

THE Evolution of Enterprise Security Systems in Cloud-native Architectures: From Perimeter-Based to Automated, Granular Models

Nagaraju Velur
Wipro Limited, USA

ARTICLE INFO

Received: 06 Jan 2026

Revised: 08 Jan 2026

ABSTRACT

Enterprise security is evolving rapidly, driven by the shift from traditional perimeter-based defenses to distributed, cloud-native environments. Classical castle-and-moat strategies are insufficient for modern architectures spanning microservices, containers, serverless functions, and multi-cloud deployments. Zero Trust Architecture eliminates implicit trust, enforcing continuous verification at every access point. Identity and Access Management must adopt Attribute-Based Access Control (ABAC) for dynamic, context-aware authorization, while machine learning enhances threat detection across complex systems. Continuous monitoring via Cloud Security Posture Management (CSPM) automates discovery, assessment, and remediation, ensuring resilient and adaptive security. Together, these approaches transform security from a reactive function into a strategic enabler for business agility, governance, and operational efficiency.

Keywords: Zero Trust Architecture, Cloud-native Security, Attribute-Based Access Control, Machine Learning Threat Detection, Cloud Security Posture Management

1. Introduction: The Paradigm Shift from Perimeter-Based to Cloud-Native Security

The present scenario in the world of enterprises sees a paradigm shift in enterprise security due to the trend of using cloud-native architecture. The classical models of enterprise security were designed around an architecture that had a distinct boundary between trusted internal networks and insecure external networks. This classical castle-and-moat concept relied on segregating security mechanisms on the edge of trusted networks to assume that threats were from external networks and not from internal networks, thus deeming the latter to be trustworthy in nature. The castle-and-moat model no longer meets the needs of today's distributed, elastic, and ephemeral cloud architectures with users spanning multiple clouds, on-premises infrastructure, and edge computing infrastructure [1]. The Zero Trust model of security radically upsets the conventional assumptions of perimeter-based security by stating that no thing should be trusted from the outset, whether it's inside or outside the perimeter. This approach understands the fact that malicious activity may come from external sources but also from inside sources. According to the National Institute of Standards and Technology, the Zero Trust concept is a body of ideas aimed at eliminating ambiguity in making exact privilege perrequest access decisions regarding information systems and services when faced with a network environment that is considered to be compromised. Zero Trust Architectures emphasize the distinction between identity and resources, always verifying the security posture before access to resources [1].

Cloud-native systems bring forth a set of complexities to the security architecture space, making the classical perimeter security approach obsolete. Microservice architecture breaks an application into several microservices communicating with each other beyond the network perimeter, making monitoring and controlling impossible for the classical firewall solution. The Cloud Security Alliance specifies the importance of addressing the following distinctive features when providing security advice for cloud computing infrastructures: "on-demand self-service, broad network access, resource pooling,

rapid elasticity, or measured service. This set of features makes it necessary to use security measures working in a distributed fashion rather than at the network perimeter. Another complicating factor for cloud security is the shared responsibility cloud service model offering, involving the responsibility for specific security controls shared between the providers and the consumers [2].

The shortcomings in perimeter security practices were highlighted by increasing breaches despite significant spending on traditional security infrastructure. However, contemporary security frameworks acknowledge that trust cannot be inferred from network location, that trust verification must be an ongoing process and not, as in initial access, and that trust must be integrated with application logic and infrastructure automation. The renewed approach presents a challenge to organizations in that it calls for breaking down traditional security architectures with granular, distributed security controls at every layer in the technology stack.

2. Architectural Foundations of Cloud-Native Security Models

Cloud-native security architectures are built around distributed, service-centric controls rather than a single, network-centric perimeter. Applications are decomposed into independently deployable microservices that communicate over APIs and service meshes, which means security must be enforced at the service interface and data flow levels. In this model, every service-to-service call can be treated as an untrusted interaction, requiring strong authentication, authorization, and encryption regardless of network location.

Transport Layer Security (TLS) is the primary mechanism for protecting data in transit between cloud-native components. Older protocol versions, such as TLS 1.0 and 1.1, have known weaknesses and are being deprecated across major cloud platforms and operating systems, driving a shift to TLS 1.2 and TLS 1.3 as the minimum standard. These newer versions support stronger cipher suites, hardened handshakes, and better protection against downgrade attacks, making them better suited to securing high-volume, automated east–west traffic in microservice environments.[3].

