American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Zero-Trust Architecture Adoption in Resource-Constrained Universities

Dr. Robert Mitchell

Associate Professor, University of Glasgow, Glasgow, United Kingdom

Abstract

Universities operate unusually complex digital ecosystems containing student information systems, learning-management platforms, research infrastructure, cloud services, administrative applications, personally owned devices, Internet of Things endpoints, laboratories, guest networks, remote users, and geographically distributed institutional resources. Conventional perimeter-oriented security becomes increasingly difficult to sustain in such environments because legitimate users and protected resources routinely operate outside a single trusted network boundary. The National Institute of Standards and Technology defines zero trust as a cybersecurity paradigm that removes implicit trust based solely on network location and places protection around users, assets, and resources through continuous and context-sensitive authorization. Although Zero-Trust Architecture offers a compelling security model for universities, institutions with limited cybersecurity staffing, aging infrastructure, decentralized administration, and restricted technology budgets cannot realistically replace their entire security stack simultaneously.

This study develops a phased, resource-sensitive framework for Zero-Trust Architecture adoption in universities operating under financial and technical constraints. The proposed approach integrates identity and access management, multifactor authentication, device posture, data classification, application-level authorization, network segmentation, visibility, analytics, governance, and legacy-system accommodation. The methodological framework is informed by NIST SP 800-207, NIST's finalized SP 1800-35 implementation guide, CISA's Zero Trust Maturity Model Version 2.0, and higher-education-specific implementation experience. NIST's 2025 implementation guide demonstrates that organizations can migrate incrementally toward zero trust while integrating legacy and cloud environments rather than requiring complete infrastructure replacement at the outset.

A transparent simulation prioritizes six adoption domains for a resource-constrained university, with identity and access management receiving the highest illustrative implementation priority score of 94, followed by data protection at 89 and device security at 84. The numerical values are methodological simulations rather than observations from actual institutions. The study concludes that financially constrained universities should approach zero trust as a maturity journey rather than a single procurement project, beginning with high-value identity, data, and device controls before progressing toward finer network segmentation, application-level policy enforcement, advanced analytics, automation, and continuous authorization.

American Journal of Advanced Multidisciplinary Innovation and Research (AJAMIR)

E-ISSN: XXXX-XXXX ● Website: www.ajamir.com ● Email: editor@ajamir.com

Keywords: Zero-Trust Architecture, university cybersecurity, higher education, identity and access management, multifactor authentication, network segmentation, cybersecurity governance, resource constraints, least privilege, cyber resilience

1. Introduction

Universities present a distinctive cybersecurity environment because openness, collaboration, mobility, research autonomy, and decentralized information technology are fundamental to the academic mission. Students move between personal and university-managed devices, researchers collaborate across institutions, faculty members access cloud applications from multiple locations, laboratories maintain specialized equipment, and administrative departments operate systems containing financial, employment, identity, and student information. The traditional assumption that activity occurring inside an institutional network boundary deserves greater trust has therefore become increasingly difficult to defend. NIST SP 800-207 defines zero trust as an approach that shifts security away from static network perimeters and toward explicit protection of users, assets, and resources, eliminating automatic trust based solely on physical or network location.

The higher-education sector is particularly compatible with zero-trust reasoning because university users, devices, applications, and data are dispersed across campus networks, homes, cloud environments, partner organizations, research facilities, and mobile infrastructure. EDUCAUSE describes zero trust for higher education through continuous explicit verification, least-privileged access, and an assumption that compromise can occur inside as well as outside institutional boundaries. EDUCAUSE has also emphasized that modern university boundaries have become difficult to define because users access institutional resources from remote networks, diverse devices, and hybrid-cloud environments. These characteristics weaken the usefulness of a cybersecurity model that relies primarily on trusted internal networks protected by a strong external perimeter.

Adoption nevertheless becomes difficult when universities lack the financial and human resources available to major technology enterprises. A zero-trust program may require stronger identity governance, multifactor authentication, device management, telemetry, application modernization, segmentation, policy enforcement, and continuous monitoring. Smaller universities may simultaneously operate legacy enterprise-resource-planning systems, specialized research equipment, unsupported laboratory devices, departmental servers, and applications whose authentication mechanisms cannot immediately support modern identity protocols. NIST's implementation guidance explicitly recognizes that organizations often need gradual migration and may integrate zero-trust capabilities with existing legacy enterprise and cloud systems rather than replacing every component simultaneously. This incremental principle is especially important for resource-constrained higher-education institutions.

