

# ZelC: A Cybersecurity-Native Programming Language for Intent-Gated, Capability-Scooped, Evidence-Generating Operations in the Agentic AI Era

Technical Report

Haja Mo  
Rocheston  
research@rocheston.com

February 10, 2026

## Abstract

Security engineering has shifted from “write software, deploy once” to continuous, adversarial operations: detect, correlate, contain, recover, and prove. General-purpose languages (GPLs) can automate these workflows, but they model security actuation as ordinary side effects of arbitrary code, leaving safety, governance, and evidentiary integrity to convention. This paper presents ZELC, a cybersecurity-native language and runtime in which *(i)* observation is separated from actuation (“kinetic execution”), *(ii)* authority is capability-scooped rather than ambient, *(iii)* actuation is constrained by machine-checkable intent contracts with formal blast-radius analysis, *(iv)* evidence artifacts are generated as a semantic output of execution and optionally anchored into verifiable append-only logs, and *(v)* information-flow control via taint tracking prevents untrusted data from reaching sensitive operational sinks.

We present the complete formal model: abstract syntax, static and dynamic semantics via typing rules and small-step operational semantics, a capability algebra with monotonic non-escalation proofs, intent contracts formalized as constraint-satisfaction problems, blast-radius bounding via abstract interpretation, a cryptographic evidence chain with Merkle-tree anchoring and inclusion-proof verification, and a conformance predicate that provides agent-agnostic safety guarantees. We include concrete ZELC program listings for ransomware containment, cloud key rotation, brute-force defense, CI/CD integrity verification, and zero-trust network gating to demonstrate that the formal model maps directly to operational reality.

We explicitly distinguish foundational primitives—capability security [1, 2], effect typing [3, 4], admission-control policy gates [5–7], information-flow control [10, 16, 17], tamper-evident logging [11–13], Merkle-tree transparency logs [14, 15], proof-carrying code [18], and abstract interpretation [9]—from ZELC’s novel contribution: an *integrated programming model* that binds these mechanisms into one coherent operational semantics for cybersecurity operations in the agentic AI era.

**Keywords:** cybersecurity operations · incident response · policy-as-code · capability security · effect systems · information-flow control · tamper-evident logs · Merkle trees · proof-carrying code · agentic AI safety · abstract interpretation · domain-specific languages

## Contents

<b>1</b>	<b>Introduction</b>	<b>5</b>
1.1	Operational Context . . . . .	5
1.2	Design Thesis: The Kinetic Gap . . . . .	5
1.3	Scope and Contributions . . . . .	5
<b>2</b>	<b>Prior Work and Positioning</b>	<b>6</b>
2.1	Capability-Based Security . . . . .	6
2.2	Effect Systems . . . . .	6
2.3	Policy-as-Code and Admission Control . . . . .	7
2.4	Information-Flow Control and Taint Tracking . . . . .	7
2.5	Tamper-Evident Logging . . . . .	7
2.6	Transparency Logs, Merkle Trees, and Ledgers . . . . .	7
2.7	Proof-Carrying Code and Static Verification . . . . .	7
2.8	Abstract Interpretation . . . . .	7
2.9	Constrained Language Model Programming . . . . .	8
<b>3</b>	<b>Core Formal Model</b>	<b>8</b>
3.1	Semantic Domains . . . . .	8
3.2	Type Universe . . . . .	8
3.3	Abstract Syntax . . . . .	9
3.4	Effect Typing: Read-Only versus Kinetic . . . . .	9
<b>4</b>	<b>Capability-Spaced Authority</b>	<b>11</b>
4.1	Capability Algebra . . . . .	11
4.2	Attenuation and Non-Escalation . . . . .	11
<b>5</b>	<b>Intent Contracts and Blast-Radius Constraints</b>	<b>12</b>
5.1	Intent as a Constrained Operational Contract . . . . .	12
5.2	Blast Radius Formalization . . . . .	12
5.3	Static Blast-Radius Bounding via Abstract Interpretation . . . . .	13
<b>6</b>	<b>Evidence-First Semantics</b>	<b>14</b>
6.1	Evidence Object Model . . . . .	14
6.2	Cryptographic Hash Chaining . . . . .	14
6.3	Signature Blocks . . . . .	15
6.4	Merkle Tree Anchoring . . . . .	15

<b>7 Taint Tracking and Information-Flow Constraints</b>	<b>15</b>
7.1 Security Label Lattice . . . . .	15
7.2 Taint Propagation Rules . . . . .	16
7.3 Sensitive Sinks and Sanitization . . . . .	16
<b>8 Agentic AI Safety: Intent-Gated Execution</b>	<b>16</b>
8.1 Threat Model for Agentic Generation . . . . .	16
8.2 Conformance Predicate . . . . .	17
8.3 Verify Before Execute: Operational Proof Obligations . . . . .	17
<b>9 Operational Semantics</b>	<b>17</b>
9.1 Runtime Configuration . . . . .	17
9.2 Small-Step Transition Rules . . . . .	18
<b>10 Concrete Operational Programs</b>	<b>19</b>
10.1 Ransomware Containment . . . . .	19
10.2 Cloud Key Leak Response . . . . .	20
10.3 Brute-Force Defense . . . . .	21
10.4 CI/CD Pipeline Integrity Verification . . . . .	21
10.5 Zero-Trust Network Gatekeeper . . . . .	23
<b>11 Architectural Overview</b>	<b>24</b>
11.1 Execution Pipeline . . . . .	24
11.2 Three-Layer Safety Architecture . . . . .	24
11.3 Memory Model . . . . .	25
11.4 Concurrency Model . . . . .	25
<b>12 Tiered Operational Model</b>	<b>25</b>
<b>13 Comparison with General-Purpose Languages</b>	<b>26</b>
<b>14 Evaluation Methodology</b>	<b>26</b>
<b>15 Contribution Claims</b>	<b>27</b>
<b>16 Conclusion</b>	<b>28</b>
<b>A ZelC Type Universe: Complete Reference</b>	<b>28</b>
<b>B ZelC Action Taxonomy</b>	<b>29</b>

**C Formal Proof: Evidence Chain Integrity****30**

## 1 Introduction

### 1.1 Operational Context

Modern security work is continuous operations under time pressure: isolate endpoints, rotate keys, revoke sessions, quarantine artifacts, snapshot disks, notify responders, and produce evidence that stands up to auditing and legal scrutiny. The correctness criterion is not only “did the action execute,” but “was it authorized, bounded, reviewable, and provable.”

General-purpose languages model execution as a sequence of steps with broad ambient authority: a process executing a script may implicitly hold credentials, filesystem rights, network access, or cloud tokens. In such environments, “rotate a key” and “delete production data” may differ only by parameters, not by language semantics. Consider the Python fragment:

```
import boto3
ec2 = boto3.client('ec2')
ec2.terminate_instances(InstanceIds=['i-prod-critical'])
# No safety gate. No evidence. No blast-radius check.
# Production is gone.
```

There is no structural distinction between this catastrophic action and a benign health-check. The language offers no mechanism to express intent, enforce blast-radius limits, or generate evidence.

### 1.2 Design Thesis: The Kinetic Gap

We define the *kinetic gap* as the mismatch between the operational demands of security actuation and how typical programming languages structure authority, effects, and evidence.

**Definition 1.1** (Kinetic Gap). The kinetic gap  $\mathcal{G}$  is the four-dimensional mismatch between operational security requirements and language-level semantics:

$$\mathcal{G} = \langle \mathcal{G}_{\text{act}}, \mathcal{G}_{\text{auth}}, \mathcal{G}_{\text{gov}}, \mathcal{G}_{\text{ev}} \rangle \quad (1)$$

where each component represents a distinct semantic deficiency:

- $\mathcal{G}_{\text{act}}$  : Side effects are not isolated from analytic computation.
- $\mathcal{G}_{\text{auth}}$  : Permissions exist as ambient privilege, not explicit objects.
- $\mathcal{G}_{\text{gov}}$  : Blast-radius constraints are external to code semantics.
- $\mathcal{G}_{\text{ev}}$  : Logs are mutable, optional, and frequently incomplete. (2)

ZELC’s core thesis: cybersecurity operations deserve a *language-level semantics* that makes actuation explicit, bounded, and evidentiary by construction. Formally, ZELC aims to close  $\mathcal{G}$  such that each component is reduced to  $\perp$  (null gap) under its operational semantics:

$$\llbracket \text{ZELC} \rrbracket \models \mathcal{G}_{\text{act}} = \mathcal{G}_{\text{auth}} = \mathcal{G}_{\text{gov}} = \mathcal{G}_{\text{ev}} = \perp \quad (3)$$

### 1.3 Scope and Contributions

This paper makes the following contributions:

**C1 Formal language model.** We define ZELC’s abstract syntax, type system with a two-point effect lattice ( $\text{RO} \sqsubseteq \text{KI}$ ), and small-step operational semantics with evidence generation as a side-product of kinetic transitions.

**C2 Capability algebra.** We formalize capability injection, attenuation, and the monotonic non-escalation invariant, proving that authority at any program point is bounded by initial injection.

**C3 Intent contracts and blast-radius analysis.** We model intents as constraint-satisfaction problems and provide a sound static blast-radius checker via abstract interpretation.

**C4 Cryptographic evidence chain.** We formalize the hash chain, signature blocks, Merkle-tree anchoring, and inclusion-proof verification, proving evidence completeness for well-typed programs.

**C5 Agent-agnostic safety.** We define the conformance predicate and prove that the same safety guarantees hold regardless of whether code is human-authored or AI-generated.

**C6 Taint tracking for operational sinks.** We specialize information-flow control for SOC actuation sinks (firewall rules, IAM revocations, process termination) with evidence-generating sanitization.

