

A Zero Trust Security Framework for Cloud-Based Enterprise Systems in Hybrid Work Environments

Kishan Raj Bellala

Independent Researcher Austin, Texas, U.S.A.

Publication Date: 2022/08/28

Abstract

The cloud-based enterprise applications, and the work-from-home model adopted by organizations today have significantly expanded the security perimeter of an organization. A layered Zero Trust security framework is proposed for continuous identity verification, multi-factor authentication, device/workload trust assessment, least-privilege access, micro-segmentation, application/API security, data security, and monitoring. The framework proposes to incorporate a contextual risk assessment based on user identity, device posture, resource sensitivity, session context, and security-related telemetry. An Enterprise Use Case has been described to highlight how Zero Trust security framework can help in restricting access, minimizing excessive privilege, restricting lateral movement, and enhanced security visibility in a cloud-hybrid environment. The paper also provides a comparison of Perimeter Security models, existing Zero Trust models, and the proposed Zero Trust security framework. The future research areas have also been outlined which include behavioral analytics, intelligent threat detection, risk assessment and adaptive security mechanisms.

Keywords: *Zero Trust Architecture; Cloud Security; Hybrid Work; Identity and Access Management; Multi-Factor Authentication; Least Privilege; Micro-Segmentation; Risk-Based Access Control; Continuous Monitoring; Policy Enforcement.*

I. INTRODUCTION

Cloud computing, enterprise applications, and hybrid work models are developing at breakneck speed and are revolutionizing the way in which companies access and manage information resources. This develops in a distributed environment where employees, freelancers, applications and services access the company resources from the company office, their home, their personal networks, their mobile devices and from cloud platforms. In such an environment, classical perimeter-based security models are less effective (Barker, 2020). These models grant more access to users and computers within a trusted network than to users and computers outside the network. However, as the enterprise resources are moving to cloud and hybrid platforms, organizations need to adjust their security strategies to grant access to information resources based on the identity, context and security state of the requestor rather than the location of the network from where the request is originated (Barker, 2020).

Zero Trust Architecture (ZTA) is a paradigm that, unlike traditional models, eliminates implicit trust and grants access to enterprise resources and applications only after verification has occurred. The core principles of Zero Trust include identity-based access, multi-factor authentication (MFA), least-privilege access, secure devices, micro-segmentation, monitoring, and securing applications and data. In a hybrid work environment, users could be connecting to enterprise-sensitive resources and applications from multiple locations, using a variety of devices, and accessing cloud-based services, APIs, and distributed systems. In order to implement Zero Trust in these scenarios, the individual security controls must function as a single architecture as opposed to standalone products (Buck, 2021).

We propose a multi-layered Zero Trust framework that can be leveraged to secure cloud-based enterprise systems operating within a hybrid work environment. The framework for securely delivering services encompasses identity and access management (IAM), device and workload trust, least-privilege access (LPA), micro-

segmentation, application and API security, data protection, continuous monitoring, and contextual risk assessment. The core characteristic of the proposed architecture is the continuous evaluation of security-related context which in turn influences access policies that adapt based on the evolving security-related context of user identity, device state, data classified at different levels, session-related parameters and the corresponding security-related telemetry. The framework has been presented, its various layers have been detailed, and it has been illustrated with an example enterprise. Comparison has also been made with perimeter-based security approaches as well as existing Zero Trust approaches. This framework serves as the necessary architectural backbone for enhancing the adaptability of enterprise security and thus paves the way for numerous research opportunities in the domains of behavioral analytics, threat detection and more importantly, risk assessment and corresponding adaptive security mechanisms (Chandramouli, 2019) (Buck, 2021).

II. ZERO TRUST SECURITY: CONCEPTS AND PRINCIPLES

Zero Trust Security is an approach to enterprise security that assumes no user, device, application or network connection within or outside of the organization's perimeter can ever be trusted. All access requests must be authenticated and authorized prior to access to resources being granted. This is in stark contrast to traditional security models that presume users both inside and outside of the organization's physical perimeter can be trusted due in large part to their physical location. As such, Zero Trust applies security across multiple dimensions including user identity, devices, applications and resources (Buck, 2021). The core philosophy of Zero Trust is based upon Identity-centric security that authenticates and authorizes users and services based upon identity-based access management, multi-factor authentication and authorizations granted to users and services. An additional core philosophy is that of Least Privilege that grants users and services only the minimum set of privileges required to complete tasks and meet objectives on an ongoing basis. All users and services function with restricted rights in order to limit potential damage should an account be compromised or access be granted to an individual or service that is not authorized to gain access. The concept of Continuous Verification extends the notion of verification beyond simple login credentials and authentications, establishing a continuous and dynamic process that monitors and determines whether access to information, applications, and resources should be continued, retracted or modified based on changing circumstances including device posture, user behavior and other security related variables on an ongoing basis. Zero Trust establishes a paradigm within which no amount of trust is considered permanent, and all aspects of access and security are determined by Identity, Authorizations and Circumstances related to the requested resource and related Variables and Context (Chandramouli R. &, 2020).

