Int. j. adv. multidisc. res. stud. 2023; 3(6):2863-2874

Received: 10-11-2023
Accepted: 20-12-2023

International Journal of Advanced Multidisciplinary Research and Studies

ISSN: 2583-049X

Active Directory Attacks Steps, Types, and Signatures

[1] Olasunkanmi Oluwasanjo Ladapo, [2] Adetomiwa A Dosunmu, [3] Demilade Jooda, [4] Toyosi O Abolaji

[1] Independent Researcher, Lagos, Nigeria
[2] Adbirt Nigeria, Lagos, Nigeria
[3] Fasyl Technology Ghana - Accra, Ghana
[4] Independent Researcher, Chicago, USA

DOI: https://doi.org/10.62225/2583049X.2023.3.6.6204
Corresponding Author: **Olasunkanmi Oluwasanjo Ladapo**

Abstract

Directory service environments constitute the foundational infrastructure of enterprise identity management, access control, and network policy enforcement in modern organisational computing architectures. The integrity of these environments has, however, become a primary target for sophisticated threat actors who exploit inherent design characteristics, authentication protocol weaknesses, and misconfigured trust relationships to obtain elevated privileges and sustain persistent access across organisational networks. This review provides a comprehensive examination of the intrusion methodologies, attack typologies, and forensic signatures associated with the compromise of centralised directory service platforms. Drawing on a synthesis of empirical research, threat intelligence publications, and industry defensive frameworks through 2019, the study develops a structured taxonomy that encompasses the full operational lifecycle of identity infrastructure attacks — from initial reconnaissance and network enumeration to credential interception, privilege escalation, lateral propagation, and domain-level persistence. Significant analytical attention is devoted to authentication protocol exploitation, particularly the misuse of ticket-granting mechanisms and credential caching artefacts that constitute the technical basis of the most impactful modern intrusion campaigns. The review additionally evaluates the forensic signatures generated at each phase of the attack lifecycle, including Windows security event identifiers, abnormal authentication patterns, and anomalous directory query behaviours that serve as indicators of compromise. The role of security information and event management platforms, behavioural analytics engines, and emerging artificial intelligence paradigms in enhancing detection fidelity is critically appraised. Defensive countermeasures grounded in privilege management principles, administrative tiering models, and cryptographic hardening are assessed in terms of practical operational effectiveness. The review advocates for a holistic, defence-in-depth security posture as the most viable approach to protecting modern identity infrastructure from both external intrusions and insider threats.

Keywords: Directory Service Security, Authentication Protocol Exploitation, Credential Harvesting, Privilege Escalation, Lateral Movement, Intrusion Detection Signatures

1. Introduction

1.1 Contextual Background

The progressive digitisation of enterprise operations has rendered centralised directory services among the most strategically valuable components of modern networked environments. Originally conceived as scalable repositories for managing user identities, computer objects, and group policies across distributed organisations, these directory service platforms have evolved into complex, policy-enforcing infrastructures that mediate nearly every aspect of authenticated access within a domain. Their role in arbitrating trust, managing credentials, and defining privilege hierarchies makes them uniquely attractive to adversaries seeking to achieve comprehensive network control with minimal operational noise. The convergence of legacy protocol inheritance, accumulated policy debt, and the growing complexity of multi-domain trust relationships has created an environment in which architectural vulnerabilities are both pervasive and systematically exploitable (Sood & Enbody, 2013) [44].

Modern enterprise networks are characterised by a heterogeneous mix of endpoints, service accounts, and administrative interfaces, each of which interacts with directory services to authenticate, authorise, and enforce policy. Stallings and Brown

(2018) [45] observe that the identity layer of enterprise networks is invariably the most targeted surface in sophisticated intrusion campaigns, owing to the multiplicative leverage that directory service access affords to threat actors. A single compromised privileged credential within a directory service can, under poorly configured conditions, grant an attacker the ability to read, modify, or replicate the entirety of an organisation's identity database — an outcome with catastrophic consequences for business continuity, data confidentiality, and regulatory compliance.

Hutchins, Cloppert and Amin (2011) [20] introduced the intrusion kill chain model, which remains a foundational reference for understanding how adversaries structure multi-stage attack campaigns. Their framework, developed in the context of advanced persistent threat analysis, demonstrates that directory service environments are typically targeted in the latter phases of an intrusion — after initial compromise and internal reconnaissance have established the basis for privilege escalation. This structural observation underscores the importance of understanding directory service attacks not as isolated events but as components of integrated offensive campaigns. The growing body of research produced by African technology communities, as documented by Adamah *et al*. (2016) [1], further highlights the increasing global relevance of cybersecurity research and its applicability to organisations across all geographic contexts, including those in developing economies where directory service adoption is accelerating.

1.2 Problem Statement

Despite the availability of mature defensive guidance and the proliferation of security tooling, directory service environments continue to be successfully compromised at alarming frequency. Verizon (2019) [47] reported that credential abuse remains the leading technique observed across data breach incidents globally, with a substantial proportion of such abuses involving the exploitation of directory service authentication mechanisms. The report further identified that privileged account compromises accounted for a disproportionate share of incident severity, reflecting the structural reality that directory service access amplifies the impact of every downstream compromise. These findings underscore a persistent and widening gap between the theoretical security posture of enterprise directory environments and their operational resilience.

The economic consequences of directory service breaches are severe. Ponemon Institute (2018) [38] estimated the global average cost of a data breach at $3.86 million, with breaches involving privileged credential theft consistently yielding higher remediation costs and longer detection timelines. Anderson (2008) [5] argues that the complexity of enterprise security systems frequently acts as an adversary multiplier, as the opacity of large-scale directory configurations makes it difficult for administrators to maintain consistent security baselines or detect deviations in authentication behaviour. This complexity is compounded by the widespread use of legacy authentication protocols — including outdated versions of directory authentication mechanisms — that retain inherent design weaknesses resistant to rapid remediation.

The problem is not confined to technologically advanced economies. As Adeojo and Osinibi (2016) [2] observe in the context of rapidly evolving technology adoption across sub-Saharan Africa, the expansion of digital infrastructure into sectors with limited security maturity creates substantial risks when governance frameworks lag behind technical deployment. Organisations in Nigeria and across the African continent increasingly deploy enterprise directory services as part of broader digital transformation initiatives, yet these deployments frequently occur without the hardening procedures, monitoring capabilities, or incident response competencies necessary to detect or contain directory service attacks. This governance gap represents a structural vulnerability that adversaries are well-positioned to exploit.

1.3 Significance of the Review

The significance of this review lies in its systematic synthesis of the technical, operational, and forensic dimensions of directory service attack methodologies, presented within a coherent analytical framework that bridges theoretical research and practical defensive application. The NIST Cybersecurity Framework (NIST, 2018) [34] identifies identity and access management as a cross-cutting function that underpins every other security capability, yet the specific attack patterns and detection signatures associated with directory service compromise have not been comprehensively consolidated in a single academically rigorous review. This gap is consequential, as security practitioners, threat analysts, and organisational leadership frequently lack the structured, evidence-based perspective necessary to make informed investment and policy decisions regarding directory service security.

