

# Revisiting Domain-Driven Design: Shared Invariants and the Domain Zero Pattern

**Alexey A. Nekludoff**

AstraVerge Research

E-mail: [an@astraverge.org](mailto:an@astraverge.org)  
ORCID: 0009-0002-7724-5762

7 March 2026

## Abstract

Domain-Driven Design (DDD) has become one of the most influential approaches to modeling complex software systems. It proposes that software architectures should be organized around domain models and bounded contexts that represent conceptual areas of a business.

This paper revisits this assumption from the perspective of real organizational operations. We show that business activity is primarily structured around end-to-end processes rather than isolated domains. These processes rely on a small set of organizational entities—such as counterparties, products, contracts, documents, employees, and financial transactions—that naturally participate in multiple processes across the enterprise.

Such entities act as *shared invariants* of business reality. Because they are simultaneously used by multiple processes, they cannot be cleanly partitioned across independent domain boundaries without introducing duplication of data, model divergence, synchronization mechanisms, and loss of a unified organizational view.

The paper argues that strict domain decomposition conflicts with the structural nature of enterprise systems. As a practical architectural response, we introduce the concept of *Domain Zero*: a foundational domain responsible for maintaining the shared informational core of the organization. Domain Zero preserves the integrity of shared invariants while allowing operational domains to remain independent.

Conceptually, Domain Zero corresponds to what has historically been known as Master Data Management (MDM), but is treated here as an explicit architectural domain rather than an external integration layer.

The central conclusion of the paper is that enterprise systems are organized around shared invariants rather than domains, and that effective enterprise architectures must explicitly preserve this shared informational core.

**Keywords:** DDD; domain model; software modeling; Domain Zero

# Contents

<b>1</b>	<b>Introduction</b>	<b>2</b>
<b>2</b>	<b>Domain-Driven Design</b>	<b>2</b>
<b>3</b>	<b>Typical Way of Business Software Development</b>	<b>3</b>
<b>4</b>	<b>It Was a Beautiful Dream</b>	<b>4</b>
<b>5</b>	<b>The Structural Reality of Business Processes</b>	<b>4</b>
<b>6</b>	<b>The Problem of Shared Invariants</b>	<b>5</b>
<b>7</b>	<b>Typical Shared Invariants</b>	<b>6</b>
<b>8</b>	<b>Shared Invariants and Process Graphs</b>	<b>6</b>
<b>9</b>	<b>Invisible Consequences of the Absence of Shared Invariants</b>	<b>9</b>
<b>10</b>	<b>Domains as an Artificial Decomposition</b>	<b>10</b>
<b>11</b>	<b>The Small Core of Business Reality</b>	<b>10</b>
<b>12</b>	<b>Regulatory Constraints and Shared Invariants</b>	<b>11</b>
<b>13</b>	<b>Domain Zero Architectural Pattern</b>	<b>12</b>
<b>14</b>	<b>Domain Zero as an Evolutionary Correction</b>	<b>14</b>
<b>15</b>	<b>Conclusion</b>	<b>14</b>
<b>A</b>	<b>Domain Zero — Foundational Shared-Invariant Domain</b>	<b>15</b>

# 1 Introduction

Modeling complex software systems has long been a central challenge in software engineering. As systems increase in scale and involve large numbers of interacting components, the primary difficulty often lies not in implementing individual algorithms but in organizing the conceptual structure of the system itself. Software must represent the entities, rules, and operational processes of real organizations, and the resulting conceptual complexity frequently dominates purely technical concerns.

Over the history of software development, multiple approaches have attempted to address this challenge. Early methodologies emphasized structured analysis and data modeling, focusing on entities, data flows, and functional decomposition. Later approaches shifted toward object-oriented modeling, where systems were organized around objects encapsulating both state and behavior. Each of these paradigms sought to provide conceptual tools for representing complex problem domains in a manageable and consistent way.

Within this broader context, Domain-Driven Design (DDD), introduced by Evans in 2003 [1], emerged as a widely influential approach to modeling complex business software. The central premise of DDD is that software architecture should be organized around domain models that capture the conceptual structure of a business domain.

DDD introduced several principles intended to improve the alignment between software and business knowledge. These include the notion of a *domain model* as the core representation of business concepts, the use of a *ubiquitous language* shared by developers and domain experts, and the division of systems into *bounded contexts* in order to maintain conceptual consistency within different parts of a system.

These ideas have had a substantial influence on contemporary software development practices and are widely applied in enterprise systems, distributed architectures, and microservice-based applications. However, the question remains whether domain-oriented decomposition accurately reflects the structural nature of real organizational systems.

