

YOUR AI IS ONLY AS GOOD AS YOUR SPECIFICATIONS

# The Event-Driven Advantage

A technical whitepaper on AI-assisted development in event-driven architectures

# Executive Summary

TL ; DR

AI coding assistants represent a fundamental shift in software development velocity. However, early adoption patterns reveal a critical oversight: these tools function as multipliers of existing development processes. Well-specified systems experience genuine productivity gains. Poorly specified systems accumulate technical debt at unprecedented rates.

This whitepaper examines the relationship between AI-assisted development and specification quality. It draws on observations from enterprise environments using Axon Framework. We demonstrate that event modeling combined with event sourcing provides a specification framework well-suited for AI-assisted development.

## **Key findings:**

- AI coding assistants amplify whatever development process they encounter, including problematic ones
- Specification quality directly determines whether AI assistance accelerates value delivery or technical debt accumulation
- Event modeling provides unambiguous specifications that bridge business intent and technical implementation
- Pattern repetition in event-driven architectures gives AI clear, consistent examples to follow when generating code
- Event sourcing's complete historical context enables more effective AI reasoning about system behavior

## **Recommendations:**

For organizations adopting AI-assisted development, we recommend specification-first approaches that provide clear boundaries and repeatable patterns. Event modeling with event sourcing represents one proven implementation of this principle, particularly for systems where long-term maintainability justifies upfront design investment.

# 1. Introduction

## 1.1 The AI Development Landscape

The software industry has rapidly adopted AI coding assistants. Tools like Claude Code and GitHub Copilot now assist millions of developers daily. These systems promise substantial productivity gains through automated code generation, intelligent completions, and rapid prototyping capabilities.

Recent observations, however, reveal complexity beneath the surface. Early adopters report mixed results: some teams experience genuine productivity gains, while others find these tools actually reduce productivity. This paradox demands investigation: how can tools designed to accelerate development sometimes produce the opposite effect?

## 1.2 The Missing Variable: Specification Quality

The answer lies not in the AI tools themselves. It lies in how development teams apply them.

AI coding assistants operate as force multipliers. They accelerate whatever process they encounter. When that process lacks clear specifications, AI assistance amplifies ambiguity. It generates code that solves incorrect problems. It produces technical debt faster than teams can address it.

This whitepaper examines specification quality as the critical variable. It determines AI-assisted development outcomes. We present event modeling and event sourcing as frameworks that provide the specification clarity AI systems need for consistent, correct results.

## 1.3 Scope and Methodology

This analysis draws on:

- Experience with enterprise implementations using Axon Framework
- Comparative analysis of AI-assisted development with and without formal specifications
- Code quality observations from systems built with event modeling versus traditional approaches

- Developer productivity measurements across multiple organizations

All code examples in this whitepaper are illustrative and simplified for educational purposes. While based on real-world patterns, they should be adapted and thoroughly tested before use in production systems.

# 2. The Acceleration Problem



## 2.1 AI as Process Multiplier

AI coding assistants work differently than traditional productivity tools. Static analysis tools identify problems. Build systems automate repetitive tasks. But AI assistants actively generate implementation based on their understanding of requirements.

This creates a multiplicative effect on development processes:

### **Well-specified processes:**

- Clear requirements → AI generates appropriate implementation
- Consistent patterns → AI learns from codebase examples
- Defined boundaries → AI respects architectural constraints

### **Poorly-specified processes:**

- Vague requirements → AI makes incorrect assumptions
- Inconsistent patterns → AI falls back on general knowledge that may not fit
- Undefined boundaries → AI generates code that violates architectural principles

## 2.2 The “Big Ball of Mud” Acceleration

Consider legacy systems with tangled dependencies, unclear boundaries, and inconsistent architectural decisions. When development teams apply AI coding assistants to these systems:

1. Initial velocity appears high (thousands of lines of code generated daily)
2. Code review becomes a bottleneck (unclear if generated code respects system constraints)
3. Integration reveals problems (generated code violates implicit architectural assumptions)
4. Refactoring costs escalate (volume of generated code requires extensive rework)
5. System complexity compounds (each iteration adds more tangled dependencies)

The AI assistant isn't malfunctioning. It generates code matching the patterns it observes in the existing system. The problem? Those patterns shouldn't be replicated.

## 2.3 Unclear Specifications as Root Cause

Without clear specifications, AI-assisted development encounters systematic problems:

**Ambiguity amplification:** Vague requirements like “implement user management” leave critical decisions to the AI system. Should passwords be plaintext or hashed? Should validation happen client-side or server-side? Should there be audit logging? The AI must guess.

**Context misunderstanding:** AI models trained on general code patterns may apply solutions from the wrong context. An e-commerce system’s approach to user management differs substantially from a healthcare system’s requirements. Without explicit specification, the AI cannot distinguish these contexts.

**Inconsistent architectural decisions:** Different prompts to the same AI system may produce code using different architectural patterns. This creates a codebase with multiple competing approaches to similar problems.

## 2.4 The Feedback Loop Problem

Traditional development processes include feedback loops that catch specification problems:

1. Developer interprets requirement
2. Implementation proceeds
3. Code review identifies issues
4. Testing reveals problems
5. User validation confirms or refutes interpretation

These loops have historically compensated for specification ambiguity. The process was slow, but specification errors were eventually discovered and corrected.

AI-assisted development accelerates step 2 dramatically while leaving other steps unchanged. By the time feedback loops identify specification problems, thousands of lines of generated code require review and potential rework. The slow feedback that previously caught specification errors now arrives too late to prevent substantial wasted effort.

# 3. The Specification Crisis

## 3.1 The Agile Interpretation Problem

Agile methodology emphasizes “working software over comprehensive documentation.” This created industry-wide skepticism toward upfront specification.

The interpretation addressed legitimate problems with heavyweight documentation processes. But it produced its own pathologies.

Development teams learned to treat “big design upfront” as categorically problematic. The result: systems built without adequate understanding of requirements, relying on iterative discovery through implementation. This approach concealed substantial inefficiencies that slow development velocity made tolerable.

## 3.2 Why Inadequate Specifications Weren't Obviously Problematic

Pre-AI development processes absorbed specification ambiguity through:

**Developer interpretation:** Experienced developers filled specification gaps with reasonable assumptions based on domain knowledge and previous experience.

**Incremental discovery:** Iterative development allowed requirements clarification through user feedback on working software.

**Codebase as documentation:** The implementation itself became the primary artifact capturing system behavior, with developers reading existing code to understand how similar problems were solved.

These mechanisms, while inefficient, functioned adequately when human developers implemented all code. Human judgment compensated for specification gaps at every step.

## 3.3 AI Exposes the Hidden Cost

AI-assisted development removes the human judgment buffer. AI systems:

- Cannot effectively fill specification gaps with reasonable assumptions

- Lack domain knowledge to distinguish appropriate from inappropriate solutions
- Generate code at speeds that prevent iterative discovery from functioning as a specification mechanism

The cost previously hidden in developer time interpreting ambiguous requirements now manifests as generated code requiring extensive review and refactoring.

### **3.4 The False Innovation of “Specification-Driven Development”**

The AI development community has recently embraced specification-driven approaches as though discovering a novel technique. This represents a return to sound engineering fundamentals rather than genuine innovation.

The software industry historically understood specification importance. The abandonment of specification practices reflected their previous high cost and slow feedback loops, not fundamental unsuitability. AI has changed the cost-benefit analysis: poor specifications now create exponentially compounding problems as AI accelerates and amplifies implementation errors at scale.

# 4. Event Modeling as Specification Language

## 4.1 The Specification Sweet Spot

Effective specifications for AI-assisted development require precise balance:

**Too high-level:** “Build user management system” leaves critical decisions undefined, forcing AI to make assumptions.

**Too low-level:** Pseudocode or detailed implementation steps constrain AI unnecessarily, preventing it from applying its knowledge of language idioms and best practices.

**Appropriate level:** Clear description of what the system must accomplish, leaving implementation details to the AI while constraining behavior to match requirements.