III. SECURITY CHALLENGES IN CLOUD-BASED HYBRID WORK ENVIRONMENTS

With more cloud-based enterprise applications and a hybrid work model, security has become more complex and distributed. Now, users, endpoints, applications, APIs, and information may reside within a company's network as well as in the cloud and outside of corporate boundaries. In order to grant access to resources within a company, Zero Trust architectures have to deal with a number of challenges when it comes to security requirements in order to grant controlled and continually monitored access (Agency, 2021).

➤ *Remote Users*

With hybrid work, employees and even contractors access enterprise resources from their homes, public hot spots and from all over the world. Because they are outside the traditional security perimeter of the organization, relying on the network location for determining trust fails. Instead, access for users and their connections must be verified based on identity and security context (Mell, 2011).

➤ *Managed and Unmanaged Devices*

Organizations managing laptops and mobile devices for their employees may find that some or all of those employees also use personal and unmanaged devices. Since these devices can vary greatly in terms of how they have been secured, their patch levels, how they are protected from malware, and their level of compliance, a single valid identity for a user is not sufficient to grant access to enterprise resources and applications (OWASP, 2019).

➤ *SaaS and Cloud Applications*

Organizations rely on a vast array of cloud-hosted applications, and in many cases, are using Software as a Service (SaaS) platform to run key parts of their organizations. By virtue of these applications being hosted in the cloud, many of these applications reside across a variety of different cloud computing environments and cloud service providers (CSPs). Ensuring proper authentication, authorization, monitoring and protection of data across these various platforms proves to be very challenging (Ward, 2014).

➤ *APIs and Distributed Resources*

As the resources of an organization are connected by APIs, access to applications, services, databases and more needs to be authenticated, authorized and monitored. Poor API security and inadequate access rights for APIs put the organization's sensitive resources at risk of attacks and extend the attack surface (Rose, 2020).

➤ *Lateral Movement and Excessive Privileges*

Once an account or device has been compromised by an unauthorized user, that user can move between applications, systems and network segments. This movement is further exacerbated by excessive user privileges, allowing the user to access resources well beyond what is required to perform his/her duties.

Applying least privilege and micro-segmentation to resources on the enterprise can help to mitigate the effects of a compromised identity and limit the amount of unauthorized movement through the enterprise (Rose, 2020).

IV. PROPOSED LAYERED ZERO TRUST SECURITY FRAMEWORK

Zero Trust Architecture (ZTA) does not assume trusted access to enterprise resources based on IP address, hence all users, whether inside or outside the organization's networks, must verify their identity, validate the health of their device or application, meet proper authorization, assess the appropriate level of access to resources based on sensitivity and monitor activity to detect any Security Incidents. The security architecture for Zero Trust is outlined by NIST SP 800-207, published in September 2020 (Rose et al., 2020). This paper extends NIST's Zero Trust Architecture (ZTA) and introduces a novel layered security architecture, specifically designed for the use in hybrid work environments for so called cloud-based enterprise systems and applications, while individual security measures are not novel. Instead, their integration in a structured fashion into a combined and unified architecture of all the single measures constitutes the core contribution of this paper (Rose, 2020) (Agency., 2021).

The proposed security process is represented as:

Request → Verify → Assess Context → Determine Risk
→ Apply Policy → Monitor → Reassess → Enforce

➤ Overall Architecture

The proposed framework consists of seven primary security layers:

- Identity and Authentication Layer
- Device and Workload Trust Layer
- Least-Privilege Authorization Layer
- Application and API Security Layer
- Micro-Segmentation and Network Security Layer
- Data Protection Layer
- Continuous Monitoring and Contextual Risk Assessment Layer

As shown in the below figure, the Policy Decision and Enforcement function spans all the layers. This function collects security information from the various layers. Based on the collected information, it decides whether to allow, restrict, challenge or deny access requests to resources based on the organization's policies. This is how Policy Decision and Enforcement functions are realized as per the Policy Decision and Enforcement concepts in NIST's Zero Trust Architecture (Rose et al., 2020).