Zero trust should therefore not be understood as a product that a university can purchase and declare complete. NIST describes Zero-Trust Architecture as a security model, system-design approach, and coordinated cybersecurity-management strategy, while CISA organizes adoption as a maturity progression across identity, devices, networks, applications and workloads, and data, supported by visibility and analytics, automation and orchestration, and governance. The architectural objective is to reduce implicit trust and evaluate access according to contextual evidence. A mature university implementation may eventually incorporate continuous authorization, risk-aware access policy, device

posture, microsegmentation, behavioral monitoring, and automated response, but the path toward that state can be staged.

This paper develops a financially realistic zero-trust adoption framework for universities that cannot implement every control simultaneously. The central argument is that scarce cybersecurity resources should first be directed toward controls with broad institutional leverage, particularly identity, privileged access, sensitive-data protection, and managed-device visibility. Later phases can introduce progressively stronger application segmentation, network controls, analytics, policy automation, and continuous verification. The paper uses a conceptual-methodological design rather than claiming completion of a field deployment. Numerical adoption scores presented later are explicitly simulated and are intended to illustrate how a university could prioritize investments under limited resources.

2. Objectives of the Study

2.1 To Identify High-Value Zero-Trust Controls for Resource-Constrained Universities

The first objective is to determine which zero-trust capabilities should receive priority when an institution cannot fund a comprehensive transformation at once. The study focuses on security controls that can reduce risk across multiple applications and services, particularly centralized identity, multifactor authentication, privileged-access management, data classification, device inventory, device posture, least-privileged authorization, and monitoring.

The objective reflects NIST's principle that access decisions should consider subject and asset information rather than granting implicit trust merely because the request originates from an internal network.

2.2 To Develop a Phased Adoption Model Compatible with Legacy Infrastructure

The second objective is to establish a migration sequence that allows zero-trust principles to coexist temporarily with existing campus infrastructure.

The framework therefore avoids the assumption that legacy applications, network architecture, and research systems can be replaced simultaneously. NIST SP 1800-35 specifically provides practical implementation examples intended to support gradual migration and integration across on-premises and cloud systems. Its final 2025 publication emerged from an NCCoE project involving 24 collaborators and 19 example ZTA implementations.

2.3 To Balance Cybersecurity Improvement with Academic Accessibility and Cost

The third objective is to prevent zero-trust implementation from undermining the legitimate accessibility requirements of teaching, learning, and research.

Security controls should be proportionate to resource sensitivity, user role, device condition, and risk. A student accessing ordinary public learning materials should not necessarily encounter the same controls as an administrator modifying payroll information or a researcher accessing restricted intellectual property.

3. Review of Literature

3.1 Zero-Trust Architecture and Continuous Verification

NIST's SP 800-207 provides the foundational federal model for Zero-Trust Architecture. Rose and colleagues describe zero trust as moving cybersecurity away from static network-based trust toward protection focused on resources and explicit authorization. Under this model, network location alone does not provide sufficient evidence that a user or device should be trusted.

The architecture relies conceptually on a policy decision process that evaluates access requests according to multiple sources of information. In practical environments, these can include authenticated identity, device posture, requested resource, time, behavior, session information, and organizational policy. The result is a model in which authenticated users are not granted unlimited lateral access merely because they have entered the institutional network.

This principle is highly relevant in universities because institutional identity populations are large and dynamic. Students enroll and graduate, faculty hold temporary appointments, visiting researchers require short-term access, external collaborators participate in research projects, and contractors support administrative systems. Access privileges can consequently accumulate unless identities and permissions are regularly reviewed.

NIST subsequently extended zero-trust guidance toward application and cloud-native access control in SP 800-207A, illustrating how zero-trust concepts can be used for fine-grained application-level access in multi-cloud and cloud-native environments. This progression is significant for universities increasingly dependent on software-as-a-service applications and externally hosted research infrastructure.

NIST's final SP 1800-35 moves from architecture to implementation. The practice guide demonstrates end-to-end architectures using commercially available technology and documents practical lessons from multiple implementation approaches. For resource-constrained universities, one of its most valuable implications is that zero trust can be approached progressively rather than through a complete infrastructure replacement.

