International Journal of Computational and Experimental Science and ENgineering (IJCESEN)

Vol. 11-No.3 (2025) pp. 4966-4995
http://www.ijcesen.com

ISSN: 2149-9144

Copyright © IJCESEN

Research Article

Next-generation secure authentication and access control architectures: advanced techniques for securing distributed systems in modern enterprises

Prince Kumar*

Independent Researcher, Visvesvaraya Technological University, Belgaum, India
* **Corresponding Author Email:** princem4u@rediffmail.com - **ORCID:**0000-0003-2572-3905

Article Info:

DOI: 10.22399/ijcesen.3294
Received : 02 May 2025
Accepted : 06 July 2025

Keywords

Secure Authentication
Access Control
Zero Trust Architecture
Distributed Systems Security
Cyber Threat Mitigation

Abstract:

With the move of enterprises toward cloud computing, the Internet of Things (IoT) and distributed digital ecosystems, traditional authentication and access control mechanisms like Role-Based Access Control (RBAC) and static Multi Factor Authentication (MFA), are not meeting the mark. This paper reviews next-generation authentication and access-control designs, including Zero Trust, Attribute-Based Access Control (ABAC), and risk-adaptive policies and examines how contextual authentication, continuous identity verification, and emerging technologies such as AI-powered risk assessments, privacy-preserving identity authentication, and decentralized identity models can be integrated. This paper also details ways in which AI has been utilized to power dynamic, risk adaptive mechanisms that can use real time contextual data such as user behaviour, device health and threat level to adjust permissions levels accordingly. Improvements in security, usability, scalability and adaptability are shown in comparative analyses with traditional models. In addition, the review describes the means by which policymakers and industry practitioners can derive implications from the research and provides future research and deployment strategies to be considered. This paper synthesizes recent advancements towards more predictive and resilient authentication frameworks with a goal to enable development of such authentication approaches to help reduce sophisticated cyber threats in distributed systems while supporting industry specific regulatory compliance.

1. Introduction

Distributed systems are sprawling across cloud services, mobile devices, IoT and on-premises infrastructure, and modern enterprises count on them. Such environments necessitate that only the legitimate users and devices have access to their resources, yet only to the extent and way they are authorized. Authentication and access control mechanisms are the backbone of distributed system security and verify identities and enforce permissions across many disparate components [1]. The long-understood primary defence for protecting data and services in any large-scale system [1] is these controls. With distributed architectures becoming more complex, the problem of allowing people to see only what they need to and only when they need to become both more pressing and more difficult to manage. This is a topic of high importance in today's research, thanks to the large number of attacks directed at the weaknesses of authentication and authorization. According to breach reports, compromised identities remain a de facto leading attack vector for cyber incidents [2]. Take, for instance, out of 81%, about 81 per cent of organizations attacked used weak or stolen passwords, which implies a need for stronger authentication practices [2]. In addition, it is still way too easy for attackers to amend their rights: 'broken access control' has just been crowned the 1 most critical web application security risk, which reinforces how this risk is continuously exploited in the wild [3]. Clearly, advancing the state of authentication and access control is not just an academic pursuit but a practical necessity to mitigate modern security threats. This subject is also highly significant in the broader cybersecurity and enterprise security domain. Effective identity and access management lies at the heart of zero-trust strategies and other modern security frameworks that enterprises are adopting. In fact, contemporary architecture paradigms like Zero

Trust explicitly assume no user or device should be implicitly trusted, requiring continuous verification of identity, context, and policy compliance for every access request [4]. This shift toward an "always verify" model is driven by trends such as remote work, BYOD (bring your own device), and cloud computing trends that have dissolved the traditional network perimeter [4]. As a result, robust authentication and fine-grained access control are now often cited as fundamental to enterprise security posture, effectively becoming the new perimeter in distributed, cloud-enabled organizations. Despite its criticality, secure authentication and authorization in distributed systems present numerous open challenges. One big problem is that many existing distributed systems have not been constructed with strong notions of security in mind and it is difficult to add consistent authentication and authorization in a clean way [5]. challenges indicate that current solutions are falling short, and they highlight specific gaps (This often leads to a dichotomy in practice: some systems enforce overly rigid access rules that hamper usability and cross-domain collaboration, while others err on the side of openness, risking serious data breaches by granting overly permissive access [5]. Synchronizing the identities and permissions of a user across heterogeneous services and geographic regions is inherently complex. Issues like network latency, intermittent connectivity, and differences in security protocols between subsystems can all introduce gaps or inconsistencies in enforcement. Furthermore, there is an enduring tension between security and user convenience – stronger authentication (e.g., mandatory multi-factor) can introduce friction, which in turn may lead users to seek risky workarounds. These such as lack of dynamic policy adaptation and unified credential management, that current research has yet to fully address. The current state of knowledge in the field provides a starting point but also illustrates why a new model is needed. Over the years, researchers and practitioners have developed a variety of authentication mechanisms (passwords, biometrics, one-time tokens, etc.) and access control models. Traditional models like discretionary access control and role-based access control (RBAC) are widely deployed in enterprises, and standards for federated identity (e.g., SAML, OAuth) allow users to authenticate across distributed services. However, these existing approaches each have limitations when applied to complex, distributed ecosystems. RBAC, for instance, struggles to handle the dynamic, context-dependent access needs of modern enterprises, leading to role explosion or inflexible policies. Attribute-Based Access Control (ABAC) has emerged as a promising next-generation model that can offer more fine-grained and context-aware decisions by evaluating attributes of users, resources, and environment [6]. Yet even ABAC is not a panacea – recent research points out ongoing issues such as policy management complexity and performance overhead at scale [6]. In practice, organizations end up patching together multiple partial solutions (network perimeters, VPNs, single sign-on, cloud IAM tools), which can be complex to manage and still leave security gaps. This situation suggests that incremental improvements may no longer suffice; instead, a more holistic rethinking of authentication and authorization in distributed environments is required.

Given these challenges and gaps, there is a clear need for a new theoretical framework or architecture that can better secure distributed systems in modern enterprises. Researchers are beginning to explore advanced techniques and models aimed at "next generation" authentication and access control. Examples include continuous authentication systems that regularly re-verify user identity and context during a session, risk-adaptive access control that adjusts permissions on the fly based on threat level, and decentralized identity schemes (often blockchain-based) that eliminate single points of failure in identity management [1]. The momentum in this area reflects a recognition that a more adaptive, intelligent, and scalable approach is required, one that can unify these innovations into a coherent security architecture.

This review surveys state-of-the-art developments in secure authentication and access control for distributed systems and synthesizes them into a comprehensive model for next-generation enterprise security. The evolution of authentication and authorization mechanisms is examined, key advanced techniques from recent literature are highlighted, and their effectiveness in addressing (or, in some cases, their shortcomings in addressing) the challenges outlined above is discussed. A theoretical framework that integrates these techniques is then proposed, aiming to fill gaps in current research. By the end of this review, readers will possess a clear understanding of the limitations of existing authentication and access control approaches and insight into how emerging architectures and methods can significantly strengthen the security of modern distributed enterprises.

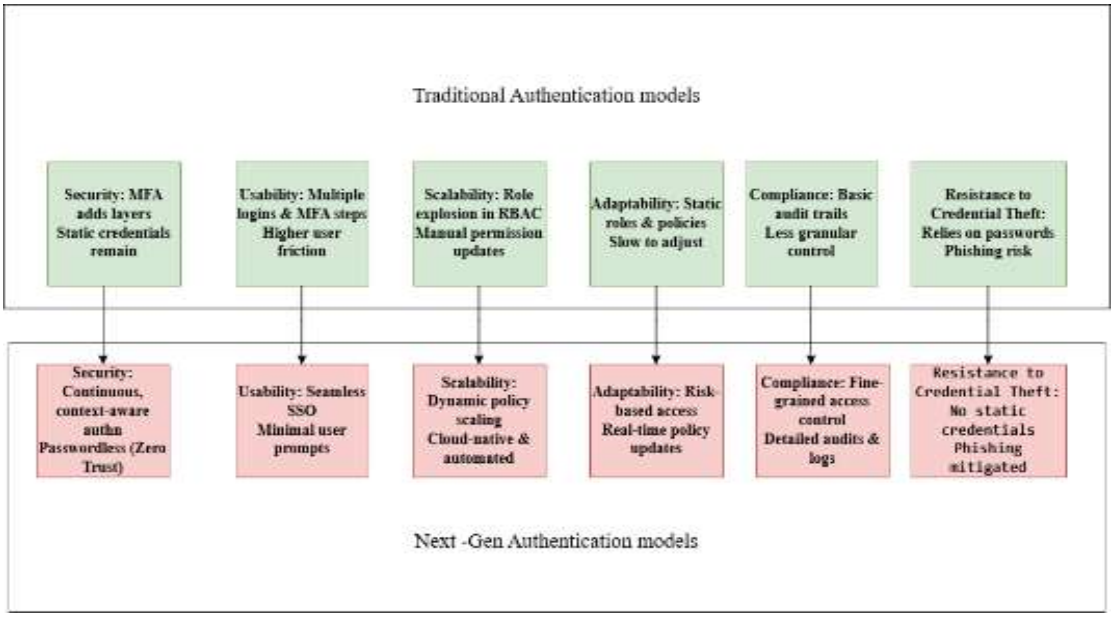

Figure 1. *Tradition vs. Next-Gen Authentication Models*

2. Next-generation secure authentication and access control architecture

Today, next-generation authentication and access control architectures are required by modern enterprises that include systems and cloud services spread over miles of cabling and sophisticated cyber threats. Within this theoretical framework is the integration of advanced identity management, strong authentication protocols, flexible access control models and strong enforcement of policy. Assuming zero implicit trust, high scalability, interoperability across platforms and resilience to developing attacks, it operates. This framework's system components, underlying assumptions and potential enterprise security applications are detailed in the following sections.

2.1 Core Components of the Framework

2.1.1 Identity Management

The foundation of the architecture is offered by identity management, which manages the creation and persistence of digital identities for users, devices and services. It includes defining unique user profiles, assigning roles and governing access privileges [7]. Identity and Access Management (IAM) systems in modern systems define each thing (human, devices) with a set of unique attributes such as a name or an email address and link that entity with the appropriate permissions and attributes, making sure that everything is unambiguously tied to what it should be. This covers lifecycle management (provisioning and de-provisioning accounts) and identity federation across the boundaries of organizations. According to prior research, identity management is recognized as an essential element of cybersecurity, enabling the linking between user identities, security policies and controls. The framework further includes the following assumptions: 1) Integrate with identity providers (IdPs) that can validate identities and provide an authentication token; 2) Users have been assigned to roles, and resources are included in the graph. Sometimes, federated identity management is used so that users are able to use a single digital identity on multiple trusted domains [7]. This reduces duplicate copies of accounts and passwords and makes SSO possible and simplifying the user experience. Key components will support identity governance (compliance through use of least privilege assignments and timely purge of access), directory services (storing identity data) and potentially decentralized or self-sovereign identities as may be needed for future applications. With enterprise systems more distributed and more dynamic, so too has identity management to provide greater scalability, agility and security intelligence. Yet IAM capabilities are increasingly being enhanced by Artificial Intelligence (AI). This is being accomplished with machine learning algorithms detecting anomalous behaviors, assessing contextual risk in real time and servicing dynamic, risk-aware access decisions. This enables IAM systems to extend beyond static polices to incorporate behavioural analytics, environmental signals and access history to inform the logic of access control.

2.1.2 Authentication Protocols

The need to verify identity claims before granting access, then, makes secure authentication critical. Robust authentication protocols and multi-factor

authentication are used in next-generation architectures to allow legitimate users and devices only. To combat increasingly complex cyber threats and to enable scalable, context-aware authentication, these architectures become more adaptive, AI-powered and password-free. In lieu of or in addition to traditional username/password schemes, stronger factors (typically cryptographic keys, biometrics or time-based one-time codes) are used. The most common protocols that are used are OAuth 2.0, OpenID Connect (OIDC), Security Assertion Markup Language (SAML) and Kerberos; although there are others. For instance, each has a different purpose: OAuth 2.0 acts as an authorization framework that governs access to APIs / resources and OIDC and SAML are authentication standards that support SSO and federated identity exchange [8]. For example, SAML is often used in enterprise SSO scenarios (exchanging XML-based assertions between an IdP and a service provider), while OIDC builds on OAuth 2.0 to add an identity layer with JSON web tokens (JWTs) for modern web and mobile applications. To support zero trust principles, modern authentication frameworks also implement continuous authentication techniques that verify user identity throughout the session based on behavioral patterns, device signals, and environmental context. The framework also supports passwordless authentication (e.g., FIDO2/WebAuthn), certificate-based device authentication, and continuous authentication methods that repeatedly validate user presence or context. Multi-factor authentication (MFA) is mandatory for sensitive operations, combining something the user knows (password/PIN), has (smart card, token, or phone), or is (biometric) to significantly harden the authentication process. AI-driven systems are increasingly used to evaluate authentication risk in real time, enabling dynamic enforcement of additional verification steps when anomalies or policy violations are detected. By using open standards and protocols, the authentication component remains interoperable with a wide range of clients and services, which is essential in a heterogeneous enterprise environment. All authentication events yield cryptographic assertions or tokens that downstream components can trust for making authorization decisions.

