

Implementation Roadmap: Application Service Framework

Version: 2.0

Date: July 29, 2025

Copyright 2025 Sebastián Samaruga

1. Introduction

This document provides a deeply technical, implementation-focused roadmap for the Application Service framework. It moves beyond high-level architecture to specify the concrete implementation details of each service, model, and pattern. The central theme is the realization of a fully reactive, functional, and behavior-driven system, ensuring a non-blocking, event-driven dataflow from data source to user interaction.

We will provide detailed explanations and code examples for the practical application of key patterns and frameworks, including:

- **Reactive & Functional Programming:** Using Spring WebFlux, Project Reactor, and functional composition.
- **Core Patterns:** DCI, DDD, Dynamic Object Model, and the Actor-Role pattern.
- **Semantic & AI Technologies:** Formal Concept Analysis (FCA), Spring AI for LLM integration, and the Model Context Protocol (MCP).
- **Enterprise Integration:** Java EE Connector Architecture (JCA) and JavaBeans Activation Framework (JAF) principles.
- **Decentralized Identity:** W3C Decentralized Identifiers (DIDs) for verifiable, interoperable identity.

Phase 1: Core Infrastructure & Reactive Data Ingestion (Months 1-3)

Objective

Establish a robust, scalable, and fully reactive microservices foundation and a versatile, non-blocking data ingestion pipeline.

1.1. Datasource Service (Java, Spring WebFlux)

This service is the entry point for all external data, built entirely on a non-blocking stack.

- **Reactive Core:** We will use Spring WebFlux's functional handler functions instead of traditional `@RestController` annotations. This provides a more explicit, functional way to define endpoints.
- **JCA Integration (for Enterprise Systems):** The Java EE Connector Architecture is critical for transactional, bidirectional communication with legacy systems like ERPs.

- **Implementation:** We will implement a `javax.resource.spi.ResourceAdapter`.
- **Inbound (Events from EIS):** The adapter's `endpointActivation` method will be given a `javax.resource.spi.endpoint.MessageEndpointFactory` by the application server. When the adapter receives an event from the EIS (e.g., an SAP IDoc), it uses the factory to get a `MessageEndpoint`. This endpoint is a proxy that, when invoked, pushes the message into a Project Reactor Sinks.Many processor, which acts as the source of a Flux. This bridges the listener-based JCA model with the reactive streams world.
- **Outbound (Write-Back):** The adapter exposes a `javax.resource.cci.ConnectionFactory`. The Activation Service will later use this to get a `Connection` and execute an Interaction (e.g., updating a record). This provides the crucial mechanism for synchronizing state back to source systems, as detailed in the [Eclipse JCA documentation](#).
- **Reactive Ingestion Adapters:**
 - **RestApiAdapter:** Uses `WebClient` to consume external APIs. It natively returns a `Flux<T>`, allowing the service to stream paginated results without holding a thread, processing each item as it arrives.
// Example: Fetching paginated items reactively

```
public Flux<Item> fetchAllItems() {
    return webClient.get()
        .uri("/items?page=0")
        .retrieve()
        .bodyToFlux(Item.class)
        .expand(lastItem -> { // expand operator for pagination
            if (lastItem.isLastPage()) {
                return Mono.empty();
            }
            return webClient.get()
                .uri("/items?page=" + (lastItem.getPageNumber() + 1))
                .retrieve()
                .bodyToFlux(Item.class);
        });
}
```
 - **R2DBCAdapter:** For supported SQL databases, it will use `R2DBC` (`spring-boot-starter-data-r2dbc`) to perform non-blocking database queries, returning a `Flux<Row>`.
- **Functional Transformation:** The transformation from source format to SPO triples will be a pure function within a reactive pipeline, aligning with functional

principles from resources like ["Functional Programming in JavaScript"](#).

// Example: Functional transformation in a reactive pipeline

```
public Flux<Statement<String, String, String, String>>
```

```
processSourceData(Flux<SourceData> sourceStream) {
```

```
    return sourceStream
```

```
        .flatMap(data -> Flux.fromIterable(transformer.toTriples(data))); // 1-to-many
```

```
transform
```

```
}
```

1.2. Augmentation Service (Java, Spring Cloud Stream)

This service is the reactive backbone, orchestrating the dataflow between all other services.

