International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

An Adaptive Data Security Architecture for Large-Scale Web and Mobile Application Workloads

Srinivas Oguri

Principal Reliability Engineer, Arcesium, Andhra Pradesh, India

ABSTRACT: Consumer-scale web and mobile applications now routinely serve hundreds of millions of users across heterogeneous devices, networks, and jurisdictions. Their data security posture must therefore deal with elastic microservices, multi-cloud storage, and fast-changing threat models, while preserving low latency and user experience. Traditional perimeter-centric or static policy-based security mechanisms are inadequate in such environments, as they cannot continuously adapt to contextual changes in user risk, data sensitivity, or infrastructure health. Building on recent work in risk-based adaptive security, context-aware access control, and zero trust architectures, this article proposes an adaptive data security architecture tailored to large-scale web and mobile workloads. The architecture integrates (1) a risk- and context-aware policy decision fabric, (2) a microservices-oriented zero trust enforcement layer, (3) data-centric protection services, and (4) continuous monitoring and feedback loops that realize a MAPE (Monitor–Analyze–Plan–Execute) cycle over security controls. We outline how this architecture supports fine-grained, policy-driven protection for identity, data-in-transit, and data-at-rest while maintaining performance and scalability requirements typical of global consumer applications. We also discuss design trade-offs, evaluation metrics, and open research challenges such as explainable adaptive controls and cross-jurisdictional data governance.

KEYWORDS: Adaptive security; zero trust; context-aware access control; cloud security; web applications; mobile applications; microservices; data privacy; risk-based security.

I. INTRODUCTION

Large-scale consumer web and mobile applications—social platforms, streaming services, digital banking, and super-app ecosystems—operate at massive scale, with globally distributed users and infrastructure. These workloads are typically deployed atop microservices and cloud-native platforms, integrating numerous third-party APIs, edge delivery networks, and multi-region data stores. This environment amplifies data security challenges: threats emerge both inside and outside organizational boundaries; identities and devices are ephemeral; and data flows traverse heterogeneous clouds and content delivery networks.

Cloud security research has shown that the growth of cloud-enabled services has outpaced the maturation of unified security models, leaving gaps across requirements, threats, vulnerabilities, and countermeasures. Kumar and Goyal (2019) argue that addressing these gaps requires a holistic taxonomy and end-to-end mapping between security requirements and countermeasures across cloud service and deployment models. Kudos Surveys of cloud computing security similarly emphasize that dynamic, distributed environments demand adaptive and risk-aware controls rather than static configurations. ISecure Journal+1

At the same time, web and mobile application workloads have moved away from tightly controlled enterprise environments toward edge-heavy ecosystems, where context (e.g., device posture, network quality, geo-location, behavioral signals) heavily influences risk. Context-aware access control mechanisms for cloud and fog networks explicitly highlight the importance of spatial, temporal, and environmental conditions in access decisions for distributed IoT and cloud platforms (Kayes et al., 2020). MDPI Similarly, privacy-preserving access control frameworks for mobile computing, such as PADEC, support expressive, context-sensitive policies and fine-grained information release to balance utility and privacy in cooperative mobile applications (Herrera et al., 2022). X-MOL

Concurrently, the security architecture paradigm has shifted toward Zero Trust (ZT). NIST SP 800-207 defines Zero Trust Architecture (ZTA) as an evolution from network perimeter defenses to a resource-centric model with continuous verification, least privilege, and strong identity, across users, devices, and workloads (Rose et al., 2020). NIST CSRC For microservices-based systems, NIST SP 800-204 specifies security strategies that rely on API gateways, service meshes, and strong authentication/authorization at every service hop. NIST CSRC

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

Despite these advances, there remains a gap: how to systematically combine Zero Trust, context-aware access control, and risk-based adaptation into a coherent data security architecture explicitly tuned for large-scale web and mobile workloads. This article addresses that gap by proposing an adaptive data security architecture whose central design goals are:

- **Adaptivity:** Security controls adjust at runtime based on changing context and risk.
- **Data-centricity:** Policies are expressed in terms of data sensitivity, purpose, and sharing constraints.
- **Scalability:** The system operates under high concurrency and low-latency constraints typical of consumer-scale applications.
- **Observability:** Continuous monitoring and feedback loops support learning and policy refinement.

