

Review Article

Application of clean architecture principles in the design of scalable server-side applications

Artem Agaev*

Department of the IT Block “People and Culture,” Sberbank, Moscow, Russia

ABSTRACT

The article examines the application of clean architecture principles to the design of scalable server-side applications, reducing architectural complexity and technical debt. The purpose of the study, based on a conceptual–empirical review, is to relate the structural decisions of clean architecture to measurable quality attributes of server systems, scalability, maintainability, testability, and resilience to operational stress. The relevance of the research is determined by the rapid growth in the complexity of cloud and microservice landscapes, in which poorly structured architecture becomes a latent constraint on organizational and technological evolution. The scientific novelty of the work lies in integrating fragmented empirical data on architectural smells, design patterns, and transformations for testability into a unified analytical framework that presents clean architecture as a scaffold unifying domain-driven design, the ports-and-adapters approach, and layered code organization. It is shown that inward-oriented dependencies, the separation of domain and application layers, and the localization of technological variability at the periphery enable simultaneous increases in maintainability and scalability with respect to load, functionality, and team size, without exponential growth in technical debt. The article is of practical value for software architects, development team leads, and engineering practitioners engaged in the design and evolution of server-side backends.

Keywords: Maintainability, microservice, scalability, testability

Submitted: 21-02-2026, **Accepted:** 15-03-2026, **Published:** 30-03-2026

INTRODUCTION

In the era of rapid digitalization, server-side applications serve as the supporting infrastructure for business processes, analytics, and user interaction. The transition to cloud environments, the widespread adoption of architectures composed of numerous small services, and the increasing degree of integration with external platforms make server systems simultaneously more flexible and substantially more complex. Empirical studies indicate that architectures based on fine-grained services are indeed used to increase horizontal scalability and fault tolerance, but are accompanied by increased complexity in interactions, orchestration, and monitoring, which directly affects operational risks and maintenance costs.^[1] The question is less about the specific infrastructure technology choices (frameworks, transport, and storage) and more about how carefully the application’s internal architecture is structured, as well as how well the application will be able to cope with increased load, data volumes, and new integrations.

If an architecture has no proper structure, architectural decay and technical debt will amass, and the internal resistance of an architecture will gradually deteriorate. Studies of architectural smells in open-source projects demonstrate a statistically significant relationship between their concentration and the deterioration of quality attributes such as modularity and testability: The more architectural defects are present, the more difficult it becomes to localize changes and ensure automated testing.^[2]

In practice, this manifests as a significant share of effort spent not on implementing new functionality but on addressing the consequences of earlier architectural decisions. A survey of more than 200 engineers, summarized in the report State of Technical Debt 2021, showed that over 60% of developers directly associate technical debt with defects, outages, and slowed development; on average, they spend about one working day per week solely on debt-related activities, and two thirds of respondents are convinced that, given a systematic

Address for correspondence: Artem Agaev, Department of the IT Block “People and Culture,” Sberbank, Moscow, Russia
E-mail: agaevartem1997@gmail.com

process for managing debt, the team could deliver functionality up to twice as fast.^[3] Taken together, this means that poorly structured architecture with implicit dependencies, tight module coupling, and a random distribution of responsibilities becomes a hidden constraint on scalability, manifesting not only in declining performance but also in the slowing of organizational change processes.

Against this background, clean architecture appears not as yet another fashionable design pattern but as a way of purposefully managing the complexity of server-side applications. Within this approach, the system is organized around stable domain rules and use cases rather than specific technologies; dependencies are oriented from outer layers, transport, storage, and integration mechanisms, toward inner layers, where business logic is concentrated in infrastructure-independent modules. The division into concentric layers of domain entities, applications, interface adapters, and infrastructure enables localized changes and prevents the infiltration of technical details into the system's core, thereby reducing the rate at which architectural debt accumulates. Empirical work confirms that such structuring produces a measurable effect: when refactoring a server application using clean architecture principles, researchers recorded improvements in maintainability metrics in the range of 21–61% across different indicators, accompanied by a reduction in complexity and a decrease in the effort required to modify code; similar results have been obtained for application systems with multilayered separation of responsibilities.^[4] In what follows, an analysis of the principles and practices of clean architecture in the context of server-side applications shows how this code organization becomes an indispensable condition for scalability, both with respect to load and to the rate of functional evolution.