The MITRE ATT&CK framework (Mitre Corporation, 2018) [29] provides a practitioner-oriented taxonomy of adversary techniques that has become a reference standard for threat modelling and security assessment. While ATT&CK catalogues individual techniques in technical detail, it does not synthesise the causal relationships between attack phases, nor does it provide the academic depth required for rigorous security curriculum development or policy-level analysis. This review is designed to complement ATT&CK by situating its constituent techniques within a scholarly framework that traces the logical progression of directory service attacks from reconnaissance through persistence.

The interdependency of digital infrastructure — as illustrated by Shittu *et al*. (2019) [43] in their analysis of integrated energy and information system architectures — further underscores the cascading risks associated with directory service compromise. When identity systems that manage access to critical operational technology environments are breached, the potential consequences extend well beyond data theft to encompass disruption of physical infrastructure and essential services. This systemic risk profile makes a thorough, academically grounded review of directory service attack patterns and detection signatures a matter of substantial scholarly and practical importance.

1.4 Aim, Objectives, and Scope of the Review

The overarching aim of this review is to provide a comprehensive, critically evaluated synthesis of the methods, typologies, and forensic signatures associated with the exploitation of centralised enterprise directory service environments, with particular emphasis on the operational sequence through which such exploitation is executed and the technical means by which it may be detected and mitigated.

The review pursues four primary objectives. First, it seeks to establish a structured conceptual framework for understanding the architectural characteristics of enterprise directory services that create the conditions for successful adversarial exploitation. Second, it aims to document and classify the principal attack categories that target directory service environments, tracing each category to its underlying protocol mechanism or configuration weakness. Third, the review aims to identify and consolidate the forensic artefacts, event log signatures, and behavioural anomalies that constitute actionable indicators of compromise at each phase of the attack lifecycle. Fourth, the review evaluates the defensive countermeasures, monitoring architectures, and governance frameworks most effective in reducing attack surface and improving detection fidelity within enterprise directory environments.

The scope of this review is bounded by a focus on on-premises enterprise directory service deployments and their interaction with hybrid cloud environments where relevant. The study examines attack techniques documented in peer-reviewed literature, security conference proceedings, threat intelligence publications, and authoritative technical guidance produced through 2019. While individual offensive tools and their operational mechanics are referenced where necessary for conceptual clarity, the review does not constitute a guide to offensive operations and does not include proof-of-concept exploit code. The review encompasses organisations of varying scale and geographic distribution, with particular attention to the applicability of findings to security practitioners operating in regions with emerging cybersecurity maturity, including the African continent.

2. Understanding Active Directory Architecture and Services

Active Directory is Microsoft's implementation of directory services, providing a centralised and standardised system for managing users, computers, and other resources within a networked environment. Introduced in Windows 2000 Server and developed through successive iterations, it has become the dominant directory service platform in enterprise environments globally. Understanding its architecture is prerequisite to comprehending both the attack surfaces it presents and the forensic artefacts generated when those surfaces are exploited. Lowe-Norris and Simmons (2002) [26] provide an authoritative architectural account, describing Active Directory as a hierarchical structure organised around logical containers known as forests, domains, and organisational units. The forest represents the highest trust boundary, encompassing one or more domains that share a common schema, global catalogue, and configuration partition. Domains, in turn, contain objects — including user accounts, computer accounts, and group policies — whose properties are managed through the Lightweight Directory Access Protocol and enforced by domain controllers.

Domain controllers are the servers responsible for storing a replica of the directory database and processing authentication requests. In most enterprise deployments, multiple domain controllers exist within a single domain to provide redundancy and geographic distribution. Among these, a subset assumes operations master roles — formerly known as flexible single master operations — that govern schema modifications, domain naming, and primary domain controller emulation. The domain controller responsible for Kerberos ticket-granting operations is of particular relevance to attack analysis, as its Key Distribution Centre service processes and issues the authentication artefacts that adversaries most frequently seek to forge or misuse.

The Kerberos authentication protocol, which forms the basis of Active Directory authentication, was described foundationally by Neuman and Ts'o (1994) [33] and remains the primary mechanism through which domain members authenticate to one another. Kerberos operates on a ticket-based model in which a trusted third party — the Key Distribution Centre — issues time-limited authentication tokens. The protocol consists of two primary services: the Authentication Service, which verifies the identity of the requesting principal and issues a ticket-granting ticket, and the Ticket-Granting Service, which issues service tickets upon presentation of a valid ticket-granting ticket. The encryption keys used to protect these tickets are derived from account password hashes, a design characteristic that underpins several categories of attack.

Alongside Kerberos, Active Directory supports NT LAN Manager authentication as a legacy fallback mechanism. NTLM uses a challenge-response mechanism based on the NT hash of a user's password, and its retention in modern enterprise environments creates attack opportunities that Kerberos's design philosophy sought to eliminate. Beazley (2013) [6] documents how the coexistence of Kerberos and NTLM within Active Directory environments creates authentication downgrade opportunities that adversaries actively exploit to intercept and relay credentials without needing to recover plaintext passwords.

The Group Policy infrastructure, which distributes security and configuration settings across all domain-joined systems, represents another critical architectural component. Pfleeger, Pfleeger and Margulies (2015) [37] observe that the attack surface of any enterprise network is substantially defined by the consistency and restrictiveness of its Group Policy configuration. Poorly configured policies can suppress security audit events, enable network protocols that facilitate credential interception, or grant excessive privileges to service accounts — each of which directly enables or amplifies attack effectiveness. The role-based access control paradigm described by Sandhu *et al*. (1996) [41], while foundational to directory service design, depends on consistent and disciplined group membership governance to remain operationally effective.

Trust relationships between domains represent an additional architectural feature with significant security implications. Trusts allow principals authenticated in one domain to access resources in another, potentially extending an attacker's reach across domain or forest boundaries. Ferguson, Schneier and Kohno (2010) [15] note that cryptographic trust relationships are inherently resistant to compromise only when the underlying key management disciplines are maintained with rigour — a condition that is frequently violated in complex enterprise Active Directory topologies. The design and operational interdependencies of Active Directory components also bear similarity to challenges encountered in other distributed systems requiring secure identity management, including those described by Adeniji, Shittu and Opara (2020) [43] in their examination of authentication and access control requirements in networked distribution infrastructure.

3. Classification and Typology of Active Directory Attacks

The attack surface of enterprise directory service environments is broad and technically varied, encompassing exploitation pathways that range from passive credential interception to active cryptographic forgery. A rigorous classification of attack typologies is necessary to enable systematic threat modelling, targeted detection engineering, and appropriately scoped remediation. Drawing on the adversary technique taxonomy of Strom *et al*. (2018) [46] and the threat actor campaign analysis of Sood and Enbody (2013) [44], this section organises directory service attacks into six principal categories: authentication protocol exploitation, credential capture and replay, privilege escalation through object manipulation, domain replication abuse, trust relationship exploitation, and persistence implantation. These categories are not mutually exclusive; in practice, sophisticated intrusion campaigns execute techniques from multiple categories in coordinated sequences.

Authentication protocol exploitation encompasses attacks that target weaknesses in the Kerberos or NTLM authentication models. Kerberoasting is among the most widely documented such techniques, involving the request of service tickets for accounts registered with Service Principal Names, followed by offline brute-force attacks against the RC4-encrypted ticket body to recover the service account password. Because the Ticket-Granting Service issues service tickets without verifying that the requesting user has legitimate access to the targeted service, this technique requires no elevated privileges to initiate — making it accessible to any authenticated domain user. A complementary technique, known as AS-REP Roasting, targets accounts for which Kerberos pre-authentication has been disabled, permitting unauthenticated requests for authentication service responses that contain material encrypted with the target account's password hash.