The rest of the article is organized as follows. Section 2 reviews relevant research on cloud and data security, adaptive and risk-based controls, and context-aware privacy. Section 3 introduces the proposed architecture and its layers. Section 4 details the core components and adaptation mechanisms. Section 5 discusses application scenarios and evaluation considerations. Section 6 outlines open research challenges, and Section 7 concludes.

II. BACKGROUND AND RELATED WORK

2.1 Cloud and application data security

Cloud security literature consistently reports that multi-tenant, on-demand environments introduce complex threat surfaces, including data breaches, misconfiguration, and insecure APIs. Surveys of cloud threats and solutions identify confidentiality, integrity, and availability risks across infrastructure, platform, and application layers, along with related countermeasures such as encryption, access control, and auditing (Butt et al., 2022; Pericherla, 2022). SpringerLink+1

Sun (2020) proposes a comprehensive privacy security protection framework for cloud computing, synthesizing technologies such as attribute-based encryption (ABE), fine-grained access control, multi-authority setups, and searchable encryption to address data confidentiality and privacy challenges. librarysearch.hillsdale.edu Their work emphasizes:

- The need for **multi-tenant** isolation in shared infrastructures.
- Combining encryption with expressive access policies.
- Handling key revocation, traceability, and multi-attribute conditions.

In parallel, higher-level surveys like Kumar and Goyal (2019) categorize security requirements, threats, vulnerabilities, and countermeasures in cloud environments and highlight open areas such as trust models and cloud-enabled big data, IoT, and SDN/NFV applications. Kudos These surveys collectively motivate architectures that operate beyond a single cloud or security layer, instead coordinating controls across network, application, and data tiers.

2.2 Zero Trust and microservices security

Zero Trust Architecture explicitly rejects implicit trust based on network location. In NIST SP 800-207, Rose et al. (2020) describe ZTA as centered on policy decision points (PDPs) that rely on rich context about user identity, device health, and environment to continuously evaluate access. NIST CSRC This aligns closely with the needs of consumer-scale applications, which serve users from untrusted networks and devices.

For microservices-based systems, Chandramouli (2019) outlines security strategies that integrate API gateways, service meshes, mTLS, service discovery, and resilience patterns (e.g., circuit breakers) to secure internal communications and manage cross-cutting concerns such as authentication, authorization, and monitoring. NIST CSRC This guidance implicitly supports ZTA by treating every service call as untrusted and subject to policy enforcement.

However, most Zero Trust and microservices guidance remains **static**: policies and enforcement typically do not adapt in real time to nuanced risk levels, performance conditions, or emergent traffic patterns.

2.3 Adaptive and risk-based security

Adaptive security research proposes architectures in which controls dynamically change behavior based on monitored conditions. Calvo and Beltrán (2022) introduce RiAS, a risk-based adaptive security model in which security controls are adjusted via a three-layer architecture and a Monitor–Analyze–Plan–Execute (MAPE) loop. searchit.libraries.wsu.edu The model supports parametric, architectural, and behavioral adaptations (e.g., changing the configuration of a web application firewall) in near real time.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

Similarly, Beer Mohamed et al. (2021) propose an adaptive security architectural model for federated identity in service-oriented computing, enhancing traditional SAML, OAuth, and OpenID mechanisms with adaptive protections for identity federation in large-scale financial data environments. ScienceDirect Their evaluation shows improved resistance to identity misuse and theft in a large-scale enterprise setting.

These works demonstrate that adaptive security can be realized in production-like environments, but they primarily focus on specific domains such as federated identity or web application filters, rather than end-to-end data flows for web and mobile workloads.