**C7 Concrete operational examples.** We present complete ZELC programs for ransomware containment, cloud key rotation, brute-force defense, CI/CD integrity, and zero-trust gating.

## 2 Prior Work and Positioning

This section enumerates major adjacent prior art to pre-empt overclaiming. In each case, we state the relationship precisely: ZELC *uses* the concept as an internal mechanism but does not claim its invention.

### 2.1 Capability-Based Security

Capability-based security was introduced by Dennis and Van Horn [1] and developed extensively in the object-capability model by Miller [2]. Authority is represented by unforgeable tokens, contrasted with ACL-driven ambient authority; capabilities implement least privilege and avoid confused-deputy attacks.

ZELC uses the capability concept as a *language invariant*: kinetic actions are only possible when the relevant capability objects are present and validated. ZELC does not claim to invent capabilities.

### 2.2 Effect Systems

Effect typing makes side effects explicit at the type level, supporting disciplined separation of pure computation from effectful computation. Lucassen and Gifford [3] introduced the first effect system; Koka [4] is a well-known example of row-polymorphic effect types.

ZELC specializes the effect boundary for security operations by introducing “kinetic” actuation as the *primary effect of interest*. ZELC does not claim to invent effect typing.

### 2.3 Policy-as-Code and Admission Control

Kubernetes admission controllers [5] evaluate create/update/delete requests and can reject unsafe changes before state is written. OPA [6] formalizes policy evaluation prior to admission. Terraform Sentinel [7] implements policy-as-code gating for infrastructure changes.

ZELC internalizes gating into the *language semantics* of kinetic blocks and applies it to multi-domain security actuation.

### 2.4 Information-Flow Control and Taint Tracking

Denning [10] introduced the lattice model for secure information flow. TaintDroid [16] is a canonical system-wide dynamic taint tracking system. Jif [17] is a security-typed language extension for Java supporting IFC.

ZELC applies taint/IFC to operational sinks (firewall rules, IAM revocations, deletions) and couples violations to intent rejection and evidence emission.

### 2.5 Tamper-Evident Logging

Schneier and Kelsey [11] describe mechanisms to detect log manipulation. Snodgrass et al. [12] address tamper detection in audit logs. RFC 5848 [13] specifies signed syslog messages.

ZELC's claim is that evidence generation is a *semantic consequence* of kinetic execution.

### 2.6 Transparency Logs, Merkle Trees, and Ledgers

Certificate Transparency [14] logs are append-only and publicly auditable. Sigstore Rekor [15] provides a transparency log with signed tree heads. Amazon QLDB [20] provides cryptographically verifiable immutable transaction history.

ZELC provides a standardized evidence object model and anchoring interface consistent with transparency-log properties.

### 2.7 Proof-Carrying Code and Static Verification

Proof-carrying code [18] attaches machine-checkable proofs to untrusted code. The eBPF verifier [8] statically analyzes programs before execution.

ZELC borrows the “verify before execute” philosophy, applying it to kinetic constraints and capability proofs.

### 2.8 Abstract Interpretation

Cousot and Cousot [9] introduced abstract interpretation as a theory of sound approximation of program semantics. ZELC uses abstract interpretation for compile-time blast-radius bounding.

## 2.9 Constrained Language Model Programming

LMQL [19] formalizes constraints over LLM output. ZELC *complements* constrained generation by enforcing constraints at the actuation layer, even if the textual output of an agent is imperfect.

## 3 Core Formal Model

### 3.1 Semantic Domains

**Definition 3.1** (Semantic Domains). The ZELC semantic model is parameterized over the following domains:

$\Sigma \in \mathbf{State}$	Machine state	(4)
$\mathcal{C} \subseteq \mathbf{Cap}$	Capability set	(5)
$\mathcal{I} \in \mathbf{Intent}$	Intent contract	(6)
$P \in \mathbf{Prog}$	ZelC program	(7)
$\mathcal{A}(P) \subseteq \mathbf{Act}$	Kinetic actions in $P$	(8)
$\mathcal{E} \subseteq \mathbf{Ev}$	Evidence trace	(9)
$H : \{0, 1\}^* \rightarrow \{0, 1\}^{256}$	Hash (SHA-256)	(10)
$\text{Sign} : \mathbf{SK} \times \{0, 1\}^* \rightarrow \mathbf{Sig}$	Digital signature	(11)
$\text{Verify} : \mathbf{PK} \times \{0, 1\}^* \times \mathbf{Sig} \rightarrow \{\text{true}, \text{false}\}$	Signature verify	(12)

We additionally define the *state transition function*  $\delta : \mathbf{State} \times \mathbf{Act} \rightarrow \mathbf{State}$  such that  $\delta(\Sigma, a(\vec{v}))$  yields the successor state after executing action  $a$  with arguments  $\vec{v}$ .

### 3.2 Type Universe

ZELC provides both standard types and a rich set of *semantic security types* that distinguish security-relevant values at the type level:

**Definition 3.2** (Type Universe). The type universe  $\mathcal{T}$  is defined by the following grammar:

$$\tau ::= \mathbf{unit} \mid \mathbf{Bool} \mid \mathbf{Int} \mid \mathbf{Float} \mid \mathbf{Text} \mid \mathbf{Duration} \quad (13)$$

$$\mid \mathbf{IP} \mid \mathbf{CIDR} \mid \mathbf{Hash} \mid \mathbf{UserId} \mid \mathbf{CVE} \mid \mathbf{ControlId} \quad (14)$$

$$\mid \mathbf{RiskScore} \mid \mathbf{IOCSet} \mid \mathbf{Incident} \mid \mathbf{Timeline} \mid \mathbf{EvidencePack} \quad (15)$$

$$\mid \mathbf{List}[\tau] \mid \mathbf{Map}[\tau_k, \tau_v] \mid \mathbf{Record}\{l_1 : \tau_1, \dots, l_n : \tau_n\} \quad (16)$$

$$\mid \tau_1 \rightarrow^{\varepsilon} \tau_2 \quad (17)$$

where  $\varepsilon \in \{\mathbf{RO}, \mathbf{KI}\}$  is the effect annotation on function types.

**Remark 3.3** (Semantic Types Prevent Confusion). Unlike GPLs where an IP address is a `String`, in ZELC the expression `firewall block ip "Hello World"` is a *compile-time type error*: the `block` action requires an argument of type `IP`, not `Text`. This eliminates an entire class of injection and parameter-confusion bugs at the type level.

### 3.3 Abstract Syntax

**Definition 3.4** (Full Abstract Syntax). The abstract syntax of ZELC is given by:

**Programs:**

$$P ::= \mathbf{check} \ Name \{ S \} \mid \mathbf{define} \ f(\vec{x}) \{ S \} \mid \mathbf{intent} \ Name \{ \mathcal{I}_{\text{body}} \} \mid P ; P \quad (18)$$

**Statements:**

$$S ::= S ; S \mid \mathbf{set} \ x = E \mid \mathbf{change} \ x = E \mid \mathbf{keep} \ x = E \quad (19)$$

$$\mid \mathbf{when} \ E \{ S \} [\mathbf{otherwise} \{ S \}] \mid \mathbf{each} \ x \ \mathbf{in} \ E \{ S \} \quad (20)$$

$$\mid \mathbf{do}[\mathcal{I}, \mathcal{C}] \{ K \} \mathbf{end} \mid \mathbf{try} \{ S \} \mathbf{catch} \ e \{ S \} \quad (21)$$

$$\mid \mathbf{evidence} \ \mathbf{record} \ title \ \mathbf{details} \ E \mid \mathbf{audit} \ \mathbf{log} \ event \ \mathbf{details} \ E \quad (22)$$

$$\mid \mathbf{alert} \ severity \ \mathbf{message} \ E \mid \mathbf{notify} \ channel \ E \mid \mathbf{stop} \mid \mathbf{wait} \ E \quad (23)$$

**Kinetic Statements (only inside do...end):**

$$K ::= \mathbf{action} \ a(\vec{v}) \mid K ; K \mid \mathbf{each} \ x \ \mathbf{in} \ E \{ K \} \quad (24)$$

$$\mid \mathbf{when} \ E \{ K \} \mid \mathbf{evidence} \ \mathbf{record} \ E \mid \mathbf{audit} \ \mathbf{log} \ E \quad (25)$$

The **do ... end** block is the *kinetic boundary*: only kinetic statements  $K$  are permitted inside. Outside **do**, only analytic (read-only) statements are admitted. This is the structural enforcement mechanism that closes  $\mathcal{G}_{\text{act}}$ .

**Definition 3.5** (Intent Contract Body). The intent body  $\mathcal{I}_{\text{body}}$  binds to a named intent declaration:

$$\mathcal{I}_{\text{body}} ::= \mathbf{category}: c \mid \mathbf{risk\_level}: \ell \mid \mathbf{constraints} \{ C_{\text{list}} \} \quad (26)$$

$$C_{\text{list}} ::= \mathbf{max\_targets}: n \mid \mathbf{time\_limit}: d \mid \mathbf{requires\_approval}: a \mid \mathbf{exclude\_scope}: S_{\text{list}} \quad (27)$$

### 3.4 Effect Typing: Read-Only versus Kinetic

**Definition 3.6** (Effect Lattice). We define a two-point effect lattice:

$$\varepsilon \in \{ \text{RO}, \text{KI} \} \quad \text{with ordering} \quad \text{RO} \sqsubseteq \text{KI} \quad (28)$$

where RO denotes *read-only* (analytic) effects and KI denotes *kinetic* (state-changing) effects. The join operator is:

$$\varepsilon_1 \sqcup \varepsilon_2 = \begin{cases} \text{RO} & \text{if } \varepsilon_1 = \varepsilon_2 = \text{RO} \\ \text{KI} & \text{otherwise} \end{cases} \quad (29)$$

Typing judgments take the form  $\Gamma \vdash E : \tau ! \varepsilon$ , meaning expression  $E$  has type  $\tau$  and effect  $\varepsilon$  under context  $\Gamma$ .