In large enterprises, business activity is primarily organized around end-to-end operational processes that span multiple departments and functional areas. These processes rely on a relatively small set of shared organizational entities—such as customers, products, contracts, documents, and financial transactions—that participate simultaneously in multiple parts of the system.

This observation raises an important architectural question: can such shared entities be cleanly partitioned across independent domain boundaries without introducing duplication, fragmentation of data, and loss of organizational consistency?

This paper revisits the conceptual foundations of Domain-Driven Design from this perspective. We analyze the role of shared organizational entities in enterprise systems and argue that many of them function as *shared invariants* that cannot be naturally confined to isolated domains. Based on this analysis, the paper introduces the concept of *Domain Zero*, a foundational domain responsible for maintaining the shared informational core of the organization.

## 2 Domain-Driven Design

Domain-Driven Design (DDD) was introduced by Eric Evans in the early 2000s as a set of principles and practices intended to improve the development of software systems whose complexity is primarily driven by the structure of a problem domain [1]. The approach emerged from practical experience in large enterprise projects where software systems were required to encode complex business rules and organizational processes.

A central observation underlying DDD is that many software projects encounter difficulties not because of implementation challenges, but because the conceptual understanding of the problem domain is poorly reflected in the structure of the software. In such cases the terminology used by developers, domain experts, and stakeholders often diverges, leading to inconsistencies between requirements, models, and code.

To address this problem, DDD proposes the use of a *domain model* as the primary conceptual representation of the problem domain. The domain model captures the key concepts, relationships, and rules that characterize the domain and serves as a shared reference for both developers and domain experts.

Closely related to the domain model is the concept of a *ubiquitous language*. This language is developed collaboratively by domain experts and developers and is used consistently in discussions, documentation, and source code. The intention is to reduce ambiguity and ensure that the terminology used in software reflects the conceptual structure of the domain.

Another important concept introduced in DDD is that of the *bounded context*. Large systems may involve multiple models of the same or related domains, each serving different purposes or reflecting different perspectives. A bounded context defines the boundary within which a particular domain model is valid and internally consistent. Interactions between contexts are explicitly defined through mappings or integration mechanisms.

Within bounded contexts, DDD introduces several modeling elements that structure the implementation of the domain model. These include *entities*, which represent objects with a persistent identity; *value objects*, which are defined by their attributes rather than identity; *aggregates*, which group related entities and enforce consistency boundaries; and *repositories*, which provide access to persistent domain objects.

Over time, the concepts of Domain-Driven Design have been widely adopted in the development of enterprise applications and have influenced approaches to modularization and service decomposition, particularly in systems that employ distributed or service-oriented architectures.

### 3 Typical Way of Business Software Development

Software systems in business environments are usually created in response to the operational needs of a real organization. Such organizations conduct activities such as sales, logistics, accounting, customer management, or production, and these activities form the basis of the processes that require automation.

Domain-Driven Design proposes to address this problem by organizing software development around domain models. Business activity is conceptually decomposed into a set of domains or subdomains, each representing a coherent area of responsibility within the overall operation.

Within this framework, developers identify domains and define bounded contexts in which corresponding domain models are constructed and implemented. Each context contains its own terminology, concepts, and internal logic reflecting a particular part of the business.

Software development then proceeds incrementally, typically focusing on one domain at a time. Domain models are refined, implemented in code, and integrated with other parts of the system through defined interfaces.

Over time, this process leads to a successful outcome. The resulting system appears to reflect the operational structure of the organization: business terminology becomes embedded in the code, domain concepts are consistently represented, and implemented functionality aligns with organizational processes.

As a result, the software system seems to faithfully support and automate the business activities of the organization.

## 4 It Was a Beautiful Dream

The description presented in the previous section may appear convincing. If one assumes that business software development proceeds in this orderly and well-structured manner, the resulting system indeed seems to faithfully reflect the structure of the organization and its processes.

However, such a picture rarely corresponds to the reality of how business organizations actually operate.

First, real businesses are not naturally organized into domains. What exists in practice are business processes that transform inputs into outcomes. A typical end-to-end process in even a relatively simple trading company involves multiple departments: sales, procurement, logistics, accounting, finance, and customer support. When production is involved, the situation becomes significantly more complex, as manufacturing, supply chains, quality control, and planning are all interwoven into the same operational processes.

Second, business processes themselves do not contain natural internal boundaries. What exists instead are organizational boundaries in the form of areas of responsibility assigned to different departments or roles. These boundaries reflect managerial structure rather than intrinsic separations within the process itself.

