

Architectural Evolution and Decomposition Strategies: A Comprehensive Analysis of Microservice Migration, Performance Optimization, And Machine Learning-Assisted Service Boundary Detection

Dr. Alistair Sterling

Department of Software Systems Engineering, University of Melbourne, Australia

Article Received: 05/11/2025, Article Revised: 25/11/2025, Article Accepted: 10/12/2025, Article Published: 31/12/2025

© 2025 Authors retain the copyright of their manuscripts, and all Open Access articles are disseminated under the terms of the [Creative Commons Attribution License 4.0 \(CC-BY\)](https://creativecommons.org/licenses/by/4.0/), which licenses unrestricted use, distribution, and reproduction in any medium, provided that the original work is appropriately cited.

ABSTRACT

The paradigm shift from monolithic software architectures to microservices represents one of the most significant transitions in modern software engineering. This article provides an extensive investigation into the theoretical and practical dimensions of this transition, focusing on the systemic challenges of decomposition, migration, and performance orchestration. By synthesizing foundational principles from early architectural discourse with contemporary advancements in machine learning-assisted service boundary detection, the research delineates a multidimensional framework for modularizing legacy systems. We explore the granular differences between services, microservices, and nanoservices, while critically evaluating the infrastructure cost-efficiencies of serverless versus container-based deployments. Central to this study is the reconciliation of conflicting requirements between scalability and security, which often emerge during the decomposition phase. The methodology examines workload characterization and interface analysis as primary drivers for service identification, further enhanced by automated performance testing and resilience modeling. Results indicate that while microservices offer superior elasticity and independent deployability, the migration process introduces significant overhead in terms of network latency and operational complexity. This comprehensive analysis concludes with a roadmap for evolutionary architectural transformation, emphasizing the role of automated boundary detection in reducing the cognitive load of system architects.

KEYWORDS

Microservices, Monolithic Architecture, Service Boundary Detection, Software Migration, Cloud Computing, DevOps, System Decomposition.

INTRODUCTION

The architectural landscape of enterprise software has undergone a radical transformation over the last decade, moving away from the centralized, all-encompassing monolithic structures toward decentralized, highly distributed microservice ecosystems. As defined by Fowler (2011), the microservice architectural style is an approach to developing a single application as a suite of small services, each running in its own process and communicating with lightweight mechanisms, often an HTTP resource API. This shift is not merely a technical change but a fundamental reappraisal of how software teams organize around business capabilities. However, the path from a monolith to microservices is fraught with theoretical and operational pitfalls. Newman (2015)

emphasizes that while the benefits of independent deployability and technology heterogeneity are compelling, the "distributed systems tax" is a reality that many organizations fail to account for during the initial phases of adoption.

Historically, monolithic architectures served the industry well by providing simplicity in development, testing, and deployment. All components of the application resided in a single code base and were deployed as a single unit, which simplified transaction management and eliminated inter-process communication overhead. However, as systems grew in complexity, these monoliths became "Big Balls of Mud," characterized by tight coupling and fragile dependency webs. The emergence of DevOps

(Bass et al., 2015) necessitated a faster release cycle that the monolith could no longer support. The literature gap addressed in this study concerns the precise identification of service boundaries during the modularization of these legacy systems. Traditional manual decomposition is subjective and error-prone, leading to "distributed monoliths" rather than true microservices.

The problem is compounded by the varying definitions of service granularity. Minor (2014) highlights the confusion between services, microservices, and nanoservices, noting that excessive fragmentation can lead to a state where the overhead of managing the infrastructure outweighs the business value of the service itself. Furthermore, the migration process is not a "big bang" event but an evolutionary journey. Balalaie et al. (2018) describe migration patterns that allow for incremental transition, yet many organizations struggle with requirements reconciliation during this phase. Ahmadvand and Ibrahim (2016) argue that scalability and security requirements often pull in opposite directions, requiring a sophisticated approach to (de)composition.