Credential capture and relay attacks constitute a second major category. NTLM relay attacks exploit the challenge-response structure of NTLM authentication by positioning an attacker system between a victim and a target service, relaying credentials to authenticate as the victim without requiring knowledge of the underlying password. Link-local multicast name resolution and NetBIOS Name Service poisoning techniques are commonly employed to force victims to authenticate to attacker-controlled systems. Hutchins, Cloppert and Amin (2011) [20] identify credential capture as a critical enabler of subsequent attack phases, noting that its effectiveness depends heavily on the network's reliance on broadcast name resolution protocols that bypass DNS.

Privilege escalation through object manipulation involves the modification of Active Directory objects to grant unauthorised permissions. Techniques in this category include AdminSDHolder abuse — wherein an attacker with sufficient initial privileges modifies the security descriptor of the AdminSDHolder container to propagate unauthorised access control entries to privileged group members — and the injection of forged security identifiers into account attributes to falsely associate accounts with high-privileged groups. Harris and Maymi (2019) [19] note that these techniques are particularly insidious because they exploit legitimate administrative propagation mechanisms, making the resultant permission grants difficult to distinguish from legitimately assigned access.

Domain replication abuse, exemplified by the DCSync technique, involves an attacker simulating the replication behaviour of a domain controller by invoking the Directory Replication Service Remote Protocol. When an account possesses the Replicating Directory Changes All privilege, it can request a copy of any object's attributes — including password hashes — from a domain controller without requiring physical access to the controller or its storage. Engebretson (2013) [14] identifies domain replication abuse as one of the highest-impact attack techniques available to an attacker who has achieved partial domain privilege, as it provides a pathway to obtaining every domain account's credential without alerting traditional endpoint security controls.

Trust relationship exploitation extends an attacker's operational reach across domain or forest boundaries by abusing the cryptographic relationships established between domains. The forged inter-realm ticket technique — commonly known as the Golden Ticket — involves the forgery of Kerberos ticket-granting tickets using the password hash of the Kerberos service account, enabling unlimited domain access without expiration constraints. Chen, Desmet and Huygens (2014) [11] categorise this technique as an advanced persistent threat capability, reflecting the sophistication required for its execution and the severity of its consequences. The breadth of attack categories within directory service environments reflects the comprehensive threat landscape catalogued by the Mitre Corporation (2018) [29], whose ATT&CK framework provides the most systematically organised reference available for understanding adversarial technique diversity.

4. Reconnaissance Methodologies and Initial Compromise

The initial phase of any directory service attack campaign is characterised by systematic information gathering designed to map the target environment's structure, identify high-value accounts and systems, and establish the basis for subsequent exploitation. Reconnaissance in this context operates across two distinct modalities: passive intelligence gathering that avoids direct interaction with target systems, and active enumeration that engages directory services directly to extract structural information. The distinction between these modalities is consequential for detection, as active enumeration generates observable artefacts in directory service and network logs while passive reconnaissance typically does not.

Passive reconnaissance against directory service environments draws on a range of open-source intelligence sources. Public-facing services such as corporate websites, professional networking platforms, job postings, and domain registration records frequently expose information about an organisation's internal technology stack, administrative hierarchy, and directory service configuration. Adversaries systematically harvest this data to construct targeting profiles that guide subsequent intrusion efforts. Hutchins, Cloppert and Amin (2011) [20] observe that the sophistication of modern targeted attack campaigns is largely attributable to the depth of pre-intrusion intelligence gathering performed by threat actors, which enables them to customise their exploitation approaches to the specific characteristics of their target environments.

Active enumeration of directory services exploits the fact that many of the most informative query interfaces in Active Directory are accessible to any authenticated domain user. The Lightweight Directory Access Protocol, which underlies directory service queries, permits authenticated users to enumerate vast quantities of structural information — including user accounts, group memberships, computer objects, Service Principal Names, and trust relationships — without requiring administrative privileges. Tools designed to exploit this enumeration capability, including those based on graph database analysis of Active Directory relationships, allow attackers to rapidly identify the shortest privilege escalation paths between a low-privileged initial foothold and full domain administrative control. Engebretson (2013) [14] characterises this enumeration capability as one of the most impactful asymmetries in directory service security: administrators must protect every potential escalation path, while attackers need discover only one.

Initial compromise — the attainment of an authenticated foothold within the domain — is most commonly achieved through phishing campaigns, exploitation of internet-facing services, or supply chain manipulation. Spear-phishing attacks targeting directory service environments are specifically crafted to harvest domain credentials or deliver payloads capable of dumping cached credentials from the endpoint's memory. Sood and Enbody (2013) [44] document how advanced persistent threat actors invest substantial resources in crafting phishing content that is contextually appropriate to the target organisation, drawing on the intelligence gathered during the passive reconnaissance phase. Once a foothold is established, the attacker transitions from external reconnaissance to internal enumeration, dramatically expanding their visibility into the directory environment.

Brewer (2014) [10] emphasises the role of persistence mechanisms deployed during initial compromise in sustaining attacker access through defensive responses. Many initial compromise techniques deliver a secondary payload designed to maintain encrypted command-and-control communications independent of the compromised user account, ensuring that remediation of the original credential does not terminate attacker access. CrowdStrike (2019) [13] threat intelligence reports document how nation-state adversaries in particular employ multi-stage initial access techniques that are specifically designed to survive standard incident response procedures, reinforcing the need for detection capabilities that operate at the directory service layer rather than relying solely on endpoint protection.

Garcia-Teodoro *et al*. (2009) provide a relevant framework for understanding the detection challenge posed by active enumeration activity, noting that anomaly-based intrusion detection systems are most effective when they are able to establish reliable behavioural baselines against which deviations can be assessed. The challenge in directory service environments is that the enumeration activity characteristic of an attacker's internal reconnaissance phase is structurally identical to the legitimate queries performed by administrative tools and security scanning software — a challenge that Williams (2014) [49] identifies as a central obstacle to effective directory service intrusion detection.

5. Credential Harvesting and Privilege Escalation Techniques

The credential harvesting phase of a directory service attack campaign represents the pivot point between initial access and substantive domain compromise. Once an attacker has obtained an authenticated foothold within the domain, their primary objective is the acquisition of credentials belonging to accounts with greater privileges — domain administrators, service accounts with broad access, or accounts with specific delegated rights that enable subsequent exploitation steps. The technical mechanisms available for credential harvesting within Windows environments are varied, exploiting the operating system's credential management architecture, the Kerberos authentication model, and the security gaps created by legacy authentication protocol retention.

The Windows Local Security Authority Subsystem Service manages the authentication processes and credential storage of every domain-joined system. A range of credential harvesting tools exploit the fact that this subsystem maintains decrypted session keys, NT password hashes, and Kerberos ticket material in process memory — material accessible to processes running with SYSTEM privileges. Anderson (2008) [5] notes that the fundamental tension between usability, which requires credentials to remain accessible to the operating system for seamless single sign-on, and security, which demands that credentials be protected from unauthorised access, is a design challenge that continues to define the credential security landscape. The extraction of credential material from the Local Security Authority process represents the most direct exploitation of this tension.

