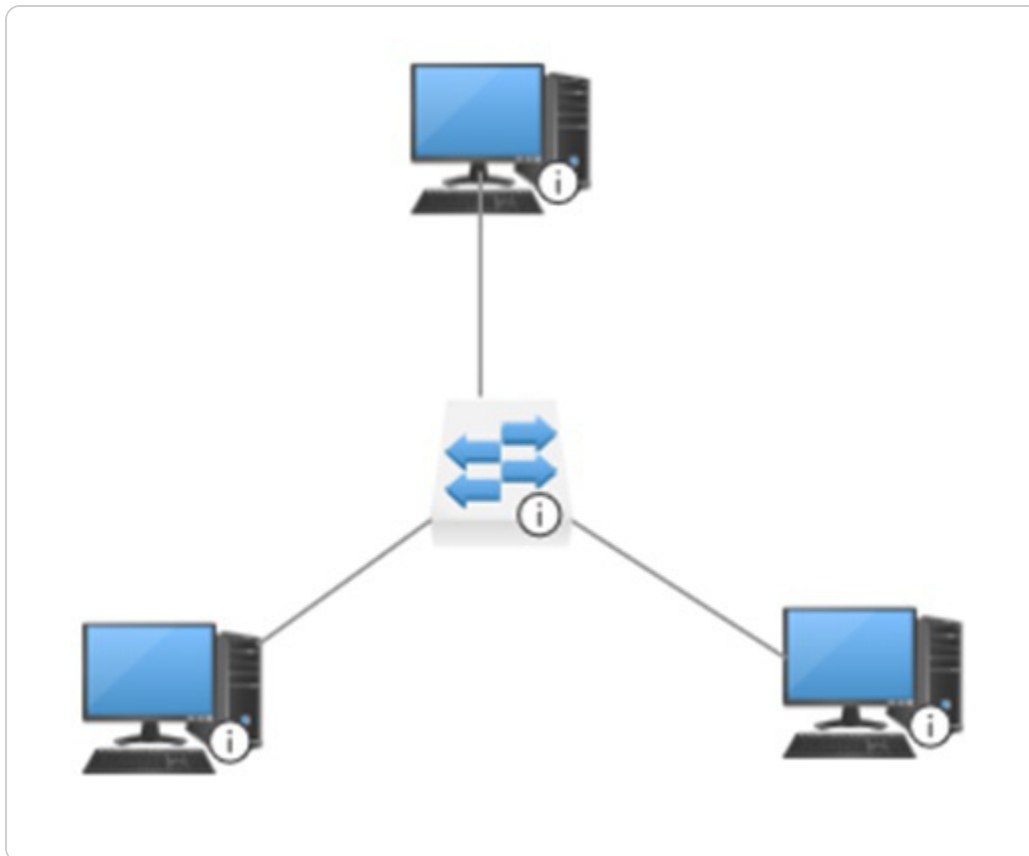


Figure 1.8 – A P2P computer network

NOTE

In the P2P network architecture, all hosts participate equally. This equality is a fundamental aspect of the model, where each host may assume the role of client or server as needed, making everyone feel included in the network.

The subsequent network architecture type is the client/server architecture.

Client/Server network architecture

In contrast, a **client/server network architecture**, presented earlier in *Figure 1.7*, designates specific roles to hosts. Clients are devices that request services, while servers are devices that provide services within the network. This structured approach to networking ensures efficient operations by clearly

defining the roles of each device. Typically, client and server responsibilities are designated to specific machines, allowing for streamlined communication and resource management across the network. This architecture is foundational in many enterprise environments where scalability and reliability are crucial.

Each component of networking, from hosts and nodes to client-server architecture, directly supports your work within Windows Server 2025. Whether configuring access permissions, setting up resource sharing, or monitoring network traffic, these concepts will recur throughout your server management tasks.

With a clearer understanding of network operations, the next section will delve into the fundamental requirement for computers to communicate within a network: the IP address.

Overview of IP addressing and subnetting

In order for a computer to effectively communicate within a network, an IP address is required, serving as its unique identifier on that network. Think of an IP address as a device’s unique identifier, similar to a postal address. In server management, IP addressing enables you to define the network structure and configure devices to communicate efficiently within and across networks.

In more complex networks, subnetting is used to define specific segments within the larger network structure, dividing it into smaller, more manageable parts. This segmentation enhances both security and performance—two key factors when managing Windows Server environments, where maintaining streamlined communication and data flow is essential.

Currently, two leading IP addressing technologies are recognized globally: IPv4 and IPv6. Despite the increasing prominence of IPv6, IPv4 remains the predominant addressing standard in internet traffic.

Let us begin by examining IPv4 network addresses.

IPv4 network addresses

IPv4, or **Internet Protocol version 4**, assigns addresses consisting of 32 bits organized into four octets, separated with dots for readability (e.g., 192.168.1.1). The designation “v4” denotes the fourth iteration of IP addressing, specified in IETF publication **RFC 791**. IPv4 addresses are categorized into classes—A, B, C, D, and E—based on their initial octet ranges, as shown in *Table 1.1*.

IPv4 classes	The IPv4 range of the first octet
A	1–127

B	128–191
C	192–223
D	224–239
E	240–255

Table 1.1 – IPv4 classes and their corresponding ranges

As you progress through this chapter, you may encounter areas that require further clarification, especially if you are new to IT. To address these challenges, it is crucial to focus on practical applications of the concepts discussed. Understanding how these foundational networking principles directly apply to Windows Server management will enhance your learning experience.

Now, let us explore the IPv6 addressing technology that was introduced to address the exhaustion of IPv4 network addresses.

IPv6 network addresses

IPv6, or **Internet Protocol version 6**, was designed to overcome the limitations posed by the exhaustion of IPv4 addresses, which are only 32 bits in length and provide around 4.3 billion unique addresses. In contrast, IPv6 employs a 128-bit address format, allowing for an astronomical total of approximately 340 undecillion unique addresses, as specified in IETF RFC 2460. This vast address space is represented in hexadecimal notation and segmented by colons, exemplified by addresses such as 2001:0DB8:85A3:0000:0000:8A2E:0370:7334.

The implementation of IPv6 not only provides an expansive range of addresses but also includes several features that enhance network efficiency and security. For instance, IPv6 supports auto-configuration, enabling devices to generate their IP addresses without the need for a DHCP server. That is particularly beneficial in dynamic environments where devices frequently connect and disconnect, such as in **Internet of Things (IoT)** applications. Additionally, IPv6 incorporates built-in security features such as IPsec, which provides end-to-end encryption and authentication, further safeguarding data transmission across networks.

To illustrate the practical benefits of IPv6, consider a smart home setup that includes various connected devices such as smart thermostats, lights, and security cameras. With the extensive address space of IPv6, each device can have a unique IP address, allowing for seamless communication and management without the complexities of address translation or the risk of address conflicts inherent in IPv4 networks. Furthermore, the ability to use multicast addressing in IPv6 facilitates efficient data

distribution, such as streaming video to multiple devices simultaneously without requiring multiple unicast streams.

IPv6 not only resolves IPv4’s address scarcity but also introduces significant advancements that improve network management, security, and scalability, making it a vital evolution in internet architecture as the number of connected devices continues to soar.

Next, we will delve into IPv4 subnetting, which plays a crucial role in identifying network addresses.

IPv4 subnetting

Subnetting involves logically partitioning a more extensive network into smaller subnetworks. A subnet mask is crucial for defining these subnetworks and identifying the network, host addresses, and broadcast addresses within each subnet. Default subnet masks, also known as classful networks, vary depending on the class of IPv4 addresses, as detailed in *Table 1.2*.

IPv4 class	Default subnet mask
A	255.0.0.0
B	255.255.0.0
C	255.255.255.0

Table 1.2 – IPv4 classful networks

NOTE

For further exploration of IPv4 addressing and related topics, visit this page on IPv4 exhaustion and classful networks:
<https://blogs.igalia.com/dpino/2017/05/25/ipv4-exhaustion/>.

The networking fundamentals covered in this section provide a critical framework for the tasks you will perform in Windows Server 2025. However, as we transition into Windows Server topics, we will focus more specifically on its functionalities and practical applications, ensuring that your understanding of these foundational concepts directly supports your server management skills. For additional support, consider leveraging resources such as online forums and targeted literature that can provide deeper insights into Windows Server administration.

Having covered fundamental networking concepts, including types, components, architectures, and addressing, it’s important to note that tools such as PowerShell and scripting are valuable for managing Windows Server environments. While this chapter doesn’t delve into these topics in detail,

they will be explored further in later chapters, providing you with essential skills for automating and streamlining server management tasks. The following section will introduce Windows Server and its core concepts.

Getting to know the server

Since we have already defined the server concept earlier, let us now focus on Windows Server in this section. Over its history, Windows Server has evolved from a primary file server to a sophisticated operating system capable of managing complex network environments such as corporate networks. It serves various network functions, such as domain controllers, web servers, print servers, and file servers, and it also acts as a platform for running enterprise applications such as **Exchange Server**, **SQL Server**, and **SharePoint Server**. With its robust performance and advanced security features, Windows Server now plays a pivotal role in shaping the landscape of cloud computing.

Understanding server hardware and software

As previously mentioned, computer hardware and software encompass both physical and logical components. Because servers are tasked with providing network services to clients, they require robust hardware. This hardware is designed to support advanced network services efficiently. Therefore, servers need high-quality components to ensure continuous service delivery and support for network operations. Servers differ from standard computers not only in their hardware but also in the specific types of services they provide. For instance, a database server necessitates significant memory capacity and storage space.

Key hardware components of servers include the **CPU**, **memory**, **disk subsystem**, and **network interfaces**, all of which significantly impact overall server performance. Monitoring the performance of these components is crucial to maintaining optimal server functionality under both regular and heavy workloads.

You will likely be familiar with these concepts already. However, the following list will jog your memory about the basics:

- **Central Processing Unit (CPU)**: Often referred to as the **processor**, this is a chip located on the server's motherboard that handles all processing and calculations. Modern CPUs are based on 64-bit architecture, which allows for more efficient data exchange between the CPU and RAM compared to older 32-bit architecture.
- **Random Access Memory (RAM)**: This serves as the server's working memory, utilized by Windows Server 2025 and its applications. The amount of RAM directly affects multitasking capabilities, enabling more applications to run simultaneously. Further details on RAM can be found in the *Understanding Memory* section of [Chapter 13, Tuning and Maintaining Windows Server 2025](#).
- **Disks**: Data storage in servers typically relies on disks, often organized into a disk subsystem. Disk performance, specifically read/write speeds, is critical as faster throughput enhances overall disk subsystem performance. Servers commonly employ **Solid**

State Drives (SSDs) and **Hard Disk Drives (HDDs)**. SSDs offer faster read and write speeds due to their lack of moving parts, while HDDs are known for durability and high-capacity storage.

- **Network interface:** This facilitates a server's connection to both LANs and the internet within an organization. Servers often feature multiple network interfaces, with higher network connection speeds enabling greater data throughput to and from the network.

Now that we have covered what constitutes a server, let us explore various server sizes, form factors, and configurations.

Understanding server sizes, form factors, and shapes

If we consider a server as essentially a computer, then similar principles governing the form factor of laptops also apply to servers. That raises the question: what exactly is a form factor? In hardware design, the form factor defines the size, shape, and technical specifications of an electronic device. Today, servers are available in three primary form factors, each tailored to specific operational needs:

- **Rack-mountable servers:** As the name suggests, these servers are designed to be mounted on racks. These servers function as general-purpose computers capable of supporting a wide range of applications and network services. Typically housed in on-premises server rooms or data centers, rack-mountable servers are secured to racks due to their weight and are depicted in *Figure 1.9*.



Figure 1.9 – An HP server in a rack

- **Blade servers:** These are modular units that enable multiple servers to be deployed within a compact space. Characterized by their slim design, blade servers typically include CPU, memory, network interfaces, and storage disks. They are commonly found in data

centers or facilities requiring high processing power, owing to their ability to house multiple servers on a single shelf.

- **Tower servers:** These resemble vertical-case PCs but are equipped with advanced hardware that provides significantly higher processing power compared to standard PCs. These servers are commonly utilized for testing purposes or to support local services in **Small Office–Home Office (SOHO)** environments.

NOTE

A 64-bit Windows Server installed on 64-bit hardware can handle twice the data compared to a 32-bit Windows Server running on 32-bit hardware.

In this section, we have explored server hardware components, such as CPU, memory, disk, and network interface, and discussed server sizes, form factors, and configurations. Furthermore, servers are equipped with an operating system that enables them to provide network services similar to conventional computers. Let us explore this aspect further.

Understanding Network Operating System

A **NOS** is specialized software designed to manage, maintain, and provide various services within a network environment. These services include file and application sharing, web services, authentication and authorization, access control, user and computer administration, configuration tools, resource management, and other network-related functions. Consequently, a NOS plays a crucial role in effectively managing network resources.

A NOS forms the foundation of server functionality, enabling centralized control of network resources and client-server interactions. Windows Server 2025 functions as a NOS, offering tools and features tailored for seamless device management, application hosting, and data handling. Understanding the role of a NOS is essential for fully leveraging Windows Server's capabilities.

Prominent examples of NOSs today are **Windows Server**, **Linux Server**, and **macOS Server**, each capable of delivering comprehensive network services. Let us delve into each of these systems individually.

Windows Server overview

Windows Server, a cornerstone of Microsoft's server product line, is renowned for its robust **Graphical User Interface (GUI)** and extensive capabilities in managing network resources. Since its inception in 1993, Windows Server has evolved to meet the demands of modern computing environments. The lineage began with Windows NT 3.5 in the early 1990s and formally started with Windows 2000 Server. Key milestones include the introduction of Windows Server 2008, which brought features such

as Server Core and Hyper-V, and the release of Windows Server 2016, which enhanced support for cloud integration, illustrating its adaptability.

Initially available for both 32-bit and 64-bit architectures, Windows Server transitioned exclusively to 64-bit architecture with the release of Windows Server 2012. The server's native filesystem remains the **New Technology File System (NTFS)**. However, Windows Server 2012 introduced the **Resilient File System (ReFS)**, primarily used in database applications due to its resilience and efficiency.

As organizations increasingly transition to cloud services, Windows Server remains relevant by providing hybrid solutions that seamlessly connect on-premises resources with cloud infrastructure. *Figure 1.10* illustrates the properties of the Windows c: drive, showcasing key attributes that support the management of file storage and system resources. With features such as Windows Admin Center and PowerShell, Windows Server empowers IT professionals to automate and streamline server management tasks, ensuring its continued relevance in an increasingly digital and cloud-centric world.



Figure 1.10 – NTFS continues to be used by Windows Server 2022

NOTE

For more details on ReFS, visit <https://docs.microsoft.com/en-us/windows-server/storage/refs/refs-overview>.

Linux Server overview

Linux is distinguished in the operating system landscape by its open source nature and extensive community support. Developed by Linus Torvalds in the early 1990s as a Unix-like system, Linux quickly gained popularity due to its robustness and flexibility. Licensed under the GNU **General Public License (GPL)**, Linux has evolved into numerous distributions tailored to various user needs. Linux servers, such as Ubuntu Server (illustrated in *Figure 1.11*), are prevalent in hosting web servers and powering supercomputers due to their security and scalability, both on-premises and in cloud environments.



Figure 1.11 – Downloading Ubuntu Server from ubuntu.com

NOTE

Learn about running the Linux subsystem on Windows Server 2025 at <https://docs.microsoft.com/en-us/windows/wsl/install-on-server>.

macOS Server overview

Although **macOS Server** has a smaller market share than Windows Server and Linux Server, it is renowned for its reliability and seamless integration with Apple’s ecosystem. As a Unix-based operating system, macOS Server adheres to Apple’s intuitive GUI design philosophy. Initially supporting both 32-bit and 64-bit platforms, macOS Server now exclusively operates on 64-bit platforms following Apple’s transition to Intel processors. Apple continues to release updates and provide support for macOS Server, maintaining its relevance in specialized environments.

NOTE

Explore more about macOS Server at <https://www.apple.com/macos/server/>.

In this section, we have gained insights into Windows Server, Linux Server, and macOS Server. In the next section, our focus will expand further into Windows Server, enhancing our understanding of its capabilities and administration.

Overview and editions of Windows Server 2025

If someone were to ask, “What is Windows Server?” you might respond with the following: Windows Server is a server operating system developed by Microsoft, part of the Windows NT family. In server environments—whether using Windows Server, Linux Server, or macOS Server—the primary objective is to ensure the system provides the necessary services to support an organization’s network. However, there are significant differences among these systems in terms of deployment processes, user interfaces, resource management, and server maintenance. These distinctions can significantly influence the overall efficiency and effectiveness of the server’s operation within an enterprise environment.

To understand its development, let’s explore the different eras of Windows Server and how it has evolved over the years.

Windows Server eras overview

For nearly 30 years, beginning with the masses era, that of Windows NT, Microsoft has consistently anticipated and integrated emerging needs within the server landscape. This foresight has driven a fascinating evolution of Windows Server, which I am eager to share with you. Pay close attention to the technological progressions and transitions over the years—they are truly impressive. The eras of Windows Server’s development are presented in the following table:

--	--	--	--

Server for the masses era (1996–2000)	Enterprise era (2000–2008)	Data center era (2009–2013)	Cloud era (2016– present)
Windows NT Server 3.5 Windows NT Server 4.0	Windows 2000 Server Windows Server 2003	Windows Server 2008 Windows Server 2012	Windows Server 2016 Windows Server 2019 Windows Server 2022 Windows Server 2025

Table 1.3 – Windows Server eras overview

In this section, you have gained an overview of Windows Server and its eras. In the next section, we will explore the steps to download Windows Server 2025.

Windows Server 2025 overview

Windows Server 2025, illustrated in *Figure 1.12*, is the latest release in Microsoft’s Windows NT family of server operating systems, generally available as of **November 2024**. Announced on 26 January 2024, Windows Server 2025 marks a shift from its predecessors—such as Windows Server 2016, 2019, and 2022—which were built on Windows 10. This new version is based on **Windows 11**, specifically **version 23H2** from the October 2023 update. Unlike Windows 11, Windows Server 2025 does not mandate TPM 2.0 for deployment, thus providing increased flexibility for varied deployment needs.

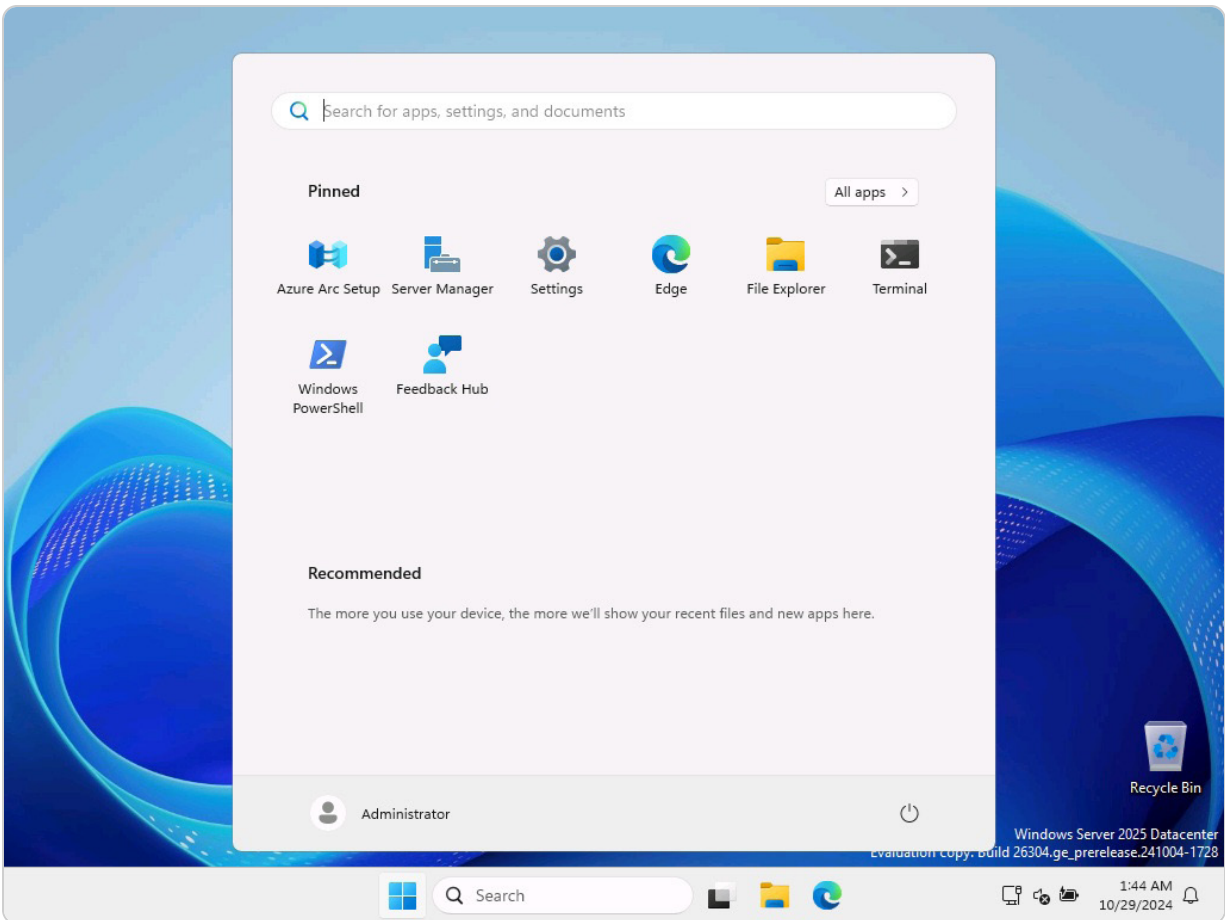


Figure 1.12 – Windows Server’s 2025 Desktop and Start menu

In today’s cloud-centric era, server operating systems must be designed with cloud capabilities in mind. This trend began with Windows Server 2016, which Microsoft aptly named **Windows Server for the cloud**. This focus has only intensified with subsequent releases, such as Windows Server 2022, which further enhanced its capabilities with improved security, flexibility, and robust support for hybrid deployments, mainly through innovations in the Windows Server 2022 Datacenter Azure edition. Each iteration, including Windows Server 2019, introduced significant features, such as System Insights, hybrid cloud tools, and enhanced security measures such as Storage Migration Service and Kubernetes support.

Microsoft has consistently evolved Windows Server to enhance security, connectivity, Azure integration, application platform capabilities, storage management, and other essential features. In today’s cloud-driven landscape, Windows Server continues to play a crucial role. Its integration with Microsoft Azure allows for the seamless management of hybrid environments, enabling organizations to leverage both on-premises and cloud resources. Features such as Windows Admin Center and PowerShell facilitate efficient cloud management, underscoring Windows Server’s ongoing relevance in a hybrid cloud world.

Building on this foundation, Windows Server 2025 introduces several new features designed to meet the evolving demands of modern cloud environments. Notably, **Active Directory Domain Services (AD DS)** has been enhanced with support for a larger 32k database page size, enhancing scalability and improving data handling for multi-valued attributes. New schema updates expand AD functionality, allowing administrators to repair objects with missing core attributes efficiently. Additionally, Windows Server 2025 introduces **Server Message Block over Quick UDP Internet Connections (SMB over QUIC)**, which allows the SMB protocol to run over QUIC, enhancing file-sharing performance and security across all editions. Security enhancements within AD DS further bolster network defenses, ensuring robust protection against contemporary threats.

Moreover, Windows Server 2025 pioneers **hot-patching** capabilities, enabling the seamless application of security patches without requiring server restarts. This feature minimizes downtime and enhances system uptime, which is critical for maintaining operational continuity. Specific requirements must be met to utilize hot-patching, including the use of **Virtualization-Based Security (VBS)** enclaves. VBS enclaves provide an isolated environment that enhances security by protecting critical processes from potential threats, ensuring that hot-patching can be performed securely and effectively.

Leveraging **AI-driven management optimizations**, Windows Server 2025 offers enhanced management capabilities that provide administrators with proactive insights and operational efficiencies. These advancements underscore Microsoft's commitment to delivering a server OS that sets new standards for performance, security, and manageability in today's IT landscapes, solidifying its position at the forefront of cloud computing innovation.

NOTE

Unlike **Remote Server Administration Tools (RSATs)**, which rely on traditional methods, Windows Admin Center is a modern server management platform built on web technologies. It offers a more streamlined and user-friendly experience, with an interface closely resembling that of Azure, enabling seamless navigation for administrators familiar with cloud environments. Windows Admin Center provides a comprehensive set of tools for managing servers and infrastructure efficiently. After installing Windows Server 2025, a dialog box for Windows Admin Center automatically appears, granting users free access to this powerful management interface, further simplifying server administration tasks. It can be downloaded from <https://www.microsoft.com/en-us/evalcenter/download-windows-admin-center>.

Next, let us explore the various editions available in Windows Server 2025.

Windows Server 2025 editions

Windows Server 2025 is available in various editions, each designed to meet specific organizational requirements. The primary editions include the following:

- **Datacenter Edition:** Ideal for extensive virtualization and cloud environments, this edition offers unlimited virtual instances and a range of advanced features. Enhancing security and performance requires VBS enclaves.

- **Standard Edition:** This edition is tailored for smaller organizations with fewer virtual instances. It includes all essential server features needed for efficient operation. VBS enclaves are also recommended for this edition to ensure robust security.
- **Azure Edition (for VM evaluation only):** Specifically designed for evaluating Windows Server within Azure VMs, this edition allows organizations to test and assess its capabilities in the cloud. VBS enclaves are utilized to provide a secure evaluation environment.
- **Annual Channel for Container Host:** Focused on container workloads, this edition ensures efficient deployment and management of containerized applications. VBS enclaves support secure and isolated container environments.

These editions clearly outline the options available within Windows Server 2025, each catering to different needs and use cases. Next, we will explore how Windows Server 2025 compares to its predecessors.

Key differences between Windows Server 2025 and Windows Server 2022

In today's rapidly evolving technological landscape, new products are often initially perceived as mere incremental updates to their predecessors. This sentiment might also apply to Windows Server 2025 at first glance, suggesting only superficial enhancements over Windows Server 2022. However, a closer examination of Windows Server 2025 reveals substantial improvements and new features. The following subsections aim to highlight some of the most significant differences and advancements between these two operating systems.

When comparing Windows Server 2025 with its predecessor, Windows Server 2022, several notable advancements and enhancements become apparent:

- **AD DS enhancements:**
 - **Windows Server 2025:** Introduces a larger 32k database page size for AD, enhancing scalability. New schema updates extend AD functionality, and enterprise administrators can repair objects with missing core attributes.
 - **Windows Server 2022:** Focused on improving AD management and schema flexibility.
- **SMB over QUIC:**
 - **Windows Server 2025:** SMB can now be configured across all editions using QUIC, enhancing file-sharing performance and security
 - **Windows Server 2022:** Introduced initial support for SMB over QUIC
- **Security enhancements:**
 - **Windows Server 2025:** Includes hypervisor-based code integrity, enhanced Secured-core server, and hardware-enforced stack protection. Default support for TLS 1.3 improves network security.
 - **Windows Server 2022:** Featured improvements in security protocols, including Secured-core Server and enhanced TLS support.

- **Hot-patching support:**
 - **Windows Server 2025:** Implements hot-patching capabilities, allowing for seamless updates without downtime
 - **Windows Server 2022:** Continued support for robust update management tools
- **AI-based management:**
 - **Windows Server 2025:** Offers enhanced management capabilities with AI-driven optimizations
 - **Windows Server 2022:** Introduced initial AI-driven management tools
- **Platform flexibility:**
 - **Windows Server 2025:** Focuses on dynamic routing and improved service account management
 - **Windows Server 2022:** Introduced **Dynamic Source Routing (DSR)** and improvements in virtualized time zones
- **Windows Admin Center enhancements:**
 - **Windows Server 2025:** Features advanced management capabilities, including automated extension life cycle management and customizable VM information views
 - **Windows Server 2022:** Supported enhanced Windows Admin Center tools
- **Kubernetes support:**
 - **Windows Server 2025:** Enhances support for Kubernetes environments with advancements in container management
 - **Windows Server 2022:** Provided initial support for Kubernetes and improvements in container orchestration

NOTE

Hotpatching requires Azure Arc connectivity, which enables the integration of on-premises servers with Azure services. Other dependencies include Azure Update Management to ensure that updates are managed and deployed efficiently, compatibility with specific server roles and features (not all roles may support hot-patching initially, so it's important to verify compatibility), and sufficient system resources to ensure that hardware and network resources are adequate to handle hot-patching processes.

Understanding these differences between Windows Server 2025 and Windows Server 2022 will help you determine which version best suits your needs. Now, let us examine the minimum and recommended system requirements.

Minimum and recommended system requirements for Windows Server 2025

Before discussing the specific system requirements, it's important to distinguish between the minimum and recommended hardware specifications. The minimum requirements allow for the basic installation and operation of the OS, while the recommended specifications ensure optimal performance and user experience. Understanding these requirements is crucial for selecting the

appropriate hardware to match the intended usage and workloads. According to Microsoft's publications, Windows Server 2025 maintains hardware requirements similar to those of its predecessors.

The following are the minimum system requirements:

- **Processor:** 1.4 GHz 64-bit processor
- **RAM:** 512 MB (2 GB for the Desktop Experience installation option)
- **Disk space:** 32 GB
- **Network:** Ethernet adapter capable of at least 1 gigabit throughput
- **Graphics device and monitor:** Capable of Super VGA (1024 x 768) or higher resolution
- **Other hardware:** DVD drive (for installations from DVD media), keyboard, mouse (or compatible pointing device), TPM, and internet access

The following are the recommended hardware requirements:

- **Processor:** 2.0 GHz 64-bit processor or higher
- **RAM:** 32 GB or more
- **Disk space:** 256 GB SSD and 1 TB HDD
- **Network:** At least 1 gigabit Ethernet NIC
- **Graphics device and monitor:** Capable of super VGA (1024 x 768) or higher resolution
- **Other hardware:** DVD drive, keyboard, mouse (or compatible pointing device), TPM, and internet access

This section provides insights into the minimum and recommended hardware requirements for Windows Server 2025, reflecting its alignment with previous versions while ensuring robust performance for modern server environments. Understanding these requirements is crucial for making informed decisions about hardware investments to meet your server deployment needs. In the following sections, we will explore the new features and enhancements introduced in Windows Server 2025.

Chapter exercise 1.1 – downloading Windows Server 2025

In this chapter’s exercise, you will be guided through the process of downloading Windows Server 2025. This version introduces several advanced features and improvements designed to enhance server performance, security, and management. By carefully following the outlined steps, you will learn how to access and download the installation files, ensuring that you are prepared to work with the latest version of Microsoft’s server operating system. This exercise is essential for IT professionals looking to stay up to date with industry standards and leverage the new capabilities of Windows Server 2025 in their environments.

Downloading Windows Server 2025

To download Windows Server 2025 on your Windows 11 computer, follow these steps:

- 1. Press the Windows key + R to open the **Run** dialog.
- 2. Type **Microsoft-edge:** and press *Enter*.
- 3. In Microsoft Edge, click on the address bar and type <https://www.microsoft.com/en-us/evalcenter>, then press *Enter*.
- 4. On the **Evaluation Center** page, click **Windows Server** on the horizontal menu at the top. Then, select the **Windows Server** option. From the list of available versions, choose **Windows Server 2025**, as shown in *Figure 1.13*.

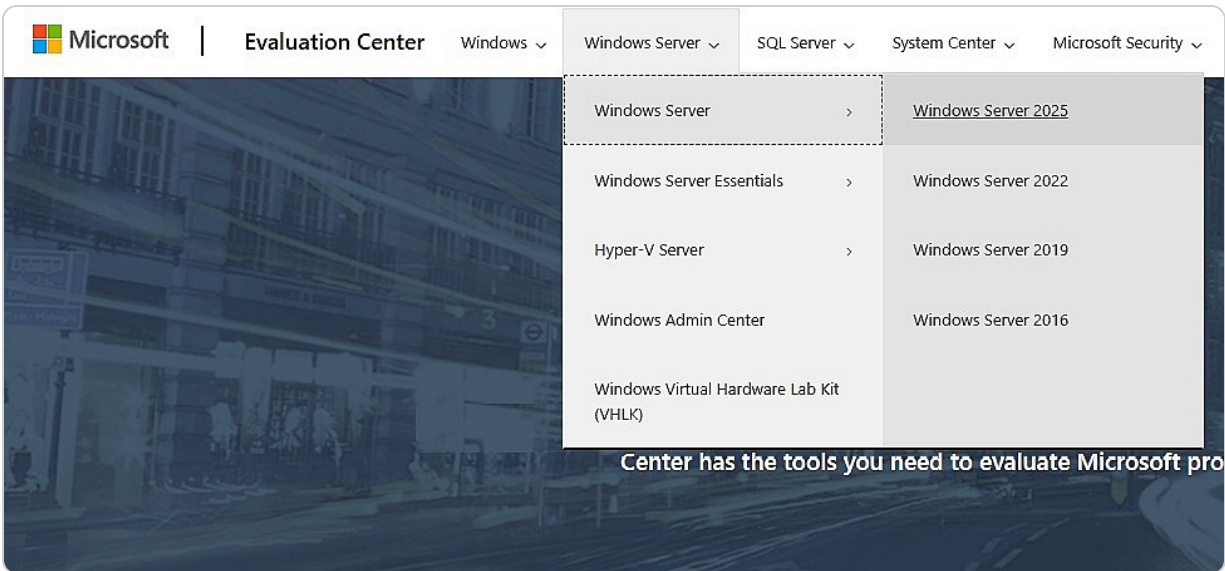


Figure 1.13 – Selecting Windows Server 2025 from the list

- 5. In the **Get started for free** section of the **Windows Server 2025** page, select **Download the ISO** as your product experience.
- 6. Complete the form as shown in *Figure 1.14*, then click **Download now**.

Evaluate Windows Server 2025

Microsoft Windows Server 2025 prepares you for tomorrow while delivering the security, performance, and flexibility you need today. Be more productive with easier networking, faster storage, and hybrid cloud capabilities that adapt to your needs. Get ahead of what's next with forward-looking security, and AI-ready compute.

Resources

- [Release notes](#) and [system requirements](#)
- [Microsoft Tech Community: Windows Server](#)
- [Windows Server technical documentation](#)

**Participation in this trial comes at no cost for the 180 day duration of the trial. Once the 180 day trial period expires, the trial instance will deactivate. No additional costs will be incurred at the end of the trial period. Customers can purchase a license and convert the license to full product after the trial period. Trial descriptions accurate as of November 2024 and are subject to change in the future.*

Register for your free trial today

Complete the form below.

* First name

* Last name

* Email

* Company name

* Country/Region

* Company size

* Job role

* Phone

Questions/Comments

Download now

Figure 1.14 – Registering for the free trial to evaluate Windows Server 2025

7. Next, please select your Windows Server 2025 download by choosing the language.
8. Shortly after, the Windows Server 2025 download will begin. If not, you may want to click the **Download** button.

NOTE

After downloading Windows Server 2025, you will need to burn the ISO file to a USB flash drive to create a bootable USB. You can follow this guide on how to burn an ISO file to a USB drive: <https://www.lifewire.com/how-to-burn-an-iso-file-to-a-usb-drive-2619270>. Once this process is complete, you will be ready to proceed with installing the Windows Server 2025 evaluation version.

Chapter exercise 1.2 – downloading Windows Admin Center

In this chapter's exercise, you will be guided through the steps required to download and set up Windows Admin Center. This tool serves as a centralized management interface for Windows Server environments, providing a streamlined approach to administering multiple servers and devices. By following the detailed instructions, you'll learn how to obtain and install Windows Admin Center, setting the foundation for efficient server management and configuration through its intuitive, web-based interface. This exercise is crucial for IT professionals seeking to enhance their ability to manage Windows Server resources with greater ease and control.

Downloading Windows Admin Center

To download Windows Admin Center on your Windows 11 computer, follow these steps:

1. Open Microsoft Edge and navigate to <https://www.microsoft.com/en-us/cloud-platform/windows-admin-center>.
2. Click the **Download Windows Admin Center** button on the Windows Admin Center site, as shown in *Figure 1.15*.

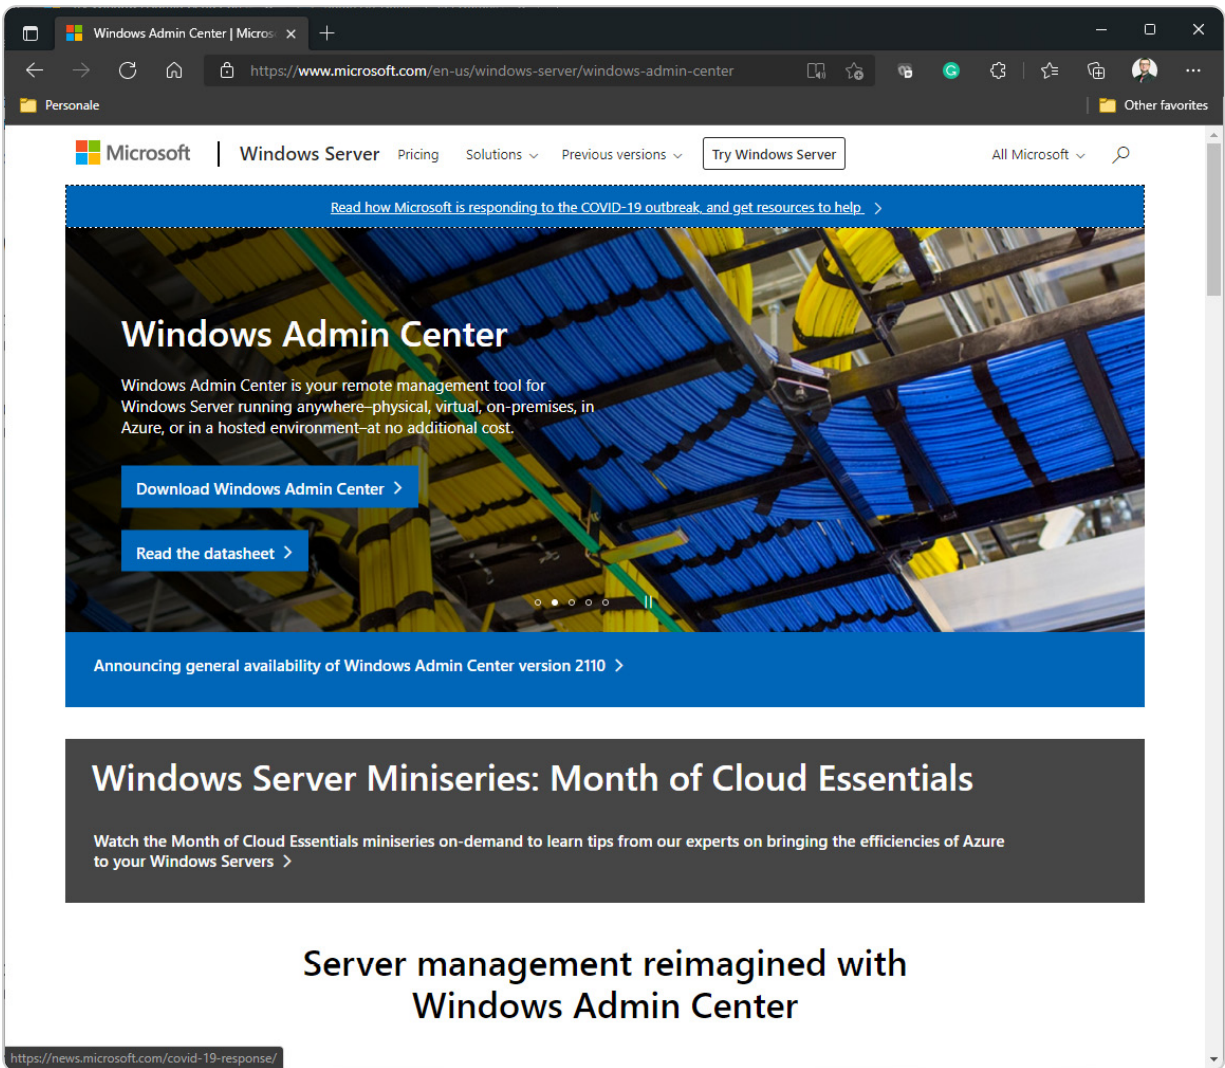


Figure 1.15 – Windows Admin Center download site

3. When prompted, choose either **Open** or **Save As** to download the file.
4. After the download is complete, proceed with the installation by following the steps in the **Windows Admin Center Setup** wizard.

These exercises were designed to reinforce the chapter's key concepts by providing practical experience. By completing them, you gained firsthand knowledge of how to install and manage server software and tools, which are crucial skills for successfully administering a modern server environment. This practical application not only solidified your understanding of the theoretical content but also prepared you for real-world scenarios in server management.

Summary

In this chapter, you have explored the foundational concepts of computer networks and gained insights into Windows Server. Specifically, you have been introduced to various types, components,

and architectures of computer networks, as well as IP addressing and subnetting. Additionally, you have learned about server hardware and software, different server sizes, form factors, and NOSs. This chapter also provided a comprehensive overview of the Windows Server timeline.

In the *Windows Server eras overview* section, you examined the evolution of the various versions of Windows Server 2025. You also compared the differences between Windows Server 2022 and Windows Server 2025, including their minimum and recommended system requirements. Furthermore, you delved into Windows Server 2025, Linux Server, and the macOS Server NOS.

The chapter also included practical exercises that guided you through downloading Windows Server 2025 from the Technet Evaluation Center portal and Windows Admin Center from the WAC portal, thus enhancing the lab-oriented learning experience. With the knowledge gained in this chapter, you should now understand what a computer network is, be able to identify different network architectures, and comprehend IP addressing and subnetting. Additionally, you will be able to recognize key hardware components, understand NOSs, and appreciate the historical evolution of Windows Server.

In the following chapter, you will learn about installing Windows Server 2025.

Questions

1. **True or false:** The computer network architecture is a design that enables computers to communicate using the request-response paradigm.
2. **Fill in the blanks:** _____ usually requests access to resources, and _____ is responsible for providing and managing access to the resources.
3. **Multiple choice:** Which of the following are considered to be types of computer networks? (Choose all that apply)
 - PAN
 - LAN
 - MAN
 - WAN
 - All of the above
4. **True or false:** Windows Server is Microsoft's operating system and is part of the Windows NT family.
5. **Fill in the blank:** _____ can provide network services such as domain controllers, web servers, print servers, and file servers.
6. **True or false:** The subnet helps to identify a specific network within the overall network.
7. **Multiple choice:** Which of the following are considered to be network architectures? (Choose two)
 - Peer-to-Peer (P2P)
 - Client/server

- Network Operating System (NOS)
 - Network topology
8. **True or false:** The CPU, memory, disk, and network are the critical system components that affect the overall performance of your servers.
9. **Fill in the blanks:** The _____ represents the physical component of a server, while the _____ represents the logical component.
10. **Multiple choice:** Which of the following are considered to be IP-addressing technologies? (Choose two)
- IPv2
 - IPv4
 - IPv6
 - IPv8
11. **True or false:** Windows Admin Center is a new server management app introduced with Windows Server 2022.
12. **Fill in the blank:** _____ technology has enabled easy-to-build, deploy, and run application images.
13. **Single choice:** What is the new server management app introduced with Windows Server 2025?
- Windows administrative tools
 - Windows PowerShell
 - Windows Admin Center
 - Active Directory Administrative Center
14. **Short answer:** Discuss the importance of a well-designed network architecture in supporting modern business operations.
15. **Short answer:** Explain the role of each critical system component (CPU, memory, disk, and network) in optimizing server performance.

Further reading

- *What is computer networking?:* <https://www.ibm.com/topics/networking>
- *What's new in Windows Server 2025:* <https://learn.microsoft.com/en-us/windows-server/get-started/whats-new-windows-server-2025>
- *Linux vs. Windows Server: The Ultimate Comparison:* <https://phoenixnap.com/blog/Linux-vs-Microsoft-windows-servers>
- *Windows Admin Center overview:* <https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/understand/windows-admin-center>

2

Installing Windows Server 2025

This chapter will guide you through the installation of **Windows Server 2025**, a powerful and versatile operating system for servers. You will learn how to perform different types of installation, such as **clean installation**, **network installation** using **Windows Deployment Service (WDS)**, **unattended installation** using the **Windows Assessment and Deployment Kit (Windows ADK)**, **Microsoft Deployment Toolkit (MDT)**, in-place upgrades, migration of network services to a new server, and testing Windows Server 2025 in **Azure**.

You will follow precise and detailed instructions, supported by helpful graphics, to complete the installation process with ease and efficiency. This chapter will help you master the skills and knowledge needed to install Windows Server 2025 on your servers quickly and effectively.

At the end of the chapter, you will practice setting up WDS, a valuable tool for deploying Windows operating systems over the network. That will give you practical experience with one of the installation methods covered in the chapter.

In this chapter, we're going to cover the following main topics:

- Understanding disk partitioning and storage options
- Exploring boot configurations and startup options
- Installation options for Windows Server 2025
- Various methods for deploying Windows Server 2025
- Perform clean installation, network deployment, in-place upgrade, migration, and Azure-based deployment

Technical requirements

To practice the skills learned in this chapter, you will need the following resources:

- A computer with **Windows 11 Pro**, a minimum of 16 GB of RAM, 1 TB of **disk space**, and an internet connection
- A virtual machine running **Windows Server 2012 R2 Standard** (Desktop Experience), with at least 2 GB of RAM, 100 GB of disk space, and an internet connection
- A virtual machine running **Windows Server 2022 Standard** (Desktop Experience), with at least 4 GB of RAM, 100 GB of disk space, and an internet connection

Understanding disk partitioning and storage options

Installing new operating systems is a routine task for a system administrator. This task involves several critical steps, such as preparing the installation media, executing the OS installation, checking the installation results, and setting up the initial server configuration. These steps are essential for laying the groundwork for further operations. While some servers may come with preinstalled operating systems, the system administrator's expertise is often required to ensure the server has the most suitable OS to meet specific needs.

Before we start the installation process, let us review the importance of partition schemes in organizing disk partitions.

Understanding partition schemes

Disk partitioning is the process of dividing a physical disk into logical sections called partitions. Each partition can have a different filesystem, such as **New Technology File System (NTFS)** or **Resilient File System (ReFS)**, and store various types of data. Partitions can also be used to create separate volumes, which are logical units of storage that can span multiple disks. The partition scheme is the technique that determines how these partitions are created and managed on the disks. There are two main partition schemes:

- **Master Boot Record (MBR)**: This is an older partitioning scheme that is now considered outdated and no longer recommended for modern systems. MBR operates on a 512-byte disk sector and supports only 4 primary partitions or 1 extended partition containing up to 26 logical partitions. It uses **Logical Block Addressing (LBA)** to manage disks with a maximum size of 2 TB. While MBR was once useful for multiboot systems, it has several limitations that make it incompatible with today's technology. Its 2 TB size restriction is insufficient for many modern storage devices, and the limited number of partitions can be a bottleneck for more complex setups. Additionally, MBR lacks the advanced redundancy and recovery features found in newer partitioning schemes. These shortcomings led to the development of the GUID Partition Table, which offers better scalability, support for larger drives, and improved reliability.
- **GUID Partition Table (GPT)**: This is a modern partition scheme that overcomes MBR's drawbacks. GPT uses a 128-bit **Global Unique Identifier (GUID)** to identify resources. It supports block sizes from 512 bytes and above, with a common default of 4,096 bytes. Each partition entry is 128 bytes. GPT is part of the **Unified Extensible Firmware Interface (UEFI)** standard, which replaces the older BIOS to work with modern hardware. GPT is resilient and can handle up to 9.4 **zettabytes (ZB)** of disk storage and 128 partitions per disk. GPT also provides better reliability and security features, such as a protective **MBR** and a **CRC32 checksum**. With MBR already explained, it's important to understand CRC32, a checksum algorithm used to detect errors in data transmission or storage. CRC32 generates a unique value, known as a checksum, derived from the data's contents. This value is then compared to the original checksum to verify the integrity of the data. If the two values do not match, it suggests possible data corruption, meaning the data may have been altered or compromised during transmission or storage, highlighting a failure in maintaining data accuracy.

To install Windows Server 2025, utilizing a GPT partition scheme is essential. This necessity arises from the requirement for UEFI, which supersedes the traditional BIOS and exclusively boots from GPT disks. UEFI not only enables quicker and more secure boot processes but also supports advanced functionalities such as **Secure Boot** and **BitLocker**, enhancing the overall system security and efficiency. Moreover, Secure Boot and BitLocker are key security features that work together to protect a system. Secure Boot safeguards the boot process by allowing only authorized and verified software to run, blocking untrusted or malicious code, such as rootkits, from compromising the system. BitLocker complements this by providing full-disk encryption, ensuring that data remains secure even if the storage device is physically removed or accessed without authorization. Together, these features enhance system security, protect sensitive information, and prevent unauthorized access, ensuring both the integrity and confidentiality of data.

In Windows Server 2025, you can create and manage disk partitions using the **Disk Management tool** or the **Diskpart command-line utility**. Alternatively, the **Windows Setup wizard** can be employed during the installation process to create and format partitions. If you need to convert an existing MBR disk to GPT, it is essential to first delete all partitions on the disk, which will result in data loss. Therefore, it is strongly advised that all data be backed up before proceeding with the conversion.

In addition to choosing a partition scheme, you can also customize the boot settings to facilitate the installation process. **Boot settings** are options that tell your computer how to start the operating system from different sources. For example, you can boot from a DVD, a USB flash drive, or a network server. To change the boot settings, you need to access the **Basic Input/Output System (BIOS)** or UEFI interface and modify the boot order or enable the boot menu. We will discuss the boot settings in more detail in the *Exploring boot configurations and startup options* section.

Note

To create a bootable USB flash drive, consider using the **Windows 7 USB/DVD Download Tool**, which you can download from Microsoft's official page at <https://www.microsoft.com/en-us/download/windows-usb-dvd-download-tool>.

One of the key aspects of setting up your server is choosing the right storage options that can deliver high performance, availability, and scalability. Windows Server 2025 provides various storage features that can help you achieve different goals and requirements. In the following section, we will discuss these storage features and how they can benefit the server environment.

Overview of storage options

Beyond disk partitioning, Windows Server 2025 offers various storage options to enhance your server's performance, availability, and scalability. Key storage features include the following:

- **Storage Spaces:** This feature allows the creation of virtual disks from a pool of physical disks, offering different resiliency levels such as *simple*, *mirror*, or *parity* to safeguard data against disk failures. It also supports tiered storage, automatically moving frequently accessed data to faster **Solid State Drives (SSDs)** and less accessed data to **slower HDDs**.
- **Storage Spaces Direct:** This feature enables the formation of a shared storage pool from local disks across a server cluster, facilitating highly available and scalable storage solutions such as **hyper-converged infrastructure (HCI)** or **software-defined storage (SDS)**. HCI represents a contemporary IT framework that merges computing, storage, and networking into a cohesive system managed through software. This integration streamlines management and scalability by consolidating hardware resources and virtualization into a single, unified platform, leading to enhanced efficiency and cost reductions. Conversely, SDS is a storage management strategy that decouples storage hardware from the controlling software. By leveraging virtualization, SDS enables dynamic management and allocation of storage resources, offering increased flexibility, scalability, and cost-effectiveness while facilitating more efficient and automated storage operations across various hardware setups.
- **Storage Replica:** This feature provides data replication between servers or clusters using synchronous or asynchronous replication, enabling robust disaster recovery solutions such as **stretch clusters** or **site-to-site replication**. Stretch clusters offer a high-availability solution by distributing a single cluster across multiple physical locations or data centers. This arrangement ensures uninterrupted operation and data redundancy, even if one site fails, as the cluster continues to function across the other locations. Site-to-site replication, meanwhile, synchronizes data between two geographically distant sites, ensuring that essential data is continuously updated and accessible at both locations. This approach bolsters disaster recovery and data resilience by maintaining an up-to-date backup, facilitating continuous data availability, and safeguarding against disruptions at any single site.

After completing the installation process, understanding advanced startup options becomes crucial. These options provide valuable functionality beyond the initial setup. Let us delve into them.

Accessing the advanced startup options

Windows Server 2025 does not have the *F8* option for restoring the server OS. Instead, you need to access the **Advanced startup** options from the **Settings** menu. To do this, you need to perform the following steps:

1. Click the **Start** button on the desktop.
2. From the **Start** menu, select the **Settings** icon.
3. On the **Settings** screen, find and click the **System** option.
4. From the list of options, select the **Recovery** option.
5. On the right-hand side of the screen, click the **Restart now** button under the **Recovery options** section, as shown in *Figure 2.1*.

 Figure 2.1 – Navigating to Advanced startup in Windows Server 2025

Figure 2.1 – Navigating to Advanced startup in Windows Server 2025

6. We will restart your device to save your work dialog box. Click again on the **Restart Now** button to confirm. Choose the reason and click **Continue**.

7. After the system restarts, select **Troubleshoot** from the *Choose an option* screen.
8. On the **Advanced options** screen, as shown in *Figure 2.2*, you can select various options to recover or repair your server OS.

Figure 2.2 – Advanced Options in Windows Server 2025

Figure 2.2 – Advanced Options in Windows Server 2025

This section provided an overview of the partition schemes, boot options, and advanced startup options in Windows Server 2025. The following section will explain the different server installation options in more detail.

Exploring boot configurations and startup options

Before a computer can load the operating system, it must go through a booting process that involves initializing hardware components and loading system software. This process is managed by firmware, either BIOS or UEFI, depending on the motherboard and hardware. Both BIOS and UEFI are responsible for configuring boot options, including boot order, boot mode, and boot device selection.

Boot options significantly impact the server's performance and its interaction with other devices and networks. Understanding the differences between BIOS and UEFI, along with their respective advantages and disadvantages, is crucial. In this section, we will delve into the boot options available in BIOS and UEFI and how they can be configured and customized in Windows Server 2025.

Understanding boot options in UEFI

To boot your system correctly, it is essential to understand the boot options available in the UEFI (Unified Extensible Firmware Interface), which has largely replaced the legacy BIOS in modern systems. UEFI is a firmware interface that initializes hardware and loads the operating system efficiently. You can access UEFI settings during startup by pressing specific keys, which may vary based on the manufacturer. Common keys include *F2*, *F10*, *Delete*, or *Esc*.

Unlike older systems, UEFI offers advanced features such as Secure Boot, faster boot times, and support for larger hard drives using GPT. To ensure a smooth installation process, configuring the boot order and verifying UEFI settings are vital for a successful installation of Windows Server 2025. Within the UEFI settings, set the primary boot device to your installation media (USB or DVD) to initiate the Windows Server setup correctly.

Additionally, enabling Secure Boot is highly recommended to enhance security. This feature allows only trusted software to load during the boot process, helping prevent malware from altering it. Secure Boot is often required to comply with modern security standards, so ensure that you're using a GPT-partitioned drive, as Secure Boot does not support MBR partitions.

By confirming these configurations before installation, you can significantly reduce the likelihood of installation failures caused by mismatched boot settings or security incompatibilities, resulting in a more secure and seamless setup for Windows Server 2025.

When you enter the BIOS, you will see various boot options. Let us look at them:

- **Installation Media (DVD):** A common type of installation media is a bootable DVD. To use this media, you need to prepare your computer to boot from the DVD drive. That requires you to access the BIOS settings and change the boot order. First, insert the bootable DVD into the drive that can read it. Then, enter the BIOS settings and select the DVD drive as the primary boot device. After saving the changes and exiting the BIOS, your computer will boot from the DVD and launch the installation process.
- **USB Flash Drive:** Another way to install an operating system is by using a bootable USB flash drive. This device must have at least 8 GB of storage space and be compatible with your computer. Before you start the installation, you need to plug in the USB flash drive and access your computer's BIOS settings. There, you

need to choose the USB flash drive as the first option in the boot sequence. Then, you need to save your settings and exit the BIOS. Your computer will then boot from the USB flash drive and start the installation process.

- **Network Boot (PXE Boot):** Another method of installing an operating system is network booting, which allows you to load the installation files from a remote server over the LAN. To use this method, you need to configure your computer to boot from the network. That involves accessing the BIOS settings and enabling the *network boot* option. You also need to set the network boot as the priority in the boot order. After saving your changes and exiting the BIOS, your computer will reboot and connect to the network server to start the installation process.

You can choose any of these methods to install an operating system, depending on your preferences and availability of resources. The installation process will vary depending on the source of the installation files.

Getting to know the startup process in BIOS

For any server technician, knowing how the hardware components and the startup process work is a valuable skill. It enables them to resolve issues related to hardware quickly and reduce downtime. To effectively diagnose and fix a server startup problem, they need to be familiar with the steps involved in starting up a server. Therefore, let us first examine the BIOS, which plays a crucial role in booting up a server.

When you power on a server, the initial activity involves a chip on the motherboard, known as ROM, which activates the BIOS program. The BIOS plays a crucial role in managing the server's hardware functionality. It detects and configures hardware components such as the CPU, memory, and disks. Additionally, the BIOS identifies bootable devices such as CD-ROMs, USB drives, and network interfaces, determining the sources from which the server can initiate the boot process. This process is illustrated in *Figure 2.3*.

Figure 2.3 – The BIOS configuration screen

Figure 2.3 – The BIOS configuration screen

However, BIOS has limitations and drawbacks that render it inadequate for modern servers. To address these issues, UEFI was developed as a replacement. UEFI provides numerous advantages over BIOS, including faster and more secure booting, support for larger disks and partitions, and an enhanced graphical interface. In the following section, we will explore UEFI in greater detail and examine its functionality in Windows Server 2025.

A different firmware program for booting modern computers

Modern computers no longer rely on the outdated BIOS system; instead, they utilize the UEFI, as depicted in *Figure 2.4*. Developed by the UEFI Consortium, UEFI addresses the limitations of BIOS in supporting contemporary hardware during startup. It can operate in both 32-bit and 64-bit processor modes and access the entire memory available in the system. UEFI employs the GPT partition scheme, which supports disks larger than 2 TB. Additionally, UEFI can be easily updated by downloading firmware updates from the manufacturer's website, offering a significant advantage over BIOS.

To access UEFI, begin by restarting your computer from the **Start** menu or by powering it off and on again. During the initial boot sequence, press the designated key or key combination to enter the UEFI settings; commonly used keys include *F2*, *F10*, *F12*, *Delete*, or *Esc*, though the specific key may vary depending on your computer's manufacturer and is often briefly shown on the screen. Once in the UEFI menu, use the arrow keys or mouse (if available) to navigate and configure various firmware settings, such as boot order, security options, and hardware configurations. Be sure to save any modifications before exiting to ensure that your changes are implemented.

Figure 2.4 – The UEFI setup utility

Figure 2.4 – The UEFI setup utility

Next, we will discuss the **Trusted Platform Module (TPM)**.

Understanding TPM

The TPM is a security chip embedded in the server motherboard designed to store encryption keys, certificates, passwords, and other sensitive data securely. TPM plays a crucial role in measuring the integrity of the boot process, ensuring that no unauthorized changes have been made to the server's firmware, bootloader, or operating system. It works in conjunction with **BitLocker**, a feature that encrypts the server's disks to prevent unauthorized data access. BitLocker leverages TPM to store the encryption key, unlocking it only if the server successfully passes the integrity check. This collaboration between TPM and BitLocker provides robust security for the server's data, protecting against tampering and theft. *Figure 2.5* shows the TPM Management console in Windows Server 2025, accessible by typing **tpm.msc** in the **Run** dialog box.

Figure 2.5 – The TPM Management console

Figure 2.5 – The TPM Management console

Next, we will discuss the **Power-On Self-Test (POST)**.

A crucial test for server hardware

For a server to start up correctly, its hardware must be in good condition. That is ensured by the POST (Power-On Self-Test), a diagnostic test that runs automatically when the server is powered on. POST checks the CPU, memory, disks, and other devices for errors or faults and communicates any issues through beep codes or error messages displayed on the screen or via the speaker. The POST process can be customized to run more or fewer tests to optimize boot time. As a vital tool for diagnosing and fixing hardware issues, POST is performed by the BIOS/UEFI, which manages server hardware operations.

Since different BIOS/UEFI manufacturers use varying beep codes, it is helpful to familiarize yourself with them. Beep codes are audio signals a computer's motherboard generates during startup to indicate hardware issues or errors. These codes contain short and long beep sequences that help diagnose problems when the system cannot provide visual error messages. Each beep pattern indicates a specific type of hardware malfunction, such as issues with memory, the graphics card, or the motherboard itself. To effectively troubleshoot and resolve these issues, users can consult the motherboard's manual or the manufacturer's documentation to decode the beep patterns and identify the underlying problems.

Particular attention should be paid to components like processors, memory, and graphics cards, as they are among the first to be tested by POST. If any of these components are defective, the server will not boot.

Note

For more information on the various beep codes used by different BIOS manufacturers, visit <https://www.computerhope.com/beep.htm>.

Next, we will discuss the GUID Partition Table (GPT) and the boot programs.

GPT and the boot programs

Once the server's hardware passes the POST, the BIOS/UEFI transfers control to the first boot device. The BIOS/UEFI then scans the boot device for the partition table, which indicates where the operating system is located. The partition table can be either MBR or GPT. GPT is a modern standard that supports larger disks and more features than MBR. It enhances reliability and recoverability by storing multiple copies of the partition table on the disk. Each partition in GPT is assigned a GUID, preventing conflicts and errors. GPT works with UEFI, a firmware program that replaces the traditional BIOS, providing a more secure and faster boot process. UEFI employs a boot loader capable of reading GPT partitions and loading the operating system from them. Depending

on the Windows OS installed on the server’s disk, the boot loader can be **NT Loader (NTLDR)**, **Boot Manager (BOOTMGR)**, or both. These programs are responsible for loading the OS into RAM. *Table 2.1* provides details on the NTLDR and BOOTMGR boot programs.

NTLDR (Windows NT to Windows Server 2003) BOOTMGR (Windows Vista to Windows Server 2025)

BOOT.INI	Boot configuration data
NTDETECT.COM	WinLoad.exe
NTOSKRNL.EXE	NTOSKRNL.EXE
HAL.DLL	Boot-class device drivers

Table 2.1 – The boot programs NTLDR and BOOTMGR

Having covered the fundamentals of GPT, it’s time to explore **Boot Configuration Data (BCD)**.

A database for booting Windows OS

BCD is a crucial database that stores settings and options for booting the Windows operating system. BCD can be managed and modified using the **bcdedit.exe** command-line tool or the graphical tool **BCDEdit**, which is part of the **Windows Recovery Environment**. It contains entries for each boot loader and operating system the server can load, such as NTLDR or BOOTMGR. It includes parameters for configuring the boot environment, such as display mode, memory limits, debugging options, and recovery settings. BCD is essential for managing the boot process and troubleshooting boot issues. It provides a standardized boot option interface for modern Windows OS versions, from Windows Vista to Windows Server 2025, regardless of the firmware, enhancing security compared to the previous **boot.ini** system. Administrators can set permissions to manage boot options, and BCD is accessible during all stages of system configuration. For example, **bcdedit.exe** (see *Figure 2.6*) is a file used to access the BCD data store located inside the disk partitions, unlike **boot.ini**.

Figure 2.6 – Running bcdedit.exe in Windows Server 2025

Figure 2.6 – Running bcdedit.exe in Windows Server 2025

In a multiboot scenario, both NTLDR and BOOTMGR may be present, with **boot.ini** and **bcdedit.exe** displaying the respective OS lists. In such cases, **bootsect.exe** (discussed later in the *Understanding boot sector* section) can be used to update the MBR for hard disk partitions requiring a switch between NTLDR and BOOTMGR.

When installing Windows Server 2025, disk partitioning and driver compatibility challenges are common. Disk partitioning issues often arise with mismatched partition styles (MBR for legacy BIOS or GPT for UEFI) or insufficient space on a partition, which can be resolved by verifying the boot mode, using compatible partition styles, and ensuring NTFS formatting. Driver compatibility challenges may occur when the installation doesn’t recognize certain hardware components, such as storage controllers. To address this, download the latest drivers from the manufacturer, load them during setup, and consider switching the boot mode if necessary to resolve conflicts. Addressing these issues before installation can streamline the setup and reduce errors.

With an established understanding of GPT and BCD, let us now explore the bootloader.

What is the bootloader?

A **bootloader**, also known as a **bootstrap loader** or **boot manager**, is a critical program responsible for initiating a computer's startup process. Once the POST confirms that the hardware is functioning correctly, the bootloader takes control. It resides in the MBR/GPT and is tasked with loading the Windows OS kernel into memory or onto disk. Windows operating systems utilize two main bootloaders, shown earlier in *Table 2.1*:

- **NTLDR**: The older bootloader used from Windows NT through Windows Server 2003
- **BOOTMGR**: The newer bootloader employed from Windows Vista to Windows Server 2025

Next, we will explore the boot sector, which contains the essential information needed to load the bootloader and start the server.

What is the boot sector?

The **boot sector** is a critical region on a disk that contains the essential information required to start the computer. Located in the first sector of the initial track on the disk, it typically includes either the MBR or the GPT. These are small programs responsible for initiating the bootloader. The bootloader, in turn, loads the Windows OS kernel into memory or onto disk. Depending on the Windows OS version, the bootloader could be NTLDR or BOOTMGR. The boot sector is fundamental to the boot process and must be compatible with the system's firmware (BIOS or UEFI) and partition scheme (MBR or GPT). Therefore, in systems using BIOS with MBR, the boot sector is made up of the MBR, which incorporates both the **Master Boot Code (MBC)** and the **Partition Table (PT)**, a common setup in older computers or those running Windows 7 and earlier. Conversely, BIOS systems using GPT rely on a **Protective MBR (PMBR)** to ensure compatibility by simulating an MBR layout, thus avoiding problems with software that does not recognize GPT. UEFI systems that utilize MBR do so with a **Compatibility Support Module (CSM)** to mimic traditional BIOS boot processes, enabling UEFI firmware to interact with the MBR as if it were a BIOS system, typically seen in configurations running Windows 7 or older Linux versions. On the other hand, UEFI systems with GPT include the **EFI System Partition (ESP)**, which stores the boot loaders and essential files needed to start the operating system, a standard setup for modern systems running Windows 10 or recent Linux distributions. These scenarios illustrate how various firmware and partition schemes interact with boot sectors to manage system startup and ensure compatibility.

Next, we will examine the boot menu, which facilitates the selection of different operating systems when multiple OSes are installed on the computer.

How to use the boot menu?

The **boot menu** is a valuable feature for selecting from multiple Windows operating systems installed on your computer, a process known as **multi-booting**. This capability is particularly beneficial for testing or troubleshooting different OS configurations. When you start your computer, the boot menu appears, allowing you to choose the OS to boot into. For older Windows versions, such as Windows NT through Windows Server 2003, this menu is managed by a text file called **boot.ini**, located in the root partition of the disk, typically **C:\boot.ini**. This file contains essential boot options, including the bootloader and available operating systems. In contrast, newer Windows versions, from Windows Vista to Windows Server 2025, utilize the BCD database to control boot settings.

The **boot.ini** file, which represents a critical configuration component, was used in earlier Windows operating systems such as Windows XP and Windows Server 2003. This file plays a key role in handling multiple operating systems on a single computer by listing all installed OSes and their corresponding boot options. It enables users to choose which operating system to boot into during system startup. Each entry within the **boot.ini** file details the path to the OS kernel files and includes parameters such as boot timeout settings and default options. Showcasing this configuration demonstrates how **boot.ini** facilitates the management and selection of different operating systems, simplifying the boot process and user experience.

The boot menu also provides access to Safe Mode, a diagnostic mode that loads only essential drivers and services. Safe Mode is valuable for identifying and resolving system issues.

How does Safe Mode operate?

When you encounter issues with booting your Windows operating system, such as improper loading or system crashes, **Safe Mode** can be a valuable diagnostic tool. Safe Mode is a troubleshooting feature that starts Windows with only the essential drivers and services necessary for basic functionality. This minimal setup can help you identify and resolve issues affecting your system. The method to access Safe Mode varies based on the Windows version you are using. For older versions such as Windows NT through Windows Server 2003, you can press the **F8** key during the boot process and select **Safe Mode** from **Windows Advanced Options Menu**, as depicted in *Figure 2.7*.

Figure 2.7 – Advanced Boot Options in Windows Server 2025

Figure 2.7 – Advanced Boot Options in Windows Server 2025

For newer versions, including Windows Vista through Windows Server 2025, you need to utilize the **Advanced startup** options. That involves holding down the **Shift** key while selecting **Restart** from the **Power** menu. Follow these steps to enter Safe Mode:

1. On the **Choose an option** screen, select **Troubleshoot**.
2. On the **Advanced options** screen, click **Startup Settings**.
3. Press the **Restart** button on the **Startup Settings** screen.
4. Once the system restarts, the **Advanced Boot Options** screen will appear, as shown in *Figure 2.8*. From there, select the **Safe Mode** option.

Windows setup and disk configuration errors

When preparing for the installation of Windows Server 2025, it's crucial to address common disk configuration errors to ensure a smooth setup process. One of the first considerations is the filesystem format requirements. Windows Server 2025 primarily utilizes the NTFS for system drives, which supports larger files and volumes compared to FAT32. If your drives are formatted in an incompatible filesystem, such as FAT32 or exFAT, you will need to reformat them to NTFS.

To format a drive or change its partition size, you can use Disk Management, a built-in Windows utility that provides a graphical interface for managing disk partitions. Here's how you can use Disk Management to resolve formatting issues:

- **Access Disk Management:** Right-click on the **Start** menu and select **Disk Management**. This will open the **Disk Management** console, where you can view all connected drives and their partitions.
- **Format a drive:** If a drive is not formatted or needs to be changed to NTFS, right-click on the partition and select **Format**. Choose **NTFS** as the filesystem and follow the prompts to complete the formatting process. Ensure that you back up any important data before formatting, as this will erase all data on the drive.
- **Resize partitions:** If you encounter errors due to insufficient space on a partition, you can resize partitions using Disk Management. Right-click on the partition you wish to resize and select **Shrink Volume** or **Extend Volume**. Shrinking a volume allows you to free up space for other partitions while extending a volume can help you add space to a partition that is running low.
- **Confirm compatibility:** After making changes, confirm that the drives are now in the NTFS format and that the sizes meet the installation requirements for Windows Server 2025.