Third, business rules and invariants rarely belong to isolated parts of an organization. In practice, constraints and policies are typically defined at the level of the entire company. Pricing policies, financial controls, compliance requirements, and operational rules often span multiple processes and departments simultaneously.

For these reasons, the conceptual decomposition suggested by domain modeling does not arise naturally from the structure of real business operations. Instead, it represents an interpretative layer imposed on top of processes that are inherently cross-cutting and organizationally entangled.

## 5 The Structural Reality of Business Processes

To understand the mismatch between domain-based decomposition and actual business operations, it is necessary to examine how real business processes are structured.

A business process is not a conceptual area of responsibility but a continuous operational chain that transforms inputs into outputs. Such processes are typically end-to-end and extend across multiple organizational units.

Consider a simple example of an order fulfillment process in a trading company. The process begins with a customer request handled by the sales department. Once the order is accepted, inventory availability must be verified. If the required goods are not available, procurement is involved to replenish stock. Logistics then prepares and executes shipment, accounting generates invoices, and finance records and processes payments. Customer support may later handle returns, complaints, or warranty claims.

From the perspective of the process itself, these activities form a single operational flow. Departmental boundaries do not divide the process into independent segments; they merely indicate which unit is responsible for performing a particular step.

The same observation applies to more complex organizations. In manufacturing environments, production planning, supply chains, production operations, quality control, maintenance, logistics, and financial accounting are tightly interconnected within the same operational workflows.

As a result, business processes form long chains of interdependent activities that cut across the internal structure of the organization. Each step in the process depends on information, decisions, or actions produced by other parts of the company.

From an operational perspective, the organization therefore behaves not as a collection of isolated domains but as a network of processes that continuously traverse departmental boundaries.

Another important observation follows directly from the structure of such processes. Although different departments participate in different steps of the workflow, they often rely on the same fundamental entities.

In the example above, every part of the process is centered around the same element: the customer. Sales interacts with the customer when the order is placed, logistics ships goods to the customer, accounting issues invoices to the customer, finance processes payments from the customer, and customer support resolves issues reported by the customer.

In other words, the customer is a shared invariant of the entire process. It is not confined to a particular department or organizational unit, nor can it be naturally assigned to a single conceptual domain. Instead, it represents a core entity that is simultaneously required by multiple parts of the organization.

This observation illustrates a structural characteristic of real business operations: certain entities are global to the organization and participate in multiple processes at once. Their role cannot be reduced to the boundaries of individual domains without introducing artificial fragmentation.

## 6 The Problem of Shared Invariants

The example discussed in the previous section illustrates a more general structural property of business organizations: many of the key entities participating in business operations are shared across multiple processes and organizational units. These entities function as company-wide invariants and cannot be naturally confined to a single domain or bounded context.

One of the most obvious examples is the customer. In most business organizations, customers interact with multiple parts of the company. Sales registers orders from customers, logistics delivers goods to customers, accounting issues invoices to customers, finance processes payments received from customers, and customer support handles service requests and complaints. Although these activities belong to different organizational units, they all refer to the same fundamental entity.

A similar situation exists for products and goods. Products are defined, purchased, stored, sold, delivered, and accounted for by different parts of the organization. Procurement departments purchase goods from suppliers, warehouse systems manage inventory, sales departments offer products to customers, logistics handles shipment, and accounting records the financial aspects of these transactions. Despite these different roles, the underlying notion of products and goods remains shared across the entire organization.

Contracts represent another example of a shared invariant. Contracts define formal agreements between the organization and its customers, suppliers, or partners. These agreements influence activities across multiple departments, including sales, legal, finance, procurement, and operations. The contractual terms therefore cannot be restricted to a single conceptual domain without affecting other parts of the organization that depend on them.

Financial resources, typically represented as money, also constitute a global invariant. Payments, invoices, budgets, and financial records connect numerous activities across the organization. Financial constraints and policies are typically defined at the level of the entire company rather than within isolated functional areas.

Taken together, these examples demonstrate that certain entities and constraints naturally span

the entire organization. They participate simultaneously in multiple processes and are referenced by many organizational units. As a consequence, they cannot be cleanly divided along domain boundaries without introducing duplication, synchronization problems, or conceptual inconsistencies.

## 7 Typical Shared Invariants

The analysis of real business processes reveals that certain entities and constraints are shared across the entire organization. These elements participate in multiple processes simultaneously and are required by many departments.