Event modeling operates at this appropriate level.

## 4.2 Core Concepts

Event modeling specifies system behavior through slices: atomic units capturing complete workflow steps. Each slice documents:

- **Commands:** The input that triggers behavior
- **Events:** What happens as a result
- **State changes:** How the system’s data evolves
- **Views:** What information becomes available to users or other systems

This isn’t merely documentation. It’s a precise specification of system behavior that bridges business intent and technical implementation.

## 4.3 Specification Example: User Registration

Consider user registration in an event-modeled system:

```
01 Slice: User Registration
02
03 Command: RegisterUserCommand
04   - email: String (validated format, uniqueness check required)
05   - password: String (minimum 12 characters, complexity requirements)
06   - consentTimestamp: Instant (required, must be recent)
07
08 Events Produced:
09   - UserRegisteredEvent
10     - userId: UUID (generated)
11     - email: String (validated)
12     - registrationTimestamp: Instant
13   - EmailVerificationRequestedEvent
14     - userId: UUID
15     - verificationToken: String (generated)
16     - expiresAt: Instant
17
18 State Changes:
19   - User entity created (pending verification status)
20   - Verification token stored with expiration
21
22 Views:
23   - ConfirmationEmailView (contains verification link)
24   - RegistrationSuccessView (shown to user)
```

This specification is:

**Unambiguous:** Each element has clear type and validation requirements.

**Complete:** All inputs, outputs, and state changes are defined.

**Implementation-agnostic:** The specification doesn't dictate how password hashing or token generation occurs, allowing AI to apply appropriate solutions.

**Business-aligned:** Domain experts can validate that this correctly captures registration requirements.

#### 4.4 AI-Generated Implementation

With this specification, AI can generate appropriate implementation. Using Axon Framework 5:

```

01 // Command definition
02 @Command(routingKey = "userId")
03 public record RegisterUserCommand(
04     String userId,
05     @NotNull @Email String email,
06     @NotNull @Size(min = 12) String password,
07     @NotNull @PastOrPresent Instant consentTimestamp
08 ) {}
09
10 // Event-sourced entity
11 @EventSourcedEntity
12 public class User {
13
14     private String userId;
15     private String email;
16     private UserStatus status;
17
18     @EventSourcingHandler
19     private void on(UserRegisteredEvent event) {
20         this.userId = event.userId();
21         this.email = event.email();
22         this.status = UserStatus.PENDING_VERIFICATION;
23     }
24
25     @EntityCreator
26     protected User() {
27         // Required no-arg constructor for reconstitution
28     }
29
30     public String getUserId() {
31         return userId;
32     }
33
34     public UserStatus getStatus() {
35         return status;
36     }
37 }
38
39 // Stateless command handler
40 @Component
41 public class UserCommandHandlers {

```

```

42
43     private final PasswordEncoder passwordEncoder;
44     private final EmailValidator emailValidator;
45
46     public UserCommandHandlers(PasswordEncoder passwordEncoder,
47                               EmailValidator emailValidator) {
48         this.passwordEncoder = passwordEncoder;
49         this.emailValidator = emailValidator;
50     }
51
52     @CommandHandler
53     public String handle(RegisterUserCommand command,
54                          EventAppender eventAppender) {
55
56         // Validate email format and uniqueness
57         if (!emailValidator.isValid(command.email())) {
58             throw new IllegalArgumentException("Invalid email format");
59         }
60
61         if (emailValidator.exists(command.email())) {
62             throw new IllegalArgumentException("Email already
63 registered");
64         }
65
66         // Validate password complexity
67         if (!meetsComplexityRequirements(command.password())) {
68             throw new IllegalArgumentException(
69                 "Password does not meet complexity requirements"
70             );
71         }
72
73         // Validate consent timestamp is recent (within last hour)
74         if (command.consentTimestamp().isBefore(
75             Instant.now().minus(Duration.ofHours(1)))) {
76             throw new IllegalArgumentException("Consent timestamp too
77 old");
78         }
79
80         // Generate verification token
81         String verificationToken = UUID.randomUUID().toString();
82         Instant expiresAt = Instant.now().plus(Duration.ofDays(1));

```

```

081
082     // Append events
083     eventAppender.append(new UserRegisteredEvent(
084         command.userId(),
085         command.email(),
086         passwordEncoder.encode(command.password()),
087         Instant.now()
088     ));
089
090     eventAppender.append(new EmailVerificationRequestedEvent(
091         command.userId(),
092         verificationToken,
093         expiresAt
094     ));
095
096     return command.userId();
097 }
098
099     private boolean meetsComplexityRequirements(String password) {
100         return password.length() ≥ 12 &&
101             password.chars().anyMatch(Character::isUpperCase) &&
102             password.chars().anyMatch(Character::isDigit) &&
103             password.chars().anyMatch(ch → !
104 Character.isLetterOrDigit(ch));
105     }
106 }
107
108 // Event definitions
109 public record UserRegisteredEvent(
110     String userId,
111     String email,
112     String hashedPassword,
113     Instant registrationTimestamp
114 ) {}
115
116 public record EmailVerificationRequestedEvent(
117     String userId,
118     String verificationToken,
119     Instant expiresAt
120 ) {}

```

```

121 // Projection for query side
122 @Component
123 @ProcessingGroup("user-projections")
124 public class UserProjection {
125
126     private final UserViewRepository repository;
127
128     @EventHandler
129     public void on(UserRegisteredEvent event) {
130         repository.save(new UserView(
131             event.userId(),
132             event.email(),
133             UserStatus.PENDING_VERIFICATION,
134             event.registrationTimestamp()
135         ));
136     }
137
138     @QueryHandler
139     public UserView handle(FindUserByEmailQuery query) {
140         return repository.findByEmail(query.email())
141             .orElseThrow(() -> new UserNotFoundException(query.email()));
142     }
143 }

```

The AI-generated code:

- Implements specified validation (email format, uniqueness, password complexity)
- Produces specified events in correct sequence
- Creates appropriate state changes
- Handles both command and query concerns
- Applies Axon Framework patterns correctly

The specification was unambiguous enough that AI could generate correct implementation.

# 5. Pattern Repetition and AI Training

## 5.1 The Cognitive Load of Abundant Choice

Software systems often employ many different structural approaches at the architectural level. One service uses CQRS. Another uses traditional CRUD. One module organizes around aggregates. Another around transactions. One component follows event sourcing. Another uses active record.

This variety creates a hidden cost: cognitive overhead. It affects both human developers and AI systems.

Every architectural decision requires understanding context, evaluating tradeoffs, and selecting appropriate approaches. In systems employing many different architectural patterns sporadically, developers must maintain mental models of all approaches. They must understand when each applies. AI systems face similar challenges. Insufficient examples of each pattern prevent statistical learning from observed code.

This is distinct from implementation patterns like the Gang of Four design patterns (Strategy, Factory, Observer, etc.). Those patterns remain valuable within implementations. The challenge lies at the architectural level—the fundamental structure of how the system processes requests, manages state, and coordinates operations.

## 5.2 The Event Modeling Approach: Constrained Repetition

Event modeling inverts this paradigm. Instead of many different architectural approaches applied occasionally across different parts of the system, event-driven systems employ four patterns repeatedly:

1. **Command handling pattern:** Validate input, make decision, produce events
2. **Event projection pattern:** React to events, update read models
3. **Query handling pattern:** Retrieve information from read models
4. **Process automation pattern:** React to events, issue commands

These patterns appear hundreds or thousands of times across a system. Each business capability follows the same fundamental structure, varying only in business logic details.