3.2 Zero-Trust Maturity and Phased Implementation

CISA's Zero Trust Maturity Model Version 2.0 provides a useful framework for evaluating organizational progress. It organizes zero trust across five principal pillars: Identity, Devices, Networks, Applications and Workloads, and Data. Visibility and analytics, automation and orchestration, and governance operate as cross-cutting capabilities.

This maturity perspective is well suited to constrained institutions because it does not imply that all capabilities must reach the most advanced state simultaneously. An organization may possess strong identity security while remaining comparatively immature in network microsegmentation or automated policy orchestration.

Identity investments often provide broad coverage because institutional accounts mediate access to email, learning systems, research repositories, cloud platforms, library systems, administrative services, and remote-access tools. Higher-education zero-trust guidance from EDUCAUSE likewise places strong emphasis on identity and access-management foundations.

Network segmentation remains valuable but can be operationally demanding in universities because academic networks historically support broad connectivity, experimentation, temporary devices, laboratories, visitors, and specialized equipment. A staged model can therefore begin by isolating particularly sensitive zones—such as finance, human resources, research data, and administrative management interfaces—before attempting fine-grained segmentation everywhere.

The same principle applies to device security. A university may be unable to manage every personally owned student device but can require stronger posture controls for institutionally managed laptops, privileged administrators, sensitive research endpoints, and devices accessing high-value resources.

3.3 Higher-Education Constraints and Security Prioritization

EDUCAUSE describes zero trust as increasingly relevant to higher education because institutional users, applications, and resources exist across distributed environments that no longer correspond cleanly to a traditional campus perimeter. Its higher-education guidance recommends beginning with important assets and use cases rather than attempting to redesign the complete institutional environment at once. This approach is particularly suitable for institutions facing restricted budgets.

A university's limited cybersecurity resources can be concentrated initially around accounts and assets whose compromise would generate substantial operational or regulatory consequences. These may include privileged administrative identities, identity-provider infrastructure, finance systems, human-resource applications, research datasets, backup environments, security-management consoles, and executive accounts.

Resource constraints also strengthen the case for reusing existing capabilities. An institution that already operates a central identity provider may extend conditional-access policies rather than acquiring an entirely new identity stack. Existing endpoint-management infrastructure may support device posture. Current firewalls or virtual-network technologies may provide initial segmentation. Existing logging infrastructure can be focused first on high-value systems.

CISA's maturity framework reinforces this incremental logic because zero-trust maturity includes governance and visibility as well as technology. Improving asset inventories, access policies, identity lifecycle management, and logging discipline may therefore advance zero-trust maturity even before major infrastructure replacement becomes financially feasible.

4. Research Methodology

4.1 Research Design

The study adopts a conceptual-methodological research design supported by simulation-based implementation prioritization.

No university network was penetrated, no institution-specific breach dataset was used, and no real university was assigned a cybersecurity maturity score.

The framework was developed through synthesis of current zero-trust standards and implementation guidance, particularly NIST SP 800-207, NIST SP 800-207A, finalized NIST SP 1800-35, CISA's Zero Trust Maturity Model Version 2.0, and higher-education implementation guidance published through EDUCAUSE.

The analytical assumption is that a hypothetical university has limited security staff, a mixed legacy/cloud environment, incomplete device management, decentralized IT administration, and insufficient funding for immediate enterprise-wide replacement.

4.2 Resource-Sensitive Zero-Trust Adoption Framework

Zero-trust domains were evaluated using five planning considerations: expected risk reduction, institutional coverage, implementation cost, dependence on legacy-system modernization, and operational staffing requirements.

Table 1: Proposed Phased Zero-Trust Adoption Framework for Resource-Constrained Universities

Adoption Phase	Priority Capability	Core Activities	Resource-Sensitive Implementation Approach	Expected Security Benefit
Phase I	Identity and Privileged Access	Central identity, MFA, account lifecycle, privileged-account control	Extend existing identity platform before acquiring new tools	Reduces account takeover and excessive privilege
Phase II	Data and Critical Asset Classification	Identify sensitive data, systems and high-value research assets	Prioritize a limited number of high-impact resources	Enables differentiated protection
Phase III	Device Visibility and Posture	Asset inventory, managed-device status, patch posture, endpoint protection	Begin with staff, administrators and critical research devices	Reduces access from unknown or compromised endpoints
Phase IV	Application-Level Least Privilege	Role-based and context-sensitive authorization	Modernize high-risk applications first	Restricts unnecessary access

Source: Author-developed framework based on NIST and CISA zero-trust guidance and higher-education implementation considerations.

