

# Integrating Microservices And Event-Driven Design For Cloud-Native Transformation

**Vinay Babu Gurram**

*Independent Researcher, USA*

## Abstract

In terms of monolithic multi-component and micro-service-based cloud-native patterns as well as event-driven design, the environment of enterprise software architecture has changed considerably. This change meets basic scalability, response, and more effective resource usage needs in current digital settings. Together with event-driven patterns in cloud-native infrastructure and a microservices architecture, companies can develop modular, fault-tolerant, and autoscaling systems. Although event-driven architecture is able to support asynchronous communication, which results in fewer service dependencies and a more responsive system, microservices subdivide complex applications into smaller, independent parts according to business capabilities. This combination of architecture has been made possible through containerization, orchestration platforms, and infrastructure as code, and cannot be matched with any other architectural combination in terms of operational efficiency and flexibility. Yet enterprises face significant challenges such as distributed data consistency, operational complexity within loosely coupled services, observability needs, and governance of asynchronous workflows. The envisioned Cloud-Native Reference Framework aligns microservice modularity with event-driven agility across several architectural levels, resolving application service, messaging infrastructure, data management, and cloud infrastructure issues. Legacy system migration involves phased transformation methodologies using patterns like the Strangler Fig pattern, domain-driven decomposition, and prudent mechanisms of data replication. Akin to resolving technical competencies, organizational transformation adopting DevOps culture, continuous delivery practices, and advanced automation is also key to success. Though their integration brings about complexity in such areas as transaction management, security, and costing optimization, the resulting architectures reflect greater flexibility in responding to fluctuating business requirements and operational needs, and hence enable enterprises to achieve long-term competitive advantage in electronic markets.

**Keywords:** Microservices, Artificial Intelligence, Autonomous Cloud Systems, Sustainability, Predictive Scaling.

## 1. Introduction

From monolithic systems to service-oriented architecture (SOA) to the current paradigms of microservices and event-driven systems, the landscape of business software architecture has evolved repeatedly over the last thirty years. This evolutionary route reflects the growing demands placed on corporate technology infrastructure: the need to meet more scalability, better agility, and better resource efficiency in the period of rapid digital transformation. The microservices architectural pattern has been developed as an answer to monolithic applications' limitations, in which tightly coupled components and centralized databases

presented great obstacles for independent deployment and scaling [1]. Research that examines how software architecture has evolved indicates that microservices are a logical extension of SOA, which addresses the granularity and coupling issues that plagued the earlier service-oriented solutions but enables organizations to separate applications into small-grained services that can be created, deployed, and scaled individually [1].

Cloud-native transformation represents a complete re-thinking of the manner in which business applications are developed, deployed, and managed. In comparison to the traditional architecture, where parts are tightly interconnected and the communication is tightly synchronized, cloud-native systems are characterized by modularity, resilience, and independent scaling. Among the most attractive ways of meeting these objectives is the microservices architecture with an event-driven design (EDA) that allows businesses to design dynamic systems that maintain balance on operations and are adaptable to the evolving needs. In contrast to the fundamental principles of monolithic architecture, the microservices architecture focuses on the division of applications into tiny, autonomous services that can be compared to specific business capabilities. Particularly, the bounded contexts defining clear boundaries for service responsibilities and data ownership, this architectural style draws heavily from domain-driven design ideas. Each microservice is a separate entity with its own data store so that teams may create, deploy, and scale services without the coordinating overhead found in monolithic systems. Empirical research has shown that the path to microservices adoption involves several aspects of technical and organizational change, such as the necessity to implement new development paradigms, operational capabilities, and governance models differing fundamentally from those used in managing monolithic systems [2]. Organizations that have made the move to microservices architectures have noted that they have experienced tremendous challenges with distributed system complexity, with special challenges arising in testing strategies, deployment orchestration, and ensuring consistency between independently evolved services [2].

Microservices and event-driven architecture congregate several imperative needs of today's enterprise systems. Microservices break up monolithic applications into separate, independently deployable components grouped around business capabilities, with each service having its own data store and sharing data with other services using clearly defined interfaces [1]. Event-driven designs enable separated asynchronous communication that enhances system responsiveness, therefore enabling services to react to state changes without direct binding to event producers [2]. This architectural synthesis allows for previously unheard-of degrees of flexibility and operational economy when applied to cloud-native infrastructure—drawing on containerization, orchestrating tools like Kubernetes, and infrastructure as code. Adoption of microservices, however, introduces great complexity in areas including service discovery, inter-service communication, distributed tracing, and monitoring, therefore requiring businesses to create sophisticated operational capabilities. Moreover, in monolithic settings [1][2] there is no need for tooling ecosystems.