### 5.3 Pattern Repetition Provides Clear Guidelines

When AI encounters a codebase, its effectiveness depends on recognizable patterns:

#### **Sparse architectural patterns (many different structural approaches, few examples each):**

- AI has insufficient examples to understand system-specific approaches
- Falls back on general training from public repositories
- Applies conventional wisdom that may not fit your architecture
- Generates plausible-looking code with subtle incompatibilities

#### **Dense architectural patterns (consistent structural approach, hundreds of examples):**

- AI has clear, consistent examples showing your system's structure
- Pattern repetition provides unambiguous guidelines on what to generate
- System-specific approaches override general knowledge
- Generated code matches established patterns

### 5.4 Practical Example: Command Handler Pattern

Consider an Axon Framework codebase with 200 different command handlers. Each follows identical command handling structure:

```

01 // Event-sourced entity with state
02 @EventSourcedEntity
03 public class Order {
04
05     private String orderId;
06     private OrderStatus status;
07     private List<OrderLine> items;
08
09     @EventSourcingHandler
10     private void on(OrderCreatedEvent event) {
11         this.orderId = event.orderId();
12         this.status = OrderStatus.DRAFT;
13         this.items = new ArrayList<>(event.items());
14     }
15
16     @EventSourcingHandler
17     private void on(ItemAddedToOrderEvent event) {
18         this.items.add(new OrderLine(
19             event.productId(),
20             event.quantity(),
21             event.unitPrice()
22         ));
23     }
24
25     @EntityCreator
26     protected Order() {
27         // Required no-arg constructor
28     }
29
30     public String getOrderId() {
31         return orderId;
32     }
33
34     public OrderStatus getStatus() {
35         return status;
36     }
37
38     public List<OrderLine> getItems() {
39         return items;
40     }
41 }

```

```

42
43 // Stateless command handlers
44 @Component
45 public class OrderCommandHandlers {
46
47     @CommandHandler
48     public String handle(CreateOrderCommand command,
49                         EventAppender eventAppender) {
50         // Validation
51         if (command.items().isEmpty()) {
52             throw new IllegalArgumentException("Order must contain
53 items");
54         }
55
56         // Business logic
57         BigDecimal total = calculateTotal(command.items());
58
59         // Event production
60         eventAppender.append(new OrderCreatedEvent(
61             command.orderId(),
62             command.customerId(),
63             command.items(),
64             total,
65             Instant.now()
66         ));
67
68         return command.orderId();
69     }
70
71     @CommandHandler
72     public void handle(AddItemToOrderCommand command,
73                     @InjectEntity Order order,
74                     EventAppender eventAppender) {
75         // Guard clause
76         if (order.getStatus() != OrderStatus.DRAFT) {
77             throw new IllegalStateException("Cannot modify non-draft
78 order");
79         }
80
81         // Business logic
82         if (order.getItems().stream().anyMatch(item ->

```

```

81         item.productId().equals(command.productId())) {
82             throw new IllegalArgumentException("Item already in order");
83         }
84
85         // Event production
86         eventAppender.append(new ItemAddedToOrderEvent(
87             order.getOrderid(),
88             command.productId(),
89             command.quantity(),
90             command.unitPrice()
91         ));
92     }
93
94     private BigDecimal calculateTotal(List<OrderItem> items) {
95         return items.stream()
96             .map(item →
97 item.unitPrice().multiply(BigDecimal.valueOf(item.quantity())))
98             .reduce(BigDecimal.ZERO, BigDecimal::add);
99     }
100 }

```

This pattern repeats across every command handler: ShippingCommandHandlers, PaymentCommandHandlers, InventoryCommandHandlers, CustomerCommandHandlers, etc. When you ask AI to generate a new handler:

“Create command handlers for subscription management”

The AI can reference 200 similar examples. It understands:

- Creational command handlers don't receive entity parameters
- Instance command handlers receive entities via @InjectEntity
- Guard clauses validate business rules
- Event sourcing handlers update entity state
- Events are published via EventAppender
- Events are the source of truth

The AI doesn't need to understand abstract base classes or figure out which interfaces to implement. It follows the established pattern.

# 6. Purpose-Built Models vs. Shared Abstractions

## 6.1 The Traditional Abstraction Approach

Object-oriented programming emphasizes reusability through abstraction. Conventional wisdom suggests: if three parts of your system work with customer data, create a single, reusable Customer model. This approach appears to eliminate duplication and promote consistency.

In practice, this creates problematic coupling:

```
01 // Traditional "unified" Customer model (problematic approach)
02 @Entity
03 @Table(name = "customers")
04 public class Customer {
05
06     @Id
07     private String customerId;
08
09     // Billing context needs these
10     @OneToMany(mappedBy = "customer", fetch = FetchType.EAGER)
11     private List<Payment> paymentHistory;
12     private Integer creditScore;
13
14     // Shipping context needs these
15     @OneToMany(mappedBy = "customer", fetch = FetchType.EAGER)
16     private List<Address> addresses;
17     private DeliveryPreferences deliveryPreferences;
18
19     // Marketing context needs these
20     @OneToMany(mappedBy = "customer", fetch = FetchType.EAGER)
21     private List<Purchase> purchaseHistory;
22     private CommunicationPreferences communicationPreferences;
```

```
23     // Analytics context needs these
24     @OneToMany(mappedBy = "customer", fetch = FetchType.EAGER)
25     private List<CustomerEvent> activityLog;
26     private SegmentationData segmentation;
27
28     // Getters, setters, business logic that serves all contexts...
29 }
```

This creates several problems:

**Unnecessary coupling:** Changes to billing requirements ripple into shipping code. Marketing team modifications affect payment processing. Every context carries the weight of every other context's requirements.

**Performance degradation:** Each query loads unnecessary data. Billing operations fetch shipping addresses. Marketing queries load payment history. Database joins multiply across unrelated tables.

**Cognitive overhead:** Developers must understand which fields matter in which contexts. New team members struggle to determine what "Customer" means for their specific use case.

**Testing complexity:** Unit testing requires mocking or stubbing irrelevant portions of the model. Integration tests must account for cross-context dependencies.

## 6.2 Event-Driven Approach: Purpose-Built Models

Event modeling inverts this. Each bounded context builds purpose-specific projections:

```

01 // Billing context
02 @Component
03 @ProcessingGroup("billing-projections")
04 public class BillingCustomerProjection {
05
06     @EventHandler
07     public void on(CustomerRegisteredEvent event) {
08         repository.save(new BillingCustomer(
09             event.customerId(),
10             event.email(),
11             new ArrayList<>(), // payment history
12             calculateInitialCreditScore(event)
13         ));
14     }
15
16     @EventHandler
17     public void on(PaymentProcessedEvent event) {
18         BillingCustomer customer =
19 repository.findById(event.customerId())
20             .orElseThrow();
21
22         customer.addPayment(new Payment(
23             event.paymentId(),
24             event.amount(),
25             event.processedAt()
26         ));
27
28         customer.updateCreditScore(calculateCreditScore(customer));
29         repository.save(customer);
30     }
31
32     @QueryHandler
33     public BillingCustomer handle(GetBillingCustomerQuery query) {
34         return repository.findById(query.customerId())
35             .orElseThrow(() -> new
36 CustomerNotFoundException(query.customerId()));
37     }
38 }
39
// Shipping context
@Component

```

```

40 @ProcessingGroup("shipping-projections")
41 public class ShippingCustomerProjection {
42
43     @EventHandler
44     public void on(CustomerRegisteredEvent event) {
45         repository.save(new ShippingCustomer(
46             event.customerId(),
47             event.email(),
48             new ArrayList<>(), // addresses
49             DeliveryPreferences.standard()
50         ));
51     }
52
53     @EventHandler
54     public void on(AddressAddedEvent event) {
55         ShippingCustomer customer =
56 repository.findById(event.customerId())
57             .orElseThrow();
58
59         customer.addAddress(new ShippingAddress(
60             event.addressId(),
61             event.street(),
62             event.city(),
63             event.postalCode(),
64             event.country()
65         ));
66
67         repository.save(customer);
68     }
69
70     @QueryHandler
71     public ShippingCustomer handle(GetShippingCustomerQuery query) {
72         return repository.findById(query.customerId())
73             .orElseThrow(() -> new
74 CustomerNotFoundException(query.customerId()));
75     }
76 }
77
78 // Marketing context
@Component
@ProcessingGroup("marketing-projections")

