



# Architectural Design Patterns for Fault-Tolerant Distributed Software Systems

Dmitrii Bezfamilnyi

Software Development Expert, Russia / Moscow

Received: 31 Mar 2026; Received in revised form: 28 Apr 2026; Accepted: 02 May 2026; Available online: 06 May 2026

**Abstract**— This study aims to develop a conceptual architectural model of a fault-tolerant distributed software system based on the integration of microservice architecture, resilience mechanisms, and adaptive control. The study is conducted as a structured systematic review with elements of qualitative analysis of scientific publications addressing architectural solutions in cloud, microservice-based, and distributed computing environments. The main focus is placed on the relationship between system architecture, failure propagation patterns, and performance and resilience indicators. Key fault-tolerance mechanisms are examined, including service decomposition, asynchronous interaction, horizontal scaling, load balancing, and the distribution of functions across system layers. It is established that the impact of architectural decisions is systemic in nature and manifests through their combined influence on latency, overload resilience, and recovery capability. The study demonstrates that fault tolerance is not the result of applying isolated mechanisms, but emerges from the coordinated interaction of architectural layers. A conceptual model of a fault-tolerant distributed system is proposed, reflecting the interconnection between infrastructure, platform, application, failure management, and adaptive intelligent layers. The results obtained make it possible to consider architecture as a tool for managing system behavior under conditions of failure and uncertainty, requiring the integration of adaptive control and predictive mechanisms. The article may be of interest to researchers in distributed systems, software architects, and practitioners involved in the design of high-load and fault-tolerant systems.

**Keywords**— distributed systems, fault tolerance, microservice architecture, cloud computing, architectural patterns, adaptive systems, resilience.

## I. INTRODUCTION

Distributed software systems today operate in environments where a single machine is no longer sufficient. These include cloud platforms, edge computing, device networks, services with continuous data exchange, and systems related to artificial intelligence [1]. As the number of nodes and connections increases, so does the load on data transmission, coordination, and recovery processes. As a result, failures cease to be rare events and become a normal part of large-scale system operation. This clearly illustrates the scale of the problem. For example, if a cluster consists of 10,000 servers and the mean time between failures of a single server is 30

years, the system will still experience, on average, one server failure per day [2].

At the same time, both the environment and the internal organization of applications become more complex. Instead of a single monolithic program, systems increasingly rely on sets of smaller services, with part of the processing moved closer to the data source while another part remains centralized [4]. This approach provides flexibility and enables faster scalability. However, its trade-offs are equally evident. A failure in one service can propagate to others, an overloaded intermediary node can disrupt the entire processing chain, and communication failures between layers may lead not to a complete

shutdown but to partial, difficult-to-detect degradation. Therefore, it is no longer sufficient to simply replicate individual nodes or restart failed processes.

The aim of the study is to develop an original architectural model of a fault-tolerant distributed software system in which microservice architecture, orchestration tools, resilience mechanisms, and adaptive management capabilities are integrated into a unified framework. To achieve this goal, the study addresses the following objectives:

- to analyze existing architectural solutions used in cloud, microservice-based, and distributed systems;
- to identify the mechanisms that affect resilience under failures, overloads, and connectivity disruptions;
- to compare the quantitative effects of architectural solutions, including latency, recovery, and stability;
- to develop an original model that integrates these solutions into a unified architectural structure.

It is assumed that the fault tolerance of a distributed system is determined not by a single mechanism, but by their coordinated combination. Service isolation, load distribution, redundancy, asynchronous interaction, system observability, and recovery mechanisms provide the greatest effect when they are initially integrated at the architectural level, rather than introduced separately after problems arise.

The scientific novelty of the study lies in shifting the focus from individual fault-tolerance mechanisms to system behavior under real failures and overload conditions. Unlike traditional approaches, where architectural solutions are considered in isolation, this study analyzes their combined impact on latency, failure propagation, and recovery time. This makes it possible to establish a relationship between system structure and its actual resilience, and to propose a model that describes fault tolerance as the result of interactions between architectural layers rather than a set of independent solutions. Unlike existing approaches, this study integrates architectural layers into a unified model

linking system structure with measurable resilience indicators.