All individual layers interact with each other. Identity and device information are used for the authorization decision on the access to application, network and data resources. Application, network and data layers protect the resources requested for access and the monitoring function provides security-related telemetry information for the contextual risk assessment. Security-related information generated by changes in the security conditions is communicated to the policy decision function (CISA, 2021).




Fig 1 Proposed Layered Zero Trust Security Framework for Cloud-Based Enterprise Systems in Hybrid Work Environments.

➤ *Framework Interaction and Continuous Feedback*

The above layers are interconnected with a continual security feedback process, which in turn integrates authentication, authorization, resource protection and monitoring with the above-described continual risk assessment process.

The interaction can be represented as:

Access Request → Identity Verification →
Device/Workload Assessment → Authorization →
Resource Protection → Monitoring → Risk Assessment
→ Policy Decision → Enforcement → Reassessment

The structural design of the framework allows for changes in security to affect subsequent access decisions. In addition to a change in device posture or other relevant security information, changes in security can be used to trigger re-evaluation and result in different behavior under the same policy. The framework provides structured architecture to implement Zero Trust in distributed cloud and hybrid enterprise environments and serves as a foundation for the more detailed control sets described in the following sections (CISA, 2021).

V. IDENTITY, DEVICE, AND ACCESS CONTROL LAYER

The Identity, Device, and Access Control Layer in the proposed framework defines the identity and security attributes to control access to enterprise resources in the cloud. This layer combines Identity and Access Management (IAM) with Multi-Factor Authentication (MFA), device posture, least-privilege, and workload identity. This layer of control defines the identity and endpoint attributes required by the policy mechanisms in the framework (Rose, 2020).

➤ *Identity and Access Management*

Identity and Access Management (IAM) enables the creation, management, authentication and control of identities of people, teams, applications and services in your company. This Identity and Access Management in a cloud-based hybrid environment includes identities of employees, guests/ visitors/ contractors, admins and more. Proper lifecycle management is key to Identity and Access Management. As a user's role changes in the course of time their access to services, applications, etc. must be updated accordingly. If a user no longer needs access to services, applications, etc. his/her account must be deprived. For example, someone who processes claims by day might need to access claims-processing applications but have no need to administer the underlying infrastructure as an administrator. IAM enables the storage of identities and roles (Grassi, 2020).

➤ *Multi-Factor Authentication*

Multi-Factor Authentication (MFA) uses multiple authentication factors to verify a user's identity. These can be something the user knows, something the user possesses and even something to do with the user's biometrics. MFA reduces the reliance on passwords for

verification of identity, as well as protects users against e-mailed or browsing-intercepted password capture. Authentication methods can be used at different levels of sensitivity of resources. Access to admin functions will need more authentication than access to low risk functions. MFA will be an additional identity-verification control to the existing device and authorization controls within the proposed framework (Hu, 2019).

➤ *Device Posture Assessment*

The device posture assessment feature determines whether a user's device complies with an organization's security requirements prior to granting access to organizational resources. This feature is highly relevant in today's hybrid work place where users connect to resources managed by the organization or personally owned devices are used to access organizational resources.

Relevant posture attributes can include:

- Operating-system version and patch status
- Endpoint protection status
- Disk encryption
- Device registration
- Security configuration
- Compliance with organizational policies

An organization-managed device that complies with the security posture of an organization will be treated differently by Policy Management components than an unmanaged or non-compliant device. The posture of the device can be used as input by policy components (Iorga, 2020).

➤ *Least-Privilege Authorization*

Least privilege ensures users and services are granted just enough permission to complete their authorized tasks. It reduces the danger of newly introduced privileged functionality by limiting the amount of power that can be lost due to system compromise. These can be Role-Based Access Control (RBAC) or Attribute-Based Access Control (ABAC), incorporating in both cases department, device status, and the sensitivity of resources as well as other environmental attributes. For example, a person would only be able to view information related to a patient or member for which they are performing work and would not be able to delete or export the information from which they are viewing. Also, for example, a program would only have access to the function(s) of the database needed for the application, rather than the entire database (Rose, 2020).

➤ *Workload Identity*

Increasingly, Cloud applications depend on APIs, Services, Containers, Virtual Machines and other automated Workloads. All these Workloads need identities to allow authentication and authorization of their interactions. Workload Identity is the association of an application or service with an identity that can be used to authenticate and authorize interactions. For example, in a company processing claims, a claims-processing

application may be required to fetch information from a database via an API call. The Workload Identity is associated to the application with the specific set of permissions required to complete the task at hand, preventing it from accessing other functions on the database. Workload identity is especially relevant in cloud-based workloads that can be deployed, scaled and even moved as necessary. Identity-based controls are crucial to manage such workloads consistently (Hu, 2019).