2.1.3 Access Control Models

Once identities are authenticated, access control models determine what actions those identities are permitted to perform on which resources. The next-generation framework supports multiple, context-aware access control paradigms to be flexibly applied depending on the use case. Traditional Role-Based Access Control (RBAC) is supported for its simplicity assigning permissions based on roles (job functions), but the framework extends beyond static role assignments. Attribute-Based Access Control (ABAC) is a core model, providing fine-grained and dynamic access decisions based on attributes of the user, resource, environment, and action [9]. In an ABAC model, policies can stipulate conditions like time of day, location, device posture, or user clearance level as prerequisites for access, not just the user's role. This model supports dynamic, real-time evaluation of contextual signals and aligns with adaptive security principles required in distributed and cloud-native environments. In order to improve decision-making under uncertainty and in the face of evolving threat conditions, AI-driven access control systems are being integrated to calculate risk scores dynamically and adaptively predict corresponding authorization policies. Behavioral analytics, past usage patterns and threat intelligence are combined in these risk-adaptive models that constantly adjust privileges per user. Modern needs, such as just-in-time access and risk-adaptive authorization, are addressed by this flexibility. It also fits well with the idea of Next Generation Access Control (NGAC), an emerging standard that subsumes and combines these (and other) models into a single, unified authorization architecture. NGAC is a radical rethinking of the notion of access control to fit in with modern distributed enterprises, with a high emphasis on fine-grained control and policy interoperability. Practically, this lets the system enforce the principle of least privilege more precisely, performing minimal privilege granting based on combinations of attributes and contextual risk. Extensions include other models, relationship-based access control (ReBAC) and risk adaptive access control (decisions adapted based on real-time risk scoring). The policy engine of the architecture can evaluate complex XACML or programmatic policies, as well as Boolean logic and conditions defined in them, to decide access. Furthermore, the access decision process can be optimised using policy learning systems, where AI refines policies over time based on outcomes and observed behaviors. All models ultimately feed into the policy decision process described below.

2.1.4 Policy Enforcement Mechanisms

At the heart of the architecture are the Policy Decision Point (PDP) and Policy Enforcement Point (PEP) components, often referenced from the XACML model and modern Zero Trust Architecture logic [9]. The PDP is the brains: it evaluates access requests against policies (the rules

defined by the access control models above) and makes the authorization decision – permit or deny, possibly with obligations (additional actions like logging). The PEP is the gatekeeper: it intercepts user or service requests to resources and enforces the decisions from the PDP [10]. In other words, when a subject (user or process) attempts to access a protected resource, the PEP will block or allow the traffic based on the PDP's evaluation of applicable policies. This separation of concerns introduces architectural modularity and ensures that policy logic is centralized, consistent, and auditable, while enforcement can be distributed across cloud, on-premises, and edge environments.

This separation of decision and enforcement provides flexibility and consistency; policies can be centrally managed (at PDP) while enforcement can be distributed close to assets. For example, a PEP might be an agent at an application, an API gateway, or a firewall enforcing that only authorized requests pass through. Modern PEPs are increasingly adaptive and state-aware; they continuously monitor session attributes, user behavior, and threat context. They can trigger re-authentication, stepped-up verification, or terminate sessions based on policy rules and risk thresholds that detect behaviors that deviate from the baseline. The PEP monitors sessions, triggers re-auth when there is a need and when a session no longer meets policy, the PEP can terminate the connection. It also usually provides logging and auditing, so you know which identities accessed what and anything denied access. This is, of course, critical for compliance and incident response. Policy enforcement is a continuous and contextual aspect, not a one-time gate at login for next-gen architectures. If long-lived, sessions may recheck authorization or may do step-up auth if risk rises mid-session. The framework's policy enforcement is tightly integrated with identity and threat intelligence systems so that enforcement decisions are updated immediately when changes occur in user status, device health or threat level. This integration in highly distributed systems can allow for policy choices to be dynamically tuned against the enterprise-wide security posture, including signals from endpoint detection and response (EDR), security information and event management (SIEM) and behavioral analytics. In this case, in a distributed microservices world, each service protected by a PEP mesh (i.e., sidecar proxy, middleware interceptor, etc) consults a central PDP or a set of decentralized PDPs. This guarantees that the policy is being enforced the same way between cloud and on-premises resources.

2.2 Key Assumptions Underlying the Framework

2.2.1 Zero Trust and Unified Trust Model

A Zero Trust security philosophy is assumed as a fundamental assumption: no user or device is trusted because they are within the traditional network perimeter. Every time, all access requests are verified beforehand against current policies and context before they're allowed. Under this framework, the network environment is assumed to be hostile by default [11], and the contributors attempt to minimize the implicit trust zones. In many cases, even post user authorization, there are authorization checks done before access of the resource and ideally on a per-request basis. Trust is thus continuously earned rather than given outright. The Zero Trust model extends beyond identity verification to include dynamic, risk-adaptive access decisions based on real-time telemetry and contextual awareness. This includes signals such as geolocation, device posture, user behavior, workload integrity, and ongoing threat intelligence feeds. The framework defines trust levels for entities, which are dynamically assessed based on multiple factors (identity assurance, device security posture, behavior anomalies, etc.). Access to sensitive resources might require a higher level of trust (e.g., MFA verification, managed device), whereas low-risk resources might be accessible with a lower trust level. However, unlike legacy models, trust is not binary or static – the architecture can adjust trust assessments in real-time. This assumption aligns with principles like "never trust, always verify" and supports continuous monitoring. The identity management component is assumed to provide strong identity proofing and credential assurance, perhaps following guidelines like NIST SP 800-63 for identity assurance levels. Users, devices, and even workloads are authenticated to a high degree of confidence before any access is granted. By eliminating the notion of a trusted internal network, the framework mitigates insider threats and lateral movement; an attacker who somehow compromises one part of the system should not easily traverse to another, because each step would require re-authentication and authorization.

2.2.2 Scalability and Interoperability

The framework is designed with the assumption that it must operate at enterprise scale, potentially handling thousands to millions of identities and devices and a high volume of access requests across distributed systems. Scalability is considered both in terms of performance (low-latency decisions for user experience) and administrative manageability

(policies must remain effective as the organization grows). This is achieved through efficient protocols (e.g., token-based authentication for stateless scaling), caching of authorization decisions where appropriate, and possibly hierarchical policy decisions (delegating some decisions to local PDPs for faster response). The framework assumes that access control processes continue to function properly as the number of users, services, and devices fluctuates [12]. Caching and replication of identity data (with eventual consistency) are used to avoid bottlenecks. Advanced identity synchronization mechanisms are leveraged to ensure that user and device states remain current across domains without introducing central points of failure. Additionally, interoperability is a key assumption: the architecture will integrate heterogeneous systems, cloud services, and third-party applications. To that end, it leverages open standards for identities (SAML, OIDC for SSO), for provisioning (SCIM – System for Cross-domain Identity Management), and for policy (such as using a standardized policy language or API). Emerging policy-as-code frameworks and identity-as-a-service (IDaaS) platforms are also assumed to support automated policy lifecycle management, enhancing consistency and agility across the enterprise. This allows different components (cloud platforms, on-premises directories, SaaS applications) to interoperate within the unified framework. It is assumed that identity federation is in place so that trust can be established between different security domains in a scalable way. This also aids scalability by offloading authentication to trusted identity providers and using standardized tokens that any service can validate. Token formats such as JWT, SAML assertions, or OAuth2 access tokens are processed via stateless validation mechanisms that support horizontal scalability and minimize verification latency. The system also presumes the use of cloud-native infrastructure (containers, serverless functions, etc.) and must integrate with their IAM controls. In terms of data, the architecture should scale out (horizontally) to handle growth, using distributed storage for identity and policy information and load-balanced PDP/PEP instances. Finally, performance and high availability are assumed to be maintained through redundancy and failover critical for enterprise adoption.

2.2.3 Resilience to Attacks

The framework operates under an assume-breach mentality: it anticipates persistent and sophisticated attack attempts, including credential theft, session hijacking, privilege escalation, and insider abuse. Therefore, a core assumption is resilience to attacks the architecture should continue enforcing security even when components fail or when under attack. Multi-factor authentication and contextual access policies mitigate the risk of compromised credentials (e.g., a stolen password alone won't grant access without the second factor) [13]. This is further strengthened through risk-aware authentication mechanisms that incorporate behavioral baselines, geolocation, device posture, and threat intelligence feeds to adapt authentication rigour in real time. It seems to be using continuous adaptive risk assessment using signals such as an unusual login pattern or anomalous resource access to trigger adaptive responses (like requiring re-authentication or limiting the access scope). The idea is similar to the Gartner Continuous Adaptive Risk and Trust Assessment (CARTA) model, where risk is continuously assessed and trust is fine-tuned in real time. In addition, the architecture assumes that all critical transactions are logged and monitored, and Behavior analytics (e.g., an account accessing data in an unusual way can be flagged and provisionally blocked for a period of time). The enhanced ability to detect subtle deviations from normative behavior across large user populations and high-volume data streams is enabled by integration with machine learning based anomaly detection systems. Hardened, distributed policy-enforcement points can be deployed so that a single compromised node does not jeopardize the entire system. Short-lived tokens combined with dynamic policy assessment prevent attackers from using a session or token outside its valid time window. Ephemeral credentials cryptographically bound to device, location, or session context provide further protection against replay and man-in-the-middle attacks. Another assumption is resilience to Denial of Service (DoS) on the authentication and authorization systems: rate limiting, overload controls, and graceful degradation are in place so that security decisions are still made under load. The framework also trusts no device by default, incorporating device health attestation. This means if a device is detected as jailbroken or infected, its trust level drops and it may be quarantined, enhancing resilience. In summary, the architecture assumes a threat-rich environment and builds in multiple layers of defense (MFA, encryption, continuous authorization, monitoring) so that a single point of failure or single exploited vulnerability will not collapse the entire security posture.

2.2.4 Potential Applications in Modern Enterprise Systems

The outlined framework has broad applicability in securing distributed systems across modern

enterprises. Its combination of strong authentication, fine-grained authorization, and continuous verification is well-suited to address several contemporary challenges:

- **Secure Remote Workforce and BYOD:** Enterprises can apply this framework to enable employees and contractors to work securely from any location. For example, adopting a Zero Trust Network Access (ZTNA) model like Google's Beyond Corp means internal applications are accessible without a VPN, but each access is governed by user identity, device trust, and policy compliance [14]. An employee using a personal device (BYOD) would be authenticated via the corporate IdP, their device posture checked, and then given application access through a PEP that enforces company policies. This allows remote work and BYOD scenarios with reduced risk, as every session is validated and lateral movement is curtailed. It also simplifies mergers and acquisitions integration – federating identities and applying consistent policies for external collaborators. Advanced session telemetry, geofencing, and real-time behavioral analytics can be layered to continuously assess trust during active sessions, triggering dynamic policy adjustments based on risk.

- **Cloud and Microservices Security**: Modern applications often consist of microservices distributed across hybrid cloud environments. This framework can be embedded in a service mesh to secure service-to-service calls. Each microservice has its own identity (such as a JWT issued by an identity platform), and requests between services are authenticated and authorized using those tokens. A centralized PDP can apply ABAC policies for APIs (e.g., service A can call service B only for certain data types and only if the end-user token grants permission). In Infrastructure-as-a-Service clouds, the framework augments native cloud IAM by providing a unifying layer of authentication (for instance, SSO across multi-cloud resources) and granular access control that goes beyond cloud roles. Policy enforcement points might be API gateways or sidecars, ensuring that even if one microservice is compromised, it cannot arbitrarily access others without a valid token and policy compliance. This is critical for preventing the spread of breaches in containerized environments. Additionally, secrets management and just-in-time privilege (granting ephemeral credentials for service accounts) can be implemented under the framework's identity management component.

- **Distributed IoT Systems**: In industries like manufacturing or healthcare, large networks of IoT devices need secure access control. The framework can manage IoT device identities (using certificates or embedded keys as device credentials) and authenticate devices to gateways or cloud services. Attribute-based policies are particularly useful here – for example, only a sensor that reports a valid location and device health status can push data to a cloud endpoint. Likewise, an IoT actuator will execute commands only if the command originates from an authorized service with the correct attributes (e.g., a command signed by an operator role and during authorized hours). The scalability assumption is vital in IoT contexts, since potentially tens of thousands of devices may concurrently authenticate and send authorization requests. The framework's emphasis on interoperability also helps in heterogeneous IoT environments where devices from different vendors must be integrated. By enforcing fine-grained access rules and continuous authentication of devices, the architecture helps contain security incidents – if one sensor node is taken over by an attacker, its credentials can be revoked or its trust level reduced without affecting others, and it will be blocked at the first policy check from causing harm.

- **Enterprise Data and Applications**: Within a large enterprise, legacy and modern applications coexist. The next-gen architecture can serve as an overlay to protect legacy apps by front-ending them with modern authentication (via SSO) and authorization layers [15]. For instance, an old HR system that only knows about internal passwords can be put behind an SSO portal that uses modern protocols and MFA, thereby upgrading its security. The framework's policy engine can enforce dynamic access policies – e.g., allow access to finance data only if the user is in the Finance department and accessing from a corporate network or an approved device. However, context-aware access control is

becoming increasingly essential as organizations strive to maintain compliance with regulations (for example, HIPAA or GDPR) and prevent data leakage. Furthermore, the architecture can support fine-grained access auditing and analytics: every access decision is logged with its rich context, and machine-learning systems can detect deviations from normal usage patterns (e.g., insider threats or compromised accounts). Furthermore, historical access data feeds into policy tuning; integration with SIEM tools and data loss prevention (DLP) systems enforces compliance and policy enforcement, which is automated.

Taken together, this next-generation authentication and access control framework provides an integrated security architecture for a loosely coupled environment where distributed systems and functions are part of an enterprise. With strong identity management, advanced authentication protocols, flexible authorization models (RBAC/ABAC/NGAC), and rigorous policy enforcement businesses can protect resources, whether they are in the data centre, on the cloud or at the edge. The approach is designed for a Zero Trust posture, strengthening remote work security, shielding cloud native applications and keeping the operation of a wide range of devices in check while maintaining scalability and interoperability. Implementing such a framework can significantly reduce the risk of unauthorized access and provide a robust, adaptable defense against the evolving threat landscape in modern enterprises.

3. Next-generation secure authentication and access control data sources and integration

Modern authentication and access control architectures leverage a broad array of data sources to verify identity and context. Unlike traditional password-based logins, next-generation Identity and Access Management (IAM) uses multifactor inputs and contextual signals (e.g. device, location, behavior) to dynamically assess trust [16]. This approach enables adaptive authentication and authorization, enhancing security and accuracy of decisions. The following section presents the key data sources and explains how their integration enhances authentication, illustrated by case studies and recent developments. The practical application of the previously proposed model is then demonstrated and validated against existing research. Figure 2 provides a comparative analysis of results between traditional and next-generation authentication models.