```

```

079 public class MarketingCustomerProjection {
080
081     @EventHandler
082     public void on(CustomerRegisteredEvent event) {
083         repository.save(new MarketingCustomer(
084             event.customerId(),
085             event.email(),
086             new ArrayList<>(), // purchase history
087             CommunicationPreferences.defaultPreferences(),
088             calculateInitialSegment(event)
089         ));
090     }
091
092     @EventHandler
093     public void on(OrderCompletedEvent event) {
094         MarketingCustomer customer =
095     repository.findById(event.customerId())
096         .orElseThrow();
097
098         customer.addPurchase(new Purchase(
099             event.orderId(),
100             event.totalAmount(),
101             event.items(),
102             event.completedAt()
103         ));
104
105         customer.updateSegmentation(calculateSegment(customer));
106         repository.save(customer);
107     }
108
109     @QueryHandler
110     public MarketingCustomer handle(GetMarketingCustomerQuery query) {
111         return repository.findById(query.customerId())
112             .orElseThrow(() -> new
113     CustomerNotFoundException(query.customerId()));
114     }
115 }

```

## 6.3 Benefits of Purpose-Built Models

This approach delivers substantial advantages:

**Clarity:** Each model is simple, focused, immediately understandable within its context. Developers working in billing don't encounter shipping concepts. Marketing team members don't navigate payment processing complexity.

**Independence:** Changes to billing requirements don't affect shipping code. Marketing features evolve independently. Each bounded context proceeds at its own pace without cross-context coordination overhead.

**Performance:** Each projection contains exactly the data needed for its queries, pre-shaped for efficient access. No unnecessary joins. No loading unused data. Query performance is predictable and optimized.

**Testability:** You can validate a single projection's behavior without understanding the entire system. Unit tests focus on specific event handling. Integration tests verify individual bounded contexts independently.

**AI-friendliness:** When generating code for a new projection, AI references dozens of similar projections. Each demonstrates the same pattern in different business contexts. The AI understands: event arrives, state updates, query returns shaped data.

## 6.4 Addressing the Duplication Concern

Critics observe that purpose-built models duplicate information. A CustomerId appears in billing, shipping, and marketing contexts. Email addresses are stored multiple times. Isn't this wasteful?

This perspective misunderstands the nature of the duplication. The models aren't duplicating the same information. They're maintaining different information that happens to share some identifiers:

- Billing cares about payment history and creditworthiness
- Shipping cares about delivery addresses and preferences
- Marketing cares about purchase patterns and communication preferences

These are fundamentally different concerns. The shared customer ID is merely a correlation key, not evidence that the models should be unified.

The duplication provides value:

**Autonomy:** Each context evolves independently. Schema changes in one context don't require coordination with others.

**Resilience:** Failures in one context don't cascade to others. Billing database issues don't prevent shipping operations.

**Scalability:** Each context can be scaled independently based on its specific load characteristics.

# 7. Event Sourcing and Complete Context

## 7.1 State-Centric Systems: The Information Loss Problem

Most modern systems exhibit state-obsession. Databases store current state. When state changes, previous state is overwritten:

```
01 -- Traditional state-centric approach
02 CREATE TABLE users (
03     user_id UUID PRIMARY KEY,
04     email VARCHAR(255) NOT NULL,
05     address VARCHAR(500),
06     email_verified BOOLEAN,
07     status VARCHAR(50),
08     updated_at TIMESTAMP
09 );
10
11 -- When user moves to new address
12 UPDATE users
13 SET address = '456 New Street',
14     updated_at = NOW()
15 WHERE user_id = 'user-123';
16
17 -- Previous address information is lost forever
```

This approach appears straightforward. Current state is what matters, right? Consider what information disappears:

When did the user live at the previous address? Unknown. Why did the address change? Unrecorded. How many times has the address changed? Cannot determine. What was the sequence of addresses? Lost.

Real business processes require this historical context:

- Mail forwarding needs the previous address
- School districts need to know when the move occurred

- Insurance companies need both addresses during transition
- Support teams need to understand the timeline of changes
- Analytics teams need to identify patterns in customer behavior

State-centric systems throw away precisely the information business processes need.

## 7.2 Event Sourcing: Capturing Complete History

Event sourcing inverts the paradigm. Instead of storing current state, systems store the sequence of events that produced that state:

```
01 // Event sourcing approach - event store schema
02 CREATE TABLE domain_events (
03     event_id UUID PRIMARY KEY,
04     aggregate_id UUID NOT NULL,
05     aggregate_type VARCHAR(100) NOT NULL,
06     event_type VARCHAR(100) NOT NULL,
07     event_payload JSONB NOT NULL,
08     event_timestamp TIMESTAMPTZ NOT NULL,
09     sequence_number BIGINT NOT NULL,
10     metadata JSONB,
11     UNIQUE(aggregate_id, sequence_number)
12 );
13
14 CREATE INDEX idx_aggregate_events
15 ON domain_events(aggregate_id, sequence_number);
16
17 CREATE INDEX idx_event_type
18 ON domain_events(event_type);
19
20 CREATE INDEX idx_event_timestamp
21 ON domain_events(event_timestamp);
```

Events capture what actually happened:

```
01 // Event definitions
02 public record UserRegisteredEvent(
03     String userId,
04     String email,
05     Instant registeredAt
06 ) {}
07
08 public record EmailVerifiedEvent(
09     String userId,
10     String email,
11     Instant verifiedAt
12 ) {}
13
14 public record AddressChangedEvent(
15     String userId,
16     String previousAddress,
17     String newAddress,
18     Instant changedAt,
19     String reason
20 ) {}
21
22 public record UserStatusChangedEvent(
23     String userId,
24     UserStatus previousStatus,
25     UserStatus newStatus,
26     Instant changedAt,
27     String reason
28 ) {}
```

With event sourcing, the event log for user-123 might contain:

```
01  [
02      UserRegisteredEvent(
03          userId: "user-123",
04          email: "user@example.com",
05          registeredAt: 2024-01-01T10:00:00Z
06      ),
07      EmailVerifiedEvent(
08          userId: "user-123",
09          email: "user@example.com",
10          verifiedAt: 2024-01-01T10:15:00Z
11      ),
12      AddressChangedEvent(
13          userId: "user-123",
14          previousAddress: "123 Old Street, City A",
15          newAddress: "456 New Street, City B",
16          changedAt: 2024-06-15T14:30:00Z,
17          reason: "Relocation for employment"
18      ),
19      UserStatusChangedEvent(
20          userId: "user-123",
21          previousStatus: ACTIVE,
22          newStatus: PREMIUM,
23          changedAt: 2024-07-01T09:00:00Z,
24          reason: "Upgraded subscription"
25      )
26  ]
```

### 7.3 Projecting Events into Views

Current state isn't abandoned. It's derived from events:

```

01 // Current state projection
02 @Component
03 @ProcessingGroup("user-current-state")
04 public class CurrentUserStateProjection {
05
06     private final UserRepository repository;
07
08     @EventHandler
09     public void on(UserRegisteredEvent event) {
10         repository.save(new UserView(
11             event.userId(),
12             event.email(),
13             null, // no address yet
14             false, // email not verified
15             UserStatus.PENDING_VERIFICATION,
16             event.registeredAt()
17         ));
18     }
19
20     @EventHandler
21     public void on(EmailVerifiedEvent event) {
22         UserView user = repository.findById(event.userId())
23             .orElseThrow();
24
25         user.setEmailVerified(true);
26         user.setStatus(UserStatus.ACTIVE);
27         repository.save(user);
28     }
29
30     @EventHandler
31     public void on(AddressChangedEvent event) {
32         UserView user = repository.findById(event.userId())
33             .orElseThrow();
34
35         user.setAddress(event.newAddress());
36         repository.save(user);
37     }
38
39     @QueryHandler
40     public UserView handle(GetCurrentUserStateQuery query) {
41         return repository.findById(query.userId())