- **Reactive Dataflow:** It will bind `java.util.function.Function<Flux<T>, Flux<R>>` beans to Kafka topics. This is a powerful feature of Spring Cloud Stream that allows for the entire service to be composed of reactive functions.

// Example: A Spring Cloud Stream function bean

```
@Bean
```

```
public Function<Flux<RawStatement>, Flux<AggregatedStatement>> aggregate() {
```

```
    return rawFlux -> rawFlux
```

```
        .log()
```

```
        .flatMap(aggregationService::processStatement);
```

```
}
```

- **Saga Pattern for Distributed Transactions:** We will implement the Saga pattern using `Flux.usingWhen` to manage long-running, distributed transactions. This operator is perfect for handling resource allocation and cleanup in a reactive chain.

// Simplified Saga example for an ingestion process

```
Flux<Void> ingestionSaga = Flux.usingWhen(
```

```
    transactionManager.begin(), // 1. Begin transaction, get a resource
```

```
    tx -> aggregationService.process(tx) // 2. Use the resource
```

```
        .then(alignmentService.process(tx)),
```

```
    transactionManager::commit, // 3. On success, commit
```

```
    (tx, err) -> transactionManager.rollback(tx), // 4. On error, rollback
```

```
    transactionManager::rollback // 5. On cancellation, rollback
```

```
);
```

1.3. Registry Service (Java, Spring Boot, Neo4j)

This is the central repository for the unified property graph.

- **Reactive API:** The REST API will be built with Spring WebFlux functional endpoints, returning `Mono<ServerResponse>` for writes and `Flux<Statement>` for reads.
- **Non-Blocking Database Interaction:** The official Neo4j Java driver is blocking. To prevent this from consuming a precious event-loop thread, we will offload the work to a dedicated scheduler, a core tenet of reactive programming as explained in ["Building Reactive Microservices with Spring WebFlux"](#).

```
public Mono<Void> saveStatement(Statement stmt) {  
    return Mono.fromRunnable(() -> {  
        // This blocking driver call runs on a different thread pool  
        try (Session session = driver.session()) {  
            session.run("MERGE (s:Resource {uri: $s_uri}) ...", parameters(...));  
        }  
    }).subscribeOn(Schedulers.boundedElastic()).then();  
}
```

Phase 2: Semantic Core & Knowledge Representation (Months 4-7)

Objective

Transform raw data into an interconnected, semantically rich knowledge graph using reactive streams, Formal Concept Analysis, and AI/ML models.

2.1. Deep Dive: The Reference Model & Prime Number Semantics

This model moves from string-based URIs to a formal, mathematically grounded identification system.

- **ID & IDOccurrence:** An ID is the canonical concept of an entity, identified by a unique `primeID`. An `IDOccurrence` is an ID appearing in a specific role within a specific context.
- **Prime Number Semantics:** We leverage the Fundamental Theorem of Arithmetic. An `IDOccurrence`'s "embedding" is a set of `primeIDs` defining its full context. Similarity is a deterministic Jaccard Index on these sets. This mathematical foundation is inspired by the lattice theory concepts discussed in sources like [Sowa's "Mathematical Foundations"](#).

2.2. Deep Dive: Formal Concept Analysis (FCA) with Prime IDs

FCA is a cornerstone of the Aggregation Service for inferring types and hierarchies.

Using primeIDs as identifiers for objects and attributes provides unique mathematical properties for inference.