**Definition 3.7** (Core Typing Rules). The ZELC type system enforces the kinetic boundary via the following rules:

**Rule (VAR)** — Variable lookup is read-only:

$$\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau ! \text{RO}} \text{ VAR} \quad (30)$$

**Rule (SET)** — Binding is read-only:

$$\frac{\Gamma \vdash E : \tau ! \text{RO}}{\Gamma \vdash \text{set } x = E : \text{unit} ! \text{RO}} \text{ SET} \quad (31)$$

**Rule (ACT)** — Kinetic actions carry effect  $\text{KI}$ :

$$\frac{a \in \mathbf{Act}_{\text{KI}} \quad \Gamma \vdash \vec{v} : \vec{\tau} ! \text{RO}}{\Gamma \vdash \mathbf{action} a(\vec{v}) : \text{unit} ! \text{KI}} \text{ ACT} \quad (32)$$

**Rule (Do)** — Kinetic blocks require authorization and intent satisfaction:

$$\frac{\Gamma \vdash K : \text{unit} ! \text{KI} \quad \text{Auth}(\Sigma, \mathcal{C}, K) \quad \text{Sat}(\mathcal{I}, \Sigma, K)}{\Gamma \vdash \mathbf{do}[\mathcal{I}, \mathcal{C}] \{K\} \mathbf{end} : \text{unit} ! \text{KI}} \text{ Do} \quad (33)$$

**Rule (OUTSIDE)** — Non-kinetic code must remain read-only:

$$\frac{S \text{ outside any } \mathbf{do} \text{ block} \quad \Gamma \vdash S : \tau ! \varepsilon}{\varepsilon = \text{RO}} \text{ OUTSIDE} \quad (34)$$

**Rule (SEQ)** — Sequential composition propagates effects:

$$\frac{\Gamma \vdash S_1 : \tau_1 ! \varepsilon_1 \quad \Gamma \vdash S_2 : \tau_2 ! \varepsilon_2}{\Gamma \vdash S_1 ; S_2 : \tau_2 ! (\varepsilon_1 \sqcup \varepsilon_2)} \text{ SEQ} \quad (35)$$

**Rule (WHEN)** — Conditional preserves effect:

$$\frac{\Gamma \vdash E : \mathbf{Bool} ! \text{RO} \quad \Gamma \vdash S_1 : \tau ! \varepsilon_1 \quad \Gamma \vdash S_2 : \tau ! \varepsilon_2}{\Gamma \vdash \mathbf{when} E \{S_1\} \mathbf{otherwise} \{S_2\} : \tau ! (\varepsilon_1 \sqcup \varepsilon_2)} \text{ WHEN} \quad (36)$$

**Rule (EACH)** — Iteration propagates body effect:

$$\frac{\Gamma \vdash E : \mathbf{List}[\tau'] ! \text{RO} \quad \Gamma, x : \tau' \vdash S : \tau ! \varepsilon}{\Gamma \vdash \mathbf{each} x \mathbf{in} E \{S\} : \text{unit} ! \varepsilon} \text{ EACH} \quad (37)$$

**Rule (EVIDENCE)** — Evidence emission is always read-only (it is a *semantic output*, not a state change):

$$\frac{\Gamma \vdash E : \mathbf{Record}\{\dots\} ! \text{RO}}{\Gamma \vdash \mathbf{evidence record} \text{ title details } E : \text{unit} ! \text{RO}} \text{ EV} \quad (38)$$

**Theorem 3.8** (Effect Soundness). *If a well-typed ZELC program  $P$  satisfies  $\Gamma \vdash P : \tau ! \text{RO}$ , then  $P$  performs no kinetic actions. Conversely, any kinetic action  $a(\vec{v})$  in  $P$  must occur inside a **do**...**end** block.*

*Proof.* By structural induction on the typing derivation. The only rule introducing effect **KI** is (ACT), which requires  $a \in \mathbf{Act}_{\mathbf{KI}}$ . By rule (OUTSIDE), any statement outside a **do** block must have effect **RO**. Since **RO**  $\neq$  **KI**, rule (ACT) cannot appear outside **do**. Rule (Do) is the only construct that promotes **KI** into the enclosing scope. Therefore, all kinetic actions are structurally confined to **do** blocks.  $\square$   $\square$

## 4 Capability-Scope Authority

### 4.1 Capability Algebra

**Definition 4.1** (Capability). A capability  $\kappa$  is an unforgeable token:

$$\kappa = \langle o, r, \text{ttl}, \text{issuer}, \text{sig} \rangle \in \mathbf{Cap} \quad (39)$$

where  $o \in \mathbf{Obj}$  is the target object (endpoint, IAM identity, firewall rule set, S3 bucket, ...),  $r \subseteq \mathbf{Right}$  is the granted right set,  $\text{ttl} \in \mathbb{R}_{\geq 0} \cup \{\infty\}$  is the time-to-live,  $\text{issuer}$  is the granting principal, and  $\text{sig}$  is a digital signature for authenticity.

**Definition 4.2** (Authorization Predicate). The authorization predicate is:

$$\text{Auth} : \mathbf{State} \times 2^{\mathbf{Cap}} \times \mathbf{Act} \longrightarrow \{\text{true}, \text{false}\} \quad (40)$$

defined as:

$$\text{Auth}(\Sigma, \mathcal{C}, a(\vec{v})) = \exists \kappa \in \mathcal{C} : [\text{obj}(\kappa) = \text{target}(a, \vec{v}) \wedge \text{reqRights}(a) \subseteq r(\kappa) \wedge \text{alive}(\kappa, \Sigma)] \quad (41)$$

where  $\text{alive}(\kappa, \Sigma) \equiv (\Sigma.\text{time} \leq \kappa.\text{ttl}) \wedge \text{Verify}(\kappa.\text{issuer}.pk, \text{serialize}(\kappa), \kappa.\text{sig})$ .

### 4.2 Attenuation and Non-Escalation

**Definition 4.3** (Attenuation Operator). The attenuation operator produces a restricted capability:

$$\text{attenuate}(\kappa, r', \text{ttl}') = \kappa' = \langle o, r', \min(\text{ttl}, \text{ttl}'), \text{issuer}, \text{Sign}(sk, \text{serialize}(\kappa')) \rangle \quad (42)$$

subject to:  $r' \subseteq r$  and  $\text{ttl}' \leq \text{ttl}$ .

**Axiom 4.4** (Monotonic Non-Escalation). For any derived capability  $\kappa'$  from  $\kappa$ :

$$\text{rights}(\kappa') \subseteq \text{rights}(\kappa) \quad \wedge \quad \text{ttl}(\kappa') \leq \text{ttl}(\kappa) \quad (43)$$

This is enforced by construction in Definition 4.3.

**Proposition 4.5** (Least-Privilege Bound). *Under Axiom 4.4, the authority available at any program point  $p$  is bounded:*

$$\forall p \in P, \forall \kappa_p \in \mathcal{C}_p : \text{rights}(\kappa_p) \subseteq \bigcup_{\kappa_0 \in \mathcal{C}_{\text{init}}} \text{rights}(\kappa_0) \quad (44)$$

*Proof.* By induction on the derivation chain from  $\mathcal{C}_{\text{init}}$  to  $\mathcal{C}_p$ . Each derivation step applies attenuate, which by Definition 4.3 can only reduce rights. Therefore the union of rights at  $p$  is bounded by the union of rights at initialization.  $\square$   $\square$

**Definition 4.6** (Capability Scope Binding). Capabilities are explicitly injected into **do** blocks via lexical scoping:

$$\mathbf{do}[\mathcal{I}, \mathcal{C}] \{K\} \mathbf{end} \implies \mathcal{C} \text{ is immutable within } K \quad (45)$$

No capability can be fabricated, copied from ambient scope, or escalated within the kinetic block.

## 5 Intent Contracts and Blast-Radius Constraints

### 5.1 Intent as a Constrained Operational Contract

**Definition 5.1** (Intent Contract). An intent contract  $\mathcal{I}$  is a seven-tuple:

$$\mathcal{I} = \langle \text{Allow}, \text{Deny}, \text{Limits}, \text{Window}, \text{Approvals}, \text{Redaction}, \text{ExcludeScope} \rangle \quad (46)$$

where:

$$\text{Allow} \subseteq \mathbf{Act} \quad \text{Permitted actions} \quad (47)$$

$$\text{Deny} \subseteq \mathbf{Act} \quad \text{Forbidden actions} \quad (48)$$

$$\text{Limits} = \langle L_{\text{tgt}}, L_{\text{rate}}, L_{\text{scope}} \rangle \in \mathbb{N}^3 \quad \text{Numerical bounds} \quad (49)$$

$$\text{Window} = [t_{\text{start}}, t_{\text{end}}] \subseteq \mathbb{R}_{\geq 0} \quad \text{Temporal bounds} \quad (50)$$

$$\text{Approvals} \in \{\mathbf{auto}, \mathbf{manual}(n), \mathbf{quorum}(k, n)\} \quad \text{HITL policy} \quad (51)$$

$$\text{Redaction} \in \mathbf{RedactionRule} \quad \text{Field masking} \quad (52)$$

$$\text{ExcludeScope} \subseteq \mathbf{Obj} \quad \text{Protected objects} \quad (53)$$

**Definition 5.2** (Intent Satisfaction). Intent satisfaction is a conjunction of five conditions:

$$\text{Sat}(\mathcal{I}, \Sigma, a(\vec{v})) \equiv \varphi_{\text{allow}} \wedge \varphi_{\text{deny}} \wedge \varphi_{\text{limits}} \wedge \varphi_{\text{window}} \wedge \varphi_{\text{scope}} \quad (54)$$

where:

$$\varphi_{\text{allow}} \equiv a \in \text{Allow} \quad (55)$$

$$\varphi_{\text{deny}} \equiv a \notin \text{Deny} \quad (56)$$

$$\varphi_{\text{limits}} \equiv \text{BR}(\Sigma, a(\vec{v})) \leq L_{\text{tgt}} \wedge \text{rate}(\Sigma, a) \leq L_{\text{rate}} \quad (57)$$

$$\varphi_{\text{window}} \equiv \Sigma.\text{time} \in [t_{\text{start}}, t_{\text{end}}] \quad (58)$$

$$\varphi_{\text{scope}} \equiv \text{Targets}(\Sigma, a(\vec{v})) \cap \text{ExcludeScope} = \emptyset \quad (59)$$

Kinetic execution is permitted **only if**  $\text{Sat} = \text{true}$ . Otherwise the runtime *fails closed* and emits a rejection evidence record.