While there is theoretical attractiveness in blending these paradigms, businesses are confronted with substantial practical implementation barriers. Issues of data consistency in distributed systems, the intricacy of coordinating asynchronous workflows, operational visibility in loosely coupled services, and the management of event-driven interactions are only partially addressed in the current literature [2]. The shift to microservices from monolithic architectures requires close attention to organizational readiness, as effective adoption relies not just on technical ability but on cultural change toward DevOps practices, continuous delivery, and cross-functional team organizations [1].

**Table 1: Cloud-Native Architectural Evolution and Auto-Scaling Characteristics [1][2]**

| Aspect                | Cloud-Native Microservices                              | Traditional Auto-Scaling Approaches |
|-----------------------|---|-------------------------------------|
| Architectural Pattern | Dominant design pattern in cloud-native implementations | Reactive threshold-based mechanisms |

| System Characteristics | Horizontal scalability, distributed state management, polyglot persistence | Fixed threshold triggers with temporal lag                     |
|------------------------|--|--|
| Operational Complexity | Distributed system complexities require sophisticated orchestration        | Over-provisioning buffers to maintain service level agreements |
| Scaling Methodology    | Independent service scalability with containerization                      | Rule-based approaches with predefined thresholds               |
| Resource Utilization   | Theoretical optimal levels requiring intelligent management                | Conventional utilization falling below optimal levels          |

## 2. Architectural Foundations and Design Paradigms

### 2.1 Microservices Architecture Principles

With contract-based communication and well-defined APIs, this freedom lets businesses adopt new technologies in stages and respond swiftly to market changes. Offering cross-cutting issues like authentication, rate limiting, and request routing, the API gateway pattern usually serves as the first point of contact between outside customers and the microservices system. Studies that looked into microservices architectural styles found that effective implementations tend to break down applications into tens to hundreds of discrete components, each of which comprises 100 to 1,000 lines of code on average, much smaller than the millions of lines common in monolithic applications [3]. The bounded context rule prevents services from having ambiguous domain boundaries, and research has shown that well-structured microservices have coupling measures 70-80% less than comparable monolithic solutions, as measured by indicators like coupling between object classes and response for class dependencies [3].

The decentralized data management principle is a foundation of the design of microservices. Instead of jointly sharing a database, each service has sole control over its data model and persistence layer. This design removes database-level coupling and allows services to choose data storage technologies tuned to their individual needs—a pattern referred to as polyglot persistence. Systematic mapping research of microservices deployments shows that companies that implement this architecture normally use several disparate database technologies, with about 60% of systems surveyed using two or more different types of databases, and 30% using three or more types of data storage solutions in their service ecosystem [4]. However, this independence brings challenges of data consistency and transactional integrity that need to be handled through other patterns like eventual consistency and distributed sagas. Evaluation of data management issues in microservices environments suggests that ensuring consistency among distributed services is among the top three technical challenges cited by practitioners, with 68% of those organizations recognizing distributed data management as a key deployment challenge deserving specialized patterns and sensitive architectural attention [4].

Autonomy for services is not limited to data management alone but also covers independent deployment lifecycles, technology stack choice, and scaling properties. Teams are able to develop individual services without affecting the overall system, allowing for continuous delivery and minimizing the risk of changes. Empirical research studying deployment practices in microservices environments has reported that organizations experience tens of times higher deployment frequencies relative to monolithic architecture, with some companies experiencing deployment rates of over 1,000 deployments per day on their entire microservices portfolio, a shift from the legacy quarterly or monthly release schedules [3]. With contract-based communication and well-defined APIs, this freedom lets businesses adopt new technologies in stages and respond swiftly to market changes. Offering cross-cutting issues like authentication, rate limiting, and request routing, the API gateway pattern usually serves as the first point of contact between outside customers and the microservices system. Empirical work on API gateway deployments shows that such components manage key integration tasks, with gateway patterns emerging in about 75% of cited microservices architectures as a vehicle for managing external communication complexity and supporting unified access points for client applications [4].