By proactively managing your disk configuration and ensuring compatibility with NTFS, you can minimize the likelihood of encountering errors during the installation process, leading to a smoother and more efficient setup of Windows Server 2025.

In this section, we have explored various elements of the Windows boot process, including BIOS, UEFI, TPM, POST, MBR, BCD, bootloaders, boot sectors, boot menus, Safe Mode, and disk configuration errors. The following section will focus on business continuity and the strategy for maintaining it.

Installation options for Windows Server 2025

When deploying Windows Server, selecting the appropriate installation option is crucial to meet your specific needs. Windows Server provides various installation options, each catering to different requirements in terms of disk space, memory usage, features, and graphical interfaces. These options also influence security, performance, management, and compatibility. This section will explore and compare the three primary installation options for Windows Server: **Desktop Experience**, **Server Core**, and **Nano Server**. We will discuss the benefits and considerations of each, along with guidance on how to switch between them based on your operational needs. By the end of this section, you will be equipped to choose the most suitable installation option for your Windows Server deployment.

Understanding the role of your server

When installing Windows Server 2025, it's crucial to consider the specific roles and responsibilities your server will fulfill. This foresight can significantly influence the configuration choices you make during installation, including hardware specifications and service selection.

- **Assessing workloads:** Determine whether your server will primarily handle read operations, write operations, or a balanced mix of both. For example, a file server that primarily serves files to multiple clients may benefit from faster disk speeds and larger RAM to accommodate high read demands. In contrast, a database server that handles extensive write operations may require optimized storage solutions and potentially more robust processing power.
- **Memory requirements:** Different server roles may have varied memory requirements. For instance, virtualization servers typically need more RAM to manage multiple virtual machines effectively. Ensuring your server has adequate memory will help maintain performance and responsiveness under load.
- **Selecting components and services:** Choosing the right components is critical. For a web server, you might prioritize **network interface cards (NICs)** for high-speed connectivity. In contrast, for a database server, you might focus on fast, high-capacity disks to handle data transactions efficiently.
- **Planning for future growth:** When planning your installation, consider future scalability. Will the server need to accommodate more users or handle increased data? Factor this into your hardware choices to avoid potential bottlenecks as demands grow.

By integrating these considerations into your installation planning, you can ensure that your Windows Server 2025 environment is tailored to effectively meet your organizational needs.

Pre-installation checks – resource compatibility checks

Before diving into the installation of Windows Server 2025, it is essential to perform thorough resource compatibility checks to ensure that your hardware meets the necessary requirements for optimal performance. This step can save time and prevent frustration during and after the installation process. Here are key points to consider:

- **System requirements:** Verify that your hardware meets or exceeds the minimum system requirements for Windows Server 2025. Key specifications include the following:
 - **CPU:** Ensure that the processor is compatible with Windows Server 2025, typically requiring a minimum of a 1.4 GHz 64-bit processor. Consider opting for a multi-core processor for better performance, especially in environments with multiple users or applications.
 - **RAM:** Confirm that the server has adequate RAM. The minimum requirement is generally 2 GB, but for better performance and to handle heavier workloads, it's advisable to have 4 GB or more, depending on the intended use and applications.
- **Resource availability:** In addition to meeting the minimum specifications, check the availability of resources on the server. This includes the following:
 - **Disk space:** Ensure that there is enough disk space for the installation, as well as for future updates and applications. A minimum of 32 GB of free space is often recommended, though more may be needed based on your configuration and usage scenarios.

- **Network resources:** Assess network bandwidth and connectivity, particularly if you are deploying in a cloud or hybrid environment. Sufficient bandwidth is necessary for downloading updates, accessing network resources, and ensuring a smooth installation process.
- **Performance considerations:** Installing Windows Server on underpowered hardware can lead to performance issues that affect not only the server's operation but also any applications or services running on it. By ensuring compatibility and availability of resources upfront, you can avoid potential performance bottlenecks post-installation.
- **Compatibility with existing applications:** If you plan to run specific applications on Windows Server 2025, double-check their compatibility with the new operating system version. Some applications may have specific resource requirements that need to be accounted for.

By conducting these resource compatibility checks before installation, you set the foundation for a successful deployment of Windows Server 2025, minimizing the likelihood of performance-related problems down the line. This proactive approach not only enhances system reliability but also optimizes the overall user experience.

Which installation option for Windows Server 2025 should I choose?

When installing Windows Server 2025, you have three distinct options to choose from, each offering unique advantages and limitations based on your server needs, hardware specifications, and management preferences. The following is an overview of each option:

- **Desktop Experience:** This option provides a complete **graphical user interface (GUI)** along with all associated tools and functionalities of Windows Server 2025. While it offers a comprehensive user experience, it requires more hardware resources and might present a higher security risk compared to the other options.
- **Server Core:** Recommended by Microsoft for its efficiency, Server Core is a minimal installation option that omits the GUI, focusing instead on core server functionalities. It consumes fewer resources and has a reduced attack surface. Management can be performed locally via **Windows PowerShell** or remotely using **Server Manager**.
- **Nano Server:** An advanced version of Server Core, Nano Server is designed to be even more lightweight and efficient. It supports only 64-bit applications and lacks local login capabilities, requiring management through remote tools such as **Windows Admin Center** or **Windows PowerShell**. Nano Server is particularly suited for cloud environments or containerized applications that demand minimal maintenance and updates.

Before proceeding with the installation of Windows Server 2025, evaluate these options to determine which best aligns with your operational requirements. In the next section, I will guide you through comparing Nano Server with Server Core.

Comparing Nano Server and Server Core

This section compares and discusses the Nano server and the server core.

Here is an overview of the Nano Server:

- **Lightweight deployment:** Nano Server is a headless (no GUI) installation option for Windows Server that is optimized for cloud environments and containers. It has a minimal footprint, making it ideal for running specific workloads efficiently.
- **Use cases:** Nano Server is particularly suited for microservices, cloud applications, and containerized environments. It's often used for hosting web services, application servers, and specialized workloads such as Hyper-V.

Here is an overview of the Server Core:

- **Reduced installation:** Server Core is a minimal installation option for Windows Server that offers a reduced GUI (no traditional desktop interface) but retains the command-line tools and Windows Management

Framework. It provides a balance between functionality and resource consumption.

- **Use cases:** Server Core is suitable for traditional server roles such as **Active Directory Domain Services (AD DS)**, DNS, and file services. It's commonly deployed in environments where administrators require more features than Nano Server but still want to limit the server's attack surface.

Comparison and scenarios

These are the considerations for choosing Nano Server:

- **Containerization:** If an organization is moving toward microservices and containerization, Nano Server is an excellent choice due to its lightweight nature and compatibility with Docker
- **Cloud-first deployments:** For enterprises focused on cloud-native applications, using Nano Server in Azure or hybrid environments can optimize resource usage and speed up deployment times
- **Web applications:** Nano Server is ideal for hosting **Internet Information Services (IIS)** and web applications that require a small footprint and high performance

These are the considerations for choosing Server Core:

- **Traditional roles:** If the server needs to support traditional roles such as AD DS or file services, Server Core is the better option. It provides the necessary functionalities without the overhead of a full GUI.
- **Compatibility requirements:** For applications that require compatibility with more traditional Windows server features or management tools, Server Core offers a more robust environment while still limiting the surface area for potential attacks.
- **Management needs:** Server Core allows for remote management via PowerShell, which is essential in larger enterprises where centralized management is critical.

By providing specific use cases for both Nano Server and Server Core, you can better understand how to select the appropriate installation option based on your organizational needs and workload requirements. This approach not only clarifies the differences but also illustrates how these technologies fit into modern IT strategies. Next, we will examine using logs to diagnose installation failures.

Using logs to diagnose installation failures

When facing installation issues with Windows Server 2025, log analysis can be an invaluable tool for diagnosing and resolving problems. Windows Setup generates several logs that contain detailed information about the installation process, helping administrators identify where the failure occurred. Two key logs to focus on are **setupact.log** and **setuperr.log**.

Key installation logs

Following are the key installation logs:

- **setupact.log:** This log records the entire installation process and contains timestamps, component states, and detailed information about actions taken during the setup. It is generally located in the **C:\Windows\Panther** directory. This log is useful for identifying the overall flow of the installation and determining which stage may have encountered issues.
- **setuperr.log:** In contrast, this log captures errors encountered during installation. It provides information about any failures that prevented components from installing successfully. This log is also found in the **C:\Windows\Panther** directory and is critical for troubleshooting specific error messages that may appear during the installation.

Locating installation logs

Here is how you locate installation logs:

- After a failed installation, access the logs by booting into the recovery environment or using a bootable USB drive to access the **C:\Windows\Panther** folder
- Copy the log files to a USB drive or a secondary drive for analysis, as these logs are essential for diagnosing issues

Interpreting the logs

The following details how you interpret the logs:

- Open **setupact.log** in a text editor and scroll through the entries to find any abnormal entries or timestamps that correspond to the installation failure. Look for keywords such as **error**, **failed**, or **warning** to identify problematic areas.
- In **setuperr.log**, review the entries for error codes or specific messages indicating what went wrong. Microsoft's documentation can provide insights into specific error codes, aiding in finding resolutions.

Action steps

The following details the steps you take as action after interpreting the logs:

- Once you identify the errors, cross-reference the error messages with Microsoft's knowledge base or community forums for potential solutions. Common issues may involve driver compatibility, disk partitioning problems, or network configuration errors.

By understanding and utilizing these logs effectively, administrators can streamline the troubleshooting process and enhance their ability to resolve installation issues with Windows Server 2025, ensuring a smoother setup experience. Next, we will look at network connectivity and domain joining, emphasizing how vital network connectivity is for domain-joining tasks.

Network connectivity and domain joining

When deploying Windows Server 2025, mainly through network installations, ensuring reliable network connectivity is crucial for successful domain joining and overall deployment. Network-related issues can prevent the server from accessing essential resources or joining the domain properly. Here are some troubleshooting steps to address these common challenges:

1. **Verify IP configuration:** Ensure that the server has a valid IP address by checking the network settings. For static IP configurations, confirm that the IP address is correctly assigned and falls within the appropriate range for your network. If using DHCP, ensure that the server successfully receives an IP address from the DHCP server.
2. **Check DNS settings:** The **Domain Name System (DNS)** is vital for locating domain controllers and other network resources. Confirm that the server is pointing to the correct DNS servers, typically the IP addresses of your domain controllers. You can test DNS resolution by using the **nslookup** command to ensure the server can resolve domain names.
3. **Examine firewall configurations:** Firewalls can block necessary communication between the server and the domain controller. Review the firewall settings to ensure that the appropriate ports are open. For domain joining, ensure that ports such as **53** (DNS), **88** (Kerberos), and **389** (LDAP) are allowed through the firewall. You may need to temporarily disable the firewall for testing purposes to see whether it's causing the issue.
4. **Ping the domain controller:** Use the **ping** command to check connectivity to the domain controller. This simple test can help identify whether the server can reach the DC. If **ping** fails, investigate network routing, cabling, or switch issues that may be affecting connectivity.
5. **Review network adapter settings:** Confirm that the network adapter settings are correctly configured. Check for issues such as network adapter being disabled, improper VLAN settings, or connectivity issues with physical network hardware.
6. **Logs and error messages:** Pay attention to any error messages during the installation or domain joining process. Windows Server logs can provide insights into what went wrong. Check the Event Viewer for any

relevant logs that might indicate network-related problems.

By following these troubleshooting steps, administrators can effectively diagnose and resolve network connectivity issues that may arise during the deployment of Windows Server 2025, ensuring a smooth domain joining process and reliable network performance. Next, activation and licenses are considered important elements of smooth operations on a Windows Server 2025.

Activation and licensing issues

When installing Windows Server 2025, it is crucial to address activation and licensing, particularly when deploying in Azure or hybrid environments. Proper activation ensures that the operating system is genuine and can receive essential updates and features. However, users may encounter common activation problems during or after the installation process. Understanding how to troubleshoot these issues can save time and ensure compliance with licensing agreements.

One common issue is activation failure due to connectivity problems. In cloud or hybrid setups, ensure that your server has a stable internet connection to reach Microsoft's activation servers. For Azure deployments, confirm that the correct Azure resource group and virtual network settings are in place. If activation fails despite proper connectivity, it may be necessary to use the **slmgr** command-line tool to troubleshoot. Running commands such as **slmgr /ato** attempts to activate the product key manually, providing feedback on any specific issues encountered.

Licensing issues can also arise when using volume licensing keys in multiple environments. Each instance of Windows Server must be properly licensed, and organizations using volume activation should verify that they are following the required licensing model. For hybrid deployments, consider using Azure Hybrid Benefit, which allows existing Windows Server licenses to be applied to Azure virtual machines, reducing costs. Ensure that the necessary licenses are linked to your Azure subscription and that you understand the implications of your licensing choice on compliance and support.

Addressing activation and licensing issues early in the installation process is vital for the smooth deployment of Windows Server 2025. By ensuring proper connectivity, utilizing command-line tools for troubleshooting, and understanding licensing models, organizations can avoid potential disruptions and maintain compliance in their cloud and hybrid environments. Next, let us examine the various deployment methods of Windows Server 2025.

Various methods for deploying Windows Server 2025

Several installation methods are available for Windows Server 2025, each suited to a different scenario. Here are some common approaches.

Clean install

This method involves setting up a fresh instance of Windows Server 2025 on a server and erasing any previous data or configurations. It is ideal for new deployments or when starting with a clean slate.

Chapter exercise 2.1 – performing a clean installation of Windows Server 2025

When installing Windows Server 2025 on a new or existing hard drive, a clean installation is an effective option. This process will remove the current operating system from the disk and replace it with a fresh installation of Windows Server 2025. Although you will need to interact with the setup, it is less intensive than performing an upgrade. Follow these steps for a clean installation:

Insert the bootable media, plug in a USB flash drive, or connect a network cable to the server.

1. Power on your computer and select your preferred boot option, such as DVD, USB flash drive, or network boot. A confirmation message will appear on the screen.

2. The installation files will load into memory (RAM).
3. Choose your language and other preferences for the installation, then click **Next**, as shown in *Figure 2.8*.

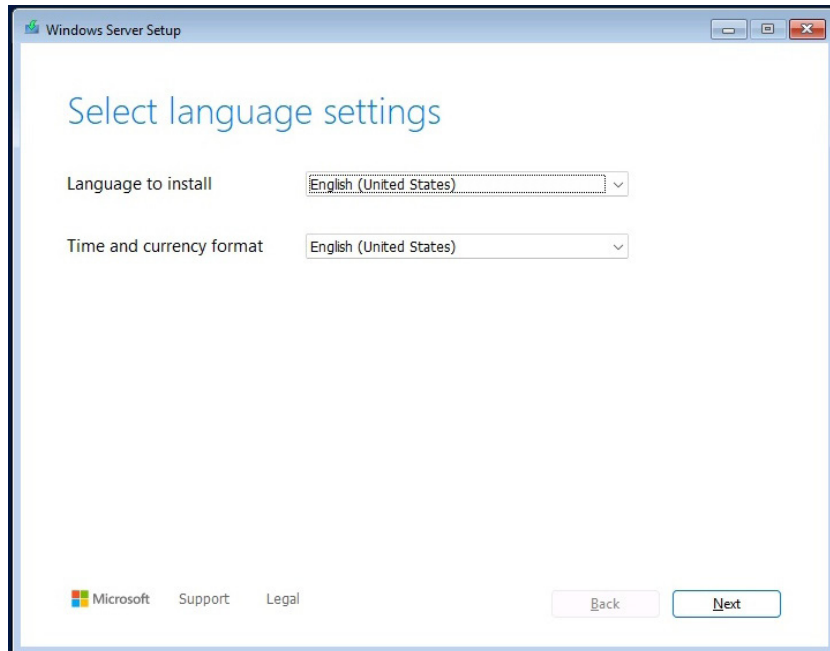


Figure 2.8 – Selecting language settings

4. Choose the keyboard or other input method and then click Next.
5. Select the **Install Windows Server** option and ensure that I agree that everything will be deleted, including files, apps, and settings box, and that it is checked, as shown in Figure 2.9. Click **Next**.

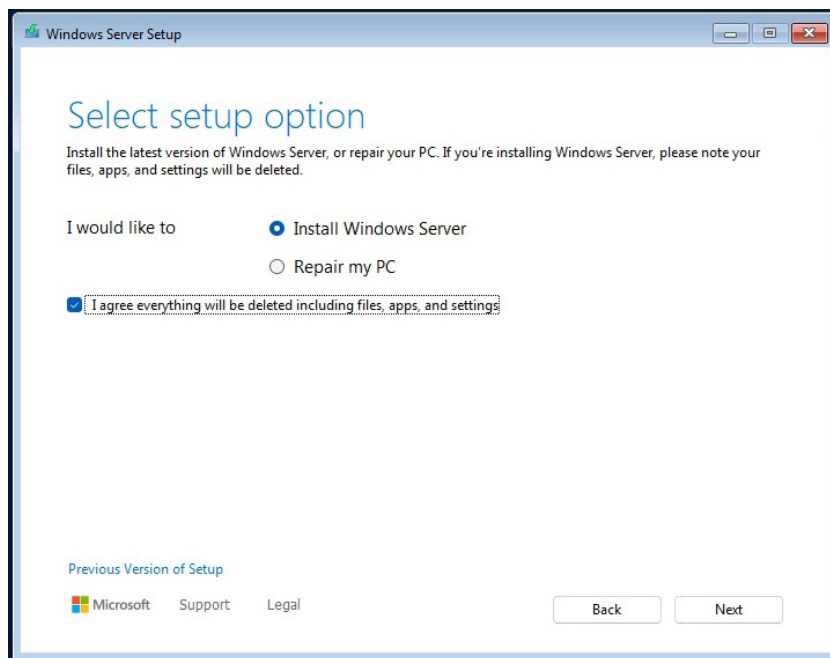


Figure 2.9 – Selecting setup option

6. Select **Windows Server 2025 Datacenter (Desktop Experience)** and click **Next**, as illustrated in *Figure 2.10*.

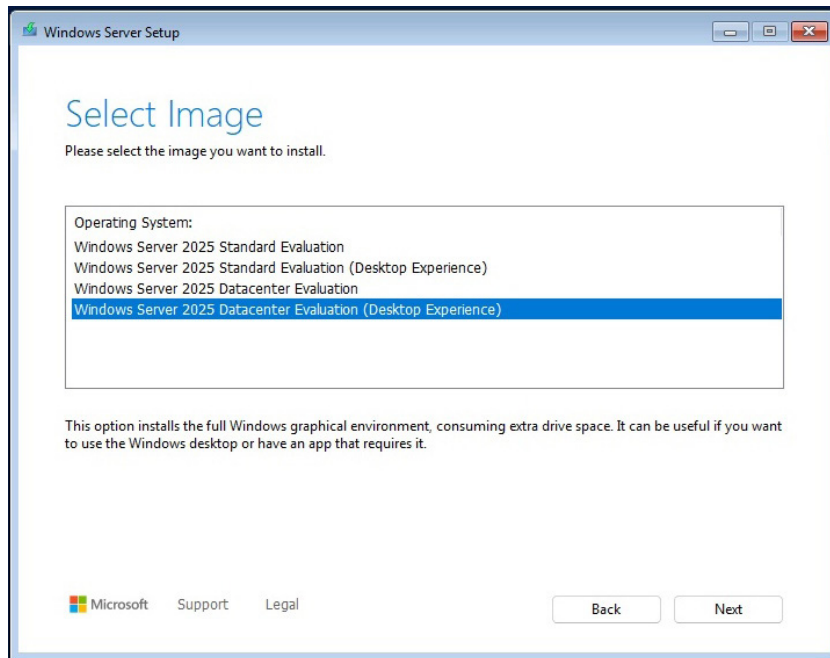


Figure 2.10 – Selecting the image to install

7. Take the time to read the Applicable notices and license terms. Click **Accept**.
8. Select a location to install Windows Server 2025, as shown in *Figure 2.11*. Click **Next**.

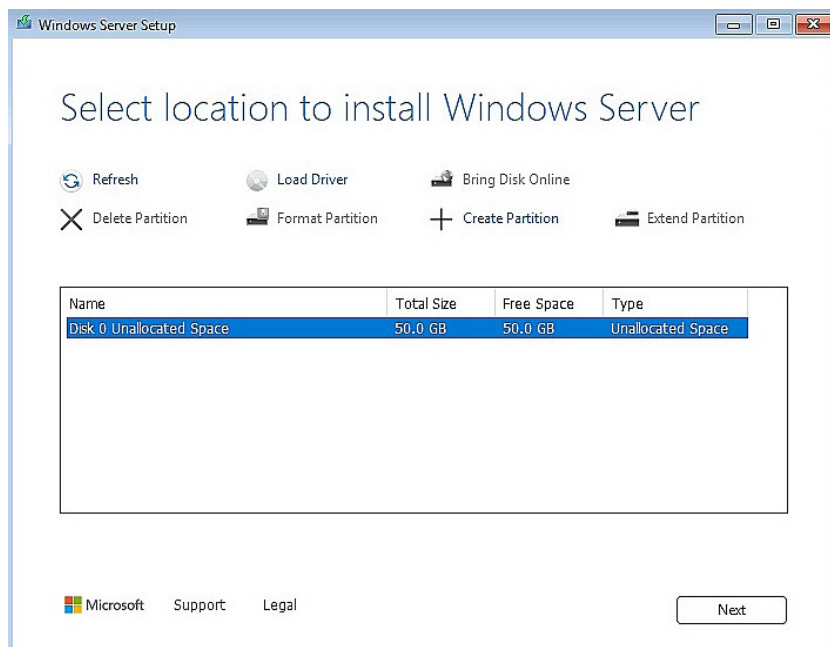


Figure 2.11 – Selecting disk or partition to install Windows Server 2025

9. On the **Ready to install** page of Windows Server Setup, click the **Install** button, as illustrated in *Figure 2.12*.

 Figure 2.12 – Starting the installation of Windows Server 2025

Figure 2.12 – Starting the installation of Windows Server 2025

10. Windows Server Setup will now begin installing Windows Server 2025. You can sit back and relax during this process.

 Figure 2.13 – Starting the installation of Windows Server 2025

Figure 2.13 – Starting the installation of Windows Server 2025

11. After the installation is complete and the system restarts several times, set up the administrator password and click **Finish**, as shown in *Figure 2.14*.


 Figure 2.14 – Setting up an administrator's password

Figure 2.14 – Setting up an administrator's password

12. Press **Ctrl + Alt + Delete** to unlock the system, as depicted in *Figure 2.15*. Enter the administrator password and press **Enter** to log in for your first login.

 Figure 2.15 - Press Ctrl+Alt+Del to unlock the system

Figure 2.15 - Press Ctrl+Alt+Del to unlock the system

13. In Send diagnostic data to Microsoft, select the preferred choice and click Accept.
14. Congratulations! You have successfully installed Windows Server 2025, as in *Figure 2.16*.


 Figure 2.16 Start menu and desktop in Windows Server 2025

Figure 2.16 Start menu and desktop in Windows Server 2025

Note

One of Microsoft's standout features is the **Windows Installer**, an **application programming interface (API)** for software installation, maintenance, and uninstallation.

With the clean installation process mastered, you are now ready to explore other installation methods. Next, we will delve into deploying Windows Server 2025 using the MDT.

Deploying with the MDT

By using the MDT, you can streamline and automate the installation process. This method is particularly useful for deploying multiple servers efficiently and consistently.

Chapter exercise 2.2 – performing an installation of Windows Server 2025 using the MDT

An unattended installation automates the deployment process, requiring minimal interaction. This method is ideal for deploying multiple servers in enterprise environments. Although Microsoft has deprecated WDS in Windows 11 and Windows Server 2025, the Windows ADK and MDT offer robust alternatives for automating installations. These tools, available for download, simplify the deployment process.

The key to an unattended installation is the **answer file**, an XML file that provides the necessary responses to installation prompts. You can create an answer file manually using Notepad or download sample files from the internet. Microsoft also offers several tools to aid in this process, such as the **Windows System Image Manager (Windows SIM)** included in the Windows Assessment and Deployment Kit (ADK) and the Microsoft Deployment Toolkit (MDT).

Note

Microsoft's official documentation offers a range of examples and in-depth explanations via the **Windows System Image Manager (WSIM)** answer file. For example, although the **TechNet Gallery** has been archived, it still provides valuable community-contributed samples for various scenarios. **GitHub** repositories often host user-shared answer files that cater to different deployment requirements. **Deployment Research** offers a collection of practical templates and samples with detailed explanations for Windows deployment. Additionally, **Microsoft Learn** provides extensive guides and examples on creating and utilizing answer files, complete with downloadable samples. These resources are essential for finding and comprehending answer file samples and facilitating efficient Windows deployment and configuration.

To install Windows Server 2025 using the MDT, follow these steps:

1. Download and install the Windows ADK on a Windows 11 computer (see *Figure 2.17*).

Figure 2.17 – Installing Windows ADK

Figure 2.17 – Installing Windows ADK

2. Download and install **Windows Preinstallation Environment (Windows PE)** by running **adkwinsesetup.exe** (refer to *Figure 2.18*).

Figure 2.18 – Installing Windows PE

Figure 2.18 – Installing Windows PE

3. Install MDT on a Windows 11 computer, as illustrated in *Figure 2.19*.

Figure 2.19 – Installing the MDT

Figure 2.19 – Installing the MDT

4. After installing the Windows ADK, PE, and MDT, run **Deployment Workbench** and select **New Deployment Share Wizard** (*Figure 2.20*).

Figure 2.20 – New Deployment Share Wizard

Figure 2.20 – New Deployment Share Wizard

5. Share the **DeploymentShare** folder through File Explorer, allowing **Everyone to read** permissions.
6. Run **Import Operating System Wizard** to import the Windows Server 2025 files.
7. Use **New Task Sequence Wizard** to create the answer file for the unattended installation.
8. On **MDT Deployment Share Properties**, within the **Platforms Supported** section, uncheck **x86**, as depicted in *Figure 2.21*, and click **OK**.


Figure 2.21 – Disabling the x86 platform on the MDT

Figure 2.21 – Disabling the x86 platform on the MDT

9. Update the deployment share to create a bootable PE image.
10. Boot the new server with the **LiteTouchPE_x64 image**, located in the **Boot** subfolder of the **DeploymentShare** folder. After a successful boot, select **Run the Deployment Wizard to install a new Operating System** (*Figure 2.22*).

Figure 2.22 – Deploying Windows Server 2025 over the MDT

Figure 2.22 – Deploying Windows Server 2025 over the MDT

11. Provide credentials to access the **DeploymentShare** folder, ensuring the user has complete control.

12. Select the task sequence and the answer file created earlier with **Deployment Workbench**, then click **Next**.
13. Enter the computer details, locale, and time settings, and specify whether to capture the image or configure BitLocker.
14. Begin the deployment of Windows Server 2025 (see *Figure 2.23*).


 Figure 2.23 – Installing the operating system over the MDT

Figure 2.23 – Installing the operating system over the MDT

15. The MDT will then handle the deployment of Windows Server 2025. Once the installation process is completed, the system will configure the devices, and **Windows Server 2025 Standard** (Desktop Experience) will be successfully deployed.

Note

You can obtain the Windows ADK by visiting this link: <https://learn.microsoft.com/en-us/windows-hardware/get-started/adk-install#download-the-adk-101261001-may-2024>. For the MDT, download it from this page: <https://www.microsoft.com/en-us/download/details.aspx?id=54259>.

Now, let us explore how to perform an in-place upgrade, which allows you to update the existing OS to a newer version while retaining user settings, applications, and data.

In-place upgrade

This approach upgrades an existing Windows Server installation to Windows Server 2025, preserving existing settings and data. It is suitable for updating from previous versions while maintaining the current configuration.

Chapter exercise 2.3 – performing an in-place upgrade of Windows Server 2022 to Windows Server 2025

To change your current operating system to a newer one, you can do an upgrade. This way, you can keep your files and settings as they are. That is also known as an **in-place upgrade** because it takes place on a machine that already has an operating system installed. Before doing an upgrade, it is advisable to back up the Windows state, files, and folders. Microsoft says that Windows Server 2025 can be upgraded directly from Windows Server 2012 R2, 2016, 2019, or 2022.

To do an in-place upgrade from Windows Server 2022 to Windows Server 2025, follow these steps:

1. Insert the Windows Server 2025 installation disk or connect the bootable USB flash drive and run the setup file.
2. The **Install Windows Server** window will show up. Click **Next** to continue, as shown in *Figure 2.24*.

 Figure 2.24 – Beginning the in-place upgrade in Windows Server 2022

Figure 2.24 – Beginning the in-place upgrade in Windows Server 2022

3. As shown in *Figure 2.25*, select the Windows Server 2025 edition you want to install and click **Next**.

 Figure 2.25 – Pick the Windows Server 2025 edition to install

Figure 2.25 – Pick the Windows Server 2025 edition to install

4. Click the **Accept** button in **Applicable notices and license terms** to agree to the license terms.
5. Select what to keep, and then click **Next** to continue.
6. Once the updates are downloaded and the Windows Server 2022 setup ensures there is enough disk space on the server, you are ready to install. Click the **Install** button to continue with the upgrade, as shown in *Figure 2.26*

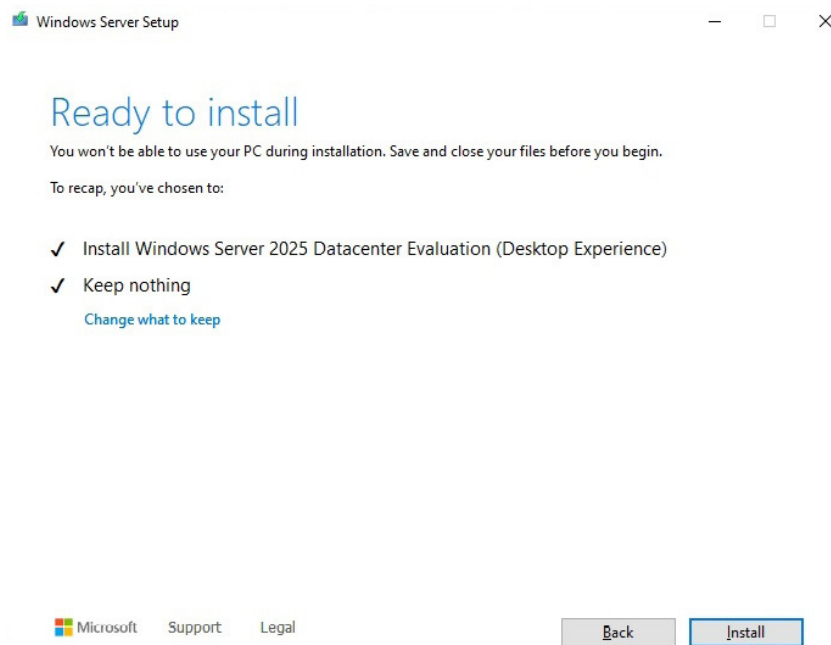


Figure 2.26 – Ready to run the in-place upgrade

7. The in-place upgrade will begin. Either relax or do other work until the upgrade is completed.
8. After several restarts, the upgrade from Windows Server 2022 to Windows Server 2025 will be completed successfully, as shown in *Figure 2.27*.



Figure 2.27 – The system properties confirm the in-place upgrade

Next, let us learn how to perform the migration using the **Windows Server Migration Tool (WSMT)**. This will help migrate the services from an old server to a new one.

Migration

Migration involves moving from an older server or environment to Windows Server 2025, often using tools or services to transfer applications, data, and settings to the new server.

Chapter exercise 2.4 – migrating network services from Windows Server 2012 R2 to Windows Server 2025

When upgrading to a new server, whether physical or virtual, you must transfer roles, features, applications, settings, and network services from the old server to the new one. This process is known as migration. Start by installing the operating system on the new server and then initiate the migration. Ensure that Windows Server 2025 is compatible with your existing applications. The WSMT feature can assist with migration, and **PowerShell cmdlets** can be used to migrate specific services.

For example, to migrate the **DHCP server** from an old server (Windows Server 2012 R2) to a new server (Windows Server 2025), follow these steps:

1. On the old server (Windows Server 2012 R2), open **Windows PowerShell** as an administrator and run the following cmdlet :

```
Export-DhcpServer -File C:\DHCPdata.xml -Leases -Force -ComputerName <OldServerName> -Verbose
```

2. Stop the DHCP service and move the **DHCPdata.xml** file to a shared folder that is accessible to everyone.
3. On the new server (Windows Server 2025), use **Server Manager** to add the **DHCP Server** role, as shown in *Figure 2.28*.

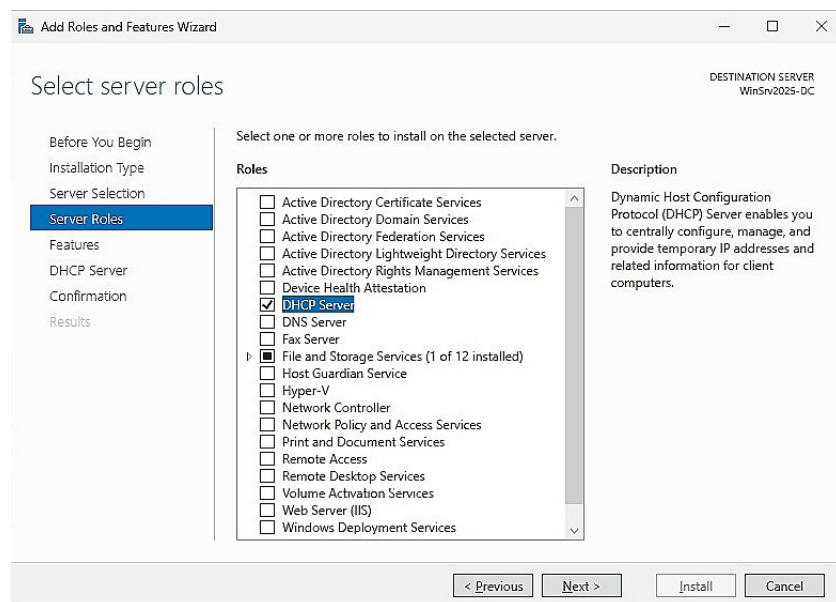


Figure 2.28 – Adding DHCP Server to a new server

4. Access the shared folder on the old server and copy the **DHCPdata.xml** file to the root directory on the new server.
5. Open Windows PowerShell as an administrator on the new server and run the following cmdlet (see *Figure 2.29*):

```
Import-DhcpServer -File C:\DHCPdata.xml -BackupPath C:\DHCP\ -Leases -ScopeOverwrite -Force .
```



```

Windows PowerShell
Copyright (C) Microsoft Corporation. All rights reserved.

Install the latest PowerShell for new features and improvements! https://aka.ms/PSWindows

PS C:\Users\Administrators> Import-DhcpServer -File C:\DHCPdata.xml -BackupPath C:\DHCP\ -Leases -ScopeOverwrite -Force -
ComputerName WinSrv2025-DC -Verbose
VERBOSE: The configuration (and leases) from the file C:\DHCPdata.xml will be imported to server WinSrv2025-DC.
VERBOSE: Dhcp Server database has been backed up at C:\DHCP\ on WinSrv2025-DC.
VERBOSE: Importing configuration on server WinSrv2025-DC from file C:\DHCPdata.xml.
VERBOSE: Importing classes on server...
VERBOSE: Class 'Default Routing and Remote Access Class' of type User already exists on server WinSrv2025-DC and will
not be changed.
VERBOSE: Class 'Default BOOTP Class' of type User already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Class 'Microsoft Windows 2000 Options' of type Vendor already exists on server WinSrv2025-DC and will not be
changed.
VERBOSE: Class 'Microsoft Windows 98 Options' of type Vendor already exists on server WinSrv2025-DC and will not be
changed.
VERBOSE: Class 'Microsoft Options' of type Vendor already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Importing option definitions on server...
VERBOSE: Option definition Classless Static Routes already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Subnet Mask already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Time Offset already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Router already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Time Server already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Name Servers already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition DNS Servers already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Log Servers already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition Cookie Servers already exists on server WinSrv2025-DC and will not be changed.
VERBOSE: Option definition LPR Servers already exists on server WinSrv2025-DC and will not be changed.

```

Figure 2.29 – Importing the DHCP Server to the new server

- Restart the DHCP service. Once the restart is complete, the DHCP server will be successfully migrated to the new server.

Next, we will explore how to install Windows Server 2025 in the cloud, such as on Microsoft Azure, to facilitate running services in a cloud environment.

Deploying in Azure

Nowadays, where businesses rely heavily on cloud services, an operating system optimized for cloud environments is essential. Multiple services from the cloud undoubtedly need an operating system optimized for the cloud environment. For about 25 years, Microsoft has offered Windows Server for midrange servers for businesses that want to entrust their network services to a cloud-based solution, ensuring scalability, reliability, and optimized performance for cloud operations.

Chapter exercise 2.5 – installing Windows Server 2025 on Azure

To explore Windows Server 2025 without impacting your existing machines, consider using **Microsoft Azure**, a cloud platform. First, you need an Azure account and subscription, which you can obtain for free on the Azure website. Once signed in to the **Azure portal**, follow these steps to create a **virtual machine (VM)** with Windows Server 2025:

- Navigate to **Virtual machines** and click **Create a virtual machine** to open the configuration wizard.
- Select your subscription and choose either an existing resource group or create a new one.

Note

In Azure resource management, you have the option to either utilize an existing resource group or establish a new one. An existing resource group serves as a predefined container that organizes and manages related resources, streamlining administration and coordination. If there is no suitable resource group already in place, you can create a new one tailored to your specific requirements, such as by project, department, or environment. This flexibility ensures that resources are systematically grouped and managed, maintaining an orderly structure within your Azure environment.

- Name your VM and select the region where you want it to run.
- For the OS image, select **Windows Server 2025 Datacenter – Gen2** (see *Figure 2.30*).

Microsoft Azure Search resources, services, and docs (G+)

Dashboard > Virtual machines >

Create a virtual machine

Basics Disks Networking Management Advanced Tags Review + create

Create a virtual machine that runs Linux or Windows. Select an image from Azure marketplace or use your own customized image. Complete the Basics tab then Review + create to provision a virtual machine with default parameters or review each tab for full customization. [Learn more](#)

Project details
Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * MSDN Platforms
Resource group * (New) WinSrv2022
[Create new](#)

Instance details

Virtual machine name * WinSrv2022-Azure
Region * (Europe) North Europe
Availability options No infrastructure redundancy required
Security type Standard

Image * Windows Server 2022 Datacenter: Azure Edition - Gen2
[See all images](#) | [Configure VM generation](#)

Azure Spot instance ☐

Size * Standard_D2s_v3 - 2 vcpus, 8 GiB memory (\$78.11/month)
[See all sizes](#)

[Review + create](#) < Previous Next : Disks >

Figure 2.30 – Setting up the VM with Windows Server 2022 in Azure

- Complete the fields in the **Basics** tab, then proceed to configure the settings in the **Disks**, **Networking**, **Management**, **Advanced**, and **Tags** tabs.
- After entering all the details, click **Review + create** to verify them (see *Figure 2.31*).

Dashboard > Virtual machines >

Create a virtual machine

Basics Disks Networking Management Advanced Tags Review + create

Validation passed

Figure 2.31 – Validating the VM's entries

- If everything is correct, click **Create** to start the deployment of the VM with Windows Server 2025 Azure edition (see *Figure 2.32*).

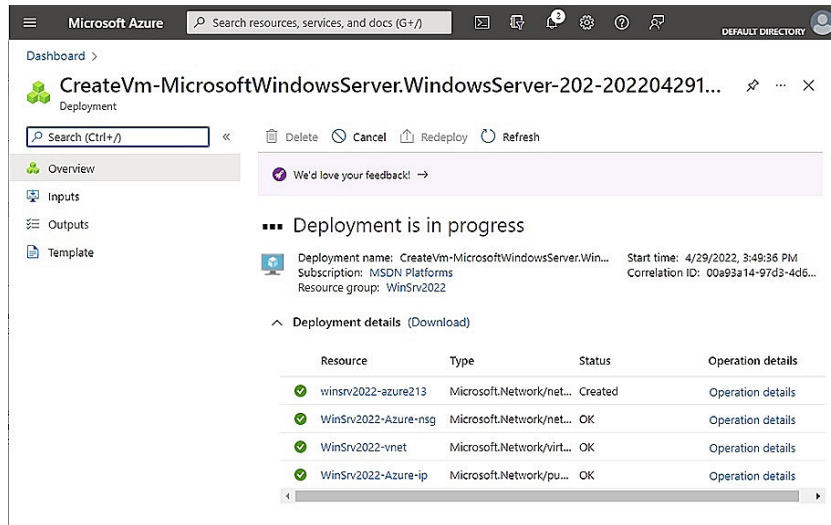


Figure 2.32 – Deploying a VM in Azure

- Once the deployment is complete, go to the resource and access the VM using the **Connect | RDP** option from the Azure portal or **Remote Desktop Connections** by entering the public IP address.

In this section, we’ve covered various methods for installing and deploying Windows Server 2025, including clean installation, network installation, unattended installation, in-place upgrades, migrations, and using Azure. These methods offer flexibility for both on-premises and cloud-based deployments.

Additionally, while **System Center Configuration Manager (SCCM)** is another option for deploying Windows Server 2025, it is more complex and requires an enterprise-level IT infrastructure. As this method involves advanced configuration, it is not covered in this book.

Summary

In this chapter, you learned about the various methods for installing and deploying Windows Server 2025. You explored how to choose the appropriate installation method based on specific scenarios and requirements, such as clean installations, network installations, in-place upgrades, migrations, and Azure-based installations. You also delved into different partition schemes, boot options, and advanced startup settings that influence the installation process. Through detailed steps and screenshots, you practiced executing each installation method. By the end of the chapter, you have gained the skills needed to install Windows Server 2025 on any physical or virtual machine, whether on-premises or in the cloud. To gain insights into its development, next, we will examine the timeline of Windows Server and explore how it has progressed over the years.

Questions

- Fill in the blank:** _____ is a new partition scheme that overcomes the limitations of the MBR partition scheme.
- True or false?** A clean installation enables automated installation over a network.
- Fill in the blank:** _____ is a replacement for Server Core that takes up far fewer hardware resources than the two other installation options, has more periodic updates, and supports only 64-bit applications.
- Which of the following tools is provided by Microsoft to automate the installation of Windows Server 2025?
Choose two:
 - Windows ADK

- MDT
 - SharePoint Server 2022
 - SQL Server 2022
5. **True or false?** An unattended installation requires interactivity during the installation of an operating system.
6. **Fill in the blank:** _____ takes place when you bring in a new machine (physical or virtual) and you want to move the roles, features, apps, and settings into it.
7. Which of these are installation options in Windows Server 2025? **Choose three:**
- Desktop Experience
 - Server Core
 - Nano Server
 - KDE and GNOME
 - Windows PowerShell
8. Discuss the pros and cons of the three boot options: installation media (DVD), USB flash drive, and network boot.
9. Discuss the installation types: clean installation, network installation, unattended or automated installation, in-place upgrade, and migration.

Further reading

- *Boot to UEFI Mode or legacy BIOS mode:* <https://docs.microsoft.com/en-us/windows-hardware/manufacture/desktop/boot-to-uefi-mode-or-legacy-bios-mode>
- *Feature update, clean install, or migrate to Windows Server:* <https://learn.microsoft.com/en-us/windows-server/get-started/install-upgrade-migrate>
- *How to Setup Windows Server 2025? Step by Step Guide:* <https://medium.com/@theusapdf/how-to-setup-windows-server-2025-step-by-step-guide-0867cf024d36>

3

What to Do After Installing Windows Server 2025

This chapter provides a comprehensive guide on the essential tasks to undertake after installing Windows Server 2025. It is divided into three main sections for ease of understanding.

The first section focuses on device drivers, the crucial software components that facilitate communication between the operating system and hardware devices. You will learn how to perform various operations on device drivers, including **installation, removal, disabling, updating, upgrading, and rollback**.

The second section delves into the **Windows Server OS** registry, a structured database that holds configuration settings and options for both the operating system and installed applications. This section also covers how to manage programs running in the background, ensuring optimal performance and stability.

The third section emphasizes the initial configuration of Windows Server 2025, a vital step post-installation. You will learn how to set up basic settings such as the **computer name, network configurations, firewall settings, and system updates**. Each topic is thoroughly explained with step-by-step instructions and illustrative graphics.

At the end of the chapter, you will find an exercise designed to reinforce your understanding of the initial Windows Server configuration process.

In this chapter, we're going to cover the following main topics:

- Understanding and managing devices and drivers, including Plug and Play, IRQs, and driver signing
- Managing and optimizing registry entries and service accounts
- Performing initial server setup for better performance and security
- Performing an initial Windows Server configuration

Technical requirements

To effectively apply the concepts covered in this chapter, you will need the following resources:

- A computer running **Windows 11 Pro**, equipped with at least 16 GB of RAM, 1 TB of HDD, and an active internet connection
- A virtual machine with **Windows Server 2025 Standard** (Desktop Experience), equipped with at least 4 GB of RAM, 100 GB of HDD, and an internet connection

Understanding and managing devices and drivers, including Plug and Play, IRQs, and driver signing

An operating system is more than just a collection of code; it also plays a crucial role in managing the physical components of a computer. The interaction between hardware and software is integral to the system's functionality. To fully grasp this relationship, it's essential to understand how the operating system identifies and interacts with various hardware components. What mechanisms enable the OS to detect and manage these hardware elements?

Understanding computer devices and device drivers

Computers come in various shapes and sizes, but they all share standard physical components essential for their operation. These components can be categorized into four main types:

- **Internal devices:** These are located inside the computer case and include crucial hardware such as the power supply, motherboard, processor, memory, storage drives, and expansion cards. These elements form the computer's core structure.
- **External devices:** These are peripherals connected to the outside of the computer case, enhancing user interaction with the system. Examples include keyboards, monitors, mice, speakers, headphones, webcams, and microphones.
- **Peripheral devices:** Positioned near the computer but not essential for its basic operation, peripheral devices include printers, scanners, projectors, and plotters, which perform specific functions and tasks.
- **Network devices:** These are peripherals that connect to the computer via network cables, facilitating communication and shared resources. Examples include network printers, scanners, backup libraries, **network-attached storage (NAS)**, and **storage area networks (SAN)**.

Additionally, devices can be classified as input or output devices. Input devices, such as keyboards, send data to the computer, while output devices, such as monitors, display information from the computer. Some devices, such as touch-enabled screens, function as both input and output devices by receiving user input and displaying visual output.

NOTE

It is important to note that external devices are sometimes referred to as peripheral devices, as they are added to the computer system to extend its capabilities. Essentially, any device outside the core computer structure can be considered a peripheral.

A **device driver** is a software program that acts as an intermediary between the computer hardware and the operating system. It enables the OS to manage and control hardware components. Device drivers are typically included with the device on a DVD or can be downloaded from the manufacturer's website. However, many modern operating systems, such as Windows 10 and 11, support **Plug and Play (PnP)** technology, which allows for the automatic detection and configuration of new devices without the need for separate drivers. This chapter will further explore PnP, IRQ, DMA, and driver signing in the subsection titled *Understanding PnP, IRQ, DMA, and driver signing*. Next, we will delve into how to manage devices and device drivers to enhance server functionality.

Managing devices and device drivers

As an IT professional, you can manage devices and their drivers in Windows using two primary tools: **Windows Settings** and **Device Manager**. Windows Settings provides a contemporary interface for configuring and customizing various system aspects, including device management. On the other hand, Device Manager, a traditional utility, allows you to view and adjust the properties of device drivers that control hardware components.

In **Device Manager**, device drivers are represented by different icons indicating their status (see *Figure 3.1*):



Figure 3.1 – Device drivers' representation in Device Manager

These icons are of the following types:

- **Generic icon:** This indicates that the device is using a default or generic driver, which may not optimize the device's performance or functionality.
- **Black exclamation point in a yellow triangle:** This signifies that the device driver is either missing or incompatible. To resolve this issue, you need to install the appropriate driver for the device to function correctly.
- **Downward black arrow:** This indicates that the device is currently disabled. Although the driver is installed, the device is not active. To enable it, right-click on the device driver and select **Enable** from the menu.

By gaining a comprehensive understanding of these tools and their icons, you are now equipped with the knowledge to manage your devices and device drivers effectively. This knowledge empowers you to make informed decisions and take appropriate actions.

Customizing the Start menu for efficient navigation

Before diving into the various configurations of Windows Server 2025, know that personalizing the **Start** menu is a helpful way to streamline your workflow on Windows Server 2025, allowing quick access to frequently used applications and tools. By organizing and pinning essential apps, administrators can optimize server navigation, reducing time spent searching for core management utilities.

Pinning applications to the Start menu

The following steps will pin applications to the Start menu:

1. **Locate the application:** Find the application you want to pin, either in the **Start** menu itself or via the search bar.
2. **Right-click to pin:** Right-click the application and select **Pin to Start**. This action places the app as a tile in the **Start** menu.

3. **Organize your pinned apps:** Drag and arrange the pinned applications to organize them by function or frequency of use. You can also resize tiles to highlight priority apps.

Grouping and organizing tiles

For a more structured layout, group related applications together by dragging tiles to create named groups. This organization enables quick access to tools based on task type, such as networking tools, monitoring utilities, or common administrative applications.

By customizing the **Start** menu, you enhance accessibility to critical tools, facilitating a more efficient and responsive server management experience.

Working with devices and Device Manager

You can use Windows Settings and Device Manager to manage the devices and device drivers on your system. Windows Settings allows you to add, remove, and update devices and device drivers, as well as change their settings and preferences. Device Manager allows you to view the details of the device drivers, disable or enable them, uninstall or reinstall them, and troubleshoot any problems. Here are the steps to access these tools.

Windows Settings

To open Windows Settings, click the **Start** button and then click the **Settings** icon. Alternatively, you can press the **Windows key + I** shortcut. In Windows Settings, click **Devices** to see the list of devices connected to your system.

Device Manager

To open Device Manager, right-click the **Start** button and then select **Device Manager**. Alternatively, you can press the **Windows key + X** shortcut and then choose **Device Manager**. You can also type `devmgmt.msc` in the **Run** dialog box and press **Enter**. In Device Manager, you will see the device drivers organized by categories, such as display adapters, network adapters, sound, video, and game controllers. You can expand each category to see the specific device drivers under it, as shown in *Figure 3.2*.

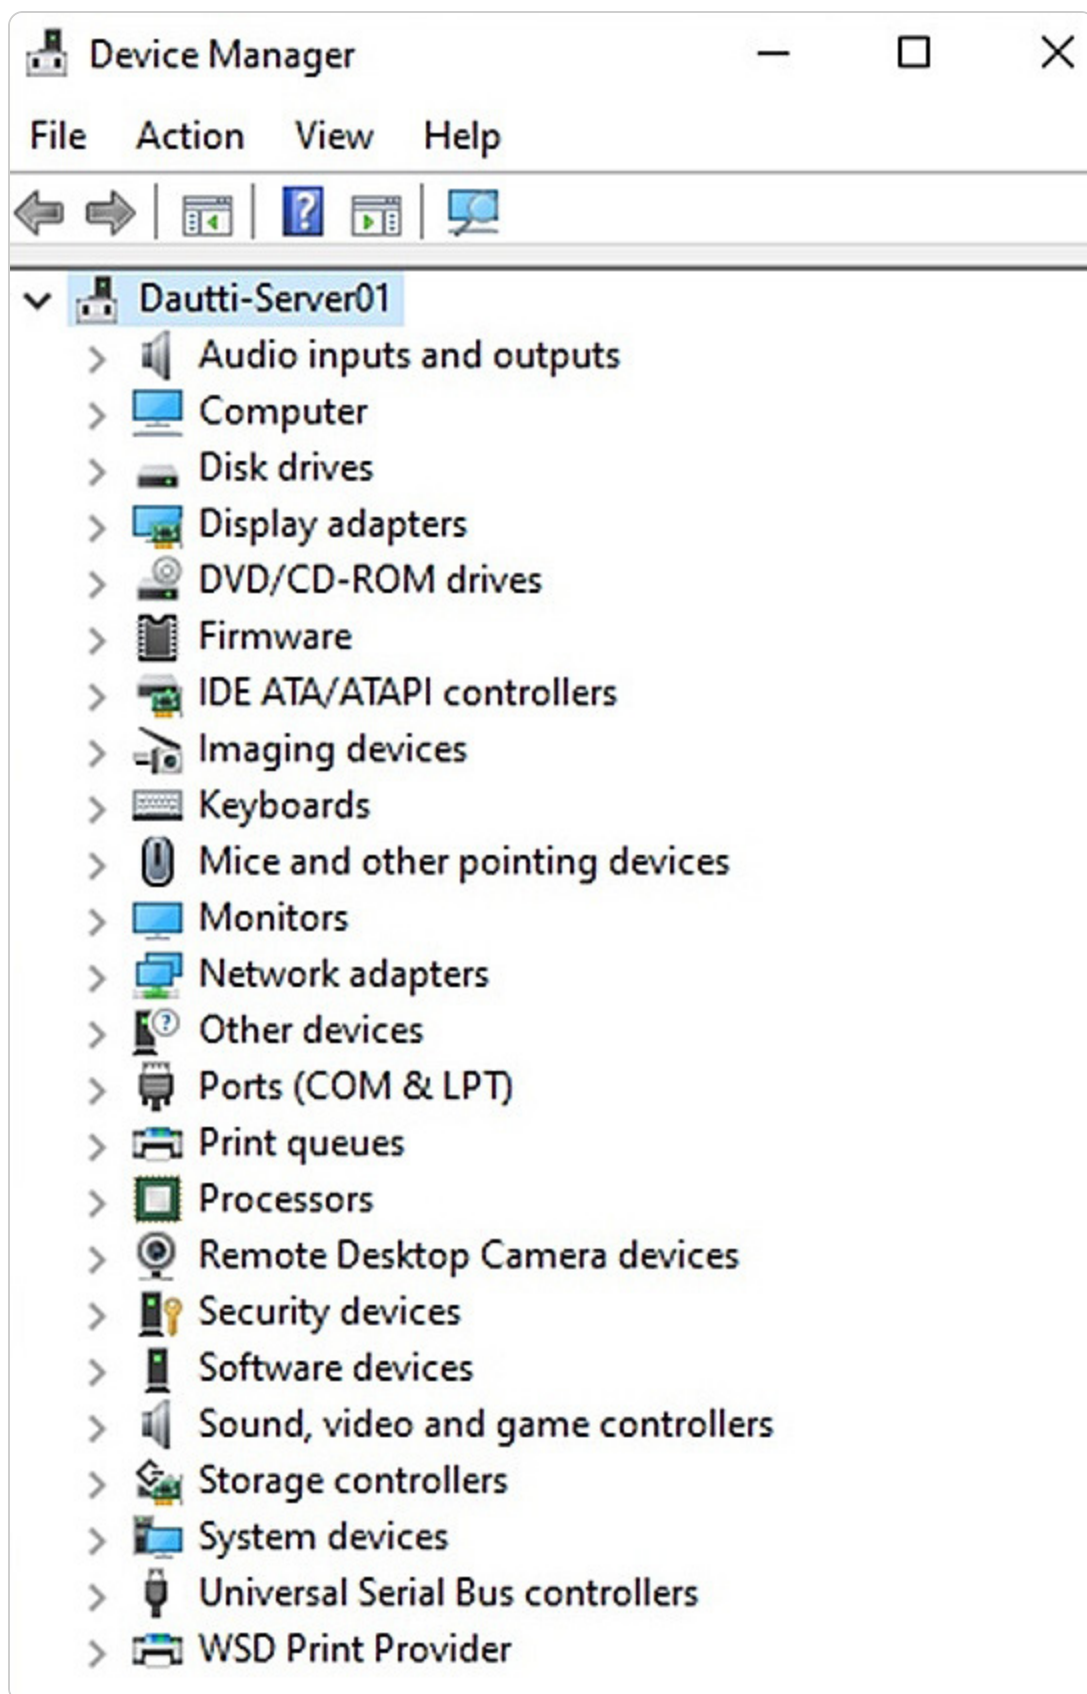


Figure 3.2 – The Device Manager

NOTE

The secret Start menu, also known as the Power User menu, is a hidden menu that provides quick access to various system tools and settings, such as Device Manager, Disk Management, Command Prompt, and Task Manager. You can open it by right-clicking the **Start** button or pressing the Windows key + X shortcut.

Adding devices and installing device drivers

To connect a new device to your system using **Windows Settings**, follow these steps:

1. Open the Start menu by clicking the **Start** button.
2. Select the **Settings** icon to open Windows Settings.
3. Click on **Devices**.
4. Navigate to **Bluetooth & other devices** in the **Devices** section.
5. Click **Add Bluetooth or other devices** to begin the connection process.

If you need to install a device driver from a file, whether it is on a DVD or downloaded from the internet, proceed with these steps:

1. Insert the DVD into the DVD drive or locate the driver file on your computer.
2. Open **File Explorer** and execute the **setup** or **install** file.

Once the device drivers are installed, you will learn how to update them in the next section.

NOTE

Windows Server offers multiple methods to access various settings and management windows. One efficient way is through **Microsoft Saved Console (MSC)** files, such as **Services.msc** for managing services. Many familiar MSC files are still supported in Windows Server 2025, including **virtmgmt.msc** for virtual machine management, **devmgmt.msc** for device management, and **services.msc** for service control. These shortcuts provide quick access to essential management consoles, streamlining server administration.

Updating device drivers

To update the device driver using **Device Manager**, take the following steps:

1. Right-click the **Start** button to open the secret Start menu.
2. In the secret **Start** menu, select **Device Manager**.
3. In the **Device Manager** window, expand the device's category.
4. Right-click the device and choose **Update driver** from the context menu.
5. Select **Browse my computer for drivers** (see Figure 3.3). If you lack a device driver, let the **Update Drivers** wizard do the work for you by clicking on **Search automatically for drivers**.

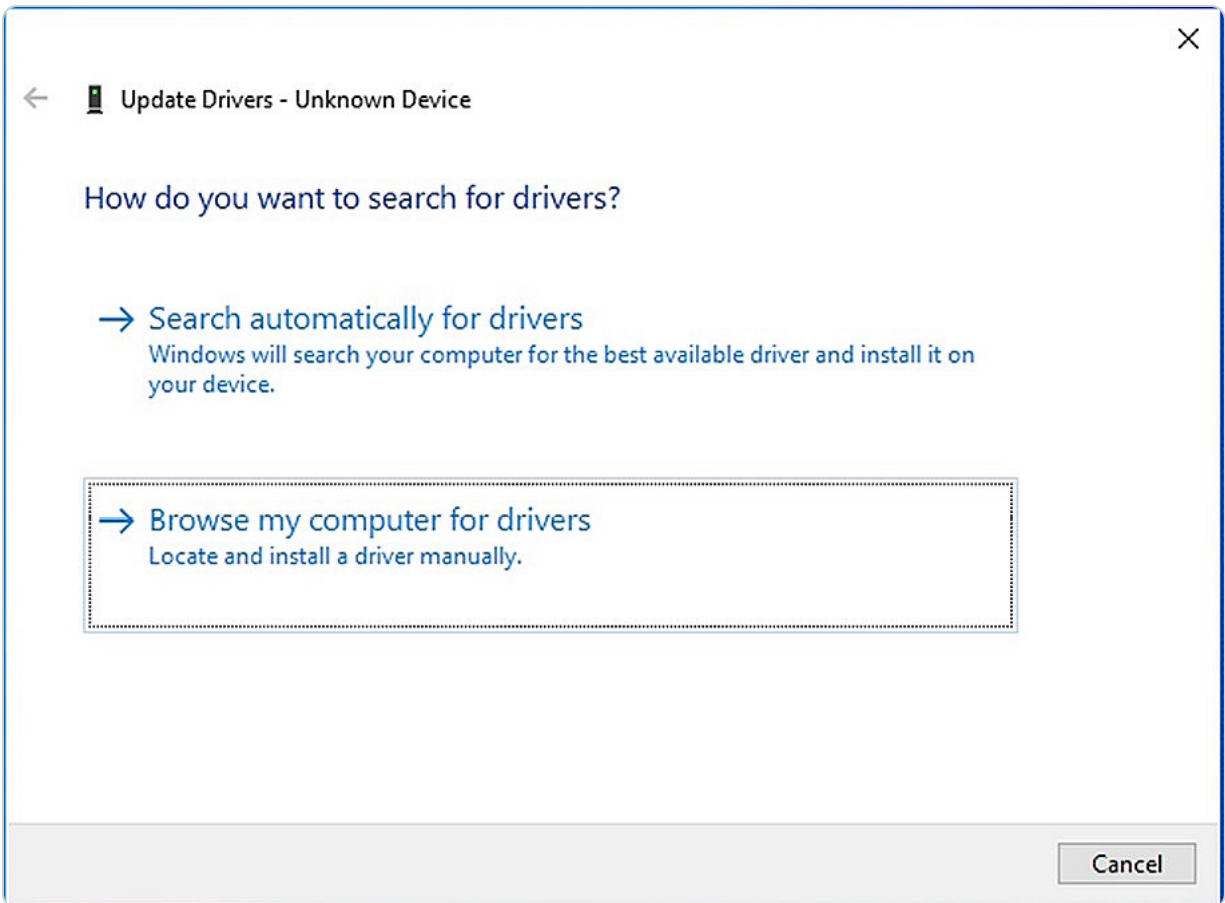


Figure 3.3 – Updating a device driver

IMPORTANT NOTE

To install or update a driver using Device Manager, select the **Update driver** option from the context menu. Additionally, it's important to prioritize security when installing drivers, especially those downloaded from the internet. Choose digitally signed drivers whenever possible, as these have been validated for authenticity and integrity, minimizing the risk of malware or corrupted files. Digitally signed drivers offer extra assurance against tampering, which is essential for keeping your server environment secure.

Now, let us learn how to remove and uninstall device drivers.

Removing devices and uninstalling device drivers

You can use **Windows Settings** to disconnect a device from your system. Here are the steps to do so:

1. Open the **Start** menu by clicking the **Start** button.
2. Click the **Settings** icon on the **Start** menu.
3. Click **Devices** in **Windows Settings**.
4. Click **Bluetooth & other devices** in the **Devices** navigation menu and choose the device you want to disconnect.
5. Click the **Remove device** button, as shown in *Figure 3.4*.



Figure 3.4 – Disconnecting a device

You can use **Device Manager** to delete a device driver from your system. Here are the steps to do so:

1. Right-click the **Start** button to open the secret **Start** menu.
2. In the secret **Start** menu, select **Device Manager**.
3. In the **Device Manager** window, expand the device's category.
4. Right-click the device and choose **Uninstall device** from the menu that appears.
5. Click the **Uninstall** button.

Next, let us learn how to manage and disable device drivers.

Managing devices and disabling device drivers

You can manage devices through **Windows Settings** by following these steps:

1. Click the **Start** button to open the **Start** menu.
2. Select the **Settings** icon.
3. Navigate to **Devices**.
4. In the **Devices** section, choose **Printers & scanners** and select the device you wish to manage.
5. Click on the **Manage** button.

Alternatively, to disable a device driver using **Device Manager**, proceed with these steps:

1. Right-click the **Start** button to open the secret **Start** menu.
2. Select **Device Manager** from the menu.
3. In the **Device Manager** window, locate the device category and click it to expand the list.
4. Right-click the specific device and select **Disable device** from the context menu.
5. Confirm your action by clicking **Yes** in the confirmation dialog, as illustrated in *Figure 3.5*.

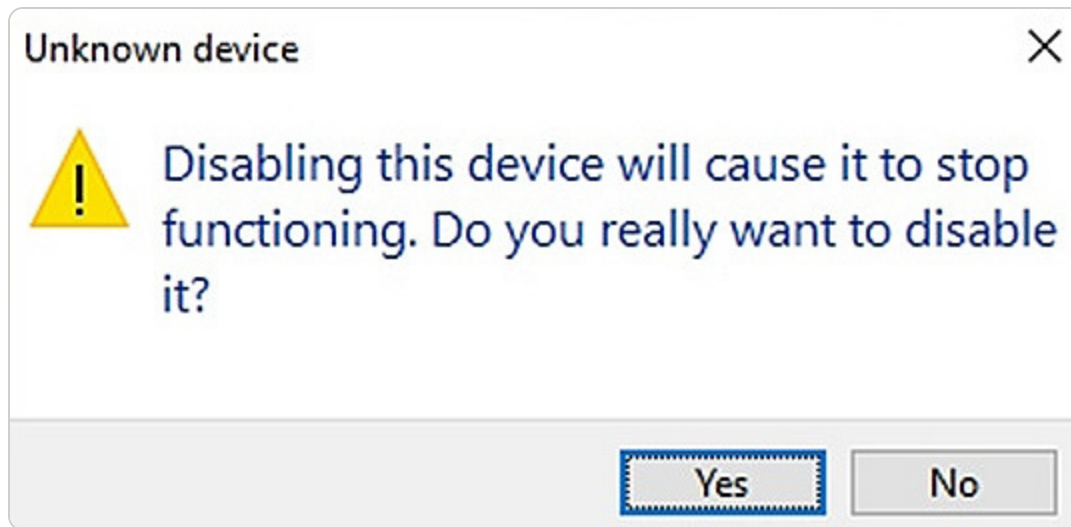


Figure 3.5 – Disabling the device driver

In the next section, we will cover how to roll back device drivers.

Rolling back device drivers

One of the ways to restore a device driver to its previous version is by using **Device Manager**. Here are the steps you need to follow for this method:

1. Right-click the **Start** button to access the secret Start menu.
2. From the secret **Start** menu, choose **Device Manager**.
3. In the **Device Manager** window, locate the device's category and click it to show the devices under it.
4. Right-click the device and select **Properties**.
5. Click the **Driver** tab and then click the **Roll Back Driver** button, as shown in *Figure 3.6*. Click **OK**.

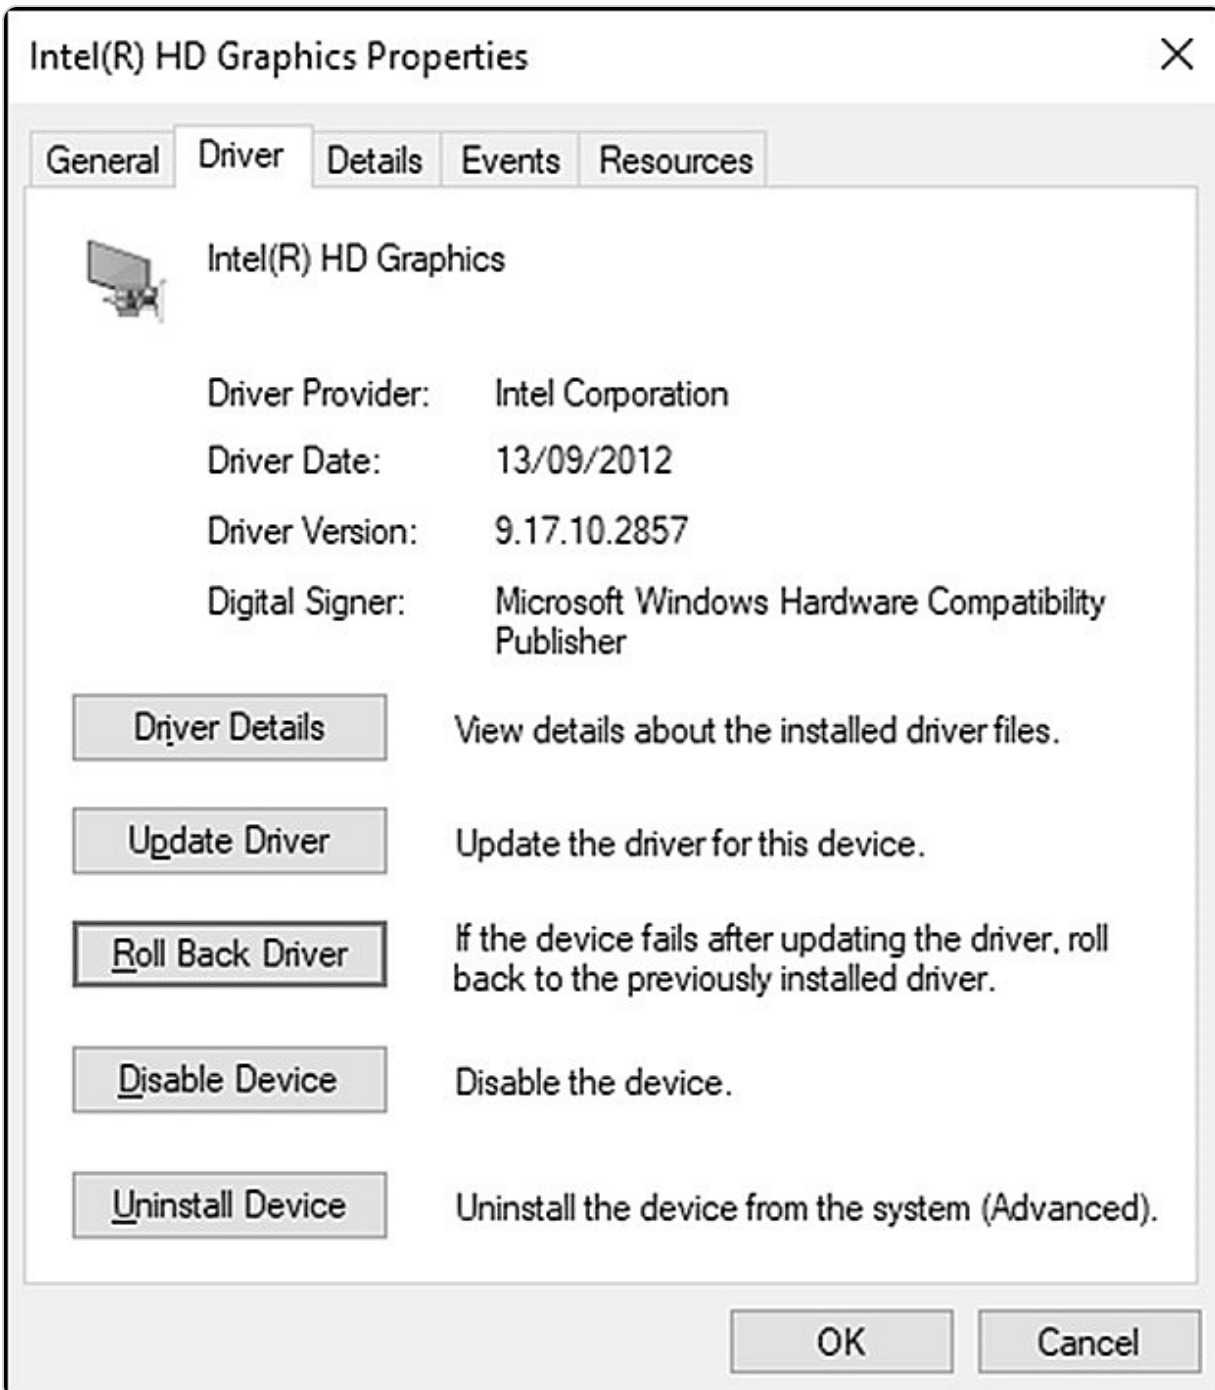


Figure 3.6 – Restoring a device driver to its previous version

NOTE

If you encounter issues after updating a device driver, rolling it back to a previous version may resolve the problem. Note, however, that if the driver you installed is the first one for that device, rollback won't be available. Also, be cautious with driver dependencies and always back up the server before rolling back to prevent potential disruptions.