Container security also brings its set of challenges and opportunities to enterprise security systems. "Containers share a kernel with the host OS, along with a level of process and filesystem isolation using Linux namespaces and control groups." Analysis studies that investigate container security concepts from a variety of viewpoints underscore that Docker containers, which offer substantial operational advantages, also introduce different concepts for security compared to virtualization. "The shared kernel model opens areas that, if exploited via escape attack vulnerabilities, could threaten the entire host system." Container security, as a practice, requires careful image provenance, runtime monitoring, and resource isolation techniques. Organizations need to practice image scanning to ensure base image and software dependencies are safe, use a protected image registry that controls access, and ensure policies are enforced to not run untrusted containers [4].

Information system security risks are divided into various levels, which are scalable to cloud-native solutions. Categorization of information system security risks, adopted in various studies, includes physical risks, natural risks, human risks, and environmental risks; human risks are further divided into malicious attacks and human errors. Cloud-native solutions differ significantly from traditional computing infrastructure in how each of these risks is realized. Malicious attacks could be on container orchestration solutions, misconfigured cloud services, and application dependencies. Human errors often cause security misconfigurations in forms like publicly accessible storage media, lax identity and access management policies, and the absence of logging and monitoring functionalities [5].

The shared responsibility model underpins all cloud-native security architectures by defining which controls are owned by the cloud service provider and which remain the customer's responsibility. The CSP generally takes responsibility for securing infrastructure such as physical environments, networks, and hypervisors, and this is complemented by the customer securing applications, data, identity, and access controls. The exact split varies across IaaS, PaaS, and SaaS offerings, but in all cases, organizations must explicitly map responsibilities and address gaps with appropriate controls,

automation, and monitoring. This model reinforces the need for granular, distributed security mechanisms that align with Zero Trust principles and are tightly integrated into cloud-native architectures.

TLS Protocol Version	Cryptographic Support	Handshake Security	Downgrade Protection	Deployment Recommendation
TLS 1.0	Limited	Weak	Vulnerable	Deprecated
TLS 1.1	Limited	Weak	Vulnerable	Deprecated
TLS 1.2	Strong	Enhanced	Protected	Recommended
TLS 1.3	Advanced	Enhanced	Protected	Recommended

Table 1: Transport Layer Security Protocol Comparison [3]

3. Zero Trust Architecture as the Cornerstone of Enterprise Security

Zero Trust Architecture(ZTA) is the complete security framework that has been developed using the concept that “No implicit trust should ever be granted to assets or user accounts based solely on their physical or network location.” The Special Publication from the NIST on Zero Trust Architecture clearly outlines that Zero Trust assumes that “No implicit trust is granted to assets or user accounts based solely on their physical or network location or based on asset ownership. Rather, authentication and authorization are separate steps that occur before the session to the enterprise resource has even been established. Zero Trust emphasizes securing resources rather than network segments, since the network location isn’t viewed as the key element contributing to the security characteristic of the resource [6].” At the core of ZTA are three logical components: the policy engine, the policy decision point, and the policy enforcement point. The policy engine acts as the trust authority, ingesting identity attributes, device health, resource sensitivity, and contextual risk indicators to decide whether to allow, deny, or revoke a session. The policy decision point coordinates evaluation and passes the outcome to the policy enforcement point, which enforces those decisions on traffic flows and resource access. Together, these components enable centralized policy logic with distributed enforcement across gateways, proxies, service meshes, and application-level controls [6].

Microsegmentation is a key implementation pattern for Zero Trust because it reduces the attack surface and limits lateral movement. Traditional network segmentation divides infrastructure into larger “trust zones,” but once an attacker enters a zone, they often have freedom to move between systems. Microsegmentation instead creates fine-grained security boundaries at the workload or application tier, enforcing policies based on identity and context rather than IP ranges alone. When combined with strong authentication, continuous posture assessment, and least-privilege access, microsegmentation can significantly constrain the blast radius of a compromise. [6].