### 5.2 Blast Radius Formalization

**Definition 5.3** (Blast Radius Function). Let  $\text{Targets} : \mathbf{State} \times \mathbf{Act} \rightarrow 2^{\mathbf{Obj}}$  return the set of assets impacted by an action. The blast radius is:

$$\text{BR}(\Sigma, a(\vec{v})) = |\text{Targets}(\Sigma, a(\vec{v}))| \quad (60)$$

For composite kinetic blocks, blast radius is additive:

$$\text{BR}(\Sigma, K_1 ; K_2) = |\text{Targets}(\Sigma, K_1) \cup \text{Targets}(\delta(\Sigma, K_1), K_2)| \quad (61)$$

For iteration over  $n$  elements:

$$\text{BR}\left(\Sigma, \text{each } x \text{ in } [v_1, \dots, v_n] \{ \text{action } a(x) \}\right) = \left| \bigcup_{i=1}^n \text{Targets}(\Sigma_i, a(v_i)) \right| \quad (62)$$

where  $\Sigma_i$  is the state after executing the first  $i-1$  iterations.

### 5.3 Static Blast-Radius Bounding via Abstract Interpretation

**Definition 5.4** (Abstract Domain for Blast Radius). Define the abstract domain  $\hat{D} = (\text{State}, \sqsubseteq_{\hat{D}})$  where abstract states  $\hat{\Sigma}$  map variables to intervals or finite sets. The abstraction and concretization functions form a Galois connection:

$$(\alpha, \gamma) : 2^{\text{State}} \rightleftharpoons \text{State} \quad \text{with} \quad \alpha \circ \gamma \sqsupseteq \text{id} \quad (63)$$

following Cousot and Cousot [9].

**Definition 5.5** (Abstract Blast Radius). The abstract blast-radius function  $\hat{U}$  over-approximates BR:

$$\begin{aligned} \hat{U} : \text{State} \times \text{Act} &\rightarrow \mathbb{N} \\ \text{satisfying} \quad \forall \Sigma \in \gamma(\hat{\Sigma}) : \text{BR}(\Sigma, a(\vec{v})) &\leq \hat{U}(\hat{\Sigma}, a(\vec{v})) \end{aligned} \quad (64)$$

For the loop construct **each**  $x$  in  $collection$  { **action**  $a(x)$  }, if the abstract analysis determines  $|collection| \leq n$  (where  $n$  may be an interval upper bound), then:

$$\hat{U}(\hat{\Sigma}, \text{each } x \text{ in } C \{ \text{action } a(x) \}) \leq n \cdot \max_v \hat{U}(\hat{\Sigma}, a(v)) \quad (65)$$

**Theorem 5.6** (Soundness of Static Blast-Radius Checking). *If  $\hat{U}$  is a sound over-approximation of BR (Definition 5.5), then:*

$$\hat{U}(\hat{\Sigma}, a(\vec{v})) \leq L_{\text{tgt}} \implies \text{BR}(\Sigma, a(\vec{v})) \leq L_{\text{tgt}} \quad \forall \Sigma \in \gamma(\hat{\Sigma}) \quad (66)$$

Conversely, if  $\hat{U}(\hat{\Sigma}, a(\vec{v})) > L_{\text{tgt}}$ , compilation is conservatively rejected.

*Proof.* By the soundness property of the Galois connection (Equation 63), for every concrete state  $\Sigma \in \gamma(\hat{\Sigma})$ , the abstract evaluation over-approximates the concrete. Therefore  $\text{BR}(\Sigma, a(\vec{v})) \leq \hat{U}(\hat{\Sigma}, a(\vec{v})) \leq L_{\text{tgt}}$ . The rejection case follows by contrapositive: if the abstract bound exceeds  $L_{\text{tgt}}$ , there *may* exist a concrete state violating the limit, so conservative rejection is sound.  $\square \quad \square$

**Example 5.7** (Compile-Time Rejection). Consider the ZELC fragment:

```

1  intent Isolate {
2    constraints { max_targets: 5 }
3  }
4
5  check MassReboot

```

```

6   set all_hosts = inventory.get_all()  -- |all_hosts| = 500
7   do
8     each host in all_hosts
9       linux service "critical-db" restart  --  $\hat{U} = 500 > 5 = L_{tgt}$ 
10    end
11  end
12 end

```

Listing 1: Blast-radius violation detected at compile time.

The abstract interpreter determines  $|all\_hosts| = 500$  and computes  $\hat{U} = 500 > 5 = L_{tgt}$ . Compilation is rejected:

$$\hat{U} = 500 > 5 = L_{tgt} \implies \text{REJECT ("Blast radius exceeds limit by 99x")} \quad (67)$$

## 6 Evidence-First Semantics

### 6.1 Evidence Object Model

**Definition 6.1** (Evidence Record). For each kinetic action instance  $i$ , ZELC emits a structured evidence record:

$$ev_i = \left\langle \underbrace{id_i, ts_i, \pi_i}_{\text{provenance}}, \underbrace{\mathcal{I}_{id}, fp(\kappa)}_{\text{auth}}, \underbrace{a_i, H(\vec{v}_i)}_{\text{action}}, \underbrace{H(\Sigma_{\text{pre}}), H(\Sigma_{\text{post}})}_{\text{state}}, \underbrace{\rho_i, H(policy_i)}_{\text{result}}, \underbrace{h_{i-1}}_{\text{chain}} \right\rangle \quad (68)$$

where  $\pi_i$  is the executing principal,  $fp(\kappa) = H(\text{serialize}(\kappa))$  is the capability fingerprint,  $\rho_i \in \{\text{success}, \text{failure}, \text{rejected}\}$  is the action result, and  $h_{i-1}$  is the hash of the previous evidence record, forming the chain link.

**Definition 6.2** (Evidence Record Schema (Formal)). Each field has a fixed type from the ZELC type universe:

$$ev : \text{Record} \left\{ \begin{array}{ll} id : \text{Hash}, & ts : \text{Int}, \\ principal : \text{UserId}, & intent_id : \text{Text}, \\ cap_fp : \text{Hash}, & action : \text{Text}, \\ args_digest : \text{Hash}, & pre_state : \text{Hash}, \\ post_state : \text{Hash}, & result : \text{Text}, \\ policy_digest : \text{Hash}, & prev_hash : \text{Hash} \end{array} \right\} \quad (69)$$

### 6.2 Cryptographic Hash Chaining

**Definition 6.3** (Hash Chain). The evidence hash chain is defined recursively:

$$h_0 = 0^{256}, \quad h_i = H(\text{serialize}(ev_i) \parallel h_{i-1}) \quad \text{for } i \geq 1 \quad (70)$$

Any modification to  $ev_j$  for  $j \leq i$  changes  $h_j$  and cascades through all subsequent hashes.

**Lemma 6.4** (Tamper Detection). *Under the collision-resistance assumption on  $H$ , if an adversary modifies  $ev_j$  to produce  $ev'_j \neq ev_j$ , then:*

$$h'_j \neq h_j \implies h'_k \neq h_k \quad \forall k \geq j \quad (71)$$

with overwhelming probability in the security parameter.

*Proof.* By the recursive definition (Equation 70),  $h_j$  depends on  $\text{serialize}(ev_j)$  and  $h_{j-1}$ . If  $ev'_j \neq ev_j$ , then  $\text{serialize}(ev'_j) \neq \text{serialize}(ev_j)$  (assuming injective serialization), so  $H(\text{serialize}(ev'_j) \| h_{j-1}) \neq H(\text{serialize}(ev_j) \| h_{j-1})$  by collision resistance. The cascade follows by induction on  $k$ .  $\square$   $\square$

### 6.3 Signature Blocks

Periodic checkpoints are signed for authenticity and sequencing:

$$sig_j = \text{Sign}(sk, h_{t_j} \| j \| ts_j) \quad (72)$$

Verification:  $\text{Verify}(pk, h_{t_j} \| j \| ts_j, sig_j) = \text{true}$ .

This is aligned with RFC 5848's signature blocks for syslog messages [13], providing: origin authentication, data integrity, replay resistance, sequencing, and missing-message detection.

### 6.4 Merkle Tree Anchoring

**Definition 6.5** (Merkle Tree Construction). For a batch of evidence hashes  $\{h_1, \dots, h_n\}$ , the Merkle tree is constructed recursively. For a complete binary tree with  $2^d$  leaves:

$$M(i, j) = \begin{cases} H(h_i) & \text{if } i = j \\ H(M(i, m) \| M(m+1, j)) & \text{where } m = \lfloor \frac{i+j}{2} \rfloor \end{cases} \quad (73)$$

The Merkle root is  $R = M(1, n)$ . The root  $R$  is anchored into an append-only transparency log [14, 15] or cryptographically verifiable ledger [20].