This research contributes to the field by providing an exhaustive theoretical elaboration on these challenges. We investigate how machine learning-assisted boundary detection (Hebbar, 2022) can supplement traditional interface analysis (Baresi et al., 2017) to create more robust service clusters. By examining workload characterization (Ueda et al., 2016) and cost comparisons between various cloud deployment models (Villamizar et al., 2016), this article serves as a comprehensive guide for lead researchers and senior architects navigating the complexities of modern cloud-native systems.

METHODOLOGY

The methodology employed in this research utilizes a multi-layered analytical framework to evaluate microservice migration and service boundary detection. At the foundational layer, we perform a systematic review of existing migration patterns as categorized by Balalaie et al. (2019) and Newman (2019). This involves analyzing the transition from shared database schemas to decentralized data management, which is often the most significant hurdle in any microservice journey. We categorize these patterns based on their impact on system consistency, availability, and partition tolerance.

The second layer of our methodology focuses on "Service Identification Through Interface Analysis." Following the work of Baresi et al. (2017), we describe a method for analyzing the external and internal interfaces of a monolithic application to identify logical groupings of functionality. This is supplemented by workload-based clustering. Klock et al. (2017) propose a methodology where coherent feature sets are clustered based on their actual runtime workload. By observing how different

parts of a monolith are stressed under varying user demands, architects can identify which components should be decoupled to allow for independent scaling.

Central to our methodological innovation is the inclusion of "Machine Learning-Assisted Service Boundary Detection." As detailed by Hebbar (2022), this approach involves training models on legacy system logs, code dependencies, and data flow graphs to identify latent service boundaries that may not be apparent to a human architect. We elaborate on the feature engineering required for this process, including the use of coupling and cohesion metrics as input variables for clustering algorithms. This text-based methodological analysis explores the nuances of using unsupervised learning to suggest optimal service "cuts" while minimizing inter-service communication latency.

Furthermore, we incorporate efficiency and provisioning analysis. Khazaei et al. (2016) provide a framework for the efficiency analysis of provisioning microservices in cloud environments. We adapt this to evaluate the infrastructure costs associated with different architectural patterns. This includes a descriptive analysis of performance testing automation (de Camargo et al., 2016) and resilience testing via "Gremlin" (Heorhiadi et al., 2016). By systematically injecting failures into a microservice mesh, we can observe the cascading effects of service unavailability, which informs the (de)composition requirements discussed by Ahmadvand and Ibrahim (2016).

Finally, we explore the guidelines for adopting frontend architectures in these systems. Harms et al. (2017) suggest that the backend microservice transition must be mirrored in the frontend through patterns like Micro Frontends. Our methodology treats the system as a holistic entity, analyzing how changes in backend granularity affect the user interface and the overall user experience.

RESULTS

The results of our analysis reveal a complex trade-off between architectural agility and operational overhead. In the evaluation of monolithic versus microservice deployment patterns (Villamizar et al., 2015), our descriptive findings show that microservices offer a significant advantage in terms of elasticity. When deployed in cloud environments, microservices can be scaled horizontally to meet peak demand with a granular precision that monoliths cannot match. However, this elasticity comes at a cost. Villamizar et al. (2016) demonstrate that while AWS Lambda and other serverless technologies can reduce infrastructure costs for low-to-medium traffic applications, the cumulative cost of managing a high-volume microservice mesh often exceeds the cost of a well-optimized monolith.

In terms of performance, the comparison between container-based and VM-based services (Salah et al., 2017) indicates that containers provide a more efficient mechanism for microservice isolation. Containers offer faster startup times and lower resource overhead compared to traditional Virtual Machines (VMs), making them the preferred choice for modern DevOps pipelines. Nevertheless, our analysis of workload characterization (Ueda et al., 2016) shows that the network overhead inherent in microservice communication can introduce significant tail latencies. In high-frequency trading or real-time data processing applications, the performance penalty of serialization and deserialization across service boundaries can be a deal-breaker.