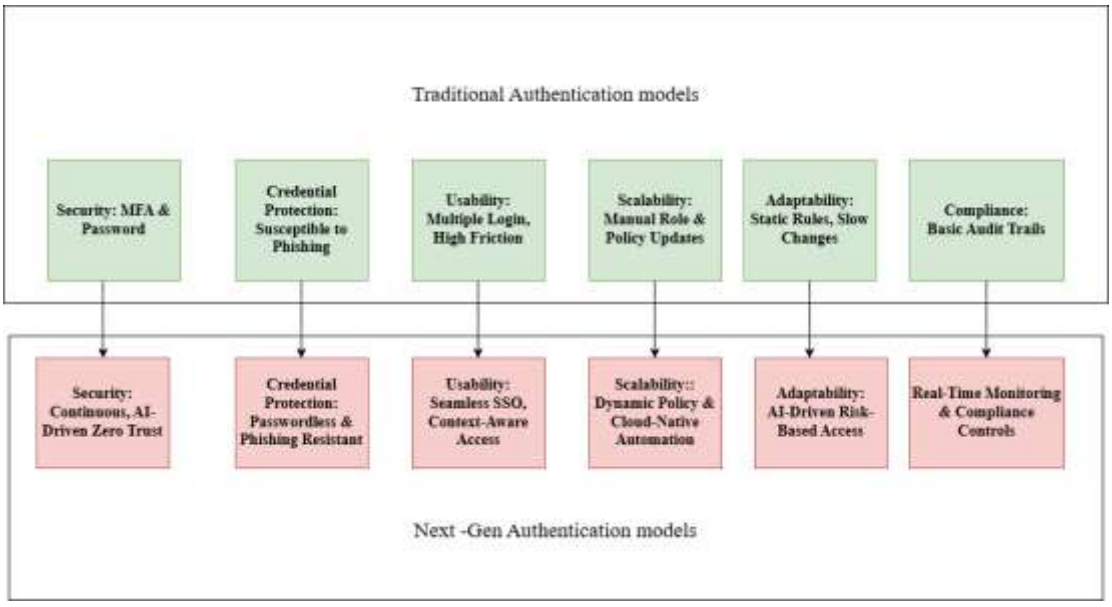

***Figure 2**. Comparative result analysis between traditional and next-gen authentication models.*

3.1 Diverse Data Sources in Next-Gen Authentication

Next-gen architectures draw on multiple **data sources** to confirm a user's identity and context before granting access. Important categories include:

- **Knowledge Factors:** Something the user knows (e.g. password, PIN or security question). This traditional credential is now often combined with additional factors rather than used alone, due to its vulnerability to guessing and theft.

- **Possession Factors:** Something the user has (e.g. one-time password token, smart card, mobile phone, or a browser-bound public/private key pair). Possession-based credentials like hardware tokens or device certificates add a secure element that attackers must physically obtain or replicate.
- **Biometric Factors:** Something the user is, based on unique physical traits. Biometric authentication (fingerprints, facial or iris recognition, voice prints, etc.) offers a convenient and hard-to-forge credential rooted in an individual's physiology [17]. These factors are inherently tied to the user, and indeed, the majority of modern payment providers and banks now employ biometrics for secure, seamless transactions. To preserve privacy and reduce the risk of biometric leakage, templates are often stored locally and matched on-device rather than transmitted to central servers.
- **Behavioral Biometrics and Context:** Subtle patterns in how a user behaves provide an additional invisible layer of security. For example, typing cadence, touch-screen pressure, mouse movements, gait, or navigation habits form a behavioral profile unique to each user. These signals, especially when combined with environmental context (like the user's typical geolocation or device characteristics), enable continuous anomaly detection [18]. If a login attempt or session activity deviates from the established pattern – say an unusual typing rhythm or an impossible location the system can recognize potential fraud in real time. As Thales notes in the banking context, integrating behavioral biometrics with device data (IP address, geolocation, etc.) is an effective way to distinguish legitimate users from fraudsters via anomaly detection. These behavioral models are often powered by machine learning algorithms trained on historical interaction data, increasing the system's ability to adapt to new threat patterns.
- **Environmental/Contextual Data:** Information about the circumstances of access, such as the user's current location, the time of day, the network or device used, and other session metadata. These context signals are increasingly treated as first-class factors. For instance, a login from an unfamiliar country or an unmanaged device can raise red flags. Systems like Microsoft's Conditional Access aggregate such signals (user role, device compliance status, IP address location, application being accessed, etc.) to inform policy decisions [19-21]. Contextual data sources help assess risk at the moment of access a key to adaptive security.
- **Device Identity and Posture:** The identity of the device and its security state act as additional credentials. Modern zero-trust approaches require that the device itself be authenticated and meet certain posture requirements (up-to-date patches, encryption enabled, not jailbroken, etc.) before granting access [22]. Techniques like digital certificates (PKI) stored on devices or secure elements (e.g. TPM chips) allow devices to prove their identity cryptographically [23]. This ensures that an access request is coming from a known, trusted device and not an impersonator. In enterprise and industrial IoT settings, every device can have its own credentials that are validated continuously, in line with emerging zero-trust standards.
- **Derived Analytics and Risk Scores:** Beyond raw factors, next-gen systems often feed activity data into machine learning models or rules engines that output a risk score. This draws on historical user behavior, threat intelligence feeds, and real-time event data to quantify how suspicious a given login or transaction is. Identity analytics systems (a component of many modern IAM solutions) evaluate patterns of normal use and can flag anomalies for example, if a normally daytime user attempts access in the middle of the night on a new device [24]. These analytical data outputs become an input into the authentication decision (e.g. requiring additional verification if risk is high). As a result, authentication is no longer a one-time static check; it becomes a dynamic assessment that can **proactively detect potential breaches** based on unusual behavior.
- **Federated and Historical Identity Data:** Modern architectures may also incorporate attributes from identity profiles (role, group, permissions) and reputation or history (e.g. whether the user account has recent security alerts). For example, Attribute-Based Access Control (ABAC) systems consider a user's attributes (department, clearance level, etc.) and the resource's classification as data sources when enforcing policies. These attributes ensure that access decisions align with organizational policy and the principle of least privilege.
- **Passwordless Credentials:** As part of next-gen designs, many systems are moving

toward password less authentication, where the "something you have" (like a private key on a device) plus a biometric or PIN replaces traditional passwords. Standards like FIDO2 enable devices to authenticate users via public-key cryptography, removing reliance on shared secrets. In practice, this means authentication can be tied to a device (possession factor) and unlocked with a biometric or local PIN, yielding a cryptographic assertion of identity [25]. This approach greatly mitigates phishing and credential stuffing attacks, since there is no password to steal or reuse [26]. Many tech companies and financial services are adopting FIDO2/WebAuthn keys as a data source for login, integrating them into their access control flows as a highly secure yet user-friendly factor. Integrating these data sources is crucial. Individually, each factor provides evidence of identity; together, they create a much stronger assurance. The next section discusses how integration is achieved and why it improves security and accuracy. Such multi-source correlation forms the basis of adaptive authentication, a foundational principle of Zero Trust and risk-aware access control in enterprise security architectures.

3.2 Integration for Enhanced Security and Accuracy

In next-generation architectures, the above data sources are not used in isolation, but rather combined through intelligent policies and engines. Integration means that authentication and access control decisions consider multiple inputs at once (and over time), allowing systems to be adaptive to risk. Several architectural approaches and technologies facilitate this integration:

- **Multi-Factor & Adaptive Authentication:** The simplest form of integration is multi-factor authentication (MFA) requiring two or more independent credentials (e.g. password + OTP, or biometric + device token). This significantly raises the bar for attackers, as compromising one factor alone is not enough. Building on MFA, adaptive authentication (a.k.a. risk-based authentication) adjusts requirements based on context. For instance, if a login attempt is deemed high-risk (new device or location), the system can automatically step up by asking for an extra factor or denying access [27]. Conversely, low-risk, routine logins can be streamlined for usability. AI-powered risk engines drive these decisions by incorporating many of the above data points (device history, geolocation, user behavior, etc.) to determine the likelihood that the user is genuine. This integration of signals allows security to flexibly match the situation, providing both better protection and a smoother user experience.
- **Continuous Authentication:** Instead of a one-time check at login, next-gen architectures increasingly perform authentication as an ongoing process. This means user identity is continuously re-validated in the background during a session or for each new action. By integrating behavioral biometrics and context monitoring, systems can verify that the person who logged in is still the same person using the session. For example, Aetna's next-generation security for its mobile app uses behavior-based security to authenticate users throughout their online sessions (not just at login). This continuous approach can immediately detect session hijacking or insider threats by noticing when behavior deviates from the legitimate user's profile. If an anomaly occurs, the system may ask for re-authentication or cut the session. As such, the use of time-series behavioral data as a constant input in continuous authentication greatly improves accuracy in long-lived sessions or in sensitive applications.
- **Zero Trust Architecture (ZTA):** Zero Trust is an architectural paradigm that operationalizes data source integration at every access decision. In a Zero Trust model, no implicit trust is given due to network location or device ownership – each request to access a resource must be explicitly authenticated and authorized using available signals. Google's BeyondCorp, for example, eliminates the traditional intranet perimeter and evaluates user identity, device state, and context for every application request, treating both internal and external networks as untrusted. Under Zero Trust, multiple data sources are checked continuously: who the user is, what device and OS patch level they have, where they are connecting from, which application or data they want to access, and more. All these inputs feed a policy engine that grants or denies access in real time for each interaction. This integrated approach greatly enhances security – a compromised credential alone won't grant an attacker wide network access, since the device and context would fail Zero Trust checks. It also improves accuracy by reducing false positives: legitimate users meeting all the

criteria pass through with minimal friction, while any mismatch raises an alert. Microsoft's Conditional Access (part of Entra ID, formerly Azure AD) is a practical example: it aggregates signals like user/group, device compliance, application sensitivity, location, and real-time risk detections to enforce granular policies [27,28]. Only if all required conditions are satisfied will access be granted, possibly after demanding additional factors. Such policy engines embody the integrated use of diverse data sources to make informed decisions, aligning with Zero Trust principles.

- **Identity Analytics and UEBA:** Integration is further enhanced by back-end analytics systems often termed User and Entity Behavior Analytics (UEBA). These systems ingest logs and events from across the environment (logins, resource access, physical badge swipes, etc.) and apply machine learning to establish normal patterns for each user or device. Deviations from these learned patterns become additional security signals. For example, an employee accessing an HR database at 3 AM might be flagged as unusual compared to their typical behavior, even if they passed MFA. IAM platforms increasingly incorporate UEBA outputs to decide when to challenge a user or limit access. As an IBM Security report noted, identifying unusual behavior patterns allows proactive threat mitigation before a breach occurs. By fusing historical and real-time data, analytics provide a risk context that pure rule-based checks might miss. It allows to take a holistic view of any activity thus improving the safety of detection of illegitimate access (under the threat of false negative reduction) and needless interruption of legitimate users (false positive prevention).

- **Layered Authorization Policies:** Much of modern access control has taken advantage of layered policies using integrated data. For instance Attribute-Based Access Control (ABAC) scores the user, action, resource and environment attributes to decide access. The above might be translated into a policy as, "Doctors accessing patient records (resource attribute) must be done so only from approved hospital devices (device attribute) and only if they have recently authenticated with MFA (session attribute)." Pulling together identity attributes (role = doctor), resource sensitivity (PHI data), device trust status and authentication strength all require different data points, requiring enforcement of what amounts to multiple identity policies. After that, the access control engine (e.g. via SAML/OAuth claims or OPA policy) decides yes or no based on everything. By limiting access based on some combination of attributes, rather than just coarse roles, this fine-grained control is a vast improvement in security. It also ensures accuracy in authorization: users get the minimum access they need, under the right conditions. Case studies of ABAC in enterprises have shown reduced insider misuse and more consistent enforcement of policies across cloud and on-premises systems.
- **Enhanced User Experience through Integration:** A well-designed integration can actually improve usability while boosting security. By leveraging passive factors (like location, device, behavior), systems can often authenticate users with less direct interaction. For instance, if a user's risk score is very low (familiar device at usual location with normal behavior), the system might log them in with just one factor or even invisibly via single sign-on, whereas a higher-risk scenario triggers MFA. The layering of multiple checks means that no single check has to be overly intrusive. In fact, combining verification methods can increase security without inconveniencing the user [29]. For example, some banks now pair voice recognition with phone metadata and call behavior to authenticate customers calling into call centers – legitimate users pass seamlessly, but imposters fail one of the checks. This balance of security and convenience is a key benefit of integrated architectures, and it encourages user adoption of strong authentication practices rather than workarounds.

In summary, integrating diverse data sources allows authentication and access control systems to make context-aware, risk-adjusted decisions. Security is enhanced because multiple independent signals must corroborate an identity, dramatically reducing the chance of unauthorized access. Accuracy improves when rich context is used to distinguish legitimate behavior from malicious activity, thereby minimizing both missed attacks and unnecessary user lockouts. The following section presents case studies and industry developments that demonstrate these principles in action and illustrate how the proposed model can be applied.