Typical examples of such shared invariants include:

- **Counterparties (customers, suppliers, partners)** — entities interacting with the organization in commercial or contractual relations. They are referenced by sales, procurement, logistics, accounting, finance, customer support, and compliance systems.
- **Products / goods / materials** — involved in production, procurement, inventory management, sales, logistics, and financial accounting, thus participating in multiple operational processes across the organization.
- **Contracts** — defining formal agreements that influence sales operations, legal obligations, procurement, finance, and operational planning.
- **Money and financial transactions** — connecting invoices, payments, budgets, financial reporting, and regulatory compliance across the entire organization.
- **Documents** — orders, invoices, delivery notes, contracts, and other formal records that accompany and coordinate business operations across sales, logistics, accounting, finance, and compliance.
- **Employees** — participants in operational processes who initiate, approve, execute, or control business activities across multiple organizational units.

These entities function as organization-wide invariants. They appear in numerous processes and cannot be naturally confined to a single domain or bounded context without introducing fragmentation or duplication.

## 8 Shared Invariants and Process Graphs

Let an organization be represented by a set of business processes

$$P = \{p_1, p_2, \dots, p_n\}$$

where each process represents an operational transformation that converts inputs into outputs.

Each process operates on a set of business entities

$$E = \{e_1, e_2, \dots, e_m\}.$$

A relation between processes and entities can therefore be defined as

$$R \subseteq P \times E$$

where  $(p_i, e_j) \in R$  means that process  $p_i$  reads, modifies, or produces entity  $e_j$ .

This relation naturally defines a bipartite graph between processes and business entities.

## Definition (Shared Invariant)

An entity  $e \in E$  is called a *shared invariant* if it participates in more than one independent business process:

$$|\{p \in P \mid (p, e) \in R\}| > 1$$

In other words, the entity is required by multiple processes within the organization.

## Observation

In real business organizations, many core entities satisfy the definition of shared invariants. Examples include counterparties, products, contracts, financial transactions, documents, and employees.

## Consequence

If system decomposition assigns entities exclusively to isolated domains, shared invariants must be replicated across domains in order to support all processes in which they participate.

## Result

Replication of shared invariants inevitably introduces one or more of the following effects:

- duplication of organizational data,
- synchronization mechanisms between domains,
- reconciliation procedures to restore consistency,
- loss of a unified organizational view of key entities.

These effects arise not from implementation errors but from the structural mismatch between domain-based decomposition and the process-oriented structure of business reality.

Thus, shared invariants are not accidental properties of enterprise systems but a direct consequence of the process structure of business operations.

## Proposition

Let an organization be represented by a set of business processes  $P = \{p_1, p_2, \dots, p_n\}$  and a set of business entities  $E = \{e_1, e_2, \dots, e_m\}$ .

Let the relation

$$R \subseteq P \times E$$

denote the participation of entities in business processes, where  $(p, e) \in R$  means that process  $p$  reads, modifies, or produces entity  $e$ .

If an entity  $e$  participates in more than one independent process

$$|\{p \in P \mid (p, e) \in R\}| > 1,$$

then  $e$  cannot be confined to a single isolated domain without introducing duplication or synchronization mechanisms between domains.

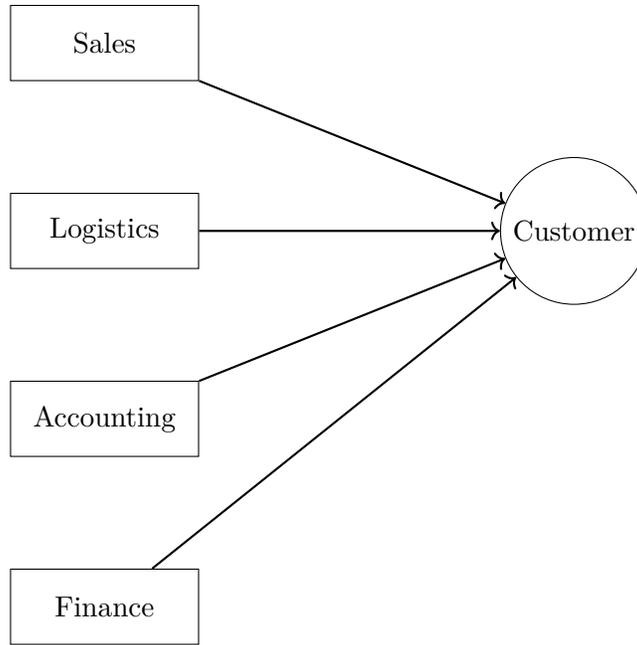


Figure 1: A shared business entity participating in multiple processes



Figure 2: Domain decomposition duplicates shared invariants and introduces synchronization between domains

### Proof (informal)

If the entity is assigned to only one domain, other domains that participate in processes involving the same entity must either access that domain directly or maintain their own representations of the entity.