Next, let us learn about some solutions when troubleshooting device drivers.

Troubleshooting a device driver

One of the challenges you may face with device drivers is that they may not work correctly or cause errors. In such cases, you have several alternatives (see *Figure 3.7*) to fix them:

- **Update Driver:** This allows you to install the latest driver automatically or manually select the driver software from the server
- **Roll Back Driver:** This allows you to restore the previous version of the driver if the current one is problematic
- **Disable Device:** This allows you to turn off the driver if it is causing significant issues, such as system instability
- **Uninstall Device:** This allows you to remove the existing driver if you have obtained the correct driver from the device manufacturer

Figure 3.7 – Troubleshooting options for the device driver

In this section, you have learned about the different methods to manage devices. Next, let us explore various system resources, such as PnP, IRQ, DMA, and driver signing.

Understanding PnP, IRQ, DMA, and driver signing

The operating system relies on system resources to manage hardware components such as the CPU, memory, disks, **Input/Output (I/O)** devices, and network connections. Key system resources include **I/O ports**, **memory addresses**, **interrupt request (IRQ) lines**, and **direct memory access (DMA) channels**, all of which are crucial for the OS to communicate with hardware effectively.

PnP

Developed through a partnership between Intel and Microsoft, PnP significantly streamlined the process of installing devices and their drivers. With PnP, you can connect a device to your computer, and the Windows OS will automatically detect and configure it. The system uses its **Driver Store** to install the necessary drivers for the device. For instance, in Windows Server 2025, the Driver Store is located at `C:\Windows\System32\DriverStore`.

Interrupt Request (IRQ) and Direct Memory Access (DMA)

In contemporary computing, IRQs are numbers ranging from 0 to 31 that signify signals sent by devices to request processor attention when they require processing. Conversely, DMA channels, numbered from 0 to 8, allow devices to access RAM directly without involving the processor.

To view IRQ and DMA resource settings using **Device Manager**, follow these steps:

1. Right-click the **Start** button to open the context menu.
2. Select **Device Manager**.

3. In the **Device Manager** window, expand the relevant device category.
4. Right-click the device and choose **Properties** from the context menu.
5. Click the **Resources** tab to review the resource settings, as depicted in *Figure 3.8*.

Figure 3.8 – Driver's Resources settings

Next, look at the driver's signature, which helps verify both the driver's integrity and identity.

Digital signature of the driver

A driver's **digital signature** serves as an electronic mark that verifies both the source and integrity of the driver's package. That indicates that Microsoft has validated and certified the driver as being both compatible with and secure for installation. This digital signature is essential for confirming that the driver remains unaltered and authentic. To verify a driver's digital signature in Windows Server 2025, follow these steps:

1. Right-click the **Start** button to open the context menu.
2. Select **Device Manager** from the menu.
3. In the **Device Manager** window, expand the category of the relevant device.
4. Right-click the device and choose **Properties** from the context menu.
5. Go to the **Driver** tab and click the **Driver Details** button to access the **Driver File Details** window, as illustrated in *Figure 3.9*.

Figure 3.9 – Digital signature information of the driver

This section has covered the essentials of computer devices and drivers, including various management techniques and the system resources used for device communication. The upcoming section will delve into the Windows Server registry and services, focusing on configuration and management practices for Windows registries and services.

Managing and optimizing registry entries and service accounts

The **Windows registry** and **Windows services** are critical components for managing and configuring the Windows operating system. The Windows registry is a structured database that keeps track of system settings, application configurations, and hardware options. Meanwhile, Windows services are background processes that handle essential functions such as networking, security, and printing. This section will guide you through the management and optimization of both the Windows registry and Windows services, utilizing various tools and techniques to ensure efficient system performance and stability.

Windows Server registry

The Windows registry is a fundamental component that tracks changes to the server's hardware and software. It functions as a hierarchical database, maintaining configuration and security settings for the operating system, applications, and hardware devices. The registry's interface displays a console tree on the left side, organized into five primary **registry keys** known as **hives (HKEYs)**. These keys follow a syntax similar to file paths, separated by backslashes. In Windows Server 2025, the five HKEYs are as follows:

- **HKEY_CLASSES_ROOT**: Contains information about installed applications and their associated file extensions
- **HKEY_CURRENT_USER**: Stores settings and data related to the currently logged-in user
- **HKEY_LOCAL_MACHINE**: Holds configuration data specific to the local computer, including details about hardware and software
- **HKEY_USERS**: Maintains data on all user profiles that have accessed the server
- **HKEY_CURRENT_CONFIG**: Records data collected during the system boot process, such as display settings

Services Control Manager and Windows Server services

The **Services Control Manager** is a crucial utility for overseeing the services operating on your Windows Server. **Services** are background processes that offer various functionalities essential for the operating system, applications, and network operations. Through the Services Control Manager, you can start, stop, restart, or pause these services as needed. Additionally, this tool allows you to configure service properties, including startup types, dependencies, and recovery options, to ensure optimal performance and reliability of your server.

Understanding service startup types

In the Services Control Manager, each service is assigned a startup type that determines how and when it is activated (see *Figure 3.10*). The available startup types are as follows:

- **Automatic**: The service is initiated automatically by the operating system during the boot process.
- **Automatic (Delayed Start)**: The service begins approximately two minutes after all other automatic services have started.
- **Manual**: The service must be manually started by a user or another service that relies on it.
- **Disabled**: The service is not initiated by the operating system, users, or dependent services.

Figure 3.10 illustrates these different service startup types in Windows Server.

Figure 3.10 – Windows Server service startup types

Next, we will explore how to access and manage both services and registry settings.

Accessing and managing the Windows registry and services

To effectively manage the Windows registry and Windows services, you can utilize two essential tools: the **Registry Editor** and the **Services Control Manager**. The Windows registry serves as a comprehensive database that holds important configuration and security settings for the operating system, applications, and hardware devices. Windows services, on the other hand, are background processes that perform various functions that are crucial to the operating system, applications, and network operations. This section will guide you through the process of accessing and managing these components using the Registry Editor to handle registry settings and the Services Control Manager to oversee service operations.

Working with registry keys and values in the Registry Editor

To manage and modify the Windows registry, which stores critical settings for the operating system, applications, and hardware, you can use the Registry Editor. Here is how you can access it:

1. Type **regedit** into the search bar on the taskbar and press *Enter*.
2. The **Registry Editor** will open, as illustrated in *Figure 3.11*.

Figure 3.11 – Windows Server Registry Editor

In the following subsection, you will learn the procedures for altering registry values.

Modifying a registry value

To alter the data of a registry value, which controls the behavior of the operating system, applications, or hardware devices, you can use the Registry Editor. Follow these steps:

1. Type **regedit** into the search box on the taskbar and press *Enter*.
2. In the Registry Editor, navigate to the appropriate registry key and sub-key(s) on the left pane that includes the value you wish to modify.
3. On the right pane, locate the value you want to change, right-click it, select **Modify**, and then adjust the **Value data** field, as shown in *Figure 3.12*.

Figure 3.12 – Changing a registry value

Next, you will learn how to rename a registry value.

Renaming a registry value

To rename a registry value that impacts the settings of the operating system, applications, or hardware devices, you can use the Registry Editor. Here is how:

1. Search for **regedit** on the taskbar and press *Enter*.
2. In the Registry Editor, navigate to the relevant registry key and sub-key(s) on the left pane.
3. On the right pane, locate the value you wish to rename, right-click on it, select **Rename**, and then enter the new name, as illustrated in *Figure 3.13*.

Figure 3.13 – Renaming a registry value

Following this, you will learn how to delete a registry value.

Removing a registry value

To delete a registry value that affects the operation of the operating system, applications, or hardware devices, you can utilize the Registry Editor. Here is how to do it:

1. Search for **regedit** in the taskbar and press *Enter*.
2. In the Registry Editor, navigate to the registry key and sub-key(s) on the left pane where the value you want to delete is located.
3. On the right pane, find the value you wish to remove, right-click it, select **Delete**, and confirm your action, as depicted in *Figure 3.14*.

Figure 3.14 – Deleting a registry value

NOTE

The procedures for managing registry values, such as deleting, renaming, and exporting, can also be applied to registry keys. That allows for consistent and efficient management of both values and keys within the Windows Registry.

Managing and accessing Windows services

To manage and access Windows services—background programs that handle various tasks for the operating system and applications—follow these steps:

1. Click the **Start** button.
2. Select the **Windows Tools** option from the **Start** menu.
3. Locate and click on **Services** from the available list.
4. That will open the **Windows Services Control Manager**, as illustrated in *Figure 3.15*.

Figure 3.15 – Windows Services Control Manager

The following subsection will guide you through configuring service recovery options using the Control Manager.

Configuring how to recover from service failures

To configure how your computer responds to service failures, such as restarting the service or rebooting the system, follow these steps using the **Control Manager**:

1. Click the **Start** button.
2. Select **Windows Tools** from the **Start** menu.
3. Find and click on **Services** from the list.
4. In the **Services** window, locate the service you want to configure and right-click on it.
5. From the context menu, select **Properties**.
6. In the **Properties** window, navigate to the **Recovery** tab. Note that while many services allow you to configure recovery options, some critical system services do not provide these settings, as they are essential for stable operation.
7. Define the actions for the first, second, and subsequent failures of the service, as depicted in *Figure 3.16*. In Windows OSs, specific actions can be defined for managing service failures at different stages to maintain stability and minimize disruption. For the first failure, the system can be set to restart the service to reduce downtime automatically. Upon a second failure, the service can be restarted again, or alternate actions, such as running a custom script or sending a notification to the administrator, can be triggered. For any subsequent failures, options include restarting the system, taking no action, or continuing to attempt restarts after a specified delay. These configurable actions provide a flexible approach to maintaining service reliability and system performance.

NOTE

In Windows operating systems, specific actions can be defined to manage service failures at different stages, ensuring stability and minimizing disruption. For the first failure, the system can be configured to restart the service, reducing downtime automatically. If a second failure occurs, the service can be restarted again, or alternative actions, such as running a custom script or sending a notification to the administrator, can be triggered. For any subsequent failures, options include restarting the system, taking no action, or continuing to attempt restarts after a specified delay. These configurable actions offer a flexible approach to maintaining service reliability and system performance.

Figure 3.16 – Configuring recovery options for service failures in Windows Server 2025

8. Click **OK** to apply your changes and close the dialog box.

In the following subsection, we will explore how to set a delay for the start of a service.

Delaying the start of a service

To delay the start of a service using Control Manager, which can enhance system performance and stability by initiating the service after other automatic services, follow these steps:

1. Click the **Start** button.
2. From the **Start** menu, select **Windows Tools**.
3. Locate and click on **Services**.
4. In the **Services** window, right-click on the service you wish to delay.
5. Choose **Properties** from the context menu.

6. In the **Properties** window, navigate to the **General** tab and click the **Startup type** drop-down list.
7. Select **Automatic (Delayed Start)** from the options, as depicted in *Figure 3.17*.

Figure 3.17 – Configuring the delayed startup of a service in Windows Server 2025

8. Click **OK** to apply the changes and close the dialog box.

In the following subsection, we will discuss how to configure the logon settings for a service.

Logon settings for a service

To configure the user account under which a service operates using Control Manager, follow these steps to enhance security and manageability:

1. Click the **Start** button.
2. From the **Start** menu, select **Windows Tools**.
3. Find and click on **Services**.
4. In the **Services** window, right-click the service you wish to configure.
5. Choose **Properties** from the context menu.
6. In the **Properties** window, navigate to the **Log On** tab.
7. Select the **This account** option under the **Log on as** section.
8. Enter the user account name, including the domain name followed by a backslash, and input the password and confirm it, as illustrated in *Figure 3.18*.

Figure 3.18 – Configuring the log-on settings for a service in Windows Server 2025

9. Click **OK** to apply the changes and close the dialog box.

In the following subsection, we will explore how to start the configured service.

Starting the service

To start a service on your computer using Control Manager, follow these steps:

1. Click the **Start** button.
2. From the **Start** menu, select **Windows Tools**.
3. Locate and click on **Services** from the list.
4. In the **Services** window, right-click the service you wish to start.
5. From the context menu, select **Start**, as depicted in *Figure 3.19*.

Figure 3.19 – Starting the service

In the upcoming subsection, we will cover how to stop a service.

Stopping a service

To stop a service on your computer using Control Manager, follow these steps:

1. Click the **Start** button.
2. Select **Windows Tools** from the **Start** menu.
3. Choose **Services** from the available options.
4. In the **Services** window, right-click the service you wish to stop.
5. Select **Stop** from the context menu, as depicted in *Figure 3.20*.

Figure 3.20 – Stopping the service

In the next section, we will cover how to restart a service.

Restarting the service

To restart a service on your computer using Control Manager, follow these steps:

1. Click the **Start** button.
2. Select **Windows Tools** from the **Start** menu.
3. Locate and choose **Services** from the list.
4. In the **Services** window, right-click the service you wish to restart.
5. From the context menu, select **Restart**, as illustrated in *Figure 3.21*.

Figure 3.21 – Restarting the service

In the upcoming section, we will explore registry entries, service accounts, and dependencies.

Understanding registry entries, service accounts, and dependencies

In managing Windows Server, you may find the need to create or modify registry keys and values to address issues or introduce new features. Given the critical nature of the Windows registry, it is essential to exercise caution during any modifications. Windows Server also relies on various services, which operate under different service accounts. These accounts can be native to Windows Server or custom ones created for managing services. The choice of service account impacts the level of access that services have to local and network resources, affecting security. Windows Server 2025 includes several native service accounts, as illustrated in *Figure 3.22*:

- **Local System:** This built-in account possesses the highest level of privileges on a Windows OS, often referred to as a superuser, and surpasses even administrator accounts in terms of power

- **NT Authority\Local Service:** This built-in account has the same privileges as members of the user group, providing a minimal level of access
- **NT Authority\Network Service:** This built-in account has more privileges than user group members, offering greater access

Furthermore, services often have dependencies on other services to operate effectively. Stopping a service with dependencies will require halting the dependent services, while starting such a service necessitates starting its dependencies as well.

Figure 3.22 – Native service accounts in Windows Server 2025

In the following section, we will cover how to create a new registry key.

Creating a New Registry Key

The Registry Editor allows you to create new registry keys or sub-keys within the Windows registry, a crucial component for storing system configuration settings and options. To add a new registry key, follow these steps:

1. Open the Registry Editor by typing **regedit** into the search box on the taskbar and pressing *Enter*.
2. In the Registry Editor, navigate to the desired location in the left pane where you want to create the new key. Right-click on the chosen registry key or sub-key.
3. Select **New** and then **Key** from the context menu, as illustrated in *Figure 3.23*.

Figure 3.23 – Adding a registry entry

NOTE

*To add a new registry value, right-click in the empty area on the right pane of the Registry Editor after selecting the relevant registry key or sub-key. Then, choose **New** from the context menu and enter the desired value.*

In the upcoming section, we will explore how to create a service account.

Creating a service account

A **service account** is a specialized account used to run background services on Windows Server, performing tasks such as managing network access, enhancing security, and monitoring performance. These accounts can be targets for malicious activity, so using **managed service accounts (MSAs)** in a domain environment is beneficial; MSAs automatically rotate passwords and support audit trails for added security. In cloud environments, managed identities offer similar security benefits by handling identity and access management. To assign a service account to a service, follow these steps:

1. Click the **Start** button and select **Windows Tools** from the **Start** menu.

2. Locate and choose **Services** from the list of tools.
3. In the **Services** window, right-click on the service you wish to assign a service account to.
4. From the context menu, select **Properties**.
5. In the **Properties** window, navigate to the **Log On** tab, as illustrated in *Figure 3.24*.

Figure 3.24 – Creating a service account

6. Choose the **This account** option and click the **Browse** button to locate the desired service account. Ensure the account is a valid user in your organization's **Active Directory** with administrator privileges, as only administrators can be set as the **Log On** account.
7. Enter and confirm the password for the selected service account.
8. Click **OK** to apply the changes and close the **Properties** window.

In the next section, we will cover how to set up service dependencies.

Adding a service dependency

To establish a service dependency, you'll need to use the Registry Editor and follow these steps:

1. Open the Registry Editor by typing **regedit** into the taskbar search box and pressing *Enter*.
2. Navigate to the service in question located under **HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services**.
3. Locate the **DependOnService** value or create a new multi-string value with this name if it doesn't already exist.
4. Enter the precise name of the service on which the current service will depend.
5. Restart the server to apply the changes.
6. After the server restarts, open the **Services** window, locate the service you configured, right-click it, and select **Properties**.
7. Go to the **Dependencies** tab to confirm that the newly added service appears as a dependency.

This process demonstrates how to use the Registry Editor and services to manage service dependencies in Windows Server. Next, we will explore configuring initial settings for both the Desktop Experience and Server Core installation options.

Performing initial server setup for better performance and security

One of the essential tasks after installing Windows Server 2025 is to configure the initial settings for the server. That involves customizing the server name and domain membership according to the role that the server will play in the network. Furthermore, it includes enabling **Remote Desktop** for remote management, setting up a static IP address for network identification, changing the time zone to match the local time, activating Windows Server 2025 with a valid license key, turning off **Internet Explorer (IE)** enhanced security for easier browsing, and checking for updates to keep the server

secure and updated. These steps will prepare the server for its intended function in the IT environment.

Initial settings for Windows Server

Before assigning roles to your server, it's crucial to configure several fundamental settings that impact both performance and security. Begin by setting a static IP address to ensure the server's network identity. Adjust the time zone to reflect the local time accurately, activate your Windows Server 2025 license, and apply the latest updates and patches. Rename the server to something meaningful, join it to your domain to facilitate resource access, enable Remote Desktop to allow for remote management, and turn off IE enhanced security to simplify web access. This section will guide you through these initial setup steps for both **Desktop Experience** and **Server Core** installation options.

Configuring the server with Server Manager

Server Manager is a graphical tool that you can use to customize the basic settings of your server in Desktop Experience. It opens automatically when you log in to Windows Server 2025 for the first time, and you can change this behavior if you want. To access the server configuration options using Server Manager, click on **Configure this local server** in the **WELCOME TO SERVER MANAGER** section, as shown in *Figure 3.25*.

Figure 3.25 – Server Manager in Windows Server 2025

Alternatively, you can use the **Server Configuration tool**, which is a command-line tool that allows you to configure the server in CLI mode. We will discuss this tool in the next section.

Initial server settings in CLI mode

When setting up a Windows Server 2025 installation in **Server Core** mode, you can use the Server Configuration tool to manage essential settings. Access this command-line interface by entering `sconfig.cmd` at the Command Prompt, as illustrated in *Figure 3.26*.

Figure 3.26 – Server Configuration tool in CLI mode

The Server Configuration tool provides options to modify various server settings, including the server name, domain membership, IP address, time zone, remote management settings, product activation, and updates.

Managing configuration drift with PowerShell Desired State Configuration

After setting up Windows Server 2025, maintaining consistent configurations across multiple servers is crucial to avoid configuration drift—a common issue where servers gradually deviate from their intended settings over time. Configuration drift can impact security, performance, and compliance, and tracking down inconsistencies can become challenging without proper tools.

PowerShell **Desired State Configuration (DSC)** provides a powerful solution to manage and automate configuration settings. DSC enables administrators to define the desired state of a server configuration through PowerShell scripts, which specify the exact setup requirements for each server. Once defined, DSC continuously monitors server configurations to ensure they align with the desired state and can automatically correct any deviations. This feature is valuable in environments with multiple servers where consistent configurations are necessary for compliance and stability.

Figure 3.27 – Querying the current state of the configuration

PowerShell DSC can be leveraged to set up configurations such as network settings, firewall rules, and installed features. Additionally, DSC can be used to enforce these configurations across various servers, making it easier to maintain alignment with organizational standards and preventing issues associated with configuration drift.

Validating hardware stability with memory testers and burn-in applications

Before deploying a new server, it's essential to ensure that the hardware can withstand the demands of a production environment. Conducting hardware validation tests—particularly on newly installed memory and other critical components—can help detect potential faults early. Memory testers and burn-in applications rigorously evaluate hardware under simulated workloads, allowing IT teams to identify weak points that may otherwise cause issues later on.

Most server manufacturers, including Dell, HP, and Cisco, provide built-in utilities designed explicitly for hardware testing. These tools are tailored to each manufacturer's systems, offering diagnostic features that run comprehensive checks on components such as memory, storage, and processors. Leveraging these utilities before deployment not only boosts confidence in the server's stability but also helps minimize unplanned downtime by catching failures that may occur shortly after initial setup.

In the following section, we will guide you through the process of configuring your server using this tool. Let us get started.

Chapter exercise – performing an initial Windows Server configuration

This exercise aims to demonstrate how to configure the fundamental settings of Windows Server using two distinct tools: Server Manager and Server Configuration. Server Manager, a graphical interface available in Desktop Experience mode, allows you to adjust various server settings, including the server name, IP address, domain, remote access, time zone, activation, and updates. Conversely, Server Configuration is a command-line utility used in Server Core mode to perform similar tasks. This chapter will guide you through the process of using both tools to customize your server's initial setup to meet your specific requirements.

Using Server Manager to configure the initial settings for Windows Server

In this section, you will learn how to use **Server Manager**, a user-friendly graphical interface, to adjust the basic settings for Windows Server 2025 Standard (Desktop Experience), such as the server name, domain, IP address, and more.

Renaming the server

To rename your server, follow these steps, as illustrated in *Figure 3.28*:

1. Begin by selecting the highlighted default computer name in the **Properties** section.
2. In the **System Properties** window, click the **Change** button to access the **Computer Name/Domain Changes** dialog.
3. Enter the desired new name for your server in the **Computer name** field and click **OK**.

Figure 3.28 – Renaming the server

4. A prompt will appear, notifying you that a restart is required for the changes to take effect. Click **OK** to proceed.
5. Close the **System Properties** window by clicking the **Close** button.
6. In the **Microsoft Windows** dialog box, choose **Restart Now** to reboot the server and apply the changes.

Following this, we will guide you through the process of joining the server to a domain.

Connecting the server to a domain

Connecting your server to an organizational domain is essential based on its intended role. For instance, if the server is to act as a **domain controller (DC)**, it will automatically join the domain upon the installation of the **Active Directory Domain Services (AD DS)** role, and no additional steps are

necessary. However, if the server will fulfill any other role, it needs to be joined to the domain manually. To connect the server to a domain, follow these steps, as illustrated in *Figure 3.29*:

1. In the **Properties** section, click on the highlighted workgroup name.
2. In the **System Properties** window, select the **Change** button.
3. In the **Computer Name/Domain Changes** dialog, opt for the **Domain** selection, enter your organization's domain name, and click **OK**.

Figure 3.29 – Connecting the server to a domain

4. In the **Windows Security** prompt, provide the credentials for an account with domain joining permissions, then click **OK**.
5. A confirmation message will appear welcoming the server to the domain; click **OK** to close this message.
6. Confirm that you want to restart the server by clicking **OK**.
7. Close the **System Properties** window by clicking **Close**.
8. When prompted by the **Microsoft Windows** dialog box, click **Restart Now** to apply the changes.

Following these steps will successfully join the server to your domain. The next topic will cover enabling Remote Desktop to allow remote access to the server.

Enabling Remote Desktop

To enable **Remote Desktop** on your server, follow these steps, as depicted in *Figure 3.30*:

1. Open the **System Properties** dialog box and navigate to the **Remote** tab.
2. Select the **Allow remote connections to this computer** option in the **System Properties** window.
3. A notification will appear indicating that the Remote Desktop firewall exception will be enabled. Click **OK** to acknowledge this and close the dialog.

Figure 3.30 – Enabling Remote Desktop

4. To designate users who can access the server via Remote Desktop, click the **Select Users...** button.
5. In the **Remote Desktop Users** window, click **Add** to choose users or groups from your AD DS. After selecting the appropriate users or groups, click **OK** to finalize your selections and close the window.
6. Click **OK** again to exit the **System Properties** window.

After enabling Remote Desktop, the next topic will cover configuring a fixed IP address for your server, which is essential for stable network operations.

Configuring the IP address

Configuring a **static IP address** for your server is essential to ensure that its address remains consistent. To set this up, follow these steps, as illustrated in *Figure 3.31*:

1. Access the Ethernet settings highlighted in the **Properties** section.

Figure 3.31 – Configuring the IP address

2. In the **Network Connections** window, right-click on your server's Ethernet connection and select **Properties** from the context menu.
3. Within the **Ethernet Properties** window, locate **Internet Protocol Version 4 (TCP/IPv4)** and click the **Properties** button.
4. In the **Internet Protocol Version 4 (TCP/IPv4) Properties** window, choose **Use the following IP address** and input the desired values for **IP Address**, **Subnet Mask**, and **Default Gateway**. Additionally, select **Use the following DNS server addresses** and enter the **Preferred DNS server** and **Alternate DNS server** details. Click **OK** to apply these settings.
5. Close the **Ethernet Properties** window by clicking **Close**.
6. Exit the **Network Connections** window by clicking the red **X** button in the top-right corner.

In the following section, we will cover how to check for updates, an important task to complete after installing Windows Server 2025.

Verifying the updates

To verify and apply updates on your server, follow these steps, as illustrated in *Figure 3.32*:

1. Select the **Last checked for updates** setting highlighted in the **Properties** section.
2. In the **Settings** window, under the **Windows Update** section on the right side, review the list of available updates, if any. If updates are available, click the **Install now** button to proceed with the installation.

Figure 3.32 – Verifying the updates

Please note that the update process may take some time, and a server restart is often required for the updates to take full effect.

NOTE

For administrators who prefer command-line tools, Windows Server 2025 supports using WinGet and PowerShell to streamline update management. Running `winget upgrade --all` provides a straightforward way to download and install available updates. Alternatively, the `PSWindowsUpdate` module in PowerShell offers comprehensive options for managing updates, including scheduling and automation, which can enhance efficiency in server maintenance tasks.

In the next section, we will guide you on how to disable IE Enhanced Security Configuration in Windows Server 2025, despite Microsoft officially discontinuing IE.

Turning off IE Enhanced Security Configuration

To disable **IE Enhanced Security Configuration**, follow these steps, as depicted in *Figure 3.33*:

1. Access the **IE Enhanced Security Configuration** setting highlighted in the **Properties** section.

2. In the **Internet Explorer Enhanced Security Configuration** window, select the **Off** option under the **Administrators** section.
3. Click **OK** to apply the changes and close the window.

Figure 3.33 – Disabling IE Enhanced Security Configuration

NOTE

Internet Explorer Enhanced Security Configuration (IE ESC) is designed to minimize the exposure of servers to potential security risks by adjusting security settings. While IE ESC primarily affects Internet Explorer, it can also impact Microsoft Edge when using IE mode for compatibility with legacy applications. Disabling IE ESC can help resolve access issues but may increase the server's vulnerability to security threats. It is essential to balance security needs with accessibility requirements when managing IE ESC settings. With this setting adjusted, the next topic will guide you through changing the time zone, which is crucial for maintaining accurate timestamps and ensuring proper operation of network services.

Adjusting the time zone

To adjust the time zone, follow these steps, as illustrated in *Figure 3.34*:

1. Click on the highlighted time zone setting in the **Properties** section.
2. In the **Date and Time** window that appears, select the **Change time zone** button.
3. In the **Time Zone Settings** window, choose your desired time zone from the drop-down menu.
4. Click **OK** to close the **Time Zone Settings** window.
5. Click **OK** again to close the **Date and Time** window.

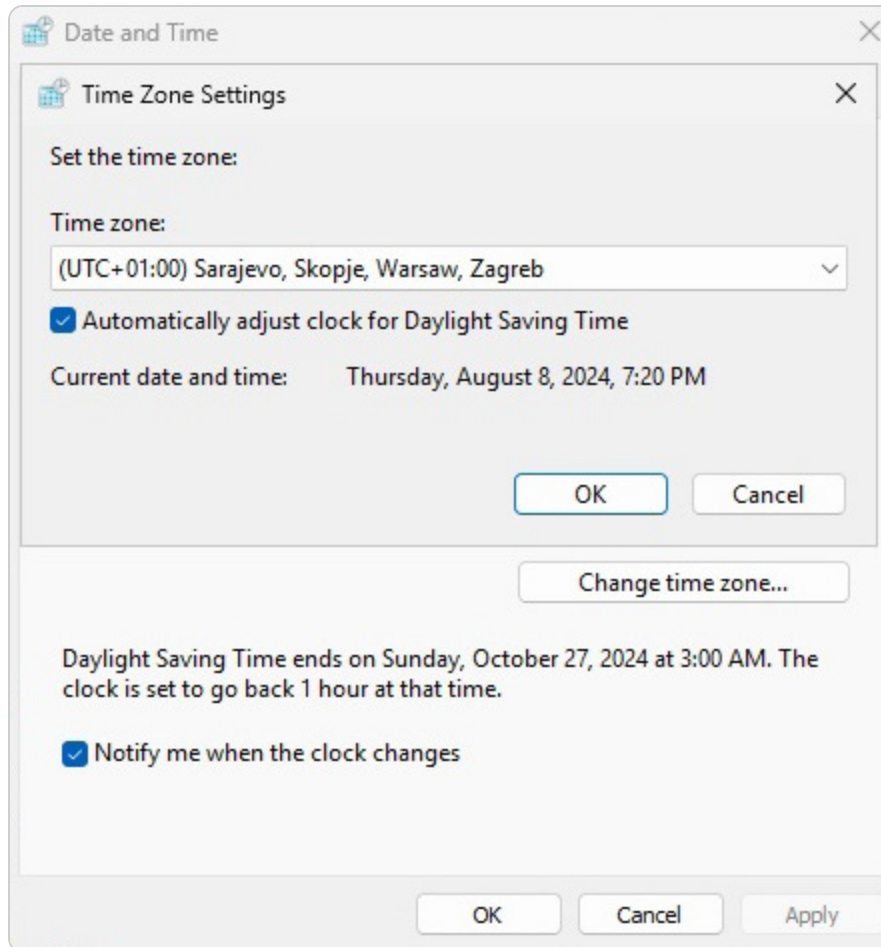


Figure 3.34 – Adjusting the time zone

Once you've adjusted the time zone, the next section will guide you through the process of activating Windows Server, a necessary step for all Windows operating systems, including Windows Server 2025.

Windows Server activation

To fully utilize the features of your Windows Server 2025 (Desktop Experience), activation is required. Follow these steps to complete the activation process, as depicted in *Figure 3.35*:

1. Click on the **Not activated** link under the product ID highlighted in the **Properties** section.

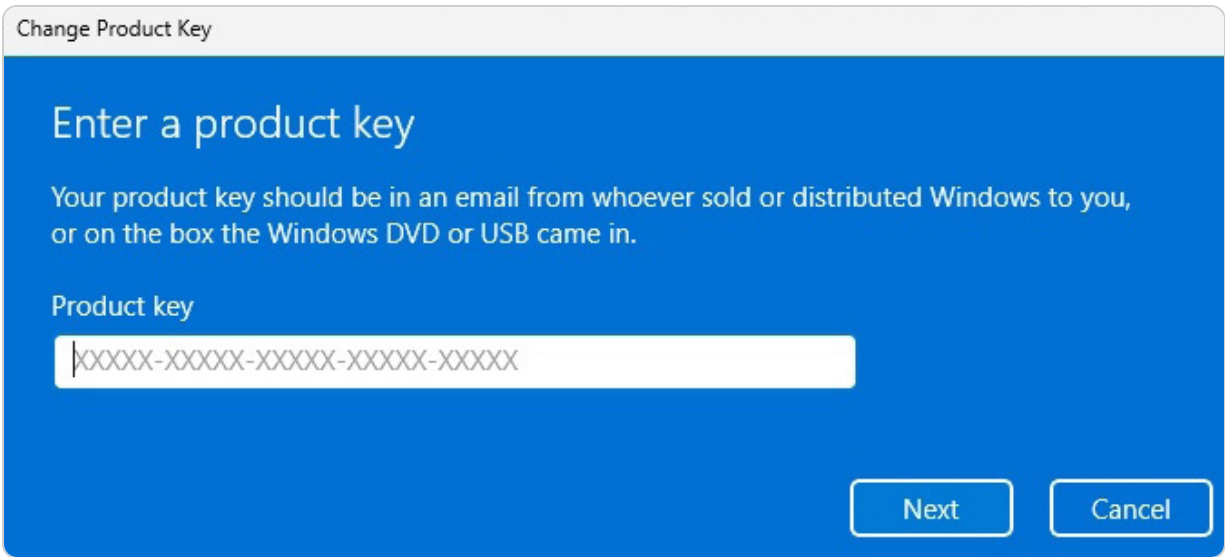


Figure 3.35 – Activating Windows Server 2025

2. In the window that appears, enter your **Windows Server 2025 product key** and click **Next**.
3. Microsoft's **Activation Server** will validate your product key. If the key is valid, click **Next** in the **Activate Windows** window.
4. Once activation is successful, click **Close** to exit the confirmation window.

This section has covered how to configure Windows Server using Server Manager initially. The following section will guide you through performing this configuration using Server Configuration.

How to use Server Configuration for Windows Server initial setup

This section will walk you through the process of configuring Windows Server 2025 Standard (Server Core) using the **Server Configuration tool**. You will learn the essential steps for performing the initial setup of the server in this command-line interface environment.

Renaming the server

You can follow these steps to rename the server, as shown in *Figure 3.36*:

1. At the **Server Configuration** menu prompt, type **2** as the option you want and press *Enter*.
2. Type the new name for the server and press *Enter*.
3. When the **Restart** dialog box appears, click **Yes** to reboot the server.


```
=====
                        Computer name
=====

Current computer name: WIN-49BBC4NBN0T

Enter new computer name (Blank=Cancel): Dautti-Server2
Changing computer name...
WARNING: The changes will take effect after you restart the computer
WIN-49BBC4NBN0T.
Restart now? (Y)es or (N)o: Y_
```

Figure 3.36 – Renaming the server

4. The server will restart to apply the new name.

In the next section, we will show you how to add the server to a domain.

Connecting to the domain

To connect your server to the domain, ensure you have completed the initial configuration using Server Manager, as outlined earlier. Then, follow these steps to establish the domain connection, as depicted in *Figure 3.37*:

1. At the **Server Configuration** prompt, select option **1** and press *Enter*.
2. To link the server to your organization's domain, type **D** and press *Enter*.
3. Input your organization's domain name and press *Enter*.
4. Provide the credentials of an authorized domain user and press *Enter*.
5. Enter the associated password and press *Enter*.
6. When prompted by the **Change Computer Name** dialog box, choose **No** if you do not wish to alter the server's name.

```
=====
                        Change domain/workgroup membership
=====

Current workgroup: WORKGROUP

Join (D)omain or (W)orkgroup? (Blank=Cancel): Dautti.local_
```

Figure 3.37 – Connecting to the domain

Next, we will guide you through enabling Remote Desktop, which is essential for remote access to the Windows Server 2025 Server Core edition.

Configuring Remote Desktop

To configure Remote Desktop, follow these steps, as outlined in *Figure 3.38*:

1. At the **Server Configuration** menu prompt, enter **7** and press *Enter* to select the **Remote Desktop configuration** option.
2. Type **E** and press *Enter* to enable Remote Desktop.
3. Choose option **1** and press *Enter* to select a more **secure access mode**.
4. In the **Remote Desktop** dialog box that appears, click **OK** to confirm that Remote Desktop has been successfully activated.

Next, we will guide you through the process of setting up the IP address, which is essential for the Windows Server 2025 Server Core edition.

```
=====
                        Remote desktop
=====

Remote desktop status: Disabled

(E)nable or (D)isable Remote Desktop? (Blank=Cancel): e

    1) Allow only clients running remote desktop with network level authentication
    2) Allow clients running any version of remote desktop (less secure)
more secure)

Enter selection (Blank=Cancel): 1
Enabling remote desktop...
Successfully configured remote desktop.
(Press ENTER to continue):
```

Figure 3.38 – Enabling Remote Desktop

Next, let us learn how to set up the IP address because even Windows Server 2025 Server Core edition requires a static IP address.

Setting up the IP address

To configure the IP address, as depicted in *Figure 3.39*, follow these instructions:

1. From the **Server Configuration** menu, choose option **8** and press *Enter*.
2. Enter the number associated with the network adapter you want to configure and press *Enter*.
3. In the next sub-menu, select option **1** to configure the network adapter's address and press *Enter*.
4. Choose **S** to set a static IP address and press *Enter*.
5. Enter the static IP address you wish to assign and press *Enter*.
6. Provide the subnet mask and press *Enter*.
7. Enter the default gateway and press *Enter*.
8. In the sub-menu, select option **2** to configure the DNS servers and press *Enter*.
9. Input the primary DNS server address and press *Enter*.
10. Confirm the settings by clicking **OK** in the **Network Settings** dialog box.
11. Enter the secondary DNS server address and press *Enter*.

12. Finally, choose option **4** to exit the sub-menu and return to the main menu.

```
=====
                        Network adapter settings
=====

NIC index:      1
Description:    Microsoft Hyper-V Network Adapter
IP address:     169.254.1.220,
                fe80::65a5:f7e3:713c:1dc
Subnet mask:    255.255.0.0
DHCP enabled:   True

Default gateway:      192.168.1.1
Preferred DNS server:
Alternate DNS server:

1) Set network adapter address
2) Set DNS servers
3) Clear DNS server settings

Enter selection (Blank=Cancel): 1
Select (D)HCP or (S)tatic IP address (Blank=Cancel): S
Enter static IP address (Blank=Cancel): 192.168.1.20
Enter subnet mask (Blank=255.255.255.0): 255.255.255.0
Enter default gateway (Blank=Cancel): 192.168.1.1
Setting NIC to static IP...
Failed to release DHCP lease.
Result code: 83
Method name: ReleaseDHCPLease
(Press ENTER to continue):
```

Figure 3.39 – Setting up the IP address

Next, let us learn how to check for updates to the Windows Server 2025 Server Core edition.

Updating Windows Server 2025 Server Core edition

To check for updates and update your Windows Server 2025 Server Core edition, follow these steps, as illustrated in *Figure 3.40*:

1. At the **Server Configuration** menu prompt, enter **5** and press *Enter* to access the update options.
2. In the **Update options** window, type **A** to install all updates or **R** to install only recommended updates, then press *Enter*.
3. **Windows Update** will proceed to search for available updates.
4. If updates are found, you can choose to install all updates by typing **A**, none by typing **N**, or select individual updates by typing **S**, then press *Enter*.
5. The selected updates will be downloaded and installed. If prompted, click **Yes** to restart the server and apply the updates.

```
=====
                          Update setting
=====

Current update configuration is: Download only

Set updates to:

    1) Automatic
    2) Download only
    3) Manual

or

    5) Opt-in to Microsoft Update

Select update configuration (Blank=Cancel): |
```

Figure 3.40 – Updating Windows Server 2025 Server Core edition

Next, we will guide you through the process of changing the time zone on your server.

Adjusting the date and time

To adjust the date and time, follow the steps shown in *Figure 3.41*:

1. Type **9** and press *Enter* at the **Server Configuration** menu prompt to choose the option.
2. In the **Date and Time** window, click the **Change date and time** button.
3. Set up the date and time by clicking on the **Date** or **Time** section.
4. Click **OK** to exit the **Date and Time** window.
5. Click **OK** again to confirm the changes.

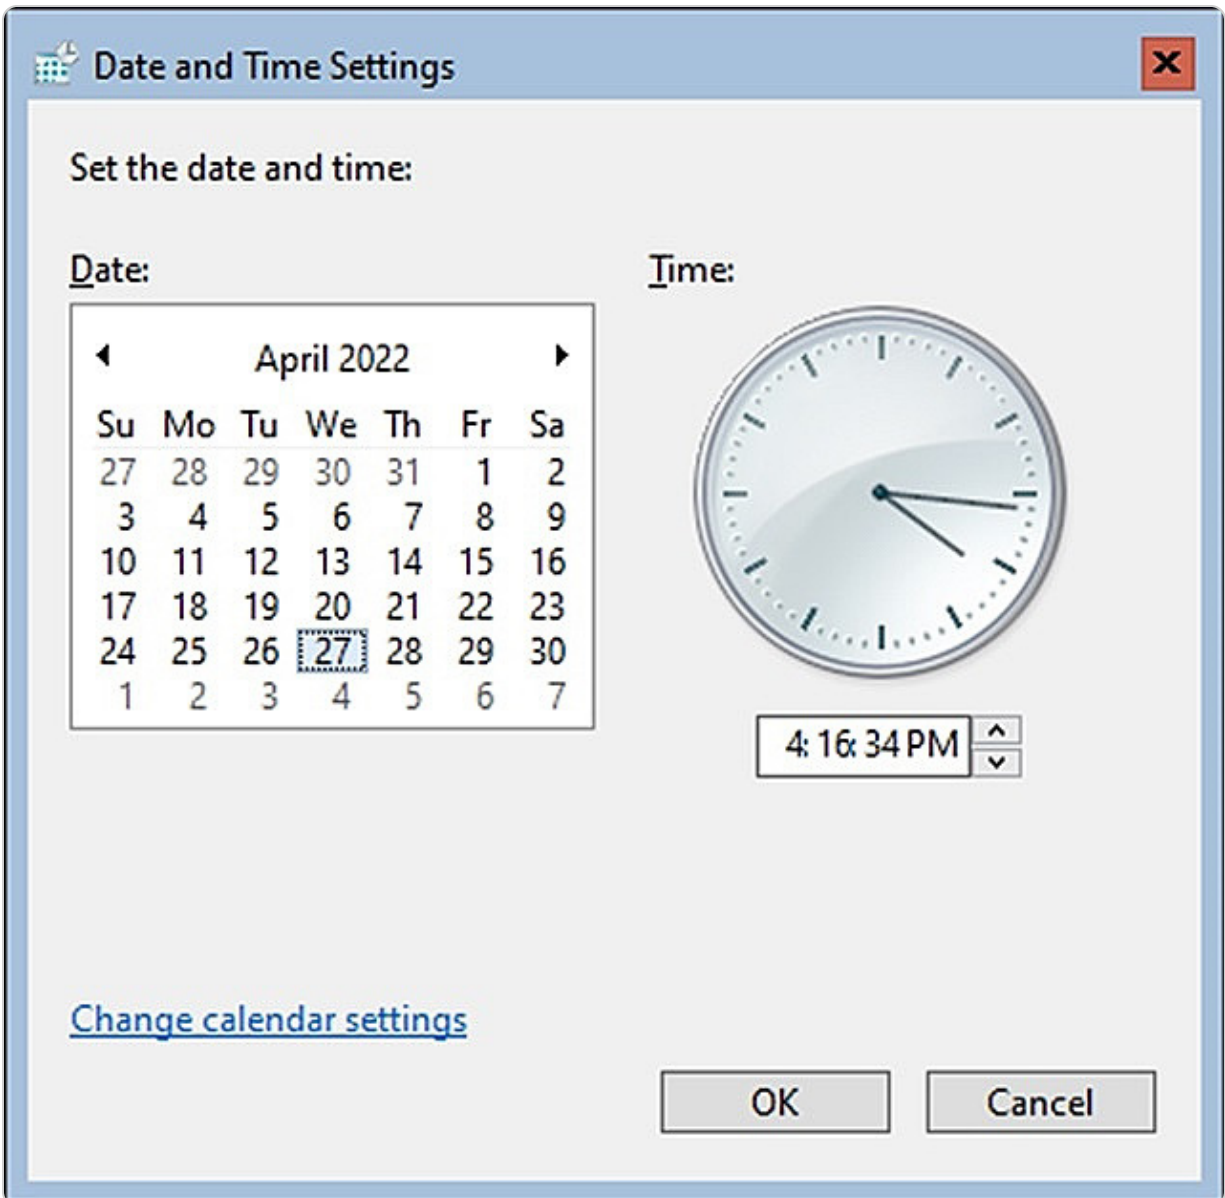


Figure 3.41 – Adjusting the date and time

Next, we will explain how to activate the Windows Server 2025 Server Core edition.

How to activate Windows Server

You need to activate Windows Server 2025 Server Core after installing it. You can do this by following the steps shown in *Figure 3.42*:

1. Type **11** and press *Enter* at the **Server Configuration** menu prompt to select the option.
2. Type **3** and press *Enter* in the sub-menu to enter the product key for Windows Server 2025.
3. In the **Enter Product Key** window, type the product key and click **OK**.
4. Type **2** and press *Enter* in the sub-menu to start the activation process.

5. Wait for a few moments until Windows Server 2025 is activated—type **Exit** to close the activation window.
6. Type **4** and press *Enter* in the sub-menu to go back to the main menu.

```
=====
                        Windows activation
=====

1) Display license information
2) Activate Windows
3) Install product key

Enter selection (Blank=Cancel): |
```

Figure 3.42 – How to activate Windows Server 2025?

This chapter exercise has taught you how to perform the initial configuration of Windows Server 2025 using Server Manager and Server Configuration.

Summary

This chapter provided a comprehensive overview of post-installation tasks in Windows Server 2025, covering key aspects such as device drivers, registry management, and service configurations. You explored how devices are organized within a computer system and the crucial role that device drivers play. Additionally, you gained insights into managing system resources, including how to handle the Windows Server registry and services effectively. Topics included the various registry keys, service startup types, and practical tasks such as installing, uninstalling, and updating drivers, as well as starting, stopping, and restarting Windows services. You also learned how to use the `regedit` tool for basic operations such as adding, modifying, and deleting registry keys. Furthermore, you were introduced to the concept of Windows Server's initial configuration, which facilitated the completion of post-installation tasks covered in this chapter. The upcoming chapter will delve into **Directory Services** within Windows Server 2025.

Questions

1. **True or false?** A device driver is a program that acts as the translator between computer hardware and an operating system.
2. **Fill in the blank:** _____ works on the principle that when a device is plugged into a computer, the device is immediately recognized by the operating system.
3. Which of the following are known as computer system resources? **Choose two:**
 - A. IRQ
 - B. DMA
 - C. SAN

D. NAS

4. **True or false?** A driver's digital signature identifies its publisher.
5. **Fill in the blank:** _____ is a hierarchical database that stores hardware and software configurations and system security information.
6. Which Windows Server tools are used to operate devices and device drivers? **Choose two:**
- A. Devices
 - B. Device Manager
 - C. Registry Editor
 - D. Control Manager
7. Which Windows Server tools are used to operate the registry and services? **Choose two:**
- A. Services Control Manager
 - B. Registry Editor
 - C. Device Manager
 - D. Devices
8. **Fill in the blank:** _____ is the Windows Server native account or an account you created to manage running services.
9. Discuss Windows registry keys.
10. Discuss Windows service startup types.

Further reading

- *How to Use the Windows Device Manager for Troubleshooting:* <https://www.howtogeek.com/167094/how-to-use-the-windows-device-manager-for-troubleshooting>
- *Structure of the Registry:* <https://learn.microsoft.com/en-us/windows/win32/sysinfo/structure-of-the-registry>
- *21 Windows Tools Explained:* <https://www.howtogeek.com/193922/21-windows-administrative-tools-explained/>
- *Introduction to Windows Service Applications:* <https://learn.microsoft.com/en-us/dotnet/framework/windows-services/introduction-to-windows-service-applications>
- *Using Device Manager:* <https://learn.microsoft.com/en-us/windows-hardware/drivers/install/using-device-manager>

Part 2: Setting Up Windows Server 2025

This part will explain how to configure essential roles and services in Windows Server 2025. Upon completion, you will be able to establish a domain environment and deploy critical network services, such as DNS, DHCP, Print Server, Web Server, WDS, and WSUS.

This part contains the following chapters:

- [Chapter 4](#), *Directory Services in Windows Server 2025*
- [Chapter 5](#), *Adding Roles to Windows Server 2025*

4

Directory Services in Windows Server 2025

In this chapter, you will progress in establishing your organization's IT infrastructure by thoroughly exploring domain services, which are pivotal for managing a Windows-based domain network. You will gain insight into **Active Directory Domain Services (AD DS)** and **Domain Name System (DNS)**, understanding their integral roles in network management. Key concepts such as **domains**, **forests**, **tree domains**, **child domains**, and **Domain Controllers (DCs)** will be covered, along with functional levels and trust relationships that facilitate network integration and resource sharing. Additionally, the chapter will delve into DNS functionalities, including forward and reverse lookup zones and DNS records, which are crucial for resolving domain names and ensuring reliable network communication.

Furthermore, you will explore how **Organizational Units (OUs)**, **default containers**, **user accounts**, and different group scopes and types are utilized to effectively manage user and **computer accounts** within a domain-based network. By understanding these components, you'll be able to streamline account management and enhance organizational structure within your network. The chapter culminates with a hands-on exercise where you will install the AD DS and DNS roles and promote a server to a DC. This practical experience will equip you with the skills needed to implement and manage a Windows Server domain, setting the stage for more advanced network configurations and administration.

In this chapter, we're going to cover the following main topics:

- Understanding the **Active Directory (AD)** infrastructure in Windows Server 2025
- Adding and configuring the AD DS role
- Exploring DNS fundamentals and configurations in Windows Server 2025
- Managing OUs and default containers
- User and group management within AD
- Installing the AD DS and DNS roles and promoting the server to a DC

Technical requirements

To complete the exercises in this chapter, you will require the following hardware:

- A **Windows 11 Pro PC** equipped with a minimum of 16 GB of RAM, a 1 TB **hard disk drive (HDD)**, and a stable internet connection
- A **Windows Server 2025 Standard** (Desktop Experience) virtual machine, designated as Virtual Machine 1, configured with a tree domain (**Dautti.local**), featuring at least 4 GB of RAM, 100 GB of HDD space, and internet access
- Another Windows Server 2025 Standard (Desktop Experience) virtual machine, designated as Virtual Machine 2, set up with a tree domain (**ITTrainings.local**), also with a minimum of 4 GB of RAM, 100 GB of HDD space, and internet connectivity

- A third Windows Server 2025 Standard (Desktop Experience) virtual machine, designated as Virtual Machine 3, configured with a child domain (**Programming.Dautti.local**), including at least 4 GB of RAM, 100 GB of HDD space, and internet access

This setup ensures that you have the necessary resources and configurations to perform all tasks effectively.

Understanding the AD infrastructure in Windows Server 2025

AD is a foundational technology from Microsoft that serves as a distributed directory service. It's essential for organizing and managing network resources in a hierarchical and secure manner. It acts as a centralized repository where critical objects—such as user accounts, computers, printers, and network services—are stored, each with its own distinct security settings.

The unique attributes of each object within AD enable granular control over resource management, allowing for precise administration across the network. For instance, each object, whether a user account, computer, printer, or network service, possesses specific attributes, including **Security Identifiers (SIDs)**, group memberships, and Access Control Lists (ACLs). These attributes empower administrators to define individual permissions, roles, and access policies, ensuring that security measures and functionalities are tailored to the requirements of each object.

The architecture of AD, as illustrated in *Figure 4.1*, is structured around three fundamental tiers:

- **Domain:** The basic unit of administration, providing a boundary for policies and security settings
- **Tree:** A collection of domains linked by a contiguous namespace, reflecting a hierarchical relationship among them
- **Forest:** The highest level of organization, which can encompass multiple trees and serves as the topmost layer that integrates the entire directory service

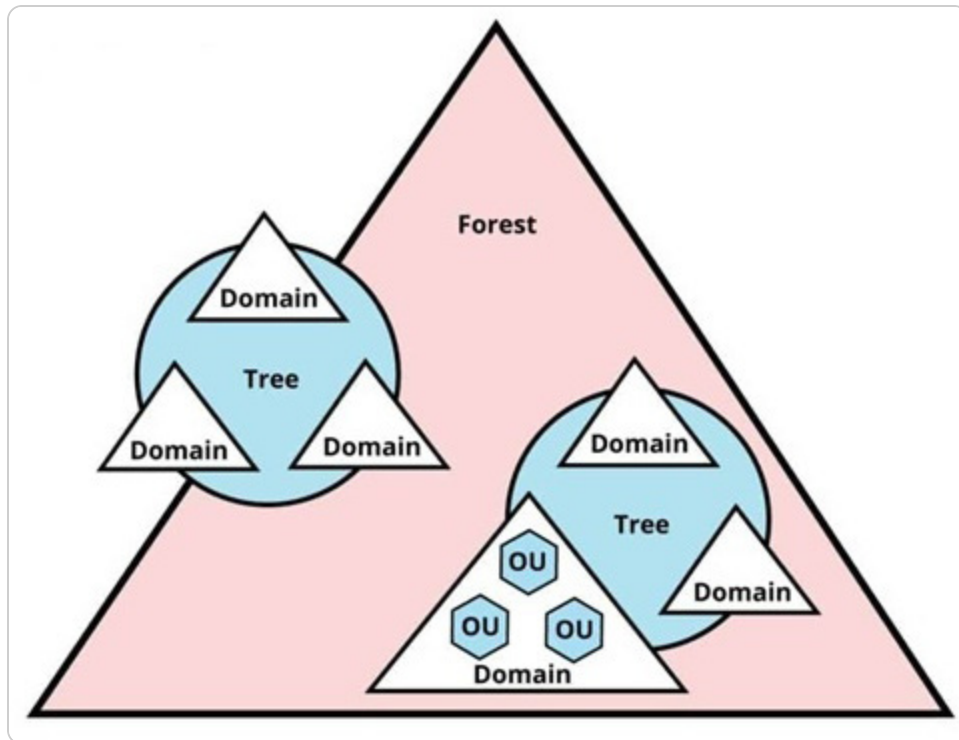


Figure 4.1 – AD architecture (source – Websentra)

This layered approach facilitates efficient resource management and scalability while supporting complex organizational structures. It allows businesses to customize their network infrastructure to align with specific operational needs while maintaining robust security and administrative oversight.

In the sections that follow, we will explore the specific features and configurations of AD, equipping you with the necessary knowledge and skills to effectively implement and manage directory services within your organization.

Addressing the importance of AD

AD is not just a directory service; it is the backbone of modern IT infrastructures, playing a critical role in the management of Windows environments. For those new to IT, understanding the significance of AD is essential for grasping its functionalities and benefits:

- **Centralized management:** One of the primary advantages of AD is its ability to centralize management. IT administrators can manage users, computers, and resources from a single location, significantly reducing complexity and administrative overhead. This centralized approach streamlines the process of user provisioning and deprovisioning, making it easier to maintain an organized and efficient network.
- **Enhanced security:** AD enhances security through its use of SIDs and ACLs. By ensuring that only authorized users can access specific resources, AD protects sensitive information and reduces the risk of unauthorized access. This security model is crucial for safeguarding organizational data and maintaining compliance with regulatory standards.

- **Scalability:** The hierarchical structure of AD supports the scalability of organizations. As businesses grow, AD allows for the seamless integration of new users and resources without compromising performance or security. This flexibility enables organizations to adapt to changing needs and expand their IT infrastructure effectively.
- **Policy enforcement:** AD facilitates the application of Group Policies across the network, allowing organizations to enforce security settings and compliance standards uniformly. This capability ensures that all users and devices adhere to the organization's policies, enhancing the overall security posture and operational efficiency.

By understanding these core principles of AD, one can appreciate its pivotal role in managing and securing network resources, laying the foundation for effective IT administration.

Core protocols and services supporting AD

AD relies on several critical protocols and services to ensure its seamless operation, each contributing to different aspects of network management and security:

- **Lightweight Directory Access Protocol (LDAP):** This is a foundational protocol that plays a crucial role in enabling users and applications to query and interact with the directory's data. LDAP provides a standardized method for accessing and managing the information stored within AD, making it a key element in directory service operations.
- **Kerberos:** This is a sophisticated authentication mechanism that underpins AD's security framework. Kerberos uses a ticketing system to securely verify the identities of users and servers on the network, preventing unauthorized access and ensuring that all entities communicating within the network are properly authenticated. This protocol is vital for maintaining the integrity and confidentiality of the network environment.
- **DNS:** This protocol is also integral to AD's functionality. DNS serves as a directory for the internet and internal networks, translating user-friendly domain names into numerical IP addresses. This translation process is essential for locating and accessing network resources efficiently. Within an AD environment, DNS not only resolves domain names but also supports AD-specific functions, such as locating DCs and ensuring that services are reachable across the network.

Together, these protocols and services form the backbone of AD, enabling it to deliver a secure, scalable, and efficient directory service that supports complex organizational needs.

Tools and roles for administering AD

AD is a robust framework that provides comprehensive services to enable centralized management of network resources, streamlining the work of system administrators in complex IT environments. To effectively manage various aspects of AD services, Microsoft offers a suite of administrative consoles within the **Microsoft Management Console (MMC)** (`mmc.exe`), each tailored to specific tasks within the directory service:

- **Active Directory Administrative Center (dsac.exe):** A key tool depicted in *Figure 4.2*, this is instrumental in managing Windows Server's directory services. This modern interface integrates several management functions, enabling administrators to oversee these services efficiently. It includes the **Active Directory Users and Computers (dsa.msc)** snap-in, which is essential for managing user accounts, computer objects, OUs, and their associated properties. This tool is foundational for daily administrative tasks, such as creating and managing users, groups, and devices and organizing them within the AD structure.

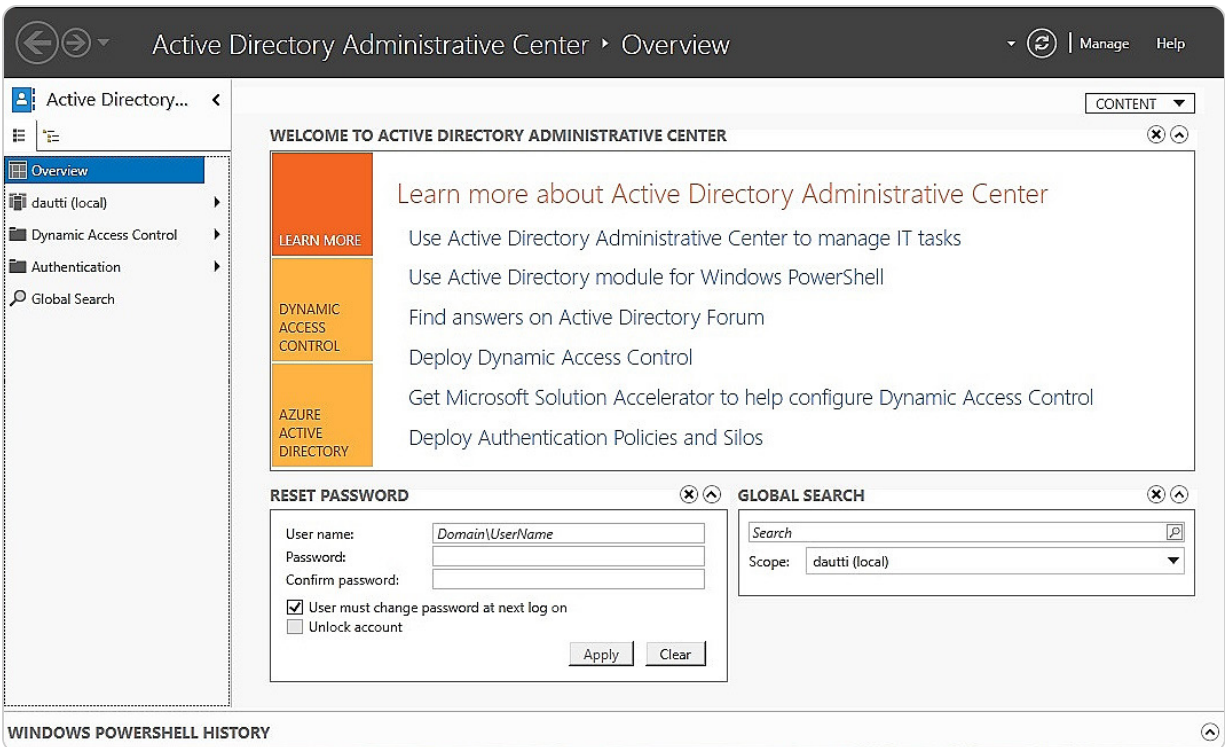


Figure 4.2 – The Active Directory Administrative Center in Windows Server 2025

- **Active Directory Domains and Trusts console (`domain.msc`):** This console is employed for tasks related to domain management. This tool allows administrators to configure and manage domain trusts, which are crucial for enabling secure communications and resource sharing between different domains within the same or different forests. It also handles the setup and management of **Domain Functional Levels (DFLs)**, which determine the features available within a domain based on the version of Windows Server being used.
- **Active Directory Sites and Services console (`dsstite.msc`):** This is a critical tool in managing replication between different AD sites. Sites in AD represent the physical structure of a network, and this tool allows administrators to optimize and control how directory information is replicated across various geographic locations, ensuring consistency and availability of data across the entire organization. It also manages services such as global catalog servers and ensures that authentication requests are routed efficiently.
- **AD module for Windows PowerShell:** Not a graphical tool, this offers a command-line interface for more advanced and automated management tasks. PowerShell cmdlets allow administrators to script complex operations, automate repetitive tasks, and manage AD objects at scale, making them invaluable for large or highly customized environments.

To deploy directory services in an organization, the AD DS role must be installed and configured on a Windows Server. AD DS is the backbone of the AD environment, enabling the storage, organization, and management of information about network resources such as users, groups, computers, and policies. It also supports advanced security features, such as centralized authentication and authorization, which are vital for maintaining the integrity and security of the network.

NOTE

You can access a wealth of free PowerShell scripts at Microsoft's Script Center (<https://technet.microsoft.com/en-us/scriptcenter/bb410849.aspx>) and PowerShell Gallery (<https://www.powershellgallery.com/>). These platforms serve as renowned repositories where IT professionals can find and share scripts for various administrative tasks. Both resources include extensive collections of scripts specifically related to AD and DNS, making them invaluable for automating and simplifying complex network management activities.

For a detailed walkthrough on setting up AD DS, including the installation of the DNS role and the promotion of a server to a DC, see [Chapter 5, Adding Roles to Windows Server 2025](#). This chapter includes practical exercises to guide you through the process, ensuring a solid understanding of the necessary steps to integrate AD DS into your infrastructure.

As we move forward, the following section will delve into the critical components of an AD infrastructure, starting with an in-depth look at DCs, which are the cornerstone of any AD environment.

Adding and configuring the AD DS role

In Windows Server environments, the role of the AD DS is crucial for providing centralized directory services, which facilitate network management and authentication. This section delves into the process of adding and configuring the AD DS role, covering key tasks such as deploying DCs, setting up and managing domains, and creating hierarchical structures such as tree and child domains. Additionally, we will explore the concept of namespaces to streamline directory organization and examine how sites enhance network **performance** and replication efficiency. By mastering these aspects, you will be equipped to implement and manage a robust and scalable AD DS infrastructure, optimizing both functionality and performance for your organization.

Understanding DCs

A DC, as depicted in *Figure 4.3*, is a server that plays a critical role in managing and verifying user identities within an organization's network. Its primary function is to authenticate users and authorize access to network resources based on the security policies defined within the domain. In earlier Windows environments, specifically Windows NT, domain management relied on a **Primary Domain Controller (PDC)** to handle main domain functions, with **Backup Domain Controllers (BDCs)** providing redundancy. However, this model was replaced with the multi-master replication model introduced in Windows 2000, allowing multiple DCs to share the responsibility of managing domain functions. This approach enhances reliability and availability, as all DCs can perform read and write operations, ensuring that authentication and directory services remain robust and accessible across the network.

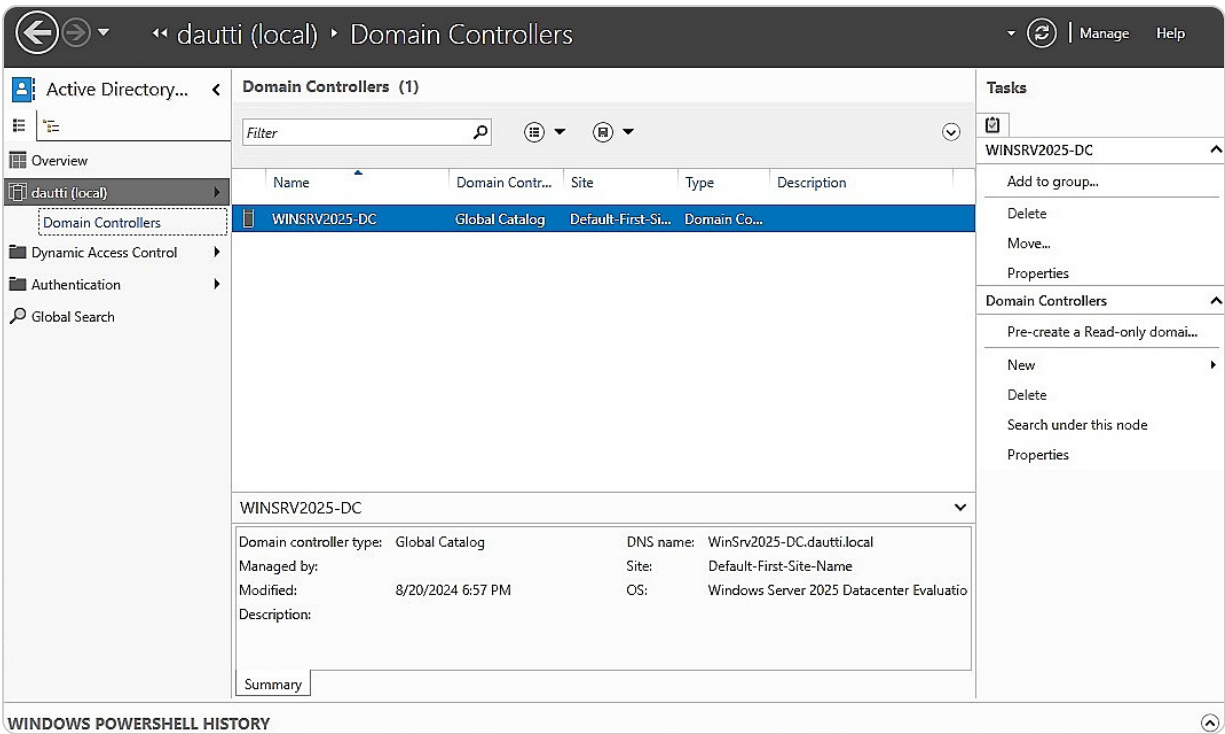


Figure 4.3 – Accessing DCs through the Active Directory Administrative Center

Windows Server 2025 has revolutionized the approach to DCs by eliminating the traditional primary and backup roles. Instead, DCs are now identified by sequential numbers, such as **DC1** and **DC2**, which denote their sequence rather than their function. This modern approach paves the way for a more flexible and scalable domain management environment, where all DCs are considered equal partners in the domain, sharing the responsibility for authentication and directory services. This evolution keeps you, as an IT professional, informed and up to date in your field.

NOTE

When a server joins a domain but does not take on the role of a DC, it is classified as a **member server**. Member servers operate under the domain's policies and access control but do not handle authentication requests or domain management tasks.

Given that DCs are central to domain access and authentication, understanding the concept of domains is not just important. It's essential to grasp the full scope of AD infrastructure. As an IT professional, your role is significant in understanding the intricacies of domain structures and their role within the network. Therefore, we will next explore these intricacies in detail.

Understanding domains

Domains are fundamental components in network management that organize and group users, computers, devices, and network services under a unified administrative framework. This logical

grouping allows for centralized management of resources and security policies. A DC is crucial in this setup, with the AD DS playing a pivotal part in establishing and maintaining domain functionality.

Figure 4.4 demonstrates the domain configuration process within the **Active Directory Domain Services Configuration Wizard** window, showcasing how domains are created and managed.