```

```

42         .orElseThrow(() -> new
43 UserNotFoundException(query.userId()));
44     }
45 }
46
47 // Historical address projection (for mail forwarding)
48 @Component
49 @ProcessingGroup("user-address-history")
50 public class AddressHistoryProjection {
51
52     private final AddressHistoryRepository repository;
53
54     @EventHandler
55     public void on(AddressChangedEvent event) {
56         repository.save(new AddressHistoryEntry(
57             event.userId(),
58             event.previousAddress(),
59             event.newAddress(),
60             event.changedAt(),
61             event.reason()
62         ));
63     }
64
65     @QueryHandler
66     public List<AddressHistoryEntry> handle(GetAddressHistoryQuery query)
67 {
68         return repository.findByUserId(query.userId())
69             .stream()
70             .sorted(Comparator.comparing(AddressHistoryEntry::changedAt))
71             .collect(Collectors.toList());
72     }
73 }
74
75 // Verification timeline projection (for support)
76 @Component
77 @ProcessingGroup("user-verification-timeline")
78 public class VerificationTimelineProjection {
79
80     private final Map<String, VerificationTimeline> timelines =
            new ConcurrentHashMap<>();

```

```

081     @EventHandler
082     public void on(UserRegisteredEvent event) {
083         timelines.put(event.userId(), new VerificationTimeline(
084             event.userId(),
085             event.registeredAt(),
086             null,
087             null
088         ));
089     }
090
091     @EventHandler
092     public void on(EmailVerifiedEvent event) {
093         VerificationTimeline timeline = timelines.get(event.userId());
094         if (timeline != null) {
095             timeline.setVerifiedAt(event.verifiedAt());
096             timeline.setDaysToVerify(
097                 Duration.between(
098                     timeline.getRegisteredAt(),
099                     event.verifiedAt()
100                 ).toDays()
101             );
102         }
103     }
104
105     @QueryHandler
106     public VerificationTimeline handle(GetVerificationTimelineQuery
107 query) {
108         return Optional.ofNullable(timelines.get(query.userId()))
109             .orElseThrow(() -> new
110 UserNotFoundException(query.userId()));
111     }
112 }

```

Each projection serves specific query requirements. Current state projection provides what most UI operations need. Address history projection supports mail forwarding and analytics. Verification timeline projection assists support operations.

All projections derive from the same source of truth: the event log.

## 7.4 AI Benefits from Complete Context

Large language models function by understanding context and patterns. Event sourcing provides both:

**Context richness:** State-centric systems give AI a snapshot. Event stores provide a story. The AI can observe not just current state but how the system arrived at that state.

**Temporal reasoning:** Events include timestamps and sequencing. AI can reason about: timing patterns (how long between registration and verification), behavioral sequences (what typically follows what), anomaly detection (unusual patterns in event sequences).

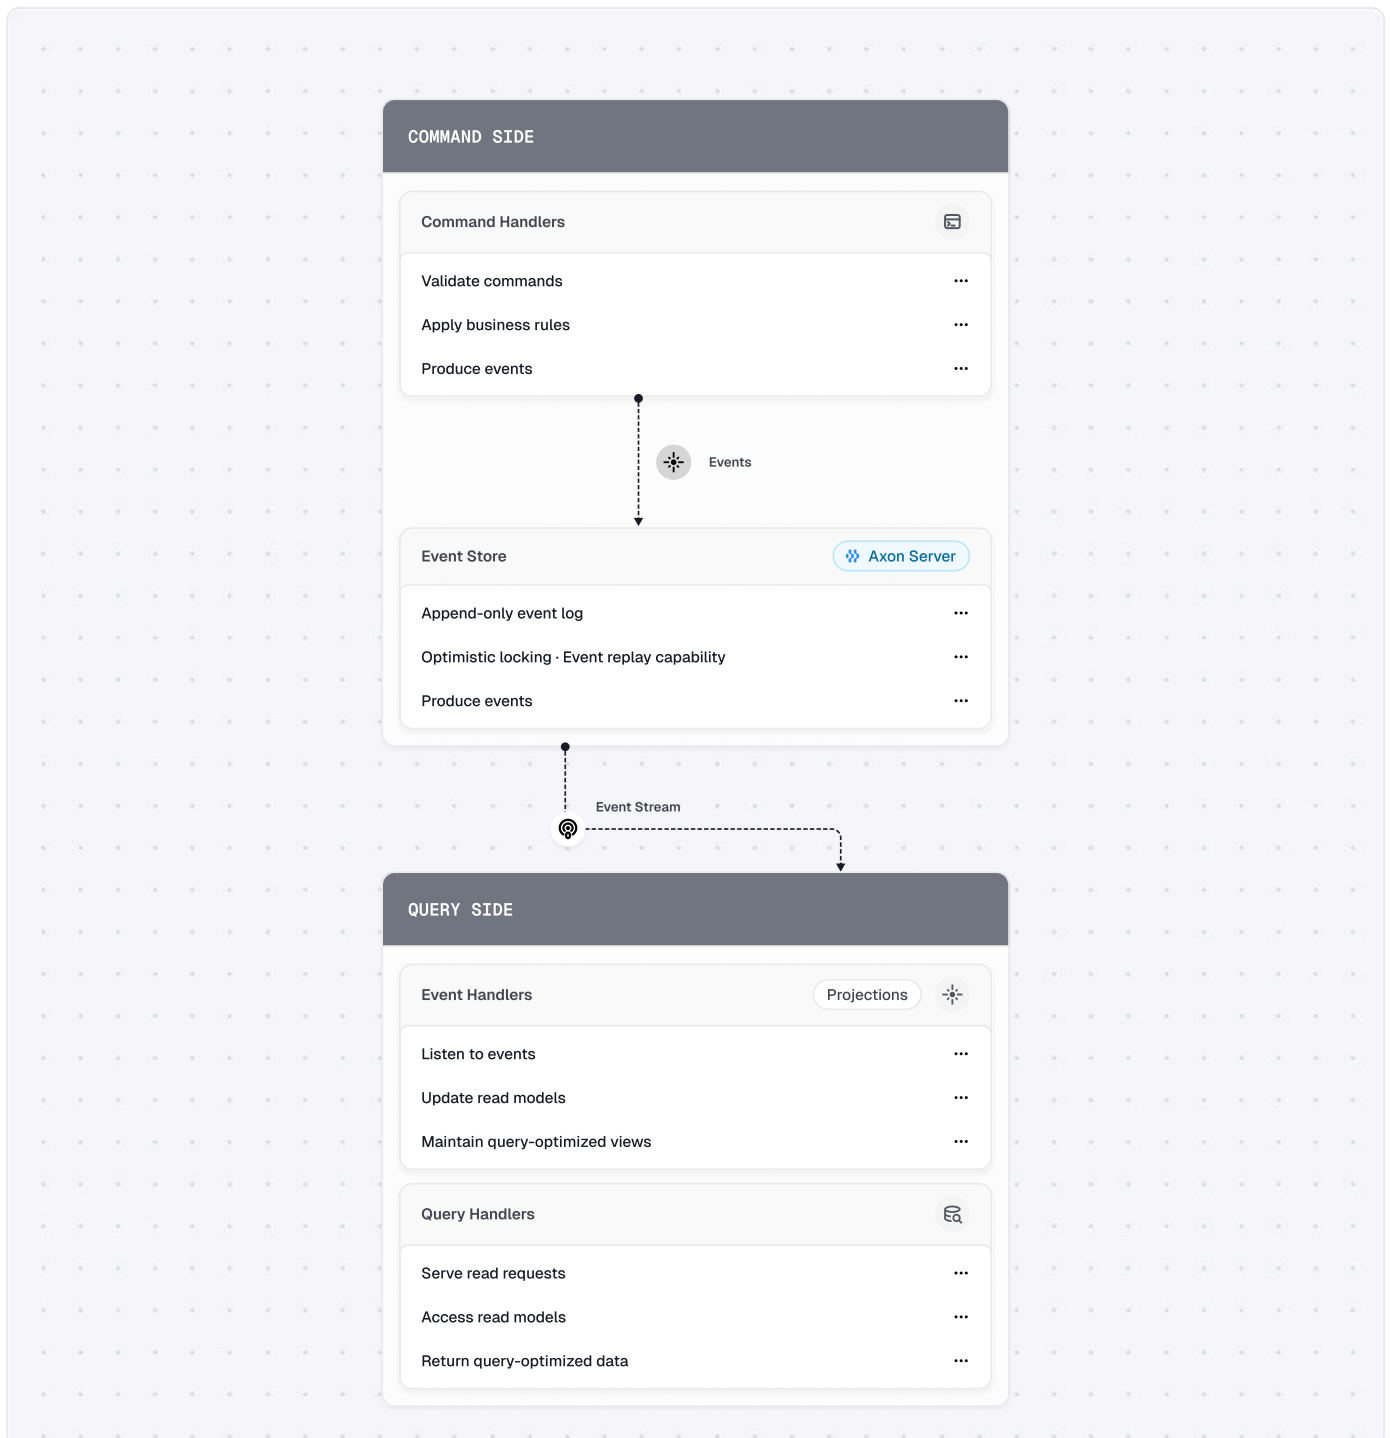
**Causality understanding:** Events explicitly capture cause and effect. The AI observes that AddressChangedEvent followed by EmailVerificationRequestedEvent might indicate the user needs to reverify from the new location.