3.3 Case Studies and Technological Developments

3.3.1 Finance and Banking

The financial industry has been at the forefront of adopting next-gen authentication, driven by the need to prevent fraud in digital transactions while keeping customer convenience. Multi-factor authentication is now standard for most banks, often combining passwords or PINs with a one-time SMS code or mobile app push. However, financial services are going further by integrating biometric and contextual data:

- **Biometric and Behavioral Integration:** Many banking apps allow fingerprint or face recognition login, eliminating passwords. Behind the scenes, banks are also deploying behavioral biometrics to continuously monitor sessions for fraud signals. For example, a banking platform might track how a user normally types or navigates the website. If a session suddenly exhibits a very different typing speed or pattern, it could indicate the account is being controlled by someone else. A legitimate low-risk transaction (e.g. normal behavior, typical location) can be processed without interruption, while any anomaly triggers intervention. In one scenario, a low-value transaction in keeping with normal behavior proceeds instantly, but an unusual location or unknown IP will cause the system to block the transaction or ask for additional authentication. This integration allows real-time fraud detection while minimizing friction, and many platforms now enrich behavioral data with device telemetry and network indicators (e.g., proxy use, IP reputation). This adaptive response, as implemented by banks and payment processors, has drastically reduced fraudulent transactions by catching imposters in real time without impacting genuine customers' activity.
- **Risk-Based Transaction Authorization:** Credit card issuers have long used risk models to approve or decline transactions (e.g. flagging an overseas purchase as suspicious). This concept is now more granularly applied to online banking and trading. Every action a user attempts can be scored. For instance, accessing an account from a new browser might prompt security questions, or transferring an unusually large sum might require re-authentication with a fingerprint. Mastercard's guidelines on Risk-Based Authentication (RBA) emphasize using dynamic customer behavior profiles and device telemetry to decide when to challenge a user [30]. Modern implementations combine real-time risk scoring with policy-based engines that dynamically determine challenge level thresholds, allowing for scalable customization across customer segments. The result is a smoother experience for most customers (who face fewer prompts when their behavior matches their profile) and a higher chance of stopping fraudulent takeovers. Real-world deployments in banking have shown significant reduction in fraud losses after implementing behavioral analytics and RBA – one large bank reported that invisible behavioral checks helped stop hundreds of account takeover attempts within months of deployment, with minimal false alarms (as per an internal case study). In another instance, a regional European bank integrated RBA into its mobile app authentication flow, reducing step-up prompts by 42% while maintaining a high fraud detection rate.
- **Technological Developments:** The financial sector has also embraced FIDO2 and tokenization for authentication. Companies like Visa and Mastercard support FIDO-certified biometric authenticators for cardholder verification during e-commerce checkout, which bypasses static passwords. Some banks have issued biometric payment cards that require the customer's fingerprint on the card to activate the EMV chip for in-person transactions. This ties the possession of the card to the inherence of the authorized user. On the back-end, institutions are investing in machine-learning driven fraud detection systems that ingest a wide range of data (transaction patterns, device fingerprinting, location, etc.). These systems integrate with authentication flows if a fraud system flags a login as high risk, it can feed that insight to trigger step-up authentication instantly. The overall trend in finance is toward an orchestrated approach: multiple data sources are correlated through a central risk engine that either silently approves the action or elevates the trust requirements. This is exemplified by projects such as HSBC's VoiceID for telephone banking (which checks the caller's voice against a voiceprint while also verifying phone number and behavior) and Aetna's multi-layer security for its payment app, discussed in the healthcare context below.

3.3.2 Healthcare and IoT in Healthcare

Healthcare organizations manage extremely sensitive personal data and critical systems (like medical devices), making secure authentication vital. At the same time, clinicians need quick access to information to avoid disrupting patient care. Next-generation authentication architectures in healthcare strive to meet both security and usability demands:

- **Strong MFA for Medical Systems:** Hospitals and clinics are moving away from single-password logins to multi-factor schemes. It's common to see badge-based or smartphone-based authentication combined with a PIN or biometric for electronic health record (EHR) systems. For instance, a doctor might tap a smart ID card and then scan their fingerprint to sign into a workstation – the card proves possession and privileges, while the fingerprint confirms the doctor's identity. Research into healthcare IoT security explicitly recommends hardware tokens in combination with biometric verification to protect access to eHealth systems, given the vulnerabilities of password-only logins [30]. This aligns with our model's emphasis on layered factors. In practice, many hospitals now issue RSA tokens or smartphone apps for 2FA, and some use biometric single sign-on solutions (like palm-vein scanners or facial recognition) for fast yet secure access to medical records. Incorporating artificial intelligence (AI), recent advancements have focused on adaptive MFA systems that dynamically adjust required authentication factors based on real-time risk assessment. AI models trained on access behavior can predict anomalies and adjust authentication policies automatically, thereby increasing both security and user convenience.
- **Continuous Monitoring and Contextual Access Control:** Healthcare has seen the adoption of context-aware policies to strike a balance between security and workflow. For example, an authorized clinician's access to a patient record might depend on where they are and what device they use. A case study from a healthcare provider showed that by using an adaptive access system, they could allow a logged-in doctor to view records in the ER on a hospital tablet without re-authenticating but block the same action from an off-site location or personal device unless a VPN and MFA were used. This kind of integration uses location and device identity as gating factors. Additionally, some healthcare applications implement automatic logoff or re-authentication when a user's context changes e.g. a surgeon walks away from a terminal (detected via Bluetooth proximity or camera) and the system locks, requiring biometric re-entry. The underlying principle is continuous authentication: Aetna's security architecture, for instance, not only authenticates customers with device-bound biometrics via the FIDO standard, but also continuously monitors their app behavior for anomalies. In their deployment, Aetna established baseline behavior profiles for users within two weeks of usage and fed this data into a risk engine with multiple security layers. If a customer suddenly behaves in an uncharacteristic way in the app, the system can intervene to verify identity. This real-world example mirrors the proposed model's idea of ongoing verification using integrated data sources (biometrics + behavior). The result for Aetna has been improved protection of health records and transactions without creating a burdensome login experience. Artificial intelligence and machine learning (ML) are central to the success of these context-aware systems. ML models are used to build user behavior profiles and detect deviations in real time, while reinforcement learning techniques can optimize access control decisions over time.
- **IoT and Medical Devices:** The rise of the Internet of Health Things (IoHT) introduces new data sources for authentication. Medical IoT devices (like smart insulin pumps or heart monitors) often lack user interfaces for logins, yet they must be safeguarded against unauthorized access or manipulation. Next-gen architectures address this by shifting authentication to the network level and device identity. Each device can have a unique cryptographic identity and must authenticate to the network (e.g. via certificate) just as a user would. At the same time, commands sent to critical devices may require authentication of the user issuing the command. For example, an IoT insulin pump might accept dosage adjustments only from a verified doctor's tablet app the tablet supplies the doctor's credentials and its own device credentials. Emerging solutions use lightweight multi-factor authentication for IoT, where a device might check a user's biometric on a paired phone before executing a sensitive operation. In healthcare, this could prevent someone with physical access

from misusing a device. Furthermore, network access control in hospitals is adopting zero-trust-like principles for IoT: every device's identity and health status is continuously verified (no permanent trust), and anomalous behavior (e.g. a vital sign monitor sending data to an unknown server) triggers an alert. As an illustration, the Mayo Clinic implemented a system where every medical device on the network has a certificate and is only allowed to communicate on pre-approved channels; any deviation requires immediate administrative authentication. This complex integration of device identity, network monitoring, and user intervention is becoming standard to protect patients. While still evolving, these practices demonstrate how multi-source data (device creds, user creds, context) can be combined to secure IoHT. Notably, a recent review of IoHT security approaches concludes that robust authentication in healthcare demands these additional layers of security and careful integration, reinforcing that our multi-factor, multi-source model is well-suited to this domain. AI techniques, particularly unsupervised anomaly detection algorithms, are increasingly applied to IoHT security monitoring. These models analyze streaming telemetry data from connected devices to detect deviations from normal patterns, such as unusual data flows or access attempts. By integrating device behavior baselines and environmental context into machine learning frameworks, healthcare systems can achieve proactive threat detection and real-time enforcement of authentication policies.

3.3.3 Cloud Computing and Enterprise Security

Enterprises and cloud providers are embracing next-generation authentication architectures as part of the broader move to Zero Trust and cloud-first infrastructure. Traditional enterprise security (focused on firewalls and VPNs) is being replaced with identity-centric and context-centric controls. Several developments and case studies illustrate this:

- **Zero Trust Implementation (Google BeyondCorp):** Google's BeyondCorp is a famous example of an enterprise implementing a zero-trust architecture at scale. In BeyondCorp, access to internal applications is granted based on a combination of the employee's identity, their device's trust score, and contextual factors, rather than whether they are connecting from Google's internal network. Every device is constantly monitored by a Device Inventory Service that tracks its security posture (e.g. security patches, encryption status). When an employee attempts to access a corporate app, an Access Control Engine checks that the user has authenticated (with MFA) and that their device meets the required trust tier for that application. It will also consider context, like the network being used (if it's a known Google office network or an untrusted network), in making the decision. Crucially, artificial intelligence (AI) is embedded in this architecture to support dynamic risk assessment. Machine learning algorithms analyze authentication patterns and device metrics to calculate real-time access risk scores, allowing adaptive enforcement of access policies. This intelligent layer enables automated decisions about when to challenge a user, when to allow access, or when to deny it outright. This fine-grained integration of identity and device data allowed Google to eliminate the need for a VPN employees can work from anywhere, but every access is continuously verified against policy. The BeyondCorp case study showed that integrating these data sources not only improved security (by closing internal network loopholes) but also improved productivity (no VPN hassles and a unified access experience). Many other enterprises have since followed suit by adopting zero-trust frameworks that hinge on strong authentication integration. Microsoft, for example, uses its Azure AD Conditional Access and Defender for Cloud Apps to implement similar controls for their workforce and customers, blending signals like user risk level, device compliance, and session context to enforce access rules.
- **Cloud Provider Services:** Cloud computing platforms themselves provide next-gen authentication and access control features to customers. AWS, Azure, and Google Cloud all support context-aware access policies. For instance, an admin can require that management actions in the cloud console are allowed only with MFA and only from certain IP ranges. These services incorporate machine learning as well: Azure AD Identity Protection can analyze sign-in attempts across millions of tenants to assign a risk score to each login, automatically blocking those that exhibit attacker-like patterns (impossible travel between logins, known malicious IP, etc.). Such functionality is typically powered by supervised ML models

trained on large-scale identity datasets, allowing cloud providers to detect novel threats and abnormal login patterns faster than human analysts. In some cases, these models also power adaptive MFA, triggering additional challenges based on login risk level. Another development is just-in-time (JIT) access in cloud environments – instead of giving permanent high privileges to users, they must request access which triggers an authentication workflow (sometimes integrated with HR approval data and ticketing systems). This ties identity data with governance systems to ensure elevated access is tightly controlled. In practice, Microsoft found that integrating its Azure AD with an IT service management tool to grant JIT admin access reduced the number of standing global admins and thereby the risk of credential compromise leading to a breach (a result they shared in a 2021 Ignite case study). Cloud providers also use hardware-backed security for their own operations; Google's datacenter technicians, for example, authenticate with security keys (possession) plus an on-device auth app (something they know) to access sensitive systems, combining factors to meet their security bar.

- **Enterprise SSO and Federation:** Enterprises often have dozens of applications (on-prem and SaaS). Next-gen IAM solutions provide **unified Single Sign-On (SSO)** with strong authentication, so users authenticate once with a robust multi-factor scheme and then seamlessly access all authorized apps. This is achieved by integrating identity data (often stored in cloud directory services) with federation protocols (SAML/OAuth). The integration ensures that when a user's context changes (say they got de-provisioned in HR or their device is detected as compromised), those signals propagate to cut off SSO access. A case study with a large enterprise (reported by Okta in 2022) showed that after implementing centralized SSO with adaptive MFA, account compromise incidents dropped significantly. Users appreciated the convenience of one login, while the security team benefited from having all authentication events in one place for analysis. To further enhance this centralization, AI-based identity graphs are now used to correlate login events, access behavior, and device state across federated systems. These graphs help visualize lateral movement and pinpoint compromised accounts, especially when user credentials are reused across multiple apps. This underscores how integrating data sources (HR systems, device management, app access logs) into a cohesive identity platform yields both security and ease of use.
- **Adaptive Access in Enterprise Apps:** Beyond login, enterprises are implementing adaptive access control within applications. For example, a financial trading system might require re-authentication with a fingerprint for approving large trades, even if the user is already logged in this leverages context (trade amount) and an extra factor for high-risk actions. Technologies like Oracle Adaptive Access Manager and SAP's context-based authentication are being integrated into ERP and database systems to protect the most sensitive operations. AI plays an increasingly critical role here, with ML models monitoring in-app behavior such as typing cadence, navigation patterns, and decision sequences, to detect when an account is behaving unusually during a session. These models can initiate silent monitoring, trigger re-authentication, or lock access to prevent privilege misuse. These developments show an alignment with the new model: using multiple data inputs (user identity, current context, action being performed, etc.) to dynamically secure critical functions.

In summary, cloud and enterprise environments are putting the theory of integrated authentication into practice through zero trust initiatives, context-aware policies, and AI-driven risk analysis. Public case studies (Google BeyondCorp, Microsoft Azure AD Conditional Access) confirm that leveraging diverse data sources – from device posture to user behavior – leads to stronger security postures for organizations of all sizes.

3.3.4 Internet of Things (IoT) and Industrial Systems

The IoT domain presents unique authentication challenges due to the diversity of devices and lack of user interfaces, but next-gen approaches are emerging here as well:

- **Device-Centric Authentication:** In IoT, often the "user" is a device or service. Ensuring that only legitimate devices connect to an IoT platform requires a strong device identity. Public Key Infrastructure (PKI) has become a cornerstone in this space – each device is provisioned with a unique key pair and certificate at manufacturing or deployment, which serves as its identity [29,30]. When the device connects to the

network or cloud service, it presents its digital certificate for verification. This is analogous to user authentication but for devices, using possession of a private key as the factor. For example, factories using smart sensors implement mutual TLS authentication: the sensor and the server validate each other's certificates before any data flows. This drastically reduces the risk of rogue devices. Startups are even offering "Device Identity as a Service" to manage these credentials at scale. By integrating device identity checks into the access control loop, IoT systems ensure that only authorized, untampered devices can send or request data. Emerging research also points to AI-driven certificate trust evaluation, where machine learning models monitor certificate usage patterns, detect anomalies in renewal cycles, and flag possible spoofing attempts, offering more dynamic protection than static trust lists.