## II. MATERIALS AND METHODS

The study was conducted in the format of a structured systematic review with elements of qualitative content analysis. The procedure for selecting and analyzing sources was formalized and reproducible, ensuring the comparability of results when examining architectural solutions in distributed systems operating under failure and overload conditions.

The literature search was carried out in international scientific databases, including Google Scholar, ScienceDirect, SpringerLink, and MDPI, over the period from January to March 2026. The following precise search queries were used: "distributed systems AND fault tolerance," "cloud architecture AND resilience AND latency," "microservices AND failure recovery AND availability," and "distributed architecture AND reliability AND MTTR." Extended combinations with logical operators AND/OR were also applied. The initial search yielded 148 publications in Google Scholar, 62 in ScienceDirect, 39 in SpringerLink, and 27 in MDPI. After merging results and removing duplicates, a dataset of 176 unique sources was formed.

The selection process was conducted in stages. At the title and abstract screening stage, 109 publications were excluded as they did not contain analysis of architectural solutions or were not related to system fault tolerance. During the full-text review stage, an additional 52 studies were excluded, as they focused on algorithms or hardware aspects without reference to system architecture or lacked quantitative indicators. The final sample included 15 sources that met the objectives of the study.

The inclusion criteria required that publications contain analysis of architectural solutions in distributed systems with measurable resilience characteristics. These characteristics included processing latency, recovery time after failure, system availability, and resistance to overload. Studies without quantitative metrics or without descriptions of system behavior under failure conditions were excluded. Source quality was also assessed, and only peer-reviewed journal articles containing either

empirical data or formalized architectural models were included in the sample.

The analytical procedure was structured as a sequential examination of publication content. At the first stage, text fragments describing types of architectures, failure-handling approaches, and system performance indicators were extracted. Recurrent solutions and approaches were then compared, allowing them to be grouped into stable categories. As a result, the following blocks were identified: architectural system decomposition, failure isolation mechanisms, recovery approaches, and adaptive load management mechanisms. A subsequent review of the sources refined the boundaries of these categories and eliminated overlaps.

The identified categories were directly used in forming the research results. Based on them, analytical tables were constructed to reflect the relationship between architectural solutions and changes in latency, recovery, and availability metrics. These same categories also formed the basis for the development of the author's model, which describes fault tolerance as the result of interactions between architectural layers.

The limitations of the study are associated with the use of only English-language publications and a limited set of databases. In addition, the sample is biased toward modern architectures, primarily cloud-based and service-oriented systems, which may limit the generalizability of the findings to other types of distributed solutions. Nevertheless, the applied approach provides sufficient analytical depth and reveals stable relationships between system architecture and its behavior under failure conditions.

### III. RESULTS

The development of distributed software systems in recent years has been accompanied by simultaneous growth in scale, increasing internal complexity, and a rising dependence on continuous interaction between components. Computational environments, including cloud platforms, microservice architectures, and distributed learning models, operate under conditions of constant dynamics, where changes in the state of nodes occur continuously [1, 4]. Under these conditions, failure is

no longer treated as an exceptional event but becomes a regular characteristic of system operation. Architecture adapts to this reality: resilience is no longer added after system construction but is embedded in the way systems are decomposed, interact, and redistribute functions.

Within this context, two stable directions of analysis emerge. The first is related to the impact of scale on the frequency and nature of failures. The second concerns how architecture redistributes the consequences of these failures and maintains system execution. These directions do not exist independently. As the number of nodes increases, the system's dependence on interaction mechanisms intensifies, and architectural decisions begin to define the limits of feasible scalability [7].

As the number of nodes grows, the system transitions into a regime where failure becomes a statistically expected condition. Even with a high level of reliability of individual servers, the aggregate system behavior changes. The frequency of failures is determined not by the quality of a single component, but by their number and interaction patterns [5]. In distributed environments with thousands of nodes, a situation arises in which, at any given moment, a portion of components is in a failed or degraded state. Table 1 presents quantitative parameters reflecting the impact of scale on failure frequency and overall system behavior.

The presented values capture a shift in system operation as scale increases. Failure ceases to be a rare event and becomes distributed over time. As the number of nodes grows, the interval between failures decreases, and the system begins to operate under conditions of continuous local disruptions. At the same time, the impact of each failure changes. In synchronous computational scenarios, even a single disruption may halt the entire process. In asynchronous environments, the effect is distributed and manifests as latency, partial request loss, or reduced throughput.

