FREE BOOK SAMPLE



CYBERSECURITY / IDENTITY ATTACK SURFACE

INSIDE THE DOMAIN

A Field Manual for Active Directory Penetration Testing

SIX PHASES • TOOL-DRIVEN • PRACTITIONER-FOCUSED

THOMAS A. ANDESRSON

EXPANDED EDITION / 2026

START HERE

Your free reader preview

Before buying the complete book, explore its writing style, practical structure and evidence-led approach to authorized Active Directory security assessments.

Included in this sample

- Copyright notice, legal disclaimer, preface and engagement safeguards.
- Executive summary and Active Directory core concepts, objects, authentication and trusts.
- Expanded foundations on identity boundaries, protocols, certificate services and attack paths.
- Overview of the six-phase assessment methodology.
- Complete Phase 1: reconnaissance, enumeration and evidence-quality guidance.
- The expanded edition source key, including linked technical references.

Available only in the complete edition

- Phases 2–6: initial access, privilege escalation, lateral movement, credential access and post-exploitation.
- Tooling directory, defensive telemetry, hardening, case studies, worksheets and laboratory design.
- Further appendices, reporting guidance and risk-based attack-path prioritization.

A note about navigation

The table of contents and page numbering are preserved from the complete edition. This preview only includes the sections listed above; later chapters remain in the paid book.

EXPANDED EDITION / OCTOBER 2026

Inside the Domain

A Field Manual for Active Directory Penetration Testing

A Six-Phase, Tool-Driven Methodology — From Reconnaissance to Domain Dominance

Author: Thomas A. Andesrson

Expanded Edition 1.1 | 2026

Copyright and publishing information

Copyright © 2026 Thomas A. Andesrson. All rights reserved, subject to the author holding the necessary rights and permissions to publish the underlying material. No reproduction, redistribution, resale, or adaptation is permitted except as allowed by applicable law or written permission.

Names of software, products, services, trademarks, and projects mentioned in this publication remain the property of their respective owners. References do not imply affiliation, endorsement, sponsorship, or certification by those parties.

This publication is an independent practitioner field manual. It is not an official Microsoft publication or a substitute for an organization’s approved security testing policies.

Legal Disclaimer

EDUCATIONAL PURPOSES AND AUTHORIZED SECURITY TESTING ONLY

All information in this book, including techniques, commands, tools, attack-path descriptions, and examples, is provided exclusively for education, training, lawful research, controlled laboratory exercises, and professional security assessments conducted with explicit authorization. It must not be used for unauthorized access or any activity that violates applicable law, contracts, or organizational policy.

Prior written authorization and scope

Before running any procedure against a real environment, the reader must obtain explicit, prior written authorization from the owner or other legally empowered authority for all affected systems, accounts, networks, and data. The permission must define the scope, acceptable techniques, testing window, operational limits, data-handling requirements, escalation contacts, and any prohibited actions. Purchasing or possessing this book does not grant authorization to test third-party systems.

Responsibility for use and misuse

The reader is solely responsible for deciding whether a proposed activity is lawful and properly authorized, for obtaining required permissions, and for complying with all relevant rules and safeguards. The author neither directs, encourages, condones, nor participates in unauthorized activities. Any misuse of the material is an independent act of the person who undertakes it and is not authorized by, endorsed by, or attributable to the author.

Operational risk and sensitive information

Some techniques can cause account lockouts, service interruption, configuration changes, loss of data, or exposure of passwords, hashes, Kerberos tickets, certificates, and other confidential information. Test only within agreed limits; first reproduce intrusive steps in an isolated laboratory, use change-control procedures and rollback plans, and securely handle, minimize, and dispose of captured evidence.

No warranty; limitation of liability

The material is provided "as is" and "as available," without any representation or warranty as to completeness, accuracy, safety, legality, suitability, fitness for a particular purpose, or compatibility with current software versions. Commands, vulnerabilities, mitigations, and tool behavior may change. To the maximum extent permitted by applicable law, the author disclaims liability for losses, damage, disruption, claims, or other consequences arising from use or misuse of the material. Nothing in this notice excludes any liability that cannot lawfully be excluded.

No legal advice or third-party affiliation

This notice is general information, not legal advice. Applicable laws differ by jurisdiction; consult qualified legal counsel where necessary. Third-party tools and trademarks are cited for identification and technical context only. The author does not claim endorsement by their maintainers or vendors.

Preface

Active Directory remains central to authentication, authorization, and administration in many enterprise Windows environments. Its interconnected identities, computers, policies, trusts, and certificate services make it essential to understand the relationships between individual findings and complete attack paths.

To the author's knowledge, there is no single universally accepted, sufficiently comprehensive, and consistently maintained penetration testing methodology that addresses the full Active Directory attack lifecycle in one place. Valuable best practices, tools, and checklists exist across community projects, vendor resources, and individual research publications, but they are often dispersed across sources, designed for specific attack types, or not updated frequently enough to account for newer techniques and evolving attack paths. This observation is a practical motivation for the manual, not a claim that existing resources have no value.

Inside the Domain was developed to bring this scattered operational knowledge into a coherent, six-phase field methodology. The workflow starts with reconnaissance, advances through initial access and credential-risk assessment, covers privilege escalation and lateral movement, and concludes with the validation of high-impact exposures, persistence risks, and domain-level compromise scenarios. The phases support assessment planning; they are not a compulsory order of operations.

The aim is a practitioner-oriented reference with task objectives, tools, executable examples, and related attack paths kept close together. Because both attacker tradecraft and defensive controls change over time, this edition should be treated as a versioned working reference that requires periodic technical review. Every example must be verified in the relevant tool version and must be used only within an explicitly authorized scope, as set out in the Legal Disclaimer.