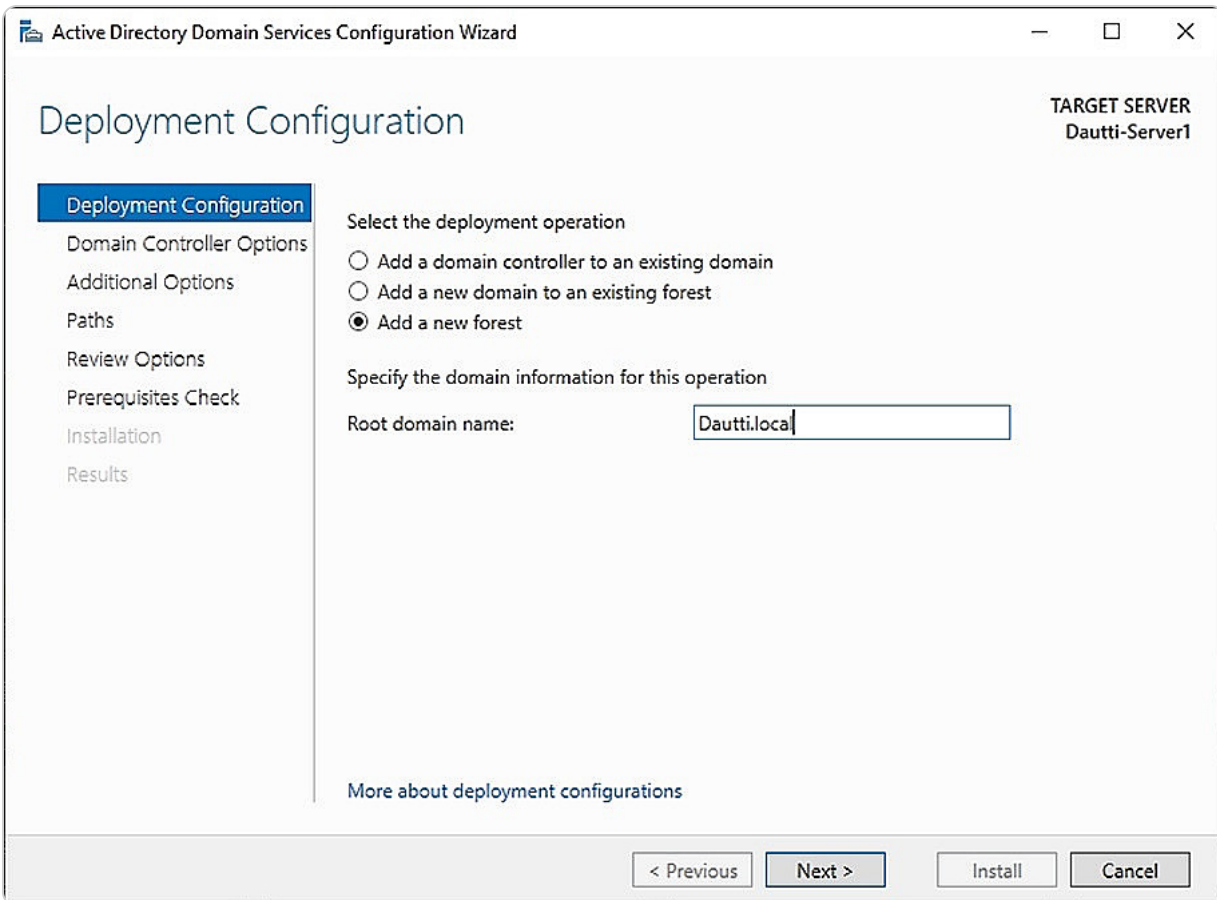


Figure 4.4 – Setting up a root domain in Windows Server 2025

NOTE

It is crucial to differentiate between a directory domain and a domain name. In the context of directory services, a domain refers to a structured database of network resources, including users, servers, and devices, that are managed collectively under specific administrative policies. This domain facilitates effective management and security within an organization's IT infrastructure. On the other hand, a domain name is part of the DNS, which is a hierarchical naming system used to identify and locate resources on the internet, such as websites and email servers.

Furthermore, domains can be organized into a **domain tree**, which represents a hierarchical structure of multiple domains. This structure allows for the organization of domains into a parent-child relationship, where each domain within the tree can inherit policies and settings from its parent domain while maintaining its distinct configuration. The following section will explore the concept of

domain trees in detail, explaining how they extend the domain structure and facilitate more complex organizational setups.

Understanding the Domain Tree

To fully understand the AD architecture, it's important to delve into the concept of a domain tree. A domain tree represents a logical structure within AD, consisting of one or more domains that share a common namespace and are arranged hierarchically. This hierarchical setup not only organizes the domains but also ensures that they inherently trust one another due to the transitive trust relationship. In AD, a trust relationship (as shown in *Figure 4.5*) allows users in one domain to authenticate and access resources in another domain without needing separate credentials. Transitive trust means that if domain A trusts domain B and domain B trusts domain C, then domain A will automatically trust domain C. This built-in trust simplifies resource sharing and authentication across domains within the same tree domain. When introducing a new domain into an existing tree, it is necessary to provide the parent domain's name during the server promotion process to establish the appropriate hierarchy.


 Figure 4.5 – Hierarchical architecture of Domain Forest (source – Websentra)

Figure 4.5 – Hierarchical architecture of Domain Forest (source – Websentra)

The process involves specifying the parent domain during the server promotion phase to integrate a new domain into an existing domain tree. This addition allows the new domain to inherit policies and settings from its parent while establishing its own unique identity within the tree. *Figure 4.6* provides a visual representation of creating a child domain in Windows Server 2025. It illustrates how a new domain is incorporated into the existing tree structure and how it aligns with the overall namespace.

 Figure 4.6 – Setting up a tree domain in Windows Server 2025

Figure 4.6 – Setting up a tree domain in Windows Server 2025

The domain tree concept becomes more expansive when multiple domain trees are combined to form a forest. A forest represents a broader organizational structure that groups all the domain trees within an enterprise, allowing for a unified directory environment. The forest serves as the highest level of the AD hierarchy, providing a framework for managing multiple trees and their interconnected domains. The detailed structure and functions of a forest will be discussed in the subsequent section.

Understanding the Forest

In AD, the concept of a **forest** is analogous to a natural forest, which is composed of multiple trees. An AD forest can consist of a single domain tree or a collection of interconnected domain trees. Each domain tree within a forest shares a common schema and global catalog, but the trees themselves do not have to share the same namespace. A root domain is the first domain created within a domain tree

and serves as the foundation for the entire domain structure. It often holds critical roles such as the schema master and domain naming master. A domain tree that operates as a root domain can exist independently within a forest, but when multiple trees are present, the forest acts as an overarching framework that integrates and manages these trees, creating a cohesive and scalable directory environment.

The forest acts as the highest level of AD structure, providing a unified directory system that enables resource sharing and administrative management across all domains within it. Although this might initially seem like circular logic, the concept of a forest as both a domain and an encompassing structure reflects its dual role in organizing and linking various domains.

To create and configure a forest in Windows Server 2025, you use the AD DS configuration wizard, which is also employed for setting up tree domains. This wizard facilitates the forest creation process, guiding you through the necessary steps and configurations, as depicted in *Figure 5.3*.

Within the framework of a domain tree, additional subdomains, referred to as child domains, can be established. These child domains function as subdivisions of the parent tree domain, allowing for a more granular organization and management of resources. The upcoming section will explore the role of child domains in detail, including their setup and functionality, as well as how they contribute to the overall structure of AD.

Understanding the Child Domain

A **child domain** is a subordinate domain within a tree domain structure in AD. For example, as shown in *Figure 4.6*, there are two tree domains: `Dautti.local` and `Training.local`. `Dautti.local` serves as the root domain of the forest, establishing the foundational namespace. In this setup, `Training.local` includes a child domain called `Administration.Dautti.local`. This child domain is an extension of the parent domain's namespace, ensuring a cohesive and organized directory hierarchy.

Creating a child domain in Windows Server 2025 is accomplished using the AD DS configuration wizard, which is also used for setting up other types of domains. The wizard guides you through the necessary steps, as illustrated in *Figure 4.7*, making the process straightforward and consistent.

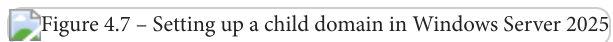


Figure 4.7 – Setting up a child domain in Windows Server 2025

The structure of tree and child domains resembles a tree data structure, where each domain functions as a node with a parent-child relationship. This hierarchical arrangement facilitates efficient resource management and delegation of administrative tasks. The parent domain oversees the child domains, while child domains inherit certain attributes and policies from their parent yet retain their own unique identity.

Understanding this hierarchical organization is fundamental for managing an AD environment effectively. The next topic will delve into the operations master roles, which are critical for maintaining the stability and functionality of the AD infrastructure.

Understanding Operations Master Roles

AD DS is a robust and intricate system that necessitates thorough planning and execution to optimize its capabilities. Once deployed, its operational benefits become increasingly evident. A key component of AD DS is the **operations master roles**, which are essential for maintaining and managing the directory services effectively.

In our previous discussion, we created the root domain, `Dautti.local`, which serves as the primary domain for the forest. This domain is hosted on a server that also functions as a DC, overseeing the network's directory services.

When the AD DS role is installed and the server is promoted to a DC, AD DS automatically assigns five critical operations master roles. These roles are split into the following two categories:

- **Forest-wide roles:** The **schema master** and **domain naming master** roles are forest-wide. The schema master oversees the directory schema, which dictates the attributes and classes of objects within the directory, ensuring consistency across the entire forest. The domain naming master, on the other hand, manages the namespace and guarantees that each domain name within the forest is unique, preventing naming conflicts.
- **Domain-wide roles:** The remaining three roles—**RID master**, **PDC emulator**, and **infrastructure master**—are domain-wide. The RID master allocates SIDs to DCs to create new security principals. The PDC emulator handles password changes and manages time synchronization within the domain, serving as a bridge for backward compatibility with older systems. The infrastructure master is responsible for maintaining and updating references to objects in other domains, ensuring that cross-domain object references remain accurate and up-to-date.

In the example provided, illustrated in *Figure 4.8*, the `Dautti.local` root domain holds the schema master and domain naming master roles for the entire Dautti forest. Each tree domain within this forest, such as `Dautti.local` and `Training.local`, possesses its own RID master, PDC emulator, and infrastructure master roles to manage domain-specific tasks and operations.



Figure 4.8 – AD DS structure

These five operations master roles are known as **Flexible Single Master Operations (FSMO)** roles. The term *flexible* reflects their ability to be transferred to other DCs if necessary. At the same time, *single* indicates that only one DC can hold each role at any given time to prevent conflicts and ensure consistency. Understanding these roles and their functions is crucial for managing and troubleshooting AD environments. The following section will delve into the differences between domains and workgroups, further expanding on network management concepts.

Understanding the difference between domains and workgroups

To effectively differentiate between a domain and a workgroup, it is important to understand the underlying network architectures they represent: **peer-to-peer (P2P) networking** and **client/server networking**, which were briefly mentioned in [Chapter 1, Network Fundamentals and Introduction to Windows Server 2025](#).

In a P2P network, commonly known as a **workgroup**, each computer operates independently and manages its resources. This architecture is ideal for smaller networks, such as those found in home or small office environments, where simplicity and direct resource sharing are key. In a workgroup, each device holds its user accounts and permissions, with no centralized management or control. This decentralized approach can be advantageous for straightforward setups but becomes increasingly challenging to manage as the number of computers grows. Without a central administration, there is a heightened risk of inconsistent security policies and user management issues, making workgroups less suitable for larger or more complex environments, particularly where sensitive data is involved.

Conversely, a client/server network, or domain, provides a more structured approach to managing resources and security. In a domain, a central server known as a DC oversees administrative tasks and enforces security policies across the entire network. This centralized management enables efficient user authentication, policy application, and resource allocation, which is crucial for larger organizations and environments such as **Metropolitan Area Networks (MANs)** or **Wide Area Networks (WANs)**. Domains support hierarchical structures and more complex security requirements, allowing for better scalability, control, and consistent policy enforcement across all networked devices. Consequently, domains are more appropriate for large, complex environments where sensitive data is processed, while workgroups may be suitable for smaller entities with less stringent security needs.

The following table, *Table 4.1*, provides a comparative summary of the key differences between domains and workgroups, outlining their respective strengths and limitations based on network scale and management needs. This comparison is instrumental in determining the most suitable network architecture for a given environment.

Domain	Workgroup
A dedicated server is used to provide services	Computers share resources equally without needing a dedicated server

Domain	Workgroup
An example is the client/server network	An example is a P2P network

Table 4.1 – Domain vs. Workgroup

Understanding these distinctions will help in making informed decisions about network design and management. The subsequent section will explore the concept of trust relationships between a computer and a DC, further enhancing our grasp of network security and administration.

Understanding trust relationships

A key concept in AD is the trust relationship, mentioned briefly in the *Understanding the Domain Tree* subsection, which plays a crucial role in how computers, DCs, and domains interact within a networked environment. When a computer is integrated into a domain, it shifts from relying on its local **Security Account Manager (SAM)** for authentication to depending on the DC's authentication system, typically **Kerberos**. This transition is significant because it centralizes user authentication, ensuring that credentials are verified by the DC rather than the local machine. That not only streamlines the authentication process but also enhances security by enforcing consistent policies across the network.

Trust relationships extend beyond individual computers to encompass entire domains within a forest—a logical grouping of multiple tree domains. Within a forest, each domain automatically trusts the authentication methods of the others, establishing a cohesive security framework. For instance, as illustrated in *Figure 5.6*, if the `Dautti.local` domain authenticates a user, this authentication is implicitly recognized and trusted by another domain in the same forest, such as `Training.local`. This trust is rooted in the shared infrastructure of the forest, where domains such as `Dautti.local` serve as the foundation for this interconnected network.

Understanding trust relationships is essential for grasping the broader administrative and communication structures that facilitate secure and efficient operations within an AD environment. These relationships ensure that users can access resources across different domains without the need for redundant authentication processes, thereby promoting seamless collaboration and resource sharing.

In the next section, we will delve into the functional levels of domains and forests, which define the capabilities and compatibility of the AD environment. We will also explore how to check and configure these functional levels to optimize the performance and security of your network.

Understanding functional levels

Functional levels in AD are critical elements that shape the functionality, compatibility, and overall behavior of the AD environment. They help define the specific features that can be utilized and ensure that all DCs within the environment are running compatible versions of Windows Server. There are two primary functional levels to understand, as illustrated in *Figure 4.9*:

- **Forest Functional Level (FFL):** This is pivotal in determining which versions of Windows Server can operate on the DCs across the entire forest, which is the topmost structure within AD that can consist of one or more domains. By setting the FFL, you not only define the minimum server version allowed but also unlock specific forest-wide features that enhance security, replication, and management capabilities across all the domains within that forest. For instance, certain advanced features such as the AD Recycle Bin or fine-grained password policies are only available at higher functional levels.
- **Domain Functional Level (DFL):** This applies to individual domains within the forest. It specifies which Windows Server versions are supported on the DCs within that domain and enables domain-specific features. Raising the DFL can unlock enhancements in domain-specific functionalities, such as improvements in authentication protocols, group policy management, and replication methods. This level of control allows administrators to gradually upgrade parts of the AD infrastructure without immediately affecting the entire forest, providing a flexible approach to modernization.



Figure 4.9 – FFL and DFL in Windows Server 2025

In the context of Windows Server 2025, the minimum FFL and DFL can be set to Windows Server 2016. This requirement ensures that the forest remains compatible with modern features while still allowing for some backward compatibility with older systems. Both the FFL and DFL can be elevated to Windows Server 2025, which enables the most current features and optimizations offered by the latest server technology. It's important to note that once a functional level is raised, it cannot be lowered. So, careful planning is essential to avoid compatibility issues with older systems that may still be in use.

Verifying and managing DFLs and FFLs

To verify and manage the forest and DFLs in Windows Server 2025, you can follow these steps:

1. Click the **Start** button, and from the **Start** menu, select **Server Manager**.
2. In the **Server Manager** window, click on **Tools** in the menu bar and choose **Active Directory Domains and Trusts**.
3. In the **Active Directory Domains and Trusts** window, right-click the root domain and select **Properties** from the context menu.
4. Within the **Properties** dialog box, under the **General** tab, the current **Domain functional level** and **Forest functional level** are displayed, as shown in *Figure 4.10*.



Figure 4.10 – Verifying the FFL and DFL

Understanding and managing FFL and DFL is essential for ensuring that your AD environment operates smoothly and securely, with all the necessary features available to support your

organizational needs. These levels also play a crucial role in the overall architecture and strategy of your AD deployment, allowing for both flexibility and control as you scale and adapt your infrastructure.

With a solid grasp of how functional levels impact your AD environment, you can explore the concept of a contiguous namespace. This concept is integral to maintaining a logical and seamless connection between child domains and their parent domains within the same tree structure, ensuring efficient and organized management of your AD hierarchy.

Exploring the concept of namespaces

In AD DS, the concept of a namespace is fundamental to the organization and management of domains and forests within a network. A namespace serves as a logical identifier that uniquely names a domain or a forest, providing structure and clarity to the AD DS environment. For example, in *Figure 4.9*, the `Training.local` domain functions as both the root domain and the overarching forest. Within this forest, `Training.local` and `ITTrainings.local` are distinct domain trees, each representing a separate branch within the same hierarchical structure. Additionally, within the `Training.local` tree domain, a child domain named `Programming` exists. The `Training.local` component, which is shared across these domains, signifies a contiguous namespace, meaning that all domains within this forest are connected through a common naming convention.

In AD DS, the concept of a namespace is fundamental to the organization and management of domains and forests within a network. A namespace serves as a logical identifier that uniquely names a domain or a forest, providing structure and clarity to the AD DS environment. For example, in *Figure 4.11*, the `Dautti.local` domain functions as both the root domain and the overarching forest. Within this forest, `ITTrainings.local` and `Administration.local` are distinct domain trees, each representing a separate branch within the same hierarchical structure. Additionally, within the `Dautti.local` tree domain, a child domain named `Programming.Dautti.local` exists. The `Dautti.local` component is shared across these domains and signifies a contiguous namespace, meaning that all domains within this forest are connected through a common naming convention.



Figure 4.11 – The namespace concept in AD DS

A contiguous namespace is crucial for maintaining a consistent and organized AD DS structure. It ensures that all domains within the forest are logically linked, facilitating easier management and navigation within the network. That shared naming structure not only simplifies the identification and management of domains but also reflects the hierarchical nature of the AD DS environment, where each domain is part of a larger, interconnected system.

To better understand namespaces, it can be helpful to draw an analogy to the **Uniform Resource Locator (URL)** system used on the internet. Just as a URL uniquely identifies and locates a specific website on a web server, a namespace in AD DS uniquely identifies and organizes domains within a forest. This logical structure allows administrators to manage resources, apply policies, and maintain security across the network efficiently.

Understanding the role of namespaces is essential for effective AD DS management, as it directly influences how domains are structured, named, and related to one another within a forest. With a clear grasp of namespaces, administrators can ensure a well-organized and navigable AD DS environment. Moving forward, the next topic to explore is the concept of a site within a domain, which represents a physical or logical location in a network. This concept is key to optimizing network traffic and replication within the AD DS infrastructure.

Sites explained

In addition to its logical structure, AD incorporates a physical structure that reflects the network's geographical or organizational layout, known as a *site*. A site represents a specific physical location within an organization's network infrastructure and can encompass one or more domains connected by high-speed links. The purpose of defining sites in AD is to optimize network traffic and enhance overall performance, particularly by managing the replication and authentication traffic between different locations more effectively.

Sites play a crucial role in reducing unnecessary network traffic, especially across **wide-area networks (WANs)** where bandwidth might be limited or expensive. By ensuring that replication traffic—used to synchronize directory data across DCs—is confined to fast, local network links whenever possible, AD helps maintain data consistency without overwhelming slower network connections. That is particularly important in large, distributed environments where DCs might be spread across multiple cities, regions, or even countries.

Moreover, sites in AD aren't just about replication efficiency; they also play a key role in authentication processes. When a user logs in, AD directs the authentication request to a DC within the same site as the user, thus speeding up the authentication process and reducing the load on remote servers. This localization of authentication and replication activities significantly improves the user experience and the reliability of network services.

Understanding how sites function, as well as their impact on replication, is essential for IT professionals tasked with managing or designing an AD infrastructure. Proper site design ensures that network resources are used efficiently, users experience minimal delays, and the network remains resilient even as it scales. Before exploring the details of how AD replication works to keep data

consistent across domains, it's important to grasp how sites contribute to the robustness and efficiency of an AD environment, particularly in large or complex networks.

Exploring replication

Replication in AD is a foundational feature that ensures data consistency and integrity across all DCs within a forest. This process is essential for maintaining an up-to-date and synchronized directory service, where any modifications—whether they involve user account details, security policies, or configuration settings—are promptly reflected across the entire network of DCs. The replication mechanism operates continuously to propagate these changes, preventing any potential conflicts or inconsistencies that could arise if different parts of the network hold divergent versions of the directory data.

The efficiency of replication is managed by the replication topology, which refers to the network of routes that replication data travels between DCs. This topology isn't random; it is meticulously generated and optimized by the **Knowledge Consistency Checker (KCC)**. The KCC assesses the network's structure and dynamics, creating a replication path that balances speed and load distribution, thus ensuring that data updates are disseminated quickly without overwhelming the network's resources. This automated process is crucial for large and complex AD environments, where manual configuration would be impractical and prone to error.

Additionally, AD supports both intra- and inter-site replication. Intra-site replication occurs within the same site, typically using high-speed connections, and is more frequent, ensuring near-real-time data synchronization. Inter-site replication, on the other hand, occurs between different sites and is less frequent. It's optimized to reduce the impact on bandwidth, particularly over slower or more expensive WAN links. Administrators can configure inter-site replication schedules and compression settings to manage this process effectively, balancing the need for up-to-date information with the constraints of network resources.

Understanding replication is not just about grasping the mechanics of how data is synchronized; it also involves recognizing its critical role in maintaining the overall health and performance of the AD environment. If replication fails or is misconfigured, it can lead to outdated or conflicting data across DCs, which can cause a host of issues, from authentication failures to incorrect policy applications.

After mastering replication, the next step in managing AD is to understand the schema—a comprehensive blueprint that defines the structure of all objects and attributes within the directory. The schema is central to how AD organizes and stores data, dictating what types of objects (such as users, groups, or computers) can exist in the directory and what attributes those objects can have. Familiarity with the schema enables administrators to extend or modify the directory to meet specific organizational needs while ensuring compatibility and stability within the AD infrastructure.

Understanding the schema

The **schema** in AD is a critical component that underpins the organization and management of all data within the directory service. It functions as a structured blueprint, dictating how data is stored, organized, and accessed across the entire AD infrastructure. The schema is composed of three core elements:

- **Objects:** These are distinct entities within the directory, such as users, computers, printers, or security groups, each representing a real-world resource or function.
- **Classes:** Objects are categorized into classes, which define the type of the object and set the framework for what it can represent within the directory. For instance, a user object might belong to a **User** class, which outlines specific attributes such as username, password, email address, and department.
- **Attributes:** These are the properties or characteristics assigned to objects, providing detailed information about them. For example, the attributes of a **user** object might include a user's full name, job title, phone number, and login credentials.

The schema not only defines what objects and attributes exist but also sets the rules for how these elements can be created, modified, and managed within the AD environment. That ensures that data integrity and consistency are maintained across the entire network. Changes to the schema are carefully controlled and replicated across the network. This controlled approach ensures that all DCs within the forest remain synchronized, maintaining a uniform structure and enabling seamless management and querying of directory data across different domains and sites. Your understanding of this process is key to ensuring the stability and reliability of the network.

In the following section, we will delve into Microsoft Passport, a modern authentication method that enhances security by allowing users to sign in without traditional passwords, relying instead on more secure alternatives such as biometrics or PINs. This approach not only strengthens security but also improves the overall user experience by simplifying the login process.

Microsoft Passport explained

In today's digital environment, managing an ever-increasing number of passwords for various applications, websites, and services presents a significant challenge. Traditional password-based authentication is not only inconvenient but also fraught with security vulnerabilities, as passwords can easily be forgotten, guessed, stolen, or compromised through phishing attacks. Recognizing these issues, Microsoft introduced *Microsoft Passport*, which is now part of Windows Hello for Business. This cutting-edge, password-less authentication system is designed to enhance both security and user convenience.

Microsoft Passport leverages the **Fast ID Online (FIDO) Alliance** standard, a widely recognized framework for secure, password-free authentication. The system operates using a two-factor authentication model that combines a single sign-in service with a wallet service. This means that

instead of relying on a password, users authenticate their identity using something they possess, such as a trusted device (for example, a smartphone or a security key), along with something unique to them, such as biometric data (fingerprint or facial recognition) or a secure PIN. The combination of these factors ensures that even if one factor is compromised, the other remains secure, significantly reducing the risk of unauthorized access.

By adopting passwordless as part of Windows Hello for Business, organizations can enhance their security posture while simplifying the user experience. Moving away from traditional passwords reduces the risk of security incidents, as passwords are often easily forgotten, guessed, or stolen, especially when not paired with **multi-factor authentication (MFA)**. Users no longer need to remember complex passwords or manage multiple credentials, which leads to fewer password-related security incidents and a decrease in support requests. This approach aligns with the broader industry shift toward more secure and user-friendly authentication methods. By embracing passwordless solutions, organizations can mitigate the vulnerabilities associated with password management and foster a more robust security framework that incorporates multi-factor and biometric-based solutions.

In this section, we have delved into the core components of the AD infrastructure, including domains, forests, trees, sites, schemas, and namespaces. We have also covered how to configure and verify the domain and FFLs in Windows Server 2025. Understanding these elements is crucial for anyone tasked with managing an AD environment. In the next section, we will explore the DNS and its pivotal role in the functioning and management of AD, ensuring seamless name resolution and directory services across the network.

Exploring DNS fundamentals and configurations in Windows Server 2025

The DNS emerged from the ARPANET project in the 1960s, addressing the need for a user-friendly way to identify network devices beyond numerical IP addresses. This concept evolved into the DNS as we know it in the early 1980s, with the release of foundational specifications documented in **Request for Comments (RFCs)**. DNS is organized into a hierarchical structure akin to a tree, where the root zone branches into various domains and sub-domains, each containing resource records that provide essential information about network resources. A domain name is constructed from multiple segments, known as labels, separated by dots—such as `packetpub.com`. This system is underpinned by a distributed database that utilizes a client-server architecture, where network hosts act as name servers. These servers are responsible for resolving domain names to their corresponding IP addresses, ensuring seamless navigation and connectivity across the internet. This hierarchical and distributed approach enhances scalability, efficiency, and reliability in managing domain names and network resources.

Understanding how DNS works

To fully understand how the DNS operates, it's helpful to follow the sequence of steps that occur when you attempt to access a website. DNS is essential for translating human-readable domain names into machine-readable IP addresses, facilitating communication between users and websites. The following describes the DNS resolution process, explaining how your browser finds the correct IP address to connect to when you enter a web address such as www.packtpub.com:

1. **Entering the URL:** When you enter www.packtpub.com into your browser's address bar and press *Enter*, your browser sends a request to connect to this domain.
2. **Recursive resolver:** This request first reaches a crucial component of the DNS infrastructure known as the **recursive resolver**. Typically managed by your **Internet Service Provider (ISP)**, this resolver is responsible for handling queries on your behalf.
3. **Root servers:** The recursive resolver then communicates with **global root servers**, which hold information about **top-level domains (TLDs)** such as **.com**. These servers do not have complete DNS information, but they direct the resolver to the appropriate TLD servers.
4. **TLD servers:** The TLD servers, in turn, respond by providing information that directs the resolver to the authoritative name servers for the specific domain, such as packtpub.com.
5. **Authoritative name servers:** The resolver queries these **authoritative name servers** to find the exact IP address associated with packtpub.com. The authoritative servers contain the actual DNS records that map domain names to IP addresses.
6. **Returning the IP address:** Once the resolver obtains the IP address of the web server hosting packtpub.com, it relays this information back to your browser.
7. **Connecting to the web server:** With the IP address, your browser can establish a connection to the web server and retrieve the website's content for you to view.

This step-by-step process illustrates the intricate workings of DNS, highlighting its role in converting domain names into IP addresses that enable seamless communication across the internet.

Understanding this process underscores the importance of correctly configuring the DNS role within your network. By doing so, you ensure efficient domain name resolution, which is critical for both internal network operations and external internet access.

Installing the DNS role

In Windows Server 2025, the DNS role is crucial for enabling the server to translate domain names into IP addresses, facilitating seamless network communication and access to resources. This role can be configured using the **Server Manager tool**, which provides a straightforward interface for managing server roles and features. As depicted in *Figure 4.12*, the process begins by accessing the Server Manager and selecting the option to add roles and features.

 Figure 4.12 – Installing the DNS role

Figure 4.12 – Installing the DNS role

You can either install the DNS role as an independent service or in conjunction with AD DS. When installed separately, the DNS role functions autonomously to handle domain name resolution. However, integrating the DNS with AD DS, as shown in *Figure 4.13*, enhances the overall functionality of the network by allowing the DNS server to support AD operations, such as DC location and service record lookups. This integration is particularly beneficial for managing large-scale networks where AD and the DNS work together to streamline operations.

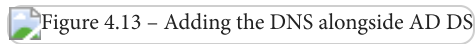


Figure 4.13 – Adding the DNS alongside AD DS

Furthermore, the DNS role is often included as part of the AD DS installation process, providing a cohesive setup that supports the resolution of domain names within the AD environment. This streamlined approach ensures that DNS services are properly configured to support AD's needs, including the automatic creation of necessary DNS records.

With the DNS role successfully installed and configured, you will be well-equipped to manage domain name resolution and enhance network functionality. The next step involves understanding how **hosts** and **LAN Manager hosts (lmhosts)** files contribute to local name resolution, which complements the role of DNS in your network setup.

Understanding the role of hosts and lmhosts files

In any network environment, effective name resolution is fundamental to ensuring seamless communication between devices, and the hosts and lmhosts files are pivotal in facilitating this process. These files are typically located in the `c:\Windows\system32\drivers\etc` directory and provide a straightforward yet powerful mechanism for resolving network names, even in the absence of dedicated name resolution services:

- **Hosts:** This file acts as a static and customizable mapping tool, linking specific IP addresses to **hostnames**. That allows for local DNS name resolution, ensuring that devices on the network can be identified by easily recognizable names rather than numerical IP addresses. That is particularly beneficial in environments where DNS services may be unavailable or unreliable, or where it could require manual overrides for testing and administrative purposes. By editing the hosts file, administrators can control and dictate how names are resolved within the local network, ensuring that critical systems are always reachable by their designated names.
- **LMHOSTS:** The LMHOSTS file maps IP addresses to **NetBIOS** computer names. While its relevance has diminished in modern networks, it may still hold value in specific scenarios, particularly in legacy environments that rely on the NetBIOS protocol for name resolution. This file provides a mechanism for resolving NetBIOS names even when a **Windows Internet Name Service (WINS) server** is not available. However, with the growing emphasis on DNS as the primary name resolution method in Windows Server 2025 and the phasing out of WINS, the practical use of LMHOSTS may be limited for most contemporary applications.

Both the hosts and lmhosts files require manual entries, with each mapping recorded on a separate line to maintain clarity and organization. This manual configuration allows network administrators to

have granular control over name resolution, ensuring that network traffic is directed correctly and efficiently.

Figure 4.14 visually represents the location and structure of the hosts and lmhosts files within a Windows Server 2025 environment, highlighting their importance in network configuration.



Figure 4.14 – The hosts and lmhosts files in Windows Server 2025

Moreover, the concept of a hostname, which serves as the domain identifier for the local computer within a network, is essential for local network identification. This concept, previously discussed in this chapter, is integral to understanding how devices are recognized and communicated within a networked environment. By effectively managing the hosts and lmhosts files, administrators can ensure consistent and reliable name resolution, which is vital for the smooth operation of network services and applications.

Understanding hostnames

As an IT professional, your familiarity with hostnames is a fundamental aspect of effective network management. Hostnames serve as the backbone for device identification and communication within a network, as shown in *Figure 4.15*. A hostname is a logical label assigned to a device, ensuring its unique recognition across a network and enabling seamless interaction, particularly within a **Local Area Network (LAN)**. This identifier is not only crucial for local network operations but also plays a significant role in broader network interactions, as it is often synonymous with a domain name. By assigning clear and meaningful hostnames, you can easily manage and troubleshoot devices, ensuring that each system is readily identifiable and can be efficiently accessed.

The assignment of hostnames is a key step in configuring devices, particularly in environments such as Windows Server 2025, where the correct identification of servers and other networked devices is critical for maintaining order and functionality. A well-chosen hostname simplifies network management by allowing administrators to locate and manage devices quickly within the network infrastructure.



Figure 4.15 – Example of assigning a hostname in Windows Server 2025

In addition to understanding hostnames, it is vital to grasp the concept of DNS zones, which act as administrative segments within a DNS. These are covered in the following subsection.

Understanding DNS zones

A deep understanding of **DNS zones** is not just theoretical but also practical. This level of understanding is crucial for mastering network management. These zones form the backbone of the hierarchical DNS structure that governs how domain names are resolved across a network. DNS zones are integral to the AD namespace, which is closely aligned with the broader DNS namespace, providing a structured and scalable approach to managing domain-related data. By segmenting DNS zones, administrators can store and manage information about specific domains more effectively, ensuring that domain name resolution is both accurate and efficient.

There are three primary types of DNS zones, each serving a distinct purpose:

- **Primary zone:** The primary zone is the authoritative source of DNS information for a domain. It holds the definitive, editable copy of the DNS database and is responsible for maintaining all DNS records within its scope. This zone is the central authority for domain name resolution, ensuring that DNS queries are answered correctly and consistently.
- **Secondary zone:** The secondary zone acts as a backup to the primary zone, containing a read-only copy of the DNS records. This zone is crucial for redundancy, as it allows DNS resolution to continue uninterrupted even if the primary zone becomes inaccessible. The secondary zone is synchronized with the primary zone, ensuring that it reflects the most current DNS information.
- **Stub zone:** A stub zone is a specialized variant of the secondary zone. Unlike the secondary zone, which contains a complete copy of the DNS database, the stub zone only holds enough information—specifically, the IP addresses of the **authoritative DNS servers** for the zone—to direct queries to the correct authoritative server. That makes stub zones useful for simplifying DNS administration and optimizing network traffic by reducing the need for full DNS data replication.

The role of DNS servers in managing these zones is crucial. An authoritative DNS server, which operates the DNS records for a specific domain, plays a critical role in this structure. This server can be configured manually by a system administrator, allowing for precise control over DNS entries or dynamically by other DNS servers through zone transfers and updates. The authoritative server is the final arbiter of DNS queries for its domain, ensuring that responses are accurate and up-to-date. In contrast, a **non-authoritative DNS server** relies on cached data from previous DNS lookups and does not hold the original DNS records. While non-authoritative servers can provide quick responses based on cached information, they are not the definitive source for DNS resolution. That can sometimes lead to outdated or inaccurate responses if the cache is not properly maintained.

Beyond DNS zones, it is also important to understand WINS, a legacy service that resolves NetBIOS names. Although DNS has largely superseded WINS in modern networks, it remains relevant in environments where older systems and applications still rely on NetBIOS for name resolution. Familiarity with WINS can be particularly important in networks that maintain legacy infrastructure, as it ensures that all systems, both old and new, can communicate effectively.

Getting to know WINS

Understanding WINS is crucial for automating the resolution of NetBIOS names to IP addresses, a task that is essential for maintaining smooth network operations. Your role in this is particularly

valuable in environments where NetBIOS names are used. WINS helps resolve these names into IP addresses, enabling seamless access to shared resources such as folders and printers across a network. By managing this name resolution process, WINS ensures that network resources are consistently and efficiently available to users.

In Windows Server 2025, WINS is available as a feature that can be installed through the Server Manager, utilizing the **Add Roles and Features Wizard**, as depicted in *Figure 4.16*. This integration into the server management tools facilitates the setup and administration of WINS, allowing network administrators to configure and maintain the service effectively. The WINS server database is not just updated but updated dynamically as NetBIOS name registrations occur, providing up-to-date resolution information in real time.



Figure 4.16 – Installing the WINS feature in Windows Server 2025

In addition to understanding WINS, it is beneficial to become familiar with the **Universal Naming Convention (UNC)**, which is covered in the following subsection.

The UNC explained

Understanding the **UNC** is crucial for effectively managing and navigating network resources. The UNC offers a consistent way to identify shared network assets across different operating systems, including Unix, by following a standardized format. This format begins with double backslashes, followed by the server name and the specific shared folder, such as `\servername\folder` (illustrated in *Figure 4.17*). By providing a clear and uniform structure, the UNC simplifies network navigation and resource access.

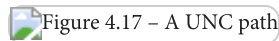


Figure 4.17 – A UNC path

Previously, we explored the DNS, examining its essential components and operations. Now, we will transition to discussing OUs and containers. OUs are critical in structuring and managing AD environments, as they can contain other containers and link to **Group Policy Objects (GPOs)**. In contrast, containers can hold AD objects but do not support GPO links. This distinction is key to ensuring effective AD management and establishing a well-organized directory hierarchy.

Managing OUs and default containers

Understanding the roles of OUs and containers is a cornerstone of effective AD management. These elements, accessible via the **AD Users and Computers** console, are integral to organizing and administering directory objects. OUs provide a flexible structure, allowing administrators to create a

hierarchical organization within the AD environment, making it easier to apply GPOs and manage permissions across different departments or user groups. In contrast, default containers serve as predefined locations for certain types of objects, such as users and computers. Still, they lack the same level of customization and policy control that OUs offer. The following sections will delve deeper into these concepts, examining how OUs can be leveraged to create an organized and secure AD infrastructure while also understanding the limitations and uses of default containers. By mastering these components, administrators can enhance their ability to efficiently manage and secure their AD environments, ensuring a well-structured and easily navigable directory.

Understanding OUs

OUs are critical components within the AD that enable a structured and efficient approach to managing users, groups, computers, and other directory entities. Functioning similarly to folders in a file system, OUs allow administrators to logically group and manage AD objects based on organizational needs. This logical grouping is pivotal in simplifying administrative tasks, such as applying GPOs and managing permissions across different departments, teams, or geographical locations within an organization.

Typically, organizations design their OU structures to reflect their internal business hierarchies, allowing for a tailored management approach that aligns with their operational framework. Each domain within an AD forest can establish its own unique OU configuration, creating a flexible and scalable system that adapts to the evolving needs of the business. This flexibility is particularly valuable in complex environments where different domains may require distinct policies and management practices, as shown in *Figure 4.18*.



Figure 4.18 – An example of OU hierarchy in Windows Server 2025

In addition to OUs, it is essential to understand the role of default containers in AD. These containers are predefined locations where users, computers, and other objects are automatically placed during their creation. Containers are discussed in the following subsection.

Default containers explained

Gaining a thorough understanding of **predefined containers** is essential when a server is promoted to a DC. This promotion automatically triggers the creation of several default containers, which are visually represented in *Figure 4.17*. These containers play a critical role in AD and are distinguished by their immutable nature—they cannot be renamed, deleted, or recreated, and they are not eligible for

linkage to any GPO. This immutability, by design, ensures that the foundational elements of AD remain consistent and secure, thereby preserving the directory's structural integrity.



Figure 4.19 – Default containers in Windows Server 2025

These default containers serve specific purposes, such as organizing users, computers, and other directory objects in a standardized manner. They provide a stable environment for core AD operations, ensuring that certain critical objects are always stored in a predictable location. Although they are not as flexible as OUs, which can be customized to fit the needs of the organization, default containers are still vital for maintaining the AD's foundational structure.

In the following subsections, we will explore the concept of hidden default containers in greater detail. These hidden containers, though not visible in the standard AD interface, play significant roles in the background processes and overall functioning of AD. Understanding both **visible and hidden default containers** will provide a comprehensive view of how AD maintains its integrity and supports the management of directory objects.

Understanding hidden default containers

Understanding the concept of hidden default containers in AD is crucial for system administrators, even though these containers might not be immediately relevant to everyday tasks. These hidden containers serve a significant purpose in maintaining a streamlined and organized view within the **AD Users and Computers** console, preventing unnecessary clutter that could complicate the management of AD objects. By keeping certain containers out of sight, AD ensures that the interface remains user-friendly and manageable, particularly in large and complex environments.

Security considerations also drive the concealment of these containers. Hidden containers protect sensitive system objects, ensuring that only users with the appropriate permissions and knowledge can access them. This layer of security helps safeguard the integrity of the directory and reduces the risk of accidental modifications or unauthorized access to critical system components.

To reveal these hidden default containers, administrators must enable the **Advanced Features** option from the **View** menu, as depicted in *Figure 4.20*. Activating this feature uncovers the hidden containers, allowing for a more comprehensive view and enhanced control over the directory's resources. This capability is particularly valuable for advanced administrative tasks, such as detailed auditing, fine-tuning security settings, or managing objects that are not typically exposed in the standard view.



Figure 4.20 – Default hidden containers in Windows Server 2025

Having gained an understanding of these hidden default containers, it's essential to delve into their practical applications and roles within the AD environment. These containers often house crucial system information, such as infrastructure objects, security principals, and replication data, which are vital for the smooth operation of AD. By understanding how to access and manage these hidden containers, administrators can ensure that they are fully equipped to maintain a secure, efficient, and well-organized directory infrastructure.

The purpose of default container types

The default containers in Windows Server 2025 are integral to the organization and management of AD objects, each serving a distinct purpose:

- **Computers:** This container is the default repository for newly created computer accounts, providing a centralized location for these objects
- **DCs:** This container is specifically designed to house all DC accounts, ensuring they are organized and easily accessible
- **ForeignSecurityPrincipals:** This container is reserved for SIDs from external domains, facilitating cross-domain security and permissions
- **Keys:** This container stores cryptographic key objects, which are essential for secure communications and encryption within the network
- **LostandFound:** This container plays a critical role in maintaining directory integrity by holding orphaned objects that have become detached from their original containers, preventing potential issues with object references
- **Managed Service Accounts:** This container is dedicated to managed service accounts, which are used to provide enhanced security and management for services running on servers
- **Users:** This container is the default location for upgraded or newly created user accounts, making it easier to manage and access user-related objects

Having established an understanding of these default containers, the next step is to delve into the concept of delegating control to an OU. Delegating control involves assigning specific administrative permissions to users or groups for particular OUs, allowing them to manage objects within that OU without granting them full administrative rights across the entire AD environment. This delegation process is crucial for maintaining a secure and organized directory, as it enables administrators to assign responsibilities to non-administrative users while limiting those users' access to only the objects and functions necessary for their roles. This approach helps in balancing administrative control with security, ensuring that users have the appropriate level of access to perform their tasks effectively without compromising the integrity of the overall AD infrastructure.

Delegating authority within an OU

Understanding the function of OUs in AD is essential for effective directory management. OUs serve as a means to organize and manage AD objects systematically. To enhance administrative efficiency,

control can be delegated to specific users or groups within an OU. This process allows for the distribution of administrative responsibilities without granting users **full administrative rights** across the entire AD environment. To delegate control, users or groups must first be moved into the designated OU, as shown in *Figure 4.21*.

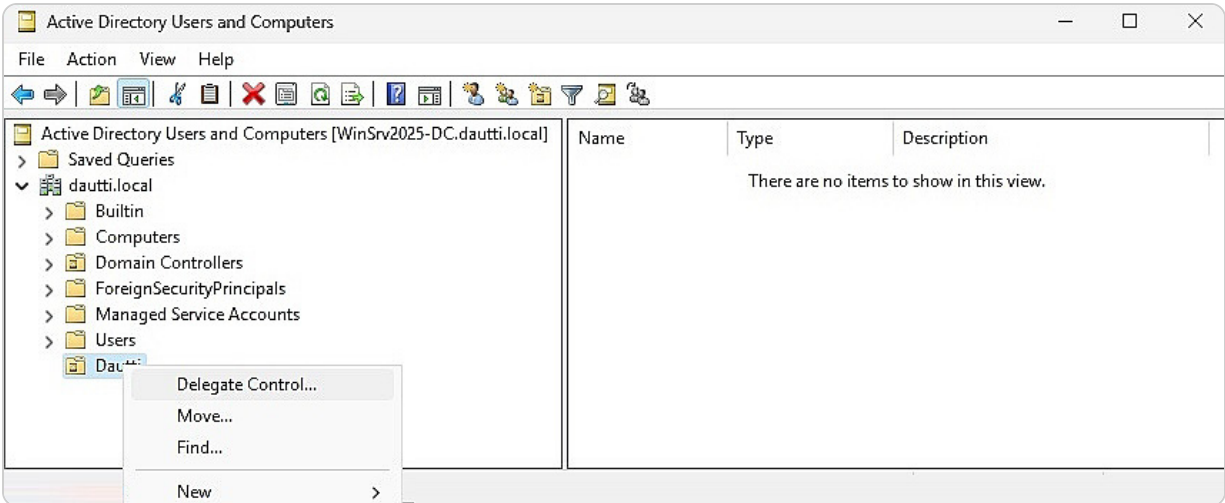


Figure 4.21 – Delegating control to an OU in Windows Server 2025

Delegating control involves assigning specific administrative permissions, such as managing user accounts, resetting passwords, or modifying group memberships within that OU. This focused delegation helps ensure that appropriate personnel perform administrative tasks while maintaining security and organization. By confining permissions to particular OUs, administrators can manage resources more effectively and reduce the risk of unauthorized access or unintended changes.

The delegation of control also enables the implementation of role-based administration, which can improve operational efficiency and accountability. Each delegated administrator can be assigned tasks that are relevant to their role, making it easier to track changes and manage directory objects according to organizational policies.

In the following section, we will further explore the management of user accounts, computer accounts, and groups within AD, delving into how these elements interact with OUs and contribute to a well-structured and secure directory environment.

User and group management within AD

Understanding user and computer accounts, along with groups, is fundamental for managing network access within a Windows-based domain environment. These accounts are crucial elements of AD, enabling both user and device authentication throughout the network. In this centralized system, groups are particularly significant as they simplify the process of assigning and managing rights and

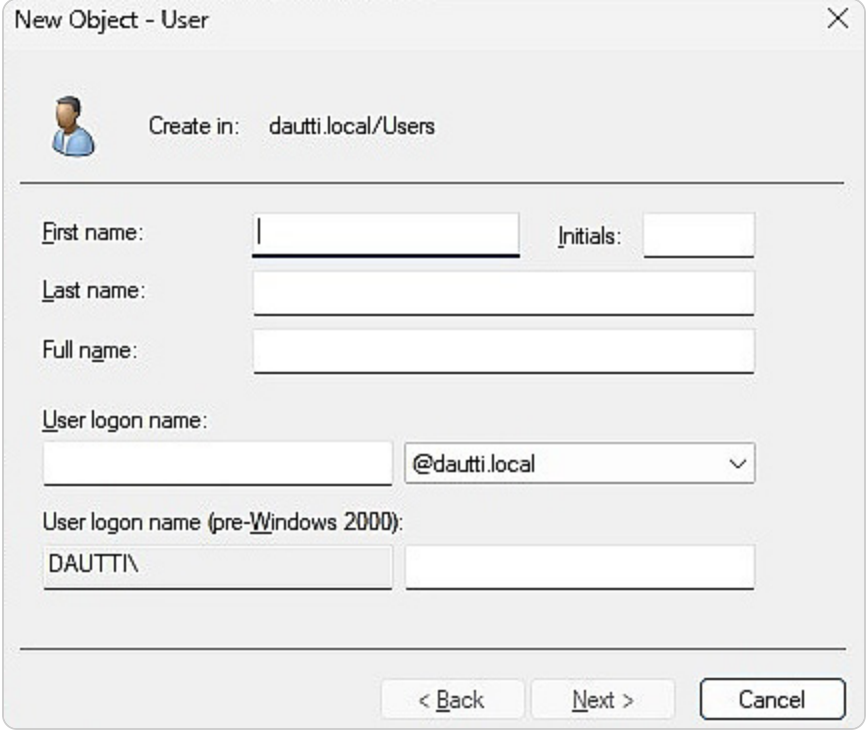
permissions. Groups aggregate multiple accounts, allowing administrators to apply policies and permissions collectively rather than individually. This streamlined approach enhances both security and efficiency. The following sections offer an in-depth exploration of the various types of accounts and groups, detailing their functions and how they are utilized within AD. By examining these components, we will gain a comprehensive understanding of their roles and applications in managing network resources effectively.

Domain accounts explained

Understanding **domain accounts** is essential for managing network access effectively within an AD environment. Domain accounts are authenticated by AD, which enables users to access both local and network resources according to the permissions assigned to the account itself or inherited from group memberships. This centralized authentication framework ensures a streamlined and secure approach to managing access across various services and applications within the network.

To create a domain account in Windows Server 2025, follow these steps:

1. Open the **Active Directory Users and Computers** console by navigating to **Windows Tools**.
2. Right-click on the **Users** container, then select **New** and choose the user.
3. Enter the required user information, as shown in *Figure 4.22*, then click **Next**.



The screenshot shows the 'New Object - User' dialog box. At the top, it says 'Create in: dautti.local/Users'. Below this are several input fields: 'First name:' with an empty text box and 'Initials:' with an empty text box; 'Last name:' with an empty text box; 'Full name:' with an empty text box; 'User logon name:' with an empty text box and a dropdown menu showing '@dautti.local'; and 'User logon name (pre-Windows 2000):' with a text box containing 'DAUTTI\'. At the bottom, there are three buttons: '< Back', 'Next >', and 'Cancel'.

Figure 4.22 – Creating a domain account in Windows Server 2025

4. Set a temporary password, confirm it, and proceed by clicking **Next**.

5. Click **Finish** to complete the creation of the domain account.
6. This process establishes the domain account and integrates it into the AD structure, providing users with access to network resources based on the assigned permissions.

In the following subsection, we will explore the creation and management of **local accounts**, which are also critical for managing user access and security on individual machines and within specific local environments.

Understanding the Local Accounts

Understanding local accounts is crucial for effective access management on individual computers. Unlike domain accounts, which are authenticated through AD and provide network-wide access, local accounts are specific to the computer where they are created and are managed by the Windows **Security Accounts Manager (SAM)**. These accounts offer access to resources on the local machine and can interact with shared resources in a P2P network without requiring additional domain-level permissions.

Local accounts are particularly useful in scenarios where a computer operates independently of a domain, or when domain connectivity is not available. They are created and managed locally, allowing for granular control over permissions and user access on a per-machine basis. That can be advantageous for managing small workgroups or standalone computers where centralized domain management is not feasible.

To create a local account in Windows Server 2025, follow these steps:

1. Access the **Computer Management** console through **Windows Tools**. This console provides a centralized interface for managing various system components, including user accounts.
2. Navigate to **System Tools**, expand the **Local Users and Groups** section, right-click on the **Users** container, and select **New**, followed by the user.
3. Enter the necessary user details, such as the user's name and password, as illustrated in *Figure 4.23*. Then click **Create** to complete the process.

New User ? X

User name:

Full name:

Description:

Password:

Confirm password:

☒ User must change password at next logon

☐ User cannot change password

☐ Password never expires

☐ Account is disabled

Figure 4.23 – Creating a local account in Windows Server 2025

IMPORTANT NOTE

An important consideration when creating a local account in Windows Server 2025 is that the server should not be operating as a DC. If the server is designated as a DC, it will handle domain-related functions and manage AD DS, which complicates local account management. By ensuring that the server is not a DC, you avoid the complexities of domain management, allowing for the straightforward setup and management of local accounts without the additional overhead of domain services.

Local accounts are stored and authenticated by SAM on the local machine, which ensures that access control and permissions are enforced independently of the network domain. These accounts are ideal for scenarios where local administration and access control are required.

In the following subsections, we will delve into **user profiles**, which are essential for storing and managing information about individual users. User profiles contain critical data that helps personalize and manage the user experience on both local and networked systems.

The User Profiles Explained

Understanding the different types of **user profiles** in Windows Server environments is fundamental for effective user management and customization. The following sections provide a brief explanation of the three types of user profiles in AD: local user profiles, which are tied to a specific machine; roaming user profiles, which offer flexibility across multiple devices; and mandatory user profiles, which maintain a fixed configuration without user modifications. Let's take a closer look:

- **Local user profile:** When a user logs into a computer for the first time, a local user profile is created and stored on that specific machine, as depicted in *Figure 4.24*. This profile includes the user's settings and documents, as well as application data tailored to that particular computer. The local profile is ideal for individual use on a single machine but lacks flexibility when users need to access their environment from multiple devices.

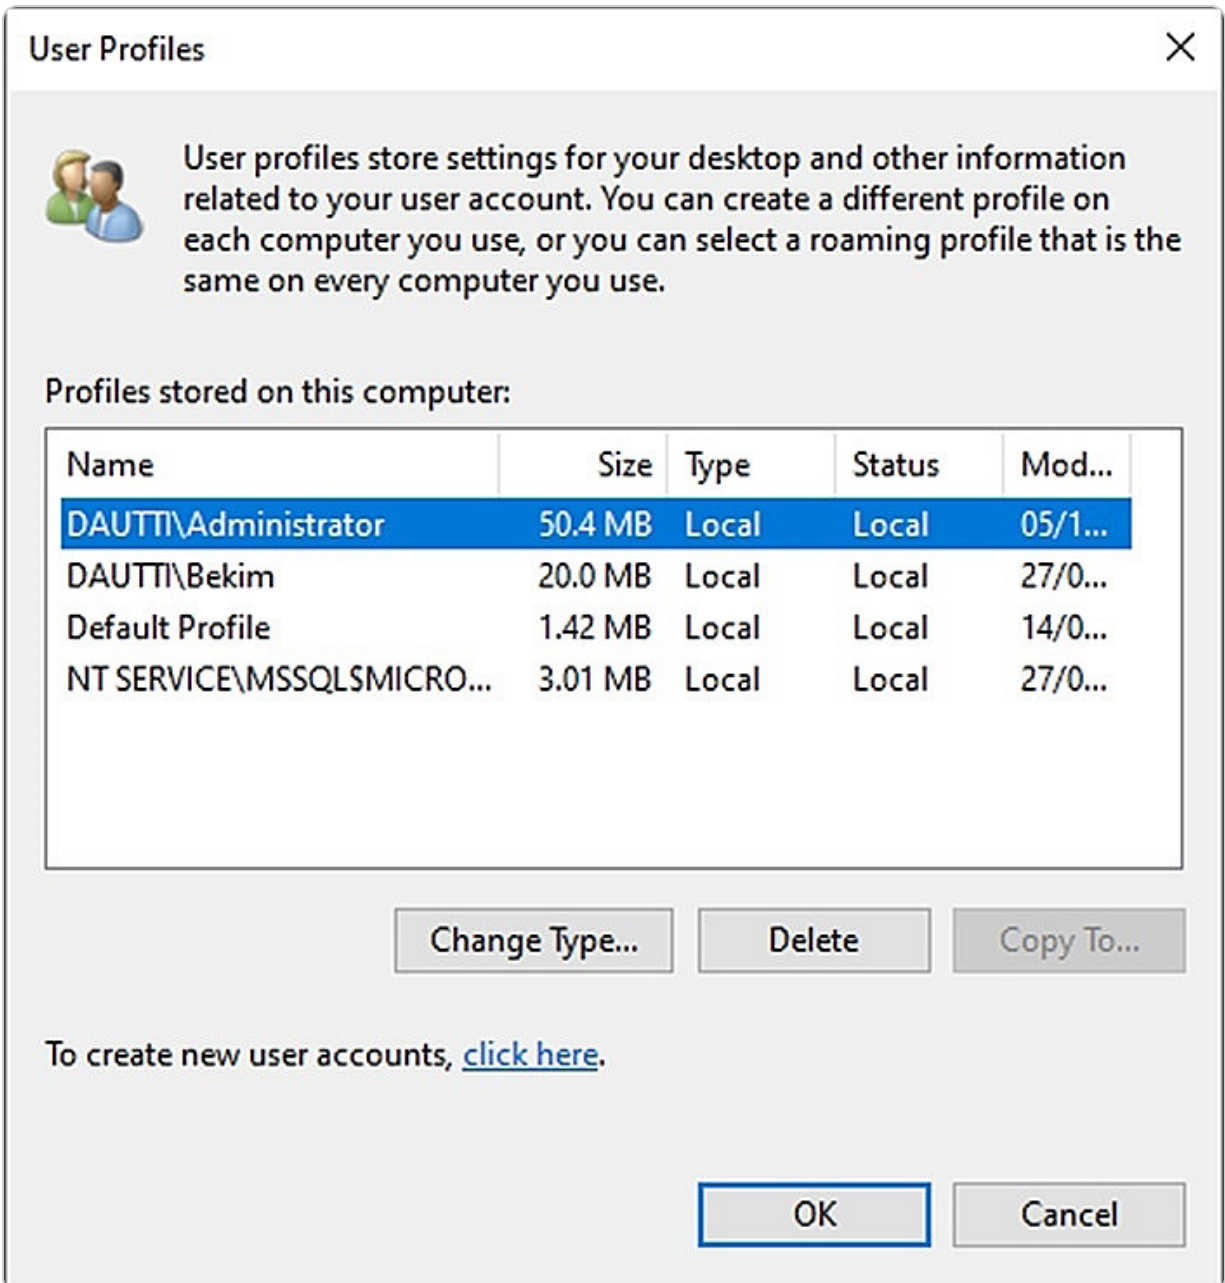


Figure 4.24 – User profiles in Windows Server 2025

- **Roaming user profile:** This kind of profile enhances this flexibility by allowing users to access their personalized settings and files from any computer within the network. This profile is essentially a copy of the local profile stored on a network share. When a user logs in from a different computer, their roaming profile is retrieved from the network, providing a consistent experience across different machines. This type of profile is especially useful in environments where users frequently switch between computers.
- **Mandatory user profile:** This type of profile enforces a fixed profile configuration. These profiles, which are also stored on a network share, are based on a pre-configured template. Any changes made by the user during their session are not saved when they log off. That ensures that every time the user logs in, they start with the same baseline configuration, which is useful in environments where uniformity is required and user customizations are not desired.

In summary, local user profiles are tied to individual computers, roaming profiles provide flexibility by being accessible from any networked machine, and mandatory profiles maintain consistency by discarding user changes and relying on a fixed template. Each type of profile serves distinct purposes, helping administrators manage user environments effectively based on organizational needs.

Next, we will delve into computer accounts, which play a critical role in identifying and managing computers within both local and centralized domain environments. These accounts are essential for maintaining network security and ensuring proper access control across the network.

Understanding Computer Accounts

In an AD environment, **computer accounts** are critical for identifying and managing computers within a domain. Before joining the domain, each computer must have a unique hostname to prevent conflicts. This unique identifier ensures that the computer can accurately be tracked and managed within the network. Once a computer is successfully added to the domain, it retains its hostname for continuous interaction with other domain resources, including files, applications, and services. This setup allows for seamless communication and integration with the domain.

The **Active Directory Users and Computers** console efficiently handles computer account administration, as illustrated in *Figure 4.25*. This console enables administrators to view and manage computer accounts, configure properties, and apply policies. Tasks such as resetting passwords, enabling or disabling accounts, and modifying account settings are performed here.

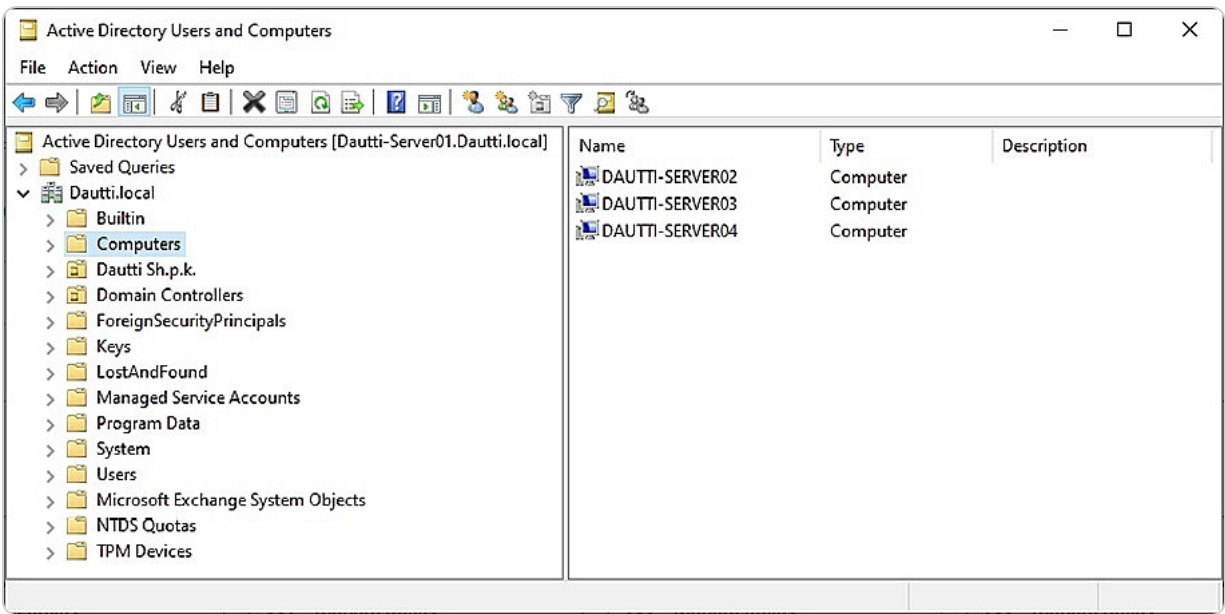


Figure 4.25 – Computer accounts in Windows Server 2025

Understanding computer accounts is essential for maintaining network integrity and ensuring proper resource access. These accounts play a vital role in authenticating and authorizing computers within the domain, thus supporting effective network management and security.

As we move forward, we will shift our focus to groups within the AD framework. Groups are integral to managing permissions and access rights, simplifying the assignment of roles and privileges, and streamlining administrative tasks. They help organize users and computers, apply consistent policies, and enhance overall network security.

Understanding Group Types

Understanding **group types** within AD is fundamental for optimizing network management and ensuring security. AD groups simplify the administration of permissions and rights by allowing administrators to manage multiple AD objects collectively rather than configuring each object individually. This approach not only enhances efficiency but also helps maintain consistent security policies across the network. Groups themselves are also AD objects and can be moved or reorganized within different OUs to align with organizational changes or administrative needs.

As presented in *Figure 4.26*, groups are administered using the **Active Directory Users and Computers** console. In AD, groups are classified into two primary categories:

- **Security groups:** These groups are essential for managing access to shared network resources such as files, folders, and printers. They apply permissions and enforce security policies across the network. Security groups can be nested within other security groups to create a hierarchical permission structure, allowing for more granular control over resource access.
- **Distribution groups:** These groups are designed to facilitate the distribution of email messages within an organization. They simplify the process of sending communications to large groups of users by acting as mailing lists. While distribution groups do not have permissions assigned to them and cannot be used to control access to resources, they play a crucial role in streamlining internal communication.


 Figure 4.26 – Group types in Windows Server 2025

Figure 4.26 – Group types in Windows Server 2025

Understanding these group types and their functions enables administrators to manage network resources and communication effectively. In the subsequent sections, we will delve into default groups—predefined groups that come with AD—and the process of creating new groups. This knowledge is essential for organizing user roles, managing access to resources, and delegating administrative tasks efficiently within an AD environment.

Getting to know default groups

Understanding default groups in AD is fundamental for effective network administration. When a server is promoted to a DC, it automatically generates a variety of default groups, as shown in *Figure*

4.27, which illustrates the default groups in Windows Server 2025. These default groups are designed to simplify administrative tasks by grouping related AD objects, thereby easing the process of assigning permissions and access rights.

 Figure 4.27 – Default groups in Windows Server 2025

Figure 4.27 – Default groups in Windows Server 2025

Default groups are pre-configured with specific roles and permissions, which can significantly streamline network management. For example, default groups such as **Domain Admins**, **Enterprise Admins**, and **Schema Admins** have predefined levels of administrative privileges that are crucial for managing different aspects of the AD environment. By leveraging these groups, administrators can efficiently manage user access and enforce security policies without the need to configure permissions for each user or object manually.

Furthermore, default groups help ensure consistent application of policies and permissions across the network, which enhances both security and operational efficiency. They also facilitate the delegation of administrative tasks by providing predefined roles that can be assigned to users based on their responsibilities. In the next section, we will delve into the concept of group scopes and explore the various types that are available. Understanding group scopes is essential for optimizing group management within an AD environment, as they determine how groups interact with AD objects and influence the scope of permissions and policies applied across the network.

Understanding group scopes

Understanding **group scopes** is a foundational aspect of managing AD environments effectively, as they directly influence how permissions and policies are applied across an organization's network. Group scopes define the reach and applicability of group memberships within the AD structure, which is crucial for maintaining security and efficiency in resource management.

In AD, there are three primary group scopes, each serving distinct purposes and contexts, as illustrated in *Figure 4.28*:

- **Domain local group scope:** This scope is designed to manage access to resources within the local domain. It allows the inclusion of accounts, domain local groups, global groups, and universal groups, enabling administrators to assign permissions to local resources efficiently. Domain local groups are particularly useful when managing access to resources such as file shares, printers, and other domain-specific resources where you want to limit access to users and groups within that domain.
- **Global group scope:** The global group scope is used to organize users and groups within the same domain that share common access requirements. This scope includes accounts and global groups specific to the parent domain's global group. Global groups are typically employed for assigning permissions to resources across different domains within the same forest, making them ideal for scenarios where users from multiple domains need access to shared resources.
- **Universal group scope:** The universal group scope is the most expansive, allowing the inclusion of accounts, global groups, and universal groups from any domain within the forest. This scope is essential for managing permissions across multiple domains,

making it highly effective in large, multi-domain environments. Universal groups are particularly useful when you need to assign permissions consistently across an entire forest, ensuring that users in different domains have appropriate access to resources regardless of their domain membership.



Figure 4.28 – Group scopes in Windows Server 2025

Each of these group scopes plays a critical role in ensuring that permissions and policies are applied appropriately and consistently within the AD environment. By understanding and utilizing these scopes correctly, administrators can enhance both the efficiency and security of their network management practices.

Moreover, the proper use of group scopes can prevent common issues such as over-permission, where users have more access than necessary, or under-permission, where legitimate access is denied. This balance is crucial for maintaining a secure and well-functioning AD environment.

In the following subsection, we will delve into the concept of group nesting. This concept builds upon the principles of group scopes by allowing administrators to create more complex and flexible group structures. Group nesting further refines the ability to manage permissions and access rights, offering a powerful tool for large-scale AD environments.

Group nesting explained

Understanding group nesting within AD is a fundamental aspect of efficient and secure permission management in complex IT environments. Group nesting allows for the hierarchical organization of groups, enabling administrators to assign permissions more effectively by leveraging a structured, tiered approach. This method not only simplifies the administration of access controls but also reduces redundancy and potential errors that could arise from individually assigning permissions to numerous user accounts.

In practice, group nesting is guided by best practices such as Microsoft's **Accounts, Global, Domain Local, Permissions (AGDLP)** and **Accounts, Global, Universal, Domain Local, Permissions (AGUDLP)** methodologies. These models offer a systematic approach to managing group memberships and permissions across a network:

- In the AGDLP model, user accounts are first assigned to a global group, which typically represents a specific role or department within the organization. This global group is then nested within a domain local group, which is responsible for managing access to specific resources within the local domain. Permissions are assigned to the domain local group, thereby granting access to all members of the global group in one step. This method is particularly effective in environments where users need consistent access to resources within a single domain.
- The AGUDLP methodology extends the AGDLP model by incorporating a universal group into the nesting structure. Here, the global group is first added to a universal group, which can span multiple domains within a forest. The universal group is then included in a domain local group, which controls access to resources. This approach is ideal for larger multi-domain environments,

where users require access to resources across different domains. By utilizing universal groups, administrators can maintain a consistent permission structure across the entire forest, ensuring that users have the necessary access regardless of the domain they are operating within.

These structured methodologies not only streamline the management of permissions but also enhance the security and scalability of the AD environment. By reducing the number of individual permissions assignments and centralizing control within well-defined group structures, administrators can more easily enforce security policies, audit access controls, and respond to organizational changes.

After gaining a solid understanding of the foundational elements of AD, such as DNS, OUs, and containers, as well as the classification of computer accounts and groups, the next step is to proceed with the installation of the AD DS and DNS roles. This phase is crucial as it lays the groundwork for configuring and managing the AD environment, ensuring that it meets the security, scalability, and administrative needs of the organization.

Chapter exercise – installing the AD DS and DNS roles and promoting the server to a DC

In this chapter's exercise, you will be guided through the essential steps to install the AD DS and DNS roles, culminating in the promotion of your server to a fully functional DC. This process is a critical aspect of establishing a secure and efficient network infrastructure within an organization.

The exercise begins with the installation of AD DS, which is the backbone of identity and access management in Windows Server environments. It's followed by the configuration of DNS, which is crucial for name resolution within the domain. You will then proceed to promote the server to a DC, a key role that manages network security, user authentication, and policy enforcement across the domain. By following the detailed instructions provided, you will gain a comprehensive understanding of how to implement and configure these roles, ensuring that your network is both robust and well-organized. This exercise not only enhances your practical skills but also deepens your theoretical knowledge, preparing you for more advanced network administration tasks.

To begin the installation process for AD DS and DNS roles and to promote the server to a DC, complete the following steps:

1. Start by accessing **Server Manager**. Click the **Start** button and select the **Server Manager** tile from the **Start** menu.
2. In the **Server Manager** window, find the **WELCOME TO SERVER MANAGER** section and click on **Add Roles and Features**, as illustrated in *Figure 4.27*. That will launch the **Add Roles and Features Wizard**; click **Next** to proceed.