**Table 2: AIOps Implementation and Microservice Orchestration Patterns [3][4]**

| Component                   | Microservice Architecture Features                        | AIOps Operational Characteristics                      |
|-----------------------------|---|--|
| Service Structure           | Decomposition according to business capabilities          | Massive operational telemetry data generation          |
| Communication Protocol      | Lightweight protocols such as REST or message queues      | Machine learning-based anomaly detection systems       |
| Data Management             | Each service maintains an independent database            | Incident detection with reduced false positive rates   |
| Infrastructure Requirements | Containerization technologies and orchestration platforms | Predictive analytics for capacity planning             |
| Adaptive Capabilities       | Dynamic behavior patterns challenge static configurations | Reinforcement learning for dynamic policy optimization |

### 3. Architectural Convergence: Integration Framework and Design Patterns

#### 3.1 Multi-Layer Integration Model

The convergence of microservices and event-driven architecture into cloud-native implementations calls for a structured integration plan addressing problems on several architectural layers. The suggested integration model classifies system building blocks into four different but interconnected layers, where each performs specialized functions while contributing to system cohesion. Systematic reviews of microservices literature trends in research have found that architectural styles and best practices account for about 23% of all publications related to microservices, reflecting heavy academic and industry interest in defining organized integration frameworks [5]. In addition, research development analysis shows that publications that tackle architectural issues have increased by 300% from 2014 to 2017, showcasing the increasing maturity and sophistication of microservices environments in production [5].

The Application Layer contains discrete microservices and their respective APIs. Services here are designed based on bounded context principles, having well-defined functional boundaries and data ownership. RESTful APIs or GraphQL endpoints are used to handle synchronous client-facing operations, and services also serve as event producers and consumers, engaging in asynchronous workflows. Studies that analyze API design patterns in microservices architectures report that RESTful interfaces continue to be the most widely chosen method for communication between services, showing up in more than 80% of reported implementations, although newer GraphQL methods show increasing usage, especially in systems that need flexible query abilities and minimal data over-fetching [5]. The duality of services—handling explicit requests and responding to events—requires diligent design to keep synchronization and decoupling of concerns and synchronization and asynchronous execution paths. Experiments comparing architectural quality attributes show that well-designed service boundaries with well-separated synchronous and asynchronous concerns lower coupling values by 40-50% in comparison to implementations where synchronous and asynchronous concerns are merged [6].

The Messaging Layer forms the nervous system of event-based architectures, which supports asynchronous communication between services. The event bus, message brokers, and streaming platforms are placed at this level and offer assurance of reliable event delivery with controllable guarantees on ordering, durability, and the provision of delivery semantics. The choice of technology at this layer has a major influence on the system properties: Kafka is better in the areas where the high throughput of event streams and replay functions are needed, RabbitMQ is more flexible with the routing features, with the good support of AMQP protocol, and cloud native solutions are characterized by managed services with less operational overhead. Comparative analysis of messaging technologies in microservices ecosystems indicates that Apache Kafka has emerged as the predominant choice for event streaming scenarios, with adoption rates exceeding 60% in large-scale deployments, while RabbitMQ maintains a strong presence in systems requiring sophisticated

message routing patterns with adoption approaching 35% of surveyed implementations [5]. Message schemas and event contracts at this layer require governance to prevent breaking changes that could disrupt consuming services, with empirical studies revealing that inadequate schema management accounts for approximately 30% of integration failures in event-driven microservices systems [6].

To guarantee data consistency without compromising service autonomy, the Data Layer uses methods including CQRS, event sourcing, and distributed caching to control states in distributed services. Event sourcing tracks the whole chronology of state changes, so one may inquire into history and determine the system's condition at any time. CQRS explicitly divides operations on the command-side, the one that changes state, and operations on the query-side, the one that reads, enabling each concern to scale independently and to be chosen independently using technology. Research examining data management patterns documents that CQRS implementation appears in approximately 15-20% of microservices architectures requiring complex read-write separation, while event sourcing adoption remains more selective at 10-15% due to increased complexity and operational overhead [5]. Distributed state management may also incorporate technologies such as Redis for caching and Apache Cassandra or MongoDB for horizontally scalable data persistence, with studies indicating that distributed caching layers reduce database query load by factors ranging from 10:1 to 50:1 in well-optimized systems [6].