MATERIALS AND METHODS

The study draws upon a corpus of sources that reflect both empirical aspects of the operation of server applications and the architectural–theoretical foundations of clean architecture principles. As an empirical basis, the research uses studies of the scalability and operational properties of microservice and server systems,^[1] work on architectural smells and their impact on modularity and testability,^[2] as well as data from industry surveys on the nature and consequences of technical debt at the level of development practices.^[3] The theoretical framework for the architectural part is formed through systematic reflection on layered and pattern-based solutions at the architectural level, including reviews of the layered pattern and analyses of the impact of architectural patterns on maintainability.^[5] As a specialized basis for examining clean architecture specifically and its influence on maintainability and the structure of server-side backends, the study draws on refactoring results from existing systems that have been introduced to clean architecture principles.^[4]

The link between architectural decisions and testability and resilience to operational stresses relies on systematic reviews of transformations aimed at increasing testability and its measurement at the design level.^[6] To substantiate the context of load scalability and the trend of increasing complexity of server applications in cloud environments, additional data are drawn on the dynamics of the global cloud computing market.^[7] This selection of sources enables connecting the levels of market, development organization, architecture, and code into a single analytical perspective.

Methodologically, the study is implemented as a conceptual–empirical review that combines the systematization of existing data with their projection onto the structural solutions of clean architecture. At the first stage, substantive mapping of the body of work is carried out along three axes: (a) quality attributes and operational effects (scalability, maintainability, testability, and resilience to defects);^[1] (b) types of architectural decisions and patterns (microservices, layered architectures, ports-and-adapters approach, and clean architecture);^[5] (c) practices of technical debt management and organizational consequences of architectural decisions.^[3] At the second stage, a comparative analysis is conducted: Empirical metrics (improvement of maintainability under refactoring, the relationship between the concentration of architectural smells and quality degradation, and the impact of architectural transformations on testability) are interpreted through the lens of clean architecture principles, namely, inward-oriented dependencies, isolation of domain logic, and concentration of technological variability at the periphery.^[4] At the third stage, a synthetic linkage between architectural principles and dimensions of scalability is constructed: By load, by functionality, and by team size, where data on the growth of cloud infrastructure and the increasing complexity of the integration landscape are used as a macro-context for interpreting the necessity of a strict structural organization of server applications.^[7]

RESULTS AND DISCUSSION

The key idea of clean architecture in the context of server-side applications is to bind system behavior not to a specific technology stack but to the domain's invariants. To that end, the principle of independence of the inner core from transport mechanisms is introduced: Code that implements business rules must not know about protocols such as Representational State Transfer or Google Remote Procedure Call, nor about delivery mechanisms such as message queues. Frameworks, web servers, and serialization libraries are treated as interchangeable details attached via abstract interfaces. The exact requirement extends to storage and integration subsystems: Domain entities and use cases do not depend on whether persistent storage is implemented using a relational database, a document store, or a distributed cache, and access to external services is encapsulated in dedicated adapters.