Who this book is for

Professional penetration testers, red-team operators, security engineers, incident responders, and administrators who need a structured approach to assessing Active Directory security. Familiarity with Windows domains, Kerberos, NTLM, networking, and command-line tools is assumed.

How to use this field manual

- Start with explicit authorization and the engagement scope; not every technique is appropriate for every environment.
- Use the six phases as an assessment map, not a mandatory attack chain.
- Review prerequisites and validate syntax and tool compatibility in a disposable lab.
- Record target scope, timestamp, observed behavior, evidence, and potential business impact for every confirmed finding.
- Apply safeguards for authentication testing, credential handling, domain-controller changes, and persistence demonstrations.
- Finish with cleanup, restoration checks, and remediation-focused reporting.

Reading conventions

Angle-bracket placeholders (for example, <DOMAIN> and <DC_IP>) indicate environment-specific values. Lines beginning with # are typically command comments. Tool names and command-line switches reflect the source manuscript and should be checked against current documentation before use.

A note on this expanded edition

The source manuscript supplied the command-oriented six-phase methodology. This revised edition retains that material and adds independently written explanations, assessment safeguards, defender-oriented analysis, case studies, worksheets, and linked source references. Statements based on general technical documentation are clearly separated from the reader's own environment-specific observations. No new offensive exploitation instructions have been added as substitutes for the original commands.

Technical review boundary: example tool invocations inherited from the first edition are not claimed to have been executed or exhaustively checked against every current release. Confirm syntax, environment impact, authorization and logging before using any command.

Engagement Safeguards

A technically valid test can still be inappropriate if it falls outside the agreement, compromises data handling, or disrupts business systems. Agree on guardrails before testing.

Pre-engagement approval checklist

- Written authorization, named points of contact, and emergency stop / escalation procedure
- Approved domains, forests, subnets, hosts, accounts, and testing windows
- Explicit approval levels for credential attacks, relay/coercion, remote execution, and high-impact configuration changes
- Lockout thresholds, volume limits, throttling, and monitoring coordination
- Data handling for passwords, hashes, tickets, certificates, dumps, logs, and screenshots
- Evidence storage, encryption, retention period, destruction plan, and clean-up verification

Evidence-first assessment workflow

For each tested control, log the hypothesis, permitted scope, prerequisites, exact tool version, observed result, validation method, and a remediation recommendation. A failed attempt is not evidence that an attack path is impossible; record environmental constraints and testing limitations.

Reporting and remediation

Distinguish verified findings from potential or untested exposures. Explain what permission or trust relationship permitted the behavior, who or what is affected, how confidently the result was reproduced, and how the organization can fix or monitor the issue. Obtain approval before retaining sensitive credentials or domain data.

Contents

Select a chapter title to jump directly to that section.

Legal Disclaimer

Preface

Engagement Safeguards

1 Executive Summary

2 Introduction to Active Directory

3 Methodology Overview

4 Phase 1 — Reconnaissance & Information Gathering

5 Phase 2 — Initial Access & Credential Harvesting

6 Phase 3 — Privilege Escalation

7 Phase 4 — Lateral Movement

8 Phase 5 — Credential Access & Dumping

9 Phase 6 — Persistence, Post-Exploitation & Domain Dominance

10 Tool Directory & Resources

11. Planning a Defensible Active Directory Engagement

12. Monitoring, Detection, and Telemetry

13. Hardening Priorities and Remediation Engineering

14. Applied Assessment Cases and Reasoning Exercises

15. Field Worksheets and Evidence Templates

16. ATT&CK Mapping and Control Review Guide

17. Attack-Path Prioritization and Risk-Based Decisions

18. Laboratory Design and Edition Maintenance

Appendix A Detection and Control Crosswalk

Appendix B Sample Finding Classifications

Appendix C Representative Reporting Language

Expanded-Edition Sources and Citation Key

References and Further Reading

Fieldwork Closeout & Edition Notes

1. Executive Summary

This document provides a comprehensive methodology for conducting penetration testing against Active Directory (AD) environments. Active Directory serves as the backbone for identity and access management in most enterprise Windows environments, providing centralized authentication, authorization, and resource control.

This methodology is structured into **six distinct phases**, progressing from initial reconnaissance through complete domain dominance. Each phase contains detailed tasks, sub-tasks, specific tools, and exact commands.

2. Introduction To Active Directory

Why a Unified Active Directory Methodology Matters

Despite the broad adoption of Active Directory, practitioners do not have a single universally agreed, end-to-end penetration testing methodology that exhaustively covers its changing attack surface. Established best practices and technical references are valuable, but the guidance is distributed across many resources; some focus on a single technique or component, and some do not yet account for newer tools, exploitation patterns, or attack-path combinations. This field manual consolidates a repeatable six-phase workflow for authorized assessment while recognizing that no static document can be permanently comprehensive.

2.1 What is Active Directory?

Active Directory (AD) is Microsoft's directory service for Windows domain networks. It provides centralized management for:

- User accounts and credentials
- Computer accounts
- Security groups
- Access permissions
- Security policies (GPOs)
- Authentication services

2.2 Core Components

Domains

The fundamental administrative and security boundary in AD, containing:

- Users, groups, computers
- Security policies
- Unique database

Trees

Collection of domains sharing:

- Contiguous namespace
- Automatic transitive trusts

Forests

Top-level container consisting of:

- One or more trees
- Common schema
- Global catalog
- Ultimate security boundary

Domain Controllers (DCs)

Servers that:

- Host AD DS role
- Store NTDS.dit database
- Handle authentication
- Replicate directory data

Organizational Units (OUs)

Containers allowing:

- Administrative delegation
- GPO application
- Logical grouping

2.3 Key Objects

Object	Description	Key Attributes
User	Person or service account	sAMAccountName, userPrincipalName, memberOf
Computer	Domain-joined machine	dNSHostName, operatingSystem, SPN
Group	Collection of objects	member, groupType
GPO	Policy settings	displayName, gPCFileSysPath

2.4 Authentication Protocols

Kerberos (Primary)

Ticket-based authentication:

- **AS-REQ/AS-REP**: Client obtains TGT
- **TGS-REQ/TGS-REP**: Client obtains service ticket
- **AP-REQ/AP-REP**: Client accesses service

Key Encryption Types:

- RC4-HMAC (0x17): Uses NT hash
- AES256 (0x12): Most secure

NTLM (Legacy)

Challenge-response protocol:

- NEGOTIATE: Client announces intent
- CHALLENGE: Server sends challenge
- AUTHENTICATE: Client responds with hash

2.5 Trust Relationships

Trust Type	Direction	Transitivity
Parent-Child	Two-way	Transitive
External	One/Two-way	Non-transitive
Forest	One/Two-way	Transitive

2.6 Common Terminology

Term	Definition
SAM	Security Accounts Manager
NTDS.dit	AD database file
LSASS	Local Security Authority Subsystem
SID	Security Identifier
ACL	Access Control List
GPO	Group Policy Object
SYSVOL	Shared GPO folder
krbtgt	Kerberos service account
SPN	Service Principal Name

2.7 Identity, Protocols, and Attack-Path Foundations

Expanded practitioner context: the following material explains the security boundaries, assessment evidence, detection opportunities and limitations surrounding the original command-oriented field

methodology. It is supplemental analysis; the technical examples in the preceding sections remain preserved from the source manuscript and require tool-version validation.

2.7.1 Forests, domains, and administrative authority

TECHNICAL CONTEXT

An Active Directory forest is a common schema, configuration and trust context, not merely a collection of machines with similar DNS names. Domains partition identity data and support administration and replication; organizational units delegate management inside a domain. A domain can be a convenient unit for authentication policy, but the forest is generally the meaningful upper boundary for traditional AD DS trust analysis. Assessors should trace which authorities can create principals, set permissions, administer certificate enrollment, and alter Group Policy rather than inferring safety from a diagram that labels a domain as isolated.

ASSESSMENT AND EVIDENCE

Build a read-only inventory with forest name, domain names, functional levels, domain controller OS versions, trust links, privileged groups and OU delegation points. For each link, identify direction, transitivity and any relevant filtering or selective authentication. Evidence should include the configuration value and where it was obtained, not a blanket claim that two domains are fully accessible to each other. Identify whether a trust represents a realistic route to a high-value target, and mark additional assumptions as unvalidated.

DETECTION AND HARDENING

Review forest-level administration, restrict privileged activity to trusted hosts, and document any cross-domain delegation. Ensure administrative boundaries match the organization’s business and recovery model. Detection should cover changes to trust objects, group membership, and domain policies with properly configured audit policy and access control lists. The useful remediation is often narrowing delegated rights, not removing every trust relationship.

INTERPRETATION AND LIMITATIONS

A transitive trust does not grant every user administrator access in a peer domain. Authentication reachability, authorization, and actual permissions are separate facts. Hybrid identity adds additional planes that a strictly on-premises methodology must explicitly scope rather than silently absorb.

> **Lab thought experiment |** Explain why a trusted external domain can authenticate a user without giving that user administrative access to a domain controller.

Source trail: [E01], [E18], [E14].

2.7.2 Domain controllers as identity control-plane assets

TECHNICAL CONTEXT

Domain controllers are hosts, but they also hold the directory database, service Kerberos and LDAP requests, apply replication logic, and participate in high-value policy decisions. A routine endpoint finding and a domain-controller finding may have similar technical names while posing very different organizational risks. Protecting a domain controller involves its OS configuration, directory permissions, replication rights, backup system, management channel and underlying virtualization or storage administrative privileges.

ASSESSMENT AND EVIDENCE

Record all controllers, site assignments, replication links, writeable versus read-only roles, backed-up configuration and relevant identity infrastructure. Map the administrative accounts and computers that can manage controllers directly or via infrastructure. Verify that logging reaches a protected collection location. Ask who can change the controller’s recovery process, hypervisor snapshot, virtual disks, or backups; control of those paths can be comparable to direct controller administration.

DETECTION AND HARDENING

Use dedicated privileged access workstations and least-privilege roles for domain administration. Patch, harden and monitor domain controllers as control-plane systems. Establish an incident-response playbook for credential compromise and identity recovery, including safeguards around the KRBTGT account. Changes to replication or directory permissions should receive extra scrutiny.

INTERPRETATION AND LIMITATIONS

Merely finding a domain controller port open is expected in a functioning domain. The issue is an unsafe configuration, exposure to an untrusted path, or an excessive identity privilege—not the existence of Kerberos, LDAP or DNS itself.

Source trail: [E01], [E18], [E27], [E10].

2.7.3 Kerberos: tickets, principals and the KDC

TECHNICAL CONTEXT

Kerberos is a symmetric-key ticket-based authentication system. A client first requests a ticket-granting ticket from the KDC, normally operating on a domain controller, and then asks for service tickets to reach servers identified by service principal names. The protocol separates proof of identity from access to a service; tickets contain time-sensitive information and authorization data. It is useful to think in terms of trust in the KDC, secrecy of service-account keys, ticket lifetimes and how services validate the authorization context.

ASSESSMENT AND EVIDENCE

In a lab, diagram the AS exchange, TGT, TGS exchange, and service authentication flow. Inventory SPNs, encryption settings, service-account ownership, and delegation configurations rather than treating every ticket request as an exploit. When explaining a suspected weakness, state which identity, key or attribute would have to be exposed or misconfigured for the threat to become plausible. Preserve only minimal event metadata from ticket exchanges in assessment evidence.