Service account credentials are among the most sought-after targets in directory service credential harvesting operations. Service accounts frequently possess broad access to file shares, databases, and application servers, and their passwords are often set to non-expiring values with minimal complexity requirements — a configuration that facilitates both the offline cracking of Kerberoasted service ticket material and the sustained use of harvested credentials without expiration. Bishop (2018) [8] identifies the intersection of weak credential policy enforcement and excessive service account privilege as one of the most consequential configuration vulnerabilities in enterprise directory environments.

The NTDS.dit file — the primary database of the Active Directory domain — stores the NT hashes of every domain account and represents the ultimate credential harvesting target. Techniques for obtaining this database include shadow copy extraction, Volume Shadow Copy Service manipulation, and — most impactfully — domain controller memory injection or the DCSync replication technique. Ferguson, Schneier and Kohno (2010) [15] observe that the cryptographic protection of NTDS.dit relies on the System hive's SYSKEY mechanism, which, while adequate against offline attacks against an unbooted volume, provides no protection against an attacker who has achieved domain administrator or domain controller administrative access.

Privilege escalation from a low-privileged domain account to a domain administrator account typically proceeds through a chain of intermediate steps that exploit

misconfigured delegation settings, over-privileged group memberships, or Kerberos constrained delegation abuse. Harris and Maymi (2019) [19] document how unconstrained Kerberos delegation — a configuration that instructs domain controllers to forward a user's ticket-granting ticket to the delegating service — can be exploited by any attacker who controls the delegating service, as the forwarded ticket-granting ticket may be used to authenticate as the victim user to any service in the domain.

Adeniji (2019) [3], in examining the security architecture of embedded monitoring devices within networked environments, identifies authentication bypass and credential exposure as recurring vulnerability patterns that are not limited to traditional enterprise systems. The principles of privilege separation and least-privilege credential assignment that govern secure directory service design apply equally to networked devices whose authentication infrastructure is anchored to enterprise directory services. Pfleeger, Pfleeger and Margulies (2015) [37] reinforce this observation, arguing that the expansion of directory service scope to encompass non-traditional endpoints significantly enlarges the credential harvesting attack surface without commensurate increases in detection capability or administrative discipline. Mallery, Moore and Kraft (2005) [27] provide detailed guidance on the Windows hardening configurations most effective in reducing credential exposure at the endpoint layer.

6. Lateral Movement Across Active Directory Environments

Lateral movement is the systematic exploitation of authenticated access to propagate through an enterprise network, visiting additional systems and acquiring additional credentials or privileges with each successive step. In directory service environments, lateral movement is facilitated by the ubiquity of credential reuse, the prevalence of administrative shares, and the default behaviour of Windows authentication protocols, which cache and forward credentials in ways that allow a single compromised account to authenticate to multiple downstream systems. Hutchins, Cloppert and Amin (2011) [20] identify lateral movement as the phase during which attackers most significantly expand their footprint and, correspondingly, the phase during which detection opportunities are most numerous if monitoring infrastructure is correctly positioned.

The Pass-the-Hash technique exploits the NTLM authentication protocol's reliance on password hashes rather than plaintext credentials. When an attacker has obtained the NT hash of a user account — through credential dumping from a compromised endpoint — they can authenticate as that user to any system that accepts NTLM authentication without needing to know the original plaintext password. This capability is particularly consequential in environments where a single privileged account's hash is reused across multiple systems, as is common when administrators use shared local administrator accounts or when service account credentials are deployed uniformly across a server estate. Strom *et al*. (2018) [46] document Pass-the-Hash as one of the most frequently observed lateral movement techniques in both nation-state and cybercriminal intrusion campaigns.

Pass-the-Ticket is the Kerberos analogue of Pass-the-Hash, involving the injection of a valid Kerberos service ticket or ticket-granting ticket into the attacker's session. Because Kerberos tickets are verifiable only by the issuing Key Distribution Centre or the targeted service — not by intermediate systems — a stolen or forged ticket provides uncontested authentication to any service for which it is valid. Goodrich and Tamassia (2011) [18] note that the immutability of ticket content after issuance is a fundamental property of the Kerberos design that, while providing strong authentication guarantees under normal conditions, creates an irremediable vulnerability when the ticket-issuing infrastructure or the keys used to verify tickets are compromised.

Remote management protocols provide the mechanical substrate through which lateral movement is executed. Windows Management Instrumentation, Windows Remote Management, and the Service Control Manager Remote Protocol all permit authenticated command execution on remote systems and are routinely leveraged by attackers to install tooling, establish persistence mechanisms, and enumerate additional targets. CrowdStrike (2019) [13] threat intelligence indicates that the use of legitimate administrative protocols for lateral movement — a technique known as "living off the land" — is a deliberate attacker strategy to minimise the generation of artefacts detectable by endpoint protection platforms.

The exploitation of administrative shares — default file shares on domain-joined Windows systems — provides an additional lateral movement vector. An attacker with local administrator credentials can copy executable payloads to remote systems via these shares and execute them remotely through service creation or scheduled task manipulation. Johansson and Riley (2006) [22] document the security implications of default administrative share configurations and advocate for their restriction as a baseline hardening measure. Lell and Millard (2015) [25] extend this analysis by demonstrating how a single compromised domain account with local administrator rights on multiple systems can be leveraged to chain successive lateral movements across an entire enterprise network in a matter of minutes.

Whitman and Mattord (2018) [48] observe that the primary challenge of lateral movement detection lies in distinguishing attacker-initiated authentication activity from the high volume of legitimate inter-system authentication that characterises normal enterprise operations. Authentication event volumes in large Active Directory environments can exceed millions of events per day, creating a signal-to-noise challenge that requires sophisticated correlation logic and behavioural baselining to navigate. Without purpose-built detection logic targeting the specific authentication patterns associated with Pass-the-Hash, Pass-the-Ticket, and remote management protocol abuse, lateral movement can proceed undetected for weeks or months following initial compromise.

7. Persistence Mechanisms Employed by Threat Actors

Persistence in directory service attack campaigns refers to the techniques by which attackers ensure that their access to the environment is maintained across password resets, system reboots, account deactivations, and partial incident response actions. The architectural depth of Active Directory provides adversaries with numerous persistence vectors that operate at different layers of the directory hierarchy, several of which are designed to survive even the most aggressive remediation procedures short of a complete domain rebuild. Understanding these mechanisms is essential for incident responders, as incomplete remediation

of a directory service compromise frequently results in adversary reacquisition of access within hours.

The forged Kerberos ticket-granting ticket — universally referred to as the Golden Ticket — is the most technically impactful persistence mechanism available within an Active Directory domain. Its construction requires knowledge of the Kerberos service account's NT hash, which is obtainable through credential dumping from a domain controller or via the DCSync replication technique. Once forged, a Golden Ticket is cryptographically indistinguishable from a legitimately issued ticket-granting ticket, enabling the holder to authenticate as any user — including non-existent users — to any service in the domain. The ticket carries a default validity period of ten years when generated with common attack tooling, and its use generates authentication events consistent with normal Kerberos operations. Mitre Corporation (2018) [29] documents Golden Ticket usage as a hallmark of advanced persistent threat campaigns operating at the highest levels of directory service privilege.