➤ *Layer Integration*

The controls in this layer provide complementary security attributes:

- IAM establishes who or what the entity is.
- MFA strengthens identity verification.
- Device posture establishes the security condition of the endpoint.
- Least privilege defines the permissions required for the entity's role or task.
- Workload identity establishes the identity of applications and services.

Together, the attributes provided by these controls allow the Zero Trust architecture to make policy decisions and apply controls in other layers, without repetition. The Identity, Device, and Access Control Layer forms the base layer for Application, API, Network, Data Protection, and Continuous Monitoring layers (Rose, 2020).

VI. APPLICATION, API, NETWORK, AND DATA PROTECTION LAYER

The Application, API, Network, and Data Protection Layer implements Zero Trust principles for the set of enterprise resources and applications that are accessed through this Layer. Zero Trust does not assign trust to applications, network segments, APIs, or data repositories just because they reside within a company's environment. Instead, they are protected by appropriate security measures and are continuously monitored and assessed on basis of their identity, sensitivity, communication needs, and security surroundings. The security measures for protection of cloud-based enterprise applications are implemented as a set of complementary measures for application security, API security, micro-segmentation, encryption, and data classification (Souppaya, 2017).

➤ *Application Security*

As part of a Zero Trust environment, application security works by treating individual applications as unique resources and doesn't grant access to the application just because access to the enterprise network has been granted. Measures such as secure configuration, vulnerability management, security testing, input validation and application-level authorization are implemented to ensure only the correct users and services are able to interact with the correct application functions. This form of protection aligns with the Zero Trust approach of protecting individual resources and removing implicit trust from the enterprise environment. Authentication to the organization is not in itself sufficient

to grant access to sensitive application functions. The application itself will impose further authorization requirements (Souppaya, 2017).

➤ *API Security*

API security as part of Zero Trust requires explicit authentication and authorization for interactions between applications, services and external systems. Although an API request might originate from an application or the company's network, it does not automatically deserve to be trusted. To secure API interactions, authentication, authorization, input validation, secure communication, API request monitoring and logging must be applied and verified individually. This also applies to interactions within a healthcare environment: An application might be authorized to read claims information via an API but does not automatically get permission to read other patients' or administrative data. Such resource-specific authorization aligns with the Zero Trust principle of verifiable and individually authorized interactions (Chandramouli R. &, 2020) (Chandramouli R. , 2019).

➤ *Micro-Segmentation*

By replacing broad network trust with small, measurable security boundaries for applications, workloads, and sensitive data, micro-segmentation supports Zero Trust architectures. Communication between segments can be restricted based on identity, application, workload, and business-related criteria. For example, a claim processing system, provider management application, and financial database could each be created as separate segments with only the required communication paths established between them. As such, micro-segmentation validates the core premise of Zero Trust: just because someone has access to one resource does not necessarily mean they should have access to others (Chandramouli R. , 2019) (Chandramouli R. &, 2020).

➤ *Encryption*

By protecting information in transit as well as at rest, encryption supports the Zero Trust model. Information stored in databases, in cloud storage services and in backups can be protected, as well as information that is exchanged between users, applications, APIs and services. The corresponding key management must, however, ensure adequate access control for the encrypted information. As a data-layer security function, encryption complements identity verification with Zero Trust-based access control to resources as well as network-based control functions (Barker, 2020).

➤ *Data Classification and Protection*

Data classification helps implement the Zero Trust model by classifying data in terms of its sensitivity. In order to better protect information, organizations don't apply identical security measures to all data. Rather, based on factors like confidentiality, business value, and compliance with regulations, sensitive information such as financial data, personal data and health-related data can be protected with stricter access rules and be encrypted. In the proposed framework, data classification is used to make

contextual decisions with respect to the necessary verification and protection for access to resources. This is also in line with the Zero Trust principle of protecting individual resources in the most adequate way possible (Barker, 2020).

VII. CONTINUOUS MONITORING AND CONTEXTUAL RISK ASSESSMENT

This phase of the framework incorporates aspects of continuous security monitoring. For the purposes of this framework, continuous security monitoring refers to the collection of security telemetry from all layers of your security controls (identity systems, endpoints, applications, APIs, network activity, etc.) and the utilization of this data within the framework's contextual risk assessment in order to continuously verify the security of active sessions and make decisions regarding policy based on user and session context, device risk, resource sensitivity and additional contextual factors such as location and behavioral signals (Rose, 2020). In this fashion, the framework's risk assessment can cause a request with acceptable security to continue to run under the existing permissions of the active session, but as security conditions change (i.e. increased risk) could cause additional authentication to be required, lower the permissions of the active session, require additional verification before granting access, or deny access completely. A user authenticates to a compliant endpoint and is granted access to resources within the organization. Later, as the user's device posture changes or there is unusual activity related to an application, security telemetry is generated that indicates increased risk. This risk is then used as part of the updated contextual risk assessment to cause the framework to re-evaluate all active sessions and make policy decisions regarding each session in turn. Therefore, as opposed to relying on security information for a one-time verification activity, the framework continuously monitors security information and uses this information to make policy decisions in real time (Dempsey, 2011).