DETECTION AND HARDENING

Prefer modern encryption supported by the estate, robust service-account secrets, managed service accounts, narrow privileges and appropriately protected administrator sessions. Audit authentication events, especially anomalous requests for service tickets or failures associated with misconfigured accounts. Establish a baseline before treating a single event as proof of compromise.

INTERPRETATION AND LIMITATIONS

Not all service tickets are encrypted with an NT hash: enctype and key-derivation details depend on account configuration and negotiation. A ticket request is usually normal behavior; risk assessment must combine the existence of a crackable offline artefact with the entropy and privilege of its underlying service identity.

Source trail: [E02], [E08], [E21], [E15].

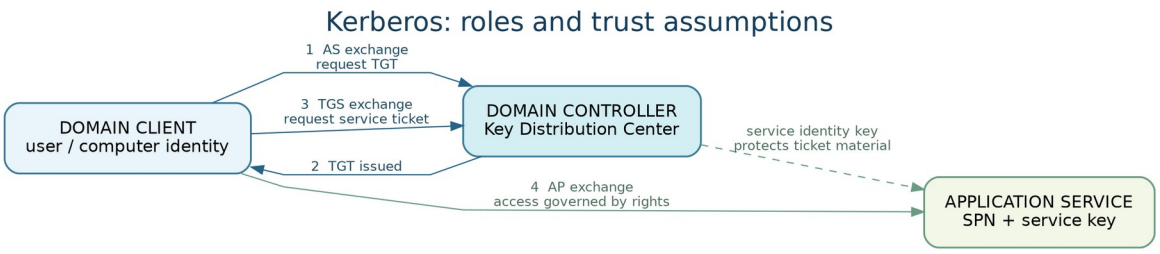


Figure 1. Kerberos roles and ticket relationships (original schematic; conceptual basis [E02]).

2.7.4 NTLM, downgrade behavior and signing

TECHNICAL CONTEXT

NTLM provides challenge-response authentication but does not provide the same mutual server authentication guarantees as Kerberos. Its continued use often reflects legacy applications, name-resolution behavior or compatibility constraints. Relay exposure is a protocol and configuration interaction: material captured in one authentication context may be forwarded to another endpoint that does not require protections sufficient to bind the message to the intended service.

ASSESSMENT AND EVIDENCE

Discover where NTLM is still negotiated and inventory SMB and LDAP protections without deliberately capturing production users' credentials. Associate every relay-relevant endpoint with its version, role, signing policy, transport protection and whether authenticated access is permitted. A complete finding identifies both sides of the trust failure and avoids claiming that every NTLM-enabled device is exploitable.

DETECTION AND HARDENING

Plan a phased reduction in NTLM dependencies, require SMB signing where supported, enforce LDAP signing and evaluate channel binding. Validate application compatibility first, because abruptly changing defaults can interrupt identity-dependent business services. Recent Windows and Windows Server versions have stronger defaults, so measure effective policy rather than projecting old defaults onto modern deployments.

INTERPRETATION AND LIMITATIONS

SMB signing addresses integrity and relay exposure for that protocol; it is not a comprehensive replacement for identity hygiene. LDAP signing and channel binding are related but distinct controls, and a secure LDAPS channel alone does not automatically imply channel binding is required.

Source trail: [E02], [E03], [E04], [E30], [E11].

2.7.5 LDAP attributes, access decisions and effective rights

TECHNICAL CONTEXT

Directory entries are objects with attributes controlled by a security descriptor. Explicit access control entries, inherited permissions, ownership, object type and the operation requested all influence whether a principal can perform a change. Reading an attribute may be routine, while permission to edit that attribute can create an indirect path to another identity. Effective privilege is therefore a relationship among an account, a target object and a capability—not a property of a user name alone.

ASSESSMENT AND EVIDENCE

Develop an inventory of delegated rights to high-value users, groups, GPOs, computer objects and certificate templates. Distinguish read access from changing membership, resetting passwords, modifying DACLs or writing service attributes. Validate suspicious rights with documented ACL review before proposing a destructive proof. Record the principal, object DN, permission, inheritance chain and expected administrative rationale.

DETECTION AND HARDENING

Remove excessive delegation, use scoped administrative roles and review inherited permissions after organizational changes. Protect identity and policy objects that control many descendants. Monitor unexpected directory modifications with appropriate audit categories and SACLs; event collection without those prerequisites gives a false sense of coverage.

INTERPRETATION AND LIMITATIONS

GenericAll, GenericWrite and WriteDACL are not interchangeable; the viable abuse depends on object class and allowed operations. A graph edge is a hypothesis about a path and should be checked against environmental policies and current privileges.

Source trail: [E01], [E14], [E19], [E24].

2.7.6 Service accounts and the identity lifecycle

TECHNICAL CONTEXT

Service accounts often combine persistent credentials, delegated resource access and SPN registrations. Human password-change workflows may not fit unattended services, producing long-lived credentials and unclear ownership. Group managed service accounts are a defensive pattern because password management is coordinated with Active Directory rather than embedded in scripts or manually rotated across servers.

ASSESSMENT AND EVIDENCE

Inventory service identities, linked workloads, SPNs, password age, privilege memberships, allowed logon types and owners. Prioritize accounts that administer servers, hold domain privileges or connect otherwise segregated services. Evidence should show the exact exposure mechanism—such as discoverable configuration, over-privileged ACL or obsolete service—rather than publishing a live secret in the report.

DETECTION AND HARDENING

Use dedicated identities, gMSAs where feasible, least privilege and strong managed secrets. Explicitly restrict interactive logon and service installation rights according to purpose. Ensure emergency rotation can be performed without disrupting production, and document service owners before changing keys or removing privileges.