- **FCA Context Model:** We will define a formal context as a triplet (G, M, I) where G is a set of objects, M is a set of attributes, and I is a binary relation $I \subseteq G \times M$. The Index Service will be able to retrieve these contexts. For example, for a statement's predicate:
 - **FCA Context Relation:** The predicate itself (e.g., worksFor).
 - **FCA Context Objects:** The set of subjects in statements with that predicate (e.g., {Alice, Bob}).
 - **FCA Context Attributes:** The set of objects in statements with that predicate (e.g., {Google, StartupX}).
- **Inference via Prime Products:** This is the model's key innovation. A "formal concept" in the resulting lattice is a pair (A, B), where A is a set of objects and B is the set of attributes they all share. The "intent" B can be uniquely identified by the product of its attribute primeIDs.
 - **Subsumption Checking:** A concept C1 is a sub-concept of C2 if C1's intent-product is cleanly divisible by C2's intent-product. This transforms expensive set logic into simple integer arithmetic, making large-scale hierarchy inference computationally feasible, a technique vital for knowledge discovery as explored in ["Formal Concept Analysis for Knowledge Discovery and Data Mining"](#).

2.3. Aggregation Service (Java, Spring AI, fcalib)

- **Reactive AI Integration:** Embeddings will be generated within the reactive stream using Spring AI's ReactiveEmbeddingClient. This ensures that the network call to an embedding model (like one from Hugging Face or a local Ollama instance via Spring AI) is non-blocking.
// Example: Generating embeddings within a reactive stream

```
.flatMap(statement ->  
    // Spring AI's reactive client  
    reactiveEmbeddingClient.embed(statement.getObject())  
        .map(embedding -> statement.withEmbedding(embedding))  
    )
```
- **FCA Implementation:** The inferTypeFromPredicates method will use the **fcalib** library. The set of predicates for a group of subjects is used to build a FormalContext. The resulting ConceptLattice provides the type hierarchy, which is then flattened back into a Flux of type assertion statements.

2.4. Alignment Service (Java, RDF4J)

- **Reactive Ontology Matching:** This service consumes the `Flux<AggregatedStatement>` and uses the **RDF4J** framework's `MemoryStore` for in-memory graph operations. It will load pre-defined upper ontologies (e.g., `Schema.org`) and use the SPARQL engine with SHACL rules to find and materialize equivalences (`owl:sameAs`, `rdfs:subClassOf`).
- **Link Completion:** Implemented with SPARQL CONSTRUCT queries within the reactive pipeline. For example, a query can find paths like (A)-[:hasRole]->(B) and (B)-[:partOf]->(C) to infer a new link (A)-[:contributesTo]->(C).

Phase 3: Activation & Behavior-Driven Interactions (Months 8-10)

Objective

Infer and enable the execution of business processes by implementing a dynamic, message-driven model based on DCI, DDD, and a JAF-inspired semantic engine.

3.1. Deep Dive: The Activation Model's Dynamic Object Model (DOM)

The runtime logic is a direct implementation of the ideas found in works like ["Dynamic Object Model"](#) and the Actor Role pattern.

- **DDD (Domain-Driven Design):** The entire Activation Service is a single Bounded Context. Its Ubiquitous Language consists of entities like Context, Role, Interaction, etc., following the principles from Eric Evans' ["Domain-Driven Design: Tackling Complexity in the Heart of Software"](#).
- **DCI (Data, Context, and Interaction):** This pattern is the blueprint for the runtime logic. An Interaction (Context) "casts" plain data Instances into Actors by dynamically injecting Roles (behavior). This avoids bloating data objects, a core tenet of DCI described by [James Coplien and Trygve Reenskaug](#).
 - **Role (Functional Interface):** A Role is a `Function<Flux<ActorState>, Flux<TransformedState>>`. It's a functional interface defining the behavior an Actor will perform.
 - **Interaction (Reactive Orchestrator):** An Interaction subscribes to the Flux streams representing its Actors' states and applies the Role functions to drive the dataflow.

3.2. The Activation Service: A JAF-Inspired Semantic Engine

The Activation Service operates like a distributed, semantic JavaBeans Activation Framework (JAF), as described in ["REST easy with the JavaBeans Activation Framework"](#).

- **ContentTypeDataHandler:** For each Content-Type (e.g., buy-able), a

corresponding Spring bean implementing this interface is registered. This handler defines the available Verbs (commands like BUY) and the Dataflow (sequence of Transforms) for each verb.