4.3 Simulation and Priority Index

The formula is not a recognized NIST or CISA metric. It is an author-developed methodological instrument for prioritization.

Six domains were assigned simulated scores:

Identity and Access = 94 Data Protection = 89 Device Security = 84 Visibility and Analytics = 81 Applications and Workloads = 79 Network Segmentation = 74

Higher scores indicate stronger justification for early implementation under assumed financial constraints, not greater intrinsic cybersecurity importance in every institution.

5. Analysis

5.1 Simulated Adoption Priorities

The priority analysis indicates that identity and access management should constitute the first major zero-trust investment for the hypothetical resource-constrained university.

Table 2: Simulated Zero-Trust Adoption Priorities

Zero-Trust Domain	Simulated Priority Score	Suggested Adoption Stage	Principal Rationale
Identity and Access	94	Immediate	Protects large numbers of applications through a shared control layer
Data Protection	89	Immediate–Early	Focuses scarce resources on the most sensitive information
Device Security	84	Early	Reduces risk from unmanaged or compromised endpoints
Visibility and Analytics	81	Early–Intermediate	Provides evidence needed for continuous verification
Applications and Workloads	79	Intermediate	Enables fine-grained authorization but may require modernization
Network Segmentation	74	Intermediate–Advanced	Limits lateral movement but can require significant redesign

Note: Scores are simulated methodological values. They do not represent measured readiness or security effectiveness at an actual university.

Identity receives the strongest priority because one identity-control improvement can influence access to numerous institutional systems. A centralized identity provider combined with MFA, privileged-access restrictions, account lifecycle management, and role-aware authorization can therefore generate wider coverage than replacing isolated security appliances.

NIST emphasizes that zero trust removes implicit trust in users and assets and requires explicit policy decisions concerning resource access. CISA similarly places identity as one of the five foundational zero-trust maturity pillars.

Data protection follows closely because the institution cannot defend every digital object at the same level. Data classification helps direct scarce funding toward information whose confidentiality, integrity, or availability is genuinely important.

Device security receives a score of 84 because authenticated identity alone cannot establish that an endpoint is trustworthy. Institutions should progressively distinguish known managed devices from unknown or high-risk devices and apply stronger requirements to sensitive access.

5.2 Lollipop Graph of Zero-Trust Priorities

Graph 1: Simulated Implementation Priorities for Zero-Trust Adoption in Resource-Constrained Universities