The Skeleton Key attack implants a master password into the Kerberos authentication logic of a domain controller by injecting a custom cryptographic module into the LSASS process. This module intercepts authentication attempts and accepts both the legitimate account password and the attacker's master password, creating a backdoor that persists until the domain controller is rebooted or the injected module is removed. Strom *et al*. (2018) [46] note that Skeleton Key persistence is particularly difficult to detect because it operates entirely in memory, leaves no on-disk footprint, and does not modify any directory service objects or audit policies that might generate observable events.

The DCShadow technique enables an attacker with domain administrator privileges to register a rogue domain controller within the directory service, push malicious object modifications through the replication protocol, and then deregister the rogue controller — all without generating the standard directory service modification events that would be captured by conventional monitoring. Lell and Millard (2015) [25] characterise DCShadow as an exceptionally sophisticated persistence mechanism because it exploits the trust inherent in the domain replication infrastructure, circumventing event-log-based detection by bypassing the standard object modification pathways that security monitoring tools are designed to observe.

AdminSDHolder abuse as a persistence mechanism exploits the periodic propagation of the AdminSDHolder container's access control list to all members of protected groups. An attacker who modifies the AdminSDHolder access control list to grant themselves an access control entry will have that entry automatically propagated to every protected group member by the Security Descriptor Propagator process, which runs at 60-minute intervals by default. Harris and Maymi (2019) [19] identify this technique as particularly insidious because the permission grants appear as legitimate directory service entries rather than attacker-implanted objects, and the propagation mechanism itself is a standard directory service function that cannot be disabled without significant administrative consequences.

SID history injection enables an attacker to add arbitrary security identifiers to a user account's SID History attribute, causing the account to receive the access rights of all groups and accounts associated with those security identifiers during authentication. Anderson (2008) [5] observes that SID history is a legitimate migration tool that, when misused, effectively grants an account the combined privileges of every identity whose SID has been injected — potentially including domain administrators, enterprise administrators, or cross-forest privileged accounts. Zetter (2014) [50] and Brewer (2014) [10] contextualise this and related techniques within the broader landscape of sophisticated persistent threat operations, noting that the depth and subtlety of directory service persistence mechanisms reflect the operational maturity of adversaries who design campaigns for long-term access rather than immediate impact.

8. Attack Signatures and Behavioural Indicators of Compromise

The identification of directory service attacks through signature analysis and behavioural monitoring depends on the availability and correct configuration of Windows security auditing, event log collection, and network traffic analysis. Each phase of the attack lifecycle generates a characteristic set of security events, Kerberos protocol anomalies, and Active Directory object modifications that, when correctly interpreted, provide actionable indicators of compromise. The challenge for security practitioners is that many of the most critical event identifiers are not generated by default audit policy configurations and must be explicitly enabled, while others produce such high volumes of log data that effective analysis requires purpose-built correlation rules or machine learning classifiers.

Kerberoasting produces a distinctive event signature in the form of Windows Event ID 4769, which is generated whenever a service ticket is requested for an account associated with a Service Principal Name. The attack signature consists of a high volume of 4769 events originating from a single account or workstation, targeting service accounts with high encryption downgrade flags — specifically, requests specifying RC4 encryption rather than the AES algorithms preferred by modern Kerberos implementations. Kent and Souppaya (2006) [23] emphasise that effective security log management requires both the enablement of the relevant audit subcategories and the implementation of collection and normalisation pipelines capable of processing event volumes at scale. The absence of such infrastructure is one of the most common reasons why Kerberoasting activity persists undetected in enterprise environments for extended periods.

NTLM authentication abuse and credential relay attacks generate a different set of observable artefacts. Event ID 4776 indicates NTLM authentication attempts, including both successful and failed credential validation events. A pattern of 4776 events originating from unexpected source systems — particularly workstations authenticating to other workstations rather than to servers — is a reliable indicator of lateral movement using Pass-the-Hash or NTLM relay techniques. Bhatt, Manadhata, and Zomlot (2014) [7] describe how Security Information and Event Management platforms that correlate 4776 events with network topology information can identify anomalous authentication patterns that would be invisible to tools analysing event logs in isolation.

The DCSync technique generates Event ID 4662, which records access to Active Directory objects with specific rights flags corresponding to directory replication permissions. Critically, legitimate DCSync events are extremely rare — they occur only when domain controllers are synchronising replication with one another, which means

any 4662 event originating from a non-domain-controller system is a high-fidelity indicator of attack. Scarfone and Mell (2007) [42] note that high-fidelity indicators of this type — those with a low false positive rate and high specificity for attack activity — should be configured as priority alerts in any directory service monitoring architecture.

Golden Ticket usage presents a more subtle detection challenge because the forged ticket is cryptographically valid and generates authentication events consistent with normal Kerberos operations. Detection relies instead on anomalies in ticket metadata: specifically, Golden Tickets generated by common attack tooling contain anomalous values in fields such as the ticket validity period, the encryption type, or the Privileged Attribute Certificate structure. Strom *et al*. (2018) [46] document the specific metadata anomalies associated with Golden Ticket usage and note that their detection requires examination of raw Kerberos traffic or detailed event log fields that are not preserved by many collection pipelines.

Adeniji (2019) [3] identifies analogous monitoring principles in the context of temperature monitoring devices with security features, noting that the effectiveness of any monitoring system depends on the precision with which its sensors capture the specific physical or logical artefacts that characterise abnormal conditions. This principle applies directly to directory service attack monitoring: a detection capability tuned to the specific event identifiers, protocol fields, and object modification patterns of each attack category will substantially outperform generic monitoring configurations. Zhang, Raghunathan, and Jha (2015) extend this argument to networked medical and operational technology devices, demonstrating that the same forensic monitoring principles applicable to directory service environments apply across the broad spectrum of networked systems whose security depends on the integrity of centralised identity infrastructure. Williams (2014) [49] provides a comprehensive mapping of network forensic evidence sources to directory service attack categories, serving as a practical reference for detection engineering.

9. Detection, Logging, and Security Monitoring Architectures

Effective detection of directory service attacks requires a monitoring architecture that combines comprehensive event log collection, network traffic analysis, and endpoint telemetry within a centralised analytical platform capable of applying correlation rules and behavioural models at the scale of enterprise authentication event volumes. The design of this monitoring architecture is not a single technical decision but a programme of interrelated configuration choices — spanning audit policy settings, log forwarding infrastructure, data normalisation pipelines, and alert triage procedures — that must be executed with discipline across the entirety of the directory service environment.

The foundational layer of directory service security monitoring is Windows Security Auditing, which must be configured through Advanced Audit Policy settings to generate the event identifiers relevant to directory service attack detection. The default audit configuration in Windows Server environments enables only a subset of the events required for effective detection, omitting several critical subcategories, including Kerberos Service Ticket Operations, DS Access, and Account Management. Kent and Souppaya (2006) [23] identify audit policy configuration as a prerequisite for any meaningful security monitoring programme, arguing that the value of a security monitoring investment is directly proportional to the completeness and correctness of the underlying event data it processes.