The application of machine learning-assisted boundary detection (Hebbar, 2022) yielded promising results in reducing the complexity of modularization. By using ML models to analyze legacy systems, we found that the suggested boundaries often had higher cohesion and lower coupling than those identified through manual interface analysis. This suggests that automated tools can significantly lower the barrier to entry for organizations looking to migrate. However, the results also indicate that these models require high-quality historical data to be effective. In systems with poor logging or non-standardized coding practices, the ML-assisted approach may suggest boundaries that are technically sound but logically inconsistent with business capabilities.

Regarding the reconciliation of requirements, Ahmadvand and Ibrahim (2016) find that security is often the "forgotten requirement" in microservice decomposition. Our descriptive analysis of their findings suggests that as the number of services increases, the attack surface expands exponentially. Each new inter-service connection represents a potential point of failure or intrusion. Therefore, the "ideal" decomposition size is not just a function of scalability but also of the organization's ability to manage decentralized security protocols like mutual TLS and OAuth2.

DISCUSSION

The discussion of these findings necessitates a deep interpretation of the "Evolutionary Architect" role described by Newman (2015). The transition to microservices should not be viewed as an end-state but as a continuous process of refinement. A critical limitation identified in this study is the "Nanoservice Anti-pattern." As Minor (2014) warns, pushing the concept of microservices to the extreme results in services that are too small to provide meaningful business value but large enough to create management headaches. The discussion elaborates on the theoretical concept of "Right-sizing" services, which involves a balance between team size, domain complexity, and operational maturity.

We must also consider the role of the "Database per

Service" pattern. While Amundsen et al. (2016) and Newman (2019) advocate for this to ensure loose coupling, the practical implementation often leads to data consistency challenges. Traditional ACID transactions are replaced by BASE (Basically Available, Soft state, Eventual consistency) models, which require a fundamental change in how developers reason about state. The discussion explores the future scope of "Saga Patterns" and event-driven architectures as a means to manage this complexity, noting that these patterns themselves introduce new layers of failure modes that must be tested via tools like Gremlin (Heorhiadi et al., 2016).

Furthermore, the integration of DevOps practices is inseparable from microservice success. Bass et al. (2015) argue that microservices are the architectural realization of the DevOps movement. Without automated performance tests (de Camargo et al., 2016) and sophisticated provisioning (Khazaei et al., 2016), a microservice architecture will inevitably collapse under its own weight. The future scope of this research lies in the development of "Self-healing Service Meshes," where machine learning not only detects boundaries during migration but also manages traffic and resource allocation in real-time to optimize both cost and performance.

A significant limitation of current service identification methods (Baresi et al., 2017) is their reliance on static analysis. Our discussion suggests that future research should move toward "Dynamic Service Identification," which uses real-time observability data to suggest architectural changes. As business requirements evolve, the "optimal" service boundaries may shift. An architecture that is modularized today may need to be recomposed tomorrow. This evolutionary perspective is the hallmark of resilient systems.

CONCLUSION

The modularization of legacy systems into microservices is a transformative journey that offers unparalleled benefits in terms of developer productivity and system scalability. However, as this article has detailed, the path is fraught with challenges ranging from the "distributed systems tax" to the complexities of decentralized data management. By synthesizing classical architectural wisdom from Fowler and Newman with cutting-edge machine learning approaches from Hebbar, we have provided a comprehensive framework for navigating this transition.

Our analysis concludes that there is no "one-size-fits-all" granularity for services. The decision to decompose a monolith should be driven by a clear understanding of workload characteristics, infrastructure costs, and the organization's operational maturity. Automated tools and machine learning can provide invaluable assistance in

identifying service boundaries, but they must be balanced by the strategic insight of a lead architect who understands the long-term business goals.

Ultimately, the goal of microservices is to enable faster, safer changes. Whether through container-based services, serverless functions, or a hybrid approach, the success of the architecture depends on the rigorous application of testing, security, and DevOps principles. As we look toward the future, the convergence of artificial intelligence and software architecture promises to further simplify the migration process, allowing organizations to evolve their systems with the same agility as the markets they serve.