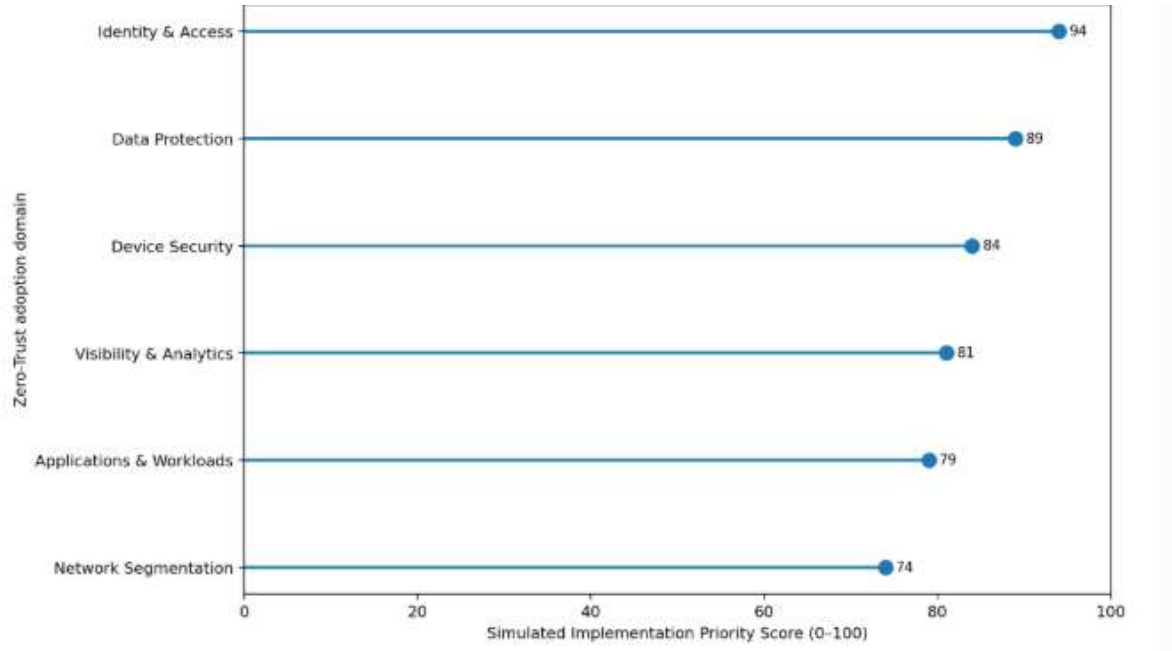

Source: Author-generated simulation.

The lollipop graph shows that identity and data controls form the strongest early investment priorities, followed by device security and visibility. Applications, workloads, and network segmentation remain important but receive slightly lower early-stage scores because they can depend upon more extensive modernization and architecture work.

The pattern should not be interpreted as evidence that segmentation is unimportant. NIST zero-trust principles explicitly shift protection toward individual resources and smaller resource groups, and

network or software-defined segmentation can help constrain lateral movement. The lower simulated priority reflects sequencing under scarcity rather than security value.

An institution with a mature identity platform but a flat and vulnerable network could reasonably assign a much higher priority to segmentation.

5.3 Cost-Aware Zero-Trust Migration Strategy

The analysis supports a four-stage university migration strategy.

Stage One: Establish Identity Control. The university should enforce MFA on privileged and high-risk accounts first, then expand to faculty, staff, and students. Dormant accounts should be removed promptly, shared administrator credentials reduced, and privileged accounts separated from ordinary user accounts. This phase can often leverage existing identity and directory infrastructure.

Stage Two: Identify High-Value Assets. The institution should document sensitive research data, financial information, human-resource information, identity infrastructure, backups, security consoles, and critical administrative services. Access policies can then be differentiated according to actual institutional risk.

Stage Three: Add Context and Segmentation. Device health, access location, application sensitivity, and user role can progressively inform authorization. Existing firewalls, network zones, cloud security groups, or virtual networking may provide initial segmentation before advanced microsegmentation tools become financially possible.

Stage Four: Build Continuous Verification. Logging, behavioral analytics, automated response, policy orchestration, and risk-based session reevaluation can be strengthened once identity, device, application, and data foundations are sufficiently mature.

This phased approach is consistent with NIST's implementation position that enterprises may migrate gradually toward greater use of zero-trust concepts while integrating legacy and cloud environments.

6. Discussion

6.1 Identity-First Adoption Provides High Security Leverage Under Limited Budgets

For resource-constrained universities, the principal question is not which zero-trust technology is theoretically most advanced but which early control improves the largest portion of the institutional security environment for an affordable level of effort. Identity frequently satisfies this requirement because the same institutional credentials are reused across email, learning systems, cloud services, research tools, finance systems, administrative applications, and remote access. Strengthening identity can therefore reduce risk across multiple services without requiring each application to be independently redesigned at the beginning of the transformation. NIST's architecture places explicit authentication and authorization at the center of access decisions, while CISA identifies identity as a principal maturity pillar rather than a peripheral supporting service.

Identity-first adoption should nevertheless extend beyond simply enabling MFA. A university can deploy MFA and still preserve excessive privileges, inactive accounts, shared administrative identities, poor contractor offboarding, or unmanaged service credentials. Zero-trust maturity requires a broader identity lifecycle. User privileges should reflect current roles, high-risk privileges should be temporary where technically feasible, access to sensitive systems should receive stronger verification,

and abnormal authentication patterns should be monitored. EDUCAUSE's higher-education guidance likewise presents identity and access management as a foundational mechanism for implementing zero-trust principles across dispersed academic environments.