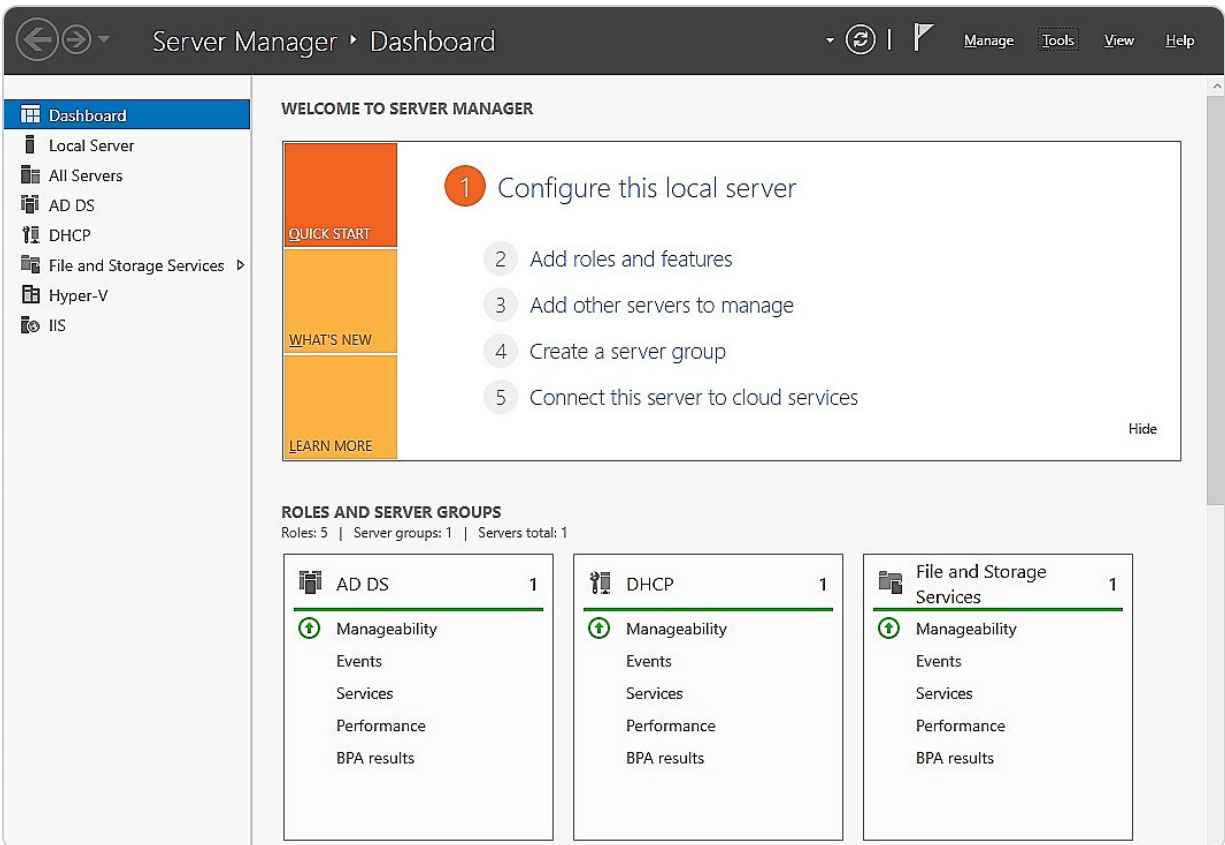


Figure 4.29 – Adding roles and features to the server using Server Manager

3. Select the **Role-based or feature-based installation** option and click **Next**.
4. Ensure that the **Select a server from the server pool** option is checked, then click **Next** again.
5. You will now choose the **Active Directory Domain Services** role from the list, as shown in *Figure 4.29*, and click **Next**.



Figure 4.30 – Installing AD DS in Windows Server 2025

6. When prompted with the **Add features that are required for Active Directory Domain Services** window, click **Add Features** and then click **Next**.
7. Leave the default settings in the **Select Features** step unchanged, and click **Next** once more.
8. Carefully read through the description and key points about AD DS installation, and then click **Next**.
9. Next, confirm the installation selections for the AD DS role and click **Install**. You can either close the wizard or wait for the installation to complete.
10. Once finished, click **Close** to exit the **Add Roles and Features Wizard**.
11. In the **Notifications** area, you will see a **Promote this server to a domain controller** option. Click on it to start the AD DS configuration wizard.
12. Choose the **Add a new forest** option, as depicted in *Figure 4.31*, and enter the desired **Root domain name** value. Click **Next** to continue.



Figure 4.31 – AD DS deployment configuration wizard

13. Accept the default settings for the forest and DFLs, and then set a **Directory Services Restore Mode (DSRM)** password. Click **Next**.
14. If your network already has a DNS server, you may need to manually create a delegation for that DNS server to ensure proper name resolution from outside your domain; otherwise, no action is needed. Click **Next**.
15. You can either accept the default **NetBIOS** name or modify it as necessary. Click **Next**.
16. Similarly, you can accept the default paths for the AD DS database, log files, and SYSVOL or change them according to your requirements. Click **Next**.
17. Review your configuration options, then click **Next**.
18. Once the wizard confirms that all prerequisites have been met, click **Install**. The server will then restart to complete its promotion to a DC.

This exercise has provided you with step-by-step instructions on installing AD DS and DNS roles and promoting your server to a DC. Completing these steps equips you with the foundational skills needed to manage and secure your network infrastructure effectively.

Summary

In this chapter, you developed a comprehensive understanding of directory services and the key principles of naming resolution, which are fundamental to managing and securing a network environment. We started by delving into the AD DS role, which serves as the backbone of identity management within a network, enabling the authentication of users and devices. Alongside this, we examined the DNS role, which plays a crucial part in translating human-friendly domain names into IP addresses, ensuring seamless communication across the network. We also explored the structure and functionality of OUs in AD, which are vital for logically organizing resources and delegating administrative control. By learning how to delegate controls within OUs, you gained the ability to distribute administrative tasks effectively, thereby enhancing the efficiency and security of your AD environment. The chapter further covered the setup and management of user accounts and groups, which are essential for maintaining an organized and secure user base by appropriately assigning users to relevant organizational groups based on their roles and responsibilities.

Moreover, we walked through a detailed chapter exercise that provided hands-on experience in installing the AD DS role and promoting a server to a DC. This practical component not only reinforced your theoretical knowledge but also equipped you with the skills necessary to implement these services in a real-world setting. As we move forward, the next chapter will focus on expanding your server's capabilities by adding and configuring additional roles in Windows Server 2025. That will further enhance your ability to manage and optimize your network infrastructure.

Questions

1. **True or False?** An AD is a distributed database that stores objects in a hierarchical, structured, and secure format.
2. **Fill in the blank:** _____ minimizes the number of individually assigned permissions to users or groups.
3. Which of the following user profiles are used mainly in Windows-based domain networks?
 - A. Domain user profile
 - B. Security user profile
 - C. Roaming user profile
 - D. Mandatory user profile
4. **True or False?** The WINS server maps the IP addresses to BIOS names.
5. **Fill in the blank:** _____ is a set of communication paths through which the DC's replication data travels.
6. Which of the following are AD's group scopes?
 - A. OU
 - B. Security group
 - C. Global group
 - D. Universal group
7. **True or False?** UNC is a standard to identify a share in a computer network.
8. **Fill in the blank:** _____ is a server responsible for securely authenticating requests to access your organization's domain resources.
9. Which snap-ins for MMC are used to manage AD?
 - A. Active Directory Administrative Center
 - B. Active Directory Users and Computers
 - C. UNC
 - D. OU
10. **True or False?** The best example of a domain is a client/server network where a dedicated server on the network is used to provide services.
11. **Fill in the blank:** _____ stores the primary copy of the DNS database and maintains all DNS zone records.
12. Which of the following are forest-wide operations master roles?
 - Master schema
 - Domain naming master
 - LAN manager hosts
 - Default containers
13. Discuss AD DS and DNS roles and their implementations.
14. Discuss Microsoft's recommendations, AGDLP and AGUDLP, for assigning permissions.

Further reading

- *AD DS Deployment*: <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/deploy/ad-ds-deployment>
- *Domain Name System (DNS)*: <https://learn.microsoft.com/en-us/windows-server/networking/dns/dns-top>
- *Creating an Organizational Unit Design*: <https://docs.microsoft.com/en-us/windows-server/identity/ad-ds/plan/creating-an-organizational-unit-design>
- *Managing Groups*: <https://docs.microsoft.com/en-us/windows/win32/ad/managing-groups>

5

Adding Roles to Windows Server 2025

This chapter aims to equip you with knowledge about the factors influencing server hardware selection and the best practices for evaluating server performance. Understanding the server's role within a network and its hardware components will enable you to choose the most suitable server hardware for your requirements and troubleshoot hardware-related issues effectively. Moreover, the chapter will cover the techniques and strategies for monitoring server performance. Effective performance monitoring involves not only identifying and mitigating potential performance issues before they escalate but also actively engaging in prompt responses to prevent further performance degradation. Establishing a performance baseline—documenting the server's performance under typical workloads—is crucial for generating comprehensive reports on overall performance and playing a vital role in the monitoring process. The chapter concludes with a practical exercise focused on analyzing performance logs and setting up alerts.

In this chapter, we're going to cover the following main topics:

- Understanding server roles and features in Windows Server 2025
- Exploring application server roles and their implementations
- Configuring web services and their roles in Windows Server 2025
- Setting up remote access roles and their functionalities
- Deploying file and print services for network environments
- Installing web server (**Internet Information Services (IIS)**) and **Print and Document Services (PDS)** roles

Technical requirements

To complete the tasks outlined in this chapter on various servers using Windows Server 2025, you will need specific hardware and software configurations:

- It would be best if you had a PC running **Windows 11 Pro**, equipped with at least 16 GB of RAM and a 1 TB hard drive, with internet connectivity.
- Additionally, you will need three **virtual machines (VMs)** running **Windows Server 2025 Standard** (Desktop Experience). Each VM should have a minimum of 4 GB of RAM and 100 GB of hard drive space and have internet access.

These VMs will serve different roles: one as a file server, another as a web server, and the third as a print server.

Understanding server roles and features in Windows Server 2025

Before assigning roles to a server, it's essential to clearly define its intended function within your organization's IT infrastructure. This chapter provides a comprehensive overview of the various roles, role services, and features available in Windows Server 2025, helping you make informed decisions about server configuration.

Roles and features overview

A **server role** defines the core function that a server performs within a network. For example, if a server's primary purpose is to store and manage shared files, then the File and Storage Services role is installed to fulfill that responsibility. Similarly, a server designated to host web applications will have the Web Server (IIS) role to handle HTTP requests securely, providing an essential platform for internet and intranet services. Servers that enable secure remote access will implement the Remote Access role, which facilitates connectivity solutions such as **virtual private networks (VPNs)** and DirectAccess, allowing users to access network resources securely from remote locations.

In most cases, assigning a single role to each server is optimal, as this ensures streamlined performance and simplifies server management. However, there are scenarios where multiple roles may be deployed on a single server. In such cases, careful planning is essential to balance hardware resources against role-specific requirements, ensuring compatibility and preventing potential conflicts or performance bottlenecks. This modular approach allows Windows Server 2025 to serve a range of purposes within an organization, with each role contributing to a reliable, responsive, and secure network infrastructure.

Role services explained

Beyond the basic roles, Windows Server 2025 offers **role services**—optional components that enhance or extend the functionality of a server role. These services allow administrators to tailor the server's capabilities to specific needs. For instance, enabling remote printing over the internet requires not only the installation of the PDS role but also the addition of the Internet Printing role service. This layered approach enables you to customize the server's functionality to meet precise operational requirements, providing flexibility in how the server supports your organization's needs.

Core native roles and features

Windows Server 2025 includes a range of built-in roles and features designed to support critical infrastructure needs without relying on additional applications. The choice to highlight three specific roles—**Active Directory Certificate Services (AD CS)**, **Rights Management Services (RMS)**, and **Network Policy Server (NPS)**—is based on their relevance to foundational aspects of Windows Server management: security, access control, and data protection. These roles provide essential infrastructure

capabilities that many organizations rely on, regardless of additional applications, such as Exchange or SQL Server:

- **AD CS:** This role provides a scalable, secure method to issue and manage digital certificates within an organization. It is fundamental for supporting secure communication, data integrity, and user authentication. It plays a crucial role in environments that prioritize security by enabling tasks such as **Secure Sockets Layer (SSL)** / **Transport Layer Security (TLS)** for websites and authenticating users and devices.
- **RMS:** RMS is vital for protecting information and safeguarding sensitive documents and communications by enforcing access and usage restrictions. This ensures that only authorized users can interact with protected content, which is especially critical in industries handling sensitive or regulated data.
- **NPS:** NPS functions as a RADIUS server, supporting centralized network access authentication, authorization, and accounting. This capability is invaluable for managing secure network access, particularly in environments where multi-site and cloud integration is essential, facilitating secure connections for VPNs, wireless networks, and other remote access solutions.

While these roles are critical, Windows Server 2025 also includes several other valuable native features that warrant exploration:

- **File and Storage Services:** Integral for centralized file sharing, storage management, and data deduplication, addressing core needs in networked environments
- **Hyper-V:** Essential for organizations leveraging virtualization, optimizing server utilization, and providing isolated virtual environments for various applications
- **DNS and DHCP:** These fundamental roles underpin network infrastructure, providing domain name resolution and **IP address management (IPAM)**
- **Windows Server Update Services (WSUS):** Critical for patch management, ensuring that servers and connected devices receive timely updates to maintain security and compliance

Expanding in a broader selection of these roles provides a more comprehensive view of Windows Server's native capabilities, equipping administrators with a solid foundation for managing network security, compliance, and accessibility. This understanding lays the groundwork for extending server functionalities through additional applications, aligning with the technical reviewer's feedback to emphasize the importance of native features in Windows Server environments.

Understanding server features

In addition to roles and role services, **server features** are supplementary components that support or enhance specific functions within the server environment. For example, installing the .NET Framework 3.5 feature might be necessary to run particular applications or services. In contrast, the IPAM feature provides advanced management capabilities for DHCP and DNS roles. Features such as WINS can be crucial in environments where resolving NetBIOS names across multiple subnets is necessary. By carefully selecting and installing the appropriate features, you can ensure that your server is fully equipped to handle its designated tasks efficiently and effectively, thus contributing to the overall stability and performance of your IT infrastructure.

In summary, understanding and strategically configuring roles, role services, and features in Windows Server 2025 is key to optimizing server performance and meeting the unique needs of your organization's IT environment. This chapter equips you with the knowledge to make informed decisions that enhance the functionality and reliability of your server deployments.

An overview of Server Manager

Server Manager is a crucial tool for adding, configuring, and managing server roles in Windows Server 2025. First introduced with Windows Server 2008, this tool has continually improved, offering a streamlined and intuitive interface that simplifies server administration. Whether you are working with a local server or managing remote servers, Server Manager allows you to install and oversee server roles efficiently. The interface is divided into two main sections: the **scope pane**, which displays all installed roles, and the **details pane**, which provides comprehensive information and management options for each selected role. This central console not only helps in monitoring the health and performance of the server but also allows for easy access to role-specific tools and settings. As shown in *Figure 5.1*, Server Manager is indispensable for performing a wide range of administrative tasks, making it an essential component of Windows Server 2025's management suite.

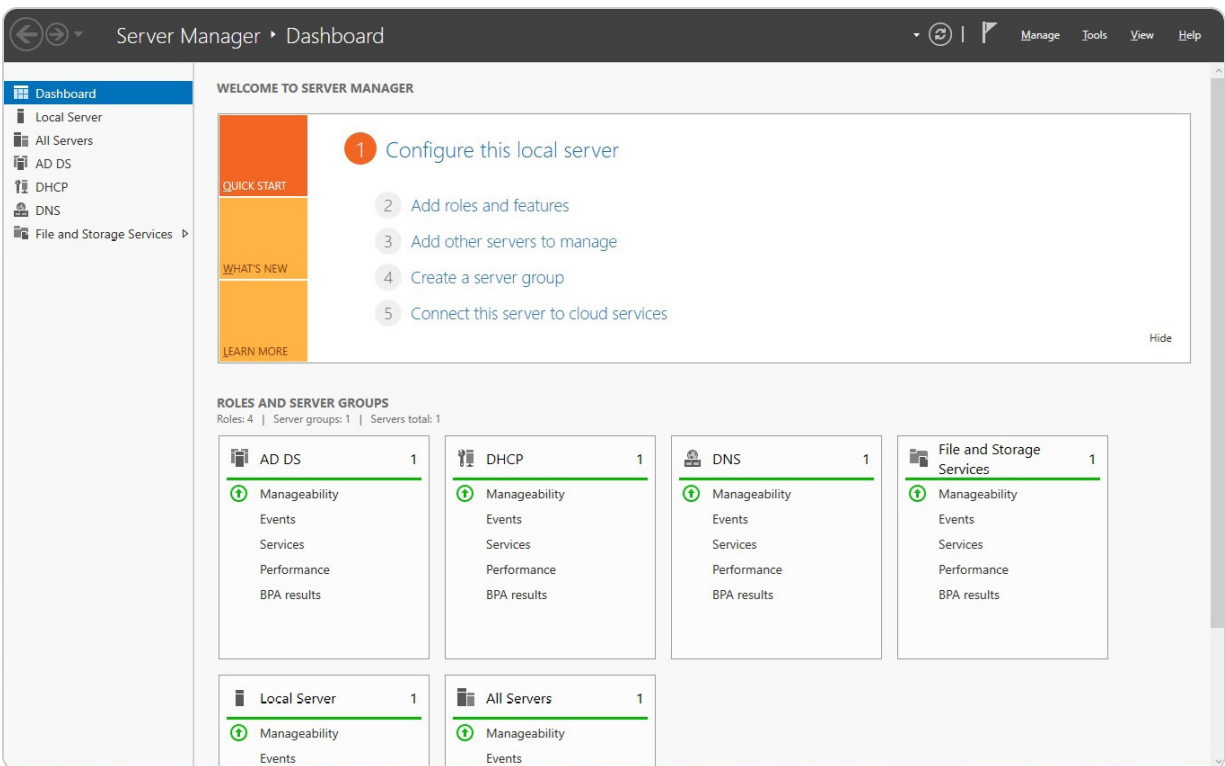


Figure 5.1 – The Server Manager interface in Windows Server 2025

With a solid understanding of server roles and features through Server Manager, the next section will delve into the concept and types of application servers, expanding your knowledge of server infrastructure and management.

Exploring application server roles and their implementations

This section delves into the various types of application servers that are integral to IT operations. Application servers are specialized network-based servers designed to provide specific services to users or other applications within an organization. They play a crucial role in handling diverse tasks, including email communication, web hosting, and database management, among others. The following subsections will explore some of the most widely used and up-to-date application servers, highlighting their importance and functionality.

We begin our exploration with the practical implementation of a mail server, which is a core element of IT infrastructure. A mail server manages the sending and receiving of emails across the internet, making it an essential tool for effective communication within any organization. By learning how to install and configure a mail server on Windows Server 2025, you will gain valuable skills that are directly applicable to real-life IT environments.

Understanding the email server in Windows Server 2025

An **email server**, often referred to as a **mail server**, plays a crucial role in handling the transmission and reception of emails across the internet. To function as an email server, a server must have specialized software installed. **Exchange Server** is commonly used for Windows-based systems. Exchange Server provides a comprehensive platform that allows system administrators to configure and manage email accounts, enabling the server to handle tasks such as sending, receiving, and storing emails.

Several key components and communication protocols underpin the functionality of an email server:

- **Mail Transport Agent (MTA):** This component is essential for transferring emails between servers. The MTA handles the routing of emails from the sender's server to the recipient's server.
- **Mail Delivery Agent (MDA):** Once emails reach their destination server, the MDA takes over, ensuring that the messages are correctly delivered to the appropriate user's inbox.
- **Mail User Agent (MUA):** This component provides the interface for users to compose, send, receive, and read emails. The MUA interacts with both the MTA and MDA to ensure seamless email communication.
- **Simple Mail Transfer Protocol (SMTP):** SMTP is the protocol used by the MTA to transfer emails between servers. It operates on port 25 and is the backbone of email delivery in most systems.

- **Post Office Protocol (POP):** POP is a protocol that operates on port **110**. It allows users to download emails from the server to their local devices, enabling offline access to their messages. POP is useful in scenarios where emails are stored locally and not synchronized across multiple devices.
- **Internet Message Access Protocol (IMAP):** IMAP, operating on port **143**, is a more flexible protocol compared to POP. It allows users to retrieve emails from the server and synchronizes them across multiple devices. With IMAP, users can access their emails from different locations without the need to download them, as they remain stored on the server.

While Exchange Server provides a robust and advanced solution for managing email communications within an organization, there are situations where a more straightforward setup might be sufficient. For example, if the goal is to establish a basic email service that focuses on sending and forwarding emails, Windows Server 2022 allows you to add the SMTP Server feature. This lightweight option is suitable for environments where full Exchange Server capabilities are not required.

NOTE

*One available option for establishing an email service within your organization's network is the upcoming **Exchange Server Subscription Edition (Exchange Server SE)**, anticipated for release in 2025. This Microsoft product enables you to set up and manage a mail server tailored to your organization's needs. To leverage its features, you must install and configure Exchange Server SE on your organization's server infrastructure. Exchange Server SE provides a comprehensive range of functionalities designed to enhance email communication and collaboration across your organization's network.*

SMTP Server feature removed in Windows Server 2025

In Windows Server 2025, the SMTP Server feature, which allows servers to send and forward emails over the internet, has been discontinued. As a result, organizations need to explore alternative solutions for managing email services. One option is to use Exchange Server, a comprehensive and advanced email platform that provides extensive functionalities for managing email communications within an organization. Exchange Server is capable of handling various email-related tasks, including hosting mailboxes, managing calendars, and ensuring secure communication.

For those looking for a cloud-based solution, Microsoft 365 offers **Exchange Online** (its admin center is illustrated in *Figure 5.2*), a part of the Microsoft 365 suite that delivers email and calendaring services without the need for on-premises infrastructure. Exchange Online operates in the cloud, hosted by Microsoft, which eliminates the need for organizations to manage physical or virtual servers directly. This cloud-based approach provides the advantage of reduced maintenance and simplifies upgrades. However, it also means relinquishing some control over configuration and system changes compared to an on-premises Exchange Server deployment.

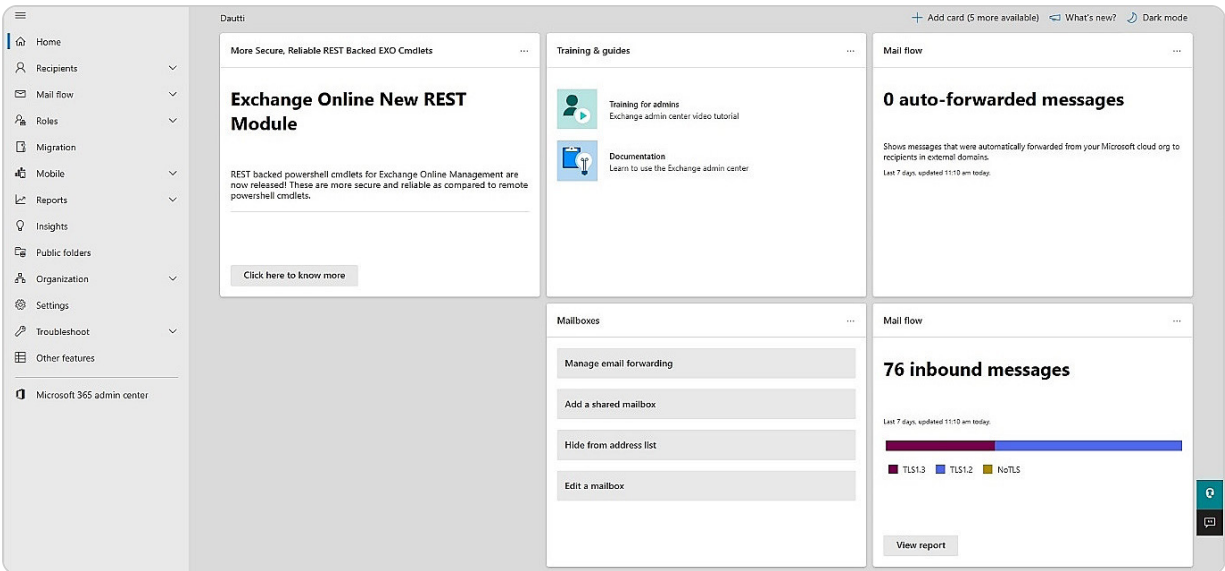


Figure 5.2- Exchange Online admin center

It is crucial to comprehend the distinctions between Exchange Server and Exchange Online to make informed decisions about which solution aligns best with your organization's needs. Exchange Server offers more control and customization options, while Exchange Online provides convenience and scalability by leveraging Microsoft's cloud infrastructure.

With your newfound knowledge of email servers and the available alternatives in Windows Server 2025, you are well prepared to delve into the next section. This section will introduce you to database servers and their access protocols, further enhancing your understanding of server roles and their implementation in Windows Server 2025.

Understanding the database server

A **database server** is a critical component in any IT infrastructure, designed to efficiently manage, store, and retrieve large volumes of data while providing secure access to authorized users. Serving as the backbone of data management, a database server centralizes data, ensuring it can be easily backed up, maintained, and shared across an organization's network. This centralization not only improves data integrity but also enhances collaboration by allowing multiple users to access and interact with the same data in real time.

To configure a Windows-based server as a database server, you can utilize the **SQL Server client/server application**. SQL Server enables the creation, management, and organization of databases, facilitating complex data operations such as querying, reporting, and analysis. For seamless data access and communication, SQL Server relies on several key protocols. These include the following:

- **Data:** The most critical asset managed by the database server. Data is the core element around which all database operations revolve. Without data, the server's functionality and purpose are nullified.
- **Database Application:** This software acts as the intermediary between users and the database server, enabling interactions such as data entry, queries, and report generation.
- **Users:** The individuals or systems that access the database server. Users range from administrators managing the database to end users retrieving data for specific purposes.
- **Open Database Connectivity (ODBC):** A widely-used protocol that allows applications to establish a connection with the database server, regardless of the **Database Management System (DBMS)** in use.
- **Java Database Connectivity (JDBC):** A protocol developed by Sun Microsystems that allows Java applications to connect and interact with the database server, ensuring that Java-based solutions can seamlessly access and manipulate data.
- **Object Linking and Embedding Database (OLEDB):** A Microsoft protocol that provides a standard method for applications to access data stored in a database, facilitating data manipulation and retrieval across various Microsoft and non-Microsoft systems.

NOTE

SQL Server 2022 (<https://www.microsoft.com/en-us/sql-server/sql-server-2022>) is a Microsoft application designed for deploying a database server within an organization's network. To set up a database server using SQL Server 2022, you will need to install and configure the application appropriately.

Understanding these components and protocols is essential for managing and optimizing database server performance, ensuring that data remains accessible, secure, and efficiently handled. With this knowledge, we can now delve into the role and functions of collaboration servers, which play a crucial part in enhancing organizational teamwork and communication.

Understanding the collaboration server

A **collaboration server** facilitates teamwork and information sharing in a digital environment, offering features that support document collaboration, chat, event scheduling, video conferencing, and more. For Windows-based servers, **SharePoint Server** is a robust solution for setting up a collaboration server. SharePoint Server allows you to create, manage, and share websites, libraries, and files within your organization. It leverages networking protocols to provide access to these resources, ensuring seamless collaboration among users.

SharePoint is available in two primary versions: **SharePoint Server** and **SharePoint Online** (see *Figure 5.3*). SharePoint Server, an on-premises solution, is installed on your servers, offering extensive customization and control over your environment. This version allows for deep integration with other on-premises systems and is particularly beneficial for organizations that prioritize complete control over their data and infrastructure.

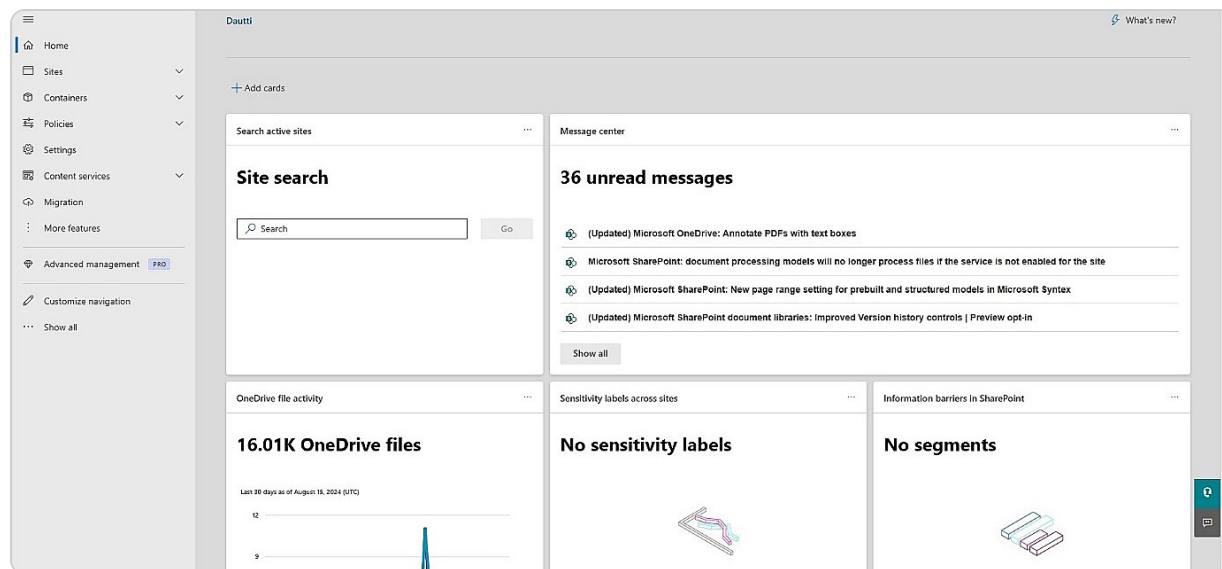


Figure 5.3 – SharePoint Online admin center

On the other hand, SharePoint Online is a cloud-based service that is included in Microsoft 365. It provides many of the same features as SharePoint Server but is hosted and maintained by Microsoft. That means that Microsoft manages updates, security, and scalability, reducing the need for in-house maintenance. SharePoint Online offers greater flexibility, allowing users to access and collaborate on documents from anywhere with an internet connection.

Understanding the differences between SharePoint Server and SharePoint Online is crucial for selecting the right solution for your organization's needs. SharePoint Server offers more control and customization for on-premises deployments, while SharePoint Online provides a more streamlined, cloud-based approach with easier maintenance and scalability.

NOTE

To establish an online collaboration platform within your organization's network, **SharePoint Server Subscription Edition (SharePoint Server SE)** is a viable option. This Microsoft product enables you to create and manage websites, libraries, files, and various resources on your server infrastructure. SharePoint Server SE offers a range of features designed to enhance teamwork and information sharing, including document collaboration, real-time chat, calendar management, video conferencing, and more. To utilize SharePoint Server SE, you will need to install and configure it on your organization's server infrastructure. The platform relies on networking protocols to grant access to its resources, ensuring smooth collaboration across your organization.

Next, we will delve into the concept of monitoring platforms and how they can be utilized to oversee and manage server performance and operations.

Understanding the monitoring server

The **monitoring server**, a cornerstone in the management of an organization's IT environment, provides a centralized view of network health and performance. This system, whether tracking on-premises, cloud-based, or hybrid infrastructure, is essential for administrators. It empowers them to monitor server performance, client/server applications, network services, IT infrastructure, and websites. The use of configurable alerts ensures that any issues are promptly detected and addressed, keeping system administrators informed and in control.

The **System Center Operations Manager (SCOM)** is a robust and comprehensive tool that can transform a server into a powerful monitoring server for Windows-based environments. With its extensive capabilities, SCOM is a reliable choice for managing and monitoring devices and services across an enterprise. However, the System Center suite is not limited to SCOM, and it encompasses several other components that further enhance its functionality:

- **System Center Configuration Manager (SCCM)**: This tool helps manage large groups of Windows-based computers. SCCM is used for software distribution, patch management, and operating system deployment. It streamlines administrative tasks by automating updates and configuration processes.
- **System Center Orchestrator**: Orchestrator focuses on automating and orchestrating IT processes and workflows. It enables the automation of repetitive tasks, integrates various systems, and ensures that complex workflows are executed efficiently.
- **System Center Virtual Machine Manager (SCVMM)**: SCVMM manages virtualized environments, providing a unified approach to deploying and managing VMs. It simplifies virtual infrastructure management and ensures optimal resource utilization.
- **System Center Data Protection Manager (System Center DPM)**: System Center DPM is dedicated to backup and recovery. It ensures that data is consistently backed up and can be quickly restored in case of data loss or system failure.

The System Center components, when used collectively, significantly enhance the monitoring and management of IT infrastructure. They provide comprehensive solutions for network service monitoring, configuration management, automation, VM management, and data protection. By leveraging these tools, organizations can effectively manage their IT environments, maintain system reliability, and improve overall operational efficiency.

NOTE

***System Center 2022** is a robust software suite from Microsoft designed to assist in establishing a monitoring server for your IT infrastructure. This suite features SCOM, a crucial tool for monitoring and managing the health and performance of network devices and applications. Utilizing System Center 2022 allows you to oversee various network metrics, ensuring optimal system performance and the rapid identification and resolution of potential issues. Additionally, Microsoft has announced that System Center 2025 is slated , released on November 1st, 2024, promising further advancements in monitoring and management capabilities.*

In the following section, we will delve into threat management servers and their role in securing IT environments.

Understanding the Data Protection Server

A **data protection server** plays a critical role in ensuring an organization's continuity and resilience by providing robust data backup and recovery solutions. This server is integral to implementing a comprehensive **business continuity and disaster recovery (BCDR)** strategy. To establish a data protection server on a Windows-based system, you would utilize the System Center DPM application. DPM offers a suite of features that cater to diverse backup and recovery needs.

With DPM, administrators can perform application-aware backups, which are essential for protecting complex applications, such as Exchange Server, SQL Server, and SharePoint Server. That means backups are not only taken of the files but also the application's metadata and configuration, ensuring that complete and functional restorations are possible. Additionally, DPM facilitates the backup of individual files, folders, and entire volumes, as well as system state data, which is crucial for recovering the operating system and its settings.

Moreover, DPM supports the backup of VMs hosted on **Hyper-V**, encompassing both Windows and Linux environments. This functionality is significant for virtualized environments where the integrity of VMs must be preserved for operational continuity. By leveraging DPM, organizations can ensure that their data is consistently protected against loss or corruption and that recovery processes are streamlined and effective.

NOTE

*To set up a data protection server for your organization's network, one practical approach is to utilize **Microsoft's System Center 2022**. This software suite features the DPM, a client/server application that must be installed and configured on a Windows-based server. DPM provides comprehensive backup and recovery capabilities, allowing you to safeguard and restore data across various applications and systems. It supports backups for Exchange Server, SQL Server, and SharePoint Server, as well as files, folders, volumes, system states, and VMs hosted on Hyper-V for both Windows and Linux environments.*

In this section, you have explored several application servers utilized in on-premises environments, providing insights into various client/server applications. This overview has equipped you with knowledge about their key features, components, protocols, and functionalities. Moving forward, the next section will shift focus to the different types of web services, expanding your understanding of their roles and implementations within IT infrastructure.

Configuring web services and their roles in Windows Server 2025

A web service is a standardized framework that enables different software applications, often operating on diverse platforms, to communicate and interact with each other efficiently. This interoperability is achieved through the use of **Extensible Markup Language (XML)**-based formats and protocols, such as **Simple Object Access Protocol (SOAP)**, **Web Services Description Language**

(WSDL), and **Universal Description, Discovery, and Integration (UDDI)**. These technologies facilitate the exchange of data and the invocation of functionalities over a network, allowing applications to work together seamlessly, regardless of their underlying systems.

Web services can be broadly classified into two types:

- **RESTful web services:** These are built on the **Representational State Transfer (REST)** architecture, which uses **Hypertext Transfer Protocol (HTTP)** methods such as **GET**, **POST**, **PUT**, and **DELETE**, along with **Uniform Resource Identifiers (URIs)**, to access and manipulate resources. This approach is known for its simplicity and scalability, making it well suited for web-based interactions.
- **SOAP-based web services:** These rely on SOAP, which uses structured XML messages and envelopes to communicate with a service endpoint. SOAP is more rigid and standardized, offering built-in error handling and security features, which makes it ideal for enterprise-level applications where reliability and security are paramount.

Next, we will delve into IIS, a crucial component for hosting web services and applications.

IIS Explained

IIS is Microsoft's robust and versatile web server platform designed to deliver scalable, manageable, and reliable web applications. IIS facilitates communication between the browser and the web server through a variety of protocols, including HTTP, **HTTP Secure (HTTPS)**, **File Transfer Protocol (FTP)**, **FTP Secure (FTPS)**, SMTP, and **Network News Transfer Protocol (NNTP)**. Additionally, Microsoft has introduced **Active Server Pages (ASP)**, a server-side scripting technology that enables the creation of dynamic web content.

With the release of **IIS version 10**, Microsoft has significantly enhanced the security and performance of the platform. This version supports longer script execution times and introduces HTTP/2 support. Moreover, in January 2020, Microsoft launched a new Chromium-based browser called **Microsoft Edge**, further complementing the IIS ecosystem. IIS 10 also brought new features in Windows Server 2025, such as improved server-side cipher suite negotiation for HTTP/3, IIS administration through PowerShell cmdlets, support for wildcard host headers, the ability to run IIS on Nano Server and within containers, and a user interface for managing **HTTPS Strict Transport Security (HSTS)**. These enhancements have collectively elevated the performance and security of IIS.

Adding IIS on Windows Server 2025

To set up a web server on Windows Server 2025, IIS must be added as a server role as follows:

1. Log in to your Windows Server 2025 with an admin account. Open **Server Manager** from the Start menu. In Server Manager, click **Manage** in the upper-right corner, then select **Add Roles and Features**. That starts **Add Roles and Features Wizard**.
2. Click **Next** on the **Before You Begin** page. Choose **Role-based or feature-based installation** and click **Next**. On the **Server Selection** page, select the local server and click **Next**.
3. Check the box for **Web Server (IIS)** on the **Server Roles** page, as shown in *Figure 5.4*. A dialog will pop up—click **Add Features**, then click **Next**.

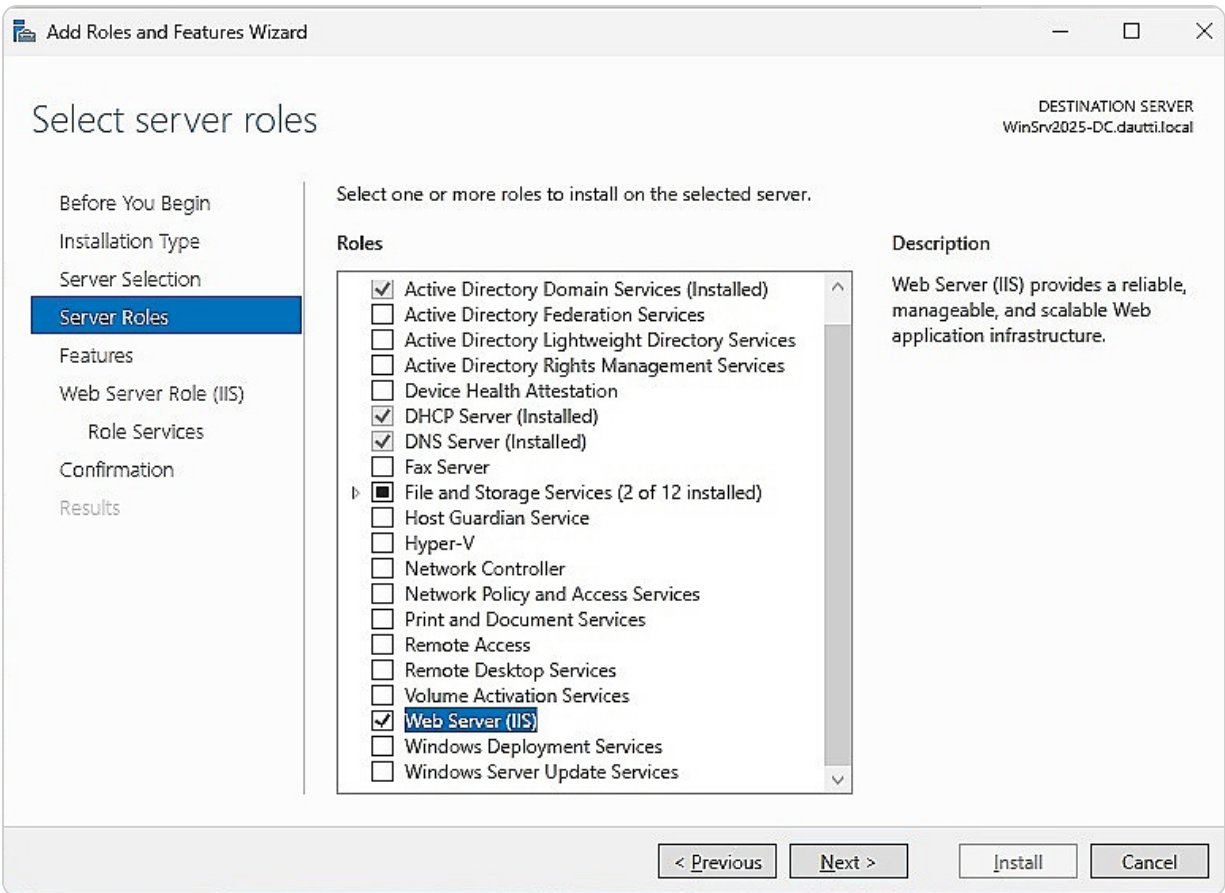


Figure 5.4- Adding Web Server (IIS) on Windows Server 2025

4. On the **Features** page, you can add extra features, but the default options are enough for most setups. Click **Next**. Review the **Web Server (IIS)** role overview and click **Next**.
5. On the **Role Services** page, you can add more IIS features, such as FTP Server or security options, if needed. Click **Next** when done.
6. Review your selections on the **Confirmation** page. You can choose to restart the server if needed automatically. Click **Install** to start the process.
7. The installation will begin, and you can watch its progress. Once it's finished, click **Close** to exit the wizard.

Once installed, IIS Manager serves as the administrative console for managing the web server. Accessible through Server Manager, Windows Administrative Tools, or by running the `inetmgr` command in the **Run** dialog box, IIS Manager allows administrators to efficiently manage their web applications, as depicted in *Figure 5.5*.

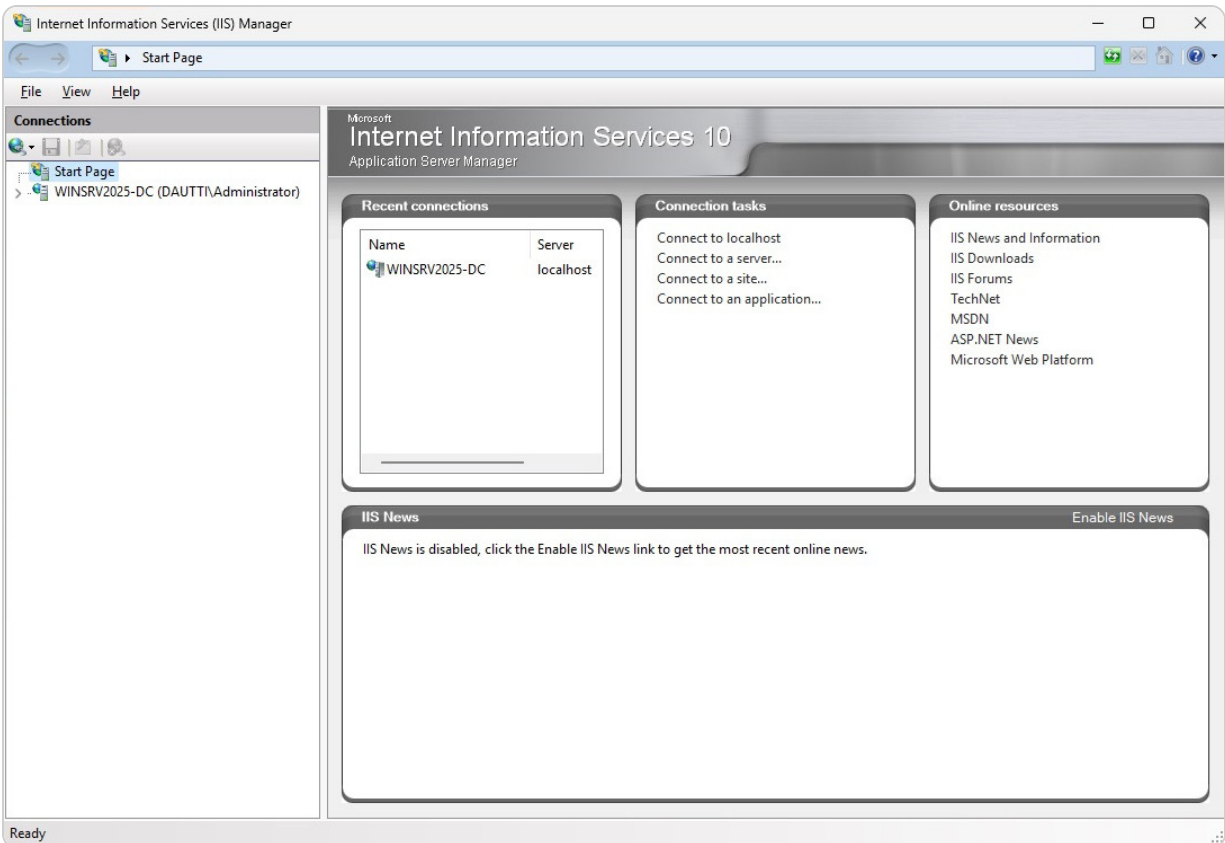


Figure 5.5- IIS Manager in Windows Server 2025

Next, we will explore the fundamentals of the **World Wide Web (WWW)**.

WWW overview

The **WWW** is a global information system that operates over the internet. It enables users to access and interact with web pages via HTTP. Web pages, typically written in **Hypertext Markup Language (HTML)**, can include text, images, videos, and other multimedia elements, making the web a rich and dynamic platform for information sharing.

The WWW was first proposed by Tim Berners-Lee in 1989 while working at CERN, and it quickly evolved into the vast network we use today. The core technologies behind the web include HTML for structuring content, HTTP for transmitting data between servers and clients, and **uniform resource locators (URLs)** for addressing resources on the web.

Initially, the web was a static medium, but it has since advanced to support interactive applications, real-time communication, and complex web services. As seen in *Figure 5.6*, which illustrates a sample web page and its underlying HTML code, the web has grown from a basic document-sharing platform into a fundamental part of modern IT infrastructure, supporting everything from e-commerce to social networking.

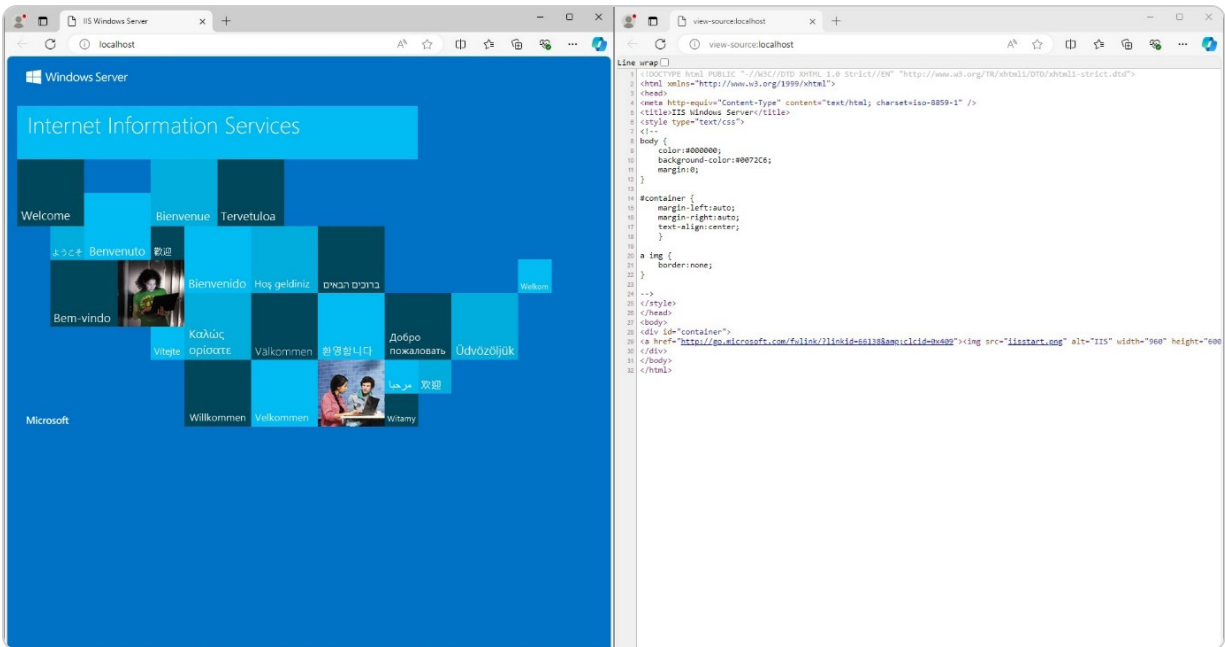


Figure 5.6- A web page and its source code

In the following subsection, we will delve into the FTP, another essential tool for managing files across the internet.

Understanding an FTP

FTP serves as a fundamental method for transferring files across the internet, offering a secure and efficient mechanism for exchanging data between computers. Initially developed in the early 1970s, FTP has remained a vital tool in networked environments due to its simplicity and reliability. It is widely employed for a range of tasks, including transmitting corporate data within internal networks, managing website content, and facilitating the upload and download of files to and from web servers.

FTP operates on a **client/server model**, utilizing two distinct ports to manage its operations. Port 21 is designated for establishing the control connection, allowing commands and responses to be exchanged between the client and server. Once this connection is established, port 20 is used for the actual data transfer, handling the transmission of files between the systems. This **dual-port system**, a testament to the protocol's efficiency, ensures that command and data traffic are kept separate, enhancing the protocol's performance.

Setting up an FTP Server on Windows Server 2025 involves several key steps. First, you must install the Web Server (IIS) role on your server, which lays the foundation for hosting various web and FTP services. After that, you add the FTP Server role service under the Web Server role, as depicted in *Figure 5.7*. This configuration enables the server to handle FTP connections, providing administrators with a robust tool for managing file transfers in a controlled, secure, and scalable manner.

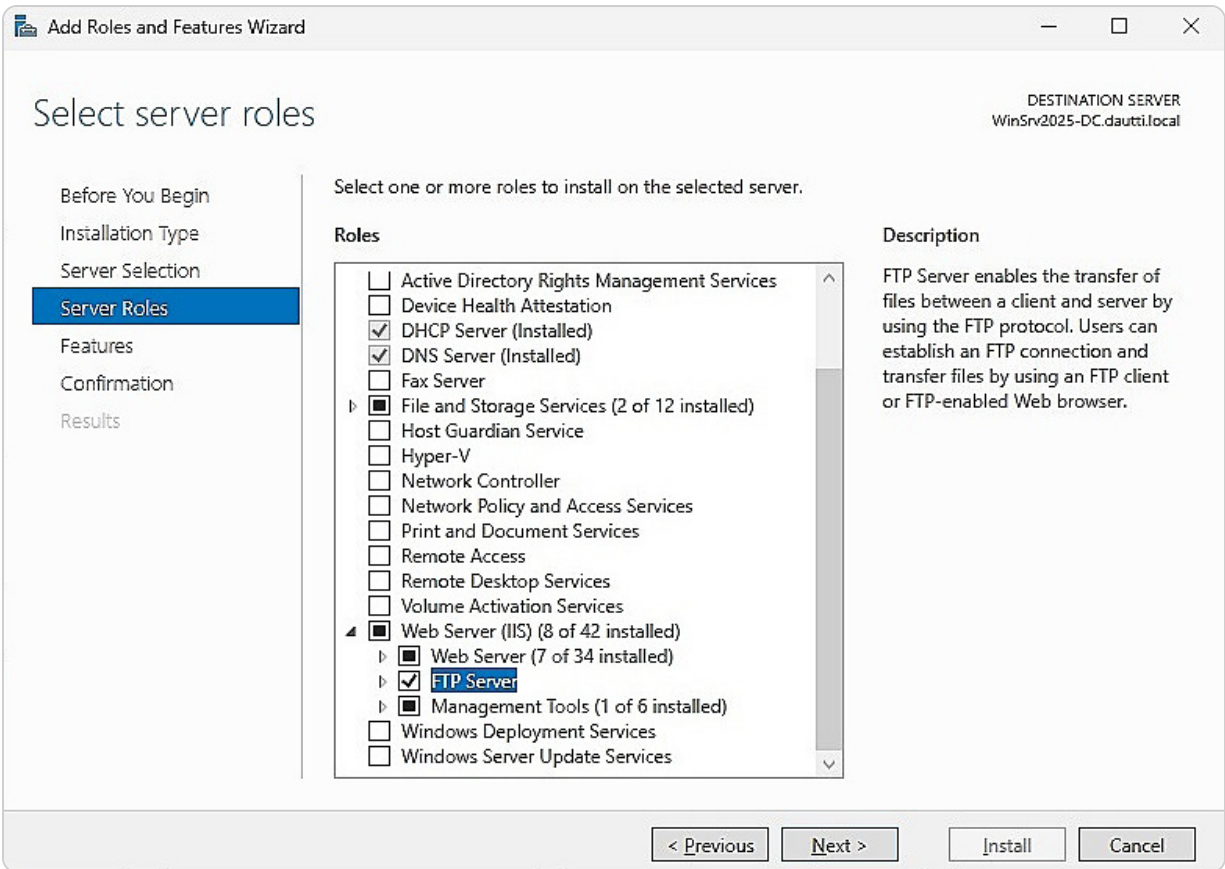


Figure 5.7- Adding an FTP Server as a role service in Windows Server 2025

Following this, we will delve into the concept of worker processes, exploring their pivotal role in managing server resources and ensuring the smooth operation of web and FTP services.

Understanding how to access and configure these processes is not just crucial, but it's your key responsibility for maintaining optimal server performance.

Worker processes and how to access them?

In IIS, each application pool is associated with its dedicated worker process. This **worker process** is a crucial component that handles the execution of web applications contained within that pool. By assigning a separate worker process to each application pool, IIS ensures isolation between different pools. This isolation is advantageous because it means that if an issue arises within one web application, it remains contained within that pool and does not interfere with the functioning of web applications in other pools. To manage and configure the worker process settings for an application pool, you can access IIS Manager, select the **Application Pools** node, and then click on **Advanced Settings** in the **Actions** pane on the right side of the interface, as depicted in *Figure 5.8*. That allows you to customize various parameters related to the worker processes, such as process recycling and idle timeout settings, to optimize performance and reliability.

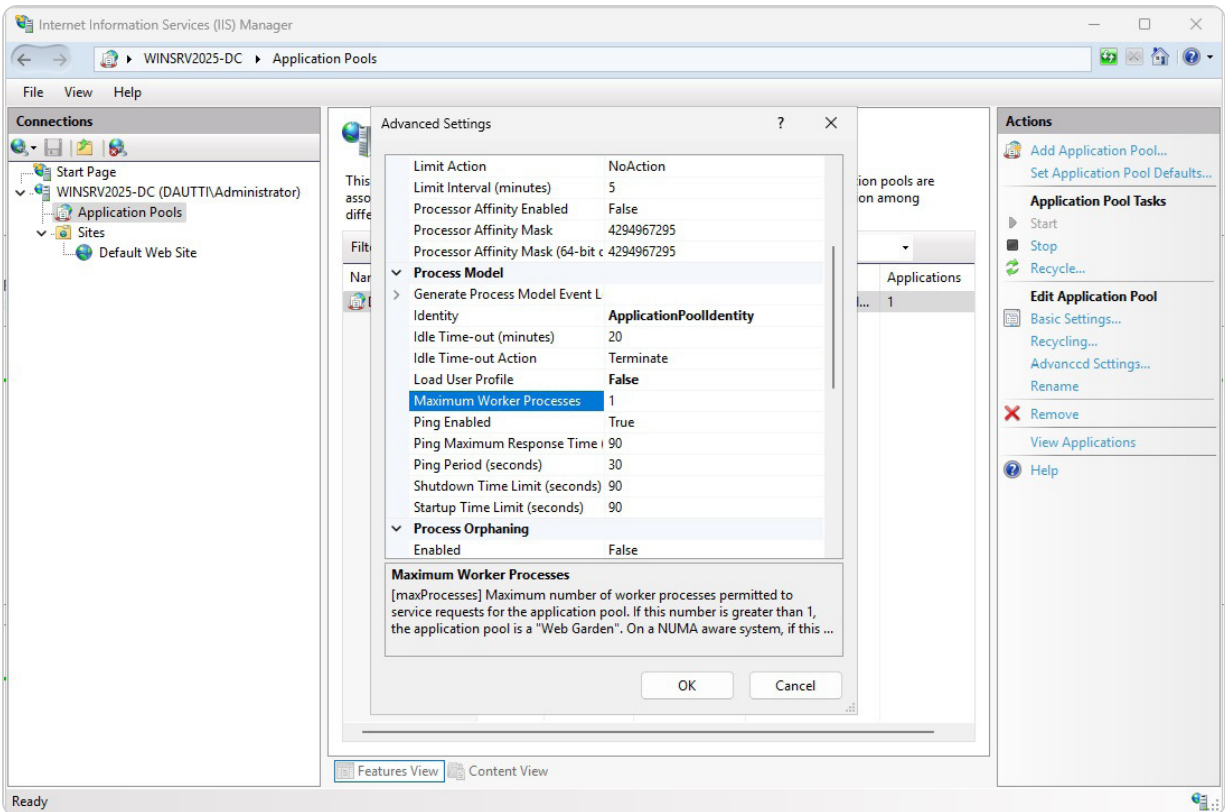


Figure 5.8- An application's pool worker process in IIS

In the subsequent section, we will delve into the procedures for installing additional features for IIS on Windows Server, which will enhance its capabilities and support a broader range of functionalities.

Installing more features for IIS

When configuring IIS on Windows Server, you start by installing the Web Server (IIS) role. During this installation, you can choose from various role services that align with your server's needs. This selection is made during a subsequent step in the installation wizard, as detailed in *Figure 5.7* of the *Understanding an FTP* subsection. These role services encompass a range of functionalities, such as FTP server capabilities, management tools, and additional modules that extend the server's features.

If you need to add or modify these Role Services after the initial installation, you can do so through Server Manager. By launching **Add Roles and Features Wizard**, you can install or adjust additional IIS features and services. This process allows you to customize your IIS environment further, adding components such as URL rewrite modules, security features, or additional management tools as required.

In the next section, we will explore the concepts of sites and websites, providing a detailed explanation of their definitions, roles, and how they function within the IIS framework.

Sites overview

A **site**, also known as a **website**, is a collection of web pages designed to present content over the internet using web services. These web pages are typically built using HTML, which provides the fundamental structure and layout of the content. However, to create more engaging and interactive experiences, additional scripting languages, such as **JavaScript**, or server-side languages, such as **PHP** or **ASP.NET**, are often utilized. This combination of technologies allows for dynamic content and user interactions.

When you install the Web Server (IIS) role on Windows Server 2025, the setup process automatically creates a default website that serves a single introductory web page. This default site serves as a placeholder and a starting point for further configuration, as illustrated in *Figure 5.9*:

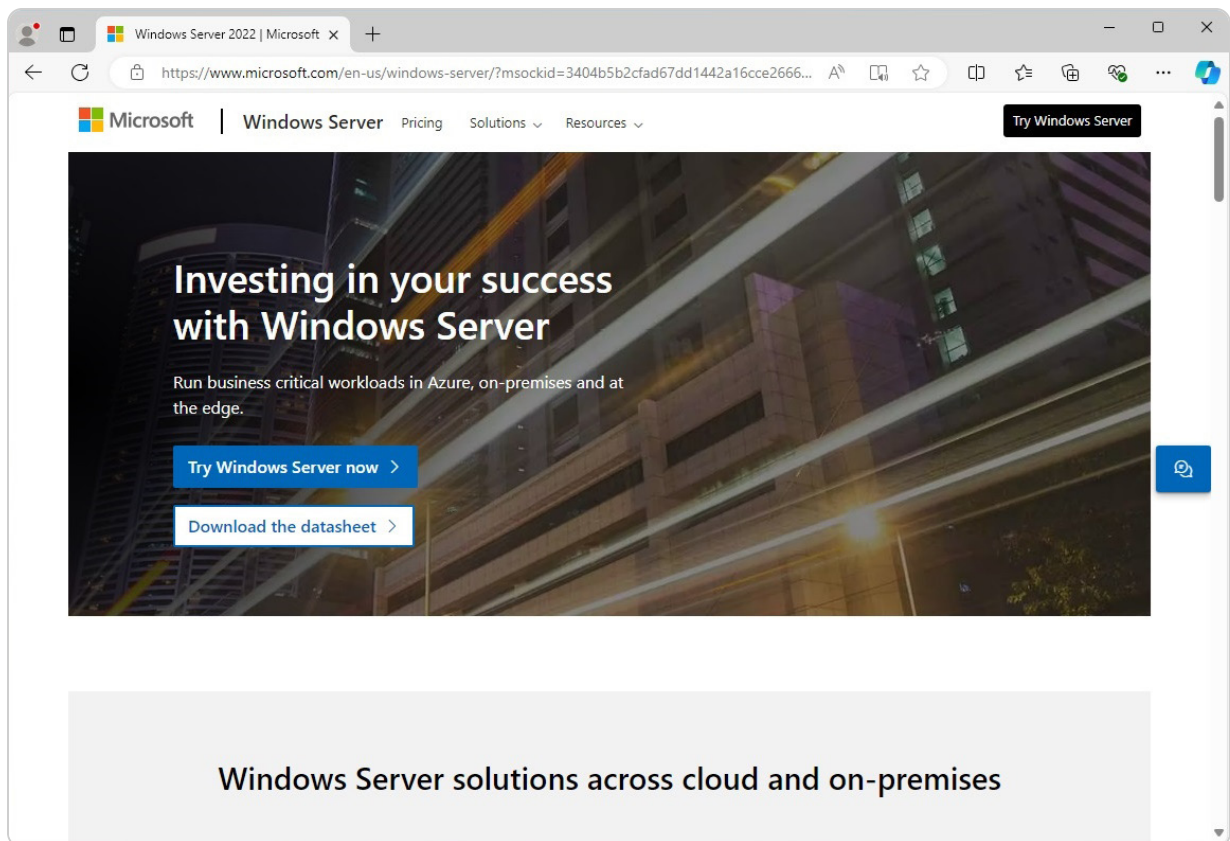


Figure 5.9- Windows Server's website powered by Microsoft

To expand your server's capabilities, you can create additional websites within the same server environment using IIS Manager. That is done by navigating to the **Sites** node, right-clicking it, and selecting **Add Website**, which is shown in *Figure 5.10*:

Add Website

Site name: Intranet

Application pool: Intranet Select...

Content Directory

Physical path: C:\Intranet ...

Pass-through authentication

Connect as... Test Settings...

Binding

Type: http IP address: All Unassigned Port: 80

Host name:

Example: www.contoso.com or marketing.contoso.com

☒ Start Website immediately

OK Cancel

Figure 5.10 – Adding a website via IIS Manager in Windows Server 2025

This feature allows you to host multiple sites on a single server, each with its unique domain or path, and manage them efficiently through IIS Manager.

In the following subsection, we will explore the concept of software ports, their role in network communications, and how they facilitate the interaction between web services and clients.

Ports overview

Ports are crucial components in computer networking, categorized into hardware ports and software ports. A **hardware port** is a physical interface on a computer or electronic device that enables data transfer and communication. Examples include USB ports for connecting peripherals, Ethernet ports for network connectivity, and HDMI ports for video output, all of which are fundamental for seamless interaction between devices.

In contrast, a **software port**, or **application port**, is a virtual endpoint used by software applications and services to manage network traffic. These logical identifiers help direct data to the appropriate process or application running on a server or network device. Software ports are divided into three main types:

- **Well-Known Ports:** Ranging from 0 to 1023, these ports are assigned to widely used protocols and services by the **Internet Assigned Numbers Authority (IANA)**. They are standardized for common services to ensure consistency across different systems. For instance, port 80 is used for HTTP traffic, port 443 for HTTPS traffic, port 21 for FTP, and port 25 for SMTP. These ports are essential for fundamental web services, secure communications, file transfers, and email transmissions.
- **Registered Ports:** These ports, spanning from 1024 to 49151, are used by software applications that are not as universally recognized as those using well-known ports. Registered with IANA, these ports avoid conflicts and ensure dedicated access for specific applications. For example, MySQL databases use port 3306, and PostgreSQL databases use port 5432.
- **Dynamic or Private Ports:** Covering the range from 49152 to 65535, these ports are typically used for ephemeral purposes. They are assigned temporarily for short-lived communication sessions and are not permanently associated with any specific service. These ports are commonly used for client-side applications to establish temporary connections with server-side services.

The following table provides a summary of some common application ports, their protocols, and transport methods:

Protocol	Port	Transportation protocol
FTP	21	TCP
SSH	22	TCP
Telnet	23	TCP
SMTP	25	TCP
HTTP	80	TCP and UDP
POP3	110	TCP
NNTP	119	TCP

Protocol	Port	Transportation protocol
NTP	123	TCP
IMAP4	143	TCP
HTTPS	443	TCP

Table 5.1- The well-known application ports

Each port and protocol combination plays a specific role in network communications. For instance, FTP uses port 21 for file transfers, **Secure Shell (SSH)** on port 22 facilitates secure remote access, and HTTP on port 80 manages standard web traffic.

In the following section, we will delve into SSL technology, which provides an additional layer of security by encrypting data transmitted over networks.

What is SSL?

SSL is a crucial technology for securing data exchanged between a web server and a web browser. By encrypting the data transmitted, SSL ensures that any communication between these two entities remains confidential and integral. When a web browser connects to a website that employs SSL, it uses the HTTPS protocol, which operates over port 443, to initiate a secure connection.

This secure connection is established through the use of digital certificates, which are cryptographic documents issued by trusted **certificate authorities (CAs)**. These certificates authenticate the identity of the website and facilitate the establishment of a secure, encrypted communication channel. During this process, the browser and the server use the certificate to agree on a shared secret key. This key is then employed to encrypt all data transmitted between them, preventing unauthorized parties from intercepting or deciphering the information, thereby ensuring a secure and protected online experience.

SSL not only ensures the confidentiality of the data but also verifies the legitimacy of the website, protecting users from potential cyber threats, such as phishing attacks. By employing encryption, SSL helps to maintain data integrity, ensuring that the data remains reliable and unchanged during transmission. This secure exchange is visually represented in *Figure 5.11*:

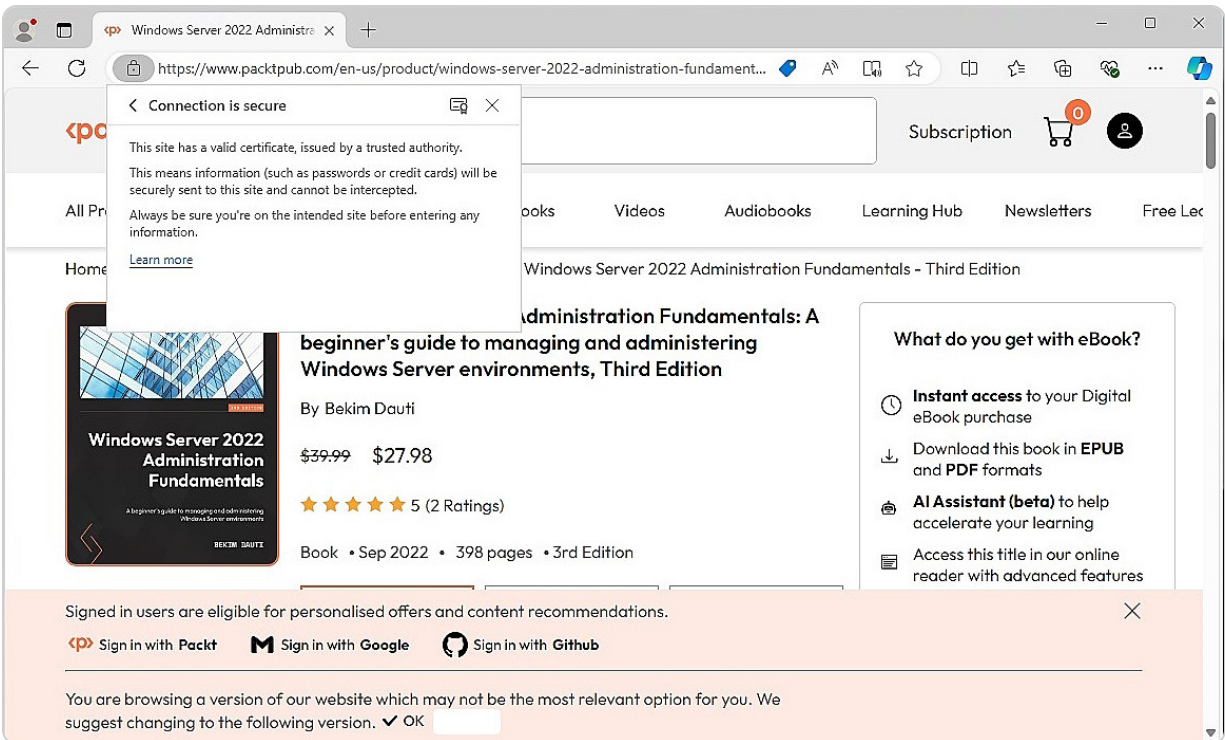


Figure 5.11- Secure communication between browser and website

NOTE

Transport Layer Security (TLS) improves upon SSL by addressing its limitations with enhanced security features. Unlike SSL, TLS supports more robust encryption algorithms and more secure cryptographic practices, offering better protection against vulnerabilities and attacks. TLS also refines the handshake process and certificate validation methods, ensuring a more secure and resilient connection. These advancements make TLS a superior choice for safeguarding data transmitted over networks, effectively overcoming the shortcomings of its predecessor.

In the following subsection, we will explore the evolution of SSL into TLS and its role in further enhancing data security.

How do certificates work?

Digital certificates are crucial for establishing secure communication over the internet, particularly between a website and a browser. These certificates, issued by a trusted entity known as a CA, authenticate the website's identity and enable encrypted data exchange. As illustrated in *Figure 5.12*, the CA is responsible for validating and issuing certificates, ensuring that they are trustworthy and accurate.