Identity-centric verification is the operational backbone of ZTA, covering both human and machine identities. Strong, phishing-resistant multi-factor authentication becomes mandatory rather than optional, and access is granted dynamically based on attributes such as role, device compliance, geolocation, and current risk level. Adaptive or risk-based authentication can step up verification in response to anomalous behavior, while continuous authorization revalidates sessions as conditions

change. In cloud-native environments, this identity-centric approach extends to workloads, APIs, and services, ensuring that every interaction—user-to-app, app-to-app, and service-to-service—is explicitly authenticated and authorized in line with Zero Trust principles

Threat Category	Threat Type	Cloud-native Manifestation	Security Control Required
Physical	Infrastructure	Data center vulnerabilities	Provider responsibility
Natural	Environmental	Service disruptions	Disaster recovery
Human - Malicious	Attacks	Container orchestration exploits	Runtime monitoring
		Cloud service misconfigurations	Configuration scanning
Human - Unintentional	Errors	Publicly accessible storage	Access control policies
		Disabled logging	Automated monitoring

Table 2: Container Security Threat Categories [4, 5]

4. Identity and Access Management within Cloud Native Ecosystems

Building on Zero Trust principles, identity and access management becomes the key enabler for enforcing continuous verification. The Identity and Access Management system needs to handle identities of human users as well as machine identities, which are exponentially increasing because of the distributed nature of computing in cloud-native architectures. Role-Based Access Control has long been a baseline model to handle user authorization by delegating access privileges based on organizational roles. The study by NIST on RBAC verifies that the access control facility links access privileges and organizational roles and allocates users to those roles corresponding to their duties and qualifications. RBAC makes access management scalable by consolidating access privileges into roles that represent organizational hierarchies and thus minimizes the complexity involved in managing access privileges of numerous users in an enterprise. Traditional methods of RBAC have been proven to have limitations when adapting to dynamically changing cloud-native applications requiring diverse context-dependent authorization [7].

Attribute-Based Access Control is an extension of Role-Based Access Control because ABAC goes beyond static roles and considers dynamic attributes. According to NIST recommendations for ABAC, this access control mechanism assesses the subject attribute, object attribute, environmental attribute, and requested operation to authorize access requests based on the subjects' and objects' characteristics independently of roles [12]. In contrast to traditional RBAC access control methods that are dependent only upon roles assigned to users and administrators, ABAC access control decisions are based on the subject's security clearance level and affiliation with business organizations, the object's information classification and sensitivity level, and environmental factors such as current risk level and the exact

time of access attempts. The flexible nature of ABAC access control makes this mechanism Cloud-ready because access requests depend on various factors that cannot be determined via static roles [8].

The connection between RBAC and ABAC is a progression in evolution and not a replacement of one paradigm by the other. ABAC can use role in addition to other identity attributes in a superset that encompasses the functionalities of RBAC. In migrating to a cloud-native deployment, corporations embracing a hybrid model use RBAC to provision basic access while supplementing with ABAC to address complex authorization rules influenced by multiple context variables [8].

The issue of machine identity management has also been growing in importance as the number of service accounts and API keys tends to far outnumber human identities in a cloud-native setup. Workload identity systems use platform-provided identity infrastructure to provide cryptographically verifiable identity for containers, serverless compute functions, and virtual machines without the need for embedded credentials. Use of certificate-based authentication with short-lived certificates facilitates two-way authentication of services with the least possible damage from a compromised identity or credential. Privileged Access Management for a cloud-native infrastructure tackles the problem of managing elevated access in a distributed infrastructure setup with just-in-time access enrollment for a temporary privilege to accomplish a task, with automatic revocation of access at the end of the task execution.

Identity Type	Management Framework	Authentication Method	Credential Lifespan	Access Provisioning	Security Benefit
Human	RBAC	Multi-factor	Extended	Role-based	Simplified management
Human	ABAC	Adaptive	Session-based	Context-aware	Granular control
Machine	Workload identity	Certificate-based	Short-lived	Cryptographic	Minimized compromise
Privileged	PAM	Just-in-time	Temporary	Task-specific	Reduced attack surface

Table 3: Identity Management Approaches [7, 8]

5. Artificial Intelligence and Machine Learning within Automated Threat Detection

The use of artificial intelligence and machine learning has significantly enhanced the capabilities for threat detection by processing large volumes of data that could not be analyzed by the human mind. The use of machine learning algorithms in creating the intrusion detection system in software-defined networks qualifies the application of machine learning algorithms in the security field. Studies involving the design and development of the ML-based intrusion detection system involve the use of various classification machine learning algorithms, such as decision tree classification, naive Bayes classification, support vector machine classification, and ensemble classification approaches, for the detection of anomalies in the network. The experimental results validate the use of ensemble classification approaches that outperform the classification approaches individually, with detection rates achieving a high degree of accuracy and reasonable FPR [9].