A resource-sensitive implementation can sequence these controls according to risk. Privileged administrators, finance users, institutional leaders, identity-management staff, and researchers handling highly sensitive data can receive stronger protections before every low-risk student application has been modernized. This prioritization does not create permanent unequal security; it provides a financially realistic migration pathway in which the university addresses its greatest exposure first and expands coverage as institutional capability improves.

6.2 Zero Trust Must Accommodate Academic Openness Without Recreating a Closed Corporate Network

Universities differ from many commercial enterprises because openness is part of their mission. Academic environments require guest connectivity, external collaboration, temporary research accounts, student-owned devices, rapidly changing project teams, experimentation, and access to global knowledge infrastructure. A security architecture that attempts to eliminate these characteristics would protect the institution partly by undermining the reason the institution exists. Zero trust is therefore particularly useful when interpreted correctly because it does not require every user and device to belong permanently to a closed trusted network. Instead, it allows access to be granted according to identity, context, resource sensitivity, and policy. NIST's model explicitly moves the security focus away from a trusted network location toward users, assets, and resources.

This means a university does not necessarily need to fully manage every student device. A personally owned device might access public learning content and low-risk services after ordinary authentication while being prevented from reaching administrative databases or sensitive research environments. Access to a high-value research repository might require MFA, an institutionally managed endpoint, current security posture, and a more specific role. Zero trust therefore enables differentiated controls rather than demanding uniform restrictions.

The same reasoning applies to research collaboration. External researchers may receive narrowly scoped authorization to a particular resource rather than broad network-level access. Application-level and resource-level policy becomes particularly useful in these circumstances because it can reduce the need to extend generalized internal trust to external collaborators. NIST SP 800-207A's focus on application-level access control in distributed cloud environments supports this finer-grained direction.

The critical implementation challenge is usability. Excessive authentication prompts, unpredictable access decisions, or poorly communicated restrictions can encourage workarounds and weaken institutional acceptance. Zero-trust policies should therefore minimize unnecessary friction through risk-sensitive authentication and clear access design. Strong verification should be concentrated where the consequence of compromise justifies it.

6.3 Governance and Maturity Are More Important Than Purchasing a “Zero-Trust Product”

The final discussion point concerns the misconception that zero trust can be acquired through a single commercial platform. NIST presents ZTA as an architecture and coordinated security-management

strategy, while CISA's maturity model spans identity, devices, networks, applications, data, visibility, automation, and governance. An institution can therefore purchase sophisticated technology and still remain at low zero-trust maturity if identity inventories are inaccurate, access privileges are not reviewed, critical assets are unknown, logs are not monitored, and responsibility for policy decisions is unclear.

This point is especially important for universities with small cybersecurity teams. Procurement can create significant recurring licensing and operational costs. Existing technologies should therefore be evaluated for their ability to support zero-trust functions before entirely new platforms are purchased. Current identity systems may support MFA and conditional access. Existing endpoint-management tools may provide posture information. Firewalls and virtual networks may support initial segmentation. Cloud platforms may already provide granular security groups, workload identities, and centralized logs. NIST SP 1800-35 was explicitly developed to provide practical examples and lessons that organizations can use to make implementation more feasible and cost effective.

Governance should determine which controls are introduced and why. A university needs named ownership for identity policy, device security, network design, research-data access, application risk, logging, incident response, and exception management. Exceptions are particularly important because legacy scientific instruments or specialized applications may not support contemporary authentication. Rather than pretending that such systems satisfy zero-trust requirements, institutions should document the exception, isolate the resource where practical, reduce unnecessary connectivity, monitor activity, and establish a modernization or retirement plan.

Zero trust should consequently be measured through reduced implicit trust and improved verification capability rather than through the number of products deployed. For constrained universities, a smaller set of well-governed controls can provide more meaningful maturity than an expensive but poorly integrated collection of security technologies.

7. Conclusion

Zero-Trust Architecture offers a strong conceptual fit for contemporary universities because academic information environments are highly distributed, dynamic, collaborative, and increasingly dependent on cloud and remote access. Traditional network location provides weak evidence of trust when faculty, students, researchers, contractors, applications, data, and devices operate across numerous networks and administrative boundaries. NIST's zero-trust model directly addresses this environment by moving protection toward users, assets, and resources and by requiring explicit authorization rather than assuming that internal network presence is trustworthy.