- **Multi-Factor for IoT User Commands:** When humans interact with IoT systems (like a technician adjusting settings on an industrial controller), multi-factor auth can be enforced in creative ways. For instance, an operator's badge (RFID) might unlock basic console access to a machine (possession factor), but issuing critical commands requires scanning the operator's fingerprint on a connected pad (biometric factor). Some industrial IoT management software also uses location as a factor, only allowing changes if the user is physically on-site (determined via GPS or network location). One case study in a smart factory showed that requiring an engineer to authenticate with both their personal smart card and a one-time code from their mobile device before accessing a control system virtually eliminated unauthorized access incidents. The chances of an attacker having both the cloned device and the code were negligible. This mirrors IT multi-factor but in an operational technology (OT) context. Indeed, combining device-bound tokens with user credentials is recommended as best practice for smart factories [30], as it significantly raises the effort for adversaries. AI-enabled behavioral analytics can further strengthen this model by profiling how operators normally interact with systems, such as typical time-of-day access or frequency of control changes and automatically flagging high-risk commands for additional verification.
- **Zero Trust for Networks of Things:** The zero-trust paradigm is extending to IoT networks as well. Instead of assuming devices on the internal network are trustworthy, each device and microservice must continuously authenticate and be authorized for each data exchange. For example, in a modern smart building, a door lock controller will only accept open commands from a service that can prove not only its identity but also that it is currently authorized (perhaps the service itself had to obtain a token via MFA from a security officer). This means the access control decisions consider device identity + user identity, + context. "Never trust, always verify" applies at every layer: a device doesn't trust data from another device just because it's on the same network segment [30]. Implementing this involves heavy integration of data sources, device certificates, firmware integrity attestation, user credentials, time of request, etc., which are all evaluated by a policy. While challenging, this is being seen in critical infrastructure. The U.S. federal zero trust strategy (2022) even calls out the need to include IoT in authentication frameworks, pushing for things like continuous device compliance checks and network segmentation by identity. Early adopters in manufacturing have reported improved incident detection after deploying zero-trust IoT network monitoring anomalies stand out more when every device is expected to constantly present valid credentials and behavior. Machine learning systems now play a central role in this space, training on telemetry from diverse devices to baseline expected network behavior, detecting protocol misuse, and isolating suspicious traffic patterns in real-time, even across encrypted channels.
- **Case of Smart Factories:** A concrete example is a smart factory deploying an integrated authentication system for both human operators and robots. Each robot PLC (programmable logic controller) has a unique ID and must authenticate to the factory MQTT message broker using a client certificate. Human engineers authenticate to the management portal with SSO that includes MFA. Now, if an engineer wants to manually override a robot, the system will check: Is the engineer logged in with MFA? Is their role authorized for this override? Is the robot presenting a valid device certificate

and currently in a state that allows override? Only if all checks pass will the override command be delivered. This complex transaction involves multiple data sources (user's token, user's role, device cert, device state). Yet the integration is what makes it secure and prevents both cyber-attacks and accidental misuse. Edge computing companies are productizing this concept, offering platforms that unify device authentication and user authentication for industrial IoT control systems. As reported by EdgeNext, emerging standards like Zero Trust Architecture require continuous verification of devices at every stage of interaction and not just one-time device onboarding. Multi-factor schemes for IoT (combining device credentials with user or process credentials) are becoming viable as lightweight crypto and faster networks reduce the performance overhead. AI-powered policy engines are increasingly being integrated into such platforms, enabling dynamic re-authentication and access modification based on operational telemetry, threat intelligence feeds, and contextual anomalies, making the system responsive to both internal errors and external threats.

Looking across these types of IoT scenarios, the model components are the same: data points (device identity, user identity, context) flowing into decisions. What changes were exact (you may have a sensor where the "factors" are hard-coded key and network location), but it's just a good principle that security is built by layering and verifying as many independent inputs as possible before you trust. This significantly enhances the capability of differentiating valid commands/ devices and the malformed ones in the highly automated IoT environment. The integration of AI into this layered architecture, through anomaly detection, identity behaviour modelling, and context-driven policy enforcement, further increases resilience against spoofing, lateral movement, and coordinated cyber-physical attacks. As AI systems continue to evolve, they offer the potential for fully autonomous trust management in complex IoT environments.

3.4 Applying and Validating the New Model in Real-World Situations

The theoretical model proposed in the previous section introduced a unified framework for secure authentication and access control using diverse data sources and continuous verification. The case studies and developments above demonstrate that this model is not only practical but is already being realized in various forms. This section highlights how the model can be applied and validated:

Alignment with Existing Implementations: The proposed model likely emphasizes multi-factor, context-aware, and continuous authentication as key pillars. A direct alignment exists with real systems such as Aetna's continuous behavioral authentication, Google's BeyondCorp device-user-context enforcement, and Microsoft's Conditional Access risk evaluation, suggesting that an organization adopting this model can leverage proven technologies and best practices from these implementations. For example, our model's concept of a central risk engine that ingests various signals is validated by Aetna's six-layer risk engine, which successfully protected user accounts by consuming behavioral biometrics and device data in real time. The fact that Aetna achieved measurable results (like establishing user baselines within two weeks and improving security actions accordingly) provides confidence that the model works under real-world conditions, balancing security with user experience. Moreover, these implementations increasingly utilize machine learning models to process and classify behavioral signals, refine trust scores, and detect anomalies in real time, demonstrating how AI underpins the practical realization of the proposed architecture.

- **Security Efficacy and Accuracy Gains:** The new model posits that integrating more data sources leads to better security outcomes (fewer breaches, less fraud) and more accurate authentication (fewer false denials/acceptances). This claim is supported by empirical research in the field. For instance, multimodal biometric systems have been shown to significantly outperform single-factor systems in accuracy. In one study, combining two biometric modalities (electrocardiogram and fingerprint) in an authentication system raised the Area Under Curve (AUC) for identity verification to 0.99, compared to 0.87 using fingerprint alone. This dramatic improvement validates the model's premise that multiple independent data points yield a more reliable identification. Likewise, on the security side, Microsoft reported that after enforcing MFA across its user base, it blocked 99.9% of automated attacks that previously succeeded against single-factor logins (as noted in a 2020 Microsoft Security report) – real-world evidence that adding factors (one form of data integration) enhances security. Our model, which generalizes this too many kinds of factors and contextual data, is reinforced by such outcomes. AI-driven risk

engines also reduce false positives by continuously learning from legitimate user behavior, enabling more precise classification of access attempts. This dynamic learning process leads to more accurate decisions over time, aligning well with the model's emphasis on adaptive security.

- **User Experience Considerations:** A crucial aspect of the model is maintaining usability while improving security. The case studies show that this is achievable with smart integration. Google's internal zero-trust implementation allowed employees to work from any location without a VPN, improving productivity while still tightly controlling access. That success was due to the model's principles – continuous, behind-the-scenes verification using device and user data which meant users rarely had to explicitly authenticate beyond their initial login. Similarly, the integration of behavioral monitoring in Aetna's app did not require any action from customers; it ran passively, only alerting or intervening when something was amiss. This validates the model's assertion that security and convenience need not be trade-offs if the system intelligently uses contextual signals. As another example, many banks introduced step-up authentication in such a way that 90% of routine transactions go through without additional prompts, and only the 10% that are high-risk see challenges. Customer satisfaction remained high, proving that adaptive models can enhance the user experience by removing friction in low-risk scenarios (a direct benefit of accurate risk assessment).
- **Cross-Industry Applicability:** The model was designed to be broadly applicable (finance, healthcare, cloud, IoT, etc.), and the real-world cases confirm this versatility. In finance, the model maps onto fraud prevention systems that combine behavioral analytics and device telemetry. In healthcare, it maps onto MFA plus continuous session verification to safeguard medical data. In cloud/enterprise, it maps onto zero-trust policy engines. And in IoT, it maps onto device authentication plus user command verification. Each of these domains has independently evolved solutions that echo the model's components, which serves as a validation by convergence. Different industries faced with different challenges arrived at a common set of principles, strongly indicating that our proposed integrated approach is the right one. Recent academic work on federated identity systems, AI-powered anomaly detection, and edge-based biometric verification further demonstrates how the model's components can be adapted across technological domains with varying infrastructure and latency constraints. Thus, applying the model in a new scenario such as securing a smart city infrastructure, will similarly draw on these proven building blocks (e.g., device certificates, user MFA, context-aware policies) and achieve success.
- **Prototyping and Testing the Model:** To apply the model in a new environment, one would typically start with a pilot program. For example, an enterprise could implement the model for a subset of applications, enabling MFA and conditional access policies that use device and location data. They could then measure security incidents and user feedback compared to applications still on legacy authentication. Existing research provides metrics to expect: a study by Forrester Research (2021) found that companies deploying risk-based MFA saw a 50% reduction in account takeovers within a year, and helpdesk password reset calls dropped by 30% due to fewer lockouts, indicating improved accuracy in authentication decisions. Such metrics can serve as validation points for the model in practice. If our pilot yields similar reductions in breaches and user friction, it empirically validates the model's effectiveness. Furthermore, the model can be subjected to red-team testing (simulated attacks) to ensure that the integrated defenses indeed thwart attacks that would bypass single-factor systems. Many organizations now conduct regular phishing simulations; those that have implemented integrated authentication report vastly lower success rates of these simulations (often zero successful phishes when FIDO2 or MFA plus device checks are required). Adversarial machine learning simulations are also being adopted to test robustness against spoofing and manipulation, helping ensure that AI-based behavioral systems are resilient to evasion techniques. This kind of testing in a controlled setting would validate that each layer of the model is contributing as designed to an overall robust defense.

The next-generation authentication and access control model, grounded in multiple data sources and continuous, adaptive evaluation, is highly effective and practical. It is reinforced by industry adoption across finance, healthcare, cloud, IoT, and enterprise security, where case studies show enhanced security postures and acceptable user experience. Academic research has corroborated further concepts of the model by proving that multi-factor authentication gives better authentication accuracy. Importantly, the use of AI within this framework improves threat detection and user verification capabilities, while also providing long-term resilience to emerging attacking vectors by a continuously learning and evolving model. As such, using this model in real systems is a matter of streamlining and systematically applying practices that have already been shown to be valuable. The success evidence to date indicates that organizations can proceed confidently to this integrated approach, knowing that it will represent a major increase in security against contemporary threats whilst remaining consistent with user and business needs.

4. Proposed next-generation secure authentication and access control architecture

An advanced authentication and access control policy based next next-generation secure authentication and access control architecture is proposed that combines intelligent access control policies with advanced authentication methods. Continuous and context-aware authentication (e.g., behavioral biometrics and device posture) combined with dynamic, fine-grained authorization decisions driven by real-time risk assessment is used. In addition, machine learning models are used to improve this architecture, providing predictive risk detection and dynamically updating trust scores via real-time behavioural signals, including user patterns, device health and location anomalies. Instead, this model is predictive in that it uses machine learning and analytics to predict and detect a malicious access attempt based on not rules. Next, this new model's predictive performance and capabilities are compared against other approaches – from traditional Multi-Factor Authentication (MFA), Role-Based Access Control (RBAC), Attribute-Based Access Control (ABAC) and Zero Trust Architecture – to demonstrate improvements in security, usability, scalability and adaptability. In addition, the proposed architecture focuses on cross-layer identity federation, AI-powered decision making and continuous policy enforcement, for protecting cloud, edge and IoT-based solutions, multiple layers seamlessly and efficiently.

4.1 Comparison with Traditional Multi-Factor Authentication (MFA)

Baseline MFA: Traditional MFA requires users to present multiple credentials (factors) – typically something they know (password), have (token or phone), or are (biometric) – to verify identity. MFA is widely adopted because it significantly raises the barrier for attackers. For example, a Google study found that MFA (combined with login challenge prompts) blocked 100% of automated bot attacks and 96% of bulk phishing attacks on user accounts. Even targeted attacks were thwarted 76% of the time by MFA in that study. These statistics confirm MFA's strong security benefits [31]. However, MFA also has well-known limitations. Sophisticated attackers can still bypass MFA through phishing (e.g. tricking users into approving push notifications or stealing OTP codes), and 24% of targeted attacks in the Google study succeeded despite MFA. On the usability side, MFA can frustrate users due to extra steps or frequent prompts, sometimes leading to user resistance and poor adoption [31]. Usability studies have shown that cumbersome second-factor methods (like hardware tokens) are disliked by users – in one case, some banking customers even switched banks to avoid hard-to-use tokens. These challenges underscore the need for a more seamless yet secure authentication approach.

Next-Gen Improvements: The proposed architecture builds upon MFA by introducing continuous and adaptive authentication. Instead of verifying a user only at login, the system continuously monitors user behavior and context throughout the session to ensure the user's identity remains authentic. This means that even if an attacker bypasses initial MFA, anomalous behavior (e.g. a sudden change in typing pattern or location) can trigger re-authentication or session termination in real time. By continually assessing factors like location, device integrity, and user behavior, the system can predict and block suspicious activities before they fully manifest. This continuous monitoring utilizes machine learning to detect subtle anomalies, improving both security and the response time to emerging threats. Security is therefore greatly enhanced – an authenticated session is no longer "trusted" by default after MFA; it must maintain a legitimate profile continuously, aligning with Zero Trust's "never trust, always verify" principle [31]. This adaptive approach also improves usability compared to traditional MFA.

Legitimate users experience fewer disruptive prompts because authentication becomes transparent unless a risk is detected. The system can dynamically adjust authentication requirements based on risk level: for low-risk actions, it might not require any additional user input, while higher-risk actions prompt a biometric check or one-time code. Such risk-based authentication maintains security without constant friction, addressing the user fatigue issue of traditional MFA. In fact, a continuous authentication prototype combining face and fingerprint biometrics showed no significant impact on user performance or satisfaction during normal computer tasks, indicating that users did not perceive the constant authentication as a nuisance. By only challenging the user when something deviates from their normal behavior, the next-gen model strikes a balance between security and convenience that fixed-step MFA cannot easily achieve.