- **Context Inference from Content Types:** The availability of use cases (Contexts) is inferred dynamically. If the system finds a set of Actors whose ContentTypes match the required Roles for a Context, that Context becomes available.
- **JCA for Transactional Write-Back:** When an Interaction's Dataflow completes, its final Transform is processed by the relevant ContentTypeDataHandler. This handler obtains a JCA Connection from the Datasource Service and invokes the outbound transaction on the backend ERP, guaranteeing data consistency.

Phase 3.5: LLM Integration & Agentic Architecture (Parallel)

Objective

Elevate the Activation Service to an intelligent, agentic system capable of communicating with LLMs and other ApplicationService instances using standardized protocols.

3.1. Deep Dive: The ApplicationService as a Model Context Protocol (MCP) Server

The ApplicationService will expose an MCP Server endpoint, allowing external clients (like LLM-based agents) to interact with its capabilities in a standardized way. We will use **Spring AI** as the primary tool to bridge our internal services with the LLM world, as discussed in the [MCP Introduction](#).

- **MCP Endpoint Implementation (Spring WebFlux):** A single REST endpoint (/mcp) will handle all MCP requests.
- **Exposing Capabilities via MCP & Spring AI:**
 1. **Resources (Aggregation/Index):** An MCP client can ask for resources. "Find me resources similar to 'a senior Java developer'."
 - **Implementation:** The MCP endpoint routes this to the Index Service. The text query is fed into Spring AI's ReactiveEmbeddingClient to get a vector. This vector is used to perform a similarity search in the vector DB.
 2. **Tools (Activation/Registry):** An MCP client can request to use a tool. "Execute the 'OnboardNewEmployee' tool for resource 'user:JohnDoe'."
 - **Implementation:** The MCP endpoint maps the tool name "OnboardNewEmployee" to an Activation Context. It then instantiates an Interaction for that Context. The LLM decides *what* to do; our framework provides the verifiable, stateful Tool to do it.
 3. **Prompt Templates (Alignment/Naming):** An MCP client can request a template for complex reasoning. "Give me the 'ConceptAlignment' prompt

template to compare 'Customer' and 'Client'."

- **Implementation:** The endpoint fetches a pre-defined prompt string from the Naming Service. The MCP client populates this and sends the completed prompt to an LLM using Spring AI's ReactiveChatClient. The LLM's response can then be fed back into the system to augment the Graph Model.

3.2. Deep Dive: COST (CONversational State Transfer) & The HAL Protocol

The communication between the Producer Service and the Activation Service will be implemented as COST, a stateful, conversational protocol built using the [Hypertext Application Language \(HAL\) specification](#).

- **HAL Link with Placeholders:** A link for a next action is a template for the next Transform message.

// Part of a HAL response for an Interaction

```
"_links": {
  "next": [
    {
      "href": "/interactions/123/transform",
      "method": "POST",
      "name": "SelectProductForPurchase",
      "schema": { // The placeholder definition
        "type": "object",
        "properties": {
          "selectedProduct": {
            "type": "string",
            "description": "The DID of the product to purchase.",
            "_links": {
              "possibleValues": { "href":
"/interactions/123/roles/Product/possibleActors" }
            }
          }
        }
      }
    }
  ]
}
```

- **Conversational Flow:**

1. The client UI sees the SelectProductForPurchase action and its schema.

2. It sees that the selectedProduct placeholder has a possibleValues link. It performs a GET on that link to retrieve a list of available products to render in a dropdown.
3. The user selects a product. The client constructs the Transform message body according to the schema and POSTs it to the href.
4. The server processes the transform and responds with the new state and the next set of possible actions. This makes the client incredibly dynamic and resilient to changes in the backend workflow.

3.3. Deep Dive: W3C Decentralized Identifiers (DIDs)

All canonical resource IDs are W3C DIDs, providing a foundation for verifiable provenance and secure interoperability, as outlined in the [W3C DID Use Cases](#).