Empirical studies of architectural patterns confirm that the selection and strict adherence to an appropriate architecture directly influence maintainability: It has been shown that the use of an unsuitable pattern can significantly worsen maintainability indicators at the architectural level, even if the rest of the code conforms to general modularity principles.^[8] For systems oriented toward high load, this means that the ability to replace transport, storage, or integration mechanisms without rebuilding the domain becomes one of the conditions for scalability: Changes are concentrated at the periphery, whereas the core evolves much more slowly, which reduces the risk of architectural debt when adapting to new requirements and to growth in the volume of interactions.^[9]

The rule of inward-oriented dependencies gives the formal expression of this isolation: Each outer layer may depend on a more inner one, but not vice versa. In terms of layered architectures, this means that the infrastructure layer is aware of the application and domain layers, the interface layer is aware of the application layer, and the domain core remains isolated from technological details. A systematic review of the layered architectural pattern shows that the expected properties, reusability, portability, and maintainability, are achieved only when a strict layer hierarchy and acyclic dependencies are preserved; extensive use of layer bypasses and callbacks leads to cycles and, in effect, to the transformation of the system into a monolith that is difficult to decompose and evolve.^[5]

In the context of clean architecture, this rule is carried to its logical extreme: Entry points (controllers and message handlers) and adapters to databases and external services implement interfaces defined within the application layer. They can be replaced independently of the domain. Practical measurements indicate that such reorientation of dependencies yields a tangible effect: In an experiment involving refactoring of a server application using clean architecture principles, all maintainability metrics improved in the range of 21–61%, accompanied by reduced complexity and lower effort for introducing changes.^[4]

It is precisely inward-oriented dependencies and the isolation of domain logic that create the preconditions for systematic testability, which serves as the foundation for the safe growth of the server system. When use cases and entities do not depend on infrastructure, they can be verified by unit tests without deploying real databases, queues, or network services; load and integration testing is concentrated on adapters and infrastructure without affecting the core. A systematic review of work on software testability and resilience, encompassing more than 30,000 initially selected publications and 27 studies analyzed in detail, shows that improved testability is closely linked to increased resilience to incorrect inputs and stressful environmental conditions: Key factors include observability, controllability, and diagnosability of components, which are primarily determined by architectural decisions.^[6]

An additional review of refactoring practices aimed at improving testability, conducted on a corpus of more than 5000 studies, reports that a substantial proportion of successful transformations is associated with changes in architectural structure, splitting modules, eliminating cyclic dependencies, and introducing explicit interfaces between subsystems.^[10] Against this background, the traditional three-layer scheme of the Controller–Service–Repository type often proves to be only a particular, technologically oriented instance of a layered architecture: Controllers are tightly bound to the web framework, services mix domain use cases with infrastructural details, and repositories embody knowledge of specific storage schemas, which leads to layer bypasses, cycles, and gradual erosion of the architecture. Clean architecture redistributes the same structural elements differently: Controllers become pure adapters that invoke use cases; use cases operate only on domain objects and abstract ports; repositories and external clients merely implement these ports. As a result, the difference between a traditional layered architecture and a clean architecture lies not in the number of layers but in the direction and rigor of dependencies, and it is this difference that determines a server application’s ability to scale both in terms of load and evolution speed. The core principles of clean architecture are shown in Figure 1.

The consistent application of clean architecture principles in a server application begins with an explicit delineation of structural layers. The inner core forms the domain layer, where the domain’s entities and invariants are concentrated. Above it lies the application layer of use cases, which describes the allowed operations on domain objects and their order. Components of the interface layer receive requests, transform external data into application-layer structures, and map results back into output formats. These take as input a command from a user or another service and convert external data into a representation that the application can understand. The infrastructure is the outermost layer, encompassing the code that implements specialized data storage technology, a message broker, an event logging framework, and other implementation-specific details. This layering arrangement meets a key requirement: inner levels do not depend on outer ones, and all technology stack variability is concentrated at the periphery.