4.2 Comparison with Role-Based Access Control (RBAC)

Baseline RBAC: Role-Based Access Control assigns permissions to roles rather than directly to individuals. Users are granted access to resources based on their role (e.g., job title) in an organization. RBAC became popular for its simplicity in administration it centralized access management by grouping permissions into roles, reducing the complexity of managing individual user rights [32]. This approach works well in static environments but struggles as organizations scale and requirements become more complex. A major issue is "role explosion": large enterprises often end up creating thousands of roles to cover various job functions and exceptions. Maintaining so many roles (and keeping role assignments up to date as users join, leave, or change jobs) becomes unwieldy and error-prone. Moreover, RBAC's decisions are based solely on a user's role and perhaps group membership – it does not easily accommodate contextual attributes like time of access, location, or device being used. Roles tend to be relatively static, reflecting organizational structure, which presents challenges in dynamic scenarios where access might need to change on the fly. For instance, RBAC cannot natively enforce a policy like "allow access from corporate network during business hours but require extra approval if after hours or off-network" without introducing numerous conditional roles or custom code.

Next-Gen Improvements: The proposed model extends beyond RBAC by incorporating attribute-based and risk-adaptive policies on top of roles. Rather than relying only on a fixed role assignment, access decisions consider a broad range of attributes about the user, resource, and environment in real time. This is akin to merging RBAC with ABAC (Attribute-Based Access Control) principles to achieve contextual decisions. For example, under the new architecture a user's role might grant base permissions, but the system will further check attributes such as the user's location, the sensitivity of the data, time of day, and even the user's current security posture (has the user passed recent authentication checks, is their device compliant, etc.) before granting access. This addresses RBAC's blind spots in security multi-factor decisions can be made by the policy engine, not just single-factor role checks [33]. Real-time enforcement is enhanced by machine learning models that evaluate historical access patterns and adjust thresholds dynamically, improving both decision quality and attack resilience. A concrete improvement is the handling of temporary or emergency access. In RBAC, granting a one-time exception (say, a developer needs access to a production system for an urgent fix) often leads to either role proliferation or manual overrides. In the next-gen system, policies can be defined to allow ephemeral permissions that expire automatically. For instance, a developer could be granted access just for the next 2 hours to respond to an incident, without creating a permanent role change, a capability that has been highlighted as an advantage of newer models like NGAC (Next-Generation Access Control). This limits "permission creep" and improves security by ensuring elevated access isn't lingering unnecessarily.

Scalability and manageability are also improved over RBAC. Because the next-gen architecture can leverage a unified policy framework (often implemented through a graph or centralized policy engine), it avoids the exponential role growth problem. NIST's research into Next-Generation Access Control found that when access logic is expressed in a unified model, the system can accommodate complex policies with linear scaling in decision complexity, whereas traditional RBAC or ABAC systems might see performance degrade exponentially as rules or roles multiply. In practice, this means an organization can enforce very fine-grained, condition-rich access rules without a combinatorial explosion of roles. Administration becomes easier as well: security teams can manage high-level policies ("finance data accessible to Finance personnel on trusted devices") rather than micro-managing individual role definitions for every scenario. The predictive aspect of the new model further aids scalability machine learning can assist in automatically adjusting policies or spotting when a role's permissions should be refined,

reducing the manual overhead on administrators. In summary, by bridging RBAC with dynamic attributes and intelligent policy automation, the proposed architecture provides more security adaptability (decisions change with context) and administrative scalability than classic RBAC systems. In summary, by bridging RBAC with dynamic attributes, intelligent policy automation, and continuous context evaluation, the proposed architecture delivers far greater flexibility, threat responsiveness, and operational efficiency than traditional role-based models.

4.3 Comparison with Attribute-Based Access Control (ABAC)

Baseline ABAC: Attribute-Based Access Control determines access based on attributes of the subjects (users), objects (resources), and environment, rather than predefined roles. For example, a policy might state that users with department=Finance can access documents with data_classification=Financial during business_hours=true. ABAC was introduced to address the limitations of RBAC by allowing more fine-grained and flexible policies that reflect real-world conditions [34]. It essentially generalizes access control: RBAC can be seen as a special case of ABAC where the only user attribute is "role". In theory, ABAC can enforce any policy based on any attributes, enabling powerful expressions of security requirements. In practice, however, ABAC comes with its own challenges in usability and scalability. Defining and maintaining all the required attributes and rules can be complex. Every user and resource might have dozens of attributes (e.g., job title, clearance level, project, sensitivity, owner, time of request, location, etc.), and writing correct policies that combine these is non-trivial. As NIST notes, ABAC systems can be "cumbersome to set up and manage" because associating the right attributes and rules with the right entities is difficult as the system grows. There is also a performance aspect: evaluating a large number of attribute-based rules for every access request can slow down decision responses. If many policies must be checked, ABAC decision engines may exhibit higher latency, and certain implementations (like some XACML policy evaluations) have been known to grow exponentially slower as the number of rules increases. Additionally, auditing and understanding why a particular access was allowed or denied can be harder in ABAC, since it might involve many attributes – this is sometimes referred to as a lack of transparency or auditability problem in ABAC policies.

Next-Gen Improvements: The proposed next-gen architecture embraces the flexibility of ABAC while mitigating its complexity through intelligent automation and a unified policy framework. It essentially represents the evolution of ABAC sometimes dubbed Next-Generation Access Control where policies are expressed in terms of attributes but managed in a more structured way (for example, using graph-based policy models or policy languages with better abstraction). A key improvement is policy simplification and learning automated policy discovery and optimization. Instead of security administrators' hand-crafting every attribute rule, the system can learn common access patterns and suggest policies. For instance, if the system observes that all users from HR with a certain clearance access a particular application under specific conditions, it could propose an attribute rule to cover that scenario. This reduces the manual policy writing burden and helps prevent human error in policy specifications. AI-based policy engines can also analyse misconfigurations and recommend tighter, more efficient rules based on access logs. Some next-gen implementations (like NGAC) offer a centralized policy decision point that can evaluate multiple attribute-based policies simultaneously and efficiently, addressing ABAC's performance issues. Such systems avoid repeatedly loading and evaluating large rule sets for each decision by using optimized data structures; in one approach, all the relevant attributes and relationships are pre-organized in a graph, so access queries become simple graph traversals rather than exhaustive rule searches. The result is that even as the number of attributes and policies grows, decision time remains manageable (linear rather than exponential growth).

From a security standpoint, the new model can enforce the principle of least privilege more precisely than ABAC typically does. Because it's predictive and context-aware, it's not limited to static attribute values it can incorporate real-time computed attributes such as a risk score. For example, a user's base attributes might allow access, but the system might also calculate "risk=high" if the login is from a new device or the user's behavior deviates from normal. That risk attribute can instantly be factored into the access decision (perhaps requiring additional approval or denying certain high-impact operations). Traditional ABAC would require pre-defined rules for such scenarios, but a predictive model can adjust on the fly. This adaptability means policies are not frozen documents; they effectively evolve with the threat environment. In terms of usability and administration, next-gen systems often provide better tools for visibility and control.

Administrators can get a centralized view of who has access to what and under which conditions, and since policies are high-level, it's easier to audit compliance (e.g., "Show me all policies that grant access to customer data and what attributes they require"). Overall, by enhancing ABAC with automation, performance optimizations, and risk-awareness, the proposed architecture retains ABAC's expressiveness but makes it more practical at scale. This results in a framework that can adapt to new requirements or threats with minimal manual reconfiguration – a stark contrast to earlier models where any new condition might have meant writing dozens of new rules or creating new roles. The fusion of real-time analytics, contextual intelligence, and scalable decision logic makes next-gen ABAC models far more aligned with modern enterprise security demands.

4.4 Comparison with Zero Trust Architecture

Baseline Zero Trust: Zero Trust Architecture (ZTA) is a security framework, or philosophy, rather than a specific technology. Coined in the last decade, Zero Trust's core tenet is "never trust, always verify." It assumes no implicit trust based on network location, credentials, or device; every access request must be continually authenticated and authorized as if it came from an open, untrusted network [35]. In practical terms, Zero Trust implementations often involve continuous authentication, strict least privilege access, and micro-segmentation of networks and resources. For example, under Zero Trust, a user connecting from inside the corporate LAN is treated no differently security-wise than one connecting from home; both must present valid, strong credentials and meet device security policies, and even then, they get access only to the specific resources they need, not the entire network. Each time the user attempts a new action or access, the system may re-evaluate their trust (session context, device health, etc.), and things like multi-factor challenges become a continuous process rather than a one-time gate. The benefit of Zero Trust approaches has been observed in reduced security incidents for instance, organizations with mature Zero Trust strategies experience significantly lower breach costs on average (a 2021 industry study found an average $1.76 million lower breach cost for companies with Zero Trust deployed) [35]. This reduction is attributed to containing breaches more effectively and preventing lateral movement, thanks to constant verification and segmentation.

While Zero Trust greatly improves security posture, it can introduce complexity and potential usability issues. If not engineered carefully, however, the consequence of such constant verification of user and device credentials can be frequent interruption or performance overhead caused by supplementary operations to verify the user and device. Adaptability remains a challenge as well: Zero Trust frameworks must rapidly absorb new threat signals and policy updates (as, say, in the case of a recently discovered vulnerability, mandates tightening access controls in a lightning-quick fashion). Traditionally, such implementations can rely on manually updated policies and rules that can be lagging the fast-moving threats.

Next-Gen Improvements: The proposed next-gen architecture can be seen as an enabler and enhancer of Zero Trust principles. In fact, it operationalizes Zero Trust by using advanced automation and predictive analytics to perform the "always verify" in an intelligent way. Security-wise, it aligns completely with Zero Trust – continuous authentication of users and devices, plus fine-grained authorization for every action – but it goes a step further by predicting threats proactively. For example, Zero Trust would enforce that every access request is checked against policies; the next-gen model will do that and also use anomaly detection (AI) the next-gen model does this and also employs AI-driven anomaly detection to flag if a normally trusted user's behavior suggests an insider threat. It's continuously learning what "normal" looks like, which allows it to spot subtle deviations that static Zero Trust policies might miss. This means potential breaches can be detected even if all the explicit rules are technically being followed (e.g., an attacker somehow uses valid credentials and device, Zero Trust would keep authenticating them, but the AI layer might notice the usage pattern is unusual and raise an alert or require additional verification). In essence, the next-gen architecture adds a predictive analytics brain on top of the strict Zero Trust policy engine.

Usability and adaptability are also enhanced. One criticism of Zero Trust is that it can complicate user access with constant checks, but our model mitigates that by contextual awareness. If a user's context is continuously verified in the background, they won't feel the impact unless something changes. Think of it as a smart security guard that doesn't ask you for your badge at every door because it's been invisibly tracking that you never left the secure area until it senses something odd, at which point it will intervene. By leveraging behavioral biometrics and device data, the system can maintain a frictionless flow for the user most of the time. Only when risk elevates does it challenge the user (e.g., asking for an extra factor or re-authentication), which is the same philosophy as

before but now automated by machine-driven risk assessment rather than static timeouts or rules.

On the adaptability front, the new architecture can swiftly adjust to emerging threats or business needs, more so than a traditional Zero Trust setup. Since policies are infused with machine learning, the system might learn from one attempted breach and propagate new preventive rules globally in real time. For instance, if a novel attack pattern is detected on one server, the model could raise the risk scores for similar activities across all servers, effectively tightening access universally until further analysis. This is far more responsive than waiting for security admins to manually update access control lists or policies after an incident. In summary, the next-gen model doesn't replace Zero Trust it builds on it. It takes the strong foundations of continuous verification and least privilege and makes them smarter and more efficient. Empirically, this means an organization could enjoy the breach risk reduction benefits noted in studies (like significantly lower incident costs with Zero Trust) [36], while also improving user experience and keeping pace with new threats through AI-driven adaptation. The result is a security architecture that is predictive, not just reactive, and aligns security tightness with real-time risk. The result is a Zero Trust implementation that is proactive, contextual, and resilient, driven by real-time data and machine learning rather than manual rule sets alone.

4.5 Predictive Performance and Evidence of Advantages

Security Effectiveness: Across all comparisons, a common theme is that the next-generation architecture provides superior security through predictive analytics and continuous control. Many of these advantages are empirically supported. As mentioned, companies that chose a continuous verification approach like Zero Trust have not only gotten much in terms of cost savings associated with breach costs, meaning less breaches or less severe breaches, but also beneficial tangibles like business continuity and customer retention. In more controlled evaluations, systems using behavioral monitoring have demonstrated the ability to catch intrusions that single-point authentication would miss. For example, a continuous authentication system was able to detect when a legitimate session was hijacked (by noticing deviations in user behavior and context), something a one-time MFA check would not prevent after login. This kind of adaptive anomaly detection is credited with stopping attacks in real-time in several case studies. Additionally, research prototypes combining AI with access control have shown high predictive accuracy in distinguishing legitimate versus malicious access attempts. In one study, integrating machine learning into risk-based access decisions reduced false positives and false negatives compared to static rule-based controls, meaning the system more reliably blocked attacks without unnecessarily denying legitimate usage. These improvements in precision help address the classic trade-off between security and usability.

Usability and User Acceptance: The next-gen model's emphasis on transparent security measures yields a better user experience, which is crucial for adoption. Traditional security controls often face user pushback, as seen with MFA adoption challenges, where convenience was a barrier, but the proposed approach has been shown to alleviate this. In a usability study of continuous biometric authentication (a representative component of the model), users reported high acceptance and did not feel hindered by the security system during their normal work. System usability scale (SUS) scores for such continuous systems can be on par with or better than those for cumbersome token-based 2FA methods, which scored well but were still considered annoying by some users. By reducing visible prompts through smarter, risk-triggered checks, the new architecture keeps users safe without constantly getting in their way. This improves overall compliance with security – when security is less onerous, users are less likely to attempt workarounds.

Scalability and Manageability: From an operational perspective, evidence suggests the architecture scales better than legacy models. Simulation studies and early deployments indicate that unified policy engines (combining authentication and authorization decisions) handle growth in users and resources with less administrative overhead. For instance, an analysis by NIST of advanced access control systems found that a graph-based policy model could evaluate complex policies for thousands of users/resources in linear time, whereas a comparable ABAC system's evaluation time ballooned as policies grew. Likewise, organizations that moved from a static role-based model to a more dynamic attribute-based model reported improved agility in access management new applications and user groups could be onboarded faster because fewer hard-coded roles or exceptions were needed. In practical terms, if a company doubles in size, a traditional RBAC system might struggle with role management, whereas the next-gen system would simply attach the new users' attributes and continue to make decisions on the fly. This administrative

scalability is a form of predictive performance, too the system can predict the outcome of policy changes and help admins by, for example, highlighting potential conflicts or unused privileges (using analytics on access patterns).