2.4 Context-aware access control and privacy

Context-aware access control (CAAC) mechanisms extend classical role- and attribute-based access control by incorporating dynamic contextual information—such as user location, time, device type, or environmental conditions—into authorization decisions. Kayes et al. (2020) provide a detailed survey and taxonomies of CAAC mechanisms for cloud and fog networks and propose a fog-based CAAC framework to reduce latency while maintaining expressive context-based policies. MDPI

In mobile settings, Herrera et al. (2022) introduce PADEC, a privacy-aware device communication framework that allows users to define rich access rules and attach different levels of information granularity to each rule, enabling fine-grained control of what data is released under which context. X-MOL

Complementarily, Jiang et al. (2021) develop Context-Aware Local Information Privacy (LIP), a privacy model that improves utility compared to pure local differential privacy by incorporating prior knowledge into perturbation mechanisms while preserving strong instance-wise privacy guarantees. Bohrium

These contributions show that **context** and **data utility–privacy trade-offs** are central to modern access control and privacy-preserving mechanisms—key ingredients for an adaptive data security architecture.

III. ARCHITECTURAL OVERVIEW

3.1 Design goals

For large-scale web and mobile workloads, the proposed adaptive data security architecture aims to:

1. **Continuously adapt to context and risk.** Security posture (e.g., access decisions, encryption strength, throttling) should change in real time based on risk indicators and operational conditions.
2. **Remain data-centric.** Policies are expressed primarily in terms of data sensitivity, purpose, retention, and residency, rather than only in terms of infrastructure topology.
3. **Scale with microservices and multi-cloud.** The architecture must support independently deployable services, polyglot persistence, multi-tenant environments, and hybrid/multi-cloud scenarios.
4. **Minimize performance overhead.** For consumer-scale applications, security components must add minimal latency and preserve high throughput.
5. **Support observability and auditability.** Decisions and adaptations should be explainable and auditable for compliance and debugging.

3.2 Layered architecture

The architecture is organized into four logical layers:

1. **Client and Edge Layer**
   - Browser and mobile clients, SDKs, and device agents.
   - Content delivery networks (CDNs), edge authentication, and DDoS protection.
2. **Access & Policy Control Plane**
   - Identity and access management (IdP, OAuth/OIDC, device identity).
   - Policy decision points (PDPs) implementing Zero Trust principles.
   - Risk and context evaluation engines.
3. **Service Mesh & Data Plane**
   - API gateways for north–south traffic.
   - Service mesh (sidecar proxies) for east–west traffic with mutual TLS, authentication, and fine-grained authorization.
   - Enforcement points for request-level and data-level policies.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

4. **Data Protection & Governance Layer**
   - Data classification and cataloging.
   - Encryption, tokenization, and key management services.
   - Data access logging, retention, and residency controls.

Cross-cutting these layers is an **Adaptive Security Orchestrator** implementing a MAPE loop over telemetry from services, clients, and infrastructure, inspired by RiAS (Calvo & Beltrán, 2022). searchit.libraries.wsu.edu

IV. CORE COMPONENTS AND ADAPTIVE MECHANISMS

4.1 Identity and access fabric

Federated and adaptive identity.

The identity fabric integrates external identity providers (e.g., social login, enterprise IdPs) with in-house identity services using standards like OAuth 2.0 and OpenID Connect. Building on adaptive federation concepts from Beer Mohamed et al. (2021), the fabric: ScienceDirect

- Continuously evaluates identity assurance (e.g., based on device posture, IP reputation, behavioral signals).
- Escalates authentication requirements (e.g., step-up MFA, re-authentication) when risk increases.
- Incorporates **device-binding** and **session risk scores** into access decisions.

Zero Trust PDPs.

Following NIST SP 800-207 (Rose et al., 2020), PDPs act as logically centralized policy engines that: NIST CSRC

- Validate user, device, and workload credentials for every request.
- Incorporate contextual attributes (geo-location, access time, network type, anomaly indicators).
- Consult data-classification-aware policies to decide on access level (allow, deny, require step-up, or restrict to redacted views).