**Definition 6.6** (Inclusion Proof). Given leaf  $h_k$  and proof path  $\pi_k = \langle s_1, s_2, \dots, s_d \rangle$  (sibling hashes along the path from leaf to root), verification proceeds by recomputing:

$$v_0 = H(h_k), \quad v_i = \begin{cases} H(v_{i-1} \| s_i) & h_k \text{ is left child at level } i \\ H(s_i \| v_{i-1}) & \text{otherwise} \end{cases} \quad (74)$$

and checking:

$$\text{VerifyInclusion}(h_k, \pi_k, R) = \text{true} \iff v_d = R \quad (75)$$

**Proposition 6.7** (Evidence Completeness). *Under ZELC's evidence-first semantics, for every kinetic action  $a_i$  executed in a well-typed program:*

$$\forall a_i \in \mathcal{A}(P) : \exists ev_i \in \mathcal{E}(P) \text{ s.t. } \text{VerifyInclusion}(H(\text{serialize}(ev_i)), \pi_i, R) = \text{true} \quad (76)$$

The design target is evidence completeness  $E_1 = 1.0$ .

## 7 Taint Tracking and Information-Flow Constraints

### 7.1 Security Label Lattice

**Definition 7.1** (Label Lattice). The taint label domain is:

$$\ell \in \mathcal{L} = \{\perp, \top\} \quad \text{with} \quad \perp \sqsubseteq \top \quad (77)$$

where  $\perp$  denotes *untainted* (trusted) and  $\top$  denotes *tainted* (untrusted). Types are annotated:  $\tau^\ell$ . The join:

$$\ell_1 \sqcup \ell_2 = \begin{cases} \perp & \text{if } \ell_1 = \ell_2 = \perp \\ \top & \text{otherwise} \end{cases} \quad (78)$$

## 7.2 Taint Propagation Rules

**Definition 7.2** (Propagation). For any  $n$ -ary operator  $f$ :

$$\frac{\Gamma \vdash x_1 : \tau_1^{\ell_1} \quad \dots \quad \Gamma \vdash x_n : \tau_n^{\ell_n}}{\Gamma \vdash f(x_1, \dots, x_n) : \tau \bigsqcup_{i=1}^n \ell_i} \text{ TAINT} \quad (79)$$

External input sources introduce taint:

$$\frac{v = \text{external\_input}()}{\Gamma \vdash v : \tau^\top} \text{ SOURCE} \quad (80)$$

## 7.3 Sensitive Sinks and Sanitization

**Definition 7.3** (Sensitive Sinks). The set of sensitive operational sinks in ZELC:

$$\mathcal{S} = \{ \text{firewall.block}, \text{iam.revoke}, \text{iam.disable}, \text{edr.isolate}, \text{edr.kill}, \text{linux.kill}, \text{delete}, \dots \} \quad (81)$$

**Definition 7.4** (Sink Rejection Rule). Tainted arguments to sensitive sinks are rejected at compile time:

$$\frac{\Gamma \vdash v : \tau^\top \quad s \in \mathcal{S}}{\Gamma \vdash s(v) : \text{error}} \text{ SINK-REJECT} \quad (82)$$

**Definition 7.5** (Evidence-Generating Sanitization). Explicit sanitization produces a clean value and an evidence trail:

$$\frac{\Gamma \vdash v : \tau^\top \quad f_{\text{san}} \in \text{Sanitizers}(\tau)}{\Gamma \vdash f_{\text{san}}(v) : \tau^\perp \wedge \text{evidence emit}(san\_rec)} \text{ SANITIZE} \quad (83)$$

The sanitization record includes: the original tainted value's hash, the sanitizer identity, timestamp, and the clean value, enabling post-hoc audit of all data-cleaning decisions.

**Theorem 7.6** (Taint Safety). *In a well-typed ZELC program, no tainted value reaches a sensitive sink without passing through a sanitizer that produces an evidence record.*

*Proof.* By the typing rules: (SINK-REJECT) prevents direct use of tainted values at sinks, and (SANITIZE) is the only rule that converts  $\top \rightarrow \perp$ . Since (SANITIZE) mandates evidence emission, every path from tainted source to sensitive sink must pass through an evidence-generating sanitization step.  $\square$   $\square$

## 8 Agentic AI Safety: Intent-Gated Execution

### 8.1 Threat Model for Agentic Generation

Agentic AI systems may propose incorrect or dangerous actions through hallucination, prompt injection, or misalignment. The threat model includes:

- (a) **Hallucinated actions:** The agent proposes `process.kill` when the intent allows only `service.restart`.

- (b) **Scope escalation:** The agent targets 500 hosts when the intent allows 10.
- (c) **Protected-object violation:** The agent targets production databases when those are in *ExcludeScope*.
- (d) **Temporal violation:** The agent acts outside maintenance windows.

ZELC's guardrails are not prompt-based; they are enforced by the runtime regardless of code provenance.

## 8.2 Conformance Predicate

**Definition 8.1** (Program Conformance). Let  $P$  be a candidate program (human-written or agent-generated). Define:

$$\text{Conforms}(P, \mathcal{I}, \Sigma, \mathcal{C}) \equiv \bigwedge_{a(\vec{v}) \in \mathcal{A}(P)} [\text{Auth}(\Sigma, \mathcal{C}, a(\vec{v})) \wedge \text{Sat}(\mathcal{I}, \Sigma, a(\vec{v})) \wedge \text{Taint-Safe}(a(\vec{v}))] \quad (84)$$

ZELC compilation and execution require Conforms to hold, or *fail closed*.

**Theorem 8.2** (Agent-Agnostic Safety). *For any program  $P$ —whether authored by a human operator or generated by an AI agent—if the ZELC runtime accepts  $P$  for execution, then for every kinetic action  $a(\vec{v}) \in \mathcal{A}(P)$ :*

$$\begin{aligned} \text{Auth}(\Sigma, \mathcal{C}, a(\vec{v})) \wedge \text{Sat}(\mathcal{I}, \Sigma, a(\vec{v})) \\ \wedge \text{Taint-Safe}(a(\vec{v})) \wedge \exists ev_i \in \mathcal{E}(P) \end{aligned} \quad (85)$$

*That is: every executed kinetic action is authorized, intent-compliant, taint-safe, and evidence-generating.*

*Proof.* By Definition 8.1, the runtime checks  $\text{Conforms}(P, \mathcal{I}, \Sigma, \mathcal{C})$  which universally quantifies over all kinetic actions. Each conjunct is verified: Auth by Definition 4.2, Sat by Definition 5.2, and Taint-Safe by Theorem 7.6. Evidence emission follows from the operational semantics (Definition 6.1 and rule KIN-ACCEPT in Section 9), which generates  $ev_i$  for every executed kinetic action.  $\square \quad \square$

## 8.3 Verify Before Execute: Operational Proof Obligations

ZELC imposes proof obligations analogous to proof-carrying code [18]:

$$\mathcal{R}(P, \mathcal{I}, \mathcal{C}, \Sigma) = \begin{cases} \text{ACCEPT} & \text{if } \text{Conforms}(P, \mathcal{I}, \Sigma, \mathcal{C}) \\ \text{REJECT} & \text{otherwise (fail closed)} \end{cases} \quad (86)$$

Note that even *rejected* programs produce evidence records documenting the rejection reason, timestamp, and the principal who attempted execution.

# 9 Operational Semantics

## 9.1 Runtime Configuration

**Definition 9.1** (Configuration). A runtime configuration is a five-tuple:

$$\sigma = \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, S \rangle \quad (87)$$

where  $\Sigma$  is the machine state,  $\mathcal{C}$  is the capability set,  $\mathcal{I}$  is the active intent,  $\mathcal{E}$  is the accumulated evidence trace, and  $S$  is the remaining statement to execute.

## 9.2 Small-Step Transition Rules

The transition relation  $\sigma \rightarrow \sigma'$  is defined by the following rules.

**Rule (SET)** — Variable binding (read-only):

$$\frac{\llbracket E \rrbracket_{\Sigma} = v}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{set } x = E; S \rangle \rightarrow \langle \Sigma[x \mapsto v], \mathcal{C}, \mathcal{I}, \mathcal{E}, S \rangle} \text{ SET} \quad (88)$$

**Rule (WHEN-TRUE):**

$$\frac{\llbracket E \rrbracket_{\Sigma} = \text{true}}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{when } E \{S_1\} \text{ otherwise } \{S_2\}; S \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, S_1; S \rangle} \text{ WHEN-T} \quad (89)$$

**Rule (WHEN-FALSE):**

$$\frac{\llbracket E \rrbracket_{\Sigma} = \text{false}}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{when } E \{S_1\} \text{ otherwise } \{S_2\}; S \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, S_2; S \rangle} \text{ WHEN-F} \quad (90)$$

**Rule (KINETIC-ACCEPT)** — The central rule; requires authorization *and* intent satisfaction:

$$\frac{\text{Auth}(\Sigma, \mathcal{C}, a(\vec{v})) \quad \text{Sat}(\mathcal{I}, \Sigma, a(\vec{v})) \quad \Sigma' = \delta(\Sigma, a(\vec{v})) \quad ev = \text{GenEvidence}(\dots)}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{action } a(\vec{v}); K \rangle \rightarrow \langle \Sigma', \mathcal{C}, \mathcal{I}, \mathcal{E} \| ev, K \rangle} \text{ KIN-ACCEPT} \quad (91)$$

**Rule (KINETIC-REJECT)** — Fail closed with evidence:

$$\frac{\neg \text{Auth}(\Sigma, \mathcal{C}, a(\vec{v})) \vee \neg \text{Sat}(\mathcal{I}, \Sigma, a(\vec{v}))}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{action } a(\vec{v}); K \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E} \| ev_{\text{rej}}, \text{halt} \rangle} \text{ KIN-REJECT} \quad (92)$$