In the latter case the entity is duplicated across domains, requiring synchronization mechanisms to maintain consistency. In the former case the entity effectively becomes shared across domains.

Therefore, entities participating in multiple processes necessarily form shared invariants of the system.

This proposition explains why shared invariants inevitably appear in enterprise systems and why strict domain isolation conflicts with the structure of real business operations.

## 9 Invisible Consequences of the Absence of Shared Invariants

If the principles of Domain-Driven Design are applied strictly and organizational units are modeled as separate domains, each domain naturally develops its own representation of key business entities. In such a situation, shared invariants disappear and are replaced by multiple local versions of the same concept.

Consider the example of a universal bank that includes, among others, a deposit department and a credit department. If each of these departments models its own clients independently, the system may contain “deposit clients” and “loan clients” rather than a single company-wide notion of a customer.

Although such separation may appear consistent within individual domains, it introduces several structural problems for the organization as a whole.

First, the bank loses the ability to evaluate the cost of acquiring a client. Marketing and acquisition expenses cannot be attributed to the overall customer relationship if clients exist only within individual product domains. As a result, cross-product pricing strategies become difficult or impossible to implement.

Second, the organization loses the ability to reuse the same client across multiple services. A person who already interacts with the bank through deposits may need to be recreated as a separate entity in the credit system, which leads to duplication and fragmentation of customer information.

Third, departments that operate across multiple products cannot access a unified view of the customer base. For example, a treasury or trading department may require access to the full list of clients for investment products or brokerage services. If customer data is confined to isolated domains, such access becomes technically complicated or organizationally restricted.

Fourth, risk management becomes significantly more difficult. Credit risk, liquidity risk, and regulatory compliance often depend on the overall exposure of the bank to a particular client. Without a unified representation of customers, it becomes challenging to assess the total risk associated with an individual or organization.

Finally, operational analytics and reporting suffer from fragmentation. Business intelligence systems require a consolidated view of customers, products, contracts, and financial flows. When these entities are distributed across isolated domain models, analytical systems must reconstruct relationships that should have existed naturally within the operational model.

These effects illustrate a broader problem: when shared invariants are eliminated from the system model, the organization loses the ability to reason about its operations at the company-wide level.

## 10 Domains as an Artificial Decomposition

The observations discussed in the previous sections suggest that the decomposition of software systems into domains does not naturally follow the structure of real business operations. Instead, domains are introduced as an interpretative layer used by software engineers to manage conceptual complexity during system design.

In practice, business organizations operate through processes, responsibilities, and shared operational entities. Activities performed by different departments are tightly connected through these processes, and many key entities — such as counterparties, products, contracts, documents, and financial transactions — are inherently shared across the organization.

Domain decomposition, however, attempts to partition this operational reality into relatively independent conceptual areas. Each domain develops its own models, terminology, and data rep-

representations, allowing development teams to reason locally about a subset of the system. From a software engineering perspective, such partitioning may simplify implementation and team coordination.

Nevertheless, this decomposition does not correspond to a natural division within the business itself. Domains rarely coincide with the actual structure of operational processes, which typically cross organizational boundaries and involve multiple functional areas.

As a result, domain boundaries are not discovered in the business structure but constructed during the software design process. They represent an engineering strategy for managing complexity rather than a faithful representation of how the organization operates.

This distinction is important. When domain boundaries are treated as if they reflect real separations in the business, the resulting models may fragment entities that are inherently shared across the organization. The system then requires additional mechanisms — integration layers, data synchronization, and reconciliation procedures — to restore relationships that existed naturally in the operational environment.

In this sense, domains should be understood not as intrinsic elements of business reality but as artifacts of software design introduced to structure the development process.

## 11 The Small Core of Business Reality

Despite the complexity of real business operations, most organizations rely on a relatively small set of fundamental informational entities. These entities form the core informational structure of the enterprise and are shared across multiple processes, departments, and systems.

Historically, enterprise information systems were designed according to several basic principles that governed the management of such data. These principles can be summarized as follows:

1. Information should not be deleted without trace.
2. Information should not be duplicated across independent systems.
3. Information should be shared across the organization whenever it represents a common business entity.

These principles formed the conceptual basis of many enterprise information architectures and were reflected in the development of Master Data Management (MDM) approaches. MDM systems were designed to maintain consistent representations of core organizational entities such as customers, counterparties, products, contracts, and other reference data used across multiple operational processes.