- **Implementation:**

1. **Creation:** In the Aggregation Service, when a new resource is first encountered, we will use a library like **did-common-java** to generate a did:ion or did:key. The resource's primeID can be part of the DID string for deterministic generation.
2. **Storage:** The generated DID Document (containing cryptographic keys and service endpoints like the resource's MCP endpoint) is stored in the Registry Service's property graph.

- **Enabled Features:**

- **Verifiable Provenance:** Any Statement can be cryptographically signed using the private key associated with the service's own DID. Downstream consumers can verify this signature, guaranteeing data integrity.
- **Secure Interoperability:** When one ApplicationService acts as an MCP Client to another, it can use DID-Auth to authenticate, eliminating the need for traditional API keys or OAuth tokens.

Phase 4: The Behavior-Driven API & UI (Months 11-12)

Objective

Expose the framework's capabilities through a fully reactive API and a real-time, behavior-driven user interface.

4.1. Producer Service (API/Frontend - Java/Spring WebFlux, React)

- **Fully Reactive API:** The entire API will be built with Spring WebFlux.
- **Server-Sent Events (SSE):** For real-time updates on long-running Interactions, the API will use SSE. A client can subscribe to an endpoint which returns a Flux<InteractionState> with the Content-Type of text/event-stream.

```

@GetMapping(value = "/interactions/{id}/stream", produces =
MediaType.TEXT_EVENT_STREAM_VALUE)
public Flux<InteractionState> streamInteractionUpdates(@PathVariable String id)
{
    return interactionService.getInteractionUpdates(id);
}

```

- **Frontend (React with RxJS):** The React frontend will use **RxJS** to manage the SSE streams, binding component state directly to an Observable so the UI updates automatically. This aligns with the "Thinking in React" and "Thinking in RxJava" mental models, as discussed in [A Complete Roadmap for Learning RxJava](#).

4.2. Complete End-to-End Example: Onboarding a New Product

1. **Ingestion:** A new row for "Pro-Lite Running Shoe" is added to a legacy ERP database. The **Datasource Service's** JCA adapter captures this event and produces a raw triple: ("sku:PL-789", "stock", "500").
2. **Aggregation:** The **Aggregation Service** receives the triple.
 - It's the first time seeing "sku:PL-789". It generates a primeID (e.g., 937) and a W3C DID: did:key:z6Mkp....
 - Using **Spring AI's** ReactiveEmbeddingClient, it generates a vector embedding for the product's name and description.
 - It runs this new entity through its **FCA** engine. Based on its attributes (stock, price, category), the FCA lattice places it within the PurchasableItem formal concept.
3. **Alignment:** The **Alignment Service** receives the aggregated statement.
 - Using its **RDF4J** store, it matches PurchasableItem to schema.org/Product.
 - It analyzes the resource's JCA provenance (knowing it comes from an ERP with an updateStock capability) and the graph patterns, inferring and adding the inventory-trackable and buy-able ContentTypes to the node in the Registry.
4. **Activation & Interaction (COST/HAL):**
 - A user browses the store via the **Producer Service** (React frontend).
 - The frontend makes a call to get available products. The **Activation Service** knows did:key:z6Mkp... is buy-able and includes it in the list.
 - The user adds it to their cart. The frontend POSTs a Transform message based on the HAL template provided by the API: { "transformName": "AddToCart", "payload": { "product": "did:key:z6Mkp..." } }.
 - The **Activation Service** instantiates a Checkout Interaction, casting the user

and the product into Buyer and Item Actors.

5. **Agentic Interaction (MCP):**

- An external marketing agent (an LLM) connects to the ApplicationService's **MCP Server**.
- It asks for "tools related to new products". The server, via the **Activation Service**, exposes a LaunchMarketingCampaign tool because it has detected a new Product Actor that is not yet part of any campaign.
- The agent invokes the tool with the product's DID. This triggers a new Interaction that could, for example, automatically generate ad copy using another LLM call via **Spring AI**.

6. **Write-Back:** When the user completes the purchase, the FinalizePurchase transform is dispatched to the buy-able ContentTypeDataHandler. This handler gets a JCA connection and invokes the createOrder interaction on the backend ERP, completing the cycle.