The PDPs expose a **Policy Decision API** used by API gateways, service meshes, and data services to obtain decisions in real time.

4.2 Context- and risk-aware access control

Building on CAAC research (Kayes et al., 2020; Herrera et al., 2022), access control policies operate over attributes such as: MDPI+1

- **Subject attributes:** user role, subscription tier, historical trust score.
- **Object attributes:** data sensitivity label (e.g., public, PII, financial), tenant, residency region.
- **Environment attributes:** network type, region, time-of-day, anomalous activity flags.
- **Session attributes:** device OS, app version, recent failed logins, transaction value.

Policies are expressed in a **high-level policy language** that supports:

- **Contextual conditions** (e.g., “allow access to PII from a new device only if user passes MFA and network is not flagged as high risk”).
- **Granularity control** as in PADEC (Herrera et al., 2022), enabling partial data disclosure (e.g., masked account number) under moderate risk, and full data only under low-risk context.
- **Obligations**, such as requiring data redaction before logging, or triggering additional monitoring.

Risk evaluation uses scores derived from external threat intelligence, anomaly detection, and internal telemetry. These scores feed into the **Analyze** phase of the MAPE loop, which may update policy parameters (e.g., thresholds) or chosen enforcement strategies.

4.3 Data-centric protection services

Following Sun (2020), the architecture combines multiple cryptographic and data-centric controls: librarysearch.hillsdale.edu

1. **Classification and tagging.** Data streams and storage objects are automatically labeled based on schema, content, and business metadata (e.g., PII, financial, health, telemetry).
2. **Encryption and tokenization.**
   - Strong encryption at rest (with separate keys per tenant and region).
   - Tokenization services for highly sensitive values (e.g., card numbers).
   - Field-level encryption for selected columns or document fields.
3. **Attribute-based access enforcement.** For particularly sensitive data sets or cross-organizational scenarios, ABE or related schemes can be used to bind decryption capabilities to attribute sets (Sun, 2020), allowing policy-based data sharing without exposing raw keys.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

4. **Privacy-preserving analytics.**
   - Context-aware privacy mechanisms such as LIP (Jiang et al., 2021) can be used to adapt noise levels based on prior knowledge and required utility. Bohrium
   - Differential privacy or local privacy schemes can be selectively enabled for analytics pipelines that process user events, especially under higher regulatory or reputational risk.

The **Adaptive Security Orchestrator** can adjust, for example, the strength of anonymization or tokenization for a dataset based on observed query patterns and risk, trading off latency and utility against privacy leakage.

4.4 Service mesh and API gateway enforcement

In line with microservices security strategies (Chandramouli, 2019), a service mesh enforces security at each hop: NIST CSRC

- **Mutual TLS (mTLS)** for all service-to-service communication, with short-lived certificates and automated rotation.
- **Fine-grained authorization** at the service proxy, consulting PDPs with workload, user, and request context.
- **Rate limiting and throttling** based on client identity, tenant, and risk level.

The API gateway at the edge performs complementary functions for external traffic:

- Validates tokens, applies coarse-grained authorization, and normalizes headers.
- Consults PDPs for cross-tenant and cross-region policy checks.
- Enforces **data egress policies**, e.g., preventing certain data classes from leaving specific regions or being forwarded to third-party APIs.

Enforcement points push **decision logs and telemetry** back into the monitoring pipeline, closing the feedback loop.

4.5 Adaptive monitoring and response

The adaptive engine implements a MAPE loop similar to RiAS (Calvo & Beltrán, 2022): searchit.libraries.wsu.edu

1. **Monitor:** Collects metrics from:
   - Access logs (denials, step-up events, anomaly flags).
   - Infrastructure (CPU, latency, error rates, deployment metadata).
   - Security tools (IDS/IPS alerts, WAF events, behavioral analytics).
2. **Analyze:**
   - Correlates events to detect patterns (e.g., credential stuffing, cross-tenant data access, high-risk geo clusters).
   - Recomputes risk scores per user, tenant, data domain, and region.
3. **Plan:**
   - Selects adaptation strategies, such as raising MFA requirements, throttling certain traffic, hardening WAF rules, or increasing anonymization levels for specific analytics.
4. **Execute:**
   - Reconfigures PDP thresholds, gateway or mesh policies, and data anonymization parameters through APIs or configuration stores.