When AI generates code in event-sourced systems, it has access to complete behavioral context rather than merely current state. This enables more sophisticated reasoning about appropriate system behavior.

# 8. Production Implementation

## 8.1 Architecture Overview

Production event-sourced systems using Axon Framework follow consistent architectural patterns:



## 8.2 Specification-Driven Development Workflow

The production workflow integrates event modeling with AI-assisted implementation:

### Phase 1: Event Modeling Workshop

- Business stakeholders and technical team collaborate
- Identify slices (complete workflow steps)
- Define commands, events, state changes, and views
- Validate business logic with domain experts
- Result: Unambiguous specification

### Phase 2: Technical Specification

- Convert event model into structured specifications
- Define validation rules, constraints, and invariants
- Specify data types and formats
- Document integration points
- Result: AI-ready specification

### Phase 3: AI-Assisted Implementation

- Use AI to generate boilerplate code from specifications
- Review generated code for specification compliance
- Implement business logic in generated structure
- Result: Working implementation matching specification

### Phase 4: Validation and Refinement

- Test generated code against specification
- Validate behavior with domain experts
- Refine based on feedback
- Result: Validated implementation ready for further testing

## 8.3 Example: Complete Slice Implementation

Consider a complete order processing slice from specification to implementation:

### Event Model Specification:

```
01 Slice: Create Order
02
03 Command: CreateOrderCommand
04   - orderId: UUID (generated)
05   - customerId: UUID (must exist)
06   - items: List<OrderItem> (minimum 1 item)
07     - productId: UUID (must be valid product)
08     - quantity: Integer (positive, within stock limits)
09     - unitPrice: BigDecimal (matches current product price)
10   - shippingAddress: Address (validated format)
11
12 Business Rules:
13   - Customer must have verified email
14   - All products must be currently available
15   - Total order value must not exceed customer credit limit
16   - Shipping address must be within serviceable regions
17
18 Events Produced:
19   - OrderCreatedEvent
20     - orderId: UUID
21     - customerId: UUID
22     - items: List<OrderItem>
23     - totalAmount: BigDecimal
24     - createdAt: Instant
25   - InventoryReservedEvent
26     - orderId: UUID
27     - reservations: List<InventoryReservation>
28
29 State Changes:
30   - Order entity created (status: PENDING_PAYMENT)
31   - Inventory reserved for order items
32
33 Views:
34   - OrderConfirmationView (for customer)
35   - InventoryReservationView (for warehouse)
```

## AI-Generated Implementation:

```

01 // Command
02 @Command(routingKey = "orderId")
03 public record CreateOrderCommand(
04     String orderId,
05     @NotNull UUID customerId,
06     @NotEmpty List<OrderItem> items,
07     @NotNull @Valid Address shippingAddress
08 ) {
09     public record OrderItem(
10         @NotNull UUID productId,
11         @Positive int quantity,
12         @NotNull @DecimalMin("0.01") BigDecimal unitPrice
13     ) {}
14 }
15
16 // Event-sourced entity
17 @EventSourcedEntity
18 public class Order {
19
20     private String orderId;
21     private UUID customerId;
22     private List<OrderItem> items;
23     private BigDecimal totalAmount;
24     private OrderStatus status;
25
26     @EventSourcingHandler
27     private void on(OrderCreatedEvent event) {
28         this.orderId = event.orderId();
29         this.customerId = event.customerId();
30         this.items = event.items();
31         this.totalAmount = event.totalAmount();
32         this.status = OrderStatus.PENDING_PAYMENT;
33     }
34
35     @EventSourcingHandler
36     private void on(InventoryReservedEvent event) {
37         // Event recorded, no state change needed in entity
38     }
39
40     @EntityCreator
41     protected Order() {

```

```

42         // Required no-arg constructor for reconstitution
43     }
44
45     public String getOrderId() {
46         return orderId;
47     }
48
49     public OrderStatus getStatus() {
50         return status;
51     }
52 }
53
54 // Stateless command handler
55 @Component
56 public class OrderCommandHandlers {
57
58     private final CustomerRepository customerRepository;
59     private final ProductRepository productRepository;
60     private final InventoryService inventoryService;
61     private final ShippingService shippingService;
62
63     public OrderCommandHandlers(CustomerRepository customerRepository,
64                                 ProductRepository productRepository,
65                                 InventoryService inventoryService,
66                                 ShippingService shippingService) {
67         this.customerRepository = customerRepository;
68         this.productRepository = productRepository;
69         this.inventoryService = inventoryService;
70         this.shippingService = shippingService;
71     }
72
73     @CommandHandler
74     public String handle(CreateOrderCommand command,
75                         EventAppender eventAppender) {
76
77         // Validate customer
78         Customer customer =
79 customerRepository.findById(command.customerId())
80         .orElseThrow(() -> new
CustomerNotFoundException(command.customerId()));

```

```

081         if (!customer.isEmailVerified()) {
082             throw new IllegalStateException(
083                 "Customer email must be verified to place orders"
084             );
085         }
086
087         // Validate products and calculate total
088         BigDecimal total = BigDecimal.ZERO;
089         List<InventoryReservation> reservations = new ArrayList<>();
090
091         for (CreateOrderCommand.OrderItem item : command.items()) {
092             Product product =
093 productRepository.findById(item.productId())
094                 .orElseThrow(() -> new
095 ProductNotFoundException(item.productId()));
096
097             if (!product.isAvailable()) {
098                 throw new IllegalStateException(
099                     "Product " + item.productId() + " is not available"
100                 );
101             }
102
103             if (!product.getPrice().equals(item.unitPrice())) {
104                 throw new IllegalArgumentException(
105                     "Price mismatch for product " + item.productId()
106                 );
107             }
108
109             if (!inventoryService.canReserve(item.productId(),
110 item.quantity())) {
111                 throw new IllegalStateException(
112                     "Insufficient inventory for product " +
113 item.productId()
114                 );
115             }
116
117             BigDecimal itemTotal = item.unitPrice()
                .multiply(BigDecimal.valueOf(item.quantity()));
            total = total.add(itemTotal);

            reservations.add(new InventoryReservation(