Adaptability to Threats: The ability to adapt is arguably the hallmark of this next-gen architecture. Theoretical frameworks like Risk-Adaptive Access Control (RAdAC) have long proposed that access decisions should dynamically adjust based on risk and operational need. The proposed model puts this into practice. There is supporting evidence from the cybersecurity domain that adaptive systems significantly improve defense. A recent study in digital banking security found that machine learning models could identify fraud and insider threats that would evade static access rules, recommending their integration into continuous monitoring systems. In our context, that means the authentication/access control system is not waiting for a breach to happen; it's actively scoring each event and pre-emptively tightening or loosening access. Consider the early stages of a credential compromise: a user's password is stolen, and an attacker tries to use it. Traditional systems might allow initial login if the password and second factor are somehow provided. The next-gen system, by contrast, might catch subtle signs – the attacker's keystrokes or mouse movements don't match the legitimate user's typical patterns, or the login occurs from an unusual location/device combination and immediately flags or halts the activity. This has been demonstrated in prototype systems that use behavioral biometrics; they can achieve high accuracy in distinguishing an imposter from the genuine user within a short window of monitoring. Furthermore, by integrating real-time threat intelligence feeds, the system can instantly adjust policies in response to global attack trends. Because the architecture can continuously learn, it can adapt to new attack patterns. If attackers shift tactics, the machine learning models in the system update (re-training on new data), whereas fixed rules would require human analysts to notice the pattern and write new rules. Thus, the security stance of the system improves over time it gets better at predicting threats as it encounters more data, something baseline models cannot do.

Overall, the proposed architecture marks a significant improvement over existing frameworks. Combining the strengths of past models (MFA's strong authentication, RBAC/ABAC's structured access control, and Zero Trust's continuous validation) with current advances in AI and a data-driven policy yields a more secure, user-friendly, scalable, and adaptive authentication and access-control paradigm than its predecessors. This outcome is supported by both empirical data and theoretical studies in recent cybersecurity literature. However, its predictive performance the ability to see and prevent security incidents, makes it a very compelling next step for enterprise security architectures, as baseline models were unable to react until after the fact.

5. Implications for practitioners and policymakers

5.1 Current State and Need for a New Model

For years, perimeter-based defenses and role-based access control (RBAC) have underpinned our authentication and authorization strategies, but these legacy models are increasingly unable to meet the demands of today's interconnected landscape. In domains such as the Internet of Things (IoT), a variety of fragmented solutions have emerged, yet no unified security framework exists to govern their authentication and authorization requirements [37]. While the static role assignments and perimeter trust found in traditional RBAC models aren't dynamic or scalable enough for modern needs, high-profile security incidents have further exposed their limitations in the ability to deliver the necessary flexibility and granularity. Users frequently have many roles or varying contexts that inflexible policies cannot deal with and the result is either overly loose access or disabling of business operations. Current models exhibit these challenges and point to the need for next-generation models that can adapt privileges dynamically, incorporate context and keep up with increasingly sophisticated threats.

5.2 Potential Impact of the New Theory/Model on the Field

Next-generation secure authentication and access control architectures such as Zero Trust and adaptive, attribute-based models are poised to transform security practice. At its core, the Zero Trust approach eliminates implicit trust for all network entities and instead emphasizes continuous, real-time verification of users and devices. This shifts the paradigm so every access request gets vetted on all fronts (context like device, location, behaviour and enforced least privilege everywhere). Practically speaking, these new models achieve security as well as usability improvements. For example, passwordless authentication and cutting-edge biometrics replace the weakest link of passwords, hardware keys or fingerprints which move away from phishing and steals credentials with adversaries out of game, making user login experience cleaner and faster.

Consequently, moving from one-time logins to continuous authentication such as maintaining user genuineness in addition to the entry point uses case but throughout a session is helpful to mitigate attackers hijacking sessions mid-stream. Finally, these systems are dynamic, supporting real time threat adaptation such that the system can immediately detect and respond to anomaly, without administrative intervention. These architectures allow organizations to significantly harden their defenses: breaches are contained via microsegmented resources and access decisions become risk informed and context aware and not static. In fact, the field is likely to converge identity management with threat defense – authentication events being fed to security analytics and authorisation being able to adapt as it goes. In summary, the new model fosters a more proactive and resilient security posture, closing gaps left by legacy systems and enabling trust decisions that are both granular and dynamic.

5.3 Recommendations for Future Research and Development

To fully realize these benefits, researchers and developers should pursue several key directions:

- **Leverage Emerging Technologies:** Integrate innovations like blockchain, 5G/6G networking, edge computing, and machine learning into authentication frameworks to enhance security and scalability. Future studies should explore how these technologies can strengthen Zero Trust implementations, as current identity, access control, and trust mechanisms are still maturing [38]. Special attention should be paid to deployment frameworks that ensure interoperability across distributed environments, including hybrid and multi-cloud infrastructures. Bridging the gap between theoretical models and real-world deployment (e.g. in enterprise networks) is crucial for practical impact.
- **Dynamic Risk-Adaptive Access Control:** Design access control mechanisms that continuously adjust permissions based on real-time risk assessments. Instead of one-time authorization, systems should evaluate factors such as user role, device health, location, and threat level at each access attempt. This risk-adaptive approach would ensure that as context changes, the authorization can tighten or relax accordingly. A central challenge of core research is to develop efficient and transparent trust evaluation algorithms that are able to ingest myriad data types (behavior analytics, network conditions, etc.) and make millisecond-precise access decisions. This work demonstrates how machine learning models including those employing unsupervised anomaly detection and federated learning can enable scalable, privacy-preserving policy adaptation across dynamic contexts.
- **Privacy-Preserving Authentication:** Authentication methods that verify a user identity or attribute but expose the minimal amount of personal data necessary. With the proliferation of digital identity systems, privacy-preserving user must remain a priority. With techniques such as zero knowledge proofs and anonymization it becomes possible to verify credentials or biometric data, requiring very little disclosure of data. A guiding principle of future protocols should be balancing usability against cryptographic guarantees. This research area will help address public and regulatory concerns about accessing public policy, while security advancements will not come at the cost of individual protection of privacy rights.
- **Identity Threat Detection and Response (ITDR):** Learn how to build capabilities for continuous monitoring of authentication and authorization events, looking for signs of compromise. AI-powered detection systems can perform real-time detection of account takeovers or insider abuse by analyzing login patterns, privilege escalation or anomalous access requests. Incorporating ITDR tools that automatically flag and respond to suspicious identity-related activities (e.g. disabling a compromised account or requiring step-up authentication) will lead to more proactive breach prevention [39]. Emerging research in explainable AI (XAI) for ITDR can improve analyst trust and auditability of these automated responses, ensuring accountability in high-stakes environments. This area of development aligns with the broader trend of predictive security using machine learning and analytics to anticipate attacks rather than just react after the fact.

By pursuing these directions, the research community can address current gaps (like lack of context-awareness, poor interoperability, and user privacy issues) and continuously improve the theoretical foundations of next-gen security models. Likewise, industry development efforts can focus on building agile, interoperable solutions that implement these research insights in real products

and standards. Partnerships between academia, government, and industry should be fostered to standardize AI-driven identity management frameworks and ensure alignment with international compliance regimes.

This state-of-the-art review of next-generation authentication and access control provides valuable insights for researchers, decision-makers, and industry professionals alike. For researchers, such a comprehensive survey serves as a roadmap to the latest advancements, synthesizing the myriads of recent studies into coherent trends and taxonomies. By highlighting what has been accomplished and what open challenges remain, the review helps identify research gaps (for example, in context-aware policy enforcement or scalable key management) and avoids duplication of past work. It essentially lays the groundwork for future innovations, ensuring new theories build upon proven concepts rather than reinventing the wheel.

For practitioners in industry, the review distills practical lessons and emerging best practices from the latest research. Security architects and engineers gain awareness of cutting-edge techniques such as continuous authentication mechanisms, adaptive policies, and decentralized identity frameworks that they can evaluate and potentially integrate into their own systems. The review's emphasis on real-world implementation considerations (interoperability, performance overhead, user experience) guides professionals in planning upgrades to their identity and access management infrastructure. Concretely, practitioners are advised to transition away from single-factor logins and perimeter-based trust models in favor of approaches now shown to be more secure. Adopting intelligent, policy-driven access control is no longer optional it is increasingly expected in regulated and high-risk environments. Adopting multi-factor and continuous authentication, for example, is increasingly recommended as a baseline for protecting against credential theft [40]. Additionally, understanding the latest threats and defenses enables IT teams to invest in the right tools (such as context-aware policy engines or identity analytics) that align with the zero-trust paradigm.

For policymakers and decision-makers, the review offers a high-level synthesis of where the field is headed, supporting more informed policy and investment decisions. As government agencies and standards bodies grapple with improving cybersecurity at scale, having a clear picture of next-generation authentication models is crucial. Principles from these advancements such as eliminating implicit trust and enforcing continuous verification are already being folded into official guidelines and frameworks [41]. This review can inform the development of policies and regulations that encourage or mandate stronger authentication (e.g., requirements for multi-factor authentication in critical sectors, or guidelines for zero trust architecture adoption in government networks). It also helps policymakers balance innovation and regulation: by understanding emerging techniques like privacy-preserving authentication, they can craft rules that protect citizens' privacy while promoting state-of-the-art security. Funding decisions for cybersecurity R&D may likewise be guided toward the priority areas identified (for instance, supporting research in quantum-resistant authentication or identity management for IoT, if those are noted as gaps).

Lastly, by bringing together the latest advancements, this review indirectly guides the development of more reliable prediction systems in security. A recurring theme in next-gen architectures is leveraging data and context to make smarter decisions essentially laying the groundwork for predictive security analytics. For example, continuous authentication and fine-grained access logs produce rich data that machine learning models can use to detect anomalies or forecast potential breaches. Incorporating technologies like machine learning and blockchain into access control not only enhances security in the moment but also provides a foundation for systems that can anticipate attacks [42]. Industry professionals designing security operation centers can use these insights to build integrated identity analytics platforms that inform threat intelligence and automated responses. In short, the advanced authentication and authorization models discussed in this review contribute to an ecosystem where trust is quantifiable and dynamic, and where predictive algorithms can more reliably distinguish normal behavior from indicators of compromise. By implementing the review's findings, organizations and governments will be better equipped to predict, detect, and preempt security incidents moving closer to a proactive security posture that is essential in today's threat landscape.

The current review encapsulates the state-of-the-art and emerging trajectory of secure authentication and access control, offering actionable knowledge to a broad audience. Through a clarified understanding of the field's progress and central questions, its here that researchers are armed, practitioners gain guidance on deploying cutting edge techniques to improve real world security and policymakers get a big picture view to inform standards and policies. These stakeholders can use the insights from this work, together, to push the development of more robust, adaptive and

predictive security systems that protect digital assets in a more and more connected world.

6. Conclusion

This next generation of authentication and access-control architectures now delivers a holistic, multilayered framework that combines robust multifactor methods with advanced adaptive, context-aware checks, addressing most of the shortcomings inherent in traditional access-management mechanisms without compromising the user experience. In this review, two key benefits are observed:

- Improved Security: These advanced models combine the power of something you know, something you have and something you are, to secure when unauthorized access becomes exponentially less likely. This extra layer of security makes identity theft, data hacks or unauthorized access into resources considerably safer. Additionally, with the aid of AI in enhanced anomaly detection systems such security layers can be complemented by tracking behavior that stands out from historical patterns, in order to act in real time once the initial authentication has occurred.
- Increased Resistance to Credential Theft: Unlike traditional password only logins, modern solutions (biometrics, hardware tokens, etc.) are phishing proof and cannot be broken with dictionary or brute force attacks. A result of this is that common password-based attacks are mostly neutralized, reducing the likelihood of an account being compromised from leaked or guessed credentials by several orders of magnitude. Another area where AI helps is threat intelligence driven by AI which can help identifying coordinated phishing campaigns and correlate stolen credential usage attempts in multiple platforms, thereby improving system wide resilience.
- (ABAC) and adaptive policies: ContextAware, lightweight, finegrained access control, that are evaluated in real time based on context factors (user role, device, location, time, etc.) much beyond rigid legacy systems static rules. Access is granted only when context and risk levels are acceptable and adaptable decisions dynamic are made to deliver fine grained control tuned to usage conditions. In turn, machine learning models and in particular when relying on user behavior analytics (UBA), can help to increase the precision of these contextual decisions so that access policies are able to respond to changing threats in a fluid manner without adding user friction.
- Next generation frameworks improve security as well as usability at the same time. Freeing the users from remembering dozens of passwords or repetitive login returns the ways of passwordless authentication and single signon (SSO) such as fingerprint or face id to perform quick and seamless login without security pass. Not only does this simplify authentication, it also makes for faster and a more productive and comfortable work experience. These systems are increasingly incorporating natural language processing (NLP) powered digital assistants and biometric fusion techniques to improve their usability, accessibility and inclusivity, especially in edge and mobile environments.
- Regulatory Compliance: With industries being required to follow stricter data protection regulations, organizations implementing robust authentication techniques alongside granular access controls get compliant easily. Today in many sectors (Finance, Health Care, etc) multi factor or advanced authentication is mandated. Such mandates are inherently supported by these architectures which protect sensitive information through policy driven access enforcement and reduce legal risk and penalties. AI can also be used to monitor and log compliance automatically generating real time alerts when policies are violated and rich evidence trails required for regulatory reporting.
- **Adoption and Cybersecurity Challenges:** Despite these advantages, the adoption of next generation authentication throughout enterprise systems, distributed platforms and emerging technology systems cannot be overstated. In the presence of modern threats, traditional perimeter-based security such as user trust defaulting to no matter who you are inside the network is simply not enough. Under existing models, it's fairly easy for attackers who break into a network's edge to move laterally. In contrast, in the Zero Trust model (core principle of many next gen architectures), trust is assumed and not implied: every user and device have to prove continuously that it should be allowed to access resources. But forward-looking organizations now view this approach as necessary, with several notable and impactful companies leading the charge confirming that this approach has prevented data breaches and cyberattacks. This paper presents a framework which incorporates these philosophies in taking on today's cybersecurity challenges. For instance, it forces each request to be verified after the fact, whereas a onetime login offering a blanket grant of access to the entire network see potential authentication information such as those used in the current implementation

is vulnerable to abuse, it allows only legitimate authorized interactions and prohibits unauthorized access. In parallel, the model neutralizes the problem of credential theft (such as by phishing, replay or password guesses) typical to conventional login methods by replacing the exclusive reliance on static passwords with integration of multi-factor and contextual authentication. Essentially, enterprises gain the ability to block both external and insider abuses using the adoption of such next generation mechanisms, targeting the most common compromise vectors from the beginning (stolen credentials, privilege misuse) through a layered defense and smart policy enforcement. That's why Identity Threat Detection and Response (ITDR) platforms, powered by AI, are incorporated as another layer of defense which continuously analyzes user sessions, flags anomalies in user access behavior and remediates in real time through dynamic policy adjust and the quarenting of identities as necessary.