**Table 3: Container Orchestration Design Patterns and Reinforcement Learning Resource Management [5][6]**

| Pattern Category         | Container Design Patterns   | Deep Reinforcement Learning Characteristics                      |
|--------------------------|---|--|
| Architectural Approach   | Single-container, single-node, multi-container, multi-node patterns | Neural network policy processing job characteristics             |
| Monitoring Strategy      | Sidecar patterns for cross-cutting concerns                         | Policy gradient algorithms learning optimal strategies           |
| Resource Consumption     | Auxiliary containers providing comprehensive observability          | Multi-objective reward functions balancing competing objectives  |
| Operational Integration  | Separation of concerns with minimal code modification               | Learning scheduling policies from operational experience         |
| Performance Optimization | Modular, composable, containerized applications                     | Convergence to near-optimal policies through training iterations |

#### 4. Cloud-Native Transformation Strategy and Implementation

##### 4.1 Modernization Roadmap and Migration Patterns

One step cannot be accomplished for the migration of old monolithic systems to cloud-based infrastructures with microservices and event-driven designs; hence, the migration must be progressive. With a balance between architectural advancement and company continuity. Any effort to rewrite everything is unacceptably risky, and it interferes with current operations, whereas established patterns of migration make it possible to effect change step by step. Empirical studies examining migration from monolithic to microservices architectures document that successful transformations typically follow iterative approaches, with organizations reporting migration durations ranging from 12 to 36 months, depending on system complexity and organizational readiness [7]. Research analyzing real-world migration experiences reveals that phased decomposition strategies significantly reduce risk, with organizations achieving functional microservices deployments handling production traffic within 3 to 6 months of initiating transformation efforts, compared to multi-year timelines associated with complete rewrites [8].

The Strangler Fig pattern provides a systematic approach to gradually replacing monolithic functionality with microservices. This pattern is named after the strangler fig plant, which grows around the host trees

and directs a particular functionality to new microservices and leaves the rest of the requests to the monolith. As microservices demonstrate their effectiveness and capability, more and more functionality is shifted over until the monolith can be switched off completely. Routing logic is implemented as an API gateway or reverse proxy, which routes requests to the monolith or microservices depending on URL patterns, headers, or other factors. Case study analysis of strangler pattern implementations demonstrates that organizations typically begin by extracting non-critical, loosely coupled functionality representing approximately 15-20% of overall system capabilities, validating the approach before migrating core business logic [8]. This approach minimizes risk by enabling incremental migration, facilitates rollback if issues arise, and allows the organization to realize benefits from completed microservices while work continues on remaining components, with documented improvements in deployment frequency increasing from quarterly to weekly cycles after extracting initial service modules [7].

Domain decomposition constitutes the analytical foundation for transformation, identifying appropriate service boundaries based on business capabilities rather than technical convenience. Domain-driven design techniques, including context mapping and bounded context identification, guide this decomposition. Each identified bounded context becomes a candidate microservice, with clear responsibilities and minimal dependencies on other services. Event storming workshops, bringing together domain experts and technical teams, effectively surface business events, commands, and aggregates that inform both service boundaries and event designs. Research on domain-driven design application in microservices contexts indicates that collaborative modeling sessions involving cross-functional teams of 6 to 12 participants prove most effective in identifying appropriate service boundaries and minimizing later refactoring needs [7]. Proper domain decomposition proves critical; poorly defined boundaries lead to excessive inter-service communication, distributed monoliths, and compromised maintainability, with studies showing that well-bounded services exhibit significantly lower coupling metrics and require fewer subsequent boundary adjustments [8].

The challenges of data migration are specific to legacy monoliths, which are normally based on shared databases with complicated schemas and referential integrity constraints. The data replication plan will use synchronization to ensure consistency between the database of the monolith and the data stores of the emerging microservice in the transition process. Database transaction logs are tracked using change data capture (CDC) tools, which propagate the changes to microservice databases or event buses. This allows micro services to have their own data stores, but with consistency with the monolith, which allows a gradual transfer of data ownership. Analysis of data migration strategies reveals that dual-write periods, where both monolithic and microservice databases receive updates, typically extend 2 to 4 months to ensure data consistency validation before complete ownership transfer [8]. As services mature and business confidence grows, writes can shift from the monolith to microservices, ultimately establishing services as authoritative sources for their domains, with monitoring and reconciliation mechanisms ensuring data integrity throughout the transition [7].