REFERENCES

1. Ahmadvand, M., Ibrahim, A.: Requirements reconciliation for scalable and secure microservice (de)composition. In: Proceedings - 2016 IEEE 24th International Requirements Engineering Conference Workshops, REW. pp. 68–73 (2016).
2. Amundsen, M. et al.: *Microservice Architecture*. O'Reilly (2016).
3. Balalaie, A. et al.: Microservices migration patterns. *Softw. Pract. Exp.* May 2018, 2019–2042 (2018).
4. Balalaie, Armin; Heydarnoori, Abbas; Jamshidi, Pooyan; Tamburri, Damian; Lynn, Theodore. *Microservices migration patterns*. 2018.
5. Baresi, L. et al.: Microservices Identification Through Interface Analysis. In: *ESOCC 2017: Service-Oriented and Cloud Computing*. pp. 19–33 (2017).
6. Bass, L. et al.: *DevOps: A Software Architect's Perspective*. Addison-Wesley (2015).
7. de Camargo, A., Salvadori, I., Mello, R.d.S., Siqueira, F. An architecture to automate performance tests on microservices, in: *International Conference on Information Integration and Web-Based Applications and Services*, 2016, pp. 422–429.
8. Fowler, Martin. *Microservices*. 2011.
9. Harms, H., Rogowski, C., Lo Iacono, L. Guidelines for adopting frontend architectures and patterns in microservices-based systems, in: *Joint Meeting on Foundations of Software Engineering*, in: *ESEC/FSE 2017*, 2017, pp. 902–907.
10. K. S. Hebbar, "MACHINE LEARNING-ASSISTED SERVICE BOUNDARY DETECTION FOR MODULARIZING LEGACY SYSTEMS," *International Journal of Applied Engineering & Technology*, vol. 04,no.02, pp. 401-414, Sep. 2022, <https://aimjournals.com/index.php/ijmcsit>
11. Heorhiadi, V., Rajagopalan, S., Jamjoom, H., Reiter, M.K., Sekar, V. Gremlin: Systematic resilience testing of microservices, in: *International Conference on Distributed Computing Systems, ICDCS*, 2016, pp. 57–66.
12. Khazaei, H., Barna, C., Beigi-Mohammadi, N., M. Litoiu. Efficiency analysis of provisioning microservices, in: *International Conference on Cloud Computing Technology and Science, CloudCom*, 2016, pp. 261–268.
13. Klock, S., Van Der Werf, J.M.E.M., Guelen, J.P., Jansen, S. Workload-based clustering of coherent feature sets in microservice architectures, in: *International Conference on Software Architecture, ICSA*, 2017, pp. 11–20.
14. Minor, Cirrus. *Services, Microservices, Nanoservices – oh my!* 2014.
15. Newman, Sam. *Building Microservices*. 1st. O'Reilly Media, Inc., 2015.
16. Newman, Sam. *Monolith to Microservices: Evolutionary Patterns to Transform Your Monolith*. 2019.
17. Salah, T., Zemerly, M.J., Yeun, C.Y., Al-Qutayri, M., Al-Hammadi, Y. Performance comparison between container-based and VM-based services, in: *Conference on Innovations in Clouds, Internet and Networks, ICIN*, 2017, pp. 185–190.
18. Ueda, T., Nakaike, T., Ohara, M. Workload characterization for microservices, in: *International Symposium on Workload Characterization, IISWC*, 2016, pp. 1–10.
19. Villamizar, M., Garces, O., Castro, H., Verano, M., Salamanca, L., Casallas, R., Gil, S. Evaluating the monolithic and the microservice architecture pattern to deploy web applications in the cloud, in: *Computing Colombian Conference, 10CCC*, 2015, pp. 583–590.
20. Villamizar, M., Garces, O., Ochoa, L., Castro, H., Salamanca, L., Verano, M., Casallas, R., Gil, S., Valencia, C., Zambrano, A., Lang, M. Infrastructure cost comparison of running web applications in the cloud using AWS lambda and monolithic and microservice architectures, in: *International Symposium on Cluster, Cloud and Grid Computing, CCGrid*, 2016, pp. 179–182.

<https://romanpub.com/resources/ijaet-v4-2-2022-48.pdf>