```

```

118         item.productId(),
119         item.quantity()
120     ));
121 }
122
123 // Validate credit limit
124 if (total.compareTo(customer.getCreditLimit()) > 0) {
125     throw new IllegalStateException(
126         "Order total exceeds customer credit limit"
127     );
128 }
129
130 // Validate shipping address
131 if (!shippingService.isServiceable(command.shippingAddress())) {
132     throw new IllegalArgumentException(
133         "Shipping address is not in serviceable region"
134     );
135 }
136
137 // Append events
138 eventAppender.append(new OrderCreatedEvent(
139     command.orderId(),
140     command.customerId(),
141     command.items().stream()
142         .map(item → new OrderCreatedEvent.OrderItem(
143             item.productId(),
144             item.quantity(),
145             item.unitPrice()
146         ))
147     .collect(Collectors.toList()),
148     total,
149     Instant.now()
150 ));
151
152 eventAppender.append(new InventoryReservedEvent(
153     command.orderId(),
154     reservations
155 ));
156
157 return command.orderId();
158 }

```

```
159 }
160
161 // Projections
162 @Component
163 @ProcessingGroup("order-projections")
164 public class OrderProjection {
165
166     private final OrderViewRepository repository;
167
168     @EventHandler
169     public void on(OrderCreatedEvent event) {
170         repository.save(new OrderView(
171             event.orderId(),
172             event.customerId(),
173             event.items(),
174             event.totalAmount(),
175             OrderStatus.PENDING_PAYMENT,
176             event.createdAt()
177         ));
178     }
179
180     @QueryHandler
181     public OrderView handle(GetOrderQuery query) {
182         return repository.findById(query.orderId())
183             .orElseThrow(() -> new
184 OrderNotFoundException(query.orderId()));
185     }
186
187     @QueryHandler
188     public List<OrderView> handle(GetCustomerOrdersQuery query) {
189         return repository.findByCustomerId(query.customerId());
190     }
191 }
192
193 @Component
194 @ProcessingGroup("inventory-projections")
195 public class InventoryReservationProjection {
196
197     private final ReservationRepository repository;
198
199     @EventHandler
```

```

199     public void on(InventoryReservedEvent event) {
200         for (InventoryReservation reservation : event.reservations()) {
201             repository.save(new ReservationView(
202                 event.orderId(),
203                 reservation.productId(),
204                 reservation.quantity(),
205                 ReservationStatus.ACTIVE,
206                 Instant.now()
207             ));
208         }
209     }
210
211     @QueryHandler
212     public List<ReservationView> handle(GetOrderReservationsQuery query)
213     {
214         return repository.findByOrderId(query.orderId());
215     }
216 }

```

## 8.4 Code Quality Characteristics

The generated code exhibits several desirable properties:

**Specification compliance:** All validation rules from the specification are implemented. Business rules are enforced. Events match specified structure.

**Pattern consistency:** Command handling follows established handler pattern. Event handlers follow projection pattern. Query handlers follow query pattern.

**Maintainability:** Each component has single responsibility. Business logic is clearly expressed. Dependencies are explicit.

**Testability:** Command handler logic can be tested with command input and event output verification. Projections can be tested with event input and query output verification.

## 8.5 Configuration and Infrastructure

Real-world systems require appropriate configuration:

```

01 // Axon configuration
02 @Configuration
03 public class AxonConfig {
04
05     @Bean
06     public Configurer configurer(
07         EventStore eventStore,
08         TokenStore tokenStore) {
09
10         return DefaultConfigurer.defaultConfiguration()
11             .configureEmbeddedEventStore(c -> eventStore)
12             .configureTokenStore(c -> tokenStore)
13             .configureEventSerializer(c -> JacksonSerializer.builder()
14                 .objectMapper(objectMapper())
15                 .build())
16             .configureMessageSerializer(c -> JacksonSerializer.builder()
17                 .objectMapper(objectMapper())
18                 .build());
19     }
20
21     @Bean
22     public EventStore eventStore(AxonServerConnectionManager
23 connectionManager,
24                                 AxonServerConfiguration configuration) {
25         return AxonServerEventStore.builder()
26             .platformConnectionManager(connectionManager)
27             .configuration(configuration)
28             .build();
29     }
30
31     @Bean
32     public TokenStore tokenStore(DataSource dataSource) {
33         return JdbcTokenStore.builder()
34             .connectionProvider(new
35 UnitOfWorkAwareConnectionProviderWrapper(
36                 new DataSourceConnectionProvider(dataSource)))
37             .schema(TokenSchema.defaultInstance())
38             .build();
39     }
40
41     private ObjectMapper objectMapper() {

```

```
40     ObjectMapper mapper = new ObjectMapper();
41     mapper.registerModule(new JavaTimeModule());
42     mapper.disable(SerializationFeature.WRITE_DATES_AS_TIMESTAMPS);
43     return mapper;
44 }
45 }
```

# 9. Cost Analysis

## 9.1 Without Clear Specifications

### Initial Phase:

- High initial velocity (AI generates code rapidly)
- Minimal specification overhead
- Rapid feature implementation

### Development Phase:

- Extended code review (3-4x normal duration due to unclear intent)
- Frequent refactoring (assumptions proved incorrect)
- Technical debt accumulation
- Performance issues emerge
- Security vulnerabilities discovered

### Maintenance Phase:

- Decreasing velocity (complexity compounds)
- Increasing bug rate (interactions not fully understood)
- Difficult onboarding (no clear system model)
- Expensive modifications (unclear boundaries and dependencies)

### Total Cost Profile:

- Low upfront investment
- Exponentially increasing maintenance costs
- Decreasing business agility over time

## 9.2 With Event Modeling and Event Sourcing

### Initial Phase:

- Moderate initial velocity (specification work required)
- Days to weeks of event modeling workshops

- Structured problem understanding

#### **Development Phase:**

- Rapid code review (generated code matches specification)
- Minimal refactoring (design validated upfront)
- Controlled technical debt
- Predictable performance
- Security validated in specification

#### **Maintenance Phase:**

- Sustained velocity (clear patterns and boundaries)
- Low bug rate (behavior well-defined)
- Efficient onboarding (clear system model)
- Manageable modifications (explicit dependencies)

#### **Total Cost Profile:**

- Higher upfront investment
- Stable maintenance costs
- Maintained business agility over time

### **9.3 Observed Patterns from Enterprise Customers**

Based on observations and feedback from enterprise customers working with event-driven architectures:

**Code Review Duration:** Customers report that code reviews for AI-generated code take substantially less time when working from clear specifications. Where teams previously spent most of a day reviewing a feature implementation, they now spend a few hours. The difference appears most pronounced in complex business logic where ambiguity would otherwise require extensive back-and-forth.

**Bug Rate (early months):** We observe significantly fewer bugs in AI-generated code when specifications are clear. Teams working without detailed specifications report finding issues regularly during initial testing phases. Teams working from event models report finding occasional issues, typically at integration points rather than in core business logic.

**Onboarding Time:** New team members appear to reach productivity much faster when event models document system behavior. Where teams previously reported months for new developers to understand the system well enough to contribute confidently, teams with event models report weeks. The explicit documentation of “what happens when” seems to eliminate much of the learning curve.

**Modification Cost (after months of operation):** Changes to existing functionality cost substantially more in systems without clear specifications. Teams report that modifications often require extensive code archaeology to understand current behavior before making changes. With event sourcing and clear specifications, the history is explicit and modifications are more straightforward. Customers indicate that changes typically cost a fraction more than the original implementation rather than several times more.

## 9.4 When Specification Overhead Is Justified

Event modeling with event sourcing makes economic sense when:

**Long-term maintenance is expected:** Systems that will be maintained for years benefit from clear specifications.

**Business logic is complex:** Domains with intricate rules and workflows justify upfront specification investment.

**Team size is substantial:** Larger teams benefit more from shared understanding through explicit specifications.

**Compliance requirements exist:** Regulated industries require auditable behavior documentation.

**AI-assisted development is standard practice:** The quality difference between specified and unspecified AI-generated code increases with system complexity.

Event modeling may not be economically justified for:

**Short-lived prototypes:** Systems expected to be rewritten or discarded don't benefit from long-term maintainability.

**Genuinely simple systems:** Systems with minimal business logic don't require elaborate specification frameworks.

**Small teams with strong tacit knowledge:** Two developers who fully understand the problem space may not need explicit specifications.