In recent years, however, the widespread adoption of domain-oriented development practices has often shifted attention away from centralized management of shared informational entities. When systems are designed strictly around domain boundaries, representations of the same organizational entities may appear independently within multiple domains, which reintroduces duplication and fragmentation of information.

This observation demonstrates that domain-based decomposition cannot serve as a primary model of business reality. However, it highlights the necessity of explicitly representing the shared informational core of the organization.

For this reason, when domain-oriented architectures are employed, it is advisable to introduce a special domain responsible for maintaining the organization-wide reference entities. This domain can be viewed as a *Domain Zero* — a foundational domain that contains shared organizational data such as counterparties, products, contracts, documents, and other master records.

In practical terms, Domain Zero corresponds to what was historically known as a Master Data Management domain. Its role is to preserve the consistency and integrity of the shared informational core of the organization while allowing other domains to operate on this common foundation.

Without a shared informational core, domain decomposition inevitably fragments the representation of the organization itself.

## Architectural Principle

Enterprise systems must preserve a shared informational core that maintains organization-wide invariants across all operational domains.

## 12 Regulatory Constraints and Shared Invariants

The importance of shared invariants becomes particularly evident in regulated industries, where organizational entities must be represented consistently across the entire system for regulatory and risk-management purposes.

A typical example can be observed in banking systems. Universal banks operate through multiple product areas, including deposits, loans, credit cards, brokerage services, payments, and treasury operations. Each of these areas may be implemented as a separate subsystem or domain.

From a domain-oriented perspective, it would be natural for each of these subsystems to maintain its own representation of a client. Deposits systems may maintain deposit customers, credit systems may maintain borrowers, and brokerage platforms may maintain trading clients.

However, banking regulation requires the institution to maintain a unified view of its exposure to each client. Risk management, compliance, and capital adequacy regulations require banks to calculate the total financial exposure associated with a single customer or counterparty.

Such exposure may include deposits, outstanding loans, credit lines, derivative positions, securities holdings, guarantees, and other financial obligations. These positions are typically distributed across multiple operational systems.

As a result, the bank must maintain a single authoritative representation of the client that can be referenced across all product areas. Without such a unified entity, it becomes impossible to compute the total risk exposure of the organization.

For this reason, banking architectures traditionally introduce a centralized customer master or counterparty master, often referred to as a Customer Information File (CIF) or Party Master. This component serves as the authoritative source of client identity across the entire system.

This architectural requirement illustrates a broader principle: in regulated industries, shared invariants are not merely architectural conveniences but structural and regulatory necessities. Systems must maintain consistent representations of key organizational entities across multiple operational domains.

From this perspective, the existence of a shared informational core is not optional. It is a direct consequence of the regulatory and operational structure of the organization.

## Structural Consequence

The observations discussed in the previous sections lead to a structural conclusion. If business processes depend on shared invariants that participate in multiple operational domains, the system must maintain a single authoritative representation of these entities.

Otherwise, each domain will introduce its own representation of the same organizational entities, which inevitably leads to duplication, synchronization mechanisms, reconciliation procedures, and loss of organizational consistency.

Therefore, any architecture that aims to faithfully represent the structure of enterprise operations must explicitly include a shared informational core responsible for maintaining organization-wide entities.

In domain-oriented architectures, this core can be modeled as a foundational domain, referred to in this paper as *Domain Zero*.

## 13 Domain Zero Architectural Pattern

The analysis presented in this paper leads to a practical architectural principle for enterprise systems. When domain-oriented development is applied to organizations whose operations rely on shared invariants, a special foundational domain must be introduced to maintain the organization-wide informational core.

This architectural solution can be described as the *Domain Zero* pattern.

### Intent

Provide a single authoritative domain responsible for maintaining organization-wide entities that participate in multiple business processes.

### Problem

Domain-oriented decomposition tends to replicate shared entities across multiple domains. This leads to duplication of organizational data, synchronization mechanisms between domains, and fragmentation of the enterprise information model.

### Solution

Introduce a foundational domain — *Domain Zero* — that maintains the shared invariants of the organization.

Typical entities managed by *Domain Zero* include:

- counterparties (customers, suppliers, partners)
- products, goods, and materials
- contracts
- financial transactions
- documents
- employees

All other domains reference these entities rather than maintaining independent copies.

### Structure

*Domain Zero* forms the informational core of the enterprise architecture. Operational domains interact with it as consumers of shared reference data.

## Consequences

The Domain Zero pattern preserves the integrity of shared organizational entities while allowing domain-oriented development to structure behavior and services around specific areas of responsibility.

It eliminates duplication of shared invariants and restores a unified organizational model across the system.