**Table 4: Production Cluster Management and Quality-of-Service-Aware Scheduling [7][8]**

| Management Aspect     | Borg Cluster Management System                            | Quasar QoS-Aware System  |
|-----------------------|---|--|
| Workload Composition  | Long-running services with opportunistic batch processing | Heterogeneous workloads, including web search and analytics        |
| Resource Optimization | Bin-packing algorithms for workload consolidation         | Collaborative filtering-based performance prediction               |
| Utilization Strategy  | Co-location of complementary workloads                    | Resource-efficient provisioning with minimal performance variation |

|                       |  |  |
|-----------------------|--|--|
| Capacity Management   | Batch workload utilizing service workload throughs | Accurate resource allocation decisions, maintaining service objectives |
| Performance Isolation | Containerization and resource monitoring           | Quality-of-service guarantees through intelligent prediction           |

## 5. Challenges, Trade-offs, and Governance Considerations

### 5.1 Data Consistency and Transaction Management

Distributed systems cannot provide the strong consistency guarantees of monolithic applications, in exchange for providing availability and partition tolerance, as defined in the CAP theorem. Microservices, which have their own data stores, are not capable of traditional ACID services across services. This fundamental constraint necessitates alternative approaches based on eventual consistency, where the system temporarily permits inconsistencies but guarantees convergence to a consistent state given sufficient time without new updates. Systematic grey literature analysis examining microservices challenges reveals that data consistency management represents one of the most frequently reported difficulties, appearing in approximately 34% of practitioner reports and technical blogs discussing microservices implementation obstacles [9]. Furthermore, research analyzing transaction management patterns documents that organizations transitioning from monolithic to microservices architectures report significant challenges in maintaining data consistency, with 68% of surveyed practitioners identifying distributed transaction management as a primary technical concern requiring substantial architectural redesign efforts [9].

The saga pattern, as previously discussed, provides a mechanism for coordinating distributed transactions through sequences of local transactions with compensating actions. However, sagas introduce complexity in error handling and state management. Services must be designed idempotently—producing identical results regardless of how many times an operation is executed—to handle message redelivery safely. Compensating transactions must carefully reverse the effects of previously completed steps, which proves straightforward for some operations (canceling a reservation) but challenging for others (reversing a completed shipment). Analysis of microservices architectural challenges indicates that implementing compensating transactions adds significant development complexity, with practitioners reporting that saga-based transaction coordination requires substantially more implementation effort compared to traditional database transactions [9]. Furthermore, sagas may leave the system in intermediate states visible to users during execution, requiring careful UX design to manage expectations and prevent confusion, with studies documenting that managing intermediate consistency states represents a recurring challenge in event-driven microservices implementations [10].

Event sourcing, by maintaining the complete history of state changes, provides an alternative consistency model. Since events represent immutable facts about past occurrences, conflicts between concurrent operations can be detected and resolved when events are applied. However, event sourcing introduces its own complexities: managing schema evolution as event definitions change over time, handling large event stores that grow continuously, and maintaining performant query capabilities over event-sourced data. Research examining microservices data management patterns reveals that event sourcing adoption remains limited due to implementation complexity, with grey literature analysis showing that only a small percentage of production systems fully implement event sourcing despite its theoretical advantages [9]. Snapshot mechanisms, which periodically persist derived state to accelerate reconstruction, help manage performance but add complexity to system design, with practitioners reporting that optimizing event replay performance requires careful consideration of snapshot frequency and storage strategies [10].

The selection of consistency models should align with business requirements rather than technical convenience. Financial transactions typically require strong consistency or carefully designed sagas with robust compensating transactions, while social media features may tolerate eventual consistency without user impact. Hybrid approaches often prove optimal, applying stronger consistency where business risks are highest while accepting eventual consistency elsewhere to maintain system responsiveness and availability. Systematic analysis of microservices challenges documents that selecting appropriate

consistency guarantees for different system components represents a critical architectural decision, with practitioners emphasizing the need to balance consistency requirements against system complexity and performance characteristics [9].