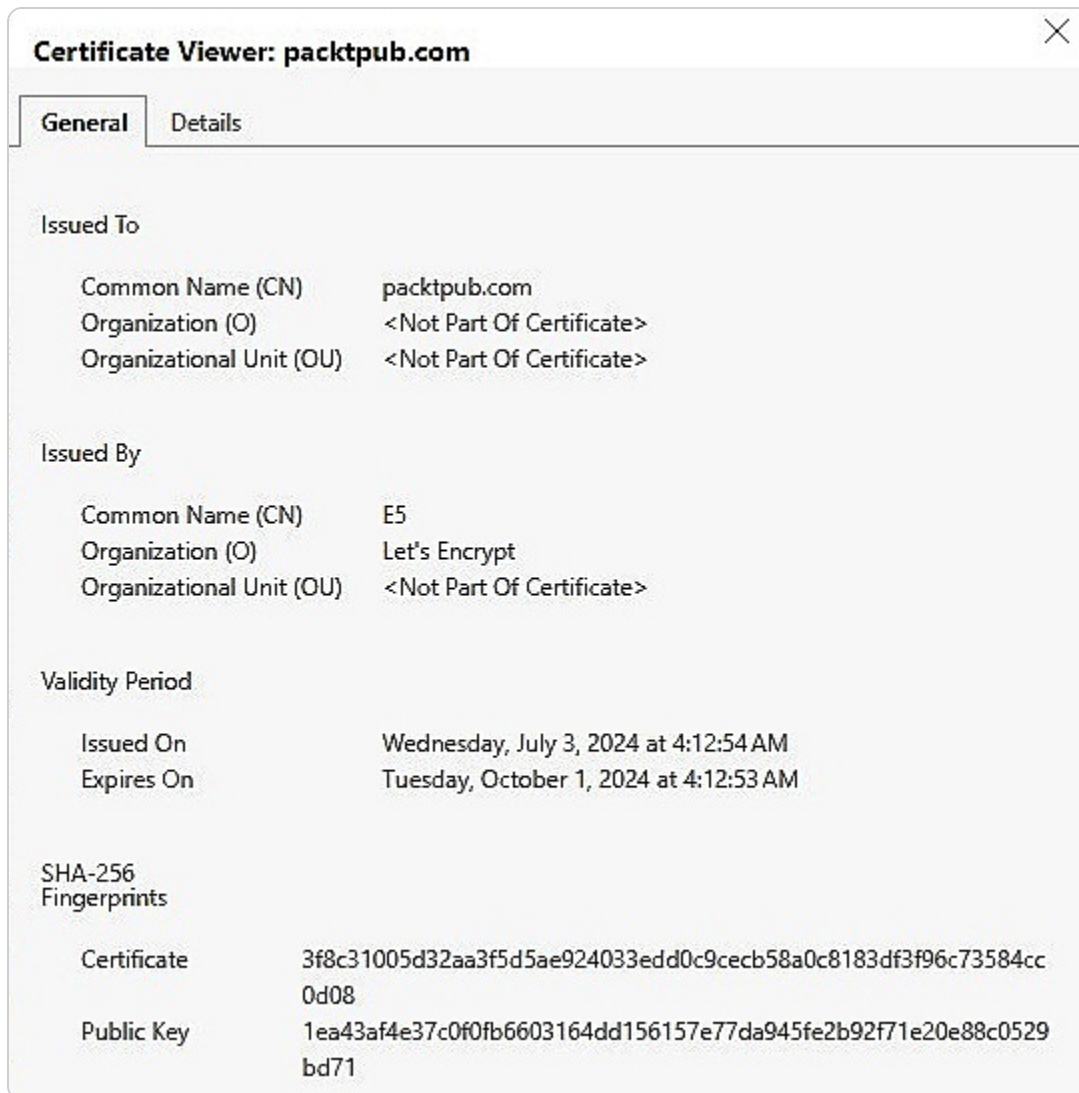


Figure 5.12- Certificate issued by a CA

Digital certificates operate within the framework of **public key infrastructure (PKI)**. PKI is a system designed to manage digital keys and certificates, providing a robust method for verifying the ownership of public keys. Each certificate contains a public key and information about the certificate holder, such as the organization's name and address. The CA's role is to confirm that the public key contained in the certificate indeed belongs to the entity it claims to represent. This validation process helps prevent fraudulent activities and ensures that the data exchanged remains secure.

TLS has evolved from SSL to address its limitations and enhance security. TLS offers more robust encryption algorithms, improved cryptographic practices, and better mechanisms for key negotiation and certificate validation. These advancements resolve SSL's vulnerabilities, providing a more secure and resilient framework for protecting data transmitted over networks.

When a browser connects to a website that uses SSL/TLS, the digital certificate is used to establish a secure connection. The website's server and the browser use the certificate to agree on a shared encryption key. This key is then used to encrypt and decrypt the data exchanged, ensuring that sensitive information remains confidential and protected from unauthorized access.

In addition to encryption, digital certificates also provide data integrity and authentication. They ensure that the data sent between the website and the browser is not tampered with during transmission and that the identity of the website is genuine.

NOTE

For a more in-depth understanding of PKI, you can refer to the comprehensive resources available at Oracle's PKI documentation: https://docs.oracle.com/cd/B10501_01/network.920/a96582/pki.html. This source provides detailed information on the principles, components, and implementation of PKI, offering valuable insights for enhancing your knowledge of secure data management and encryption technologies.

The following section will introduce the role of remote access in Windows Server 2025 and detail how it supports secure remote connectivity and access management.

Setting up remote access roles and their functionalities

Remote Access in Windows Server 2025 is an essential role that provides users with the flexibility to access network resources from virtually any location and at any time while also allowing administrators to manage and secure these remote connections efficiently. This capability is precious in today's increasingly remote work environments, where maintaining secure and reliable access to corporate networks is critical. The Remote Access role integrates several key technologies that support different types of network connections, each tailored to meet specific remote access requirements:

- **DirectAccess:** This is one of the core components of Remote Access, which was initially introduced in Windows Server 2008 R2. It offers a seamless and secure connection to corporate networks without the need for a traditional VPN. By encrypting communications between DirectAccess clients and servers using IPsec, DirectAccess ensures that data remains protected as it travels across the internet. Moreover, it utilizes IPv6 over IPv4 tunneling to facilitate connectivity to the intranet, making it easier for users to access internal resources securely from remote locations.
- **The Routing and Remote Access Service (RRAS):** This succeeded the **Remote Access Service (RAS)** from the Windows NT era, which was launched in Windows 2000. RRAS is a versatile service that combines the capabilities of VPN and dial-up connections, enabling the creation of secure links between remote sites. Additionally, RRAS supports routing traffic between different sub-networks, making it an essential tool for managing complex network environments and ensuring that remote users can connect to the resources they need.
- **Web Application Proxy:** This is available on Windows Server 2025. Acting as a reverse proxy, the Web Application Proxy enables secure access to web applications hosted on the intranet from external networks. That is achieved through integration with **Active Directory Federation Services (AD FS)**, which authenticates corporate users and ensures that only authorized individuals can access sensitive internal applications via an extranet. This technology is beneficial for organizations that need to provide secure remote access to web-based resources without compromising security.

Setting up a remote access server in Windows Server 2025 involves adding the Remote Access role to the server, as depicted in *Figure 5.13*. This role configuration is the first step in enabling a range of remote access technologies that enhance both the connectivity and security of remote users. By implementing these technologies, organizations can ensure that their employees have the tools they need to work effectively from any location. At the same time, IT administrators can maintain control over the security and management of these connections.

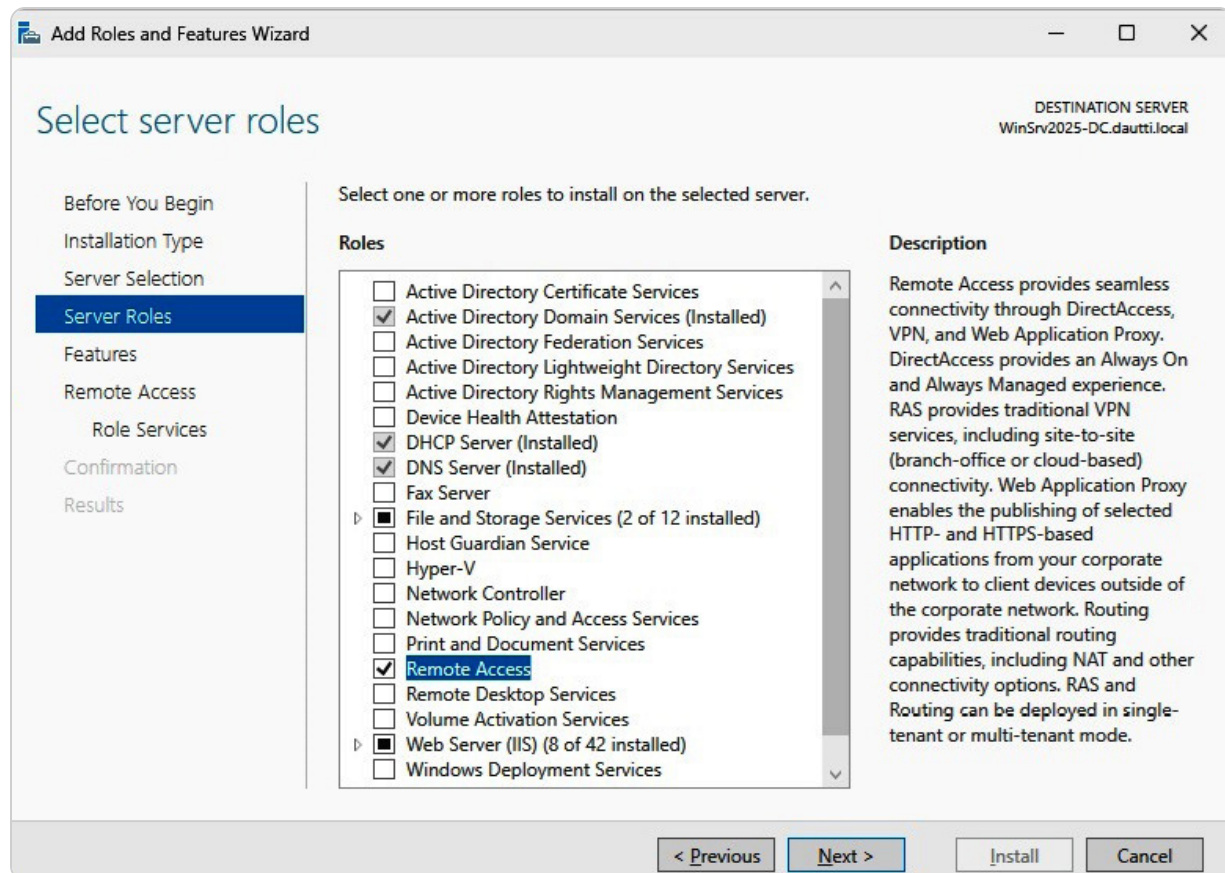


Figure 5.13- Adding the Remote Access role in Windows Server 2025

In the upcoming subsection, we will delve into the **Remote Assistance feature**, providing detailed instructions on how to enable and configure this feature on your server, further enhancing your organization's remote support capabilities.

How to use Remote Assistance

Remote Assistance in Windows Server 2025 is a valuable feature that enables a trusted individual, typically an IT support professional (the initiator), to remotely view and control a user's desktop (the invitee) to assist with diagnosing and resolving technical issues. This capability is particularly beneficial for organizations with remote employees or distributed networks, as it allows support

personnel to provide real-time help without being physically present, thereby reducing downtime and increasing efficiency. The process begins when the invitee requests assistance, granting the initiator permission to access their system. That ensures that the invitee remains in control and can oversee the troubleshooting process, enhancing both security and user comfort. To activate the Remote Assistance feature on a server, administrators need to use **Add Roles and Features Wizard**, a straightforward process illustrated in *Figure 5.14*.

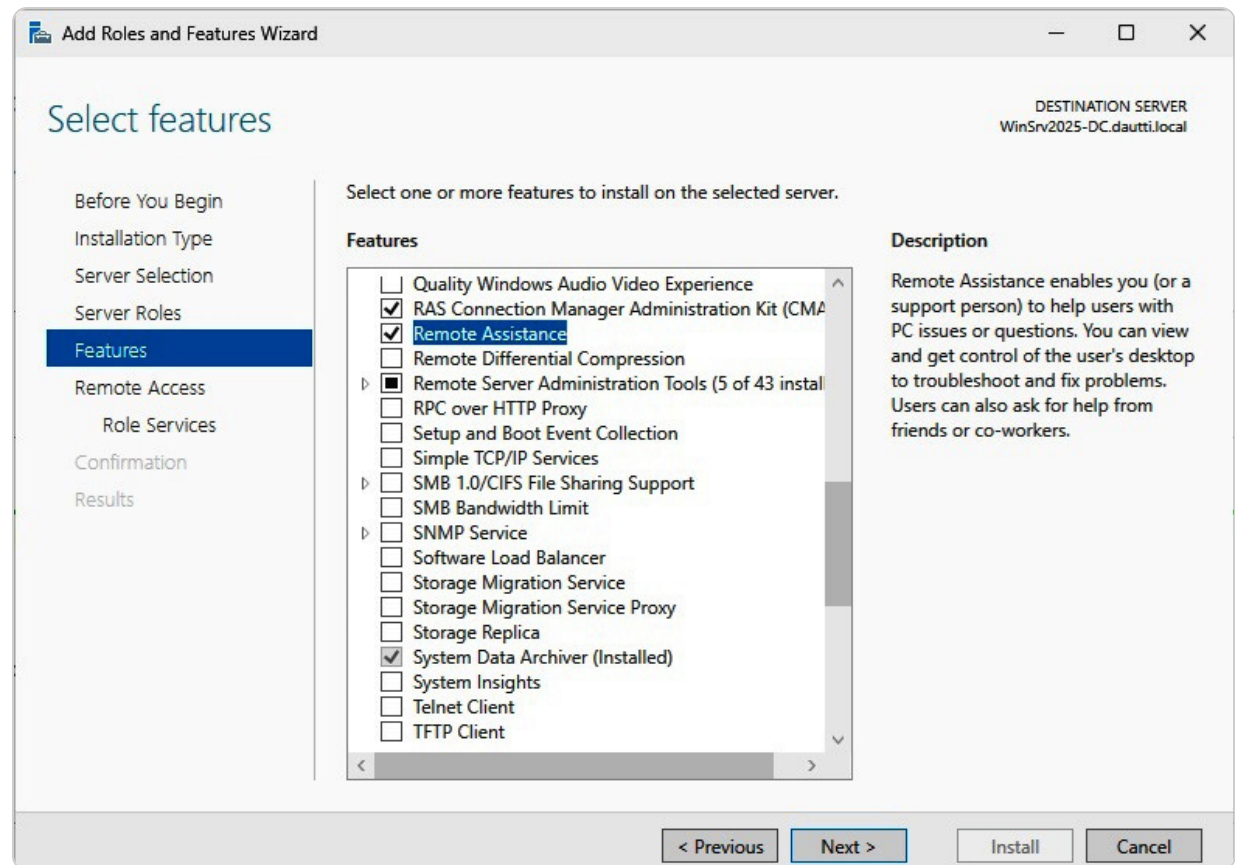


Figure 5.14- Adding the RSAT feature in Windows Server 2025

In the upcoming subsection, we will explore the **Remote Server Administration Tools (RSAT)** feature, a robust set of tools that further enhances remote management capabilities by allowing administrators to manage roles and features on other servers from a single workstation, streamlining the management of Windows Server environments.

How does RSAT work?

RSAT in Windows Server 2025 empowers system administrators to remotely manage server roles and features across other servers running Windows Server 2025. RSAT provides both graphical and command-line interfaces, giving administrators flexibility in how they interact with their server

environments. Additionally, RSAT is compatible with client computers running Windows 10 or 11, enabling centralized management from a broader range of devices.

Administrators can use **Add Roles and Features Wizard**, depicted in *Figure 5.15*, to enable the RSAT feature. Once activated, RSAT facilitates a wide array of administrative tasks, from configuring Active Directory settings to managing Group Policy and monitoring server performance.

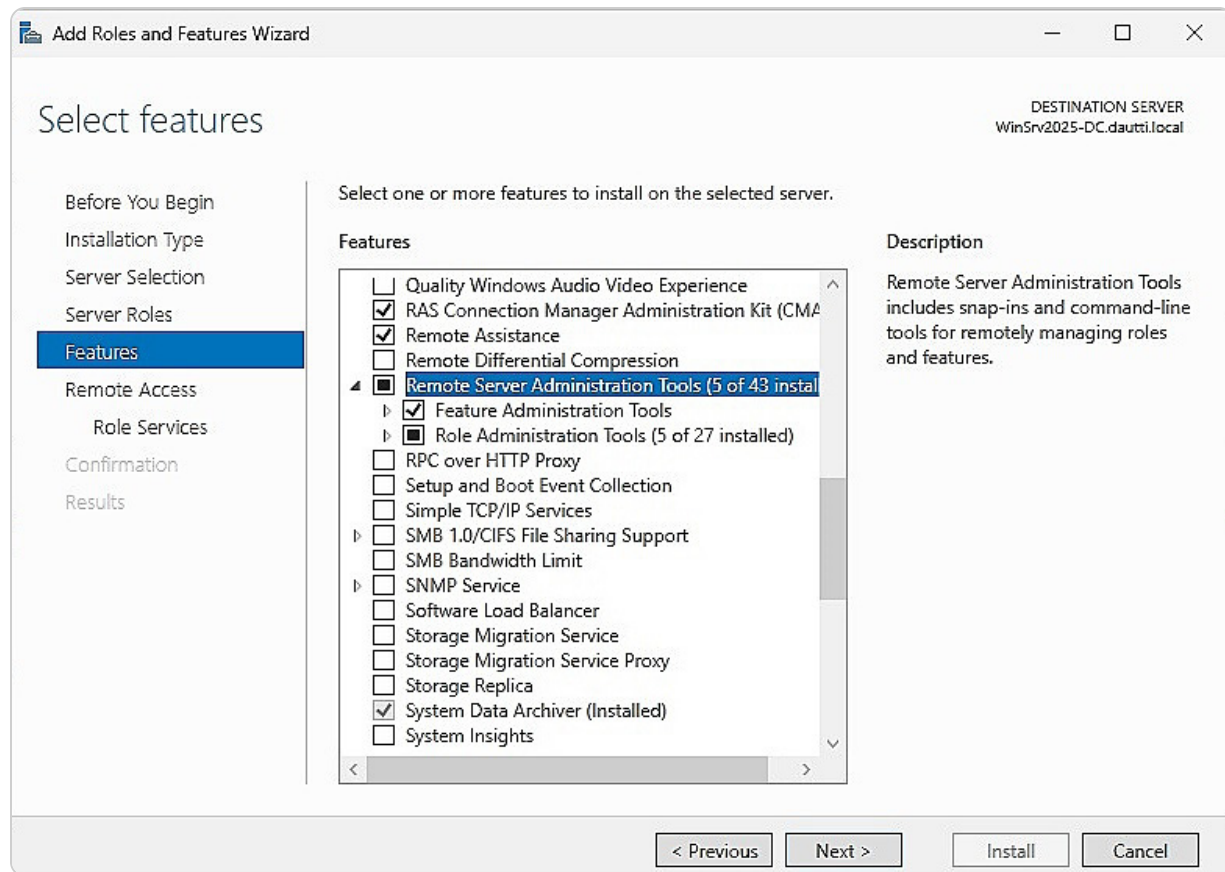


Figure 5.15- Adding the RSAT feature in Windows Server 2025

In comparison, **Windows Admin Center** offers a modern, browser-based alternative to RSAT. While RSAT provides a more traditional approach to remote server management, Windows Admin Center integrates multiple management tools into a single, unified interface, streamlining the management experience. This tool allows administrators to manage servers, clusters, hyper-converged infrastructure, and Windows 10 or 11 devices with ease, all through a web-based interface that can be accessed from anywhere.

In the next section, we will discuss the **Remote Desktop Services (RDS)** server and its configuration.

Explaining RDS

RDS plays a key role in Windows Servers, enabling users to access the graphical interfaces of computers and applications from remote locations, whether they are on the same network or connecting over the internet. Previously known as **Terminal Services (TS)** until its rebranding in Windows Server 2008 R2, RDS provides robust functionality, allowing users to run virtualized applications directly on their desktops. This feature is precious in enterprise environments where centralized management of applications and desktops is critical.

By default, Windows Server permits two concurrent Remote Desktop sessions without requiring additional licensing. This limitation is sufficient for basic administrative tasks, but when more than two users need to connect simultaneously, an RDS licensing server must be configured to manage additional licenses. To enable and adequately configure RDS on a Windows Server 2025 environment, the RDS role must be added through the server's management interface, as depicted in *Figure 5.16*.

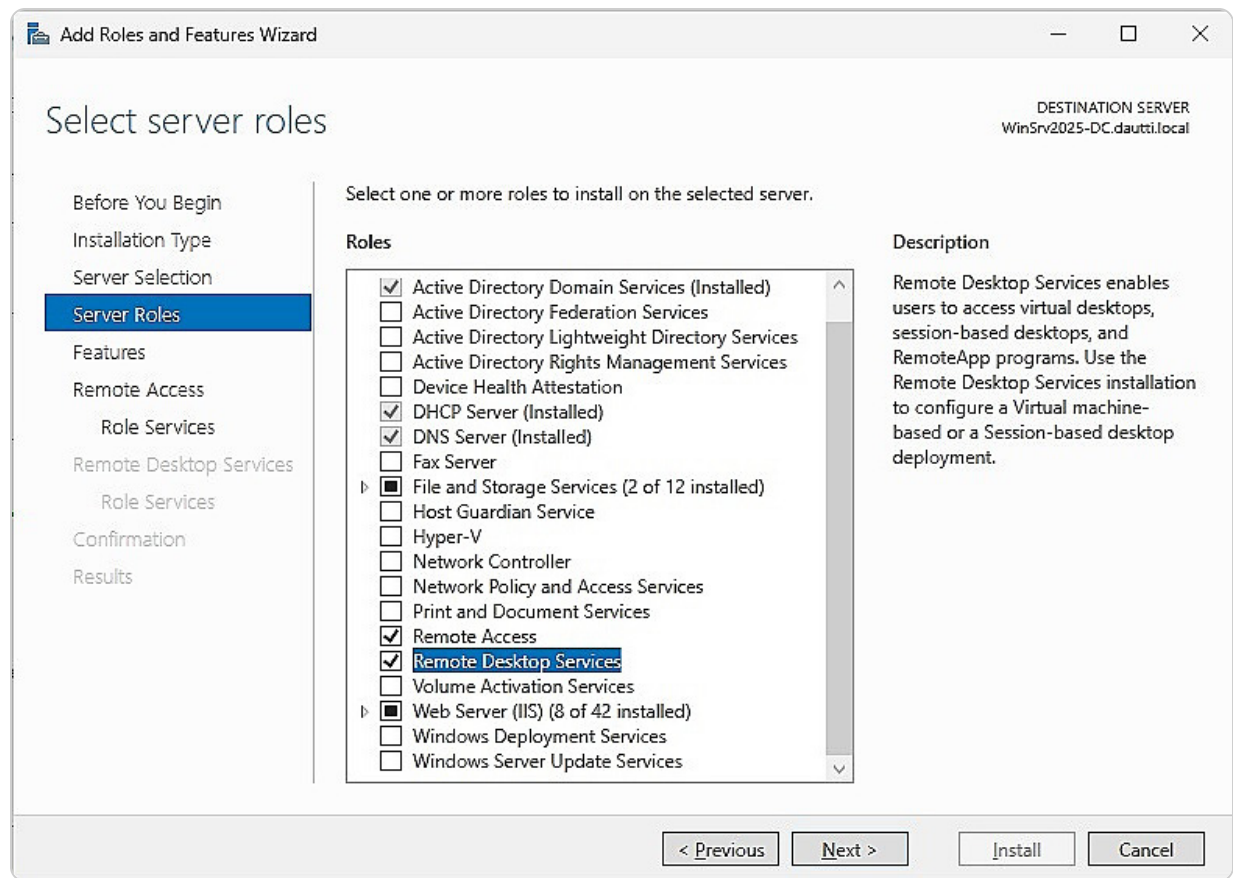


Figure 5.16- Adding the RDS role in Windows Server 2025

In the upcoming section, we'll delve into the installation and configuration of the RDS licensing server, which is essential for expanding beyond the default session limitations.

How to manage RDS CALs

Remote Desktop Services Client Access Licenses (RDS CALs) are required for users and devices to access a **Remote Desktop Session Host (RDSH)** server, which allows remote connections to desktops and applications. The RDS licensing server is pivotal in managing these licenses and issuing and tracking their usage within the network. By default, the RDS licensing server supports up to two concurrent Remote Desktop sessions without additional licensing. For organizations that need more than these two free connections, purchasing extra RDS CALs is necessary to comply with licensing requirements and support additional users.

To configure an RDS licensing server on Windows Server 2025, you must first install the RDS role on the server. Afterward, select **Remote Desktop Licensing** as a role service, as depicted in *Figure 5.17*. This setup ensures that the server can issue the appropriate number of licenses to accommodate your organization's remote access needs.

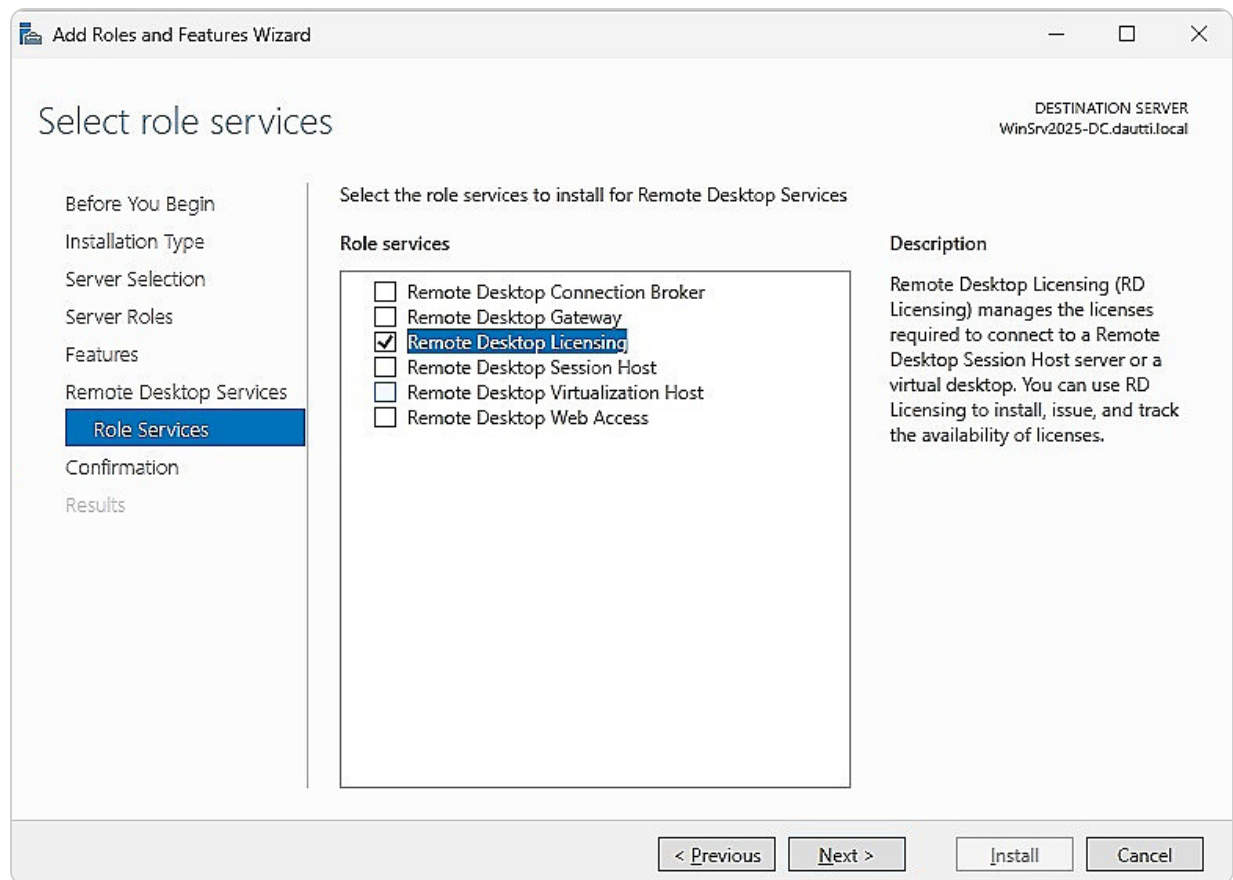


Figure 5.17- Adding Remote Desktop Licensing Role Services in Windows Server 2025

Additionally, it is crucial to regularly review and manage your RDS CALs to ensure compliance with licensing agreements and avoid service interruptions. In the subsequent section, we will delve into the **Remote Desktop Gateway (RDG)** server, focusing on how to configure it to provide secure remote access to your network resources.

Setting up RDG

The **RDG** is an essential role service in Windows Server that facilitates secure remote access to computers within a private network from anywhere on the internet. It acts as an intermediary, or proxy, between the Remote Desktop client and the internal target computer, ensuring that all communications are encrypted and secure. By implementing RDG, organizations can provide remote users with access to network resources without exposing internal systems directly to the internet.

To install and configure an RDG server on Windows Server 2025, begin by adding the RDS role to the server. After installing this role, select **Remote Desktop Gateway** as the specific role service, as shown in *Figure 5.18*. This setup allows the RDG server to handle incoming Remote Desktop requests, authenticate users, and establish a secure, encrypted connection between the client and the network resources.

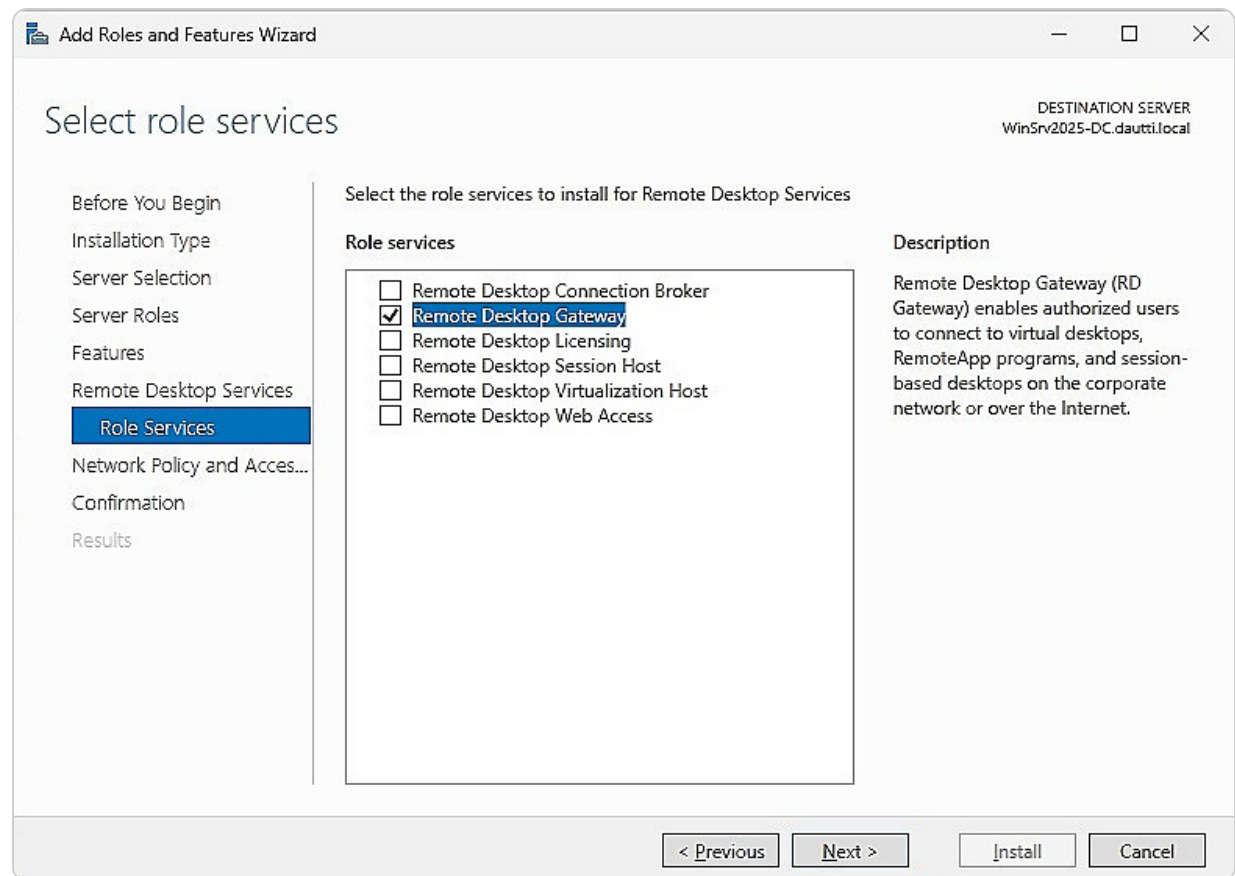


Figure 5.18- Adding RDG role services in Windows Server 2025

Additionally, RDG supports advanced features, such as centralizing user access management and enforcing security policies. For example, RDG can integrate with NPS to enforce conditional access policies, ensuring that only compliant devices and users can connect.

In the following subsection, we will introduce VPNs, explain their role in enhancing remote access security, and detail the steps for implementing a VPN solution.

What is a VPN?

A **VPN** provides a secure way to connect and transmit data over the internet between computers located on different networks. By leveraging tunneling protocols and encryption techniques, a VPN creates a virtual link that simulates a direct connection between the endpoints. This technology is essential for connecting remote users to their corporate networks or for linking different organizational networks over the internet. There are two primary types of VPNs:

- **Remote-Access VPN:** This is designed to grant telecommuters or remote workers access to their company's private network from any location. This type of VPN creates a secure tunnel from the user's device to the corporate network, ensuring data privacy and integrity during transmission.
- **Site-to-Site VPN:** This connects entire networks from different locations, allowing seamless communication between two distinct networks over the internet. This type of VPN is ideal for organizations with multiple branches or offices that need to share resources and data securely.

To set up a VPN server in Windows Server 2025, you need to install the Remote Access role and select the **DirectAccess and VPN (RAS)** role services, as depicted in *Figure 5.19*.

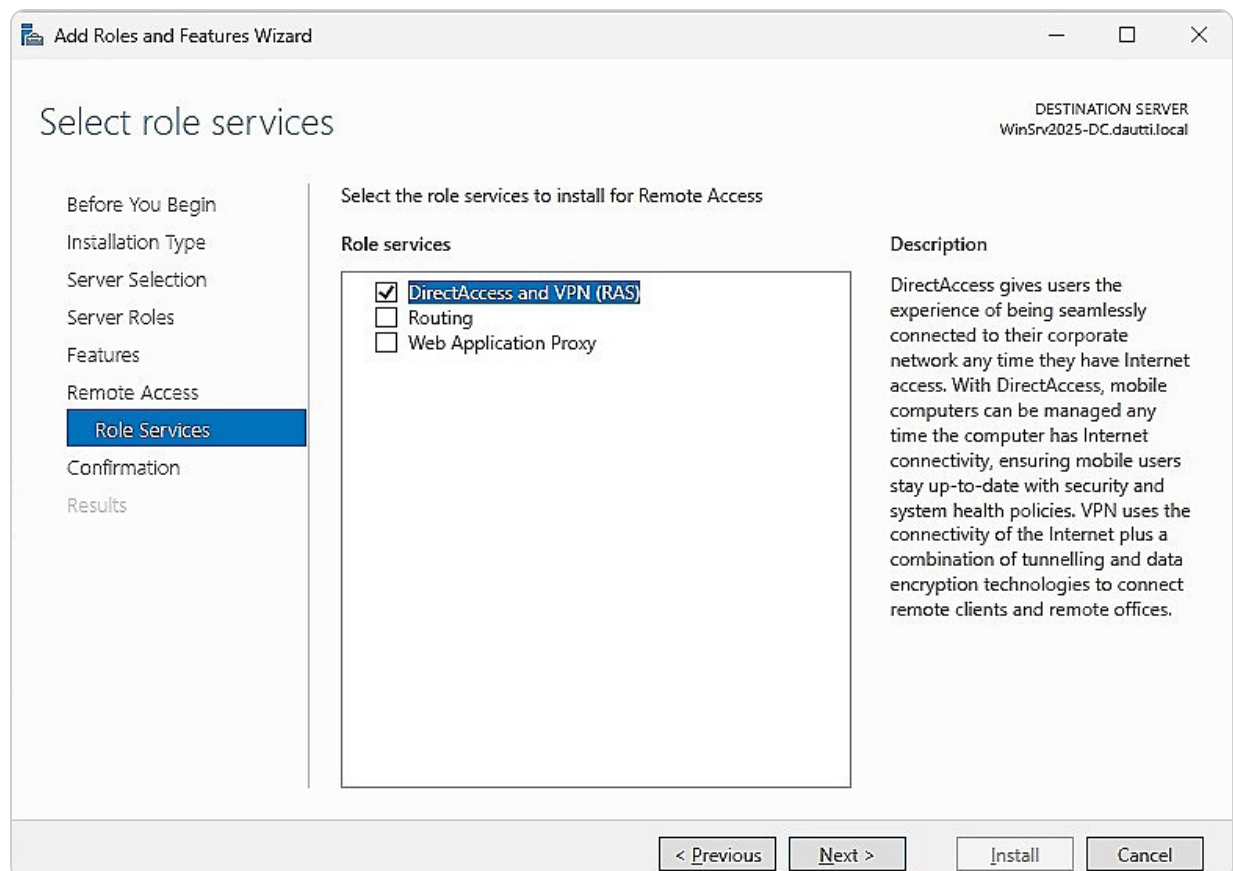


Figure 5.19- Adding the DirectAccess and VPN (RAS) role services in Windows Server 2025

DirectAccess and VPN, while similar in their aim to provide secure remote connectivity, differ significantly in their implementation and use cases. DirectAccess provides a more seamless experience by automatically connecting remote users to the corporate network without requiring a separate VPN client. It integrates directly with the Windows operating system and offers features such as **always-on connectivity** and **automatic access** to internal resources. On the other hand, traditional VPN solutions require users to initiate a connection manually and often involve additional client software. VPNs are generally more versatile and can be used across various operating systems and devices, whereas DirectAccess is tailored explicitly for environments with a homogeneous Windows infrastructure.

In the upcoming subsection, we will delve into the **Application Virtualization (App-V)** service. App-V enables the execution of applications without the need for local installation, leveraging virtualization technology to achieve this. By isolating applications from the local environment, App-V helps minimize conflicts and compatibility issues that can arise from varying application requirements and settings. This service proves to be an invaluable tool for IT professionals seeking to streamline application management, deployment, and security, enhancing overall efficiency and reducing potential disruptions.

Explaining App-V

Microsoft App-V enables users to access applications hosted on a server rather than installed locally, offering a seamless experience as though the applications were running directly on their devices. This service simplifies application management by reducing conflicts and compatibility issues that may occur with traditional installations. App-V allows system administrators to manage application access centrally, thereby controlling which applications users can interact with. To implement App-V, you would first need to obtain the **Microsoft Desktop Optimization Pack (MDOP)** from Microsoft's website.

In modern environments, App-V is increasingly utilized in cloud-based settings rather than traditional on-premises installations. **Cloud-based App-V** deployment provides several advantages, including scalability, reduced infrastructure costs, and easier updates and maintenance. By leveraging cloud resources, organizations can deploy and manage virtualized applications more flexibly, ensuring that users have access to the latest applications without the need for local installation or extensive hardware. This approach enhances the efficiency and security of application management, particularly in dynamic and distributed work environments.

NOTE

For detailed guidance on deploying Microsoft App-V on a local server, please visit the official documentation at Microsoft's App-V deployment guide <https://learn.microsoft.com/en-us/microsoft-desktop-optimization-pack/app-v/appv-deploy-the-appv-server>. This resource provides comprehensive instructions and best practices for successfully setting up and configuring the App-V server in your local environment.

Next, we will delve into the topic of managing multiple ports and their applications.

Understanding multiple ports

In this subsection, we will explore the concept of using **multiple ports**, which is essential for managing network communications effectively. A key example is the **App-V** service, which allows applications to run on a server rather than being installed locally on each user's device. This virtualization technology helps prevent software conflicts and compatibility issues by isolating applications from the local operating system. For IT professionals, App-V offers a streamlined approach to application management and deployment, enhancing security and operational efficiency.

Previously, we discussed RDS and the use of port 3389 for establishing Remote Desktop connections. However, RDS, by default, only supports one active remote session per port. To facilitate concurrent access to multiple computers, RDS employs additional port numbers, beginning with 3390 and increasing sequentially. Each port number corresponds to a different remote session, allowing multiple users to connect simultaneously within a **local area network (LAN)**.

Furthermore, remote access to multiple systems can be managed using IP sockets, which combine an IP address with a port number to specify unique communication channels. An IP socket helps direct data traffic accurately between the client and the server. The format for an IP socket is as follows:

Syntax: `Public_IP_address:Port_number`

Example: `113.79.43.133:3389`

Understanding and managing multiple ports is crucial for optimizing network performance and ensuring efficient remote access. This section covered the fundamentals of port usage with RDS and App-V, setting the stage for our following discussion on file and print services in Windows Server 2025.

Deploying file and print services for network environments

File and print services trace their origins to the early days of computer networking when the primary goal was to enable resource sharing among users. These services have significantly evolved over the years and are now fundamental to both home and business networks. They encompass the management of file storage, allowing users to store, retrieve, and share files efficiently, as well as print

services, which enable the sharing of printers across a network, facilitating access to printing resources from multiple devices. Modern **network operating systems (NOSSs)**, including Windows Server 2025, provide advanced file and print services that support enhanced features such as network security, access control, and integration with cloud storage solutions. These services are designed to streamline data management and enhance productivity by ensuring that files and printers are readily accessible, secure, and efficiently managed across various network environments. As networking technologies continue to advance, file and print services adapt to offer improved performance, scalability, and compatibility with emerging technologies.

File Services overview

The **File Services role** in Windows Server 2025 is an integral network component that is automatically included during the installation of the operating system (see *Figure 5.20*). This role provides essential functionalities for data management and accessibility across a network. Key features encompassed within the File Services role include the following:

- **File Sharing:** This feature allows users to share files and folders across the network, making them accessible to authorized users from various devices and locations. It simplifies collaboration and data exchange within organizations.
- **Work Folders:** This functionality enables users to synchronize their work files between their local devices and the server, ensuring that their data is consistently updated and available across multiple endpoints.
- **Distributed File System (DFS) Namespaces:** DFS namespaces offer a unified view of file shares, allowing users to access files and folders from a single namespace, regardless of their physical location on the network. That simplifies file access and enhances user experience.
- **BranchCache:** This feature optimizes network performance by caching frequently accessed files locally at branch offices. It reduces bandwidth consumption and improves access speed, making it particularly useful for remote locations with limited connectivity.

Together, these features provide a comprehensive suite of tools for managing and accessing data efficiently within a networked environment. They support various data access scenarios, from local file sharing to remote synchronization and efficient caching.

In the following subsection, we will explore the role of PDS, including its various components and installation procedures, to further enhance network resource management.

PDS Role overview

The **PDS role** in Windows Server 2025 plays a crucial role in centralizing the management of network printing and scanning functions. By supporting network printers and scanners, the PDS role streamlines the process of sending and receiving print jobs, as well as managing scanned documents within the network. To configure a print server, the PDS role must first be installed on Windows

Server 2025, as depicted in *Figure 5.20*. Once installed, various additional role services can be added to extend the PDS capabilities.

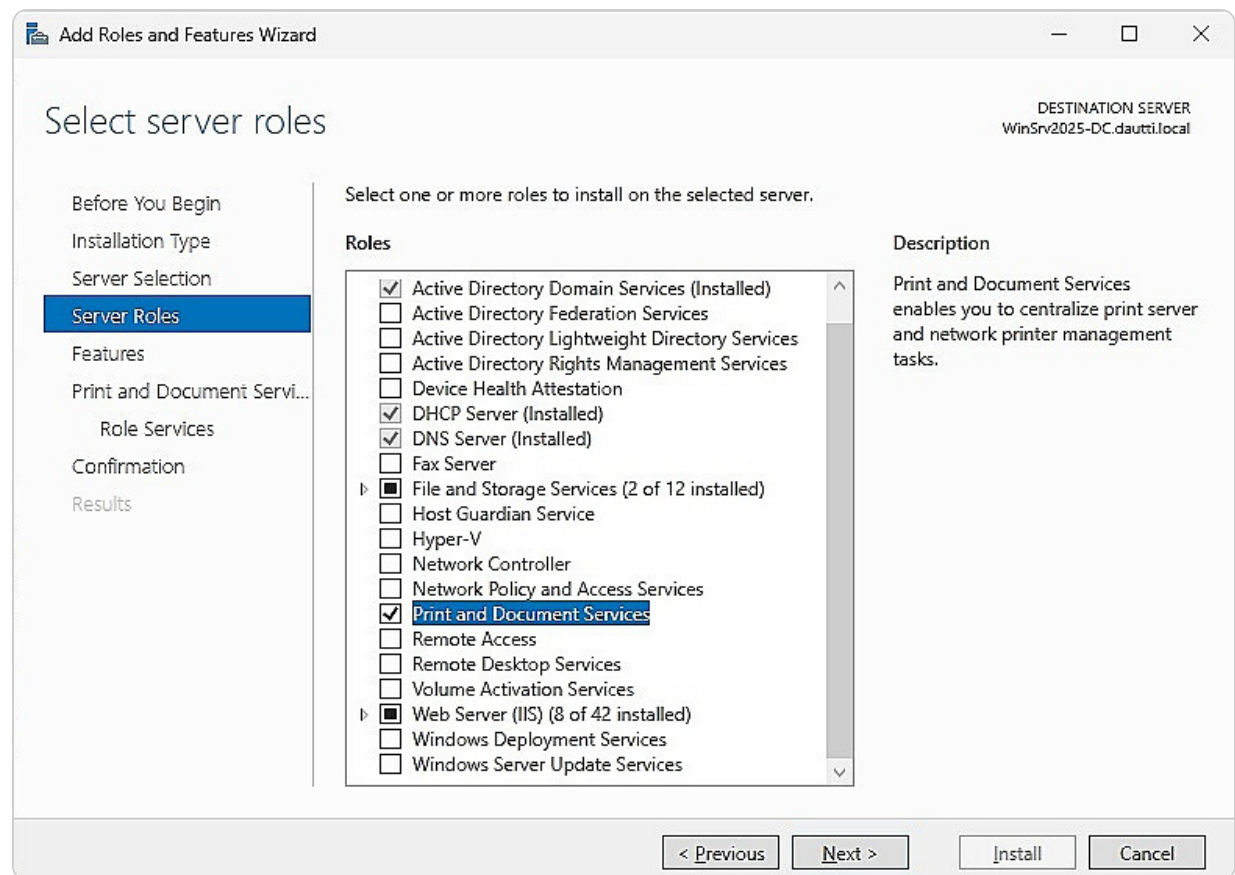


Figure 5.20- Adding PDS role in Windows Server 2025

Key role services available with PDS include the following:

- **Print Server:** This service manages print queues, handles the deployment of printers, and facilitates the migration of print servers, ensuring that print jobs are processed efficiently and that printer resources are utilized optimally.
- **Internet Printing:** This feature allows users to print documents over the web using the **Internet Printing Protocol (IPP)**. It provides a convenient method for remote printing, enabling users to submit print jobs from anywhere with internet access.
- **Line Printer Daemon (LPD) Service:** The LPD service enables printing from Unix-based systems and other non-Windows operating systems using the **Line Printer Remote (LPR)** protocol. This service ensures interoperability across different platforms, making it easier to integrate diverse printing environments.

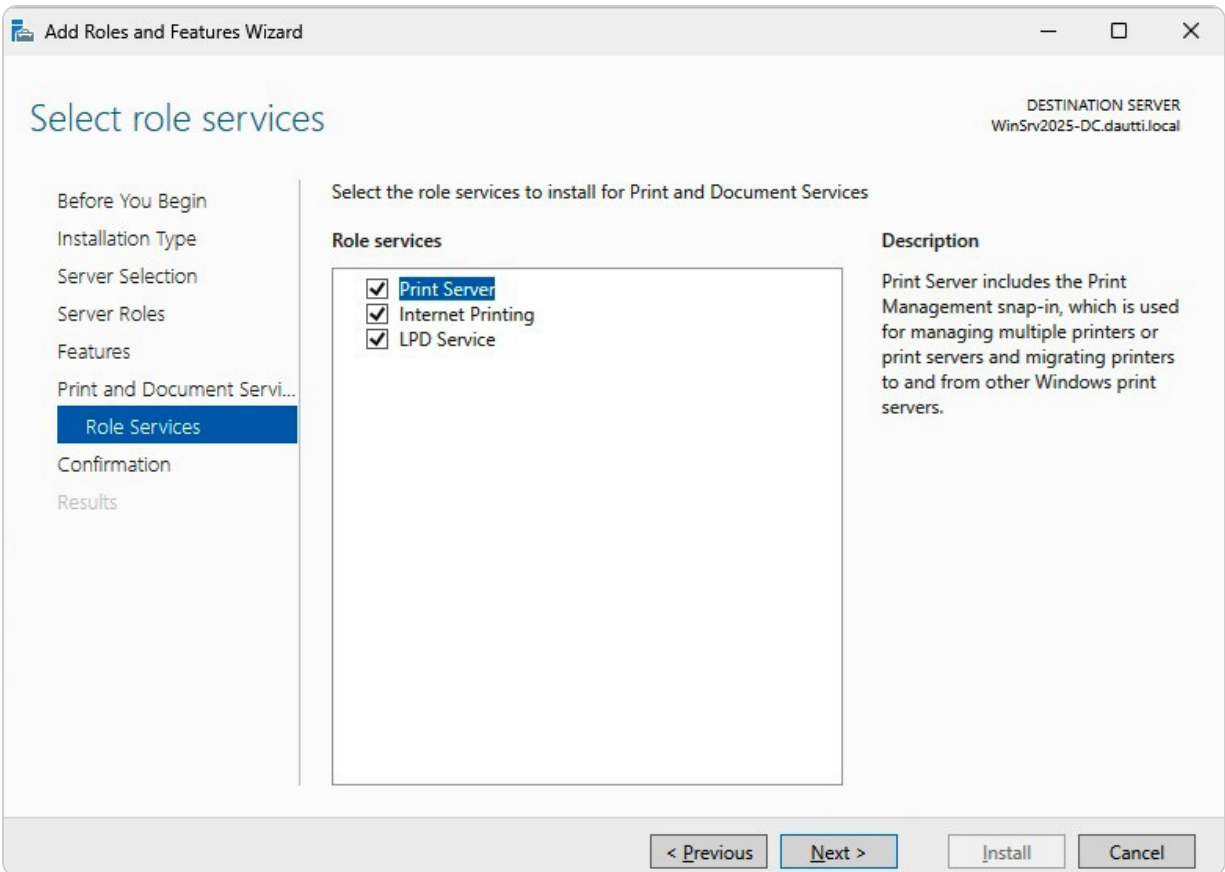


Figure 5.21- Installing PDS Role Services in Windows Server 2025

Additionally, the PDS role supports advanced features, such as print management, auditing, and secure print release, which enhance the control and security of printing tasks. In the following subsection, we will explore various printer-related concepts, starting with the setup and management of local printers, and expand on how these concepts fit into a broader network printing strategy.

What is a local printer?

A **local printer** is directly connected to a single computer via a parallel port or USB port, functioning primarily to print documents for that specific computer. This direct connection limits its accessibility to the host machine unless printer sharing is enabled, allowing other networked computers to use it. Local printers are straightforward to set up and manage but may become a bottleneck in larger environments where multiple users need access.

In contrast, **network printers** are connected to a network and can be accessed by multiple computers simultaneously without requiring a dedicated connection to each one. These printers are often equipped with their **network interface card (NIC)** or are connected through a print server, allowing them to handle print jobs from any authorized user on the network. This setup centralizes printing,

improves efficiency, and simplifies management by eliminating the need for individual printer connections on each workstation.

Similarly, **internet printers**, which can be a type of network printer, enable printing over the internet. These printers are accessible via **web protocols**, allowing users to print documents from remote locations. Internet printers often utilize specialized software or cloud-based services to facilitate printing from anywhere, providing greater flexibility and accessibility compared to local and network printers.

In summary, while local printers are suitable for individual use or small setups with shared access, network, and internet printers offer scalable solutions for larger organizations by providing centralized, remote, and more flexible printing options.


Network printer explained

A **network printer** differs from a local printer in that it does not rely on a direct connection to a single computer. Instead, it features its network interface, either wired or wireless, allowing it to be accessed by multiple devices on the same network. This capability eliminates the need for individual connections to each workstation, making network printers ideal for shared environments.

Additionally, with proper configuration, network printers can be accessed remotely over the internet, extending their functionality beyond the local network. While the concept of network printers was introduced in a previous subsection, this section focuses on the deployment and management of network printers within a Windows Server 2025 environment. Examining the web address displayed in the browser in *Figure 5.22* illustrates how network printer services effectively manage network printers.

EPSON

WF-2520 Series



Printer Information

Information1	Information2
Printer Name :	EPSON51D283
Connection Status :	100BASE-TX Full Duplex
Obtain IP Address :	Manual
IP Address :	192.168.0.150
Subnet Mask :	255.255.255.0
Default Gateway :	192.168.0.1
MAC Address :	B0:E8:92:51:D2:83

Refresh

Figure 5.22- Adding a network printer

In the subsequent section, we will explore printer pooling, which provides further enhancements for managing multiple printers and optimizing print resources.

Understanding printer pooling

Printer pooling in Windows Server 2025 is a critical feature for optimizing print management within a corporate environment. By combining multiple physical printers into a single virtual printer, printer pooling helps streamline printing operations and increase efficiency. In a large organization, this setup is particularly beneficial as it balances the print load across several devices, minimizing the risk of any single printer becoming a bottleneck or experiencing downtime.

In practical terms, printer pooling allows for more reliable and scalable printing services by handling higher volumes of print requests and reducing the likelihood of printer congestion. When users send print jobs to the virtual printer, these jobs are automatically distributed among the available printers in the pool. This distribution ensures that print tasks are completed more quickly and that no single printer is overwhelmed by excessive demand.

Additionally, printer pooling simplifies printer management by presenting a unified printer interface to end users, even though multiple physical devices are in operation. That means users only need to interact with one printer queue, making the printing process more intuitive and less prone to errors.

To set up printer pooling, you need to install the PDS role along with the Print Server role services on your Windows Server 2025. Once installed, you can add printers to the pool and configure the pooling options via the **Print Management** console, as shown in *Figure 5.23*. This configuration not

only enhances print service reliability but also contributes to overall organizational productivity by ensuring efficient and consistent printing capabilities.

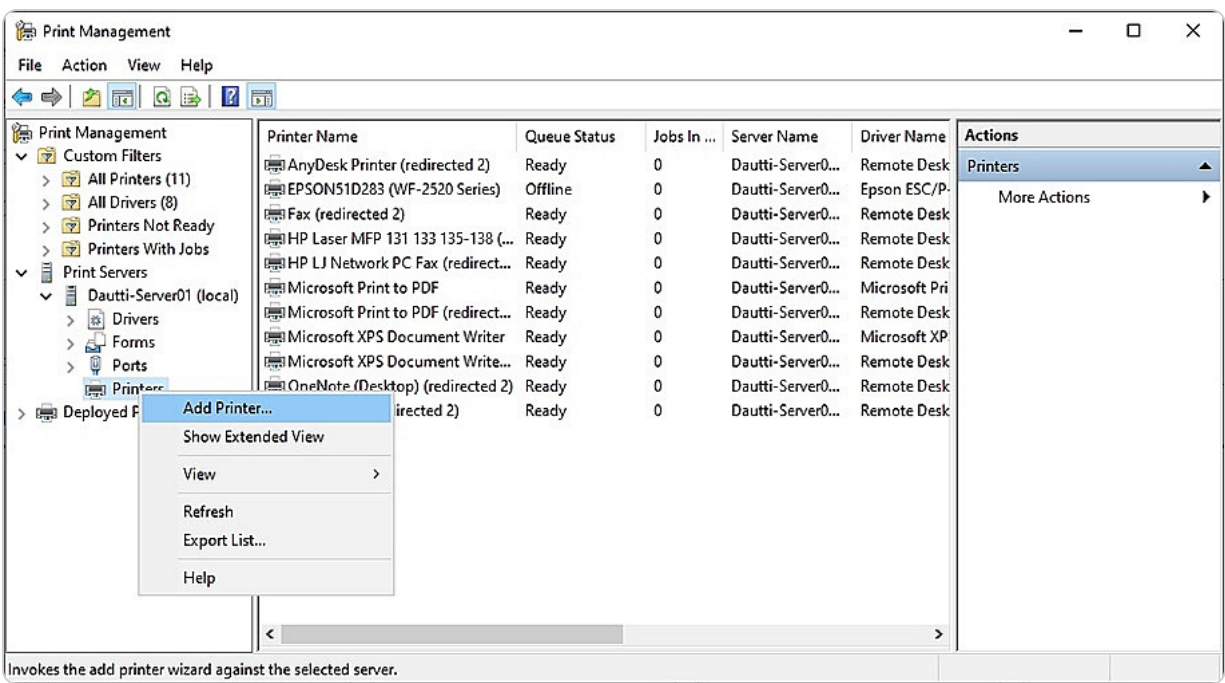


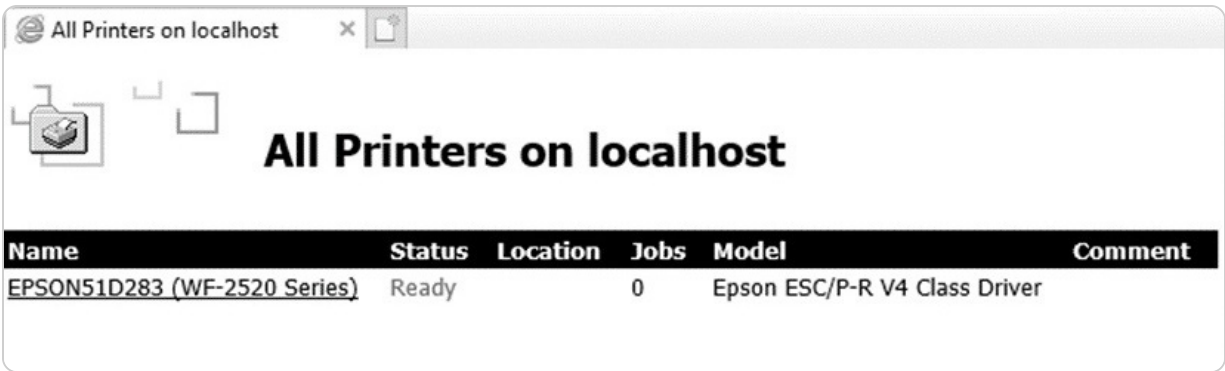
Figure 5.23 – Adding printers in a printer pooling in Windows Server 2025

In the following section, we will delve into how to integrate and manage shared network resources, which further supports effective IT infrastructure management.

Internet printing overview

Internet printing allows users to send print jobs to network printers via a web browser, offering a convenient and remote method to manage printing tasks. Previously, we discussed the concept of internet printing, which enables users to access and utilize network printers through web-based interfaces. This section will focus on a practical example of deploying internet printing within your organization. To set up internet printing on a Windows Server 2025, you must install the PDS role and select the Internet Printing role service. Additionally, configuring the web server role (IIS) is necessary to support this feature.

Once configured, users can access available printers by navigating to <http://servername/printers> in their web browser, as shown in *Figure 5.24*. This setup allows for remote printing, making it easier for users to manage print jobs from various locations within the network, thus enhancing overall efficiency and flexibility. In the next section, we will delve into managing web printing services, focusing on configuration and optimization to better serve organizational needs.



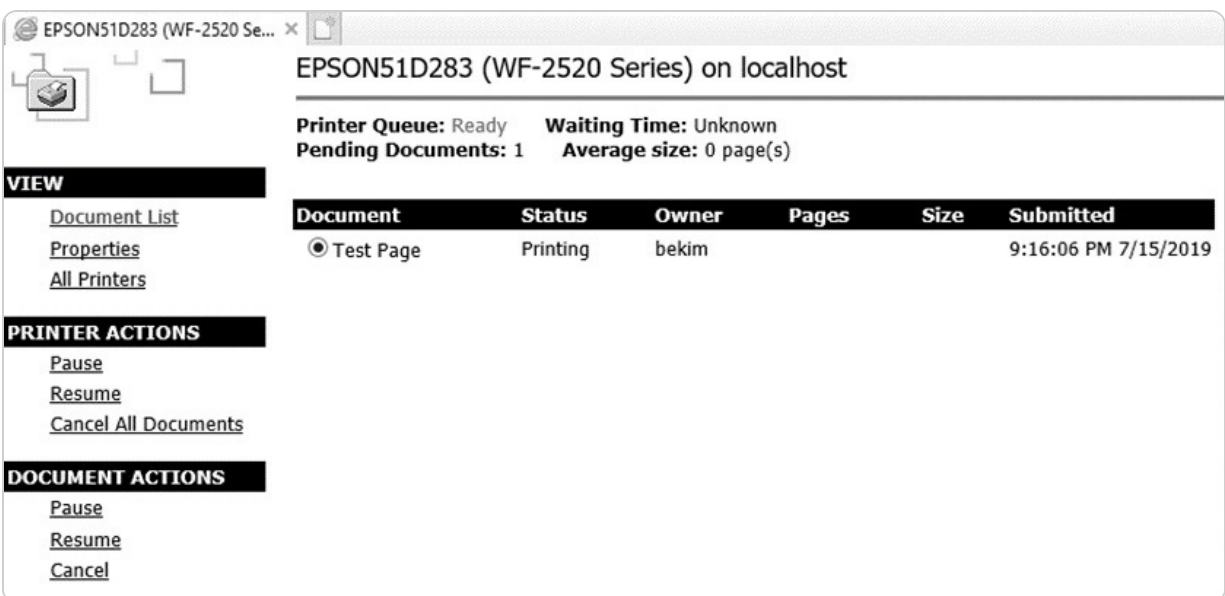
Name	Status	Location	Jobs	Model	Comment
EPSON51D283 (WF-2520 Series)	Ready		0	Epson ESC/P-R V4 Class Driver	

Figure 5.24- Web printing

Next, we will explore the management of web printing services. That involves configuring and optimizing web-based printing solutions to ensure efficient operation and ease of use within an organization. We'll cover key aspects such as setting up and maintaining web print servers, managing user access, and troubleshooting common issues to ensure a smooth and effective printing experience.

Understanding Web Printing Management

Web printing management, often referred to as **internet printing**, offers more than just the ability to print documents over the web; it also provides comprehensive tools for managing print jobs remotely. This functionality enables users to interact with their print tasks from any web browser, offering a similar level of control and convenience as if they were using a local or network printer directly. To utilize this feature, users must enter `http://servername/printers` into their browser's address bar, which will display a management interface (see *Figure 5.25*) where they can view and control their print jobs.



VIEW

- [Document List](#)
- [Properties](#)
- [All Printers](#)

PRINTER ACTIONS

- [Pause](#)
- [Resume](#)
- [Cancel All Documents](#)

DOCUMENT ACTIONS

- [Pause](#)
- [Resume](#)
- [Cancel](#)

EPSON51D283 (WF-2520 Series) on localhost

Printer Queue: Ready Waiting Time: Unknown
Pending Documents: 1 Average size: 0 page(s)

Document	Status	Owner	Pages	Size	Submitted
Test Page	Printing	bekim			9:16:06 PM 7/15/2019

Figure 5.25- Web printing management

This capability is particularly beneficial for environments where users need to manage print tasks from different locations or when centralized control over printing resources is required. However, before you can take advantage of web printing management, ensure that the Internet Printing role service is installed. That involves first adding the PDS role, followed by the Internet Printing role service.

In the following subsection, we will delve into the deployment and management of printer drivers, which are essential for ensuring that printers operate correctly and efficiently within your network.

Understanding Printer Driver Deployment

Understanding **printer driver deployment** is a critical component of effective print management within a networked environment. The Print Management administrative console in Windows Server 2025 facilitates this process by offering a centralized platform for managing printers and their drivers. Deploying printer drivers through this console involves selecting the appropriate driver for a specific printer and then pushing this driver to one or multiple computers on the network. That ensures that the printer operates with the latest driver version, which is crucial for maintaining compatibility, functionality, and performance across different systems.

In addition to installing or upgrading drivers, the console provides tools for managing print queues, monitoring print jobs, and configuring printer settings. This centralized approach simplifies the administration of print resources and helps address any potential issues quickly.

Figure 5.26 demonstrates how to deploy print drivers using the **Print Management** console in Windows Server 2025, highlighting the steps involved in ensuring that all networked printers are equipped with the correct and most up-to-date drivers.

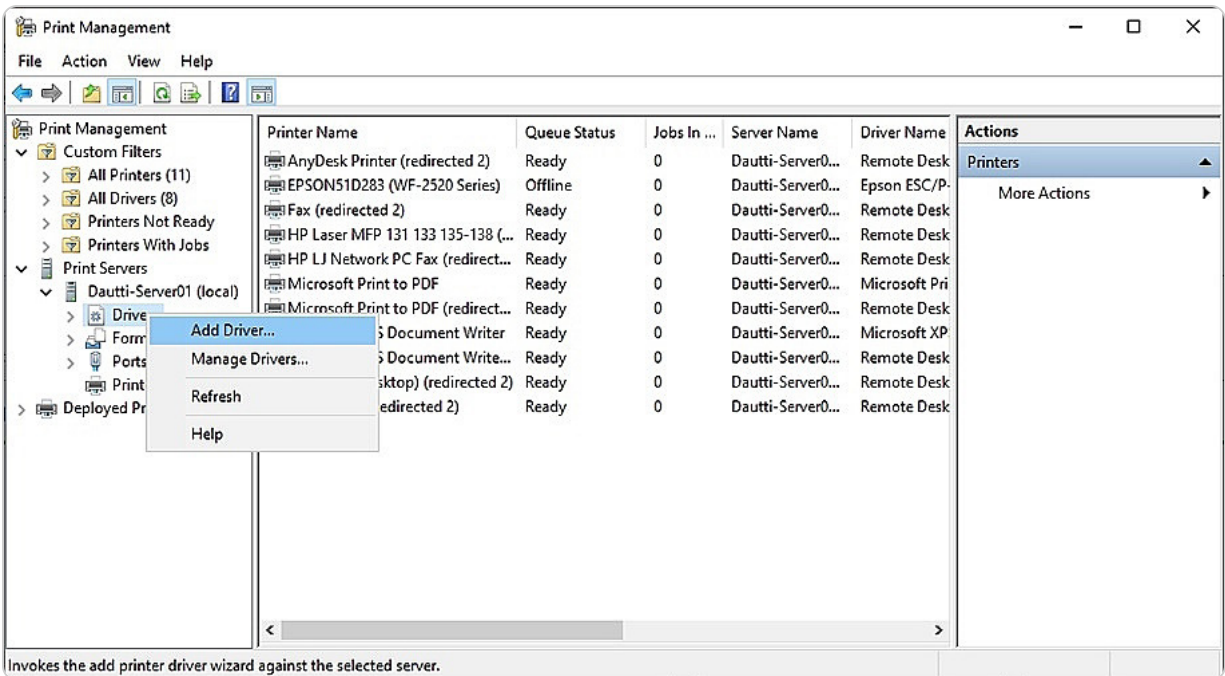


Figure 5.26- Deploying print drivers with the Print Management console

In the subsequent section, we will explore user rights, **New Technology File System (NTFS)** permissions, and share permissions. Understanding these concepts is essential for managing access control and ensuring the security of shared resources in a networked environment.

Understanding User Rights and Permissions Management

Effective access management and security are fundamental to maintaining the integrity of any networked environment. To ensure robust protection and proper resource management, administrators must have a comprehensive understanding of **user rights**, **NTFS permissions**, and **share permissions**. User rights specify the capabilities and restrictions of users regarding system operations and resources. NTFS permissions control access to files and directories, whether on a local machine or a network drive. In contrast, share permissions govern access to network-shared resources. Mastering these concepts enables administrators to enforce appropriate access levels, safeguard data, and prevent unauthorized access to sensitive information. In the following sections, we will explore each of these elements in detail, offering practical guidance on configuration and management.

NTFS permissions explained

To effectively manage access and security within a Windows environment, it's essential to grasp the nuances of **NTFS permissions**. When managing a folder's security settings in Windows, you can access

a detailed view of permissions by right-clicking the folder, selecting **Properties**, and navigating to the **Security** tab. Here, you'll encounter several types of permissions that control how users interact with files and folders:

- **Full Control:** This permission provides comprehensive access, enabling users to perform all actions, including reading, writing, modifying, executing, changing attributes and permissions, and deleting files and subfolders. It grants complete administrative capabilities over the folder's contents.
- **Modify:** This allows users to view and alter the contents, which includes adding, editing, and deleting files and subfolders, but does not permit changing permissions or attributes.
- **Read & Execute:** Users can run executable files and manage their execution. This permission also allows viewing files and their contents but does not permit modifications.
- **List Folder Contents:** This enables users to see the names and properties of files and subfolders within the folder but does not grant access to open or modify the files.
- **Read:** This provides the ability to view the files and their properties but does not allow any changes to the file contents.
- **Write:** This permits users to add new files and modify the contents of existing ones but does not allow reading or executing files.
- **Special Permissions:** This offers advanced, granular control over specific file or folder operations, such as creating files, deleting subfolders, and more nuanced actions.

Permissions are applied in an *allowed* or *denied* manner (refer to *Figure 5.27*), which directly impacts user access levels. The configuration of these permissions determines how users can interact with system resources, thereby influencing data security and operational efficiency.

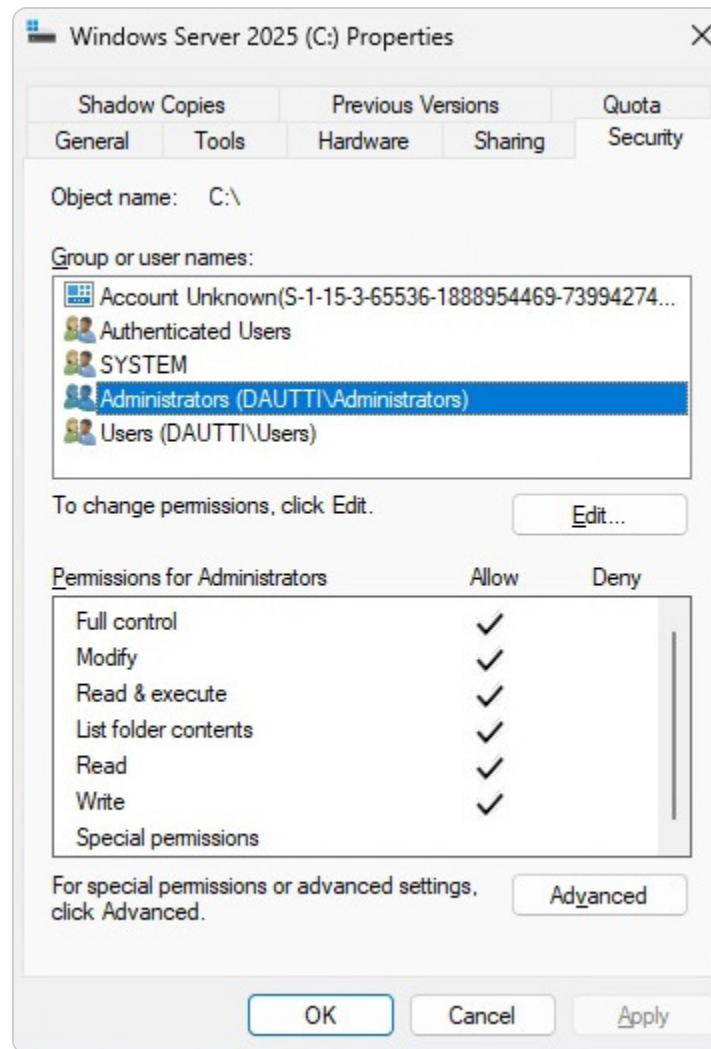


Figure 5.27- NTFS permissions in Windows Server 2025

By understanding and appropriately assigning these permissions, administrators can effectively manage user access and maintain a secure network environment.