This logical structure naturally transitions into the modular organization of code. The domain layer is extracted into a separate module containing domain models, value objects, and domain services that use no external libraries beyond the basic ones. The application layer is represented by its own module of use cases that depends only on domain code and a set of abstract ports for accessing storage, queues, and external services. The interface layer groups modules responsible for receiving and sending requests, transforming external representations into application-layer structures, and

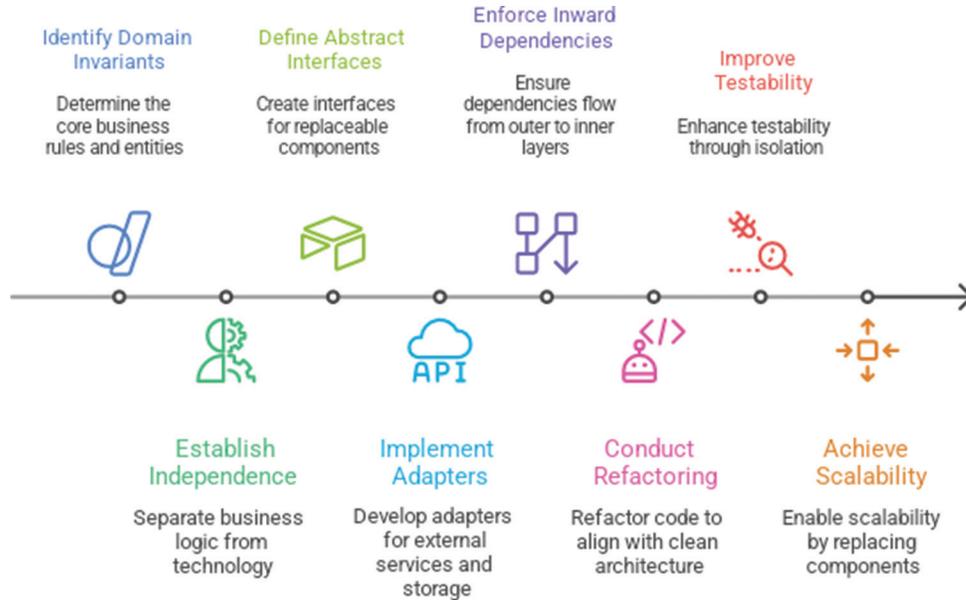


Figure 1: Basic principles of clean architecture

mapping results back. The infrastructure layer implements the application-level ports: Concrete repositories built on the chosen storage systems, adapters to messaging systems, and clients for external services. The request-processing flow in such an architecture appears as a transparent chain: an entry component in the interface layer receives the request, performs minimal validation, and converts it into a use-case structure, which works with domain objects and abstract ports; the result is then returned, transformed into an output format, and sent to the consumer. Ports and adapters connect the application logic and infrastructure. Concrete implementations can be changed without affecting use cases or the domain model.

These patterns relate naturally with domain-driven design, command, query responsibility segregation, and the ports-and-adapters architecture. In the DDD sense, the domain layer is the domain model and the bounded contexts. In the application layer, use cases such as user stories and system stories exist. Interface and infrastructure form the ports and adapters that represent interfaces separately from their implementation. Separation can happen at the application layer, where the use cases and models are different for read queries and write commands, without affecting overall composition. These choices allow for an architecture that can be decoupled from the underlying technologies, so that the load can be distributed, the data store or interaction protocols can change, affecting only the adapters and infrastructure locally. This makes such a structure a convenient foundation for scalable server applications that can evolve without continuous, radical code rework. The structure of a server application based on clean architecture principles is shown in Figure 2.

Practical scalability of a server-side application primarily manifests in how the system handles increasing load. This issue of scalability fits into a broader macro-trend: as shown in Figure 3, the global cloud computing market had already reached an estimated value of approximately USD 574 billion in 2023 and, according to forecasts, will grow to around USD 1.9 trillion by 2030, which corresponds to a compound annual growth rate of about 18.7%.^[7]

With a multilayered organization based on clean architecture, the bulk of changes is concentrated at the infrastructure and adapter levels. In contrast, the core containing domain rules and use cases remains unchanged. This enables the introduction of intermediate fast-data stores, queues for smoothing peaks, deferred processing mechanisms, and sharding without affecting domain entities and application algorithms. Bottlenecks are identified as concrete external components: Slow storage, overloaded network interfaces, or constrained compute resources. Since these components are represented in the system as implementations of abstract ports, their replacement or configuration change occurs locally, through modifications to the adapters; the domain core perceives these changes as interactions with the same concepts through different carriers. Thus, the architecture enables scaling system throughput and load redistribution without requiring radical restructuring of the logic.