# 10. Conclusions and Recommendations

## 10.1 Key Findings

AI coding assistants represent a fundamental shift in software development economics. These tools function as process multipliers. They accelerate whatever development approach they encounter. This makes specification quality the critical variable. It determines whether AI assistance delivers genuine productivity gains or accelerates technical debt accumulation.

Traditional development processes absorbed specification ambiguity through human judgment at every implementation step. AI removes this buffer. It exposes the hidden costs of inadequate specifications. The slow feedback loops that previously caught specification problems now arrive too late to prevent substantial wasted effort.

Event modeling provides a specification framework well-suited for AI-assisted development. It operates at the appropriate level of abstraction. Detailed enough for implementation. High-level enough for stakeholder validation. Event modeling produces unambiguous specifications that enable AI systems to generate correct implementations.

Pattern repetition in event-driven architectures gives AI clear, consistent examples to follow when generating code. Rather than many different architectural approaches applied sporadically, event-driven systems employ four patterns repeatedly. Hundreds or thousands of implementations of the same patterns. With a limited number of patterns repeated consistently, AI has clear guidelines on what to implement. There is little left to hallucinate about. The AI doesn't fall back on general knowledge that may not fit your architecture.

Event sourcing's complete historical context enables more effective AI reasoning about system behavior. Instead of current state snapshots, AI systems have access to complete behavioral stories. This enables temporal reasoning, causality understanding, and pattern recognition. All of which improves generated code quality.

## 10.2 Practical Recommendations

**For Organizations Adopting AI-Assisted Development:**

1. **Invest in specification practices:** The cost of specification development is substantially lower than the cost of reviewing and refactoring AI-generated code that solved the wrong problem.
2. **Evaluate event modeling:** For systems where long-term maintainability matters, event modeling provides proven specification approaches compatible with AI-assisted development.
3. **Prioritize pattern consistency:** Establishing repeatable patterns throughout your codebase improves AI-generated code quality more than sophisticated abstractions.
4. **Measure the right metrics:** Optimize for code quality and maintainability rather than raw generation velocity.

#### **For Development Teams:**

1. **Resist the temptation to skip specification:** Initial velocity gains from immediate code generation create technical debt that compounds rapidly.
2. **Learn event modeling techniques:** Even if not implementing full event sourcing, event modeling's specification approaches improve AI-assisted development outcomes.
3. **Build purpose-specific models:** Resist the urge to create shared abstractions. Purpose-built models improve clarity, independence, and AI-friendliness.
4. **Leverage AI for appropriate tasks:** Use AI for implementation details and boilerplate generation, not for fundamental architectural decisions.

#### **For Technical Leaders:**

1. **Set specification standards:** Establish clear expectations for specification quality before AI-assisted implementation begins.
2. **Create pattern libraries:** Document and enforce consistent patterns that AI systems can learn from.
3. **Monitor technical debt:** Track indicators of specification quality problems before they compound.
4. **Budget for specification work:** Allocate time for event modeling workshops and specification development in project plans.

### **10.3 Future Directions**

AI-assisted development will continue evolving. Several trends warrant attention:

**Improved context handling:** As AI models' context windows expand, they'll better handle complex system specifications. This strengthens the case for comprehensive event modeling.

**Agentic systems:** More autonomous AI systems will require even clearer specifications and guardrails. Event modeling's explicit boundaries become increasingly important.

**Specification as code:** Tools that represent event models in machine-readable formats will enable more sophisticated AI assistance.

**Pattern learning:** AI systems that actively learn from your codebase patterns will benefit more from consistent event-driven architectures.

## 10.4 Final Assessment

The productivity paradox of AI-assisted development resolves when we recognize a simple truth: these tools optimize whatever we ask them to optimize.

Without clear specifications, they optimize code generation velocity. At the expense of correctness and maintainability.

With clear specifications, they optimize implementation of well-defined behavior.

Event modeling and event sourcing don't make AI coding assistants necessary. They make them productive. The combination works: human domain expertise (captured in event models) plus AI implementation capability (guided by clear specifications). This delivers genuine productivity improvements without sacrificing code quality or long-term maintainability.

For organizations serious about AI-assisted development, specification practices aren't optional overhead. They're the foundation. They determine whether AI assistance accelerates value delivery or technical debt accumulation.

The choice is clear: invest in specifications upfront, or pay exponentially more in maintenance costs later.

# References

1. Fowler, M. (2005). Event Sourcing. MartinFowler.com. Available at: <https://martinfowler.com/eaaDev/EventSourcing.html>
2. Dymitruk, A. (2020). Event Modeling: What is it? EventModeling.org. Available at: <https://eventmodeling.org/posts/what-is-event-modeling/>
3. Dilger. (2024). Understanding Event Sourcing: Planning and Implementing Scalable Systems with Event Modeling and Event Sourcing. Available at: (<https://leanpub.com/eventmodeling-and-eventsourcing>)
4. Young, G. (2010). CQRS and Event Sourcing. Retrieved from various conference presentations and documentation.
5. Vernon, V. (2013). Implementing Domain-Driven Design. Addison-Wesley Professional.
6. Axon Framework Documentation (2024). Axon Framework Reference Guide. AxonIQ. Available at: <https://docs.axoniq.io/>
7. Evans, E. (2003). Domain-Driven Design: Tackling Complexity in the Heart of Software. Addison-Wesley Professional.
8. Stopford, B. (2018). Designing Event-Driven Systems. O'Reilly Media.
9. Kleppmann, M. (2017). Designing Data-Intensive Applications. O'Reilly Media.
10. Richardson, C. (2018). Microservices Patterns. Manning Publications.

# About the Authors

**Allard Buijze** is the founder of Axoniq and creator of the Axon Framework. With over 15 years of experience in event-driven architecture, he has guided numerous organizations through event sourcing and CQRS implementations. He regularly speaks at conferences worldwide on topics related to event-driven design, domain-driven design, and software architecture.

**Martin Dilger** is the founder of Nebulit GmbH and author of *Understanding Eventsourcing*, the first German-language book on Event Modeling. With over 20 years of experience in software architecture, he has helped teams across the DACH region align business and engineering, close requirements gaps, and deliver event-driven systems with confidence. He is a recognized core contributor to event modeling and regularly speaks and trains on event-driven design, event sourcing, and CQRS.

**Adam Dymitruk** is the CEO of Adaptech Group and the creator of event modeling. With over three decades of experience in software development and architecture, Adam is an expert in event-driven systems and event sourcing. He has been instrumental in advancing these methodologies, particularly through his work on event modeling, which provides a visual and collaborative approach to system design. Adam is a frequent speaker at international conferences and has contributed significantly to the evolution of responsible, efficient system design and implementation in the software industry.

# About Axoniq

Axoniq provides products and services for building event-driven applications using event sourcing and CQRS. The company's flagship products include Axon Framework, a Java framework for building event-driven applications, and Axon Server, a purpose-built event store and message routing solution.

For more information: <https://www.axoniq.io>

## ABOUT Nebulit

Nebulit helps software teams in Europe close requirements gaps, align business and engineering, and deliver event-driven and event-sourced systems with confidence - through workshops, training, and architecture consulting built on industry proven methods and technologies.

Learn more: <https://nebulit.de/en/index.html>

## ABOUT Adaptech Group

Adaptech Group is a collective of top technologists focusing around microservices and event-centric architecture, involving concepts like event modeling and event sourcing. They provide teams of highly skilled developers that can jump-start projects, services, or products. They provide consultation, outsourcing, and even staffing; following a method of enabling a unified vision across the organization to enable the business and technical staff to use the same language and share the same vision. Adaptech Group embraces the understanding of Conway's Law as a means to offer solutions with the best architecture, enabling scalability concerns to be adjusted easily in the future.

Learn more: <https://adaptechgroup.com/>

This whitepaper represents analysis and opinions based on real-world experience and established software engineering principles. Code examples are illustrative and simplified for educational purposes. Implementation results may vary based on specific organizational context, team capabilities, and system requirements. Always thoroughly test and adapt approaches based on your specific needs and constraints before deploying to production.