This adaptivity is constrained by **guardrails** specifying allowable action ranges, avoiding uncontrolled oscillations or overly aggressive restrictions that could degrade availability.

V. APPLICATION SCENARIOS AND EVALUATION

5.1 Example end-to-end flow

Consider a mobile banking application used by millions of customers across regions:

1. A user logs in from a new mobile device over a cellular network in a high-risk region.
2. The IdP authenticates credentials and issues an access token with initial risk attributes based on device fingerprinting and geo-IP reputation.
3. The API gateway validates the token and queries the PDP with context (device, location, account type, requested operation).
4. For a **balance inquiry**, the PDP may allow access after step-up MFA but enforce **masked account numbers** via a data transformation policy, leveraging mechanisms similar to PADEC's granularity control (Herrera et al., 2022). X-MOL
5. For a **large fund transfer** from the same context, the PDP may deny or require additional verification (e.g., call center validation).
6. All decisions and contextual features are sent to the monitoring pipeline. If many high-risk events are observed from similar IP ranges, the adaptive engine might:

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

- Increase friction (e.g., always require MFA for that region).
- Tighten rate limits in the API gateway.
- Temporarily degrade data utility for analytics involving IPs from that region by increasing privacy perturbation (Jiang et al., 2021). Bohrium

This illustrates how the architecture balances security, usability, and performance through context- and risk-aware adaptation.

5.2 Evaluation criteria

Evaluating such an architecture in large-scale web and mobile deployments involves multi-dimensional metrics:

- **Security effectiveness:**
  - Reduction in successful data breaches or unauthorized accesses.
  - Time-to-detect and time-to-respond for incidents.
  - Coverage of identified threats and vulnerabilities, aligned with taxonomies from cloud security surveys (Kumar & Goyal, 2019; Pericherla, 2022). Kudos+1
- **Performance and scalability:**
  - Latency overhead added by PDP and enforcement points.
  - Throughput under peak load across regions.
  - Impact of adaptive actions (e.g., additional cryptographic operations) on resource utilization.
- **Privacy and compliance:**
  - Measured privacy guarantees for protected datasets (e.g., via information-theoretic metrics as in LIP). Bohrium
  - Regulatory alignment with data residency and access rules.
- **Operational usability:**
  - Ease of authoring and auditing policies.
  - Explainability of access decisions and adaptive actions for developers and security analysts.

VI. DISCUSSION AND OPEN CHALLENGES

While the proposed architecture draws heavily on prior work, several research and engineering challenges remain, especially for consumer-scale deployments:

1. **Policy complexity and manageability.**

Expressive, context-rich policies can quickly become difficult to reason about. Automated tooling is needed for policy verification, conflict detection, and impact analysis, especially when adaptive changes are made automatically.

2. **Explainability of adaptive decisions.**

When risk engines use machine learning for anomaly detection or risk scoring, explaining why certain controls were tightened or relaxed becomes non-trivial. This is critical for trust, debugging, and regulatory scrutiny.

3. **Cross-jurisdictional governance.**

Large-scale applications often operate across regions with different privacy regulations. Data-centric policies must encode residency, retention, and purpose limitations in a way that can be enforced across multiple clouds and legal domains.

4. **Coordination with application developers.**

Some security obligations, such as proper input validation or output encoding, remain in application code. The architecture must provide clear contracts and libraries so that developers can integrate with identity, policy, and data protection services without bespoke, error-prone implementations.

5. **Resilience against adaptive adversaries.**

As adversaries probe and learn adaptation strategies, they may attempt to shape the risk models or exploit blind spots. Complementary research on game-theoretic modeling of adaptive attackers and defensive strategies is relevant here, even though much of that literature has matured post-2022.