## Relationship to MDM

Conceptually, Domain Zero corresponds to what has historically been known as Master Data Management (MDM). However, instead of existing as a separate integration layer, it is explicitly modeled as a domain within the system architecture.

## Architectural Law

Let an enterprise system be represented by a set of processes  $P$  and a set of entities  $E$ , connected by a relation

$$R \subseteq P \times E.$$

For any entity  $e \in E$ , define its process participation degree

$$\text{deg}(e) = |\{p \in P \mid (p, e) \in R\}|.$$

If

$$\text{deg}(e) > 1,$$

then the entity necessarily belongs to the shared informational core of the system.

Such entities form the structural foundation of enterprise architectures and must be maintained within a common domain (denoted in this paper as *Domain Zero*).

Entities with participation degree greater than one therefore constitute the shared invariants of the enterprise system.

## 14 Domain Zero as an Evolutionary Correction

The analysis presented in this paper demonstrates that strict domain-based decomposition does not reflect the structural nature of real business organizations. Enterprise systems operate around shared informational entities that naturally span multiple processes and organizational units.

When such entities are artificially confined to isolated domains, their representations inevitably fragment across the system, introducing duplication, synchronization mechanisms, and organizational inconsistency.

At the same time, domain-oriented development practices have become deeply embedded in contemporary software engineering. Large numbers of existing systems, frameworks, and development methodologies rely on domain-based decomposition.

For this reason, a complete rejection of domain-oriented architectures is rarely feasible in practice.

A more realistic approach is to introduce a corrective architectural mechanism that restores the shared informational core of the organization while preserving the operational structure of domain-oriented systems.

This mechanism can be described as the *Domain Zero* pattern.

Domain Zero represents a foundational domain responsible for maintaining organization-wide informational entities that participate in multiple processes.

Typical entities maintained within Domain Zero include:

- counterparties (customers, suppliers, partners)
- products, goods, and materials
- contracts
- financial transactions
- documents
- employees

Operational domains interact with Domain Zero as consumers of shared organizational data rather than maintaining independent representations of these entities.

Conceptually, Domain Zero corresponds to the role historically played by Master Data Management systems. However, instead of being treated as an external integration layer, it is explicitly modeled as a foundational domain within the system architecture.

In this sense, Domain Zero should be understood as an evolutionary correction that allows domain-oriented architectures to coexist with the structural realities of enterprise systems.

Business reality is organized around shared invariants, not domains. Domains are tools of software engineering, whereas shared invariants define the structure of the enterprise itself.

## 15 Conclusion

This paper examined the relationship between domain-oriented software design and the operational structure of real business organizations. The analysis demonstrated that business activity is organized around processes and shared informational entities rather than isolated conceptual domains.

Entities such as counterparties, products, contracts, financial transactions, documents, and employees naturally span the entire organization and participate in multiple operational processes. Treating these entities as domain-specific objects leads to fragmentation, duplication of information, and loss of organizational consistency.

For this reason, domain decomposition cannot serve as a primary model of business reality. Enterprise systems require a shared informational core that preserves the integrity of organization-wide entities.

In domain-oriented architectures this role should be implemented as a foundational *Domain Zero* responsible for maintaining master organizational data.

Business reality is organized around shared invariants, not domains.

## A Domain Zero — Foundational Shared-Invariant Domain

This appendix provides a practical architectural pattern intended primarily for practitioners who currently employ Domain-Driven Design (DDD). While the main body of the paper presents a structural critique of strict domain-based decomposition, the Domain Zero pattern offers an evolutionary mechanism that allows existing domain-oriented systems to address the problem of shared invariants.

The concept introduced here should therefore be understood as a transitional architectural solution. In future work, the Domain Zero idea will be further developed into a broader architectural paradigm for modeling enterprise systems.

### Intent

Provide a single authoritative domain responsible for maintaining organization-wide entities that participate in multiple independent business processes. Domain Zero forms the informational core of the enterprise upon which all other domains depend.

### Motivation

In real organizations, key entities such as customers, counterparties, products, contracts, employees, documents, and financial transactions participate simultaneously in multiple business processes.

Such entities function as *shared invariants* of business reality and cannot be cleanly partitioned across domain boundaries without introducing duplication, model divergence, and synchronization mechanisms.

Strict decomposition into bounded contexts therefore tends to produce multiple local representations of the same organizational entities, leading to:

- duplication of data,
- divergence of models,
- synchronization mechanisms between domains,
- loss of a unified organizational view,
- incorrect analytics and risk management.

Domain Zero addresses this issue by introducing a single layer of authoritative shared entities.