Scale affects not only the frequency of failures but also the structure of dependencies between components. In compact systems, disruptions tend to remain localized. In large-scale systems, interaction chains become longer, and a failure begins to affect a greater number of elements. Under these conditions,

scale increases system sensitivity to coordination and raises the cost of any delay or node loss [13].

Table 1 – Quantitative impact of scale on system failures (Compiled by the author based on source: [2])

Scenario	Parameters	Observed Effect
General-purpose cluster	10,000 servers, MTBF per server: 30 years	On average, 1 server failure per day
Synchronous AI training on TPU	9,000 TPUs, MTBF per TPU: 5 years	Computation interruption approximately every 4 hours
Service on mature servers	2,000 servers, annual crash rate: 5%	Expected 1 machine failure every 3–4 days
Llama3 training	54 days, 16K GPUs	Approximately 320 hardware failures and 10% throughput degradation

A shift is observed from the problem of component reliability to the problem of interaction organization. Improving the MTBF of an individual node does not eliminate the regular occurrence of failures at the system level. At sufficient scale, even extremely rare events begin to occur continuously. In this context, resilience is determined not by hardware characteristics but by how the system responds to inevitable disruptions and how it redistributes load at the moment they occur [11].

Architectural solutions used in distributed systems form a stable structure in which different

mechanisms play distinct roles with respect to failure [6]. Some solutions reduce coupling between components and limit the scope of failure propagation. Others ensure continued execution by redistributing load and altering processing routes. These mechanisms operate simultaneously and reinforce one another, creating a system capable of functioning under conditions of continuous local disruptions [12]. Table 2 presents the main architectural mechanisms used in microservice-based systems.

Table 2 – Architectural mechanisms of fault tolerance in microservice-based systems (Compiled by the author based on source: [4])

Architectural Mechanism	Implementation	Impact on Fault Tolerance	Limitations
Microservice decomposition	Independent services with API-based interaction	Failure isolation, elimination of a single point of failure	Increased management complexity
Horizontal scaling	Service replication (Kubernetes, Docker Swarm)	Resilience to overloads and node failures	Requires orchestration and load balancing
Load balancing	Load balancer across service instances	Prevents overload of individual services	Additional infrastructure required
Asynchronous interaction	Event-driven architecture (Kafka, RabbitMQ)	Reduces cascading failures and coupling	Increased system complexity
Edge–cloud segmentation	Separation of processing across layers	Local resilience under network/cloud failures	Requires data synchronization
API Gateway	Centralized entry and control point	Traffic control, overload protection	Potential bottleneck

Decomposition into independent services transforms the structure of the system. Failures no longer propagate across the entire architecture and remain confined within a single functional block. However, decomposition alone does not sustain execution. In the absence of additional mechanisms, the system isolates the failure but does not compensate for its consequences. Service replication and horizontal scaling alter this situation by introducing redundancy, enabling the redistribution of load to operational instances and maintaining availability even when part of the nodes are lost [14].

Asynchronous interaction breaks rigid dependencies between components. A request no longer requires an immediate response from a specific service, reducing the likelihood of cascading failures in which a single disruption triggers a chain reaction. Instead, the system shifts toward a mode of event accumulation and subsequent processing. Interaction becomes more flexible, but at the same time, the complexity of managing state and operation sequencing increases [8].

Load balancing redistributes incoming requests among available instances and smooths local overloads. Without it, replication does not achieve its full effect, as load may still concentrate on specific nodes. The combined use of scaling and load balancing forms a feedback loop in which the system adapts to changes in request intensity and to partial loss of resources.

The distribution of processing between edge and central layers alters the form of failure. When connectivity with the central layer is lost, part of the functionality continues to operate locally. The system does not completely stop but transitions into a constrained mode of operation.

Under such conditions, architecture functions not as a set of isolated solutions but as an interconnected system of mechanisms. Decomposition limits the scope of failure. Asynchrony reduces interdependencies between components. Scaling and load balancing sustain execution under resource loss. Segmentation distributes functions

across layers. The combined effect of these mechanisms determines system behavior under failure. In one configuration, a disruption leads to a complete halt of computation; in another, it results in load redistribution and continued execution with modified performance characteristics.