Understanding Share Permissions

In addition to NTFS permissions, which manage access to files and folders on a local server, **share permissions** are essential for controlling access to resources shared across a network. Share permissions determine how users can interact with these network-shared resources and are categorized into three distinct levels:

- **Full Control:** This grants users complete authority over the shared resource, enabling them to read, modify, and alter permissions and ownership of files and subfolders
- **Change:** This permission allows users to read, execute, write, and delete files and subfolders but does not grant the ability to modify permissions

- **Read:** This permission limits users to viewing and listing the content of the shared resource without any capability to make changes

These permissions work in tandem with NTFS permissions to provide a layered security model. For instance, if NTFS permissions restrict access to a file, share permissions cannot override this restriction, and vice versa. Proper configuration of share permissions ensures that users have appropriate access based on their roles and needs while maintaining the security of network resources. *Figure 5.28* provides a visual representation of how to set and manage these permissions in Windows Server 2025, illustrating the steps necessary to configure network access and ensure effective resource management.

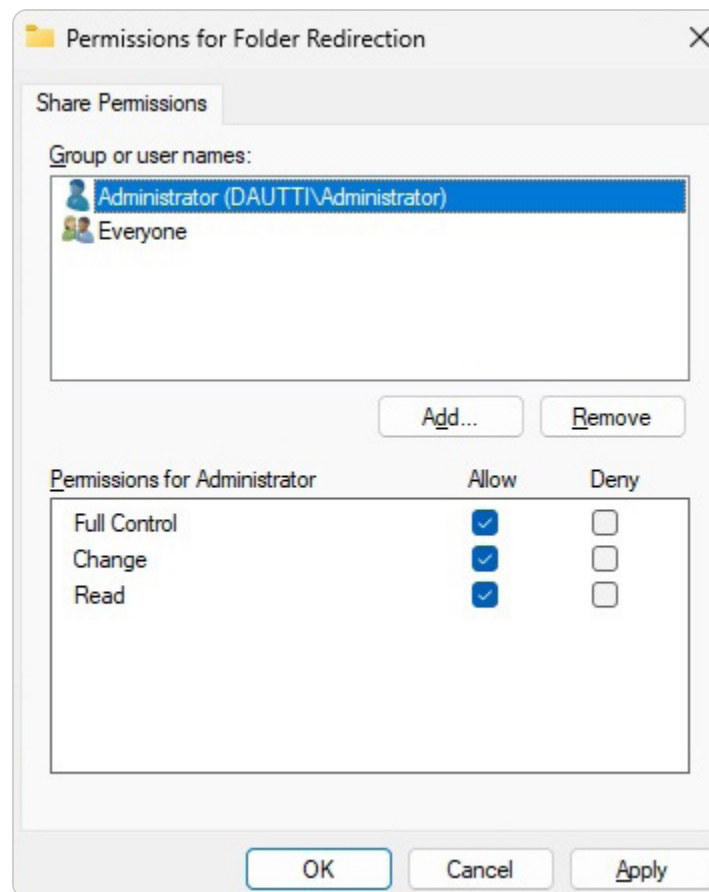


Figure 5.28- Shared permissions in Windows Server 2025

Configuring User Rights

Effective management of user rights is fundamental to ensuring both system security and proper operational efficiency. User rights can be configured using tools such as **Local Group Policy Editor** (`gpedit.msc`), **Local Security Policy**, or **Default Domain Policy**. To access these settings, navigate to **Computer Configuration | Windows Settings | Security Settings | Policies | User Rights Assignment**. These policies allow administrators to define specific system privileges for user accounts,

including who can log in locally, access the network, or manage system services. In a domain environment, some user rights may already be established according to the organization’s default domain policies (see *Figure 5.29*).

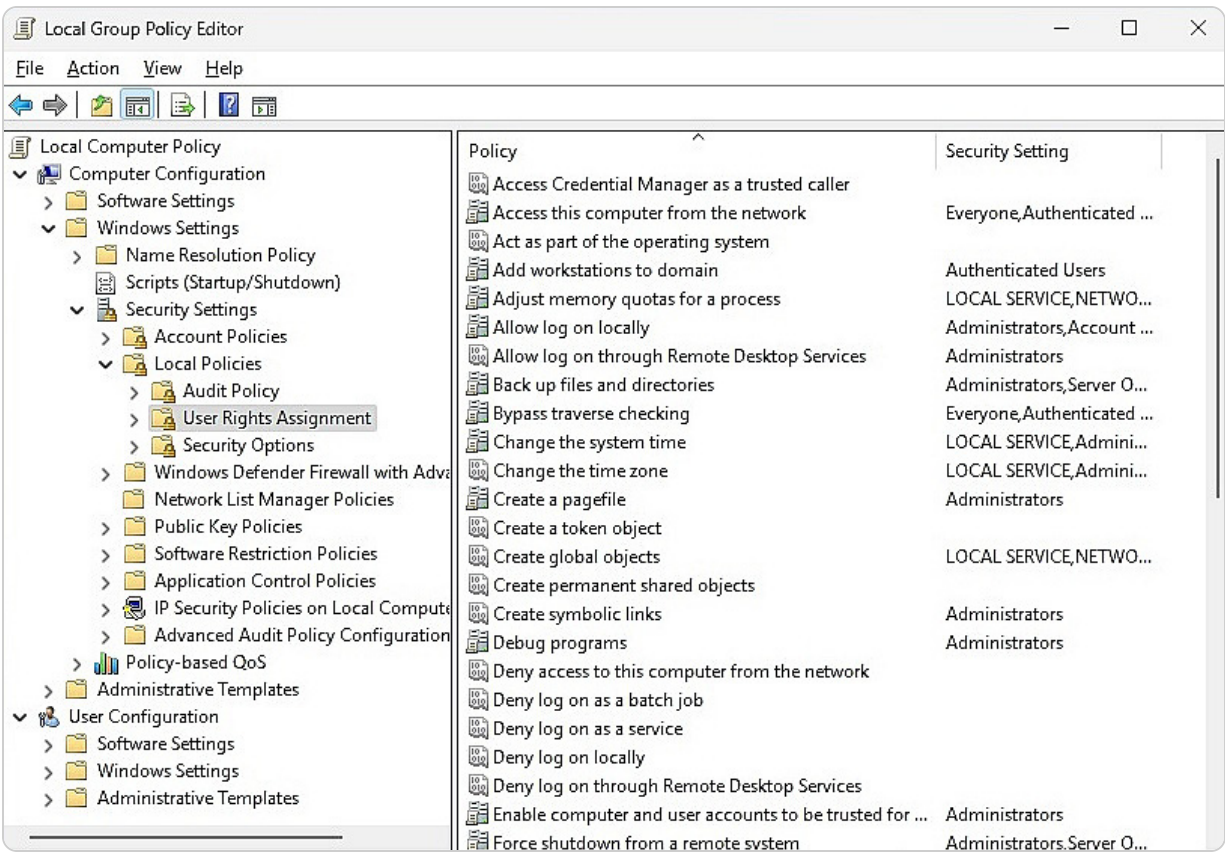


Figure 5.29- User Rights Assignment in Windows Server 2025

Understanding the distinction between user rights and permissions is essential for effective system management. *User rights* dictate what actions a user can perform on the system as a whole, such as shutting down the server or managing system services. In contrast, *permissions* are more granular and control access to specific resources, such as files and folders.

Proper configuration of both user rights and permissions ensures that users have the necessary access to perform their roles while maintaining the security and integrity of the server environment. This comprehensive approach helps prevent unauthorized access and ensures that security policies align with organizational requirements.

Monitoring file server activities

Given the critical role that file servers play in an organization’s IT infrastructure, maintaining rigorous oversight through auditing is essential for ensuring data security and compliance. **Auditing** enables

administrators to track and document all activities related to file access, modifications, and deletions. That includes capturing detailed records of who accessed specific files, what changes were made, and when these actions occurred. Such comprehensive monitoring is indispensable for detecting unauthorized access, troubleshooting issues, and ensuring adherence to organizational policies and regulatory requirements.

To set up and manage auditing in Windows Server 2025, you should use Local Group Policy Editor (`gpedit.msc`), Local Security Policy, or Default Domain Policy. Navigate to **Computer Configuration | Windows Settings | Security Settings | Local Policies | Audit Policy** to configure the audit settings according to your needs (see *Figure 5.30*). This setup allows you to define which activities are logged and how they are reported, enabling detailed and actionable insights into file server operations.

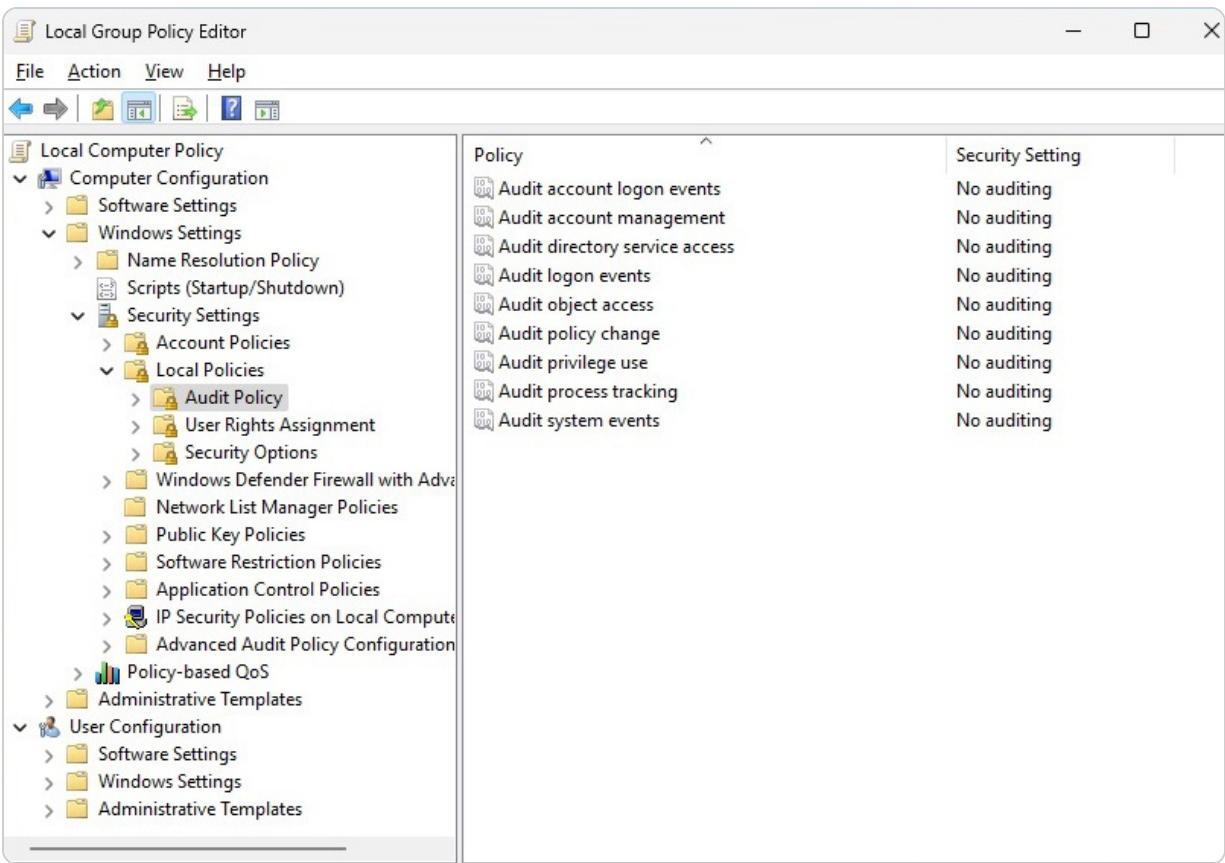


Figure 5.30- Audit policies in Windows Server 2025

Effective auditing not only helps safeguard sensitive data but also plays a pivotal role in maintaining the overall health of the IT infrastructure. It provides transparency, aids in compliance, and supports proactive management. In the next section, we will shift our focus to installing and configuring the Web Server (IIS) and PDS roles, which are essential for expanding and optimizing server functionalities.

Chapter exercise – installing webserver (IIS) and PDS roles

In this chapter's exercise, you will be guided through two critical tasks that are essential for enhancing your server's capabilities:

- First, you'll install the Web Server (IIS) role, which allows your server to host and manage web applications, websites, and a variety of web-based services, making it a central point for internet or intranet activities. This step is crucial for any environment where web services play a key role in business operations.
- Next, you'll install the PDS role, a vital component for any networked environment that relies on centralized management of printers and print jobs. This role enables your server to efficiently oversee printing tasks, ensuring that networked printers are effectively managed, print jobs are handled with priority, and resources are utilized optimally.

Together, these tasks will equip your server with the foundational tools needed to support both web hosting and document services in a networked setting, making it a more versatile and powerful asset in your IT infrastructure.

Setting up a Web Server (IIS) role

Setting up the Web Server (IIS) role in Windows Server 2025 allows you to host and manage web applications and services efficiently. To install this role, follow these steps:

1. Begin by opening **Server Manager** from the Start menu.
2. Click on **Add Roles and Features** to launch **Add Roles and Features Wizard**, as shown in *Figure 5.31*.
3. Choose the **Role-based or feature-based installation** option, and then click **Next**.
4. Ensure the correct server is selected from the server pool by choosing **Select a server from the server pool** and then clicking **Next**.

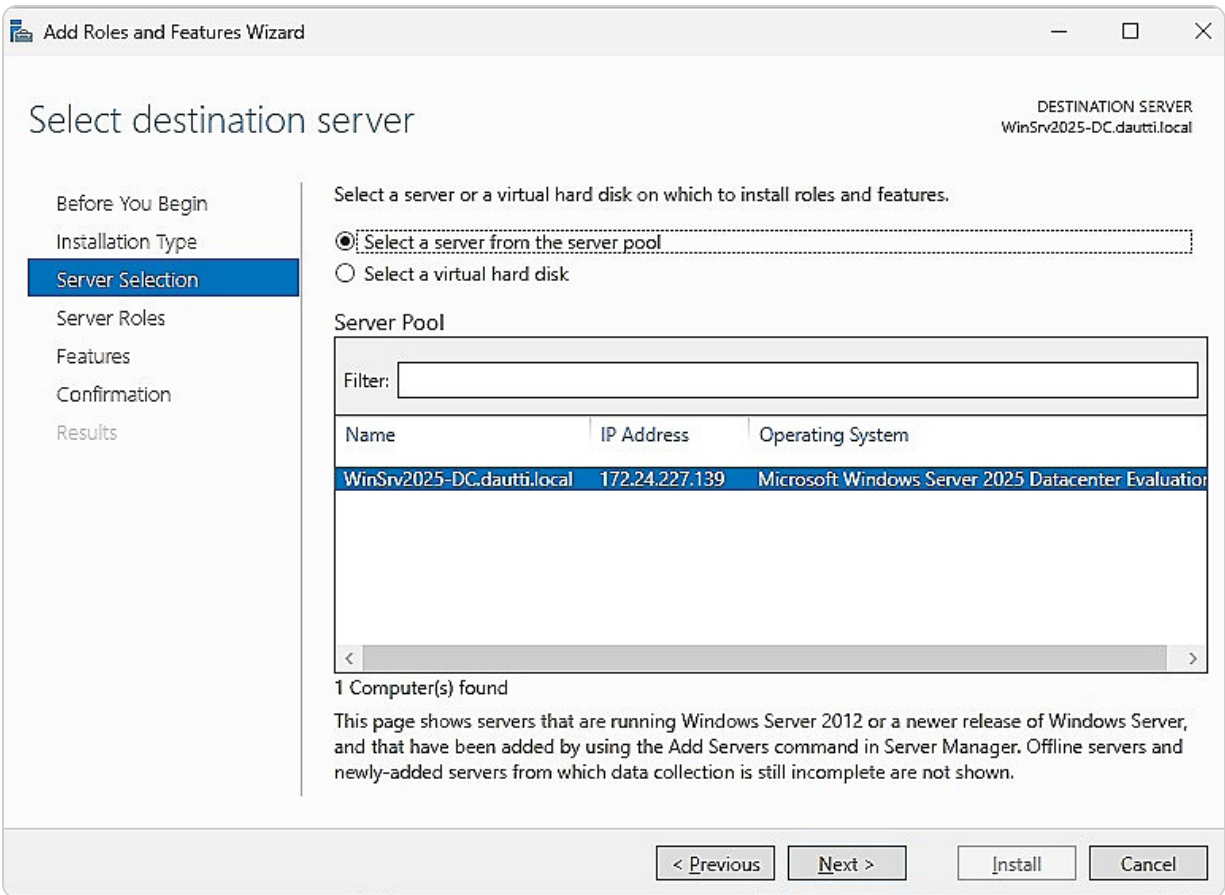


Figure 5.31- Selecting the destination server

5. In the list of available roles, check the box next to **Web Server (IIS)**.
6. When prompted with a message to add features required for Web Server (IIS), click the **Add Features** button.
7. At this point, you do not need to add any additional features, so click **Next**.
8. Please review the description and installation notes for the Web Server (IIS) role before proceeding. Then, click **Next**.
9. You can accept the default role services for Web Server (IIS) or customize them based on your specific needs.
10. Once you have reviewed your selections, click the **Install** button to begin the installation process.
11. After the installation is complete, click **Close** to exit the wizard. At this point, a server restart is not necessary.

By following these steps, you will have successfully installed the Web Server (IIS) role on your server, equipping it to host and manage a wide range of web applications and services. This setup is foundational for any server that will be serving web content or providing web-based services in a networked environment.

Installing a PDS role

Setting up the PDS role in Windows Server 2025 enables your server to manage printers and print jobs effectively across the network. To install this role, follow these steps:

1. Begin by opening **Server Manager** from the Start menu.
2. Click on **Add Roles and Features** to launch **Add Roles and Features Wizard**.
3. Choose the **Role-based or feature-based installation** option, and then click **Next**.
4. Ensure the correct server is selected from the server pool by choosing **Select a server from the server pool** and then clicking **Next**.
5. In the list of available roles, check the box next to **Print and Document Services**.
6. When prompted with a message to add features required for PDS, click the **Add Features** button, as shown in *Figure 5.32*.

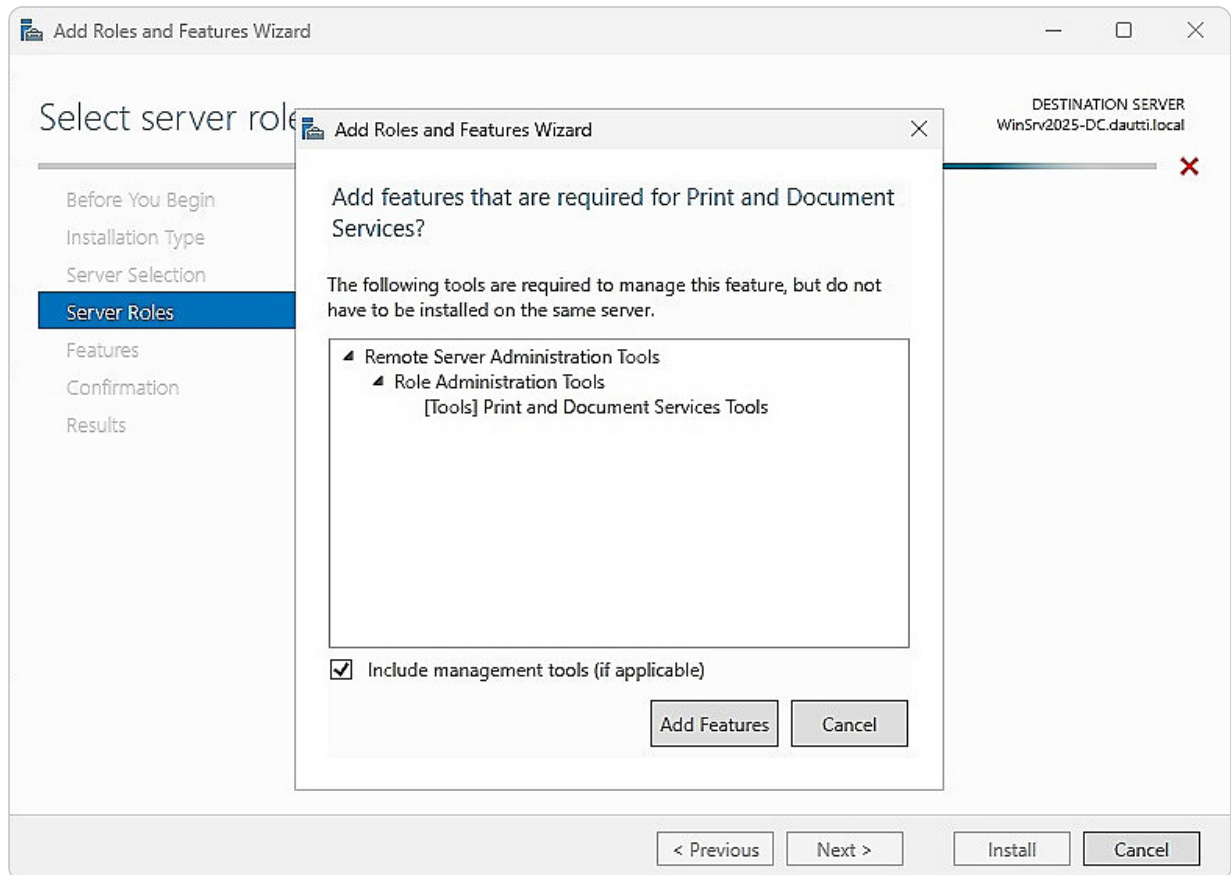


Figure 5.32- Adding features that are required for the PDS role

7. At this point, you do not need to add any additional features, so click **Next**.
8. Please review the description and installation notes for the PDS role before proceeding, then click **Next**.
9. You can accept the default role services for PDS or customize them based on your specific needs.
10. Once you've reviewed your selections, click the **Install** button to begin the installation process.
11. After the installation is complete, click **Close** to exit the wizard. A server restart is not necessary at this point.

The server now has the PDS role installed, enabling efficient management of printing and document services on the network.

By following these steps, you will have successfully installed and configured the Web Server (IIS) role, enabling your server to host and manage web applications and services effectively. In the next section, we will delve into advanced IIS configurations and management practices to optimize your server's web hosting capabilities. Similarly, by completing the steps outlined earlier, you have installed and set up the PDS role, allowing your server to manage network printing tasks efficiently.

Summary

In this chapter, you were introduced to the foundational concepts of roles, role services, and features within Windows Server 2025. We delved into critical aspects of file server management, including user rights, NTFS permissions, share permissions, and the importance of auditing to ensure secure and efficient server operations. Additionally, you gained insights into various common application servers, such as email servers, database servers, collaboration servers, monitoring servers, and data protection servers, along with the installation and configuration of web services, such as IIS, FTP, SSL, and digital certificates. The chapter also covered essential remote access services, including Remote Assistance, RSAT, RDS, RDS licensing, RDG, VPN, App-V, and the management of multiple ports. These topics are crucial for effectively adding and managing roles in Windows Server 2025 and securing your file server through the proper application of user rights, NTFS permissions, and share permissions.

Furthermore, this chapter included practical exercises on installing the Web Server (IIS) and PDS roles, equipping your server to host web applications and manage network printing efficiently. As we move forward, the next chapter will focus on Group Policy in Windows Server 2025, providing you with the knowledge to apply more granular and precise settings to both user and computer accounts.

Questions

1. **True or False:** A server role defines the primary function that a server performs within a network.
2. **Fill in the Blank:** _____ transfers files from computer to computer, computer to server, or vice versa, both on a LAN and WAN.
3. **Multiple Choice:** Which of the following are NTFS permissions in Windows Server 2025? (*Choose three.*)
 - Modify
 - Write
 - Change
 - Read

4. **True or False:** A web service is a communication method between two devices based on the request/response methodology using the FTP protocol.
5. **Fill in the Blank:** _____ is any logical endpoint where applications on your computer communicate with applications on other computers, both on a LAN and WAN.
6. **Multiple Choice:** Which of the following protocols are utilized by mail servers? (*Choose two.*)
- FTP
 - HTTP
 - SMTP
 - POP
7. **True or False:** Remote Assistance is a feature that enables a helper to access the host's desktop remotely to assist with resolving issues.
8. **Fill in the Blank:** _____ is responsible for securing the communication channel between a website and a browser.
9. **Single Choice:** Which of the following ports is used by RDS?
- 25
 - 110
 - 443
 - 3389
10. **True or False:** Web printing enables users to print files to network printers through Windows Explorer.
11. **Fill in the Blank:** _____ have to do with user access to shared folders and drives on the network.
12. **Multiple Choice:** Which of the following are share permissions? (*Choose two.*)
- Read
 - Change
 - Write
 - Modify
13. **Short Answer:** Discuss the Remote Access and RDS roles.
14. **Short Answer:** Explain the differences between user rights, NTFS permissions, and share permissions.

Further reading

- *IIS Web Server Overview:* <https://docs.microsoft.com/en-us/iis/get-started/introduction-to-iis/iis-web-server-overview>
- *DirectAccess:* <https://docs.microsoft.com/en-us/windows-server/remote/remote-access/directaccess/directaccess>
- *User Rights Assignment:* <https://docs.microsoft.com/en-us/windows/security/threat-protection/security-policy-settings/user-rights-assignment>
- *Exchange documentation:* <https://learn.microsoft.com/en-us/exchange/>

- *Microsoft SQL documentation:* <https://learn.microsoft.com/en-us/sql/?view=sql-server-ver16>

Part 3: Configuring Windows Server 2025

This part focuses on essential configurations in Windows Server 2025, encompassing Group Policy, virtualization, and storage technologies. By the end of this section, you will be proficient in managing Group Policy Objects (GPOs), configuring virtual machines (VMs), and effectively utilizing storage technologies.

This part contains the following chapters:

- [Chapter 6](#), *Group Policy in Windows Server 2025*
- [Chapter 7](#), *Virtualization with Windows Server 2025*
- [Chapter 8](#), *Storing Data in Windows Server 2025*

6

Group Policy in Windows Server 2025

In the previous chapters, you learned how to install, configure, and add roles and features to Windows Server 2025. Building on that foundation, this chapter delves into **Group Policy (GP)** in Windows Server 2025, a powerful tool for controlling and restricting user and computer settings across your network. You'll explore how to effectively manage GP on both local servers and domain controllers, gaining a deeper understanding of the various settings and options available within **Group Policy Objects (GPOs)** and how they are implemented.

The latter part of this chapter will focus on **Local Group Policy Editor**, where you'll learn to create and modify GPOs on individual servers. Additionally, you'll discover how to refresh and update local GPOs to help you differentiate between computer and user configurations.

To solidify your understanding, the chapter concludes with practical exercises featuring GPOs that are particularly useful for system administrators. These hands-on examples will enhance your proficiency in managing GP within Windows Server 2025.

In this chapter, we're going to cover the following main topics:

- Understanding GP fundamentals in Windows Server 2025
- Exploring GP processing mechanisms and order of precedence
- GP editors overview
- Examples of GPOs for system administrators

Technical requirements

To effectively work through the content in this chapter, it's essential to have the following hardware configuration in place.

- First, you'll need a personal computer running **Windows 11 Pro**, equipped with at least 16 GB of RAM, 1 TB of hard disk space, and a stable internet connection
- Additionally, you'll need a virtual machine configured with **Windows Server 2025 Standard** (Desktop Experience), featuring a minimum of 4 GB of RAM, 100 GB of hard disk space, and internet access

This setup ensures that you have the necessary resources to follow along with the exercises and concepts discussed in the chapter.

Understanding GP fundamentals in Windows Server 2025

System administrators often need to enforce specific configurations across an organization's network to ensure consistency and security. For example, they may set the company's website as the default home page on all browsers across the organization's computers and restrict access to removable media drives. Additionally, they might need to disable the use of Microsoft accounts on Windows 10 and 11 systems. These tasks can be efficiently managed using GP within a Windows Server environment, providing a centralized way to apply and enforce policies without relying on third-party tools or utilities.

GPO's default location

GP is a crucial feature in Windows Server that allows administrators to enforce policies at both the user and computer levels. Through GPOs, administrators can define and enforce settings that control user and computer behavior across the network. GPOs provide administrative templates that specify permissible actions and configurations for users and devices, ensuring that organizational standards and security measures are

uniformly applied. Additionally, GPOs can be leveraged as a security mechanism, applying critical security settings to users and computers in a domain-controlled network, thus enhancing the overall security posture of the organization.

Note

It is important to briefly differentiate between **Group Policy Preferences (GPP)** and **Group Policy settings (GPS)**. While GPS enforces specific configurations and settings on users and computers, GPP provides more flexibility by allowing users to modify their settings without administrative intervention. Understanding this distinction can help readers leverage both tools effectively in managing their environments.

By default, GPOs are stored in the **C:\Windows\SYSVOL\sysvol\<domain>\Policies** directory on the domain controller, as depicted in *Figure 6.1*. This default location ensures that all configured policies are systematically managed and replicated across domain controllers.

 Figure 6.1 – GPOs’ default location in Windows Server 2025

Figure 6.1 – GPOs’ default location in Windows Server 2025

With a solid understanding of GP and GPOs, the next phase involves learning how to effectively manage and configure these policies to align with organizational needs and security requirements. That will include practical exercises on creating, modifying, and troubleshooting GPOs to ensure they meet the intended administrative and security objectives.

Note

Effective troubleshooting of GPs is crucial for maintaining a well-functioning Windows Server environment. Utilizing tools such as **Resultant Set of Policy (RSOP)**, **gpresult**, and Group Policy Log Viewer can significantly enhance your ability to diagnose and resolve issues. These tools provide valuable insights into policy application, allowing administrators to identify conflicts, assess policy precedence, and ensure that configurations are applied as intended. In the upcoming sections, we

will delve into these tools, equipping you with practical strategies for effective troubleshooting within your organization.

Managing Group Policy Objects (GPOs)

Efficient management of GPOs is crucial for system administrators to enforce and standardize configurations across a network. The **Group Policy Management Console (GPMC)** is an essential tool for this task, offering a centralized interface to create, configure, and apply GPOs within a domain-based network. The GPMC is divided into two main panes: the **Forest** pane and the GPOs pane. The **Forest** pane displays the hierarchical structure of the domain, allowing administrators to navigate through **domains** and **organizational units (OUs)** effectively. The GPOs pane provides detailed tabs, including **Status**, **Linked Group Policy Objects**, **Group Policy Inheritance**, and **Delegation**. These tabs enable administrators to view the current state of GPOs, understand how policies are applied and inherited, and manage delegation settings to control who has the authority to modify policies. *Figure 6.2* demonstrates the GPMC interface as it appears on a domain controller, highlighting its key components and layout.

 Figure 6.2 – The GPM console in the domain controller

Figure 6.2 – The GPM console in the domain controller

Accessing the GPM console in Windows Server 2025 can be achieved through multiple methods, each offering different ways to launch the tool based on administrative needs. These methods will be explored in the following sections, providing a comprehensive understanding of how to efficiently use the GPMC to manage GP settings and ensure consistent policy enforcement across the network.

Note

Providing clear guidance on how to exclude or filter specific users and devices from the GP application is essential. Many environments require tailored policy configurations to meet organizational needs effectively.

Discussing the nuances of GP processing, including filtering and security settings, will empower administrators to navigate complex environments with confidence. For more insights, refer to the comprehensive resources on GP processing on Microsoft Learn: <https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/group-policy/group-policy-processing>.

Managing administrative templates

Administrative templates are a critical component of GP management in Windows Server 2025. They provide a structured way to configure and enforce policies across the environment. Understanding how to install and update these templates is essential for effective policy implementation. Here are key considerations:

Installing administrative templates

Administrative templates for Windows Server come as **.ADMX** files and are typically included in the Windows Server installation. However, to access the latest settings, especially for new features or updates, administrators should periodically download the most recent versions from the Microsoft Download Center. Once obtained, the **.ADMX** files can be placed in the Central Store for GP in the **SYSVOL** folder of the domain controller. This ensures that all GPOs can reference the latest administrative template settings.

Updating administrative templates

Keeping administrative templates up to date is crucial as Microsoft releases new templates with updates and service packs. To update templates, do the following:

1. Download the latest versions from the Microsoft website.
2. Replace the existing **.ADMX** files in the Central Store with the new versions.

3. Update the associated **.ADML** language files, which are also stored in the Central Store, reflect any new or changed settings.
4. After updating, ensure that all GPMC instances are refreshed to recognize the changes.

Best practices for template management

Effective management of administrative templates is essential for optimizing GP performance and ensuring consistent configuration across your organization:

- Regularly check for updates to administrative templates, particularly after significant Windows updates or feature releases
- Test new templates in a non-production environment to assess their impact before widespread deployment
- Maintain documentation of any changes made to the templates, including the rationale and date of updates, to aid in future audits and troubleshooting

By incorporating these practices, administrators can ensure that they are leveraging the full capabilities of GP and staying aligned with the latest configurations available in Windows Server 2025. This will enhance their ability to manage settings effectively while minimizing the risk of configuration errors.

Accessing the GPM Console via Administrative Tools

To open the GPM console, you can use the **Administrative Tools** option available in the Start menu, which provides access to a variety of management tools for configuring Windows Server features and roles:

1. Begin by clicking the Start button, then navigate to **Windows Tools** within the Start menu.
2. In the list of available tools, select **Group Policy Management**, as illustrated in *Figure 6.3*.

 Figure 6.3 – Accessing the GPM console from Windows Tools

Figure 6.3 – Accessing the GPM console from Windows Tools

Alternatively, you can quickly access the GPMC using the **Run** dialog box. This method provides a convenient shortcut for accessing **Group Policy Management** directly, which is explained in the following subsection.

Launching the GPM Console via the Run Dialog Box

Another efficient method to open the GPM console is through the **Run** dialog box. This approach allows you to access the console quickly using a simple command. Here's how to proceed:

1. Begin by pressing the Windows key + *R* simultaneously to open the **Run** dialog box.
2. In the text field of the **Run** dialog, enter **gpmc.msc** and then click **OK**, as illustrated in *Figure 6.4*. This command will immediately launch the GPMC, providing you with a centralized interface for managing GP settings.


 Figure 6.4 – Accessing the GPM console from the Run dialog box

Figure 6.4 – Accessing the GPM console from the Run dialog box

In the upcoming section, we will explore how to access the GPM console through the Server Manager menu, offering you an additional method to manage GP within your Windows Server environment efficiently.

Accessing the GPM Console from the Server Manager

The GPM console can also be launched through the **Server Manager** option found in the Start menu. This method provides access to various Windows Server management tools. Follow these steps to open the GPM console:

1. Begin by clicking the **Start** button. From the Start menu, select **Server Manager**.

2. Once the Server Manager window is open, navigate to the **Tools** menu and select **Group Policy Management**, as depicted in *Figure 6.5*. This action will open the GPMC, enabling you to manage GP settings efficiently.

 Figure 6.5 – Accessing the GPM console from the Server Manager

Figure 6.5 – Accessing the GPM console from the Server Manager

Best practices for Group Policy Management

Effective management of GP is essential for maintaining a secure and efficient Windows Server environment. By adhering to best practices, administrators can optimize GP implementation, ensure compliance, and minimize potential issues. Here are key recommendations:

- **Limit the Use of Default GPs:** Default GPs can be comprehensive and complex. Instead of relying heavily on these out-of-the-box policies, create custom GPOs tailored to your organization's specific needs. This approach reduces the risk of unintended consequences that may arise from default settings and helps streamline policy management.
- **Establish Clear Naming Conventions:** Implement a consistent naming scheme for GPOs that reflects their purpose and scope. For example, you might use a prefix to indicate the department (e.g., **HR-PasswordPolicy**) followed by a brief description of the policy. Clear naming not only simplifies identification but also aids in troubleshooting and auditing.
- **Regularly Review and Clean Up GPOs:** Periodically audit your GPOs to assess their relevance and effectiveness. Remove any obsolete or redundant GPOs to maintain a tidy and manageable policy landscape. This practice can also enhance system performance by reducing processing time during GP updates.
- **Test GPOs Before Deployment:** Always test new or modified GPOs in a controlled environment before applying them to production systems. This testing phase allows you to identify potential conflicts

and ensure that the policies behave as expected, safeguarding your network from unintended disruptions.

- **Document GPO Settings and Changes:** Maintain thorough documentation for all GPOs, including their settings, intended purpose, and any modifications made over time. Documentation serves as a reference point for future administrators and aids in understanding the policy landscape, particularly during audits or compliance reviews.
- **Implement Security Filtering and Windows Management Instrumentation (WMI) Filtering:** Utilize security filtering and WMI filtering to target GPOs to specific users, groups, or computers. This approach helps ensure that policies are only applied where necessary, reducing clutter and the risk of unintended policy application.
- **Leverage GP Modeling:** Use the GP modeling feature to simulate how policies will affect users and computers in different scenarios. This tool can help you predict the impact of changes before implementation, allowing for informed decision-making.

By following these best practices, administrators can enhance the effectiveness of GP management, ensuring a more secure, efficient, and compliant Windows Server environment.

Real-life applications of Group Policy

GP is more than just a tool for managing settings in Windows Server; it plays a crucial role in driving organizational efficiency and security. By implementing effective GP strategies, organizations can align their IT practices with business objectives, streamline operations, and enhance data protection. This subsection explores real-life scenarios where GP has been instrumental in helping organizations meet their security and operational goals, demonstrating the tangible benefits of mastering this powerful management tool.

Enhancing security through Group Policy in a financial institution

A mid-sized financial institution faced challenges with ensuring consistent security settings across its diverse network of users. To mitigate risks, the IT team implemented GPOs to enforce strict password policies, requiring all employees to use complex passwords and change them every 90 days. Additionally, they configured GPOs to disable local administrative rights on workstations, significantly reducing the potential attack surface.

The results were profound. The organization experienced a 40% decrease in security incidents within the first year. By monitoring compliance through GP reports, the IT department could demonstrate adherence to regulatory requirements, ultimately fostering a culture of security awareness among employees.

Streamlining user experience in a university

A large university wanted to streamline the user experience for its faculty and students, who often faced challenges with varying desktop configurations across campus. The IT department employed GP to establish standardized desktop settings, including specific applications pinned to the taskbar, default printer configurations, and access to shared network drives.

By implementing these GPOs, the university not only reduced help desk calls related to configuration issues by 30% but also enhanced user satisfaction. Faculty and students reported an improved workflow, as they could rely on consistent settings regardless of which computer they logged into on campus.

Managing software updates in a healthcare organization

In a healthcare organization, timely updates to medical software and systems are critical for maintaining compliance with health regulations. The IT team utilized GP to configure **Windows Server Update Services (WSUS)** settings, ensuring that all devices automatically received the latest updates and security patches.

This proactive approach minimized downtime caused by software vulnerabilities, enabling healthcare professionals to access up-to-date tools

for patient care. As a result, the organization not only achieved a higher compliance rate with health standards but also gained recognition for its commitment to patient safety and data integrity.

These examples underscore how GP can be leveraged to enhance security, improve user experience, and ensure compliance in various organizational contexts. By providing such scenarios in the chapter, readers can better understand the practical applications of GP in real-life situations.

In the following section, we will delve into the various GPO configuration settings, which will help you understand how to configure and manage GPOs in your Windows Server environment effectively.

Exploring GP processing mechanisms and order of precedence

In Windows Server administration, effectively applying GP to users and computers within a domain is essential for managing various aspects of the system, including security and application settings. To ensure that GP functions smoothly and reliably, it's important to grasp how GP is processed and the order in which policies are applied. Understanding these processes will help you prevent or address any conflicts that may arise when multiple policies are assigned to the same users or devices. In this section, we'll explore the principles behind GP processing, the hierarchy of policy application, and the rules for resolving conflicts. With this knowledge, you will be equipped to design and implement GP strategies that offer robust control over your IT environment.

Configuring GPO settings

Configuring GPO settings is a crucial task for system administrators, as it empowers them to control and manage various user and computer

behaviors within a networked environment. GPOs are essentially predefined templates that can be customized to enforce specific policies, with their settings directly corresponding to registry keys in the **Windows Registry Editor**. These settings dictate how policies are applied across the network, giving administrators control over everything from security configurations to application management.

There are three main configurations for GPO settings, as illustrated in *Figure 6.6*:

- **Not Configured:** This is the default state for a GPO. In this setting, no specific registry value is assigned, meaning the policy is inactive and has no effect on the targeted users or computers. It essentially leaves the system in its default state.
- **Enabled:** Activating this setting changes the corresponding registry value to **0x1**. That means the GPO is now active, and the policy it represents will be enforced across the network, ensuring that users and computers adhere to the defined rules.
- **Disabled:** When this setting is applied, the registry value is altered to **0x0**, deactivating the GPO. In this state, the policy is turned off, and its effects are nullified, allowing the system to operate without the constraints imposed by the policy.

 Figure 6.6 – GPO settings configuration values

Figure 6.6 – GPO settings configuration values

After configuring the desired GPO settings, the policy is ready for deployment within the network. However, to ensure that GPOs function as intended and to avoid potential conflicts, it's crucial to have a deep understanding of the order in which these policies are processed. This order of precedence plays a significant role in how GPOs interact with each other, mainly when multiple policies are applied to the same users or computers. In the following section, we will delve into the mechanisms behind GPO processing and the hierarchy that governs their application, equipping you with the knowledge needed to implement your GPO strategies effectively.

GPO application

GPOs are powerful tools that system administrators use to enforce specific configurations and behaviors across users and computers within a domain environment. These policies, which are stored within the system registry, serve as authoritative guidelines that cannot be overridden by individual users, ensuring consistency and control across the network. GPOs can be implemented at various levels, either locally on individual machines through **Local Group Policy Editor** or across the entire domain using the GPMC.

The application of GPOs follows a meticulously structured hierarchy, a key element for system administrators in maintaining an orderly and predictable environment:

- **Local GPOs:** These are the first to be applied and are the bedrock of GP enforcement. They affect all users and settings on the local machine, providing specific policies for the individual computer.
- **Site GPOs:** These follow in the order of application. These policies apply to all computers within a specific geographical site. That allows administrators to enforce policies relevant to the physical location of the computers, which can be useful in large organizations with multiple sites.
- **Domain GPOs:** These are applied after site policies. These policies affect all computers that belong to the same domain, providing a centralized way to manage and enforce settings across a broader network. This centralized management ensures that policies are consistently applied, making administrators more efficient in their roles.
- **OU GPOs:** These represent the pinnacle of policy application. These policies target specific OUs within the domain, providing a level of control that is tailored to the unique needs of each unit. This granular control empowers administrators to enforce policies that are most effective for different departments or teams within an organization.

Methods to configure GPOs

Administrators have the flexibility to configure GPOs from two perspectives:

- **Local Computer GPOs:** These are set directly on the individual computer. These policies influence local settings and are particularly useful for standalone machines or those not joined to a domain.
- **Domain Computer GPOs:** These are managed by the domain controller and affect all computers within the domain. This centralized approach ensures that policies are uniformly applied across all devices, making it easier to manage large networks.

The timing of the GPO application is a crucial consideration. GPOs targeting user accounts are applied at user logon, ensuring user-specific settings are in place before the session begins. Conversely, GPOs targeting computer accounts are enforced during system startup, ensuring machine-specific policies are active as soon as the computer boots up. This understanding of GPO timing empowers administrators to apply policies exactly when needed, minimizing disruptions and ensuring a seamless user experience.

Note

Microsoft has introduced a valuable resource for administrators working with Windows Server 2022: a GP settings reference spreadsheet. This comprehensive spreadsheet includes detailed information on administrative templates for both computer and user configurations that can be managed through GP. By using this tool, administrators can easily stay up to date with the latest GP settings, ensuring they apply the most current configurations across their network. The spreadsheet can be accessed directly from the following URL: <https://www.microsoft.com/en-us/download/details.aspx?id=104005>. This resource is indispensable for maintaining effective and consistent policy management in a Windows Server environment.

In this section, we have covered the foundational aspects of GP in Windows Server 2025, including how to manage GPOs using the GPMC, the different types of configuration settings, and the hierarchical order in which GPOs are applied. As we move forward, the next section will focus

on Local Group Policy Editor, providing detailed guidance on how to update and manage local GPOs and settings for both users and computers, further enhancing your ability to maintain a well-managed and secure IT environment.

GP editors overview

System administrators managing Windows Server 2025 have several tools at their disposal for configuring GPO settings, depending on whether the policies need to be applied locally or across a domain. Local Group Policy Editor is the **Microsoft Management Console (MMC)** snap-in that allows administrators to manage local GPOs for individual users and computers on a single machine. This tool comes pre-installed with Windows Server 2025, making it readily accessible for local policy management. In contrast, the **Group Policy Management Editor**, another MMC snap-in, requires the installation of **Active Directory Domain Services (AD DS)** and the setup of a **domain controller**. This tool is essential for administrators who need to create and modify domain-based GPOs, linking them to specific sites, domains, or OUs. Together, these editors provide a comprehensive framework for effective GPO management across both local and domain environments.

Local Group Policy Editor

Local Group Policy Editor is not just a tool but an essential component for system administrators who manage Windows Server 2025, especially when dealing with servers that are not part of a domain. This tool provides a straightforward way to configure and enforce GPO settings on individual servers, ensuring that specific policies are applied consistently across user and computer configurations. Unlike domain-based GPOs, which are managed centrally, local GPOs are edited directly on the server, making Local Group Policy Editor particularly useful for standalone or isolated systems.

Accessing Local Group Policy Editor is a quick and easy process for administrators, adding to the convenience and efficiency of their tasks.

By pressing the Windows key + *R*, they open the **Run** dialog box, where typing **gpedit.msc** and clicking **OK** will launch the editor.

Once open, Local Group Policy Editor presents a straightforward interface, as depicted in *Figure 6.7*, where administrators can navigate and modify both user and computer policies.

 Figure 6.7 – The Local Group Policy Editor in Windows Server 2025

Figure 6.7 – The Local Group Policy Editor in Windows Server 2025

This tool allows for granular control over the server's settings, making it possible to enforce security measures, restrict access to certain features, or configure specific applications, all without requiring a domain controller. Additionally, understanding how to update these local GPOs is crucial for maintaining compliance and ensuring that the latest policies are enforced.

The following subsection will delve into the process of updating GPOs, ensuring that administrators can keep their servers aligned with organizational requirements.

Applying local GPOs

To ensure that your configured local GPOs are applied effectively on your server, you can update them by following these steps:

1. Begin by pressing the Windows key + *R* to open the **Run** dialog box.
2. In the dialog box, type the **gpupdate /force** command and click **OK**.
This command, as depicted in *Figure 6.8*, forces an immediate update of the GPOs.


 Figure 6.8 – Running the gpupdate /force command via the Run dialog box

Figure 6.8 – Running the gpupdate /force command via the Run dialog box

3. Once executed, a Command Prompt window will appear, displaying the **Updating policy...** message, as illustrated in *Figure 6.9*. That indicates that the system is processing the updates to apply the new or modified policies.

 Figure 6.9 – The process of deploying the policy

Figure 6.9 – The process of deploying the policy

The **gpupdate /force** command is particularly useful because it ensures that all GPOs, including those that have yet to be applied, are enforced without delay. This process includes both computer and user configurations, ensuring that any changes you've made to the local policies are immediately reflected across the system.

During this process, the Command Prompt window remains open to show the progress of the policy update. Once the update is completed and the system has successfully applied the policies, the window will close automatically, signaling that your local GPOs have been fully updated and are now active.

Note

On Windows Server 2025, the Local Group Policy Editor tool can be accessed through various methods, including using the search box on the Start menu or by typing **gpedit.msc** in Windows PowerShell or the Command Prompt. While these methods are convenient for opening the editor, special consideration should be given when updating GPO settings. Although the **gpupdate /force** command is commonly used to apply GP updates, it re-applies all policies, which can create unnecessary system load. Instead, it is often more efficient to use the **gpupdate** command without the **/force** option, as this will only update the GPOs that have been modified, reducing the strain on system resources.

This procedure is crucial for administrators who need to quickly enforce new policies or troubleshoot existing configurations on a local server, ensuring that all settings are properly applied and effective without the need for a system reboot.

Exploring GPO settings categories

GPOs consist of two primary categories of settings that specifically target users and computers. These settings play a crucial role in shaping the behavior and security of the server environment. In the following sections, we will delve deeper into these categories to understand their impact in greater detail. The settings are not randomly scattered, but they are systematically arranged in a hierarchical structure of folders and subfolders, which can be navigated using Local Group Policy Editor or the Group Policy Management Editor. This structured approach allows administrators to efficiently manage and configure policies that enhance the overall security and performance of the server.

Configuring GPOs for computers

To configure GPOs at the computer level, which apply settings across the entire machine regardless of the user logged in, follow these steps:

1. Open the **Run** dialog by pressing the Windows key + *R*.
2. Enter **gpedit.msc** in the **Run** dialog box and click **OK** to access the Group Policy Editor.
3. In the GPMC's **Forest** pane, right-click the domain, select **Create a GPO in this domain, and Link it here...**, as illustrated in *Figure 6.10*.

 Figure 6.10 – Creating a GPO in the domain controller

Figure 6.10 – Creating a GPO in the domain controller

4. Assign a name to the new GPO in the **New GPO** window and click **OK** to proceed.
5. Navigate to the GPOs pane, and under the **Linked Group Policy Objects** tab, right-click the newly created GPO and choose **Edit**.
6. In the Group Policy Management Editor, expand **Policies** under **Computer Configuration** and select the administrative template you wish to configure, as shown in *Figure 6.11*.


 Figure 6.11 – Configuring the GPO for computers

Figure 6.11 – Configuring the GPO for computers

7. After making the necessary changes, click **Close** to exit the Group Policy Management Editor.
8. Return to the GPOs pane, right-click the newly created GPO, and select **Enforced**.
9. Confirm your selection by clicking **OK** in the **Group Policy Management** dialog box.

These steps ensure that computer-level GPOs are applied during the system's boot-up process, affecting both local and domain-based machines as required.

Configuring GPOs for users

User configuration GPOs are settings specifically designed to impact the user level, regardless of the computer the user logs into. These settings are consistently applied to the user account across different devices. To configure GPOs at the user level, follow these steps:

1. Begin by following the initial steps in the *Configuring GPOs for computers* section.
2. In the Group Policy Management Editor, navigate to **Policies** under **User Configuration**, and then select the appropriate user administrative template you wish to configure, as illustrated in *Figure 6.12*.

 Figure 6.12 – Configuring the GPO for users

Figure 6.12 – Configuring the GPO for users

3. Once the configuration is complete, click **Close** to exit the Group Policy Management Editor.
4. In the GPOs pane, right-click on the newly created GPO and choose **Enforced**.

5. Confirm by clicking **OK** in the **Group Policy Management** dialog box.

Through these steps, you have learned how to manage user-level GPOs using Local Group Policy Editor. Additionally, you have gained insight into the different types of configuration policies. In the following section, we will explore examples of GPOs that are particularly useful for system administrators.

Note

The Resultant Set of Policy (RSOP) tool, accessed via **rsop.msc**, is invaluable for users looking to understand which GPs are applied or missing on their systems. RSOP allows administrators to simulate and analyze policy settings, providing clarity on how policies interact and the resultant configurations for users and computers. This tool enhances troubleshooting and management by ensuring that the intended policies are effectively enforced in the environment.

Chapter exercise – examples of GPOs for system administrators

In this chapter's exercise, you will be introduced to several key GPOs that are highly beneficial for system administration. You will gain hands-on experience in applying GPOs to perform tasks such as renaming the administrator account, disabling the guest account, blocking the use of Microsoft accounts, restricting access to the Control Panel and PC settings, and preventing the use of removable media drives. These GPOs are crucial for improving the security and manageability of your network. Additionally, you will learn how to use the GPMC to create, link, and enforce these GPOs within your domain, ensuring a more secure and controlled environment.

Renaming the administrator account

Renaming the default administrator account is not just a step but a strategic move in bolstering network security within Windows Server 2025. By changing the name of this account through a GPO, you obscure the default administrator identity, which significantly reduces the risk of unauthorized access and targeted attacks. Default administrative accounts are often a prime target for attackers, as they are widely recognized and expected. By renaming this account, you make it less obvious and, therefore, less vulnerable to brute force or other attack methods aimed at exploiting default credentials:

1. Access the GPO you wish to modify and navigate to **Computer Configuration | Policies | Windows Settings | Security Settings | Local Policies | Security Options**.
2. Locate and double-click on **Accounts: Rename the administrator account**, as illustrated in *Figure 6.13*.


Figure 6.13 – Renaming the administrator account

Figure 6.13 – Renaming the administrator account

3. In the **Properties** dialog box, enable the policy setting by selecting **Define this policy setting** and entering the new name for the administrator account.
4. Click **OK** to apply the changes and close the **Properties** dialog box.

Applying this policy ensures that the new name is enforced across all relevant machines, thereby providing a layer of obscurity and defense against potential intruders who may seek to exploit well-known default account names. This practice is part of a broader strategy to protect your network's infrastructure and resources, ensuring they are not easily accessible to unauthorized users or malicious actors. Following this, we will address how to similarly enhance security by renaming the guest account through a GPO, further securing your network environment.

Renaming the guest account

Renaming the default guest account is an important security measure to enhance network protection within Windows Server 2025. By changing the name of the guest account through a GPO, you help mitigate risks associated with unauthorized access and potential misuse of this account. The guest account is often targeted because its default status can be easily identified and exploited. Therefore, modifying its name adds an extra layer of security by making it less recognizable to potential attackers:

1. Access the GPO you wish to modify and navigate to **Computer Configuration | Policies | Windows Settings | Security Settings | Local Policies | Security Options**.
2. Locate and double-click on **Accounts: Rename guest account**, as shown in *Figure 6.14*.

 Figure 6.14 – Renaming the guest account

Figure 6.14 – Renaming the guest account

3. In the **Properties** dialog box, enable the policy setting by selecting **Define this policy setting** and entering the new name for the administrator account.
4. Click **OK** to apply the changes and close the **Properties** dialog box.

Implementing this policy helps obscure the guest account's identity, thereby reducing the likelihood of it being targeted for unauthorized access. By effectively managing and renaming such accounts, you strengthen your network's overall security posture. In the next section, we will explore how to block Microsoft accounts using GPOs to secure your system further.

Blocking the Microsoft accounts

You can also use a GPO in Windows Server 2025 to prevent users from using Microsoft accounts to access or add to the network computers. That can enhance the security of your network by restricting the use of personal accounts that the organization does not manage. To implement this policy, you need to do the following:

1. Access the GPO you wish to modify and navigate to **Computer Configuration | Policies | Windows Settings | Security Settings | Local Policies | Security Options**.
2. Locate and double-click on **Accounts: Block Microsoft accounts** (see *Figure 6.15*).


 Figure 6.15 – Blocking the Microsoft accounts

Figure 6.15 – Blocking the Microsoft accounts

3. In the **Properties** dialog box, enable the policy setting by selecting **Define this policy setting** and entering the new name for the administrator account.
4. Click **OK** to apply the changes and close the **Properties** dialog box.

This policy applies to the user configuration settings and blocks the use of Microsoft accounts on the network computers. By doing this, you can prevent users from accessing or adding their accounts, which may compromise the security of your network services and resources. Next, we will see how to use a GPO to deny access to the Control Panel and PC settings.

Prohibiting access to the Control Panel and PC settings

To bolster network security and control system settings, you can use a GPO in Windows Server 2025 to restrict user access to the Control Panel and PC settings. This measure helps prevent unauthorized or unintended alterations to system configurations, thus enhancing overall security and maintainability:

1. Access the GPO you wish to modify and navigate to **User Configuration | Policies | Administrative Templates | Control Panel**.
2. Locate and double-click on **Prohibit access to the Control Panel and PC settings** (depicted in *Figure 6.16*).


 Figure 6.16 – Prohibiting access to the Control Panel and PC settings

Figure 6.16 – Prohibiting access to the Control Panel and PC settings

3. In the **Properties** dialog box that appears, select the **Enable** option to activate the policy. Confirm your changes by clicking **OK**.

This configuration ensures that only authorized administrators can access or modify critical system settings, thereby protecting the integrity of your network environment. The policy will be applied to all users linked to the GPO, effectively preventing them from accessing or altering the Control Panel and PC settings. In the subsequent section, we will explore how to use a GPO to restrict access to removable media drives, further securing your network against potential threats.

Denying access to all removable storage classes

Another way to improve the security and management of your network is to use a GPO in Windows Server 2025 to disable access to all removable storage devices. That can prevent users from transferring data to or from external drives, which may pose a security risk. To enable this policy, you need to do the following:

1. Access the GPO that you want to modify and navigate to this location: **User Configuration | Policies | Administrative Templates | System | Removable Storage Access**.
2. Select **All the removable storage classes: Deny all access**, as shown in *Figure 6.17*, and double-click it to open it.


 Figure 6.17 – Denying access to all removable storage classes

Figure 6.17 – Denying access to all removable storage classes

3. In the **Properties** dialog box, mark the **Enable** option and click **OK** to exit the dialog box.

This policy applies to user configuration settings and blocks the use of removable storage devices on network computers. You can also enhance this policy by enabling the **Prohibit access to the Control Panel and PC settings** policy, which prevents users from changing system settings.

These are just some of the GPO configuration examples that you can use in Windows Server 2025. There are many more GPOs that you can explore and customize according to your needs. These GPOs can help you manage and secure your network services and resources more effectively. That concludes this chapter exercise.

Summary

In this chapter, you were introduced to the fundamental concept of GP in Windows Server 2025, a crucial tool for managing and enforcing user and computer settings across your network. You gained insights into the GPMC, which is essential for creating, editing, and managing GPOs at the domain level. Additionally, you explored Local Group Policy Editor, which provides a method for configuring policies on individual machines that is useful for scenarios where domain-based GPOs are not applicable. The chapter provided an in-depth look at various computer and user configuration settings that can be applied through GPOs, including security policies, account settings, and restrictions. You also learned how to refresh GPOs, ensuring that updates and changes are applied correctly on both local servers and domain controllers. This skill is vital for maintaining up-to-date configurations and policies across your network.

Furthermore, practical examples were provided to illustrate the application of GPOs in real-life scenarios, such as renaming accounts, restricting access to system settings, and managing device usage. These examples are designed to help you understand how to implement and enforce security measures and administrative controls effectively. With these foundational skills, you are now equipped to deploy and manage GPOs on both individual computers and across a domain, enhancing the security and efficiency of your network. The next chapter will build on this knowledge by exploring the topic of virtualization within Windows Server 2025,

offering insights into how virtualization technologies can further optimize and streamline your IT infrastructure.

Questions

1. **True or False:** GPOs are processed in the following order: local, site, domain, and OUs.
2. **Fill in the Blank:** The _____ are administrative templates that enable system administrators to configure what users can and cannot do on computers, peripheral devices, and network applications across the organization's network.
3. **Multiple Choice:** Which of the following represent GPO configuration values? (*Choose two.*)
 1. Enabled
 2. Disabled
 3. Allow
 4. Deny
4. **True or False:** The GPM is a console in the domain controller that enables configuring and deploying GPOs across an organization.
5. **Fill in the Blank:** The _____ displays the hierarchical structure of the domain, whereas the _____ contains the Status, Linked Group Policy Objects, Group Policy Inheritance, and Delegation tabs.
6. **Single Choice:** Which of the following commands is used to update GPOs?
 1. **gpupdate /enforce**
 2. **gpupdate /setup**
 3. **gpupdate /run**
 4. **gpupdate /force**
7. **True or False:** Not configured is the default setting for GPOs, meaning that the registry value has not been manipulated.
8. **Fill in the Blank:** The _____ is another MMC snap-in that enables you to manage GPO settings on a local computer.
9. **Single Choice:** GPOs assigned to the computer level are applied when computers are in which of the following states?
 1. Turned on

2. Turned off
3. In hibernate mode
4. In sleep mode

Further reading

- *Group Policy Objects:* <https://docs.microsoft.com/en-us/previous-versions/windows/desktop/policy/group-policy-objects>
- *Linking GPOs to Active Directory Containers:* <https://docs.microsoft.com/en-us/previous-versions/windows/desktop/policy/linking-gpos-to-active-directory-containers>
- *Applying Group Policy:* <https://docs.microsoft.com/en-us/previous-versions/windows/desktop/policy/applying-group-policy>.
- *Group Policy Best Practices:* https://www.netwrix.com/group_policy_best_practices.html

Virtualization with Windows Server 2025

Cloud computing has become a prominent and rapidly evolving trend in the IT industry. It enables the delivery of services and resources over the internet from a network of interconnected servers. At the heart of cloud computing lies a critical technology known as **virtualization**. This technology allows multiple **virtual machines (VMs)** to operate on a single physical server or across a cluster of servers, effectively creating a virtualized infrastructure that underpins cloud services.

In this chapter, you will delve into the concept of virtualization and its role in cloud computing. You will focus on using **Hyper-V**, a Microsoft technology that facilitates virtualization for both Windows clients and servers. The chapter will guide you through the process of installing the Hyper-V role on Windows Server 2025, utilizing the **Hyper-V Manager** for managing virtual environments and creating and configuring VMs. These steps will provide you with a solid understanding of virtualization principles and practical skills for implementing Hyper-V.

By the end of this chapter, you will be equipped with the knowledge to activate the Hyper-V role, set up VMs, and manage virtual environments effectively. The chapter concludes with a hands-on exercise designed to reinforce your learning by walking you through the installation of the Hyper-V role on Windows Server 2025. This practical experience will be crucial in understanding how virtualization enhances cloud computing and IT infrastructure.

In this chapter, we're going to cover the following main topics:

- Understanding virtualization fundamentals in Windows Server 2025
- Adding and configuring the Hyper-V role on Windows Server 2025
- Exploring Hyper-V Manager for VM administration
- Installing the Hyper-V role on Windows Server 2025

Technical requirements

To practice the skills learned in this chapter, you will need the following resources:

- A computer with Windows 11 Pro, a minimum of 16 GB of RAM, 1 TB of disk space, and an Internet connection

Understanding virtualization fundamentals in Windows Server 2025

Virtualization is a transformative technology that allows for the creation and operation of multiple VMs on a single physical server or a network of interconnected servers, known as a cluster. Each VM acts as an independent computer, complete with its own operating system, applications, and allocated resources, operating in isolation from other VMs. This technology also extends to virtualizing storage devices and network resources, which enhances the flexibility and efficiency of the virtualized environment. By consolidating workloads onto fewer physical servers, virtualization can significantly reduce hardware costs, lower power consumption, and minimize the physical space required for data centers.

Windows Server 2025 incorporates Hyper-V, a powerful virtualization feature that enables effective deployment and management of VMs on both Windows client systems and server environments. Hyper-V, a successor to the earlier **Windows Virtual PC**, has evolved from its inception in Windows Server 2008 to become a widely adopted

and highly regarded platform among system administrators. Its robust suite of services and tools supports the creation, configuration, and administration of VMs, providing a comprehensive solution for managing virtual environments.

In the sections that follow, you will gain insights into various aspects of server virtualization, starting with an examination of different virtualization modes. That will include an exploration of how Hyper-V integrates into the broader virtualization landscape, the benefits of its features, and practical applications to optimize your IT infrastructure.

Emphasizing the connection between Hyper-V and cloud computing

Hyper-V, Microsoft's robust virtualization technology, serves as a critical building block for cloud computing infrastructure. As organizations increasingly adopt cloud services, understanding Hyper-V equips IT professionals with the foundational skills necessary to design, deploy, and manage virtualized environments efficiently:

- **Foundation of Cloud Infrastructure:** Hyper-V facilitates the creation and management of VMs on a physical server, enabling efficient resource utilization. This virtualization capability is essential for cloud environments, where resources must be dynamically allocated to meet varying workloads. By mastering Hyper-V, IT professionals can leverage this technology to create scalable, flexible cloud infrastructures.
- **Integration with Microsoft Azure:** Microsoft Azure, one of the leading cloud platforms, heavily relies on virtualization principles similar to those employed by Hyper-V. Understanding Hyper-V enables professionals to seamlessly transition their on-premises virtualization skills to Azure, where they can deploy Azure Virtual Machines and utilize features such as Azure Site Recovery for disaster recovery. This familiarity fosters a smoother migration process and enhances overall cloud management capabilities.
- **Cost Efficiency and Resource Management:** The skills learned through Hyper-V, such as VM configuration and resource allocation, directly translate to cost management in cloud computing. Professionals can apply these skills to optimize cloud resource usage, ensuring that organizations only pay for what they need. This financial acumen is invaluable as companies strive to balance performance and costs in their cloud strategies.
- **Enhanced Disaster Recovery Solutions:** Hyper-V offers features such as virtual machine replication, which are integral to establishing disaster recovery plans. Understanding these features not only prepares IT professionals to implement reliable backup and recovery solutions but also enhances their ability to design resilient cloud architectures that can withstand failures, ensuring business continuity.
- **Skills in Hybrid Environments:** As businesses increasingly adopt hybrid cloud models, where on-premises infrastructure coexists with cloud resources, expertise in Hyper-V becomes even more critical. Knowledge of how to manage VMs locally can simplify the integration of cloud services, enabling professionals to create a unified environment that optimally leverages both on-premises and cloud resources.

By developing foundational skills in Hyper-V, IT professionals not only enhance their technical proficiency but also position themselves as valuable assets to organizations navigating the complexities of cloud computing. The knowledge gained from mastering Hyper-V will catalyze exploring advanced cloud technologies and strategies, ultimately leading to more effective IT management in an increasingly cloud-centric world.

Virtualization modes

Virtualization allows for the operation of multiple **operating systems (OSes)** on a single physical server or across a cluster of servers by leveraging different modes. Each mode offers distinct features and benefits tailored to various needs in virtualized environments:

- **Fully Virtualized Mode:** In this mode, each OS runs in its own isolated and secure virtual environment, as if it were on a separate physical machine. The **virtualization layer**, or **hypervisor**, manages the resources of the host server and allocates them to each VM. This approach provides robust isolation between VMs, ensuring they can operate independently without altering their configurations. Fully virtualized environments are ideal for scenarios requiring robust security and separation, as the guest OSes remain unaware of the underlying virtualization infrastructure. *Figure 7.1* depicts how Windows Server 2025 operates within such an isolated environment, emphasizing the separation of resources and processes between different VMs.

Figure 7.1 – Windows Server 2025 running in an isolated and secure virtual environment

- **Paravirtualized Mode:** Paravirtualization offers a more integrated approach by allowing the guest OS to communicate directly with the hypervisor. Unlike fully virtualized mode, where the guest OS is unaware of the virtualization layer, paravirtualized systems require the guest OS to be modified to interact efficiently with the hypervisor. This mode employs an **Application Program Interface (API)** to enable direct communication, which reduces the overhead typically associated with hardware emulation. The result is a significant boost in performance and resource utilization, making it a compelling choice for environments where efficiency and speed are paramount.
- **Containerization Mode:** Containerization focuses on encapsulating applications along with their runtime environments, system tools, and settings into self-contained units known as **containers**. Unlike VMs, which virtualize entire operating systems, containers virtualize at the application layer, providing a lightweight and portable solution. Each container operates independently but shares the host OS kernel, making it an efficient choice for deploying and managing applications across different environments. Containers enhance scalability, streamline application deployment, and ensure consistency by packaging all necessary components together. This approach is particularly useful for developing, testing, and consistently deploying applications across various platforms.

Note:

A physical server operates with an operating system known as the host OS, which manages the server's hardware resources. In contrast, a VM runs an operating system referred to as the guest OS, which operates within a virtual environment created by the host OS or a hypervisor. For example, in my configuration, my laptop is equipped with Windows 11 Pro as the host OS, while a VM on the same machine runs Windows Server 2025 Standard as the guest OS. The host OS plays a critical role in allocating and controlling hardware resources, enabling the guest OS to function seamlessly within the virtualized environment.

Performance considerations in virtualization

In virtualization, performance is a critical factor that can significantly influence the effectiveness of deployed VMs. Understanding the impact of different storage types and infrastructure configurations is essential for optimizing VM performance.

Storage types – SANs vs. Local Disks

When designing a virtualized environment, understanding the differences between **Storage Area Networks (SANs)** and local disks is crucial, as each storage solution presents unique performance characteristics and implications for VM operations:

- **Storage Area Networks (SANs):** SANs provide a centralized storage solution that allows multiple servers to access a shared pool of storage resources. While SANs offer benefits such as high availability and scalability, they can introduce latency that affects performance. This latency arises from network overhead, as data must travel over the network fabric to reach the storage array. In high-demand environments where rapid data access is crucial, this delay can result in slower VM performance, particularly for I/O-intensive applications.
- **Local Disks:** In contrast, local disks are directly attached to the physical server hosting the VM. This configuration typically yields lower latency, as data does not need to traverse a network. Local storage is ideal for applications requiring fast data access and high throughput, such as databases and transaction processing systems. However, local disks limit redundancy and scalability compared to SANs, making them less suitable for larger virtualized environments.