## Conclusion

The convergence of microservices and event-driven design in cloud-native architecture is a paradigm shift in the development of enterprise systems and satisfies the key user demands of scalability, resilience, and rapid evolution required by modern digital enterprises. This integration helps organizations to build dynamic systems that can be responsive to different needs, and at the same time, their stability can be maintained in an environment of operational stability of modular and independently deployable services and communicating through asynchronous event patterns. The path to cloud native microservices requires more than simply a technical reorganization, but organizational change, adoption of DevOps culture, product-focused groups, and platform engineering skills, which are entirely new as compared to traditional IT operations. Leadership dedication to incremental migration plans, investment in observability and automation facilities, and tolerance of the higher operational complexity of distributed systems are the keys to success. The issues that have been noted during this exposition, such as the data consistency in distributed environments, complexity in the operation, security, and governance issues, are not barriers to adoption but critical design factors that need to be properly addressed and architectural patterns that have been tested. Canonical solutions, such as sagas to distributed transactions, circuit breakers to resiliency, event sourcing to auditability, and the outbox pattern to reliable event publication, offer solution patterns that have proven effective and practical when used intelligently based on required organizational contexts. The proposed Cloud-Native Reference Framework summarizes these aspects as an organized strategy on the scale of scalability, resilience, observability, and sustainability facets, and is a conceptual framework shaping architecture decisions to fit specific situations instead of blueprints. Companies undertaking cloud-native transformation must focus on capabilities of platform engineering that put in place shared infrastructure and tooling, domain-oriented design to thoughtfully identify service boundaries, adopt eventual consistency where business needs allow, have explicit governance of APIs and event definitions, maintain team independence, and do transformation in small steps through validated pilot projects before going enterprise-wide. The architecture concepts that have resulted in this convergence, such as modularity, loose coupling, asynchronous communication, and automated operations, offer sustainable baselines of enterprise architecture, which are likely to persist as the particular implementation technologies keep changing, but allow organizations to build systems based on the requirements of their ongoing digital transformation in an era of constant change.

## References

- [1] Nicola Dragoni, et al., "Microservices: Yesterday, Today, and Tomorrow," SpringerNature Link, 2017. [Online]. Available: [https://link.springer.com/chapter/10.1007/978-3-319-67425-4\\_12](https://link.springer.com/chapter/10.1007/978-3-319-67425-4_12)
- [2] Pooyan Jamshidi, et al., "Microservices: The Journey So Far and Challenges Ahead," IEEE, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8354433>
- [3] Cesare Pautasso, et al., "Microservices in Practice, Part 1: Reality Check and Service Design," IEEE, 2017. [Online]. Available: <https://ieeexplore.ieee.org/document/7819415>
- [4] Claus Pahl, Pooyan Jamshidi, "Microservices: A Systematic Mapping Study," ResearchGate, 2016. [Online]. Available: [https://www.researchgate.net/publication/302973857\\_Microservices\\_A\\_Systematic\\_Mapping\\_Study](https://www.researchgate.net/publication/302973857_Microservices_A_Systematic_Mapping_Study)
- [5] Paolo Di Francesco, et al., "Research on Architecting Microservices: Trends, Focus, and Potential for Industrial Adoption," IEEE, 2017[Online]. Available: <https://ieeexplore.ieee.org/document/7930195>
- [6] Davide Taibi; Valentina Lenarduzzi, "On the Definition of Microservice Bad Smells," IEEE, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8354414>
- [7] Gerald Schermann, et al., "Continuous Experimentation: Challenges, Implementation Techniques, and Current Research," IEEE, 2018. [Online]. Available: <https://ieeexplore.ieee.org/document/8255793>

- [8] Armin Balalaie et al., "Microservices Architecture Enables DevOps: an Experience Report on Migration to a Cloud-Native Architecture," ResearchGate, 2016. [Online]. Available: [https://www.researchgate.net/publication/298902672\\_Microservices\\_Architecture\\_Enables\\_DevOps\\_an\\_Experience\\_Report\\_on\\_Migration\\_to\\_a\\_Cloud-Native\\_Architecture](https://www.researchgate.net/publication/298902672_Microservices_Architecture_Enables_DevOps_an_Experience_Report_on_Migration_to_a_Cloud-Native_Architecture)
- [9] Jacopo Soldani, et al., "The pains and gains of microservices: A Systematic grey literature review," ScienceDirect, 2018. [Online]. Available: <https://www.sciencedirect.com/science/article/abs/pii/S0164121218302139>
- [10] Hui Kang, et al., "Container and Microservice Driven Design for Cloud Infrastructure DevOps," IEEE, 2016. [Online]. Available: <https://ieeexplore.ieee.org/document/7484185>