VII. CONCLUSION

Large-scale web and mobile application workloads demand data security architectures that are **adaptive, data-centric, and cloud-native**. Static perimeter defences and coarse-grained policies are no longer sufficient in environments characterized by microservices, elastic multi-cloud deployments, and highly dynamic user contexts.

By synthesizing research on cloud security and privacy, Zero Trust architectures, risk-based adaptive security, and context-aware access control and privacy, this article has outlined an adaptive data security architecture with four key

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 4, Issue 6, November-December 2022 |

DOI:10.15662/IJEETR.2022.0406011

layers and a unifying MAPE-based orchestration loop. While significant implementation and research challenges remain—particularly around policy complexity, explainability, and cross-jurisdictional governance—the architectural principles presented here provide a blueprint for securing consumer-scale digital applications in an increasingly complex threat and regulatory landscape.

REFERENCES

1. Beer Mohamed, M. I., Hassan, M. F., Safdar, S., & Saleem, M. Q. (2021). Adaptive security architectural model for protecting identity federation in service oriented computing. *Journal of King Saud University – Computer and Information Sciences, 33*(5), 580–592. https://doi.org/10.1016/j.jksuci.2019.03.004
2. Butt, U. A., Amin, R., Mehmood, M., Aldabbas, H., Alharbi, M. T., & Albaqami, N. (2022). Cloud security threats and solutions: A survey. *Wireless Personal Communications*. https://doi.org/10.1007/s11277-022-09960-z
3. Kolla, S. (2022). Effects of OpenAI on Databases. International Journal Of Multidisciplinary Research In Science, Engineering and Technology, 05(10), 1531-1535. https://doi.org/10.15680/IJMRSET.2022.0510001
4. Calvo, M., & Beltrán, M. (2022). A model for risk-based adaptive security controls. *Computers & Security, 115*, 102612. https://doi.org/10.1016/j.cose.2022.102612
5. 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 NIST CSRC
6. Herrera, J. L., Chen, H., Berrocal, J., Murillo, J. M., & Julien, C. (2022). Context-aware privacy-preserving access control for mobile computing. *Pervasive and Mobile Computing, 87*, 101725. https://doi.org/10.1016/j.pmcj.2022.101725
7. Vangavolu, S. V. (2022). IMPLEMENTING MICROSERVICES ARCHITECTURE WITH NODE.JS AND EXPRESS IN MEAN APPLICATIONS. International Journal of Advanced Research in Engineering and Technology (IJARET), 13(08), 56-65. https://doi.org/10.34218/IJARET_13_08_007
8. Jiang, B., Seif, M., Tandon, R., & Li, M. (2021). Context-aware local information privacy. *IEEE Transactions on Information Forensics and Security*. https://doi.org/10.1109/TIFS.2021.3087350
9. Kayes, A. S. M., Kalaria, R., Sarker, I. H., Islam, M. S., Watters, P. A., Ng, A., Hammoudeh, M., Badsha, S., & Kumara, I. (2020). A survey of context-aware access control mechanisms for cloud and fog networks: Taxonomy and open research issues. *Sensors, 20*(9), 2464. https://doi.org/10.3390/s20092464
10. Kumar, R., & Goyal, R. (2019). On cloud security requirements, threats, vulnerabilities and countermeasures: A survey. *Computer Science Review, 33*, 1–48. https://doi.org/10.1016/j.cosrev.2019.05.002
11. Pericherla, S. (2022). Cloud computing threats, vulnerabilities and countermeasures: A state-of-the-art. *ISeCure*. https://doi.org/10.22042/ISECURE.2022.312328.718
12. 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 NIST CSRC
13. Sun, P. (2020). Security and privacy protection in cloud computing: Discussions and challenges. *Journal of Network and Computer Applications, 160*, 102642. https://doi.org/10.1016/j.jnca.2020.102642