Hyper-Converged Infrastructure (HCI)

Hyper-converged infrastructure (HCI) integrates computing, storage, and networking into a single software-driven solution, offering a unified approach to virtualization. HCI can enhance performance in several ways:

- **Improved Data Locality:** By utilizing local storage within each node of the HCI cluster, data access times are significantly reduced. This locality minimizes latency, providing faster VM performance, especially for applications with high I/O requirements.
- **Scalability and Elasticity:** HCI allows organizations to scale their resources seamlessly. As demand increases, additional nodes can be added to the cluster, effectively distributing workloads and enhancing performance without compromising the integrity of the storage system.
- **Intelligent Resource Management:** Many HCI solutions incorporate advanced analytics and resource management features that optimize workload placement based on performance metrics. This capability ensures that VMs are allocated to the most appropriate resources, further enhancing overall performance.

The choice of storage type—whether SANs or local disks—directly impacts the performance of virtualized environments. Additionally, leveraging HCI can provide significant advantages in managing resources efficiently and ensuring optimal performance for VMs. Understanding these factors enables IT professionals to make informed decisions that align with their organization's performance requirements and business objectives.

Understanding these virtualization modes allows you to select the most appropriate method for your infrastructure, depending on your performance, security, and scalability requirements.

Adding and configuring the Hyper-V role on Windows Server 2025

A significant advantage of server virtualization is its ability to run multiple VMs on a single physical server while maximizing performance and resource efficiency. Hyper-V, a robust and versatile technology, allows you to create, manage, and operate VMs seamlessly on Windows Server 2025. In this section, we will guide you through the process of adding the Hyper-V role to your server and configuring it to meet your organization's specific needs. Whether your goal is to reduce operational costs, enhance system efficiency, or establish a virtualized testing environment, mastering the installation and configuration of Hyper-V on Windows Server 2025 is essential for modern IT infrastructure management.

Hyper-V architecture

To fully comprehend the Hyper-V architecture, it's helpful to picture it as a tree structure, with the hypervisor acting as the root and deeply integrated into the hardware foundation, which serves as the soil. The **hypervisor** is the fundamental component of the Hyper-V virtual platform, with direct access to the physical server's hardware resources, including CPU, memory, storage, and networking. This direct control allows the hypervisor to manage and allocate these resources efficiently among multiple VMs.

From the hypervisor, branches extend to form separate execution environments known as **partitions**. Each partition is isolated, meaning it operates independently without interference from others. This isolation is crucial for security and stability, as it ensures that an issue in one partition does not affect others. The partitions themselves do not have direct access to the physical hardware; instead, they interact with a virtualized layer provided by the hypervisor or the root partition. This virtualized layer abstracts the hardware, presenting the guest operating systems with a consistent and manageable environment in which to operate.

The **root partition** is the first and most privileged partition, running both the host operating system and the Hyper-V role. It acts as the central hub, managing hardware interactions and overseeing the creation and operation of other partitions, known as child partitions. These child partitions host the guest operating systems, which can include various versions of Windows or Linux, allowing for a diverse and flexible computing environment.

Communication between the root and **child partitions** is facilitated by specialized components known as the **Virtualization Service Provider (VSP)** and the **Virtualization Service Consumer (VSC)**, as illustrated in

Figure 7.2. These components use a logical communication channel called the **Virtual Machine Bus (VMBus)** to exchange data and commands efficiently. This setup ensures that the guest operating systems in the child partitions can perform necessary operations, such as accessing storage or network resources, without direct hardware access.

Figure 7.2 – Hyper-V architecture

Understanding this architecture is key to appreciating how Hyper-V supports different types of virtualization, including full virtualization, where the guest operating system operates unchanged within the virtual environment. This capability is critical for scenarios requiring compatibility and minimal modification of existing software. However, before deploying Hyper-V, it's essential to be aware of the specific hardware and software requirements, as well as the prerequisites, to ensure a successful implementation and optimal performance.

Hyper-V installation requirements

Before installing and utilizing Hyper-V, it is crucial to ensure that your server meets the specific prerequisites necessary for enabling the hypervisor. The primary requirement is that your server must support **hardware-level virtualization**, which is the backbone of Hyper-V's functionality. That involves having a processor with virtualization technology enabled, such as **Intel VT-x** or **AMD-V**. These technologies are essential as they allow the processor to efficiently manage multiple VMs by allocating resources dynamically and securely without significant overhead.

In addition to virtualization support, your server must also have other features enabled, such as **Data Execution Prevention (DEP)**, which provides an additional layer of security by preventing malicious code from executing in protected memory regions. Furthermore, the server's BIOS or UEFI firmware must be configured correctly to support these technologies, with virtualization options such as **Intel VT** or **AMD-V** turned on.

Another important consideration is **nested virtualization**, particularly in environments where you intend to run Hyper-V inside a VM. This advanced feature enables you to create virtual environments within VMs, providing flexibility for testing, development, and training scenarios without needing additional physical hardware.

Beyond these hardware requirements, it's also vital to ensure that your operating system is compatible with Hyper-V. For instance, Hyper-V is only available on specific editions of Windows Server and Windows client operating systems, such as **Windows 11 Pro** or **Enterprise**. Ensuring that your server meets these conditions is key to achieving a smooth and efficient Hyper-V deployment, allowing you to leverage the benefits of virtualization in your IT environment fully.

Nested virtualization

Nested virtualization is an advanced feature that allows you to run a VM inside another VM. Essentially, this means that the hardware of the host machine is capable of running Hyper-V within a guest operating system, enabling the guest OS to create and manage additional VMs, just as if it were running directly on physical hardware. This capability, though it might initially seem theoretical, has been supported by Microsoft since Windows Server 2016 and has become an invaluable tool in various scenarios, such as testing complex configurations, training environments, or running virtual labs where multiple layers of virtualization are required.

To put it simply, nested virtualization allows you to treat a guest operating system as if it were the host operating system, running Hyper-V and creating a virtualized environment within an already virtualized environment. This nested setup can be particularly useful for scenarios that require multiple isolated environments or for developers and IT professionals who need to test deployments and configurations in a safe and controlled manner.

Setting up nested virtualization in Windows Server 2025 is straightforward using **Windows PowerShell**. Begin by right-clicking the **Start** button and selecting **Windows PowerShell (Administrator)** from the admin menu. Once in the PowerShell window, execute the following command:

```
Set-VMProcessor -VMName <YourVMName> -ExposeVirtualizationExtensions $true
```

This command enables the guest VM to expose the necessary virtualization extensions, allowing it to run Hyper-V.

```
Get-VMNetworkAdapter -VMName <YourVMName> | Set-VMNetworkAdapter -MacAddressSpoofing On
```

This command enables MAC address spoofing on the VM's network adapter, which is required for networking functionality in a nested virtualization scenario.

After these configurations, you can proceed with installing Hyper-V within the guest VM, following the instructions provided later in the chapter that discuss installing Hyper-V on Windows Server 2025.

In summary, this section has covered various aspects of virtualization, including virtualization modes, the architecture of Hyper-V, installation requirements, and the concept of nested virtualization. With this knowledge, you're now prepared to learn how to use Hyper-V Manager effectively, which will be covered in the next section.

Exploring Hyper-V Manager for VM administration

Hyper-V Manager is a versatile and essential tool for administering VMs within a Windows Server 2025 environment. It provides a centralized interface for managing a variety of VM-related tasks, streamlining the administration of virtualized resources. With Hyper-V Manager, you can efficiently create new VMs, import existing ones, and delete those that are no longer needed, thus offering flexibility in managing your virtual infrastructure. The tool also allows you to set up and manage virtual switches, which are critical for connecting VMs to your network and ensuring they can communicate effectively with other network resources. Additionally, Hyper-V Manager facilitates the creation of a **Storage Area Network (SAN) manager**, enabling VMs to connect to shared storage solutions. This capability is vital for maintaining high availability and performance across your virtual environment.

Furthermore, Hyper-V Manager includes features for inspecting and optimizing virtual disks, allowing you to adjust disk space allocation and improve performance based on your requirements. You can also manage the state of your VMs by stopping them or shutting them down, which is helpful for maintenance and troubleshooting purposes.

The interface of Hyper-V Manager in Windows Server 2025, depicted in *Figure 7.3*, is organized into five main sections: **the server pane**, which displays the list of servers; **the VM pane**, which shows the VMs on the selected server; **the checkpoint pane**, which provides access to VM checkpoints for restoring states; **selected VM details**, which offers information and settings for the currently selected VM; and **the Actions pane**, which provides access to various management actions. Each of these components plays a crucial role in managing and configuring your virtual environment efficiently.

Figure 7.3 – The Hyper-V Manager in Windows Server 2025

Understanding key Hyper-V Manager functions

As users navigate through Hyper-V Manager, it is essential to familiarize themselves with specific **user interface (UI)** elements that play a critical role in managing virtual environments. This section focuses on two important features: **Replication Health** and the ability to export VMs.

Replication Health

The **Replication Health** column within the **Virtual Machines** window provides valuable insights into the status of VM replication. This feature is crucial for maintaining data integrity and availability, especially in disaster

recovery scenarios:

- **Importance of Replication:** Replication allows for the duplication of VMs to a secondary host, ensuring that a backup is readily available in the event of a failure. The **Replication Health** status can display indicators such as **Normal**, **Warning**, or **Critical**, each representing the current health of the replication process.
- **Monitoring and Action:** Regularly monitoring **Replication Health** can help identify and address potential issues proactively. A **Normal** status indicates that the replication is functioning correctly, while warnings or critical alerts necessitate immediate investigation to safeguard data availability in failover situations.

Exporting VMs

Another vital function available in Hyper-V Manager is the capability to export a virtual machine to another host. This feature is essential for resource management and operational continuity:

- **Export Process Overview:** Exporting a VM involves creating a complete copy, including its configuration settings, virtual hard disks, and any associated snapshots. This functionality allows for seamless migration of VMs between different Hyper-V hosts.
- **Benefits for Administrators:** Understanding how to export VMs equips IT professionals with the skills necessary to adapt their virtual environments to evolving organizational needs. This flexibility ensures minimal disruption during maintenance or when balancing workloads across hosts.

By grasping the significance of these UI functions—**Replication Health** and VM exporting, IT professionals can enhance their proficiency in using Hyper-V Manager. This knowledge not only improves the ability to manage virtualized environments effectively but also prepares professionals for real-world challenges in cloud computing and IT infrastructure management.

In the following section, we will delve into the process of configuring Hyper-V settings that will influence the VMs you deploy, ensuring that your virtual infrastructure is optimized to meet your specific operational needs and organizational goals.

Configuration settings in Hyper-V

Before you embark on creating and managing virtual machines, it's essential to configure the Hyper-V settings on your server effectively. These settings are accessible through the **Hyper-V Settings...** option found in the **Actions pane** of Hyper-V Manager. As depicted in *Figure 7.4*, several key configuration areas can be tailored to optimize your virtual environment:

- **Virtual Hard Disks:** This setting allows you to specify a default directory for storing virtual complex disk files. Proper configuration here is vital for maintaining organized storage and ensuring that disk space is used efficiently across your VMs.
- **Virtual Machines:** This option enables you to define the default location for VM configuration files. By setting this location, you ensure that all VM settings and metadata are managed in a centralized and easily accessible manner.
- **Physical GPUs:** Here, you can designate which **graphical processing unit (GPU)** will be assigned to your VMs. That is particularly important for workloads requiring high graphical performance, such as graphic-intensive applications or virtual desktops.
- **NUMA Spanning:** This setting controls whether VMs can span multiple **Non-Uniform Memory Access (NUMA)** nodes. NUMA spanning can enhance virtual machine performance by allowing them to access a broader range of computing resources, which is especially useful for environments running multiple VMs or high-performance applications.
- **Storage Migrations:** This configuration allows you to set the maximum number of concurrent storage migrations that your server can handle. Efficient storage migration management is crucial for maintaining performance and minimizing downtime during data movement operations.
- **Enhanced Session Mode Policy:** This option enables or disables the ability to redirect local devices and resources (such as printers, USB drives, or local disks) from the host machine to **Virtual Machine**

Connection (VMConnect). This feature enhances the user experience by allowing seamless interaction between the host and the VM.

Figure 7.4 – Hyper-V Settings in Windows Server 2025

Note:

In addition to the core functionalities of Hyper-V Manager, it is essential to acknowledge the significance of the **Replication Configuration** feature. This functionality allows administrators to set up Hyper-V Replica, which provides disaster recovery capabilities by enabling the replication of VMs to another host. Additionally, the **Live Migrations** option facilitates the seamless movement of running VMs from one host to another without downtime, enhancing flexibility and resource management in virtualized environments. Understanding and utilizing these features can significantly improve your disaster recovery strategies and optimize resource utilization in Hyper-V.

By thoroughly configuring these settings, you set a solid foundation for effectively managing and optimizing your virtual machines. Once these initial configurations are complete, you will be well-prepared to dive into the creation and management of virtual hard disks, which will be covered in the next section.

How to make and adjust VHDs

To create and manage a **virtual hard disk (VHD)** on Windows Server 2025 using Hyper-V Manager, follow these detailed steps:

1. Click the **Start** button, then navigate to **Windows Tools** and select **Hyper-V Manager** from the menu to launch the application.
2. In Hyper-V Manager, open the **Actions** pane on the right side of the interface. Click **New** and select **Hard Disk...** to initiate the VHD creation process, as depicted in *Figure 7.5*.

Figure 7.5 – Creating a virtual hard disk

3. The **New Virtual Hard Disk Wizard** will appear. On the first page, click **Next** to begin the setup.
4. Choose the format for the VHD. You can select **VHDX** for enhanced performance and larger capacity or **VHD** for compatibility with older systems. Click **Next** to proceed.
5. Determine the type of VHD that best fits your needs. Options include **fixed size** (where the disk size is set and remains constant), **dynamically expanding** (where the disk size grows as data is added up to a maximum size), or **differencing** (which tracks changes from a base VHD, useful for snapshots and testing). Click **Next** to continue.
6. Specify the name and location of the VHD file. Ensure that the location has sufficient space for the VHD's intended use. Click **Next** to proceed.
7. Decide whether to create an empty VHD or to import data from an existing physical disk. If you choose to import, you'll need to follow additional steps to select the physical disk and configure the import settings. Click **Next** to move forward.
8. Review your selections, and click **Finish** to complete the wizard and create the VHD. The process will generate a new VHD that you can use on your VMs.

Once your VHD is created, managing the VM's memory allocation is crucial, as it directly impacts performance. Properly configuring virtual RAM ensures efficient operation and stability of your virtual environments.

Adjusting the RAM of a VM

To adjust the RAM allocation for a VM in Windows Server 2025 using Hyper-V Manager, follow these comprehensive steps to ensure optimal performance and resource management. Begin by ensuring that the VM is

powered off to prevent any conflicts or issues during the configuration process. Here is a detailed guide:

1. Click the **Start** button, then navigate to **Windows Tools** and select **Hyper-V Manager** to launch the application.
2. In Hyper-V Manager, locate the **VM** you want to configure from the list of available VMs. Right-click on the selected VM (make sure it is turned off) and choose **Settings...** from the context menu, as shown in *Figure 7.6*.

Figure 7.6 – VM settings in Hyper-V

3. The **Settings** window will open. In the left-hand pane, under the **Hardware** section, select **Memory** to access the memory configuration options.
4. You now have two options for configuring the VM's memory (refer to *Figure 7.7*):

Figure 7.7 – Managing virtual memory in Hyper-V

5. To allocate a fixed amount of memory, input the desired memory size in **megabytes (MB)** in the **RAM** box. This setting ensures that the VM always has access to this amount of memory, regardless of its current needs.
6. To use dynamic memory, check the **Enable Dynamic Memory** box. This feature allows Hyper-V to allocate memory based on the VM's demand. You can set a **Minimum RAM** value, which is the least amount of memory the VM will always have, and a **Maximum RAM** value, which is the maximum amount of memory the VM can use. Dynamic memory can improve resource utilization by adjusting memory allocation as needed.
7. After configuring the desired memory settings, click **OK** to apply the changes and close the VM settings window.

Adjusting RAM allocation is crucial for optimizing VM performance and ensuring efficient resource management. Once the memory settings are configured, you can proceed to set up a virtual network to enable communication between your VMs, allowing them to interact and share resources effectively. Another interesting perspective is that **overprovisioning** refers to the practice of allocating more virtual resources—such as CPU, memory, and storage—than physical hardware can support, which can enhance flexibility and resource availability in a virtualized environment. While this approach can facilitate the deployment of multiple VMs and improve utilization rates, it also carries risks, such as performance degradation and increased contention for resources, if the physical server reaches its limits. Therefore, it is crucial to balance resource allocation carefully, considering workload requirements and physical capacity, to avoid overprovisioning pitfalls while maximizing the benefits of virtualization. Understanding overprovisioning is essential for anyone venturing into virtualization, as it directly impacts performance, scalability, and the overall efficiency of IT operations.

Virtual networks in Hyper-V Manager

Setting up a virtual network is crucial for ensuring seamless communication between VMs and between VMs and the external physical network. In Hyper-V, this is accomplished by configuring a virtual switch, which acts as a bridge for network traffic. There are three primary types of virtual switches, each serving different networking needs:

- **External Switch:** This type connects VMs directly to the host server's physical network adapter, allowing VMs to communicate with the physical network and other devices on it. This configuration is useful for scenarios where VMs need to interact with external networks or the internet.
- **Internal Switch:** An internal switch connects VMs to the host server but does not provide access to the external network. This setup is ideal for scenarios where VMs need to communicate with the host and with each other but do not require external network access.

- **Private Switch:** A private switch allows communication solely between VMs on the same host. It does not connect to the host server or any external network. This is useful for isolated environments where VMs need to interact without external interference.

Setting up a virtual network

To set up a virtual switch in Windows Server 2025 using Hyper-V Manager, follow these steps to ensure proper network configuration:

1. Click the **Start** button, navigate to **Windows Tools**, and select **Hyper-V Manager** to launch the application.
2. In Hyper-V Manager, access the **Actions** pane and click on **Virtual Switch Manager...** as shown in *Figure 7.8*. That opens the Virtual Switch Manager interface.

Figure 7.8 – Creating a virtual switch

3. In the Virtual Switch Manager window, select the type of virtual switch you wish to create (**External**, **Internal**, or **Private**), and then click **Create Virtual Switch** as indicated in *Figure 7.9*.

Figure 7.9 – Virtual switch properties

4. Provide a descriptive name and optional description for the virtual switch to identify its purpose and functionality easily.
5. Choose the appropriate connection type that aligns with the switch type selected. For an external switch, select the **physical network adapter**; for internal or private switches, configuration options will vary accordingly.
6. If needed, enable **Virtual LAN (VLAN) Identification** to manage network traffic and enhance security by segregating network traffic within the virtual network.
7. Click **OK** to apply the settings and close the Virtual Switch Manager window.

With the virtual switch set up, you can now manage network traffic and connectivity for your VMs based on the switch type chosen. This configuration is foundational for effective VM management. In the following section, we will discuss checkpoints, a valuable feature that enables you to capture and restore the state of VMs, facilitating recovery from issues that may arise during updates, installations, or configuration changes.

Understanding checkpoints

Checkpoints in Hyper-V are a pivotal feature that allows you to capture and preserve the state of a VM at a specific point in time. This feature is invaluable for maintaining system stability and mitigating risk during critical operations such as updates, installations, or configuration changes. By creating a checkpoint, you can revert a VM to its previous state if something goes awry, ensuring that you can recover from issues without significant downtime or data loss.

To create a checkpoint for a VM, follow these detailed steps:

1. Click the **Start** button, navigate to **Windows Tools**, and select **Hyper-V Manager** to open it.
2. In Hyper-V Manager, locate the **VM** for which you want to create a checkpoint. Right-click on this VM and choose **Checkpoint** from the context menu, as shown in *Figure 7.10*.

Figure 7.10 – Creating a checkpoint

3. The checkpoint will now appear in the **Checkpoints** section of the **VM details** pane. This section lists all existing checkpoints, allowing you to manage them effectively.
4. If the VM is running when you create the checkpoint, you will receive a notification confirming the successful creation of the checkpoint, as depicted in *Figure 7.11*. This confirmation ensures that the checkpoint has been established correctly and is ready for use.

Figure 7.11 – Checkpoint creation confirmation

Types of checkpoints

Hyper-V offers two types of checkpoints, as illustrated in *Figure 7.12*, each catering to different needs:

- **Production Checkpoint:** This type focuses on capturing the VM's state from the operating system's perspective without including the state of running applications. It is optimized for scenarios where application consistency is not as critical, making it suitable for production environments where minimal disruption is desired.
- **Standard Checkpoint:** This type captures the complete state of the VM, including all running applications and their states. It provides a full snapshot of the VM, which is essential for scenarios where complete restoration is required, including the applications and their configurations.

Figure 7.12 – Checkpoint types

Understanding these checkpoints is crucial for maintaining both data and application consistency. Data consistency ensures that all users see a unified view of the data, reflecting all changes and transactions made. Application consistency, on the other hand, ensures that the entire application state is preserved, allowing for coordinated backups and restorations across all components. In Hyper-V, checkpoints play a key role in backing up virtual machines, containers, and cloud services, providing a reliable method for recovery and minimizing potential disruptions.

In the following section, we will explore VHD and VHDX files, which are fundamental to understanding virtual disk management within Hyper-V. These files represent the virtual storage where the operating system, applications, and data are stored and managed, forming the backbone of VM functionality.

VHD and VHDX formats

Since its debut in Windows Server 2008, Hyper-V has significantly evolved in terms of virtual disk storage capabilities. Initially, Hyper-V used the **Virtual Hard Disk (VHD)** format, which had a maximum disk size limitation of 2 TB. This format served as the standard for virtual machine storage, providing basic functionality for many virtualized environments. However, with the introduction of Windows Server 2012, Microsoft enhanced Hyper-V's storage capabilities by introducing the **Virtual Hard Disk Extended (VHDX)** format. VHDX not only increased the maximum disk size limit to 64 TB, accommodating larger and more data-intensive applications, but also introduced additional features such as improved resilience to power failures, better performance, and support for larger block sizes. VHDX also includes enhancements such as protection against data corruption and more efficient disk space utilization.

Despite these advancements, the VHD format remains supported in Windows Server 2025 for compatibility with existing systems and legacy environments. Both VHD and VHDX formats allow administrators to choose the most appropriate option based on their storage needs and compatibility requirements. Understanding the differences and capabilities of these formats is crucial for effective virtual machine management and ensuring optimal performance and data integrity in virtualized environments.

In the next section, we will explore the process of converting physical machines to virtual machines, an essential technique for migrating workloads from on-premises infrastructure to the cloud, thereby facilitating more flexible and scalable IT operations.

Transitioning from physical to virtual servers

Virtualization is increasingly favored by organizations seeking to optimize resources, reduce costs, and improve scalability. Converting physical servers to virtual machines is a key part of this process, achieved by creating VHDs. **Microsoft's Disk2vhd** tool simplifies this transition by converting physical disk drives into VHD files that can be imported into virtual environments. The Disk2vhd tool, as depicted in *Figure 7.13*, enables the seamless transformation of physical storage into a virtual format. After generating the VHD, you can use Hyper-V Manager to set up a new virtual machine and attach the VHD to it, completing the **physical-to-virtual (P2V)** migration. This migration process not only modernizes IT infrastructure but also enhances flexibility and manageability by consolidating multiple physical servers into fewer virtual ones. By adopting this approach, organizations can better allocate resources, improve disaster recovery capabilities, and scale their operations more efficiently.

Figure 7.13 – The Disk2vhd app facilitates the conversion of a physical disk drive into a VHD

Note:

The Disk2Vhd tool, essential for converting physical servers to virtual machines, can be downloaded from the official Sysinternals website at <https://docs.microsoft.com/en-us/sysinternals/downloads/disk2vhd>. This tool facilitates the seamless migration of physical disk drives into VHDs, enabling efficient P2V conversions. As we proceed, we will explore the process of converting virtual machines back to physical servers, known as **virtual-to-physical (V2P)** conversions.

Reverting from virtual to physical servers

Although virtualization is increasingly adopted for its benefits, such as cost savings, enhanced resource utilization, and scalability, there are specific scenarios where organizations might need to transition their virtual servers back to physical hardware. This reverse process, known as V2P conversion, is less straightforward and generally not as well-supported as P2V migrations. Many hypervisor vendors, including Microsoft, do not provide dedicated tools for V2P conversions, which makes this task more challenging. As a result, organizations often need to explore third-party solutions that can facilitate this process, though these tools may vary in functionality and reliability.

Another approach to managing this transition involves a manual migration process. That entails setting up a new physical server with Windows Server 2025 and then manually transferring the virtual server's settings, applications, and data to the physical server. This method can be more labor-intensive but offers the advantage of tailored configuration and optimization for the new physical environment. During this process, it is crucial to carefully plan and execute the migration to ensure that all necessary data and settings are accurately transferred and to minimize downtime or disruption to services.

Note:

To carry out a V2P conversion, you can utilize the **EZ Gig IV** cloning software, which is available for download on Apricorn's website: <https://www.apricorn.com/upgrades/ezgig>. This tool simplifies the process by following three straightforward steps: first, choose the source drive from which the data will be cloned; second, select the destination drive where the data will be transferred; and finally, click the **Start Clone** button to begin the cloning process. This approach helps facilitate the transfer of data from a virtual environment to physical hardware.

Migrating from VMware to Hyper-V

As organizations increasingly recognize the benefits of Hyper-V, transitioning from VMware to Hyper-V has become a strategic move for many. This process involves several critical steps to ensure a smooth migration while

minimizing downtime and maintaining data integrity.

Key considerations before migration

Before initiating the migration from VMware to Hyper-V, it is essential to consider several key factors that will influence the success of the transition and ensure a seamless integration into the new environment.

- **Assessment of Current Environment:** Begin by evaluating the existing VMware infrastructure. Identify the number of VMs, their resource allocations, and the applications they host. This assessment will inform the planning and execution of the migration process.
- **Compatibility and Licensing:** Verify that the applications running on VMware are compatible with Hyper-V. Additionally, compliance with licensing agreements for both VMware and Hyper-V should be ensured to avoid legal issues.
- **Data Backup:** Prior to initiating the migration, back up all data associated with the VMs. This backup serves as a safeguard against any potential data loss during the migration process.

Migration tools

Microsoft provides several tools to facilitate the migration from VMware to Hyper-V, including these:

- **Microsoft Virtual Machine Converter (MVMC):** This tool allows for the conversion of VMware VMs to Hyper-V format. It supports both physical and virtual machine migrations, making it versatile for different environments.
- **Disk2VHD:** A free utility from Sysinternals, Disk2VHD enables the creation of VHD files from physical disks, allowing for easy transfer of workloads into Hyper-V.

Migration process overview

The migration process from VMware to Hyper-V involves a systematic approach that encompasses careful planning, execution, and validation to ensure that virtual machines are transferred effectively while minimizing downtime and maintaining data integrity:

1. **Prepare the Hyper-V Environment:** Set up Hyper-V on the target host, ensuring that the server is configured correctly and updated. This includes enabling virtualization features in the BIOS and installing the Hyper-V role through Server Manager or PowerShell.
2. **Convert VMs:** Utilize the chosen migration tool to convert VMware VMs to Hyper-V. The process may involve exporting the VM from VMware and importing it into Hyper-V. Ensure that the virtual hardware settings are appropriately configured to match the capabilities of Hyper-V.
3. **Testing:** After the migration, thoroughly test the VMs in the Hyper-V environment. Check for application performance, connectivity, and any dependencies that may have been affected during the migration.
4. **Finalization:** Once testing is successful, update DNS records and make necessary changes to network configurations to reflect the new Hyper-V environment. Finally, the VMware infrastructure can be decommissioned if it is no longer needed.

Migrating from VMware to Hyper-V can offer organizations significant advantages, including cost savings and enhanced integration with other Microsoft services. By following a structured approach and leveraging the right tools, IT professionals can ensure a successful transition that aligns with their organization's operational goals. For further details on migration strategies and tools, refer to the resources provided in the *Further reading* section of this chapter.

Next, we will delve into configuring settings specific to individual VMs. Understanding these VM-specific settings is crucial as it highlights the distinctions between general Hyper-V configuration and settings applied directly to each VM.

Adjusting VM settings

To effectively manage a VM in Hyper-V, you can access its settings by right-clicking on the VM's name in Hyper-V Manager and selecting **Settings** from the context menu. This action opens a comprehensive configuration window where you can fine-tune various aspects of the VM, as detailed in *Figure 7.14*. Here is an overview of the settings you can adjust:

- **Add Hardware:** This option allows you to attach additional devices to the VM, such as network adapters or memory controllers, to expand its capabilities.
- **Firmware:** Configure the virtual firmware settings for the VM, ensuring compatibility with the underlying hardware and enabling features such as secure boot or boot from virtual devices.
- **BIOS:** Configure the boot order to determine the sequence in which the VM's virtual devices are used during startup. This setting is crucial for ensuring the VM boots from the correct device.
- **Security:** Enable encryption to secure the VM's state files and data during migration. That ensures that sensitive information remains protected as the VM moves across different environments.
- **Memory:** Adjust the amount of memory allocated to the VM, either by setting a fixed amount or enabling **Dynamic Memory** to allow the VM to use memory more flexibly based on its workload requirements.
- **Processor:** Specify the number of virtual processors assigned to the VM, which can influence its performance and ability to handle multiple tasks simultaneously.
- **IDE Controller 0 and 1:** Manage the hard drives and CD/DVD drives connected to the IDE controllers. That includes adding or removing storage devices and configuring their settings for optimal performance.
- **SCSI Controller:** Configure hard drives connected to the SCSI controller, which can be useful for high-performance storage solutions and attaching multiple drives.
- **Network Adapter:** Set up and customize the network adapter settings to control how the VM connects to the network, including configuring VLANs and network bandwidth.
- **COM 1 and COM 2:** Set up virtual COM ports for serial communication, which can be useful for legacy applications or specialized hardware interactions.
- **Diskette Drive:** Select and configure a virtual floppy disk file, which may be necessary for specific legacy applications or to provide additional storage options.

Figure 7.14 – Establishing VM settings

By carefully adjusting these settings, you can optimize the VM's performance, ensure proper connectivity, and tailor it to meet specific operational requirements.

In the next section, we will delve into techniques for managing multiple VMs simultaneously in Hyper-V, including best practices for resource allocation and automation.

Working with VMs

When managing VMs in Hyper-V, leveraging both the **Actions** pane and the VM's context menu can significantly streamline administrative tasks. The **Actions** pane, depicted in *Figure 7.15*, is an essential component of Hyper-V Manager that facilitates comprehensive VM management. It provides options to create new virtual machines, configure Hyper-V settings, set up virtual switches, and establish virtual **Storage Area Networks (SANs)**. This pane also allows for modifications and examinations of virtual disks, stopping and starting services, deleting VMs, and refreshing the list of available VMs. Its role as a central management tool ensures that administrators can perform a wide range of tasks from a single interface, enhancing efficiency and ease of use.

Figure 7.15 – Actions pane in Hyper-V Manager

The VM's context menu, illustrated in *Figure 7.16*, complements the **Actions** pane by offering options specific to the selected VM. This menu includes essential functions such as **Connect...** for accessing the VM's console, **Rename...** for updating the VM's name, and various other management options tailored to the individual VM. For instance, you can manage the VM's settings, checkpoints, and snapshots or even control its power state (e.g., start,

stop, pause) directly from this menu. The context menu's specificity allows for precise, focused management of individual VMs, making it a valuable tool for handling VM-specific tasks.

Figure 7.16 – Context menu in Hyper-V Manager

Understanding how to use both the **Actions** pane and the context menu equips you with a robust set of tools for efficient VM management in Hyper-V.

Best practices for VM startup and recovery settings

Configuring VM startup and recovery settings is essential for ensuring smooth operations and system resilience, particularly after a host reboot. These settings allow administrators to control the behavior of VMs in response to host restarts, minimizing potential downtime and optimizing resource allocation:

- **Startup Action:** Configuring VMs based on operational needs is recommended. Options include the following:
 - **Do Nothing:** Ideal for non-critical VMs, which can help preserve host resources upon restart.
 - **Automatically Start if Running:** For essential VMs, this setting ensures they resume their prior state without requiring manual intervention, which is helpful for consistent service continuity.
 - **Always Start:** Suitable for critical systems that must be operational immediately after a host restart, regardless of their prior state.
- **Automatic Start Delay:** Staggering VM startups can prevent performance bottlenecks by reducing the immediate load on CPU and memory. This feature is beneficial in environments with multiple VMs on a single host.
- **Automatic Stop Action:** Configuring VMs to gracefully shut down when the host is shut down or rebooted helps prevent data loss and maintains VM integrity.

Hyper-V provides options in the **Automatic Start Action** and **Automatic Stop Action** settings for each virtual machine, allowing administrators to configure how VMs behave when the Hyper-V host starts or shuts down. These settings are essential in Hyper-V environments for the following:

- Ensuring uptime for critical workloads by automatically restarting essential VMs
- Managing resource allocation during host reboots by setting delays for non-critical VMs, helping prevent performance bottlenecks
- Preserving VM data integrity by configuring the shutdown action to power down VMs before the host shuts down gracefully

These startup and recovery configurations are integral to Hyper-V's management options and play a key role in maintaining reliable, resilient VM operations. Adhering to these practices ensures a well-orchestrated VM environment, with minimized impact from unexpected reboots and optimized performance across workloads.

Real-world applications of Hyper-V for modern IT environments

To enhance understanding and provide practical insights, this section highlights real-world scenarios where Hyper-V plays a crucial role in modern IT operations. These examples not only demonstrate Hyper-V's versatility but also offer actionable steps for common industry practices, such as migration, disaster recovery, automation, and backup.

Migrating from VMware to Hyper-V

For many organizations, migrating from VMware to Hyper-V represents a strategic shift toward consolidating IT resources within Microsoft's ecosystem. This migration requires thoughtful planning, starting with a comprehensive compatibility assessment of existing VMs. Utilizing tools such as **Microsoft Virtual Machine**

Converter (MVMC) or **System Center Virtual Machine Manager (SCVMM)** can streamline the migration process by automating certain stages, such as VM disk conversion and network configuration. Testing each VM in a staging environment prior to live deployment ensures optimal functionality and mitigates potential issues. With careful preparation, organizations can transition to Hyper-V while maintaining high performance and minimizing service disruptions.

Disaster recovery with Hyper-V Replica

Hyper-V Replica is a powerful feature for disaster recovery, enabling organizations to replicate VMs to a secondary site, either on-premises or in the cloud. This setup provides a critical safety net, ensuring rapid recovery and minimal data loss in the event of a primary site failure. Configuring Hyper-V Replica involves setting up replication at the VM level, defining the frequency of replication based on **recovery point objectives (RPOs)**, and configuring the network connections between primary and replica sites. By regularly replicating VMs to a standby site, businesses can safeguard their data and reduce downtime, supporting a resilient infrastructure.

Automating VM checkpoints with PowerShell

Routine maintenance, such as system updates, can introduce changes that impact VM stability. Automating VM checkpoints with PowerShell before each update is a best practice to facilitate rollback if issues arise. The following PowerShell script captures a checkpoint for each running VM, labeling each snapshot with a timestamp for easy identification:

```
Get-VM | Where-Object { $_.State -eq 'Running' } | ForEach-Object {Checkpoint-VM -VMName $_.Name
```

This script helps administrators save time while implementing a safety mechanism across multiple VMs, promoting operational consistency and minimizing the risk associated with updates.

Hyper-V VM backups using Windows Server 2025

Regular backups of Hyper-V VMs are essential for business continuity, regulatory compliance, and data protection. Windows Server 2025 provides tools such as Windows Server Backup and System Center **Data Protection Manager (DPM)** to schedule automated backups or create on-demand snapshots. Configuring routine backups ensures that VM state, configuration, and data are securely saved, supporting rapid recovery in case of accidental data loss or cyber incidents. By implementing regular backups, organizations not only protect critical data but also build a resilient and compliant IT infrastructure capable of meeting modern business demands.

These real-world examples emphasize Hyper-V's capabilities in achieving reliable and efficient virtualized environments, helping IT professionals apply these best practices directly in their organizations.

With this foundation, you are prepared to move on to practical exercises, such as installing the Hyper-V role in Windows Server 2025, to enhance your skills further and apply your knowledge in real-world scenarios.

Chapter exercise – installing the Hyper-V role on Windows Server 2025

To install the Hyper-V role on Windows Server 2025, begin by preparing your server environment for the role installation. Hyper-V is a powerful virtualization platform that allows you to create and manage virtual machines on Windows Server. This process ensures that your server is configured to support virtualization, enabling you to optimize resources and run multiple virtualized workloads efficiently.

The installation process involves several straightforward steps, which are outlined in detail next. Once you have completed these steps, you will have successfully installed Hyper-V and can start leveraging its capabilities to deploy and manage virtual machines.

To install the Hyper-V role on Windows Server 2025, follow these detailed steps using Server Manager:

1. Open Server Manager by clicking the **Start** button and selecting **Server Manager** from the **Start** menu.
2. In the **Server Manager** window, click on the **Add Roles and Features** hyperlink to begin the installation process.
3. On the **Before You Begin** screen, click **Next** to proceed.
4. On the **Installation Type** screen, select **Role-based or Feature-based Installation** and click **Next**.
5. On the **Server Selection** screen, choose the appropriate server from the server pool and click **Next**.
6. On the **Server Roles** screen, check the box for the **Hyper-V** role, as illustrated in *Figure 7.17*.

Figure 7.17 – Selecting the Hyper-V role

7. Click the **Add Features** button to include the necessary features for Hyper-V.
8. If no additional features are required, click **Next**.
9. On the Hyper-V screen, review the information and click **Next**.
10. Select the network adapter(s) to be used for Hyper-V and click **Next**.
11. Check the box labeled **Allow this server to send and receive live migrations of virtual machines**. If you want to enable live migration, then click **Next**.
12. Specify the location where virtual machines will be stored and click **Next**.
13. Review the installation selections for the Hyper-V role and click **Install** to begin the installation process.
14. Once the installation is complete, click **Close** (refer to *Figure 7.18*). The server will automatically reboot to apply the changes.

Figure 7.18 – Installing the Hyper-V role on Windows Server 2025

15. After the server reboots, the Hyper-V role will be installed and ready for use.

Ensure that you verify the installation and configure Hyper-V settings to suit your specific requirements. This setup marks the beginning of utilizing virtualization to enhance server management and operational efficiency in your environment.

Note:

For those looking to gain experience with PowerShell and streamline the installation process, Hyper-V can also be installed using the following PowerShell command:

Install-WindowsFeature -Name Hyper-V -ComputerName <computer_name> -IncludeManagementTools -Restart

This command installs the Hyper-V role and the management tools and automatically restarts the computer to complete the installation.

Summary

In this chapter, you were introduced to Hyper-V, a powerful virtualization platform designed to enhance server efficiency and flexibility by enabling the creation and management of multiple virtual machines on a single physical server. You gained insights into the core concepts, essential components, and various features of Hyper-V, which allow for the simultaneous operation of different operating systems and applications, optimizing resource utilization. The chapter included a hands-on exercise where you learned to install the Hyper-V role on Windows Server 2025, providing you with practical experience in configuring and deploying virtual environments. As you progress, the next chapter will shift focus to data storage solutions in Windows Server 2025. This upcoming

section will offer a detailed exploration of the various data storage options available, including their advantages and best practices for managing and safeguarding your data within the server infrastructure.

Questions

1. **True or False:** Hyper-V provides services you can use to create and manage VMs and their resources.
2. **Fill in the Blank:** _____ is based on a hierarchical format where the first level represents the hypervisor as the main element of the Hyper-V virtual platform.
3. **Multiple Choice:** Which of the following are virtualization modes in Hyper-V? (*Choose two*)
 1. Fully virtualized mode
 2. Paravirtualized mode
 3. Production checkpoints
 4. Standard checkpoints
4. **True or False:** Checkpoints enable you to make a backup of the disk image at a specific time so that when unexpected situations arise, you can revert your VM to a previous state.
5. **Fill in the Blank:** Components such as _____ and _____, through a logical channel for communication known as VMBus, enable communication between the root portion and the branch OSes.
6. **Multiple Choice:** Which of the following are checkpoint types in Hyper-V? (*Choose two*)
 1. Production checkpoints
 2. Standard checkpoints
 3. Inspect disk
 4. Edit disk
7. **True or False:** Organizations migrate their physical servers to virtual servers (P2V) for cost, ease of management, and future expansion.
8. **Fill in the Blank:** _____ is an administration tool that you can use to manage VMs.
9. **Multiple Choice:** Which of the following are elements of the Hyper-V architecture? (*Choose two*)
 1. Hypervisor
 2. Root
 3. Branch
 4. Snapshot
10. **Short Answer:** Discuss nested virtualization.
11. **Short Answer:** Discuss P2V conversion.
12. **Short Answer:** Discuss V2P conversion.

Further reading

- Virtualization: <https://docs.microsoft.com/en-us/windows-server/virtualization/virtualization>
- How to Work with Hyper-V VHD and VHDX Files: Essential Basics: <https://www.nakivo.com/blog/work-hyper-v-vhd-vhdx-files-essential-basics/>
- Disk2vhd v2.01: <https://docs.microsoft.com/en-us/sysinternals/downloads/disk2vhd>
- Set up Hyper-V Replica: <https://learn.microsoft.com/en-us/windows-server/virtualization/hyper-v/manage/set-up-hyper-v-replica>
- Working with Hyper-V and Windows PowerShell: <https://learn.microsoft.com/en-us/virtualization/hyper-v-on-windows/quick-start/try-hyper-v-powershell>

8

Storing Data in Windows Server 2025

In this chapter, you will explore the fundamentals of **storage technologies** and their critical roles in server operations. **Disks** are crucial hardware components of servers, responsible for various tasks related to digital data, including storage, retrieval, management, and the provision of files and services. Understanding these technologies is vital for effective computer management and data handling.

The chapter delves into essential topics such as physical interfaces, disk controllers, and data storage methods within various storage media. You will gain insights into network-based storage systems and key storage concepts and protocols, including **Data Deduplication**, **Storage Spaces Direct (S2D)**, **Software-Defined Storage (SDS)**, **Small Computer System Interface (SCSI)**, **Internet Small Computer System Interface (iSCSI)**, **Fiber Channel (FC)**, and **Fiber Channel over Ethernet (FCoE)**. Furthermore, you will learn to manage server storage using both **Server Manager** and **Windows PowerShell**.

Additionally, the chapter introduces the concept of **Redundant Array of Independent Disks (RAID)** and explores different storage technologies based on their volatility. These include RAM, ROM, **hard disk drives (HDDs)**, **solid-state drives (SSDs)**, optical drives, and flash memory drives and cards. To consolidate your learning, you will complete an exercise on enabling Data Deduplication in Windows Server 2025, reinforcing your understanding and practical skills in server storage management.

In this chapter, we're going to cover the following main topics:

- Understanding storage technologies and their evolution in Windows Server 2025
- Exploring storage architectures and their implications for network environments
- Overview of storage protocols and their roles in data transmission and access
- Managing server storage using Server Manager and Windows PowerShell
- Understanding RAID principles and configurations
- Understanding primary storage concepts and optimizing storage solutions in Windows Server 2025
- Enabling Dedup on Windows Server 2025

Technical requirements

To effectively practice and refine the skills discussed in this chapter, it is essential to have the appropriate technical setup:

- Ensure you have access to a PC equipped with **Windows 11 Pro**, featuring a minimum of 16 GB of RAM, 1 TB of hard drive space, and a stable internet connection
- Additionally, you will need a **virtual machine (VM)** running **Windows Server 2025 Standard** (Desktop Experience) edition, configured with at least 4 GB of RAM, 100 GB of hard drive space, and internet connectivity

These resources will enable you to configure and manage the various storage technologies within Windows Server 2025, providing a practical environment to apply your learning.

Understanding storage technologies and their evolution in Windows Server 2025

Storage technologies are foundational to the operation of any computer system, especially servers, which are tasked with storing, managing, and processing extensive volumes of data. These technologies are diverse, each offering unique characteristics, functions, and designs based on their intended application. For ICT professionals and those preparing for certification exams, understanding these differences is crucial. In this section, we will introduce you to the key storage technologies that are vital in the field, such as **Integrated Drive Electronics (IDE)**, **Serial Attached SCSI (SAS)**, **SCSI**, **Direct Attached Storage (DAS)**, **Network Attached Storage (NAS)**, **Storage Area Network (SAN)**, and **RAID**. Each of these technologies serves distinct purposes, ranging from direct data storage solutions to complex networked storage systems, and plays a pivotal role in ensuring data availability, reliability, and performance in enterprise environments.

We will also discuss the significance of each technology in various scenarios, including how they impact system performance, scalability, and data protection strategies. By understanding these technologies, you'll be better equipped to make informed decisions when configuring, managing, or troubleshooting storage solutions. Following this overview, we will explore each storage type in greater detail, examining their specific features, advantages, and use cases to provide a comprehensive understanding of how they contribute to the overall efficiency and reliability of IT infrastructures.

Exploring different storage types

Different storage technologies bring unique features and functionalities to the table, each serving specific roles in the IT landscape. A thorough understanding of these technologies is crucial for IT professionals who need to select the appropriate solutions for different scenarios. Here is how the main categories of storage technologies can be classified:

- **Optical disks:** Often used for archival and backup purposes, optical disks offer substantial storage capacity and reliable read-and-write speeds. Although they were once more prevalent, especially for media distribution and data backup, their role as primary storage solutions has diminished with the rise of more advanced technology.

- **HDDs:** For many years, HDDs were the standard choice for operating systems and general data storage due to their large capacity and robust read-and-write performance. However, as technology has evolved, HDDs have become less favored as primary storage options. However, they remain widely used for storing large volumes of data in client/server applications where cost-effectiveness and capacity are prioritized over speed.
- **SSDs:** SSDs have rapidly become the preferred storage medium in modern computing environments. With their superior read-and-write speeds, enhanced durability, and continuously increasing capacity, SSDs are often chosen for operating systems and applications that require fast access to data. Their lack of moving parts also makes them more reliable and less prone to physical damage than HDDs.
- **Non-Volatile Memory Express (NVMe):** NVMe is a modern interface designed specifically for SSDs. It allows for faster data transfer speeds by leveraging the PCIe bus. This technology significantly reduces latency and enhances throughput compared to older storage protocols, making it ideal for high-performance computing environments. NVMe has become increasingly popular for applications that demand rapid data access, such as databases, virtualization, and intensive computing tasks.

Understanding these core storage types is essential for making informed decisions about data management and storage strategies. With this foundation, we can now delve into the various interfaces used to connect and manage these storage technologies within computer systems, including **Advanced Technology Attachment (ATA)**, **Parallel ATA (PATA)**, **Serial ATA (SATA)**, and **SCSI**. Each interface has its advantages, compatibility considerations, and use cases, making it important to understand how they influence storage performance and system architecture.

Understanding ATA and SCSI interfaces

Various interfaces are essential for connecting storage devices such as HDDs, optical drives, and even legacy floppy drives to computer systems. These interfaces are often recognized by their acronyms—**ATA**, **PATA**,

SATA, and **SCSI**—which each signify specific characteristics and functions that define their roles in data management and system architecture:

- **ATA:** Also known as **IDE (Integrated Drive Electronics)**, ATA is a foundational interface developed to connect storage devices to computers. Despite its age, ATA laid the groundwork for subsequent storage interfaces. The ATA interface has evolved into two main subtypes, as shown in Figure 8.1:
 - **PATA:** PATA, which uses a **40-pin ribbon cable** for data transfer and a **Molex connector** for power, was widely used in earlier PCs. This interface connects the storage device directly to the motherboard and power supply. A key feature of PATA is that the disk controller is built directly into the drive, simplifying the connection process. However, the parallel nature of data transmission, with multiple bits sent simultaneously, can result in slower data transfer rates compared to more modern interfaces.
 - **SATA:** SATA succeeded PATA and quickly became the standard in modern computing due to its superior speed and efficiency. It uses a **7-pin cable** for data transfer and a **15-pin power connector**, which are both smaller and more flexible than their PATA counterparts. Like PATA, the disk controller is integrated into the drive. Still, SATA's serial data transmission method—sending bits one after another—allows for faster data transfer speeds and better airflow inside the case, which is crucial for maintaining optimal system performance.


 Figure 8.1 – PATA (left) and SATA (right) data cables

Figure 8.1 – PATA (left) and SATA (right) data cables

- **SCSI:** Pronounced *scuzzy*, SCSI is a more advanced interface designed to connect a variety of peripherals, including storage devices, to computer systems. SCSI has evolved with two prominent variants:
 - **SCSI Parallel Interface (SPI):** SPI is the older parallel version of SCSI. It was once widely used in enterprise environments for its ability to connect multiple devices on a single bus. However, due to limitations in speed and scalability, SPI has been largely phased out in favor of newer technologies.

- **SAS:** SAS represents the modern evolution of SCSI, offering significantly higher data transfer speeds and enhanced reliability. SAS uses point-to-point serial connections, similar to SATA, but it is designed to meet the demanding needs of enterprise environments. It supports a more significant number of devices and offers features like dual-port access for redundancy, making it a preferred choice for servers and high-performance workstations.

Understanding these interfaces is critical for IT professionals tasked with managing storage solutions, as each interface offers unique benefits and is suited to specific use cases. As we move forward, we will delve into the role of **Peripheral Component Interconnect (PCI)** and **PCI Express (PCIe)** expansion cards, which further enhance system capabilities by providing additional connectivity options and improving overall performance in various computing environments.

PCI and PCIe overview

PCI is an interface standard that emerged in the 1990s, becoming a foundational technology for connecting various hardware components in personal computers. Developed by *Intel*, PCI represented a significant upgrade from the older **Industry Standard Architecture (ISA)** interface, which *IBM* initially created. The ISA interface, with its 16-bit internal bus specification, had been the standard for a time. Still, it could not keep pace with the increasing demands for faster data processing and more efficient system performance. PCI addressed these limitations by offering both 32-bit and 64-bit internal bus specifications, allowing more robust and quicker data transfer between the computer's components and its memory (RAM). That made PCI a vital element in the development of more powerful and versatile computing systems throughout the 1990s and early 2000s.

As technology continued to evolve, however, the need for even greater speed and efficiency led to the development of a new interface standard: PCIe. PCIe was designed to replace PCI and surpass its capabilities, offering a significant performance boost. Unlike PCI, which uses a parallel bus architecture, PCIe utilizes a serial communication protocol. This serial bus standard supports various connection types, including **PCIe x1**, **PCIe x4**,

PCIe x8, and **PCIe x16**, with each kind differing in the number of lanes available for data transfer. These lanes are crucial because they enable full-duplex communication, meaning data can be sent and received simultaneously, thereby maximizing the speed and efficiency of data transmission within the system. PCIe has become the standard interface for high-speed components such as graphics cards, SSDs, and network cards, providing the necessary bandwidth for today's demanding applications.

PCIe slots, depicted in *Figure 8.2*, prominently featured on modern motherboards, accommodate these connections and play a critical role in a computer system's overall performance. PCIe's flexibility and scalability make it ideal for a wide range of applications, from high-performance gaming and professional workstations to enterprise servers and data centers.



Figure 8.2 – The PCIe slot

With this understanding of PCI and PCIe interfaces, we are now equipped to delve into the concept of local storage, explicitly focusing on the internal disks that store data within a computer. This exploration will further enhance our comprehension of how these storage solutions integrate with the interfaces we've discussed, ensuring optimal system performance and reliability.

Explaining local storage

Local storage pertains to the internal disk drives within a server or computer, encompassing both HDDs and SSDs. These drives are directly connected to the system via cables, a configuration known as DAS. DAS refers to any storage device that is physically attached to a computer or server through a direct connection, such as internal drives housed within the system unit or external drives connected via interfaces such as SATA, USB, or PCIe. For

instance, the HDD installed in your computer is a classic example of DAS, as is any external hard drive or SSD connected to the system.

DAS systems offer the advantage of high-speed access to data, as the storage devices are directly linked to the system's internal buses. This direct connection typically results in lower latency and faster data transfer rates compared to network-based storage solutions. However, DAS is limited by its lack of scalability and flexibility. Adding more storage or upgrading existing drives often requires physical changes to the system, which can be cumbersome in large-scale environments.

To illustrate, consider a single HDD or a collection of disks connected to a server—each disk in this setup is a part of the DAS system (see *Figure 8.3*), directly interfacing with the server's internal architecture. This setup allows straightforward data storage and retrieval, with all data management occurring within the confines of the individual server or computer.


 Figure 8.3 – DAS system

Figure 8.3 – DAS system

As we have established a clear understanding of local storage and its direct-attachment configuration, the next step is to explore network storage solutions. Network storage contrasts with DAS by offering remote access to storage resources over a network, and we will examine how it differs in terms of connectivity, scalability, and operational efficiency.

Exploring storage architectures and their implications for network environments

Network storage encompasses systems that connect to a network to provide centralized data services to multiple users and devices, distinguishing it from local storage, which is confined to the internal disk of a single computer or a directly attached storage device. Network storage solutions leverage network

interfaces and protocols to enable seamless data access and sharing across a networked environment. Two primary types of network storage are NAS and SAN, each offering unique features and benefits.

Network-Attached Storage

Network-Attached Storage (NAS) is a specialized storage device that connects to a network via a switch, functioning as a dedicated file server. NAS devices use standard Ethernet connections to provide file-sharing capabilities directly over the network. They offer a simple and scalable solution for file storage, allowing multiple users and devices to access, share, and manage files concurrently. NAS systems are particularly useful for environments that require centralized data management and collaborative file access without the need for additional server infrastructure. Their ease of deployment and cost-effectiveness make them an attractive choice for small to medium-sized businesses and home networks, providing practical solutions for their file storage needs (see *Figure 8.4*).

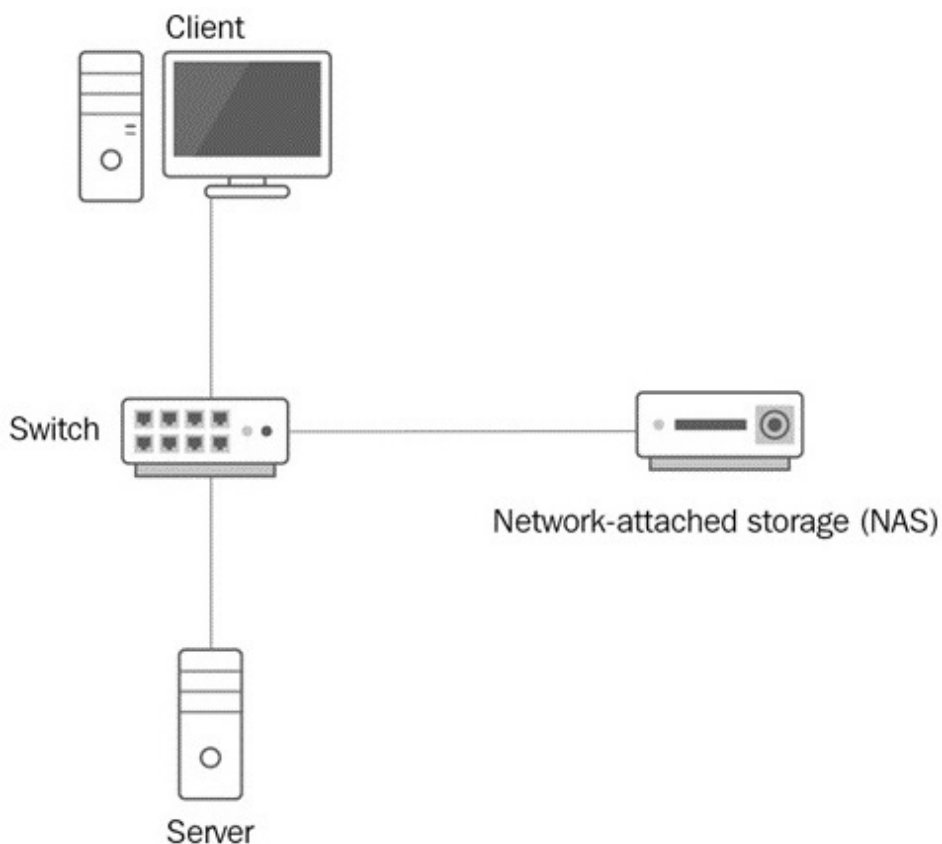


Figure 8.4 – NAS system

Storage Area Network

Storage Area Network (SAN), on the other hand, is a more advanced and high-performance storage solution designed for more extensive and more demanding environments. A SAN consists of a dedicated network that interconnects multiple storage devices, creating a shared pool of storage resources accessible to servers and clients. SANs use either Ethernet or FC connections to link storage devices with servers, with FC providing higher data transfer speeds. To facilitate this, servers in a SAN environment are equipped with **Host Bus Adapters (HBAs)**, which are specialized expansion cards that enable the servers to connect to the SAN. SANs often utilize proprietary protocols or **Simple Network Management Protocol (SNMP)** for managing storage resources, offering advanced features such as high availability, redundancy, and efficient data management. They are ideal for high-performance computing environments, data centers, and enterprise applications where speed and reliability are critical, introducing you to cutting-edge storage technology (see *Figure 8.5*).


 Figure 8.5 – SAN system

Figure 8.5 – SAN system

With a comprehensive understanding of both NAS and SAN systems, we can now turn our attention to the concepts of block-level and file-level storage. These two approaches represent different methods of organizing and managing data within various storage technologies, each suited to different use cases and performance requirements.

Understanding block-level and file-level storage

When managing data on storage devices, understanding the distinctions between file-level and block-level storage is crucial for optimizing performance and meeting specific storage needs. File-level storage and block-level storage are two fundamental approaches to organizing and accessing data, each with unique characteristics and advantages:

- **File-level storage:** This type organizes data into a hierarchical structure of files and folders, which can be easily accessed and modified by users and applications through file systems. This method presents data in a familiar format, making it straightforward for users to manage and interact with their data. File-level storage is commonly used in NAS systems, which are designed to provide centralized file access and sharing across a network. Its simplicity and ease of use make it well-suited for scenarios where file access and sharing are primary needs, such as in collaborative environments or for general-purpose storage.
- **Block-level storage:** This manages data by dividing it into fixed-size blocks, each with a unique identifier, and storing these blocks within a logical volume. This approach abstracts the data storage process from the file system, allowing for more detailed control over how data is written and retrieved. Block-level storage is ideal for applications requiring high performance and low latency, such as databases, VMs, and high-throughput applications. It is commonly employed in SANs, where it provides a high-speed, low-overhead storage solution that supports complex data operations. The block-level method's granularity and flexibility enable efficient data management, including the ability to perform advanced features such as snapshots, cloning, and replication.

The following table (*Table 8.1*) provides a comparative overview of file-level and block-level storage, summarizing their key differences and similarities. This comparison highlights the strengths and limitations of each method, offering insights into their appropriate use cases.

File-level storage

Data is stored and accessed in the form of files and folders

Used by NAS

Block-level storage

Data is stored in blocks representing volumes, which the operating system then manages

Used by SAN

Table 8.1 – File-level versus block-level storage

Before delving into the next topic, it's essential to understand the role of adapters and controllers in the data transfer process. These hardware components are critical for facilitating communication between storage devices and computer systems, ensuring efficient data exchange and system performance. Adapters and controllers manage the data flow, interface with various storage technologies, and support the connectivity required for effective data management.

Understanding how adapters and controllers operate

Disk controllers are crucial electronic components embedded within storage devices, responsible for managing essential functions such as rotating the platters, positioning the read/write heads, and facilitating data transfer between the disk and system memory. These controllers are located underneath the sealed enclosure of the disk, as depicted in *Figure 8.6*. The controller ensures that the disk operates efficiently by regulating its mechanical movements and handling data read and write requests. That includes converting digital signals into physical operations and vice versa, ensuring accurate data retrieval and storage.

 Figure 8.6 – Disk controller in an HDD

Figure 8.6 – Disk controller in an HDD

To fully comprehend how data is processed and communicated by storage devices, it is also important to explore the underlying technologies of serial bus systems and data transmission protocols. Serial buses, such as those used in modern interfaces such as SATA and USB, transmit data one bit at a time over a single channel, offering streamlined and efficient data transfer. Data transmission technologies dictate the speed and reliability of these transfers, influencing overall system performance. In the upcoming section, we will examine these technologies in detail, providing insights into their role in optimizing data communication and enhancing the functionality of storage systems. This deeper understanding will help in selecting and

configuring storage solutions that meet specific performance and reliability requirements.

Data transmission in storage devices

Storage devices utilize two main communication methods for data transfer: parallel and serial:

- **Parallel communication:** This type of communication transmits multiple bits simultaneously, typically handling 8 bits or 1 byte at a time. This method can theoretically offer high-speed data transfer, as multiple bits are sent at once. However, parallel communication can face challenges such as signal degradation and timing issues, known as data skew, especially over long distances.
- **Serial communication:** This type of communication sends data one bit at a time sequentially. Although it transmits data more slowly than parallel communication in raw bit terms, serial communication often proves more efficient and reliable for storage devices. That is because serial communication reduces the complexity associated with managing multiple data paths simultaneously and minimizes issues such as interference and signal degradation. By processing bits in a continuous stream, serial interfaces such as SATA, SAS, FC, and USB allow stable and high-performance data transfer. These interfaces ensure that data transmission is streamlined and less susceptible to the errors that can arise from parallel communication's complexity.

In the upcoming section, we will compare storage protocols and network protocols. This analysis will elucidate how these protocols interact with data transmission, highlighting their respective roles and examining their similarities and differences to provide a comprehensive understanding of data communication in various contexts.

Overview of storage protocols and their roles in data transmission and

access

Storage protocols play a pivotal role in IT infrastructure, serving as the foundation for data transfer, access, and management across networks. They are key to determining the efficiency, scalability, and security of storage systems, directly impacting the quality of data transactions. The choice of protocol varies based on the type and scale of storage required, with each protocol offering unique benefits to optimize performance and minimize latency. As enterprise systems continue to evolve, the importance of selecting the appropriate storage protocol to meet growing demands cannot be overstated. In this section, we will delve into the core storage protocols that enable data transmission and access, exploring their functions, applications, and the technical considerations that guide their deployment in diverse IT environments.

Communication protocols in storage devices

Storage devices rely on various protocols to manage data storage and retrieval efficiently. These protocols establish the rules and formats that enable seamless communication between storage devices and other system components. Key protocols used in storage technologies, such as NAS and SAN, are as follows:

- **SCSI:** SCSI is a well-established protocol used for block-level storage systems, allowing devices to exchange data efficiently. It facilitates communication between the operating system and storage devices by utilizing SCSI commands to manage data transfers. SCSI supports a range of devices and offers various command sets, making it versatile for different storage configurations.

Note

For a comprehensive understanding of SCSI, including its history, functionality, and applications, you can refer to detailed resources available at Lifewire's SCSI Guide <https://www.lifewire.com/small-computer-system-interface-scsi-2626002>. This source offers valuable insights into how SCSI

operates, its various command sets, and its relevance in modern storage solutions.

- **iSCSI:** iSCSI extends SCSI commands over IP networks by encapsulating them within IP packets. This protocol allows organizations to utilize standard network infrastructure to connect storage devices, providing a cost-effective solution for remote storage access. Its ability to integrate seamlessly with existing IP networks makes iSCSI suitable for both small-scale and large-scale deployments, enhancing flexibility and scalability in modern data centers.
- **FC:** FC is a high-speed network protocol designed for block-level storage, transferring data with low latency and high reliability. Ideal for enterprise environments requiring robust performance, FC networks facilitate the efficient consolidation of storage resources. Its capability to cover long distances makes it essential for large data centers where high availability and speed are critical.
- **FCoE:** FCoE integrates the FC protocol with Ethernet frames, allowing FC traffic to traverse Ethernet networks. This convergence enables organizations to leverage their existing Ethernet infrastructure while still enjoying the performance benefits of FC. FCoE simplifies network management by unifying storage and network traffic onto a single fabric, supporting high-speed data transfers crucial for modern enterprise applications.

Each of these protocols plays a crucial role in optimizing storage systems, offering varying levels of performance, scalability, and integration. By understanding these protocols, IT professionals can select the most appropriate technology for their specific storage needs and ensure efficient data management across their infrastructure.

File-sharing protocols

File-sharing protocols play a crucial role in enabling the efficient transfer and access of data across various networks by defining specific rules and formats for requesting and exchanging files between clients and servers. These protocols not only ensure that data can be shared seamlessly but also

address security and compatibility issues across different systems. Here is a more detailed look at some of the most common file-sharing protocols:

- **Server Message Block (SMB):** This protocol facilitates file sharing over **local area networks (LANs)** and is widely used in Windows environments. SMB allows applications to read and write to files and request services from server programs. An extension of SMB, **Common Internet File System (CIFS)**, developed by Microsoft, provides enhanced functionality and is used for sharing files over the Internet as well.
- **Network File System (NFS):** NFS supports file sharing over both local and **wide area networks (WANs)**, particularly within Unix and Linux systems. It allows users to mount remote directories on their local systems, allowing them to access files stored on remote servers as if they were on the local machine. NFS is known for its ability to integrate seamlessly with Unix-based file systems.
- **File Transfer Protocol (FTP):** FTP is designed to transfer files over the Internet between servers and clients. It operates in a client-server model where the client initiates a connection to the server to upload or download files. FTP supports various commands and responses that facilitate efficient file transfer, making it suitable for transferring large files or batches of files.
- **Hypertext Transfer Protocol (HTTP):** HTTP is the foundation of data communication on the web. It is used to deliver files through web browsers and other web-based applications. HTTP enables the retrieval of web pages, images, and other resources from web servers, facilitating a wide range of online activities.
- **Secure Shell (SSH):** SSH provides a secure channel for remote file sharing by encrypting the data during transmission. It is commonly used for accessing remote servers securely and transferring files via protocols such as **Secure Copy Protocol (SCP)** or **Secure File Transfer Protocol (SFTP)**. SSH ensures that data integrity and confidentiality are maintained during file exchanges.

Note:

For further information about SSH and its applications, please visit the official SSH website at <https://www.ssh.com/ssh/>. This resource offers

comprehensive details on SSH protocols, including their secure communication capabilities, encryption methods, and various use cases in remote server management and file transfer.

Each protocol has distinct features that cater to different networking environments and security requirements, making it integral to managing and sharing files efficiently across diverse platforms and networks.

Next, we will explore how HBAs and FC switches facilitate the connection and operation of network storage technologies, such as SANs. HBAs are specialized hardware components that enable servers to interface with storage networks. At the same time, FC switches manage the data traffic within the SAN, ensuring efficient and reliable data transfer between servers and storage devices. Understanding these components is crucial for optimizing the performance and scalability of networked storage solutions.

HBA and FC switches

In this subsection, we will briefly discuss the essential functions of HBAs and FC switches within network storage infrastructures. Understanding these components is crucial for understanding how data storage and retrieval are efficiently managed in modern IT environments.

- An **HBA switch** is a crucial component that provides fiber connectivity between a server and the storage network, enabling efficient data transfer. It essentially acts as an interface that allows the server to communicate with the storage array over high-speed fiber links.
- An **FC switch** operates at Layer 3 of the OSI model, managing and directing data traffic within the network. The FC switch plays a pivotal role in the SAN by facilitating communication between multiple HBAs and storage devices.

The two components are connected by high-speed FC cables, which form the FC fabric. The FC fabric comprises one or more FC switches interconnected to create a network topology that supports high-performance data transfer. This network fabric is essential for maintaining the efficiency and reliability of data storage and retrieval operations within a SAN.

Understanding the interaction between HBAs and FC switches is fundamental for grasping how SANs are structured and managed. In the following section, we will shift our focus to iSCSI hardware, another technology used for connecting network storage systems, and compare its functionalities with those of FC-based systems.

iSCSI hardware

iSCSI is a protocol that enables block-level storage communication over IP networks, providing a means for transferring data between servers and storage devices over long distances using standard networking infrastructure. iSCSI operates by encapsulating SCSI commands—known as **command descriptor blocks (CDBs)**—into IP packets, which are then transmitted over TCP/IP networks. This encapsulation allows the integration of storage devices into IP networks, making it possible to use existing network equipment and infrastructure for storage purposes.

In the iSCSI architecture, storage devices are identified as targets, while clients that initiate communication are referred to as initiators. Each logical disk within a SAN is assigned a unique identifier known as a **Logical Unit Number (LUN)**, which helps in organizing and accessing storage resources effectively. iSCSI uses two key **TCP ports**: **port 860** for the iSCSI system port, which is used for management and configuration, and **port 3260** for the iSCSI default port, which handles data transfers. This dual-port setup facilitates both control and data paths, ensuring efficient communication between initiators and targets.

iSCSI's ability to leverage existing IP networks and its flexibility in integration make it a cost-effective solution for enterprises looking to expand their storage capabilities without the need for specialized hardware. By using standard Ethernet infrastructure, iSCSI enables scalable and manageable storage solutions, enhancing data availability and accessibility across various network environments.

In the next section, we will explore S2D, a technology that enables the creation and management of storage pools. S2D is designed to optimize

storage efficiency and performance, providing a scalable solution for modern data storage needs.

Explaining S2D

S2D is a robust feature introduced in Windows Server 2016 and carried forward into Windows Server 2025, designed to streamline and enhance storage management in data centers. S2D enables the construction of high-availability storage infrastructure by leveraging locally attached disks within a cluster of servers. One of its key capabilities is storage tiering. It improves performance by utilizing faster storage media such as SSDs or NVMe drives as a cache for frequently accessed data, while slower but higher-capacity disks handle less critical data. This tiered approach ensures that data retrieval times are minimized, thereby optimizing the overall performance of the storage system.

Additionally, S2D supports the creation and management of **SDS pools**. By consolidating multiple physical disks into these pools, administrators can create flexible and scalable virtualized storage spaces. These storage pools can be dynamically expanded or reconfigured to meet changing storage needs without the need for physical hardware changes. S2D also incorporates features such as resilience and automated healing, which enhance data protection and ensure continuous availability in case of hardware failures.

Figure 8.7 illustrates the process of creating a new storage pool in Windows Server 2025, highlighting the user-friendly interface and options available for setting up and managing these pools.

 **Figure 8.7 – Creating a new storage pool in Windows Server 2025**

Figure 8.7 – Creating a new storage pool in Windows Server 2025

In the following section, we will explore the deduplication feature, which complements S2D by reducing storage costs and improving efficiency through the elimination of redundant data.