#### IV. DISCUSSION

Fault tolerance in distributed systems is formed as a result of the entire architecture rather than a single mechanism. It emerges from the coordinated interaction of layers. Attempts to address it through isolated tools such as replication or load balancing lead to fragmented solutions. In complex systems, such approaches quickly lose effectiveness.

At the infrastructure level, resilience is established through replication, distributed storage, and geographic distribution of resources. This level provides the basic ability to preserve data and maintain availability during failures of individual nodes or regions. Redundancy alone is insufficient for overall system resilience. Coordination at higher levels is required.

The platform layer is responsible for resource management, orchestration, and load redistribution. Here, resilience appears as adaptation to changing operating conditions. Container orchestration, scaling, and task scheduling maintain continuity under fluctuating workloads. The system adjusts to its current state.

At the application level, resilience takes on an architectural form. Microservice decomposition and asynchronous interactions change the nature of failures. A failure becomes a localized incident rather than a critical event. Its impact is limited by service boundaries.

As a result, resilience is formed at the intersection of all architectural layers. It emerges from their combined operation. Figure 1 shows how architectural decisions affect efficiency, convergence, latency, and load balancing. The relationship becomes measurable.




Fig.1 – Quantitative effects of microservices-based fault-tolerant optimization framework (Compiled by the author based on source: [14])

Optimization based on microservice architecture and fault-tolerance mechanisms leads to improved energy efficiency, faster convergence processes, reduced latency, and more effective load distribution. These changes are not side effects but direct consequences of the system's architectural organization. Importantly, the impact of architecture manifests simultaneously across multiple metrics, indicating its systemic nature. Consequently, architecture has evolved from a descriptive model into a performance management instrument, where structural decisions directly determine system behavior in measurable terms.

In this context, there is a need to move from empirically observed effects to their architectural interpretation and formalization. Figure 2 presents the author's model of a fault-tolerant distributed system architecture, which generalizes the identified relationships and translates them into a structured multi-layer conceptual framework.

The model describes a system where different architectural layers interact through the exchange of data such as performance metrics, failure events, and resource utilization. This data is transmitted upwards, allowing the adaptive layer to detect anomalies and

predict potential failures. At the same time, control actions flow downwards, influencing orchestration, service behavior, and infrastructure configuration.

The proposed model focuses on fault tolerance as a result of layer interaction rather than isolated solutions. It includes the infrastructure layer (with replication and distributed storage mechanisms), the platform layer (with container orchestration and scaling), the application layer (with microservice architecture and API gateways), the resilience layer (with failure management mechanisms), and the adaptive intelligence layer (with failure prediction and resource optimization methods).

The main advantage of the model is in explaining improvements in fault tolerance as a result of the coordinated work of all layers. Unlike traditional approaches, where fault tolerance is achieved through redundancy or recovery alone, this model shows that resilience can be effectively managed and optimized. The model proposes a shift from a reactive approach to an adaptive one, where the system not only maintains operability but also actively optimizes its behavior under uncertainty.