VIII. RISK-BASED POLICY DECISION AND AUTOMATED ENFORCEMENT

The Risk-Based Policy Decision and Automated Enforcement layer translates contextual information and risk scored by the monitoring layer into an access decision. Information such as verified identity, the authentication status of a user, the device posture, user and session information, the sensitivity of resources requested and information about the scored risk are input for the central policy engine. The central policy engine compares this information against security policies of an organization and translates this information into an access decision based on the security required for the resources requested and the current access context of a user (Rose, 2020).

➤ Risk Thresholds and Policy Decisions

All risk assessments are then compared by the policy engine to organization-defined risk thresholds. The thresholds in turn determine the amount of access to be

granted for the access request. The risk values and policy decisions are organization-specific and therefore cannot be defined in this framework as they will vary greatly from one organization to another based on a variety of factors including risk tolerance, resource sensitivities, and even regulatory requirements (Rose, 2020).

The framework makes four types of access decisions: allow, restrict, challenge, and deny.

- *Allow:*

For an identity, authentication, device, authorization, and contextual (IADAC) set of conditions to be satisfied for access to be granted to a user accessing a system to perform a function within an organization for which there is an acceptable level of risk in accordance with the organization's risk threshold. For example, an employee authenticating using their organization-managed device to access a claims-processing application (Grassi, 2020).

- *Restrict:*

Access with restricted permissions is granted to the requesting entity. Restrictions on access and functionality may include view only access to resources, restrict access to specific resources, limit the functionality provided by an application, and reduce the session duration of an active session. An employee requesting view claims may be granted access to view claims but not to modify or export the information for elevated contextual risk (Grassi, 2020).

- *Challenge:*

More authentication or a verification step must be completed by the requesting identity, for example reauthentication, additional MFA or other Step-Up authentication step(s) as defined by policy. This decision is taken where identity is known but more assurance is required (Grassi, 2020).

- *Deny:*

Mandatory security conditions for access have not been met. This may be due to the user not being authorized for access to the requested resource, or the assessed risk having exceeded the organization's defined risk tolerance for access. Examples include unknown or unauthorized users, insufficient permissions, non-compliant devices, or other conditions that violate one or more of an organization's defined security policies (Grassi, 2020).

The relationship can be represented as:

Risk Assessment → Threshold Evaluation → Allow /
Restrict / Challenge / Deny

This way the organization can map different security states to corresponding access answers instead of having to choose between allow or deny.

➤ Policy Enforcement

Once a decision has been taken by the policy engine regarding access to a resource, that decision can be enforced on the requested resource. Enforcement decisions include Grant (or allow) Restrict, Challenge, or Deny. For

example, if a user's device has become non-compliant whilst connected, a new risk Assessment is completed. If the resulting risk is greater than the organization's defined threshold, for example, the policy could be configured to require additional authentication, to restrict the user's access to certain resources, or to end the user's session entirely. This ensures that initial authorization of a user does not result in permanent access, but rather that enforcement of the resulting access state can change as relevant security-related information becomes available (Hu, 2019).

➤ *Continuous Policy Enforcement and Feedback*

The policy decision process will remain integrated with the monitoring and risk-assessment capabilities described in Section 7. Updated security-related information (security-(telemetry)-information) generated by the Security-Primitive during the execution of an application (during an active session) will enable reevaluation of risk and corresponding reevaluation of the access-policy in effect for that user (Barker, 2020).

The resulting process can be represented as:

Monitor → Assess Risk → Evaluate Threshold → Decide
Policy → Enforce → Monitor Again

This feedback mechanism is a central integration component in the proposed framework. By linking the security measures for the access control (Section 5 and 6) with the continuous monitoring and the corresponding risk assessment (Section 7), this access control can be updated while the security situation is changing, as opposed to by using the data collected at the time of the user's login. Our approach enables the effective translation of contextual security information into appropriate and continuously enforceable access decisions for use across an enterprise's cloud-based resources (Iorga, 2020) (Agency., 2021).