### Applicability

Use Domain Zero when:

- an entity participates in two or more independent business processes;
- the organization requires a unified representation of customers, products, or contracts;
- regulatory or accounting requirements demand consistent organizational records;
- enterprise analytics requires a complete organizational view;
- microservice or domain decomposition leads to fragmentation of core data.

In structural terms, Domain Zero becomes necessary whenever an entity participates in multiple independent processes.

## Structure

Domain Zero acts as a foundational bounded context located beneath all operational domains.

It maintains shared invariants of the enterprise, including:

- counterparties (customers, suppliers, partners),
- products and goods,
- contracts,
- documents,
- employees,
- financial transactions.

Operational domains reference these entities but do not maintain independent copies.

Domain Zero therefore forms the informational foundation of the enterprise architecture.

## Participants

### Domain Zero

- maintains master records for shared invariants;
- provides stable and immutable identifiers;
- manages lifecycle and integrity of core entities;
- ensures analytical and regulatory consistency.

### Operational Domains

- implement business behavior and process logic;
- reference Domain Zero entities;
- avoid duplicating shared invariants.

### Anti-Corruption Layer

- protects Domain Zero from inconsistent external models;
- normalizes incoming data structures.

### Event Consumers

- receive updates from Domain Zero;
- maintain local read models if required.

## Consequences

### Positive

- single authoritative source of organizational entities;
- elimination of duplicated master data;
- consistent analytics and reporting;
- simplified integration between domains;
- support for regulatory compliance.

### Negative

- Domain Zero becomes a critical architectural component;
- strong governance of shared data is required;
- higher data-quality requirements;
- event propagation and integration must be carefully designed.

## Examples

### Banking systems

- Customer Information File (CIF) — unified client identity;
- Product Master — unified product catalog;
- Contract Master — unified contract registry;
- Employee Master — unified staff registry.

### Retail

- unified product catalog;
- unified supplier registry;
- unified customer loyalty system.

### Manufacturing

- material master;
- bill-of-materials registry;
- enterprise counterparty registry.

## Known Uses

- Master Data Management (MDM) systems;
- Customer Information File (CIF) in banking;
- Product Information Management (PIM) in retail;
- Enterprise Party Models in insurance platforms;
- enterprise data models in traditional ERP architectures.

## Related Patterns

- Shared Kernel
- Anti-Corruption Layer
- Canonical Data Model
- Event-Driven Master Data Distribution
- Process-Centric Architecture

## Summary

Domain Zero represents an evolutionary architectural correction that restores the shared informational core of enterprise systems within domain-oriented architectures.

Business reality is organized around shared invariants rather than isolated domains. Domain Zero provides a practical mechanism for preserving these invariants while allowing domain-oriented development to structure operational behavior.

## References

- [1] E. Evans, *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Boston: Addison-Wesley, 2003, ISBN: 978-0321125217.
- [2] V. Vernon, *Implementing Domain-Driven Design*. Boston: Addison-Wesley, 2013, ISBN: 978-0321834577.
- [3] V. Vernon, *Domain-Driven Design Distilled*. Boston: Addison-Wesley, 2016, ISBN: 978-0134434421.
- [4] S. Newman, *Building Microservices: Designing Fine-Grained Systems*. Sebastopol: O'Reilly Media, 2015, ISBN: 978-1491950357.
- [5] R. C. Martin, *Clean Architecture: A Craftsman's Guide to Software Structure and Design*. Boston: Prentice Hall, 2017, ISBN: 978-0134494166.
- [6] M. Kleppmann, *Designing Data-Intensive Applications*. Sebastopol: O'Reilly Media, 2017, ISBN: 978-1449373320.
- [7] L. Bass, P. Clements, and R. Kazman, *Software Architecture in Practice*, 3rd ed. Boston: Addison-Wesley, 2012, ISBN: 978-0321815736.
- [8] M. Shaw and D. Garlan, *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.

- [9] D. Perry and A. Wolf, “Foundations for the study of software architecture,” *ACM SIGSOFT Software Engineering Notes*, vol. 17, no. 4, pp. 40–52, 1992.
- [10] A. A. Nekludoff, “Architecture of complex systems,” 2026. DOI: 10.5281/zenodo.18277986 [Online]. Available: <https://doi.org/10.5281/zenodo.18277986>
- [11] A. A. Nekludoff, “Architecture as a directed object–relation graph: A minimal and complete model of complex systems,” 2026. DOI: 10.5281/zenodo.18768371 [Online]. Available: <https://doi.org/10.5281/zenodo.18768371>