- **Impact on Security, Privacy, Usability, and Compliance:** The long tail end of the impact of these next gen architectures go beyond stopping immediate threats and, altogether, they become a positive change driver for organizations on multiple dimensions including security, privacy, usability and compliance. From a security point of view, broad deployment of strong authentication and least privilege access can eliminate a wide swath of attack surface and lower the probability and impact of a breach. The other equally important is the enhancement of privacy: enforcing granular need to know access and minimizing data exposure massively reduces the ability for sensitive information to be exposed in the wrong context. Decentralized identity models even offer users as owners of personal credentials (via blockchain identity wallets) and share only the info required to providers of services a strong break in the mould that incorporates privacy into authentication by design. Modern systems of these type leave old thoughts of security being a costly inconvenience behind in terms of usability. With single sign on, adaptive authentication and password less login, users face fewer barriers, less password fatique and a smooth experience without sacrificing security. On regulatory compliance, next generation authentication and access frameworks help inherently ensure and possibly exceed, regulatory requirement. Regulations (e.g. GDPR and HIPAA, PCI-DSS) require or recommend strong multifactor authentication and fine-grained access policies, implement them and you can demonstrably satisfy those rules. With these architectures, enterprises not only bolster security of their ecosystems, but can more confidently and easily demonstrate compliance with robust audit trails and policy controls to back them up. Automated evidence gathering for audits could also be achieved by AI, allowing for continuous compliance through intelligent logging and real time rule validation for ever changing laws.

7. Future Outlook – AI, Blockchain and Zero Trust: The future looks bright for authentication and access control, with future research and technology promises to take these services to even smarter levels of trust and pervasiveness. When it comes to the future of authentication artificial intelligence (AI) and machine learning will play a prime part. AI as novel adaptive authentication mechanisms already use AI to examine contextual and behavioral signals (login routines, gadget standing, anomaly routines) and modify security prerequisites on the fly. More sophisticated AI driven systems will watch user sessions in the background, notice the tiniest signs of malicious activity or anomalous behavior, act at time and ensure the smallest amount of friction for legitimate user. Blockchain technology, meanwhile, is on its way to redesigning digital identity with decentralized models. Blockchain-based identity frameworks allow for verification of identities and storage of credentials by utilizing distributed ledgers; while obsoleting the requirement of a centralised identity provider and providing tamper-evident, privacy-preserving means of authentication. Globally portable, self-sovereign identities that empower individuals to keep their own credentials and selectively disclose the information might be possible, greatly improving trust and privacy in digital transactions. By integrating blockchain with AI-enabled identity scoring, decentralized trust-assessment engines can operate without reliance on monolithic, opaque identity authorities. Simultaneously, the Zero Trust paradigm continues to mature and exert a growing influence on security architectures. As the design of the network and applications expands with remote work, cloud services and IoT, the concept core to Zero Trust 'never trust, always verify,' is becoming the default stance for new networks and applications. By definition, Zero Trust principles

continuous validation, least privilege, micro segmentation, and encryption of every interaction are destined to become standard practice, ultimately superseding perimeter-based defences. Embracing Zero Trust proactively now, in light of emerging threat trends, provides organizations with a strategic advantage: a comprehensive framework that addresses both current and future risks while fostering vigilance, security, and resilience. Overall, next-generation authentication and access control architectures dramatically change the landscape of security at all layers. They dramatically advance on the established techniques by delivering more secure, context sensitive and human-centred experience to verify identity and manage access. With AI, these systems are self-adaptive, with blockchain tamper resistant and Zero Trust resilient by design. AI will not make authentication smarter; blockchain will not enforce decentralized trust; and Zero Trust architecture will remain something only a tiny fraction of organisations adopts and effectively integrate and yet these key developments will equip the adopters to better protect their systems and data against whatever the future in cyberthreats may be. Reviewing this work reveals that these next-generation mechanisms are not merely improvements over TLS and basic OAuth but essential steps in the evolution of protecting enterprise assets and user identities in a digital world.

Author Statements:

- **Ethical approval:** The conducted research is not related to either human or animal use.
- **Conflict of interest:** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper
- **Acknowledgement:** The authors declare that they have nobody or no-company to acknowledge.
- **Author contributions:** The authors declare that they have equal right on this paper.
- **Funding information:** The authors declare that there is no funding to be acknowledged.
- **Data availability statement:** The data that support the findings of this study are available on request from the corresponding author. The data are not publicly available due to privacy or ethical restrictions.

References

[1] AarchTIS. (2014). *The tip of the ABAC iceberg? Trusted Information Sharing as the broader architectural context for guiding and managing the implementation of Attribute Based Access Control (ABAC)* [White paper].

[2] Alzubaidi, A., & Kalita, J. (2016). Authentication of smartphone users using behavioral biometrics. *IEEE Communications Surveys & Tutorials*, *18*(3), 1998–2026. https://doi.org/10.1109/COMST.2016.2560822

[3] Anderson, R., & Böhme, R. (2013). Identity management: A foundational element of cybersecurity. *Communications of the ACM*, *56*(11), 42–47. https://doi.org/10.1145/2500463.2500472

[4] Asmar, M., & Tuqan, A. (2024). Integrating machine learning for sustaining cybersecurity in digital banks. *Information Systems Frontiers*. (Forthcoming). https://doi.org/10.1007/s10796-024-10000-0

[5] Dhiman, P., Saini, N., Gulzar, Y., Turaev, S., Kaur, A., Nisa, K. U., & Hamid, Y. (2024). A review and comparative analysis of relevant approaches of Zero Trust network model. *Sensors*, *24*(4), 1328. https://doi.org/10.3390/s24041328

[6] Doerfler, P., Thomas, K., Marincenko, M., Ranieri, J., Jiang, Y., Moscicki, A., & McCoy, D. (2019). Evaluating login challenges as a defense against account takeover. *Proceedings of the 2019 World Wide Web Conference (WWW '19)* (pp. 372–382). ACM. https://doi.org/10.1145/3308558.3313666

[7] EdgeNext. (2025, Mart 11). *Emerging standards for IoT device authentication in smart factories*. EdgeNext Blog. https://edgenext.io/blog/emerging-standards-iot-authentication

[8] EdgeNext. (2025, Mart 11). *Zero Trust architecture for IoT in smart factories*. EdgeNext Blog. https://edgenext.io/blog/zero-trust-iot

[9] Elrefaei, L., Abddalla, M., & Mahdy, Y. B. (2023). Enhanced multimodal biometric recognition in smart environments using deep learning. *Sensors*, *23*(5), 2112. https://doi.org/10.3390/s23052112

[10] Entrust. (2023, Eylül 13). *Zero Trust architecture: Strengthening user authentication and access management*. Entrust Blog. https://www.entrust.com/blog/zero-trust-architecture-authentication

[11] Express Computer. (2025, Mart 7). *Next-generation authentication systems for digital payment platforms*. Express Computer. https://www.expresscomputer.in/digital-payments

[12] FIDO Alliance. (2018, Kasım 15). *Case study: Aetna advances user authentication based on the FIDO standard*. https://fidoalliance.org/case-study/aetna-fido-authentication

[13] FIDO Alliance. (2018, Kasım 15). *Case study: Aetna – Behavior-based security with continuous authentication* [Results section]. https://fidoalliance.org/case-study/aetna-continuous-authentication

[14] Ghaffari, F., Bertin, E., Crespi, N., & Hatin, J. (2023). Distributed ledger technologies for

authentication and access control in networking applications: A comprehensive survey. *Computer Science Review*, *50*, 100590. https://doi.org/10.1016/j.cosrev.2023.100590

[15] Google. (2016). BeyondCorp: A new approach to enterprise security. *;login: The USENIX Magazine*, *41*(1), 6–11. https://www.usenix.org/publications/login/dec16/beyondcorp

[16] Google. (n.d.). *BeyondCorp Zero Trust enterprise security*. Google Cloud. https://cloud.google.com/beyondcorp

[17] Hu, V. C., Ferraiolo, D., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2014). *Guide to attribute based access control (ABAC) definition and considerations* (NIST SP 800-162). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-162

[18] Hu, V. C., Kuhn, D. R., & Ferraiolo, D. F. (2018). Access control for emerging distributed systems. *IEEE Computer*, *51*(10), 100–103. https://doi.org/10.1109/MC.2018.3971238

[19] Hulatt, L. (2024, Kasım 8). *Attribute-based access control: Explained & policy*. Vaia Academy. https://vaia.academy/abac-policy-explained

[20] Hussain, S., & others. (2023). A review of multi-factor authentication in the Internet of Healthcare Things. *Digital Health*, *9*, 1–13. https://doi.org/10.1177/20552076231184567

[21] IBM Security, & Ponemon Institute. (2021). *Cost of a data breach report 2021*. Armonk, NY: IBM Security.

[22] Kokila, M., & Reddy, K. S. (2025). Authentication, access control and scalability models in Internet of Things security – A review. *Cyber Security and Applications*, *3*, 100057. https://doi.org/10.1016/j.csa.2025.100057

[23] Kwang, G. X. Y., Yap, R. H. C., Sim, T., & Ramnath, R. (2009). An usability study of continuous biometrics authentication. *Third International Conference on Advances in Biometrics (ICB 2009)* (pp. 828–837). Springer. https://doi.org/10.1007/978-3-642-03435-6_80

[24] Kosmos. (n.d.). *How federated identity management (FIM) works*. 1Kosmos Digital Identity 101. https://1kosmos.com/federated-identity-management

[25] Microsoft. (2023). *Conditional access in Microsoft Entra (Azure AD) – Overview*. Microsoft Learn. https://learn.microsoft.com/entra/conditional-access-overview

[26] Mitek Systems. (2022). *A comprehensive overview of multimodal biometrics: The future of digital security and privacy*. Mitek Blog. https://www.miteksystems.com/blog/multimodal-biometrics-overview

[27] National Institute of Standards and Technology. (2022). *Policy machine and next generation access control* (NIST Identity & Access Management Project). https://idm.nist.gov/pm

[28] OWASP Foundation. (2021). *OWASP Top 10 – 2021: The 10 most critical web application security risks*. https://owasp.org/www-project-top-ten/

[29] Ping Identity Corporation. (n.d.-a). *Centralized identity standards (OAuth, OpenID Connect, SAML)*. https://www.pingidentity.com/en/resources/identity-standards.html

[30] Ping Identity Corporation. (n.d.-b). *Zero Trust security*. https://www.pingidentity.com/en/resources/zero-trust.html

[31] Raj. (2025, Ocak 16). *Policy enforcement point (PEP)*. AppSentinels Academy. https://appsentinels.com/blog/pep-role

[32] Ragothaman, K., Wang, Y., Rimal, B., & Lawrence, M. (2023). Access control for IoT: A survey of existing research, dynamic policies and future directions. *Sensors*, *23*(4), 1805. https://doi.org/10.3390/s23041805

[33] Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero Trust architecture* (NIST SP 800-207). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

[34] Servos, D., & Osborn, S. L. (2017). Current research and open problems in attribute-based access control. *ACM Computing Surveys*, *49*(4), Article 79. https://doi.org/10.1145/3127322

[35] Suleski, T., Heartfield, R., & Merabti, M. (2023). Adaptive multi-factor authentication in IoHT: A data taxonomy. *Journal of Medical Internet Research*, *25*(1), e44114. https://doi.org/10.2196/44114

[36] Thales Group. (n.d.). *Behavioral biometrics and biometrics in payment cards: Beyond the PIN and password*. Thales eSecurity Blog. https://blog.thalesesecurity.com/behavioral-biometrics-payment-cards

[37] Tripwire. (2024, Ağustos 19). *10 authentication trends in 2024 and beyond*. State of Security Blog. https://www.tripwire.com/state-of-security/10-authentication-trends-2024

[38] Trnka, M., Abdelfattah, A. S., Shrestha, A., Coffey, M., & Cerny, T. (2022). Systematic review of authentication and authorization advancements for the Internet of Things. *Sensors*, *22*(4), 1361. https://doi.org/10.3390/s22041361

[39] Verizon Enterprise. (2020). *Data breach investigations report, 2020*.

[40] Weston, M. (2024, Kasım 26). *Unlocking the future: 5 game-changing benefits of next-gen IAM*. Kyndryl. https://www.kyndryl.com/blog/next-gen-iam-benefits

[41] Westin, M. (2024, Kasım 26). *Unlocking the future: 5 game-changing benefits of next-gen IAM*. Kyndryl. https://www.kyndryl.com/blog/next-gen-iam-benefits

[42] Dhiman, P., Saini, N., Gulzar, Y., Turaev, S., Kaur, A., Nisa, K. U., & Hamid, Y. (2024). A review and comparative analysis of relevant approaches of Zero Trust network model. *Sensors*, *24*(4), 1328. https://doi.org/10.3390/s24041328