**Rule (EVIDENCE-EMIT)** — Evidence emission within kinetic blocks:

$$\frac{ev = \text{MkEvidence}(title, \llbracket E \rrbracket_{\Sigma}, \Sigma, \mathcal{I})}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{evidence record } title \text{ det } E; S \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E} \| ev, S \rangle} \text{ EV-EMIT} \quad (93)$$

**Rule (EACH-STEP)** — Iteration unfolds one element:

$$\frac{\llbracket E \rrbracket_{\Sigma} = [v_1, v_2, \dots, v_n] \quad n \geq 1}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{each } x \text{ in } E \{S_b\}; S \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, S_b[v_1/x]; \text{each } x \text{ in } [v_2 \dots]; S \rangle} \text{ EACH} \quad (94)$$

**Rule (STOP)** — Hard halt:

$$\frac{}{\langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E}, \text{stop}; S \rangle \rightarrow \langle \Sigma, \mathcal{C}, \mathcal{I}, \mathcal{E} \| ev_{\text{halt}}, \text{halt} \rangle} \text{ STOP} \quad (95)$$

**Invariant 9.2** (Evidence Monotonicity). The evidence trace  $\mathcal{E}$  is monotonically non-decreasing across all transitions:

$$\sigma \rightarrow \sigma' \implies \mathcal{E}' \supseteq \mathcal{E} \quad (96)$$

Evidence is *append-only*; no transition removes or modifies existing evidence records.

## 10 Concrete Operational Programs

We present complete ZELC programs demonstrating that the formal model maps directly to operational reality.

### 10.1 Ransomware Containment

```

1  intent RansomwareContainment {
2      category: containment
3      risk_level: critical
4      constraints {
5          max_targets: 20
6          time_limit: 5m
7          requires_approval: auto
8          exclude_scope: [gateway, dns_primary]
9      }
10 }
11
12 check RansomwareContainment
13     keep critical_hosts = ["prod-db", "prod-web", "prod-api"]
14
15 when ransomware_detected
16     alert critical message "RANSOMWARE DETECTED"
17
18 do
19     set infected = event.hostname
20     set hash = event.file_hash
21
22     -- Isolate: cuts network, leaves mgmt tunnel
23     edr isolate host infected
24
25     -- Kill malicious processes
26     each proc in event.suspicious_processes
27         edr kill process proc.pid on infected
28     end
29
30     -- Quarantine the payload
31     edr quarantine file event.file_path on infected
32
33     -- Forensic snapshot before remediation
34     aws ec2 snapshot instance infected
35
36     -- Escalate if production
37     when infected in critical_hosts
38         pager trigger message "CRITICAL: {infected}"
39     end
40
41     -- Threat intel enrichment
42     set intel = threat lookup hash hash
43
44     -- Evidence with blockchain anchor
45     evidence record "ransomware-incident" details {

```

```

46     host: infected,
47     malware_hash: hash,
48     processes_killed: event.suspicious_processes,
49     threat_intel: intel,
50     containment_time: now()
51   }
52
53   rosecoin anchor evidence_pack "ransomware-{infected}"
54 end
55 end
56 end

```

Listing 2: Ransomware containment playbook with evidence anchoring.

## 10.2 Cloud Key Leak Response

```

1  check CloudKeyLeakResponse
2   keep repo = "rochesteron/zelfire"
3   keep channel = "#soc"
4
5   when access_key_leak or secret_leak
6     alert critical message "Possible cloud key leak detected"
7     notify slack channel channel message "Rotating keys"
8
9   do
10    -- Rotate compromised credentials
11    aws rotate keys
12
13    -- Scan source for additional leaks
14    github scan repo repo for secret_leak
15
16    -- Generate and anchor evidence
17    noodles generate evidence_pack
18    rosecoin anchor evidence_pack "latest"
19
20    evidence record "Key rotation and repo scan" details {
21      repo: repo,
22      timestamp: now()
23    }
24
25    audit log "ZelC run completed" details {
26      action: "rotate+scan",
27      repo: repo
28    }
29  end
30 end
31 end

```

Listing 3: Cloud key leak response with credential rotation.

### 10.3 Brute-Force Defense

```

1  intent BruteForceDefense {
2      category: containment
3      constraints {
4          max_targets: 5
5          time_limit: 1h
6          exclude_scope: [internal_subnets, vpn_gateway]
7      }
8  }
9
10 check BruteForceShield
11     keep threshold = 12
12     keep block_duration = 2 hours
13
14 when bruteforce_login
15     alert high message "Brute force attack detected"
16
17 do
18     set attacker_ip = event.source_ip
19     set target_user = event.username
20
21     -- Block at network edge
22     firewall block ip attacker_ip for block_duration
23
24     -- Lock target account
25     iam lock user target_user
26
27     -- Notify and track
28     notify slack channel "#security-alerts"
29         message "Brute force from {attacker_ip}"
30
31     ticket open title "Brute Force - {attacker_ip}"
32         severity "high"
33
34     evidence record "brute-force-containment" details {
35         attacker_ip: attacker_ip,
36         target_user: target_user,
37         attempts: event.attempt_count,
38         actions_taken: ["ip_blocked", "user_locked"]
39     }
40     end
41 end
42 end

```

Listing 4: Brute-force detection and containment.

### 10.4 CI/CD Pipeline Integrity Verification

```

1  check PipelineGuard(repository_path)
2      set repo = git.open(repository_path)

```

```

3   set head = repo.get_head()
4
5   -- Verify commit signature (GPG)
6   set sig_valid = crypto gpg verify head.signature
7
8   when sig_valid is false
9     alert critical message "UNVERIFIED AUTHOR"
10    do
11      github block merge reason "Unsigned Commit"
12      evidence record "Bad_Actor_Attempt" details {
13        author: head.author,
14        email: head.email
15      }
16      stop
17    end
18  end
19
20  -- Scan for leaked secrets
21  set scan = aina code scan_secrets repo.files
22
23  when scan.secrets_found > 0
24    alert warning message "LEAK: {scan.count} secrets"
25    do
26      notify slack user head.author
27        message "You committed an API Key"
28      stop
29    end
30  end
31
32  otherwise
33    do
34      -- Generate SBOM
35      set sbom = sbom generate format "spdx"
36
37      -- Build distroless container
38      set image = docker build repo base "distroless"
39
40      -- Vulnerability scan
41      set vulns = docker image scan image
42      when vulns.critical_count > 0
43        alert critical message "VULNERABLE DEPENDENCY"
44        stop
45      end
46
47      -- Sign with cosign
48      crypto sign artifact image
49
50      -- Anchor provenance to blockchain
51      rosecoin anchor {
52        type: "Supply_Chain_Release",
53        commit: head.hash,
54        image_hash: image.sha256,
55        sbom_hash: crypto hash sha256 sbom
56      }

```

```

57
58      -- Deploy
59      kube apply image.tag namespace "production"
60      notify slack channel "#releases"
61      message "Deployed {repo.tag} securely"
62  end
63 end
64 end

```

Listing 5: Supply-chain integrity daemon with SLSA attestation.

## 10.5 Zero-Trust Network Gatekeeper

```

1  check ZeroTrustGatekeeper
2  keep MAX_ENTROPY = 7.85
3  keep REQ_PER_SEC = 5000
4  keep GEO_RISK = 85
5
6  set stream = net capture interface "eth0" filter "tcp port 443"
7
8  while true
9    set pkt = stream.next()
10
11   -- Entropy analysis for encrypted payloads
12   set entropy = crypto entropy pkt.payload
13   set verdict = aina net inspect pkt
14
15   when entropy > MAX_ENTROPY or verdict.is_malicious
16     alert critical message "MALWARE: {pkt.src_ip}"
17     do
18       linux firewall block ip pkt.src_ip duration "permanent"
19       net kill_connection pkt.session_id
20
21     evidence record "Payload_Capture" details {
22       src_ip: pkt.src_ip,
23       entropy: entropy,
24       ai_confidence: verdict.confidence
25     }
26
27     rosecoin anchor {
28       type: "Threat_Neutralized",
29       threat_hash: crypto hash sha256 pkt.payload
30     }
31   end
32 end
33
34   -- DDoS rate limiting
35   set rate = net get_flow_rate pkt.src_ip
36   when rate > REQ_PER_SEC
37     alert warning message "DDoS: {pkt.src_ip} @ {rate} rps"
38     do
39       aws waf update_ip_set action "INSERT"

```

```

40         ip pkt.src_ip list "Global_Blocklist"
41     end
42 end
43
44 -- Geofencing
45 set geo = net geoip pkt.src_ip
46 when geo.risk_score > GEO_RISK
47 do
48     net drop pkt
49     audit log "Geo_Block" details {
50         country: geo.country,
51         risk: geo.risk_score
52     }
53 end
54 end
55 end
56 end

```

Listing 6: Zero-trust gatekeeper with entropy analysis and geofencing.

## 11 Architectural Overview

### 11.1 Execution Pipeline

The ZELC runtime processes programs through a multi-stage pipeline, depicted in Figure 1.