Supervised learning algorithms, which are trained on labeled data sets containing malicious as well as legitimate activity, are used to categorize novel behaviors by identifying patterns within existing data. The performance of these models largely depends on the quality and representation of data used for

training. Generally, machine learning-based intrusion detection systems need data that contains various attack scenarios as well as normal data, which helps to learn well for production scenarios. “Feature engineering is a key component of ML-based security systems, as identifying important features right out of raw network traffic or system logs has a direct influence on accuracy” [9]. “The choice of algorithm can significantly affect accuracy, as some models are good for a small number of samples but poor for a larger set, while others work well for a larger set but are poor for a small number of samples” [10].

The comparison of network anomaly detection systems needs to be done through highly accurate statistical analysis and needs to be compared with existing benchmarks. Research studies on network anomaly detection on the UNSW-NB15 dataset offer extensive analysis on existing network anomaly detection datasets. In this study, it has been shown that new datasets are a better projection of current network scenarios compared to old benchmarks. The study process includes analysis in terms of characteristics and distributions in classifiers, which use several machine learning algorithms. The study has found that network detection system performance differs largely in terms of characteristics and attack types, with some attacks being more difficult to detect than others. This study asserts that it is necessary to model updated attacks on evolving networks to efficiently perform intrusion detection [10]. Unsupervised machine learning algorithms are very effective in anomaly detection techniques, which attempt to detect those points in the data that deviate considerably from anticipated patterns. Unsupervised learning algorithms are very useful in finding emerging attack patterns that are not associated with any defined signatures. Cluster analysis algorithms are used to categorize similar patterns and point out those that deviate from these patterns, which are then further examined, whereas autoencoder algorithms attempt to reconstruct input patterns and quantify their differences to detect deviating patterns. Recurrent learning algorithms, which are basically deep, examine network traffic patterns and attempt to detect intricate patterns of attacks that are performed over time. The models are also effective in finding multi-step attacks, in which apparently harmless steps are involved. The combination of various explainable AI models overcomes the challenges involved in using neural networks.

Learning Approach	Training Data Requirement	Pattern Recognition	Novel Attack Detection	Implementation Method	Interpretability
Supervised	Labeled	Historical	Limited	Classification models	Moderate
Unsupervised	Unlabeled	Deviations	Effective	Clustering, autoencoders	Challenging
Deep	Large	Sequential	Multi-stage attacks	Recurrent neural networks	Complex
Explainable AI	Varies	Transparent	Varies	Interpretable models	Enhanced

Table 4: Learning Algorithm Approaches for Threat Detection [9, 10]

6. Continuous Security Monitoring & Cloud Security Posture Management

Cloud native infrastructure needs a constant monitoring process to ensure adequate visibility. Best practice guidelines in cloud security define a structured set of guidelines to provide a multicategory layer

of defense to cloud infrastructure. Best practice guidelines for cloud security highlight the necessity to build a defense-in-depth strategy for cloud infrastructure based on a variety of approaches related to identity and access management, infrastructure, data, logging and monitoring, and incident response. Cloud native infrastructure needs constant monitoring to provide adequate visibility. Cloud native infrastructure continuously monitors data to provide immediate notification to the organization for any potential security breach, thereby providing a reduced time frame to allow an attacker to gain a foothold in the infrastructure or steal the sensitive data [11].

A Cloud Security Posture Management solution enables organizations to have the ability to undertake automated discovery and inventory of the cloud resources. This enables organizations to have the ability to detect all the compute resources, storage, database, networking, and identity components within the multi-cloud environment. Cloud Security Posture Management tools enable continuous assessment against security configurations. Security best practices highlight the need for organizations to have the ability to implement monitoring of security configurations on an automated basis. This enables organizations to have the ability to maintain consistent security configurations during the evolution of the infrastructure. This should be achieved through the implementation of infrastructure as code [11].

The Center for Internet Security Critical Security Controls are a prioritized series of actions to address common internet attacks in order to secure information and operations through defense actions aimed against popular internet attacks. CIS Controls Version 8 sets up eighteen control groups categorized according to implementation groups that help organizations build appropriate security programs. These controls interact with asset control and protection of enterprise resources, protection of data, secure configuration of resources, vulnerability management processes, management of audit logs, email and web browser protection mechanisms against attacks, protection of resources against malware attacks, data recovery processes for institutions against attacks, security awareness training processes for employees of institutions and management of incident response processes of an organization against attacks [12].