INTERPRETATION AND LIMITATIONS

A long password age is not alone proof of compromise or weak entropy. Some managed identities rotate automatically, while legacy applications may require staged migrations; recommend a controlled plan rather than a generic immediate password reset.

Source trail: [E21], [E08], [E18].

2.7.7 Certificate-based identity and PKI trust

TECHNICAL CONTEXT

Active Directory Certificate Services can issue certificates used for authentication, encryption and signing. A certificate template determines enrollment policy, intended usages, identity information and governance requirements. When a template allows a low-privileged requester to obtain a certificate for a more privileged identity, a PKI configuration issue can undermine otherwise strong password controls. The trust decision extends across template ACLs, CA configuration, enrollment endpoints and certificate mapping.

ASSESSMENT AND EVIDENCE

Record deployed CAs, issued templates, enrollment rights, template ownership, EKUs, subject name supply options, issuance requirements, enrollment agents and web enrollment endpoints. Pair any potentially dangerous combination with its actual scope and compensating controls. Inspect CA backup and private key protections, and identify whether certificate-management privileges are separated from ordinary domain administration.

DETECTION AND HARDENING

Harden enrollment templates, constrain subject identity selection, require appropriate approvals and protect CA administration. Audit CA configuration and certificate issuance as security-sensitive operations. Review Defender for Identity posture assessments as one source of findings, but validate their applicability to the installed Windows and template versions.

INTERPRETATION AND LIMITATIONS

The presence of a template option does not necessarily mean an attack is exploitable: authorized signatures, enrollment policy, EKUs and mapping requirements may change the outcome. State what was observed and any missing verification steps.

Source trail: [E22], [E06], [E07].

2.7.8 Attack graphs: from isolated errors to paths

TECHNICAL CONTEXT

An attack path is a chain of conditions or capabilities that can connect a low-privilege starting point with a protected outcome. A graph model represents users, groups, computers, sessions, domains and control relationships as nodes and edges. This is more explanatory than a flat list of vulnerabilities because a single permission becomes significant when joined to another exposure. The graph should be treated as a model of the assessed environment, not as proof that a chain was executed.

ASSESSMENT AND EVIDENCE

Identify the initial identity, target asset, required privilege and each intermediate condition. Keep evidence for each edge, including timestamp and collection scope. Mark edges as confirmed, inferred, blocked or untested; stale group membership or session data can misrepresent viable paths. During a report, use a small path diagram to show the logic without exposing reusable secrets.

DETECTION AND HARDENING

Defenders can break paths by removing one risky delegation, protecting an administrative workstation, enforcing signing or rotating a weak service secret. Prioritize chokepoints that reduce many paths rather than simply counting vulnerable nodes. Recompute the graph after remediation to show whether reachability actually changed.

INTERPRETATION AND LIMITATIONS

Shortest does not mean easiest or most likely. Edge reliability, monitoring, permissions, business constraints and defensive protections affect practicality. BloodHound is one source of graph analytics, and output should be interpreted alongside direct configuration evidence.

Source trail: [E19], [E18], [E20].

2.7.9 Security boundaries and hybrid environments

TECHNICAL CONTEXT

Modern enterprise identity may include on-premises AD DS, Microsoft Entra ID, synchronization services, federation and SaaS applications. The trust and privilege assumptions for a traditional domain-controller attack do not automatically apply to cloud identities. A sync connector, administrative bridge or privileged service may span the environments, however, so a purely on-premises assessment should explicitly identify the boundary it did and did not cover.

ASSESSMENT AND EVIDENCE

During scoping, distinguish AD DS domains, federation services, tenant roles, synchronized identities and device registration. Document which components are in scope, which are third-party managed and which require separate consent. Map only high-level dependencies for out-of-scope components; testing them without authorization is not an extension of the approved on-premises domain assessment.

DETECTION AND HARDENING

Apply control-plane protection to every administrative bridge, not just domain controllers. Review the privilege of synchronization services, emergency access and identity lifecycle workflows. Separate administrator and user access paths and avoid treating a cloud identity as protected merely because on-premises tiering has been implemented.

INTERPRETATION AND LIMITATIONS

This manual primarily addresses AD DS penetration testing. Entra ID, AD FS and hybrid-specific exploitation need separate evidence standards and assessment authorizations; brief mentions here are context, not a comprehensive cloud security methodology.

Source trail: [E18], [E27], [E28].

3. Methodology Overview



4. Phase 1: Reconnaissance & Information Gathering

4.1 Objectives

- Identify network topology and AD infrastructure
- Discover Domain Controllers and key services
- Enumerate valid usernames without credentials
- Identify potential attack vectors

4.2 Network Discovery

4.2.1 Port Scanning

Task: Identify AD-related services

```
# Basic AD port scan
nmap -p 53,88,135,139,389,445,464,593,636,3268,3269,3389,5985,9389 -sV -O <target_ip>
# Comprehensive scan with scripts
nmap -p 53,88,135,139,389,445,464,593,636,3268,3269,3389,5985,9389,49152-65535 \
--script smb-enum-shares,smb-enum-users,ldap-rootdse,ldap-search,\
krb5-enum-users,smb-os-discovery,smb-vuln-ms17-010,smb2-security-mode \
-A <target_ip>
# Quick ping sweep
nmap -sn <network>/24
# SMB version enumeration
nmap --script smb-protocols -p 445 <target_ip>
```

AD-Related Ports:

Port	Service	Attack Surface
53	DNS	Zone transfer, poisoning
88	Kerberos	Roasting, ticket attacks
135	MS-RPC	DCOM exploitation
137-139	NetBIOS	SMB relay, enumeration
389	LDAP	Injection, anonymous bind
445	SMB	Relay, EternalBlue, enumeration
464	Kerberos Password	Password spraying
636	LDAPS	Certificate attacks
3268-3269	Global Catalog	LDAP enumeration
3389	RDP	BlueKeep, credential forwarding
5985-5986	WinRM	Remote execution

4.2.2 Service Identification

```
# Identify Domain Controller
nmap -p 88,389 -sV <target_ip>
# Check SMB signing (critical for relay attacks)
nmap --script smb2-security-mode -p 445 <target_ip>
# NetExec/CrackMapExec for AD identification
nxc smb <target_ip>
crackmapexec smb <target_ip>
```

4.3 DNS Enumeration

4.3.1 DNS Record Discovery

```
# Enumerate DNS records
dig @<dc_ip> <domain> any
dig @<dc_ip> <domain> axfr  # Zone transfer attempt
# SRV records for AD services
dig @<dc_ip> _ldap._tcp.<domain> SRV
dig @<dc_ip> _kerberos._tcp.<domain> SRV
dig @<dc_ip> _gc._tcp.<domain> SRV
# nslookup queries
nslookup -type=SRV _ldap._tcp.dc._msdcs.<domain> <dc_ip>
nslookup -type=any <domain> <dc_ip>
# dnsrecon for comprehensive enumeration
dnsrecon -d <domain> -n <dc_ip> -t std
dnsrecon -d <domain> -n <dc_ip> -t axfr
```

4.3.2 DNS Enumeration via LDAP

```
# Using ldapsearch
ldapsearch -x -h <dc_ip> -s base namingContexts
# Using nxc
nxc ldap <dc_ip> -u '' -p '' --dns-records
```

4.4 Anonymous Access Testing

4.4.1 LDAP Anonymous Bind

```
# Test anonymous LDAP access
ldapsearch -x -h <dc_ip> -b "DC=domain,DC=com" -s sub "(objectclass=*)"
# Using nxc
nxc ldap <dc_ip> -u '' -p ''
# enum4linux comprehensive enumeration
enum4linux -a <dc_ip>
enum4linux -a -u '' -p '' <dc_ip>
```

4.4.2 SMB Anonymous Access

```
# Check null session
smbclient -L //<dc_ip> -N
smbmap -H <dc_ip>
smbmap -u '' -p '' -H <dc_ip>
# CrackMapExec/NetExec
nxc smb <dc_ip> -u '' -p '' --shares
nxc smb <dc_ip> -u 'guest' -p '' --shares
# enum4linux
enum4linux -U <dc_ip>  # Users
enum4linux -G <dc_ip>  # Groups
enum4linux -P <dc_ip>  # Password policy
```

4.4.3 RPC Enumeration

```
# rpcclient anonymous connection
rpcclient -U "" <dc_ip>
rpcclient> enumdomusers
rpcclient> enumdomgroups
rpcclient> querydominfo
# Using impacket
impacket-lookupsid <domain>/''@<dc_ip>
```

4.5 User Enumeration

4.5.1 Kerbrute User Enumeration

```
# Install kerbrute
go install github.com/ropnop/kerbrute@latest
# User enumeration (no authentication required)
kerbrute userenum --dc <dc_ip> -d <domain> <userlist.txt>
# Brute force with known password
```

```
kerbrute bruteuser --dc <dc_ip> -d <domain> <password_list.txt> <username>
# Password spray (single password, many users)
kerbrute passwordspray --dc <dc_ip> -d <domain> <userlist.txt> <password>
# User-as-pass spray
kerbrute passwordspray --dc <dc_ip> -d <domain> <userlist.txt> --user-as-pass
```

4.5.2 LDAP User Enumeration

```
# With credentials
ldapsearch -x -H ldap://<dc_ip> -D "<user>@<domain>" -w '<password>' \
-b "DC=domain,DC=com" "(objectClass=user)" sAMAccountName
# Using GetADUsers (Impacket)
impacket-GetADUsers <domain>/<user>:<password> -dc-ip <dc_ip> -all
```

4.5.3 RID Cycling

```
# Impacket lookupsid
impacket-lookupsid <domain>/<user>:<password>@<dc_ip>
# enum4linux
enum4linux -r -u <user> -p <password> <dc_ip>
```

4.6 IPv6 Reconnaissance

```
# Check for IPv6-enabled hosts
passivediscovery6 <interface>
# IPv6 scanning
nmap -6 -sn <ipv6_network>
```

4.7 Reconnaissance Interpretation and Evidence Quality

Expanded practitioner context: the following material explains the security boundaries, assessment evidence, detection opportunities and limitations surrounding the original command-oriented field methodology. It is supplemental analysis; the technical examples in the preceding sections remain preserved from the source manuscript and require tool-version validation.

4.7.1 Scoping network exposure without over-scanning

TECHNICAL CONTEXT

Network discovery has two questions: which services are genuinely reachable from the approved test position, and which security-relevant roles do they represent? A port scan is a measurement at a specific time from a specific vantage point. A closed or filtered port is not a claim that the service is absent elsewhere. Segmented environments and VPN routes often yield different answers for the same host, so discovery reports should preserve the source subnet and other access conditions.

ASSESSMENT AND EVIDENCE

Build an authorized target inventory before running checks. Agree on connection volume and avoid testing domain controllers with intrusive vulnerability scripts unless explicitly approved. Record host/IP mapping, reachable service, service banner where available, protocol version, test node and time. Supplement scanning with server-side configuration export or approved asset-management records if remote observations are ambiguous.

DETECTION AND HARDENING

Compare exposures to the intended segmentation design: authentication services may have to be available to domain clients, but management services should be narrower. Investigate unnecessary remote administration access, outdated protocol versions and unexpected cross-segment reachability. Monitoring teams can use the approved scan window to distinguish test activity from unrelated traffic.