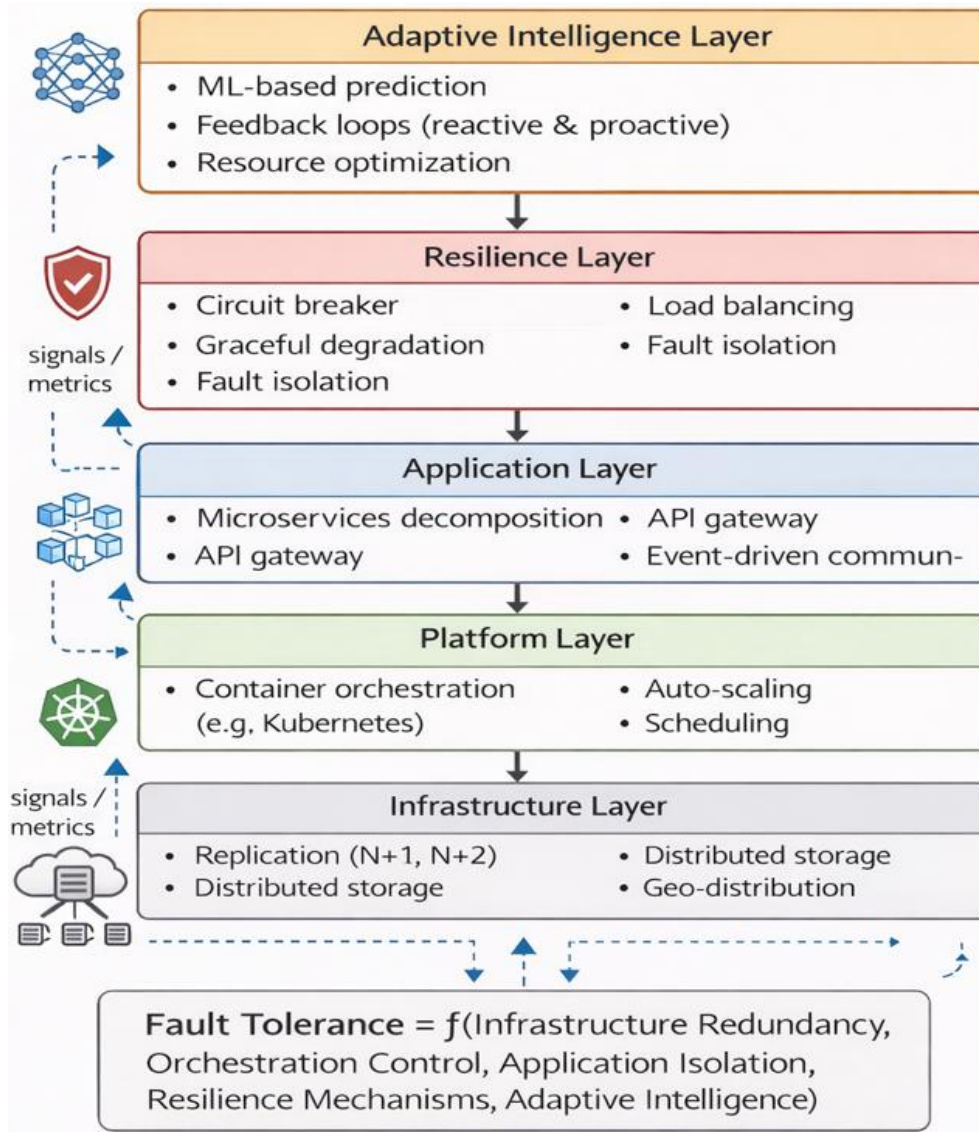


Fig.2 – Conceptual model of fault-tolerant distributed system architecture(Author's own development)

## V. CONCLUSION

The results obtained make it possible to consider the fault tolerance of distributed software systems not as a consequence of implementing individual mechanisms, but as a process formed within the architectural organization of the system. System behavior under failure conditions, recovery patterns, and resilience to overloads are determined less by the presence of specific tools and more by the configuration of interactions between infrastructure, platform, and application layers. The influence of architecture is manifested through the structure of connections within which computation, data transfer, and workload redistribution occur.

The consistency of architectural layers and management mechanisms becomes decisive. Even with effective solutions at each individual level, the overall effect remains unstable in the absence of coordination between replication mechanisms, orchestration, service decomposition, and failure management. When such coordination is achieved, a stable operational logic emerges, within which failures become localizable and recovery becomes a controllable process that does not lead to critical system disruptions.

The practical significance of the results lies in the need to reconsider approaches to the design of distributed systems. System resilience is determined by the architecture's ability to provide a continuous

management loop spanning all layers, from infrastructure to the adaptive intelligent layer. In this context, the priority shifts from implementing isolated fault-tolerance mechanisms to their initial integration into the system architecture, enabling measurable improvements in performance, latency, and load resilience.

Future research directions are associated with further formalization of the proposed model and the analysis of interaction mechanisms between architectural layers. Particular interest lies in developing methods for quantitatively assessing the degree of architectural coherence and its impact on system behavior over time. Another promising direction is the study of adaptive intelligent mechanisms in enabling the transition from reactive to proactive failure management, as well as evaluating their influence on the resilience of distributed systems under conditions of high uncertainty.