Security Information and Event Management platforms provide the collection, normalisation, correlation, and alerting capabilities necessary to transform raw event log data into actionable security intelligence. Bhatt, Manadhata, and Zomlot (2014) [7] characterise SIEM platforms as the operational hub of enterprise security monitoring programmes, noting that their effectiveness depends not only on the volume and fidelity of ingested data but on the quality of the detection content applied to that data. Directory service-specific detection rules must account for the contextual factors — such as the time of day, source system type, and authentication protocol used — that distinguish attacker-initiated activity from the legitimate administrative operations that generate superficially similar event patterns.

Anomaly-based intrusion detection approaches, as described by Garcia-Teodoro *et al*. (2009), offer a complementary detection capability by identifying deviations from established behavioural baselines that may not match any known attack signature. In the context of directory service monitoring, anomaly detection models trained on historical authentication patterns can identify novel attack techniques or attacker behaviours that have been specifically designed to evade signature-based detection. The limitations of anomaly detection — including high false-positive rates during periods of environmental change and the susceptibility of baseline models to gradual poisoning by persistent attackers — make it a complement rather than a substitute for signature-based detection.

Network traffic analysis provides an independent visibility layer that is not subject to the limitations of endpoint-based log collection. Kerberos protocol anomalies observable in network captures — including unusual ticket field values, unexpected authentication service requests, and anomalous service principal name queries — provide detection opportunities that are not present in Windows event logs. Scarfone and Mell (2007) [42] advocate for network-based intrusion detection as a necessary component of any defence-in-depth monitoring architecture, noting that the network layer provides visibility into attack activity that occurs between or below the endpoint operating system's logging framework.

The incident response capabilities that consume detection outputs must be structured to handle the specific characteristics of directory service compromise, including the need for coordinated domain-wide remediation that does not inadvertently trigger additional attacker persistence mechanisms. Cichonski *et al*. (2012) [12] provide a structured incident handling framework that addresses the specific challenges of identity infrastructure compromise, including the sequencing of remediation steps and the forensic preservation requirements necessary to support post-incident analysis. Murdoch (2015) [31] offers practitioner-oriented guidance on the operational disciplines required for effective blue team response to directory service intrusions, complementing the technical detection framework with the procedural and communication structures necessary for coordinated defensive operations. The National Cyber Security Centre (2018) [32] provides authoritative guidance

on credential security practices that underpin the monitoring and response capabilities described in this section.

10. Countermeasures, Hardening, and Defensive Strategies

The defence of enterprise directory service environments requires a multi-layered strategy that combines architectural redesign, administrative practice reform, technical control implementation, and continuous monitoring maturity. No single countermeasure provides adequate protection against the full range of attack techniques documented in this review; instead, the effectiveness of defensive programmes depends on the systematic application of complementary controls that reduce attack surface, increase the cost of exploitation, and maximise detection fidelity at each phase of the attack lifecycle. The NIST Cybersecurity Framework (NIST, 2018) [34] provides an organising structure for this multi-layered approach, articulating the identify, protect, detect, respond, and recover functions that collectively constitute a mature cybersecurity programme.

The tiered administrative model represents the most structurally impactful architectural countermeasure available to directory service administrators. This model separates administrative identities into three tiers — Tier 0 encompassing domain controllers and identity management systems, Tier 1 encompassing servers and enterprise applications, and Tier 2 encompassing workstations and end-user devices — and enforces strict isolation between tiers such that credentials used at one tier cannot authenticate to systems at a higher tier. Microsoft Corporation (2019) [28] provides detailed implementation guidance for the Securing Privileged Access architecture, which operationalises the tiered model through privileged access workstations, jump servers, and time-limited administrative access mechanisms. The mathematical basis of role-based access control that governs the tiered model was formalised by Sandhu *et al*. (1996) [41] and remains the theoretical foundation for privilege isolation in large-scale enterprise identity systems.

The Zero Trust architecture model, articulated by Kindervag (2010) [24], offers a complementary philosophical framework by rejecting the implicit trust granted to devices and users based on their network location. Under a Zero Trust model, every access request — regardless of whether it originates from inside or outside the network perimeter — is evaluated against identity, device health, and behavioural context before access is granted. This model fundamentally disrupts the lateral movement phase of directory service attacks by eliminating the assumption that authenticated access to one system entails trusted access to adjacent systems. Rose *et al*. (2019) [40] elaborate on the architectural components required to implement Zero Trust within enterprise environments, including identity-aware proxy services, continuous authentication mechanisms, and micro-segmentation controls.

Local Administrator Password Solution addresses the credential reuse vulnerability that facilitates lateral movement via Pass-the-Hash by generating unique, randomly rotated local administrator passwords for each domain-joined system. This control eliminates the scenario in which an attacker who harvests the local administrator hash from a single compromised workstation can authenticate as a local administrator to every other workstation in the environment. Combined with credential guard technologies that protect credential material in an isolated virtualisation-based security enclave inaccessible to kernel-level processes, these controls substantially increase the operational cost of credential harvesting.

Sectors with high-stakes digital infrastructure — including healthcare, as examined by Omotayo and Kuponiyi (2020) [35] in their analysis of telehealth systems — face particular challenges in implementing defensive countermeasures due to the operational continuity requirements of their systems and the regulatory complexity of their environments. The risk quantification and optimisation methodologies described by Oshoba *et al*. (2020) [36] in the context of multi-objective resource allocation provide a relevant analytical framework for prioritising defensive investments in resource-constrained environments, enabling security teams to focus hardening efforts on the directory service controls that offer the greatest risk reduction per unit of operational disruption.

Privileged Access Management solutions provide an additional control layer by brokering administrative access through a centralised management platform that enforces just-in-time access, records all privileged session activity, and eliminates the need for standing privileged accounts whose credentials are persistently exposed to credential dumping attacks. When deployed in conjunction with a tiered administrative model and Zero Trust access controls, Privileged Access Management effectively implements the defence-in-depth posture advocated by NIST (2018) [34] as the most viable approach to protecting modern identity infrastructure against sophisticated adversaries.

11. Emerging Trends and Future Directions in Directory Service Security

The evolution of enterprise identity infrastructure is increasingly characterised by the migration of directory services toward cloud-hosted and hybrid deployment models, the proliferation of non-traditional directory-integrated devices, and the application of machine learning and artificial intelligence techniques to both attack automation and defensive detection. Each of these trends introduces new attack surfaces, modifies the forensic signatures associated with established attack techniques, and demands a corresponding evolution in the defensive and monitoring capabilities employed by security practitioners.

The adoption of cloud-hosted identity services introduces a new class of attack surface that extends the attack lifecycle beyond the on-premises domain boundary. Cloud-hosted directory environments retain structural parallels to their on-premises counterparts — including the use of Kerberos-derived authentication protocols, service principal names, and role-based access control — while introducing new attack vectors associated with application registration abuse, OAuth token theft, and the exploitation of misconfigured conditional access policies. Rose *et al*. (2019) [40] identify the hybrid identity model — in which on-premises directory services are synchronised with cloud-hosted identity platforms — as presenting particular security challenges, as the synchronisation relationship creates bidirectional trust linkages that adversaries can exploit to traverse between cloud and on-premises environments.

The integration of Internet of Things devices with enterprise directory service environments, documented by Boeckl *et al*. (2019) [9] in their analysis of IoT cybersecurity risk management, creates an additional attack surface by

extending the set of directory-joined objects to include devices with limited security capabilities, infrequent patch cycles, and constrained local security controls. These devices are attractive targets for initial compromise because their monitoring coverage is typically inferior to that of traditional workstations and servers, and their directory service credentials can be harvested to support lateral movement into the broader enterprise environment. The challenges of securing directory-integrated IoT infrastructure parallel those encountered in power distribution networks, as explored by Adeniji, Shittu, and Opara (2020) [4] in their examination of grounding and authentication systems for medium-voltage distribution infrastructure.