INTERPRETATION AND LIMITATIONS

Presence is not vulnerability; version fingerprinting can be wrong, and active probing can produce partial results. Assign conclusions only after correlating multiple evidence sources. When access restrictions prevent validation, document a coverage limitation rather than a pass.

Source trail: [E20], [E01], [E04].

4.7.2 DNS and SRV records as infrastructure evidence

TECHNICAL CONTEXT

AD-integrated DNS enables clients to locate domain controllers, Kerberos services and global catalogs using service records. An assessment uses those records to describe the identity topology. Misconfigured zone-transfer permissions can expose inventories, but simply resolving SRV records is normal and expected: clients need those records to function. Distinguish information disclosure from a DNS availability or integrity problem.

ASSESSMENT AND EVIDENCE

Document authoritative zones, controller discovery records, name-server delegation and the intended exposure of records. Verify that zone-transfer policy matches the authorized secondaries; do not request or retain entire zone contents when a configuration-level control check suffices. Correlate names with approved configuration inventories, recognizing that decommissioned records and split-horizon DNS can mislead results.

DETECTION AND HARDENING

Restrict zone transfers to approved servers and monitor unexpected changes to critical DNS records. Protect DNS administration because service discovery and identity availability depend on it. Recommend cleanup of stale objects and consistent change tracking rather than hiding standard AD SRV records required for normal authentication.

INTERPRETATION AND LIMITATIONS

An unexpected DNS record can signal error, migration or malicious tampering; context matters. Avoid equating visibility of a domain controller name with a security flaw. Document what an unauthenticated observer can actually learn and whether that learning changes the threat model.

Source trail: [E01], [E20].

4.7.3 Anonymous enumeration and information exposure

TECHNICAL CONTEXT

Unauthenticated LDAP, SMB or RPC interactions may disclose directory or operating-system metadata depending on policy and service design. A finding should identify the actual object classes and fields returned and the sensitivity of those attributes. Anonymous access itself is not synonymous with unauthorized modification or credential leakage: read-level disclosure and write-level control are different security outcomes.

ASSESSMENT AND EVIDENCE

Use the smallest read-only query needed to demonstrate exposure. Describe whether the test used null credentials, guest access or ordinary authenticated access. Capture redacted example output that shows the attribute class rather than full employee lists. Verify server-side intended policy with an administrator before treating broad access as an accidental misconfiguration.

DETECTION AND HARDENING

Restrict unnecessary anonymous directory binds and exposed shares, but account for legitimate discovery and application requirements. Audit guest access and review object-level ACLs. A proportional mitigation might narrow only sensitive attributes or network reachability while leaving required public metadata accessible to business applications.

INTERPRETATION AND LIMITATIONS

Some tools blur anonymous and guest authentication; record the negotiated identity. A failed connection may be caused by firewall or incompatible tooling. Use an explicit evidence confidence level and avoid severity inflation from enumeration alone.

Source trail: [E03], [E01], [E20].

4.7.4 Users, service identities and privilege hints

TECHNICAL CONTEXT

Usernames are not passwords, and enumeration does not by itself prove an exploitable account. Nevertheless, knowing valid principal names can lower uncertainty for later authentication attacks and reveals aspects of organization structure. A useful reconnaissance inventory differentiates employees, service accounts, managed identities, computers and built-in groups rather than aggregating them into one undifferentiated list.

ASSESSMENT AND EVIDENCE

When permitted, identify naming patterns, disabled and stale accounts, high-value groups, SPN assignments and service principals. Avoid generating broad dictionaries from personal data if the engagement does not require them. Use unique synthetic placeholders in sample evidence, and do not replicate production usernames in the published book or publicly shared workbooks.

DETECTION AND HARDENING

Combine directory hygiene with phishing-resistant authentication where supported, managed service accounts, lockout-aware protections and monitoring for abnormal authentication patterns. Privilege classification should be based on effective rights and business ownership; being a member of a group with an impressive name is not necessarily a control-plane privilege.

INTERPRETATION AND LIMITATIONS

An account marked disabled today might be re-enabled for business reasons; finding currency must be recorded. User-enumeration responses and lockout behavior vary by protocol and domain configuration, and should be interpreted only under documented test constraints.

Source trail: [E01], [E08], [E15], [E21].

4.7.5 Reconnaissance deliverable: asset and exposure map

TECHNICAL CONTEXT

The phase output should be a reconciled view of identity-relevant assets, not a collection of unfiltered terminal logs. Organize the map by domain, trust, controller, service, certificate authority, administrative entry point, privileged group and evidence owner. This structure lets the next phase concentrate on plausible trust failures rather than trying every technique against every host.

ASSESSMENT AND EVIDENCE

For each observed endpoint capture an anonymized identifier, role, source location, protocol, effective hardening state, method of verification, collection time and confidence. Include a separate table of assets excluded by scope and assumptions awaiting confirmation. Where two sources disagree, retain both observations with explanation rather than silently choosing the more dramatic result.

DETECTION AND HARDENING

Deliver a network-reachable surface summary and a short prioritized action list: unexpected management exposure, risky LDAP/SMB settings, uncontrolled identity services and unowned service accounts. Provide defenders with sufficient metadata to locate the object or setting and to verify a change later.

INTERPRETATION AND LIMITATIONS

An accurate negative finding is valuable: explicitly recording that signing is required, or that a sensitive port is not accessible from the test position, helps illustrate which paths were ruled out. It should not be generalized beyond that scope or time window.

Source trail: [E20], [E03], [E04], [E19].