The Automated Compliance Validation Engines check security controls and produce evidence that proves their compliance with regulations and their own internal security policies. The tools check configuration snapshots, access logs, the execution of encryptions, and tests of incident response readiness. The Automated Remediation feature fixes security policy deviations through corrective configurations, such as turning on encryption on unencrypted storage resources or removing overly broad access control rules. The combined use of CSPM solutions and Security Information and Event Management solutions enables a comprehensive security operation center, where CSPM solutions offer context on configurations of cloud resources, and Security Information and Event Management solutions enable combining that context with security events. This enables a substantial reduction in Mean Time to Detection and Response to security events, thereby confining security incident impacts to a significantly limited scope through immediate containment and remediation activities.

Conclusion

The transition from perimeter-based defenses to cloud-native security represents more than a tooling refresh; it requires a re-architecture of trust, identity, and control across the enterprise technology stack. Zero Trust Architecture operationalizes the principle of “never trust, always verify” by enforcing continuous, identity- and context-driven access decisions for users, devices, and workloads, independent of network location. Cloud-native environments demand fine-grained controls at the microservice, container, and serverless layers, strengthened by modern cryptographic protections such as TLS 1.2 and TLS 1.3 for service-to-service communication. Evolving IAM from purely rolebased models toward attribute-based policies and strong workload identities enables more precise, adaptive authorization decisions in highly dynamic environments. Machine learning-based analytics and anomaly detection extend visibility and detection capabilities across high-volume, high-velocity

telemetry, uncovering subtle or previously unseen attack patterns. Finally, Cloud Security Posture Management and automated compliance validation close the loop by continuously assessing configurations, prioritizing risk, and triggering remediation, thereby maintaining a defensible, auditable security posture across multi-cloud deployments. Organizations that integrate these capabilities into a cohesive, automated security architecture can turn security from a reactive control function into a strategic enabler of agility, resilience, and governance in cloud-native ecosystems

References

- [1] S. Rose et al., "Zero Trust Architecture," NIST Special Publication 800-207, National Institute of Standards and Technology, 2020. [Online]. Available: <https://csrc.nist.gov/pubs/sp/800/207/final>
- [2] Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing v4.0," 2017. [Online]. Available: <https://cloudsecurityalliance.org/artifacts/security-guidance-v4/>
- [3] Microsoft, "TLS 1.0 and TLS 1.1 deprecation in Windows," [Online]. Available: <https://learn.microsoft.com/en-us/windows/win32/secauthn/tls-10-11-deprecation-in-windows>
- [4] T. Combe, A. Martin, and R. Di Pietro, "To Docker or Not to Docker: A Security Perspective," IEEE Cloud Computing, vol. 3, no. 5, 2016. [Online]. Available: <https://ieeexplore.ieee.org/document/7742298>
- [5] M. Jouini et al., "Classification of Security Threats in Information Systems," Procedia Computer Science, vol. 32, 2014. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1877050914006528>
- [6] S. Rose et al., "Zero Trust Architecture," NIST Special Publication 800-207, 2020. [Online]. Available: <https://nvlpubs.nist.gov/nistpubs/specialpublications/NIST.SP.800-207.pdf>
- [7] NIST, "Role-Based Access Control," 2025. [Online]. Available: <https://csrc.nist.gov/projects/rolebased-access-control>
- [8] V. Hu et al., "Guide to Attribute-Based Access Control (ABAC) Definition and Considerations," NIST SP 800-162, 2014. [Online]. Available: <https://csrc.nist.gov/pubs/sp/800/162/upd2/final>
- [9] A. Abubakar and B. Pranggono, "Machine Learning-Based Intrusion Detection System for Software Defined Networks," Seventh International Conference on Emerging Security Technologies (EST), 2017. [Online]. Available: <https://ieeexplore.ieee.org/document/8090413>
- [10] N. Moustafa and J. Slay, "The evaluation of Network Anomaly Detection Systems: Statistical analysis of the UNSW-NB15 data set and the comparison with the KDD99 data set," 2016. [Online]. Available: <https://www.tandfonline.com/doi/full/10.1080/19393555.2015.1125974>
- [11] Google Cloud, "Google Cloud security best practices center," [Online]. Available: <https://cloud.google.com/security/best-practices>
- [12] Center for Internet Security, "CIS Critical Security Controls Version 8," [Online]. Available: <https://www.cisecurity.org/controls/v8>