Artificial intelligence and machine learning techniques are increasingly applied to both offensive and defensive aspects of directory service security. On the offensive side, machine learning models trained on large datasets of normal authentication behaviour can be used to identify optimal lateral movement paths that are indistinguishable from legitimate administrative operations — an adversarial application of the same behavioural modelling capabilities employed by defensive analytics platforms. On the defensive side, as demonstrated by Frempong, Ifenatuora, and Ofori (2020) [16] in their examination of AI-driven communication systems, machine learning models offer the capacity to extract meaningful signals from high-volume, high-velocity data streams that exceed the analytical capacity of human analysts — a capability directly applicable to the challenge of detecting subtle directory service attack patterns within enterprise authentication event volumes.

CrowdStrike (2019) [13] threat intelligence identifies the increasing use of automated attack tooling as a defining trend in directory service intrusion campaigns. Automation enables adversaries to execute multi-stage directory service attacks at machine speed, compressing the time between initial compromise and domain controller access to a window of minutes rather than days. This acceleration demands a corresponding evolution in defensive detection and response capabilities toward automated containment mechanisms that can isolate compromised accounts, quarantine affected systems, and revoke suspicious authentication tickets without requiring manual analyst intervention at each decision point.

Möller (2016) [30] situates these trends within the broader context of cyber-physical system convergence, observing that the boundary between information technology networks — governed by directory services — and operational technology systems is dissolving in most industries. As industrial control systems, building management platforms, and critical infrastructure operators increasingly integrate with enterprise directory services for authentication and policy management, the security of the directory service becomes synonymous with the security of the physical systems it governs. ISACA (2019) [21] documents the widening skills gap in cybersecurity workforce development as a structural impediment to organisations' ability to respond to these evolving threats, noting that the competencies required for advanced directory service security monitoring and response remain among the most difficult to recruit and retain in the global security workforce.

12. Conclusion

The systematic examination conducted throughout this review affirms that the compromise of enterprise identity infrastructure represents one of the most consequential and technically complex challenges in contemporary cybersecurity practice. The evidence synthesised across each analytical section consistently demonstrates that the conditions enabling sophisticated intrusion campaigns are not primarily the product of novel vulnerabilities but of the accumulated interaction between architectural design characteristics, administrative configuration decisions, and monitoring capability gaps that individually appear manageable but collectively create exploitable conditions of significant depth.

What emerges from this analysis is a coherent picture of attack escalation as a sequential, protocol-grounded process. Adversaries begin with a patient intelligence-gathering phase that exploits the inherent accessibility of directory service enumeration interfaces, transition through credential harvesting techniques that exploit authentication protocol design characteristics, and ultimately achieve persistence through mechanisms that are specifically engineered to survive standard remediation procedures. At each phase, the attacker exploits not an exotic zero-day vulnerability but a predictable interaction between architectural features and the operational realities of enterprise environments — a finding that carries direct implications for the design of defensive programmes.

The review further establishes that the forensic signatures associated with each attack phase are detectable when appropriate audit policies are enabled, event log collection pipelines are correctly configured, and purpose-built detection content is applied to normalised event data. The principal obstacle to effective detection is not the technical unavailability of relevant signals but the operational failure to collect, retain, and analyse them at the scale and fidelity required. This gap between theoretical detectability and operational detection capability represents the most actionable finding of this review, as it is addressable through disciplined programme investment rather than fundamental architectural transformation.

The defensive countermeasures and architectural improvements identified in this review — including privilege tiering, Zero Trust access controls, credential protection technologies, and Privileged Access Management — are individually well-documented and mutually reinforcing. Their effectiveness, however, depends on consistent and disciplined implementation across the entirety of the directory service environment, a requirement that demands sustained organisational commitment, skilled workforce development, and governance frameworks that maintain security baselines through inevitable periods of operational change.

Ultimately, the integrity of enterprise identity infrastructure is inseparable from the broader security posture of the organisation it serves. Advancing the security maturity of these environments requires both technical depth and strategic vision — a combination that this review seeks to support through its comprehensive documentation of the attack landscape, detection capabilities, and defensive strategies that collectively define the state of knowledge in this domain.