IX. ENTERPRISE IMPLEMENTATION SCENARIO AND COMPARATIVE ANALYSIS

➤ *Enterprise Scenario*

Consider a hybrid enterprise environment in a healthcare organization for claims processing, provider management, patient information system access and associated databases via APIs. Employees can access tools and information to do their jobs efficiently while protecting health and financial information in a workplace setting that includes a corporate office and remote locations. Employees can be issued organization managed laptops to do their work. A Zero Trust framework can be implemented to integrate Identity Verification, Device Health Assessment, Least Privilege, Application/ API Security, Micro-segmentation, Data Protection, Continuous Monitoring, Contextualized Risk, Policy Enforcement (Moyle, 2022).

➤ *Access Request Walkthrough*

We start by assuming the user to be granted remote access to a cloud-based claims application. The identity of the user is first verified followed by Multi-Factor Authentication (MFA). Next the posture of the user's remote devices are assessed, followed by an assessment of the user's role, the functions of the application that the user needs to perform and the level of sensitivity of the resources that the user will be accessing. Other security-related factors such as session context can also be taken into consideration.

If the request is approved, then the employee will have access to the claim's functions for their job role. The application and API layers then manage which functions and services within the application that the user can access. Micro-segmentation is also in place to restrict communication to only required resources for processing claims. Healthcare sensitive data is also protected through encryption and data-specific security functions. Security telemetry is also collected during the session to monitor changes in device posture and unusual activity. If this information results in a sufficient change in risk, then the policy engine can apply the appropriate access response (e.g. restrict the session or request further verification) (NIST, 2020).

➤ *Traditional Security Model*

Typical access decisions for a perimeter-based security model are based on the network location where access is requested and whether or not the user has authenticated to the organization's resources. Once a user has gained access to the organization's network, they have access to all resources on the internal network based on their network permissions. As a result of this model, a compromised employee account or device could potentially allow an attacker to access resources above and beyond what is required to perform the employee's job if additional security measures are not put into place. Even with the use of network segmentation, authentication, and endpoint controls, the traditional perimeter-based security model is focused on building and maintaining a trusted internal boundary (NIST, Systems security engineering: Considerations for a multidisciplinary life cycle engineering process, 2017).

➤ *Existing Zero Trust Approach*

An existing Zero Trust implementation uses the following key elements: explicit verification, least privilege, identity-based access, device security, resource protection, and ongoing monitoring. This is opposed to a broad network location-based approach of trust. The healthcare example uses verification of an employee's identity and the condition of their device before granting access. The authorization policies for that employee determine what claims are allowed for what resources. Also, micro-segmentation and application-level controls are used to reduce unnecessary connections between resources and restrict the ability to move laterally (Bellala, 2025).

➤ *Proposed Framework*

The established Zero Trust controls are organized into layered architecture. All layers are connected through a contextual risk assessment and the corresponding policy feedback. This structured relationship between the different security layers is unique. Identity and device information are used for authorization of identity. The corresponding application, API, network and data controls are protecting the resources for which access has been requested. The security monitoring and analytics layer is generating the corresponding security telemetry. The contextual risk assessment is analyzing changes in the security related situation. The corresponding policy engine is translating the security situation into appropriate access

decisions. The architecture does not introduce MFA or micro-segmentation as new technologies but organizes established Zero Trust controls in a structured fashion, including monitoring for security-related events, contextual risk assessment, and policy enforcement. This results in a common decision process within the enterprise for accessing resources and services. The architecture represents this process as a sequence of steps: Request → Verify → Assess Context → Determine Risk → Apply Policy → Access Resource → Monitor → Reassess (Barker, 2020).

➤ Comparative Analysis

Table 1 Key Security Dimensions Compared Across Three Enterprise Security Approaches

Security Aspect	Traditional Security	Existing Zero Trust	Proposed Framework
Trust basis	Primarily network/perimeter-based	Identity, device, resource, and context	Layered identity, device, resource, and contextual evaluation
Authentication	Primarily initial authentication	Explicit authentication and MFA	Explicit authentication and MFA integrated with risk evaluation
Authorization	Often broader network or role-based access	Least-privilege resource access	Least privilege combined with contextual risk and resource sensitivity
Network protection	Perimeter and network segmentation	Micro-segmentation and resource-level controls	Micro-segmentation integrated with application, API, and data controls
Monitoring	Primarily perimeter and system monitoring	Continuous security monitoring	Telemetry linked directly to contextual risk assessment and policy decisions
Risk evaluation	Often limited or event-based	Context-aware evaluation	Contextual risk evaluation connected to policy thresholds
Access response	Generally, allow or block based on established rules	Policy-based access decisions	Allow, restrict, challenge, or deny based on assessed conditions
Policy feedback	Limited integration between controls	Continuous evaluation	Structured feedback from monitoring and risk assessment to policy enforcement