Expanded-Edition Sources and Citation Key

Sources marked [E01] through [E30] support the expanded explanatory sections. The technical command material predating this edition remains part of the original manual and has not been independently validated against every current tool release. The links below identify primary documentation; readers should verify new updates before testing.

[E01] Microsoft. Understanding the Active Directory Logical Model.

https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/plan/understanding-the-active-directory-logical-model

[E02] Microsoft. Kerberos Authentication Overview in Windows Server.

https://learn.microsoft.com/en-us/windows-server/security/kerberos/kerberos-authentication-overview

[E03] Microsoft. LDAP Signing for Active Directory Domain Services.

https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/ldap-signing

[E04] Microsoft. SMB Security Hardening in Windows Server and Windows Client.

https://learn.microsoft.com/windows-server/storage/file-server/smb-security-hardening

[E05] Microsoft. Get Started with Windows LAPS and Windows Server Active Directory.

https://learn.microsoft.com/en-us/windows-server/identity/laps/laps-scenarios-windows-server-active-directory

[E06] Microsoft. Secure, Delegate, and Audit Certification Authority Administration.

https://learn.microsoft.com/en-us/windows-server/identity/ad-cs/secure-delegate-audit-certification-authority

[E07] Microsoft. Security Assessment: Certificates (Microsoft Defender for Identity).

https://learn.microsoft.com/en-us/defender-for-identity/security-posture-assessments/certificates

[E08] MITRE. ATT&CK T1558.003: Kerberoasting.

https://attack.mitre.org/techniques/T1558/003/

[E09] MITRE. ATT&CK T1558.004: AS-REP Roasting.

https://attack.mitre.org/techniques/T1558/004/

[E10] MITRE. ATT&CK T1003.006: DCSync.

https://attack.mitre.org/techniques/T1003/006/

[E11] MITRE. ATT&CK T1557.001: Name Resolution Poisoning and SMB Relay.

https://attack.mitre.org/techniques/T1557/001/

[E12] MITRE. ATT&CK T1550.002: Pass the Hash.

https://attack.mitre.org/techniques/T1550/002/

[E13] MITRE. ATT&CK T1484.001: Group Policy Modification.

https://attack.mitre.org/techniques/T1484/001/

[E14] Microsoft. Audit Directory Service Changes.

https://learn.microsoft.com/windows/security/threat-protection/auditing/audit-directory-service-changes

[E15] Microsoft. Audit Kerberos Authentication Service.

https://learn.microsoft.com/windows/security/threat-protection/auditing/audit-kerberos-authentication-service

[E16] Microsoft. How Credential Guard Works.

https://learn.microsoft.com/en-us/windows/security/identity-protection/credential-guard/credential-guard-how-it-works

[E17] Microsoft. Protected Users Security Group.

https://learn.microsoft.com/en-us/windows-server/security/credentials-protection-and-management/protected-users-security-group

[E18] Microsoft. Enterprise Access Model.

https://learn.microsoft.com/security/privileged-access-workstations/privileged-access-access-model

[E19] SpecterOps. BloodHound Community Edition Documentation.

https://github.com/SpecterOps/bloodhound-docs

[E20] NIST. SP 800-115: Technical Guide to Information Security Testing and Assessment.

https://www.nist.gov/publications/technical-guide-information-security-testing-and-assessment

[E21] Microsoft. Group Managed Service Accounts Overview.

https://learn.microsoft.com/en-us/windows-server/identity/ad-ds/manage/group-managed-service-accounts/group-managed-service-accounts/group-managed-service-accounts-overview

[E22] Microsoft. Active Directory Certificate Services Overview.

https://learn.microsoft.com/en-us/windows-server/identity/ad-cs/active-directory-certificate-services-overview

[E23] Microsoft. Windows Security Event 4662: Directory Object Operation.

https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-10/security/threat-protection/auditing/event-4662

[E24] Microsoft. Windows Security Event 5136: Directory Object Modified.

https://learn.microsoft.com/en-us/previous-versions/windows/it-pro/windows-10/security/threat-protection/auditing/event-5136

[E25] MITRE. ATT&CK T1558.001: Golden Ticket.

https://attack.mitre.org/techniques/T1558/001/

[E26] MITRE. ATT&CK T1558.002: Silver Ticket.

https://attack.mitre.org/techniques/T1558/002/

[E27] Microsoft. Privileged Access Overview.

https://learn.microsoft.com/en-us/security/zero-trust/security-concept-privileged-access

[E28] MITRE. ATT&CK T1606.002: Forge SAML Tokens.

https://attack.mitre.org/techniques/T1606/002/

[E29] Fortra. Impacket project README and release information.

https://github.com/fortra/impacket/blob/master/README.md

[E30] Microsoft. LDAP Channel Binding in Active Directory.

https://learn.microsoft.com/en-gb/windows-server/identity/ad-ds/ldap-channel-binding

GO BEYOND THE PREVIEW

Continue past reconnaissance.

The complete edition follows an end-to-end assessment from initial discovery to domain-level risk analysis, with context on how to document evidence, reduce operational impact and translate findings into fixes.

The six-phase workflow

01 **Reconnaissance**
Surface and identity discovery

02 **Initial access**
Authentication and exposure

03 **Privilege escalation**
Permissions and delegation

04 **Lateral movement**
Routes and trust boundaries

05 **Credential access**
Identity-material exposure

06 **Post-exploitation**
Persistence and recovery risks

What makes the expanded edition useful?

- Attack mechanics with prerequisites, interpretation, detection and hardening context.
- Thirty numbered technical references, focused diagrams and fieldwork worksheets.
- Risk-based assessment, reporting, ATT&CK mapping and safe laboratory exercises.

INSIDE THE DOMAIN

A Field Manual for Active Directory Penetration Testing

Find the full edition on its Leanpub book landing page.