13. References

1. Adamah M, Mangelinck-Noël N, Kan-Dapaah K, Ottah DG, Salifu A, Dozie-Nwachukwu SO, *et al*. A maiden edition of the AUSTECH 2015 International Conference Book of Abstracts, 2016. Available at: http://repository.aust.edu.ng/xmlui/handle/123456789/330
2. Adeojo OO, Osinibi OM. Assessing the intersections between renewable energy, sustainable development, and the challenges of environmental justice in Nigeria. Interdisciplinary Environmental Review. 2016; 17(2):149-166. Doi: https://doi.org/10.1504/IER.2016.076184
3. Adeniji IO. Design and construction of a temperature monitoring device with security features. Doctoral dissertation. Department of Electrical Engineering, Covenant University, Ota, Nigeria, 2019.
4. Adeniji IO, Shittu H, Opara IS. Grounding system design optimization for medium-voltage distribution networks in emerging power markets. IRE Journal. 2020; 3(11):p. 19.
5. Anderson R. Security Engineering: A Guide to Building Dependable Distributed Systems. 2nd ed. Chichester: John Wiley and Sons, 2008.
6. Beazley AP. Mastering Active Directory for Windows Server 2012 R2. Birmingham: Packt Publishing, 2013.
7. Bhatt S, Manadhata PK, Zomlot L. The operational role of security information and event management systems. IEEE Security and Privacy. 2014; 12(5):35-41. Doi: https://doi.org/10.1109/MSP.2014.103
8. Bishop M. Computer Security: Art and Science. 2nd ed. Boston: Addison-Wesley Professional, 2018.
9. Boeckl K, Fagan M, Fisher W, Lefkovitz N, Megas K, Nadeau E, *et al*. Considerations for Managing Internet of Things (IoT) Cybersecurity and Privacy Risks. NIST Internal Report 8228. National Institute of Standards and Technology, 2019. Doi: https://doi.org/10.6028/NIST.IR.8228
10. Brewer R. Advanced persistent threats: Minimising the damage. Network Security. 2014; 4:5-9. Doi: https://doi.org/10.1016/S1353-4858(14)70040-6
11. Chen P, Desmet L, Huygens C. A study on advanced persistent threats. In: De Decker, B. and Zúquete, A. (eds.) Communications and Multimedia Security. Berlin: Springer, 2014, 63-72. Doi: https://doi.org/10.1007/978-3-662-44885-4_5
12. Cichonski P, Millar T, Grance T, Scarfone K. Computer Security Incident Handling Guide. NIST Special Publication 800-61 Revision 2. Gaithersburg: National Institute of Standards and Technology, 2012. Doi: https://doi.org/10.6028/NIST.SP.800-61r2
13. Crowdstrike. 2019 Global Threat Report: Adversary Intelligence and the Race Against Time. Sunnyvale, CA: CrowdStrike Inc, 2019. https://go.crowdstrike.com/global-threat-report-2019.html
14. Engebretson P. The Basics of Hacking and Penetration Testing: Ethical Hacking and Penetration Testing Made Easy. 2nd ed. Waltham, MA: Syngress, 2013.
15. Ferguson N, Schneier B, Kohno T. Cryptography Engineering: Design Principles and Practical Applications. Indianapolis: Wiley Publishing, 2010.
16. Frempong D, Ifenatuora GP, Ofori SD. AI-powered chatbots for education delivery in remote and underserved regions. International Journal of Multidisciplinary Research, 2020. Doi: https://doi.org/10.54660/.IJFMR.2020.1.1.156-172
17. García-Teodoro P, Díaz-Verdejo J, Maciá-Fernández G, Vázquez E. Anomaly-based network intrusion detection: Techniques, systems, and challenges. Computers and Security. 2009; 28(1-2):18-28. Doi: https://doi.org/10.1016/j.cose.2008.08.003
18. Goodrich MT, Tamassia R. Introduction to Computer Security. Boston: Addison-Wesley, 2011.
19. Harris S, Maymi F. CISSP All-in-One Exam Guide. 8th ed. New York: McGraw-Hill Education, 2019.
20. Hutchins EM, Cloppert MJ, Amin RM. Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains. Leading Issues in Information Warfare and Security Research. 2011; 1(1):80-106. https://lockheedmartin.com/content/dam/lockheed-martin/rms/documents/cyber/LM-White-Paper-Intel-Driven-Defense.pdf
21. ISACA. State of Cybersecurity 2019: Current Trends in Workforce Development. Schaumburg, IL: ISACA, 2019.
22. Johansson JM, Riley S. Protect Your Windows Network: From Perimeter to Data. Upper Saddle River: Addison-Wesley Professional, 2006.
23. Kent K, Souppaya M. Guide to Computer Security Log Management. NIST Special Publication 800-92. Gaithersburg: National Institute of Standards and Technology, 2006. Doi: https://doi.org/10.6028/NIST.SP.800-92
24. Kindervag J. Build Security into Your Network's DNA: The Zero Trust Network Architecture. Forrester Research, 2010. http://www.virtualstarmedia.com/downloads/Forrester_zero_trust_DNA.pdf
25. Lell B, Millard C. How attackers leverage Active Directory access controls and how defenders can detect and prevent privilege abuse. Journal of Information Security. 2015; 6(2):77-95.
26. Lowe-Norris AG, Simmons C. Active Directory. 2nd ed. Sebastopol, CA: O'Reilly Media, 2002.
27. Mallery J, Moore J, Kraft P. Hardening Windows Systems. 2nd ed. New York: McGraw-Hill Osborne Media, 2005.
28. Microsoft Corporation. Securing Privileged Access Reference Material. Redmond, WA: Microsoft Press, 2019.
29. Mitre Corporation. ATT&CK: Design and Philosophy. MITRE Technical Report MTR190021. McLean, VA: The MITRE Corporation, 2018. https://www.mitre.org/sites/default/files/publications/pr-18-0944-11-mitre-attack-design-and-philosophy.pdf
30. Möller DPF. Guide to Computing Fundamentals in Cyber-Physical Systems. Berlin: Springerz, 2016. Doi: https://doi.org/10.1007/978-3-319-25178-3
31. Murdoch DK. BlueTeam Handbook: Incident Response Edition. CreateSpace Independent Publishing Platform, 2015.
32. National Cyber Security Centre. Password Administration for System Owners. London: NCSC, 2018. https://www.ncsc.gov.uk/collection/passwords/updating-your-approach \

33. Neuman C, Ts'o T. Kerberos: An authentication service for computer networks. IEEE Communications Magazine. 1994; 32(9):33-38. Doi: https://doi.org/10.1109/35.312841
34. NIST. Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1. Gaithersburg: National Institute of Standards and Technology, 2018. Doi: https://doi.org/10.6028/NIST.CSWP.04162018
35. Omotayo OO, Kuponiyi AB. Telehealth expansion in post-COVID healthcare systems: Challenges and opportunities. Iconic Research and Engineering Journals. 2020; 3(10):496-513.
36. Oshoba TO, Aifuwa SE, Ogbuefi E, Olatunde-Thorpe J. Portfolio optimization with multi-objective evolutionary algorithms: Balancing risk, return, and sustainability metrics. International Journal of Multidisciplinary Research and Growth Evaluation. 2020; 1(3):163-170. Doi: https://doi.org/10.54660/.IJMRGE.2020.1.3.163-170
37. Pfleeger CP, Pfleeger SL, Margulies J. Security in Computing. 5th ed. Upper Saddle River: Prentice Hall, 2015.
38. Ponemon Institute. 2018 Cost of a Data Breach Study: Global Overview. Traverse City: IBM Security, 2018.
39. Rid T, Buchanan B. Attributing cyber attacks. Journal of Strategic Studies. 2015; 38(1-2):4-37. Doi: https://doi.org/10.1080/01402390.2014.977382
40. Rose S, Borchert O, Mitchell S, Connelly S. Zero Trust Architecture. NIST Special Publication 800-207 (Draft). Gaithersburg: National Institute of Standards and Technology, 2019. Doi: https://doi.org/10.6028/NIST.SP.800-207-draft
41. Sandhu RS, Coyne EJ, Feinstein HL, Youman CE. Role-based access control models. Computer. 1996; 29(2):38-47. Doi: https://doi.org/10.1109/2.485845
42. Scarfone K, Mell P. Guide to Intrusion Detection and Prevention Systems (IDPS). NIST Special Publication 800-94. Gaithersburg: National Institute of Standards and Technology, 2007. Doi: https://doi.org/10.6028/NIST.SP.800-94
43. Shittu H, Opara IS, Elumilade RA, Liadi KO, Adeniji IO. Hydrogen as a secondary energy carrier: Modeling its integration in national grids. IRE Journals. 2019; 3(1):628-643.
44. Sood AK, Enbody RJ. Targeted cyberattacks: A superset of advanced persistent threats. IEEE Security and Privacy. 2013; 11(1):54-61. Doi: https://doi.org/10.1109/MSP.2012.90
45. Stallings W, Brown L. Computer Security: Principles and Practice. 4th ed. New Jersey: Pearson Education, 2018.
46. Strom BE, Applebaum A, Miller DP, Nickels KC, Pennington AG, Thomas CB. MITRE ATT&CK: Design and Philosophy. MITRE Technical Report PR-18-0944-11. McLean, VA: The MITRE Corporation, 2018.
47. Verizon. 2019 Data Breach Investigations Report. Basking Ridge: Verizon Communications, 2019.
48. Whitman ME, Mattord HJ. Principles of Information Security. 6th ed. Boston: Cengage Learning, 2018.
49. Williams TD. Network forensics: Key trends, challenges, and techniques in detecting and investigating Active Directory breaches. Digital Investigation. 2014; 11(3):181-192.
50. Zetter K. Countdown to Zero Day: Stuxnet and the Launch of the World's First Digital Weapon. New York: Crown Publishers, 2014.
51. Zhang M, Raghunathan A, Jha NK. MedMon: Securing medical devices through wireless monitoring and anomaly detection. IEEE Transactions on Biomedical Circuits and Systems. 2013; 7(6):871-881. Doi: 10.1109/TBCAS.2013.2245664