

---

# Adaptive Architecture: Performance Through Evolutionary Design

Author: Pankaj Jain

Software Architect | Enterprise Systems & Performance Engineering

---

## Abstract

Software architectures are designed to deliver consistency, performance, and maintainability across enterprise applications. However, excessive rigidity—manifested through hardcoded assumptions, inflexible data structures, and generic access patterns—can stifle innovation and degrade performance. This whitepaper explores strategies for designing adaptive software systems that evolve in response to real-world query patterns and business needs.

---

## Table of Contents

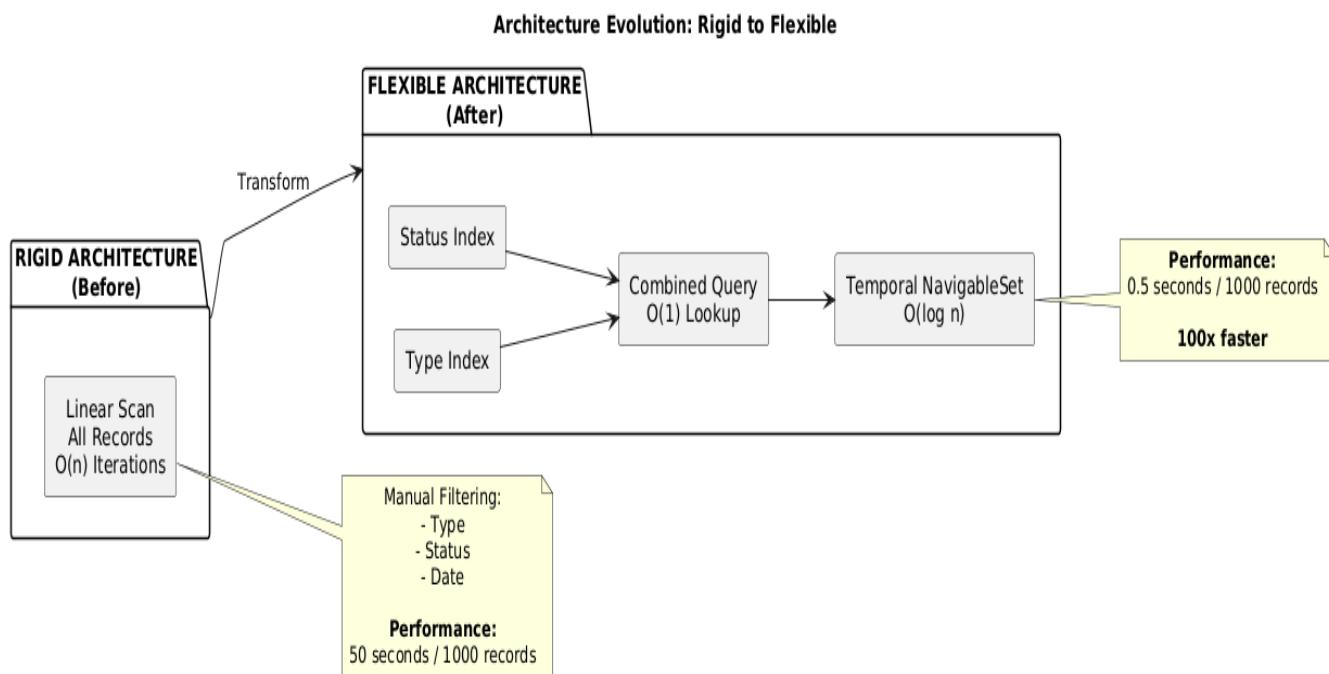
1. Introduction
2. Research & Profiling
3. Building Modular Indexes
4. Temporal Queries with NavigableSet
5. Handling Status Changes & Updates
6. Two-Tier Caching for Hot Paths
7. Realising the Benefits of Modular Architecture
8. Conclusion
9. References

---

## 1. Introduction

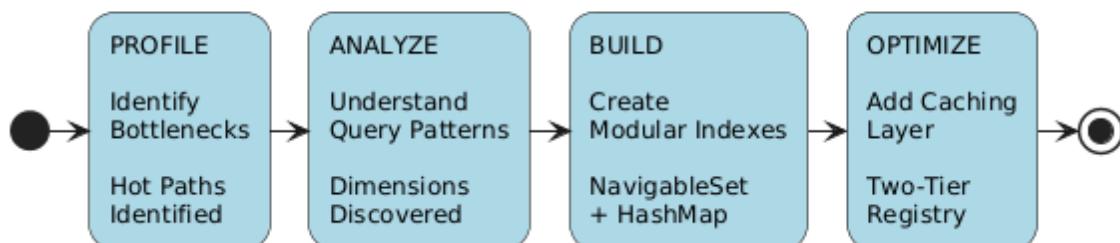
Software architectures typically begin with noble intentions: consistent data access, optimised performance, and clear separation of concerns. Yet, many falter when they become overly strict, morphing into rigid patterns that hinder growth and performance. Sustainable architectures must be adaptive—structured for consistency, yet open to new query patterns and evolving business requirements without necessitating complete rewrites.