➤ Expected Security and Operational Improvements

Implementation of this framework will result in a number of architectural and operational improvements. From a security perspective, least privilege authentication and micro-segmentation will ensure sufficient restriction of access to resources whilst attempting to limit potential scope of a dataset breach. In addition to basic security controls on applications, APIs and data, contextual risk assessment allows security policy to be adaptive to constantly changing access conditions. Operationally, the integration of security telemetry, risk assessment, and security policy decisions will form a more deterministic process to access requests and aid in handling of changed access conditions. Improved policy enforcement will also be a byproduct of this architecture in that access responses are explicit and associated with corresponding risk and resource access requirements. However, improvements will largely be dependent upon effective implementation, associated policy, individual technologies, and underlying organizational processes (Agency., 2021).

X. CONCLUSION, LIMITATIONS, AND FUTURE RESEARCH

The architecture is structured into different layers of security controls for Zero Trust security in cloud-based enterprise systems. It helps in addressing security challenges for distributed and hybrid work environments.

Each layer comprises of already existing controls of Zero Trust, which are integrated in a structured manner. The layers comprise of identity and authentication, device and workload trust, least-privilege authorization, application and API security, micro-segmentation, data protection, and continuous monitoring. In order to connect all the different layers, contextual risk assessment and policy-based access decisions are used, enabling security decisions based on changing security conditions and enabling access decisions such as allow, challenge, restrict, or deny (NIST, Computer security incident handling guide, 2012).

One limitation of the framework as presented is that it is conceptual architecture. It has not been subject to a production deployment, or a controlled experimental evaluation to assess the security and operational benefits that the framework is expected to offer. Secondly, the framework does not include risk scores or thresholds for several reasons. First, risk scores and thresholds are often organizational dependent. Secondly, the level of risk that can be tolerated for sensitive information and resources can vary from organization to organization (Rose, 2020) (Barker, 2020). Thirdly, the type and extent of threats that organizations face can vary significantly. Lastly, there are many security-related information sources that can be used to calculate a risk score. Implementing the framework requires the integration of identity management systems, endpoint security, applications, APIs, network security

functions, data protection functions, and monitoring and policy enforcement systems. Implementing the security functions can vary depending on the enterprise architecture and the capabilities of the technologies used in the implementation. Furthermore, the framework does not describe behavioral analysis algorithms or automated threat classification algorithms (Teerakanok, 2021).

The author would like to continue developing this framework. First of all, future research could investigate the application of behavioral analytics in order to determine typical user and workload behavior as well as unusual behavior. Then, intelligent threat detection could be a further point of interest for future research, including a possible correlation of all types of security-relevant events in identity management, endpoint security, application security, network security, and data protection. Furthermore, automated risk assessment and corresponding dynamic policy management could be further developed, for example by determining adequate thresholds for different policies on the basis of a large number of different security-relevant contextual information. These possibilities could lead to adaptive security architectures, in which access rights and / or security measures can be changed during run time depending on the actual security situation (Kerman, 2020).