## REFERENCES

- [1] Angelis, A., & Kousiouris, G. (2025). An overview on the landscape of self-adaptive cloud design and operation patterns: Goals, strategies, tooling, evaluation, and dataset perspectives. *Future Internet*, 17(10), 434. <https://doi.org/10.3390/fi17100434>
- [2] Barroso, L. A., Hölzle, U., & Ranganathan, P. (2026). Trusted computing: Reliability, availability, security, privacy. In *The data center as a computer (Synthesis Lectures on Computer Architecture)*. Springer. [https://doi.org/10.1007/978-3-031-99489-0\\_10](https://doi.org/10.1007/978-3-031-99489-0_10)
- [3] Dai, F., Hossain, M. A., & Wang, Y. (2025). State of the art in parallel and distributed systems: Emerging trends and challenges. *Electronics*, 14(4), 677. <https://doi.org/10.3390/electronics14040677>
- [4] El Akhdar, A., Baidada, C., Kartit, A., Hanine, M., García, C. O., Lara, R. G., & Ashraf, I. (2024). Exploring the potential of microservices in Internet of Things: A systematic review of security and prospects. *Sensors*, 24(20), 6771. <https://doi.org/10.3390/s24206771>
- [5] Harshith, M., Ansari, Z. A., Fatima, S., et al. (2026). Federated microservices architecture with blockchain for privacy-preserving and scalable healthcare analytics. *Scientific Reports*, 16, 9023. <https://doi.org/10.1038/s41598-026-39837-1>
- [6] Ileana, M., Petrov, P., & Milev, V. (2025). AI-enabled secure and scalable distributed web architecture for medical informatics. *Applied Sciences*, 15(19), 10710. <https://doi.org/10.3390/app151910710>
- [7] Miao, T., Shaafi, A. I., & Song, E. (2025). Systematic mapping study of test generation for microservices: Approaches, challenges, and impact on system quality. *Electronics*, 14(7), 1397. <https://doi.org/10.3390/electronics14071397>
- [8] Mukhiyadin, A., Akmoldina, A., Tainova, K., et al. (2026). Design and implementation of an interface architecture for automated IoT device testing platforms. *Energy Informatics*, 9, 28. <https://doi.org/10.1186/s42162-026-00629-6>
- [9] Muñoz-Muñoz, D., Garcia-Carballeira, F., Camarmas-Alonso, D., et al. (2025). Malleability and fault tolerance in ad-hoc parallel file systems. *Cluster Computing*, 28, 860. <https://doi.org/10.1007/s10586-025-05575-8>
- [10] Murala, D. K., Prasada Rao, K. V., Vuyyuru, V. A., et al. (2025). A service-oriented microservice framework for differential privacy-based protection in industrial IoT smart applications. *Scientific Reports*, 15, 29230. <https://doi.org/10.1038/s41598-025-15077-7>
- [11] Pournazari, J., Ullah, A., Al-Dubai, A., et al. (2025). Computation offloading in the edge-to-cloud compute continuum: A survey of federated architectural solutions. *Cluster Computing*, 28, 839. <https://doi.org/10.1007/s10586-025-05577-6>
- [12] Sabuhi, M., Musilek, P., & Bezemer, C.-P. (2024). Micro-FL: A fault-tolerant scalable microservice-based platform for federated learning. *Future Internet*, 16(3), 70. <https://doi.org/10.3390/fi16030070>
- [13] Samala, J., Mekala, B. B., & J, S. (2025). Enhancing fault-tolerant application mapping in network-on-chip through transformer network based reinforcement learning approach. *Discover Computing*, 28, 235. <https://doi.org/10.1007/s10791-025-09704-0>
- [14] Sheng, S. (2023). Research on multi-dimensional reconstruction mechanism of cloud native full link in the metaverse scenario. *Scientific Reports*, 13, 22059. <https://doi.org/10.1038/s41598-023-48724-y>
- [15] Wang, Y., Wan, Q., Wu, Y., et al. (2025). Grouped Byzantine fault tolerant consensus algorithm based on aggregated signatures. *Cybersecurity*, 8, 60. <https://doi.org/10.1186/s42400-025-00362-9>