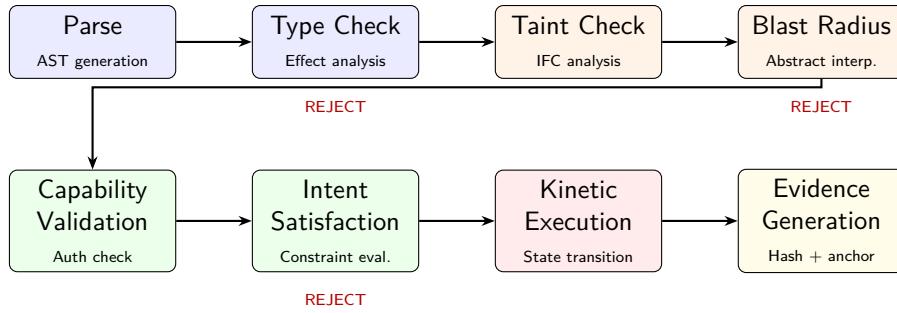


Figure 1: The ZELC execution pipeline. Static phases (top row) precede runtime phases (bottom row). Rejections at any stage produce evidence records and halt execution.

### 11.2 Three-Layer Safety Architecture

ZELC implements defense-in-depth through three orthogonal safety layers:

Table 1: The three layers of safety in ZELC.

Layer	Mechanism	Phase	Prevents
Type Safety	Effect-typed compiler	Compile	Type confusion, kinetic actions outside <b>do</b> blocks
Memory Safety	Scope-bound managed refs	Runtime	Buffer overflows, use-after-free, memory leaks
Kinetic Safety	Intent-gated guard	Runtime	Mass deletion, scope violation, unauthorized escalation

### 11.3 Memory Model

ZELC uses a *scope-bound* memory model without raw pointers:

**Definition 11.1** (Scope-Bound Memory). Memory allocations within a **check** or **do** block are automatically reclaimed when the block exits:

$$\text{exit}(\mathbf{do}) \implies \forall \text{ref} \in \text{locals}(\mathbf{do}) : \text{dealloc}(\text{ref}) \quad (97)$$

No pointer arithmetic, no `malloc/free`, no garbage-collector pauses. This eliminates buffer overflows and memory leaks for operational tasks.

### 11.4 Concurrency Model

ZELC employs an actor model for concurrent execution:

**Definition 11.2** (Actor Isolation). Each **check** block executes in an isolated lightweight actor with:

$$\text{Actor}(\mathbf{check} N) = \langle \Sigma_{\text{local}}, \mathcal{C}_{\text{local}}, \mathcal{E}_{\text{local}}, \text{mailbox} \rangle \quad (98)$$

Actors share no mutable state. Communication is via message passing through the *mailbox*. This enables 1000+ concurrent security checks without CPU thrash or race conditions.

## 12 Tiered Operational Model

ZELC is designed to operate within a four-layer defense stack, depicted in Figure 2.

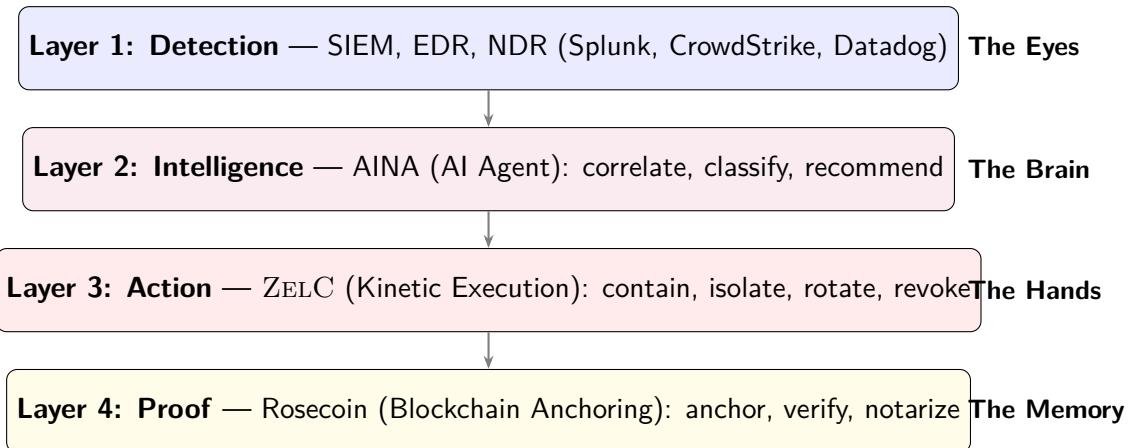


Figure 2: The tiered defense model. ZELC operates at Layer 3, receiving intelligence from Layer 2 and producing evidence for Layer 4.

## 13 Comparison with General-Purpose Languages

Table 2 provides a systematic comparison across the dimensions of the kinetic gap.

Table 2: Systematic comparison: ZELC vs. general-purpose languages along kinetic gap dimensions.

Dimension	GPL (Python/Bash)	ZelC
Execution Model	Unrestricted; any line can mutate state	Read-only by default; mutations only in <code>do</code> blocks
Safety Architecture	None; developer writes all checks	Three-layer: type safety, memory safety, kinetic guard
Permission Model	OS-level ambient authority (root = God)	Capability tokens: explicit, attenuated, time-bounded
Audit Trail	Text logs (mutable, deletable)	Evidence objects (immutable, hash-chained, signed)
AI Safety	Prompt engineering (hope and pray)	Intent validation (compiler + runtime contracts)
Blast Radius	No language-level concept	Compile-time abstract interpretation + runtime re-check
Cloud Integration	External libraries (boto3, azure-sdk)	Native primitives ( <code>cloud.aws</code> , <code>cloud.azure</code> )
Compliance	Manual reports after the fact	Real-time proof generated at runtime
Data Types	Strings, integers, lists	Security-semantic: <b>IP</b> , <b>CIDR</b> , <b>CVE</b> , <b>RiskScore</b> , <b>IOCSet</b>
Taint Tracking	None (developer responsibility)	Language-level IFC with evidence-generating sanitization

## 14 Evaluation Methodology

We propose the following metrics organized by evaluation dimension.

Table 3: Proposed evaluation metrics for ZELC.

ID	Dimension	Description	Formula
$S_1$	Safety	Fraction of unsafe requests rejected	$ \{a : \neg \text{Sat}\}  /  \{a : \text{unsafe}\} $
$S_2$	Safety	False-reject rate	$ \{a : \text{safe} \wedge \neg \text{Sat}\}  /  \{a : \text{safe}\} $
$S_3$	Safety	Time-to-containment (ms)	$t_{\text{contain}} - t_{\text{detect}}$
$E_1$	Evidence	Evidence completeness	See Eq. 99
$E_2$	Evidence	Tamper-detection rate	$ \{m : \text{detected}\}  /  \{m : \text{tampered}\} $
$E_3$	Evidence	Verification time (ms)	$t_{\text{verify}}(\text{inclusion proof})$
$O_1$	Operator	Time-to-understand playbook	$t_{\text{comprehend}} \text{ (sec)}$
$O_2$	Operator	Review error rate	$ \text{missed dangers}  /  \text{total dangers} $
$A_1$	Agentic	Hallucinated actions blocked	$ \{a : \text{hallucinated} \wedge \text{blocked}\}  /  \{a : \text{hallucinated}\} $
$A_2$	Agentic	Permitted task success rate	$ \{t : \text{success}\}  /  \{t : \text{attempted}\} $

The evidence completeness metric is formally:

$$E_1 = \frac{|\{\{a_i \in \mathcal{A}(P) \mid \exists ev_i \in \mathcal{E}(P)\}\}|}{|\mathcal{A}(P)|} \quad (99)$$

Under ZELC’s evidence-first semantics (Proposition 6.7), the design target is  $E_1 = 1.0$ .

The containment-time metric  $S_3$  captures the end-to-end operational velocity. For the ransomware scenario (Listing 2), the target is:

$$S_3 \leq 400 \text{ ms} \quad (\text{detect} \rightarrow \text{isolate} \rightarrow \text{evidence anchored}) \quad (100)$$

## 15 Contribution Claims

To maintain academic rigor, we state ZELC’s contributions as *integration-level* mechanisms. Each claim identifies the novel binding and cites the prior art it builds upon.

**Claim 15.1** (C1: Kinetic Semantics as Primary Operational Effect). A language/runtime model enforcing structural separation between read-only reasoning (RO) and state-changing cybersecurity actuation (KI), with actuation permitted only inside explicit kinetic blocks (**do**...**end**) and prohibited elsewhere, via effect checks (Theorem 3.8) and runtime gating (Rule KIN-ACCEPT). *Prior art*: effect systems [3, 4].

**Claim 15.2** (C2: Capability-Scope Actuation with Formal Non-Escalation). A language invariant binding each kinetic action to explicit capability tokens with proven monotonic non-escalation (Proposition 4.5). *Prior art*: capability security [1, 2].

**Claim 15.3** (C3: Intent Contracts as Enforced Blast-Radius Constraints). A compiler/runtime mechanism interpreting intent declarations as constraint-satisfaction problems (Definition 5.2) with compile-time blast-radius bounding via abstract interpretation (Theorem 5.6). *Prior art*: admission control [5–7], abstract interpretation [9], eBPF verifier [8].

**Claim 15.4** (C4: Evidence-First Semantics with Cryptographic Anchoring). A runtime emitting structured evidence records as a semantic consequence of kinetic execution (Rule KIN-ACCEPT), with hash chaining (Lemma 6.4), signature blocks, and Merkle-tree anchoring (Definition 6.5). *Prior art*: tamper-evident logging [11–13], transparency logs [14, 15].

**Claim 15.5** (C5: Agent-Agnostic Enforcement with Taint Safety). A safety model in which identical constraints are enforced regardless of code authorship (Theorem 8.2), with taint tracking that prevents untrusted data from reaching sensitive sinks without evidence-generating sanitization (Theorem 7.6). *Prior art:* PCC [18], LMQL [19], IFC [10, 16, 17].

## 16 Conclusion

ZELC proposes a cybersecurity-native programming model that treats actuation, authority, governance, and evidence as *first-class semantics*. While each underlying primitive has significant prior art, ZELC’s contribution is an integrated, operator-centered language/runtime that binds these mechanisms into a coherent operational semantics for the security operations center and incident response—especially in an agentic era where “prompt discipline” is not a safety strategy.