REFERENCES

- [1]. Barker, E. (2020). Recommendation for key management: Part 1—General (NIST Special Publication 800-57 Part 1 Revision 5). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-57pt1r5>
- [2]. Buck, C., Olenberger, C., Schweizer, A., Völter, F., & Eymann, T. (2021). Never trust, always verify: A multivocal literature review on current knowledge and research gaps of zero-trust. *Computers & Security*, 110, 102436. <https://doi.org/10.1016/j.cose.2021.102436>
- [3]. Chandramouli, R. (2019). Security strategies for microservices-based application systems (NIST Special Publication 800-204). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-204>
- [4]. Chandramouli, R., & Butcher, Z. (2020). Building secure microservices-based applications using service-mesh architecture (NIST Special Publication 800-204A). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-204A>
- [5]. Bellala, K. R. (2025). AI at the edge: Cloud-edge synergy. *International Journal of Innovative Science and Research Technology*, 10(5), 2561–2567. <https://doi.org/10.38124/ijisrt/25may967>
- [6]. Cybersecurity and Infrastructure Security Agency. (2021). Zero trust maturity model (Version 1.0). U.S. Department of Homeland Security.
- [7]. Dempsey, K. L., Chawla, N. S., Johnson, L. A., Johnston, R., Jones, A. C., Orebaugh, A. D., Scholl, M. A., & Stine, K. M. (2011). Information security continuous monitoring (ISCM) for federal information systems and organizations (NIST Special Publication 800-137). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-137>
- [8]. Ferretti, L., Manganin, F., Andreolini, M., & Colajanni, M. (2021). Survivable zero trust for cloud computing environments. *Computers & Security*, 110, 102419. <https://doi.org/10.1016/j.cose.2021.102419>
- [9]. Grassi, P. A., Fenton, J. L., Newton, E. M., Perlner, R. A., Regenscheid, A. R., Burr, W. E., Richer, J. P., Lefkovitz, N. B., Danker, J. M., Choong, Y.-Y., Greene, K. K., & Theofanos, M. F. (2020). Digital identity guidelines: Authentication and lifecycle management (NIST Special Publication 800-63B). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-63B>
- [10]. Bellala, K. R. (2025). Data privacy and compliance in the cloud (GDPR, HIPAA, CCPA). *International Journal of Innovative Science and Research Technology*, 10(7), 2091–2097. <https://doi.org/10.38124/ijisrt/25jul1264>
- [11]. Hu, V. C., Ferraiolo, D. F., Kuhn, D. R., Schnitzer, A., Sandlin, K., Miller, R., & Scarfone, K. (2019). Guide to attribute-based access control (ABAC) definition and considerations (NIST Special Publication 800-162). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-162>
- [12]. Hu, V. C., Iorga, M., Bao, W., Li, A., Li, Q., & Gouglidis, A. (2020). General access control guidance for cloud systems (NIST Special Publication 800-210). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-210>
- [13]. Kerman, A., Borchert, O., Rose, S., Division, E., & Tan, A. (2020). Implementing zero-trust architecture. National Institute of Standards and Technology.
- [14]. Mell, P., & Grance, T. (2011). The NIST definition of cloud computing (NIST Special Publication 800-145). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-145>
- [15]. Microsoft. (2021). How to secure your hybrid work world with a Zero Trust approach. Microsoft Security.
- [16]. National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations (NIST Special Publication 800-53 Revision 5). U.S. Department of Commerce.
- [17]. Open Worldwide Application Security Project. (2019). OWASP Top 10 API security risks—2019. <https://owasp.org/API-Security/>
- [18]. Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). Zero trust architecture (NIST Special Publication 800-207). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-207>
- [19]. Souppaya, M. P., Morello, J., & Scarfone, K. (2017). Application container security guide (NIST Special

- Publication 800-190). National Institute of Standards and Technology.
<https://doi.org/10.6028/NIST.SP.800-190>
- [20]. Teerakanok, S., Beam, R., & Lee, S. (2021). Migrating to zero trust architecture: Reviews and challenges. *Security and Communication Networks*, 2021, Article 9947347.
<https://doi.org/10.1155/2021/9947347>
 - [21]. Ward, R., & Beyer, B. (2014). BeyondCorp: A new approach to enterprise security. *login*, 39(6), 6–11.
 - [22]. National Institute of Standards and Technology. (2019). Attribute considerations for access control systems (NIST Special Publication 800-205). U.S. Department of Commerce.
 - [23]. National Institute of Standards and Technology. (2020). Security and privacy controls for information systems and organizations: Control baselines for information systems and organizations (NIST Special Publication 800-53B). U.S. Department of Commerce.
 - [24]. National Institute of Standards and Technology. (2020). Risk management framework for information systems and organizations: A system life cycle approach for security and privacy (NIST Special Publication 800-37 Revision 2). U.S. Department of Commerce.
 - [25]. Bellala, K. R. (2025). AI-assisted cloud migration. *International Journal of Innovative Science and Research Technology*, 10(8), 1396–1403.
<https://doi.org/10.38124/ijisrt/25aug947>
 - [26]. National Institute of Standards and Technology. (2018). Framework for improving critical infrastructure cybersecurity (Version 1.1). U.S. Department of Commerce.
 - [27]. National Institute of Standards and Technology. (2012). Computer security incident handling guide (NIST Special Publication 800-61 Revision 2). U.S. Department of Commerce.
 - [28]. National Institute of Standards and Technology. (2012). Guide for conducting risk assessments (NIST Special Publication 800-30 Revision 1). U.S. Department of Commerce.
 - [29]. National Institute of Standards and Technology. (2017). Systems security engineering: Considerations for a multidisciplinary life cycle engineering process (NIST Special Publication 800-160 Volume 1). U.S. Department of Commerce.
 - [30]. National Institute of Standards and Technology. (2020). Security strategies for microservices-based application systems (NIST Special Publication 800-204). U.S. Department of Commerce.
<https://doi.org/10.6028/NIST.SP.800-204>
 - [31]. Moyle, E. (2022). Case study: Cloud-native security using Zero Trust. *ISACA Journal*, 2022(3).