Equally important is scalability in terms of functionality, that is, the capacity of the application to expand its behavior without breaking existing capabilities. In clean architecture, new use cases are realized as separate application components built upon already defined domain entities and values. When necessary, new domain objects are introduced, integrated

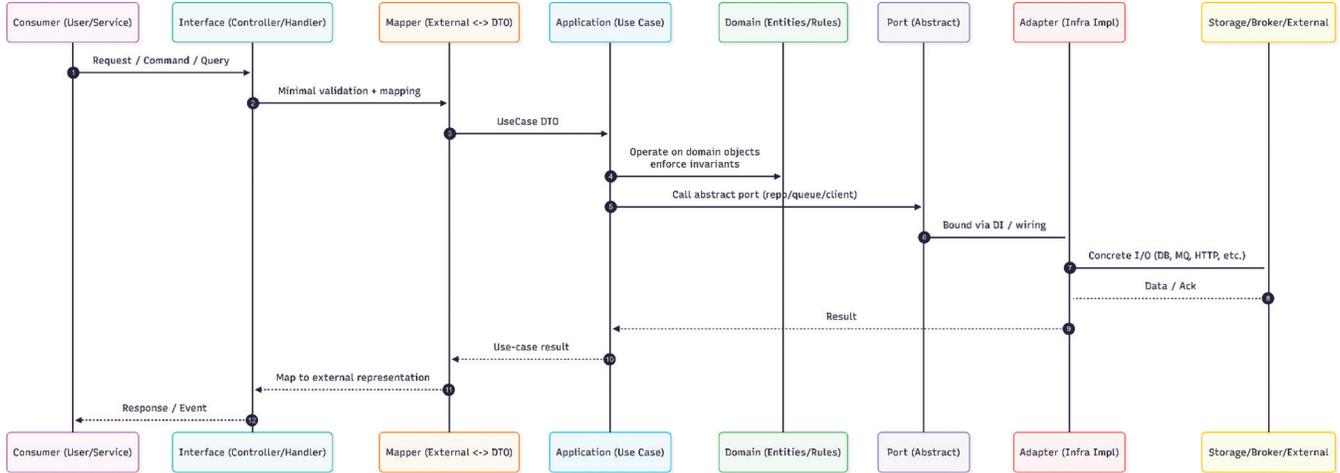


Figure 2: Building a scalable server application

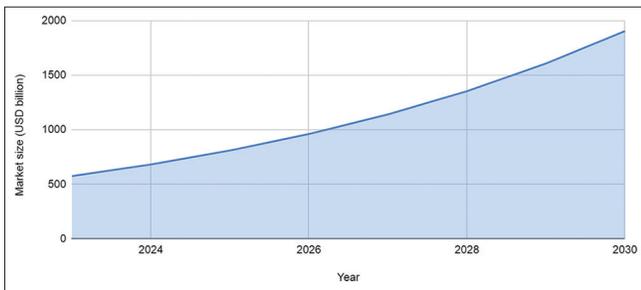


Figure 3: Projected growth of the global cloud computing market^[7]

into the existing structure, and do not require changes to the infrastructure if the data representation remains unchanged. Because domain rules and use cases that interact with them are separated from storage and transport details, the developer can focus on formalizing new operations and invariants without being distracted by the organization of network communication or specifics of a particular storage system. Such localization of changes reduces the risk of unforeseen side effects, facilitates regression testing, and makes the addition of functions a repeatable and predictable process.