The formal model presented herein comprises: (i) an effect-typed kinetic boundary with soundness proof, (ii) a capability algebra with monotonic non-escalation, (iii) intent contracts formalized as constraint-satisfaction problems with blast-radius bounding via abstract interpretation, (iv) cryptographic evidence chains with Merkle-tree anchoring and inclusion-proof verification, (v) taint tracking with evidence-generating sanitization for operational sinks, and (vi) a conformance predicate providing agent-agnostic safety guarantees.

The concrete program listings demonstrate that these formal mechanisms are not abstract constructs but directly executable security operations: ransomware containment, credential rotation, brute-force defense, supply-chain integrity, and zero-trust network gating.

By making kinetic operations explicit, capability-scoped, intent-constrained, and evidentiary by construction, ZELC aims to reduce catastrophic operator error, improve auditability, and enable safe automation at scale—closing the kinetic gap that has separated security intent from operational execution for decades.

## A ZelC Type Universe: Complete Reference

Table 4 enumerates the complete type universe.

Table 4: Complete ZELC type universe.

Category	Type	Description
Primitive	<code>unit</code> , <code>Bool</code> , <code>Int</code> , <code>Float</code> , <code>Text</code>	Standard value types
Temporal	<code>Duration</code> , <code>Timestamp</code>	Time-aware operations
Network	<code>IP</code> , <code>CIDR</code> , <code>Port</code> , <code>MAC</code> , <code>URL</code>	Network address types
Identity	<code>UserId</code> , <code>GroupId</code> , <code>ServiceAccount</code>	Identity primitives
Crypto	<code>Hash</code> , <code>Signature</code> , <code>PublicKey</code> , <code>Certificate</code>	Cryptographic types
Threat	<code>CVE</code> , <code>IOCSet</code> , <code>YaraRule</code> , <code>SigmaRule</code>	Threat intelligence
Risk	<code>RiskScore</code> , <code>Severity</code> , <code>Priority</code>	Risk quantification
Incident	<code>Incident</code> , <code>Ticket</code> , <code>Timeline</code>	Case management
Evidence	<code>EvidencePack</code> , <code>AuditRecord</code> , <code>ChainOfCustody</code>	Forensic types
Compliance	<code>ControlId</code> , <code>Baseline</code> , <code>GapReport</code>	GRC mapping
Cloud	<code>Instance</code> , <code>Bucket</code> , <code>SecurityGroup</code> , <code>VPC</code>	Cloud resource types
Container	<code>Pod</code> , <code>Container</code> , <code>Image</code> , <code>SBOM</code>	Container ecosystem
Compound	<code>List</code> [ $\tau$ ], <code>Map</code> [ $\tau_k$ , $\tau_v$ ], <code>Record</code> {...}	Generic containers
Function	$\tau_1 \rightarrow^{\epsilon} \tau_2$	Effect-annotated functions

## B ZelC Action Taxonomy

Table 5 classifies the kinetic action vocabulary by domain.

Table 5: Kinetic action taxonomy (representative subset).

Domain	Action	Signature
Firewall	<code>firewall.block</code>	$IP \times Duration \rightarrow^{KI} unit$
EDR	<code>edr.isolate</code>	$Text \rightarrow^{KI} unit$
EDR	<code>edr.kill</code>	$Int \times Text \rightarrow^{KI} unit$
EDR	<code>edr.quarantine</code>	$Text \times Text \rightarrow^{KI} unit$
IAM	<code>iam.revoke</code>	$UserId \rightarrow^{KI} unit$
IAM	<code>iam.lock</code>	$UserId \rightarrow^{KI} unit$
AWS	<code>aws.rotate_keys</code>	$UserId \rightarrow^{KI} unit$
AWS	<code>aws.ec2.snapshot</code>	$Instance \rightarrow^{KI} Hash$
AWS	<code>aws.s3.block_public</code>	$Bucket \rightarrow^{KI} unit$
Azure	<code>azure.ad.block_user</code>	$UserId \rightarrow^{KI} unit$
Azure	<code>azure.nsg.deny</code>	$Text \times IP \rightarrow^{KI} unit$
GCP	<code>gcloud.iam.disable</code>	$ServiceAccount \rightarrow^{KI} unit$
K8s	<code>kube.isolate_pod</code>	$Pod \rightarrow^{KI} unit$
Docker	<code>docker.stop</code>	$Container \rightarrow^{KI} unit$
Linux	<code>linux.kill</code>	$Int \rightarrow^{KI} unit$
Linux	<code>linux.service.restart</code>	$Text \rightarrow^{KI} unit$
Crypto	<code>crypto.sign</code>	$Text \rightarrow^{KI} Signature$
Rosecoin	<code>rosecoin.anchor</code>	$EvidencePack \rightarrow^{KI} Hash$

Note that *every* action in this table carries effect  $KI$  and therefore can only appear inside a `do` block, subject to capability authorization and intent satisfaction.

## C Formal Proof: Evidence Chain Integrity

**Theorem C.1** (Evidence Chain Integrity). *Let  $\mathcal{E} = \langle ev_1, \dots, ev_n \rangle$  be the evidence trace. Let  $\langle h_0, \dots, h_n \rangle$  be the hash chain (Definition 6.3),  $R = \text{MerkleRoot}(h_1, \dots, h_n)$ , and  $\langle sig_1, \dots, sig_m \rangle$  the signature checkpoints. The following properties hold under standard cryptographic assumptions:*

- (i) **Completeness:** Every kinetic action has a corresponding evidence record ( $E_1 = 1.0$ ).
- (ii) **Tamper evidence:** Modification of any  $ev_j$  is detectable (Lemma 6.4).
- (iii) **Non-repudiation:** For each signed checkpoint  $sig_j$ , the signer cannot deny having seen  $h_{t_j}$ .
- (iv) **Append-only:** Deletion of any  $ev_j$  is detectable via the hash chain break.
- (v) **Verifiable inclusion:** For any  $ev_k$ , an auditor can verify inclusion via `VerifyInclusion` in  $O(\log n)$  time.

*Proof.* (i) follows from Proposition 6.7 and the operational semantics (rules KIN-ACCEPT and KIN-REJECT both emit evidence). (ii) follows from Lemma 6.4. (iii) follows from the unforgeability of digital signatures under the chosen scheme. (iv) follows because deletion of  $ev_j$  removes  $h_j$  from the chain, causing  $h_{j+1}$  verification to fail. (v) follows from the standard Merkle-tree inclusion proof: the proof path  $\pi_k$  has length  $\lceil \log_2 n \rceil$ , and verification requires  $\lceil \log_2 n \rceil$  hash computations.  $\square \quad \square$

## References

- [1] J. B. Dennis and E. C. Van Horn. Programming semantics for multiprogrammed computations. *Communications of the ACM*, 9(3):143–155, 1966.
- [2] M. S. Miller. *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control*. Ph.D. thesis, Johns Hopkins University, 2006.
- [3] J. M. Lucassen and D. K. Gifford. Polymorphic effect systems. In *Proceedings of the 15th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 47–57, 1988.
- [4] D. Leijen. Koka: Programming with row polymorphic effect types. *Electronic Proceedings in Theoretical Computer Science*, 153:100–126, 2014.
- [5] Kubernetes Contributors. Admission controllers reference. <https://kubernetes.io/docs/reference/access-authn-authz/admission-controllers/>, 2024.
- [6] Open Policy Agent Contributors. OPA Gatekeeper: Policy and governance for Kubernetes. <https://open-policy-agent.github.io/gatekeeper/>, 2024.
- [7] HashiCorp. Sentinel: Policy as code framework. <https://www.hashicorp.com/sentinel>, 2024.
- [8] Linux Kernel Contributors. eBPF verifier. <https://docs.kernel.org/bpf/verifier.html>, 2024.
- [9] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*, pages 238–252, 1977.

- [10] D. E. Denning. A lattice model of secure information flow. *Communications of the ACM*, 19(5):236–243, 1976.
- [11] B. Schneier and J. Kelsey. Secure audit logs to support computer forensics. *ACM Transactions on Information and System Security*, 2(2):159–176, 1999.
- [12] R. T. Snodgrass, S. S. Yao, and C. Collberg. Tamper detection in audit logs. In *Proceedings of the 30th International Conference on Very Large Data Bases (VLDB)*, pages 504–515, 2004.
- [13] J. Kelsey, J. Callas, and A. Clemm. Signed syslog messages. RFC 5848, Internet Engineering Task Force, 2007.
- [14] B. Laurie, A. Langley, and E. Kasper. Certificate Transparency. RFC 6962, Internet Engineering Task Force, 2013.
- [15] Sigstore Contributors. Rekor: Transparency log for software supply chains. <https://docs.sigstore.dev/logging/overview/>, 2024.
- [16] W. Enck, P. Gilbert, S. Han, V. Tendulkar, B.-G. Chun, L. P. Cox, J. Jung, P. McDaniel, and A. N. Sheth. TaintDroid: An information-flow tracking system for realtime privacy monitoring on smartphones. *ACM Transactions on Computer Systems*, 32(2):5:1–5:29, 2014.
- [17] A. C. Myers. JFlow: Practical mostly-static information flow control. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 228–241, 1999.
- [18] G. C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 106–119, 1997.
- [19] L. Beurer-Kellner, M. Fischer, and M. Vechev. Prompting is programming: A query language for large language models. In *Proceedings of the ACM SIGPLAN 2023 Conference on Programming Language Design and Implementation (PLDI)*, pages 1507–1520, 2023.
- [20] Amazon Web Services. Amazon Quantum Ledger Database (QLDB). <https://aws.amazon.com/qldb/>, 2024.