## ARCHITECTURE EVOLUTION DIAGRAM



## KEY TRANSFORMATION STAGES

### Key Transformation Stages



## 2. Research & Profiling

“Premature optimisation is the root of all evil.” – Donald Knuth

“But profiling reveals the 20% of code consuming 80% of execution time.”

Performance profiling should be a continuous process, not a one-off task. As user behaviour and data volumes shift, so do access patterns. Without ongoing profiling, architectures risk becoming disconnected from actual usage. Begin with foundational profiling to understand user queries and frequently called operations, grounding data structures in real behaviour. Regular performance testing and clear feedback loops ensure the architecture remains performant, flexible, and aligned with evolving needs.

### Real-World Example:

In an insurance illustration system, profiling revealed a single method was called 10,000+ times per policy calculation, executing nested iterations through hundreds of records on every call. This linear  $O(n)$  scan accounted for 50 seconds of

processing per 1,000 policies in batch operations. Profiling identified key query dimensions: type-based lookups, status-based filtering, combined queries, and temporal queries. Without this data, optimisation efforts might have targeted the wrong components.

### 3. Building Modular Indexes

Data access structures should be modular, composable, and adaptable—building blocks rather than final, unchangeable implementations. Each index should optimise a specific query pattern. Composability allows teams to add new query dimensions without disrupting existing access patterns, avoiding rigid schemas and enabling ongoing optimisation.

#### Modular Index Architecture Example:

Java

```
// Master index of all records
private NavigableSet<RecordKey> allRecords = new TreeSet<>();

// Single-dimension indexes
private Map<Integer, NavigableSet<RecordKey>> recordsByType;
private Map<String, NavigableSet<RecordKey>> recordsByStatus;

// Multi-dimension index: Type + Status
private Map<Integer, Map<String, NavigableSet<RecordKey>>> recordsByTypeAndStatus;
Show more lines
```

#### Benefits:

- Single responsibility per index
- Composability for new dimensions
- $O(1)$  hash lookups replace  $O(n)$  scans
- Efficient date range queries with NavigableSet

#### Querying Example:

- **Before (Linear Scan –  $O(n)$ ):**

Iterates through all records, filtering by type, status, and date.

- **After (Indexed Lookup –  $O(1)$ ):**

Uses pre-filtered sets for rapid retrieval, reducing call time from 10–50 ms to 0.01–0.1 ms.

### 4. Temporal Queries with NavigableSet

NavigableSet enables efficient date-based queries, providing  $O(\log n)$  binary search and  $O(1)$  range view operations. Records are sorted by date and sequence, allowing clear separation of temporal filtering from other query dimensions.

#### Benefits for Date Range Queries:

- **Before:** Linear scan through all records (10–20 ms)
- **After:** NavigableSet headSet for instant view (0.01 ms, 400x faster)

### Use Case:

Combining type/status indexes with temporal filtering enables powerful multi-dimensional queries, dramatically faster than nested loops.

---

## 5. Handling Status Changes & Updates

Adaptive architectures must handle record updates without breaking indexes. When a record's status changes, all relevant indexes must be updated consistently using a “delete-then-insert” strategy. This ensures synchronisation with current data and prevents memory bloat by removing empty index structures.

---

## 6. Two-Tier Caching for Hot Paths

Frequently accessed computed values benefit from a two-tier caching strategy:

- **Tier 1:** Active Registry (pre-filter) instantly rejects invalid queries.
- **Tier 2:** Result Cache (memoisation) stores computed results for rapid retrieval.

This approach prevents wasted computation and delivers dramatic performance improvements—reducing total call time from 500 ms to 0.15 ms for 10,000 calls (3,333x faster).

---

## 7. Realising the Benefits of Modular Architecture

The goal is not just performance, but adaptability. Modular architectures enable systems to evolve with query patterns, data growth, and business requirements. Key outcomes include:

- Query performance improvements (100x–5,000x faster)
- Scalable batch processing
- Minimal memory overhead
- Maintainability for new query dimensions

### Architectural Principles Applied:

1. Profile first
2. Modular indexes
3. Composability
4. Temporal support
5. Two-tier caching
6. Update strategy

---

## 8. Conclusion

Effective software architecture demands more than predefined patterns—it requires continuous profiling, modular design, and evolutionary adaptation. By focusing on these principles, architectures remain adaptable, performant, and relevant.

---

Modular indexes, NavigableSet for temporal queries, two-tier caching, and robust update strategies ensure systems grow and evolve alongside their applications, enhancing performance, flexibility, and maintainability.

#### Performance Summary:

- **Before:** Linear scans, slow batch processing, poor scalability
- **After:** Indexed lookups, rapid batch processing, constant scalability, minimal memory overhead

---

#### 9. References

1. Donald Knuth, "Structured Programming with go to Statements" (1974)
2. Martin Fowler, "Refactoring: Improving the Design of Existing Code"
3. Joshua Bloch, "Effective Java" – Performance optimisation patterns
4. Robert C. Martin, "Clean Architecture" – Principles of software design

---