Finally, structured architecture directly affects the scalability of the development team. Clearly delineated boundaries between domain, application, interface, and infrastructure layers translate into boundaries of responsibility between specialist groups. One group can be engaged in evolving the domain model and use cases; another can work on the efficiency of storage and communication mechanisms; a third can be responsible for network interfaces and user-facing protocols. This reduces overlaps and conflicts in changes, as interaction between layers proceeds through stable abstractions. At the organizational level, such a structure simplifies the transition from a single modular application to a distributed solution: independent areas are first identified within a single execution context, and the most heavily loaded or rapidly changing

subsystems are then extracted into separate processes. Since their boundaries have already been defined in terms of ports and adapters, decomposition into autonomous services is a technical step rather than a complete system reconstruction. As a result, architectural decisions made at early stages become an instrument for controlled growth in terms of load and functionality, as well as the number of participants in development.

CONCLUSION

This study synthesizes conceptual and empirical evidence to substantiate clean architecture as a rigorously grounded approach for structuring scalable server-side applications under the pressures of cloud adoption and microservice proliferation. By consolidating findings on architectural smells, architectural patterns, refactoring outcomes, and transformations for testability into a unified analytical lens, the article explicates how inward-oriented dependencies, the strict separation of domain entities and application use cases, and the confinement of technological volatility to peripheral adapters jointly correlate with improved scalability, maintainability, testability, and operational robustness. The reviewed empirical results indicate measurable maintainability gains following clean-architecture-oriented refactoring and reinforce the broader claim that disciplined layer hierarchies and acyclic dependencies reduce the propensity for architectural decay and the accumulation of technical debt, thereby sustaining evolutionary capacity as load, functionality, and team size expand. In this framing, clean architecture operates as a methodological scaffold aligning domain-driven design with ports-and-adapters and layered organization, enabling localized infrastructural substitutions and controlled decomposition pathways that support long-term backend evolution amid increasing integration complexity and market-driven growth in cloud computing.

REFERENCES

1. Wang Y, Kadiyala H, Rubin J. Promises and challenges of microservices: An exploratory study. *Empir Softw Eng* 2021;26:63.
2. Jolak R, Karlsson S, Dobsław F. An empirical investigation of the impact of architectural smells on software maintainability. *J Syst Softw* 2025;225:112382.
3. Omeyer A, Carlo N. Technical Debt Isn't Technical: What Companies Can Do to Reduce Technical Debt," *InfoQ*. Available from: <https://www.infoq.com/articles/reduce-technical-debt> [Last accessed on 2021 Sep 08].
4. Nugroho YN, Kusumo DS, Alibasa MJ. Clean Architecture Implementation Impacts on Maintainability Aspect for Backend System Code Base. In: *Proceedings of International Conference on Information and Communication Technology*; 2022.
5. Belle AB, Boussaidi GE, Lethbridge TC, Kpodjedo S, Mili H, Paz A. Systematically Reviewing the Layered Architectural Pattern Principles and their use to Reconstruct Software Architectures. [arXiv Preprint]; 2021.
6. Zakeri-Nasrabadi M, Parsa S, Jafari S. Measuring and improving software testability at the design level. *Inform Softw Technol* 2024;174:107511.
7. Next MSC. Cloud Computing Market Size and Share. Next MSC; 2025. Available from: <https://www.nextmsc.com/report/cloud-computing-market> [Last accessed on 2025 Dec 11].
8. Rahmati Z, Tanhaei M. Ensuring software maintainability at software architecture level using architectural patterns. *AUT J Math Comput* 2021;2:81-102.
9. Li S, Zhang H, Jia Z, Zhong C, Zhang C, Shan Z, *et al*. Understanding and addressing quality attributes of microservices architecture: A systematic literature review. *Inform Softw Technol* 2020;131:106449.
10. Bagheri-Galle F, Parsa S, Zakeri M. A systematic literature review on transformation for testability techniques in software systems. *Inform Softw Technol* 2025;185:107788.



This work is licensed under a Creative Commons Attribution Non-Commercial 4.0 International License.