The principal difficulty for resource-constrained universities is not whether the zero-trust concept is desirable but how it can be implemented without requiring an unaffordable enterprise-wide transformation. The analysis presented in this paper supports a phased strategy. Identity and access management should generally receive early attention because centralized authentication, MFA, privileged-access controls, and improved account lifecycle management can influence numerous institutional services simultaneously. Data classification should follow closely so that scarce protection resources are directed toward the information whose compromise would create the greatest harm. Device security and visibility can then provide stronger contextual evidence for access decisions, while

application-level authorization and network segmentation can be expanded as modernization resources become available.

The simulated priority analysis reinforces this sequencing by assigning identity and access the highest illustrative priority score of 94, followed by data protection at 89 and device security at 84. These scores are not empirical measurements and should not be interpreted as universal rankings for every university. Institutions with different infrastructure, threat exposure, staffing, or technical maturity may generate different priorities. Their value lies in demonstrating that zero-trust migration should be explicitly prioritized rather than approached as an undifferentiated list of technologies.

The broader conclusion is that resource constraints do not make zero trust impossible; they make disciplined sequencing essential. A university should not postpone all progress until it can afford an ideal architecture. It can progressively reduce implicit trust through better identity governance, stronger authentication, critical-asset classification, device visibility, least privilege, segmentation, logging, and continuous verification. NIST's implementation guidance and CISA's maturity framework support precisely this evolutionary interpretation of zero trust.

Future empirical research should validate the proposed framework across small and medium-sized universities, compare security gains against implementation costs, measure user friction and help-desk effects, assess legacy-system integration, and develop higher-education-specific zero-trust maturity benchmarks. The ultimate objective should not be to imitate the cybersecurity architecture of a large commercial enterprise but to create a university security environment in which academic openness is preserved while access to valuable resources is continuously, proportionately, and explicitly justified.

References

1. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology.
2. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2019). *Zero Trust Architecture* (Initial Public Draft, NIST SP 800-207). National Institute of Standards and Technology.
3. Kindervag, J. (2010). *No More Chewy Centers: The Zero Trust Model of Information Security*. Forrester Research.
4. Kindervag, J., & Balaouras, S. (2017). *The Zero Trust Model of Information Security*. Forrester Research.
5. Ward, R., & Beyer, B. (2014). BeyondCorp: A new approach to enterprise security. *;login: The USENIX Magazine, 39*(6), 6–11.
6. Google. (2019). *How Google adopted BeyondCorp*. Google Online Security Blog.
7. Google. (2014). *BeyondCorp: A new approach to enterprise security*. Google Security Research.
8. Gilman, E., & Barth, D. (2017). *Zero Trust Networks: Building Secure Systems in Untrusted Networks*. O'Reilly Media.
9. CISA. (2020). *Defend Today, Secure Tomorrow: Federal Zero Trust Strategy*. Cybersecurity and Infrastructure Security Agency.
10. NIST. (2020). *Cybersecurity Framework for Improving Critical Infrastructure Cybersecurity, Version 1.1*. National Institute of Standards and Technology.

11. NIST. (2017). *Digital Identity Guidelines: Authentication and Lifecycle Management (SP 800-63B).* National Institute of Standards and Technology.
12. NIST. (2018). *Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy* (SP 800-37 Rev. 2). National Institute of Standards and Technology.
13. NIST. (2013). *Security and Privacy Controls for Federal Information Systems and Organizations* (SP 800-53 Rev. 4). National Institute of Standards and Technology.
14. Cloud Security Alliance. (2019). *Software-Defined Perimeter and Zero Trust: A Practical Guide*. Cloud Security Alliance.
15. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust Architecture: General Deployment Models and Use Cases*. National Institute of Standards and Technology.
16. EDUCAUSE. (2020). *2020 EDUCAUSE Horizon Report: Teaching and Learning Edition.* EDUCAUSE.
17. EDUCAUSE. (2019). *Top 10 IT Issues, 2019: Higher Education and Digital Transformation*. EDUCAUSE.
18. EDUCAUSE. (2018). *2018 Top 10 IT Issues: The Higher Education IT Landscape*. EDUCAUSE.
19. EDUCAUSE. (2017). *2017 Top 10 IT Issues: Be the Change You Seek*. EDUCAUSE.
20. ENISA. (2020). *Cybersecurity Guidelines for SMEs on the Security of Cloud Services*. European Union Agency for Cybersecurity.