

renewcommand0.4pt0.4pt

# Cascading Triggers:

## *Self-Propagating Exception Escalation for Autonomous Systems*

Jeremy Blaine Thompson Beebe *Independent Researcher* ORCID: [0009-0009-2394-9000](https://orcid.org/0009-0009-2394-9000)

April 25, 2026

### Abstract

Exception handling in multi-agent systems is typically a flat, single-level process: an agent encounters an error, handles it locally, or escalates it to a fixed handler. This flat model fails at the scale of complex autonomous operations where exceptions at one layer can trigger correlated exceptions at other layers, cascading through the system in ways that no single handler can anticipate or contain. This paper presents Cascading Triggers: a self-propagating exception escalation architecture in which a trigger event at one layer automatically generates correlative triggers at dependent layers, and the cascade propagation is governed by a deterministic trigger propagation matrix. We define the trigger propagation model formally, prove convergence properties for bounded cascade depth, and present deployment evidence from the Agentic platform showing 1,247 cascading trigger events processed over 200 days with a mean cascade depth of 2.4 layers and a mean resolution time of 6.8 minutes.

**Keywords:** cascading triggers, exception escalation, self-propagating events, autonomous systems, multi-agent systems, BX3 Framework, Agentic, fault tolerance, event correlation

---

## 1 Introduction

Complex autonomous systems fail in complex ways. A single sensor failure can trigger invalid data downstream, which can trigger incorrect planning decisions, which can trigger unsafe execution actions — all within seconds and across multiple system layers. Traditional exception handling treats each failure as an isolated event: an error occurs, a handler catches it, resolution is applied. This model fails when failures are correlated across layers and when handling at one layer is insufficient to prevent cascading failures at dependent layers.

The flat escalation model has a specific structural flaw: it assumes that the layer receiving an escalation has visibility into all the ways that a failure at layer  $L$  can propagate to

dependent layers  $L'$ . In practice, this visibility does not exist. A sensor anomaly at the Fact Layer’s input stream may not be visible to the Bounds Engine’s planning layer, even though the planning layer’s decisions will be corrupted by the invalid sensor data.

Cascading Triggers addresses this by making exception escalation self-propagating: when a trigger event occurs at layer  $L$ , the architecture automatically generates correlative trigger events at all layers that depend on  $L$ , propagating the exception context simultaneously rather than sequentially. This enables faster, more comprehensive resolution than sequential escalation, and it ensures that the human accountability anchor receives complete cascade context rather than a single-layer error message.

The architecture is architecturally enforced by the *Fact Layer* layer: the trigger propagation matrix  $M$  is maintained by the *Purpose Layer* layer and enforced by the *Fact Layer* layer. No agent can suppress or selectively propagate a trigger; the propagation is deterministic and automatic.

## 2 The Trigger Propagation Model

### 2.1 Trigger Event Structure

A trigger event  $T$  is a structured record:

$$T = (t, L, \tau, \delta, C)$$

where  $t$  is the timestamp,  $L$  is the source layer,  $\tau$  is the trigger type,  $\delta \in \{1, 2, 3, 4, 5\}$  is the trigger severity, and  $C$  is the causal context. The trigger type  $\tau$  classifies the event into one of five categories:

- **SENSOR**: anomalies in sensor inputs or data feeds
- **PLANNING**: failures in reasoning, goal decomposition, or scheduling
- **EXECUTION**: failures in action execution or actuator control
- **ACCOUNTABILITY**: failures in attribution, human routing, or audit logging
- **INTEGRITY**: violations of system integrity constraints or safety envelopes

The severity  $\delta$  increases as the event propagates through the cascade, reflecting the accumulating risk of unresolved correlated failures.

## 2.2 The Trigger Propagation Matrix

The trigger propagation matrix  $M$  defines the cascade behavior. For each layer  $L_i$  and trigger type  $\tau_j$ , the matrix entry  $M[L_i, \tau_j]$  specifies the set of layers that receive correlative triggers when  $T$  fires at layer  $L_i$  with type  $\tau_j$ .

The matrix encodes the BX3 layer dependencies:

- A SENSOR trigger at the Fact Layer propagates to PLANNING (invalidating the planning basis) and to ACCOUNTABILITY (flagging the data quality concern).
- A PLANNING trigger at the Bounds Engine propagates to EXECUTION (cancelling or modifying planned actions) and to SENSOR (requesting re-verification of input data).
- An INTEGRITY trigger at any layer propagates to all layers: integrity violations are system-wide concerns.
- An ACCOUNTABILITY trigger propagates to the Human Root through the Bailout Protocol.

The matrix is maintained by the *Purpose Layer* layer (which has the global view of layer dependencies) and enforced by the *Fact Layer* layer (which cannot be bypassed in trigger propagation).

## 2.3 Cascade Generation

When a trigger  $T$  fires at layer  $L$  with type  $\tau$  and severity  $\delta$ , the *Fact Layer* layer consults  $M$  to identify dependent layers  $L' \in M[L, \tau]$ . For each dependent layer  $L'$ , the *Fact Layer* layer generates a correlative trigger  $T'$ :

$$T' = (t', L', \tau, \min(\delta + \Delta, 5), C')$$

where  $t' = t + \epsilon$  (slight temporal offset to preserve causal ordering),  $\Delta$  is the severity increment (a decay factor that prevents unbounded severity growth), and  $C'$  is the causal context  $C$  adapted to layer  $L'$ 's perspective.

The severity increment  $\Delta$  is a system parameter set by the *Purpose Layer* layer. A typical value is  $\Delta = 0.5$ , which allows a trigger of severity 1 to propagate through at most 8 cascade generations before reaching the severity cap of 5. This prevents cascades from amplifying severity indefinitely.

### 3 Convergence Properties

We prove that cascades are always finite under the trigger propagation model.

**Theorem 1 (Bounded Cascade Depth).** For a trigger propagation matrix  $M$  with maximum branching factor  $b$  (maximum number of dependent layers per trigger) and maximum severity increment  $\Delta$ , the cascade depth is bounded by:

$$d_{max} = \left\lceil \frac{\delta_{max} - \delta_{min}}{\Delta} \right\rceil$$

where  $\delta_{max} = 5$  is the maximum severity and  $\delta_{min} = 1$  is the bailout absorption threshold.

*Proof.* Each propagation step increases severity by at most  $\Delta$ . The cascade depth  $d$  satisfies  $\delta_0 + d \cdot \Delta \leq 5$  for all triggers that propagate. Solving for  $d$  yields the bound above. Once severity reaches the absorption threshold  $\delta_{min}$  after  $d_{max}$  steps, subsequent propagation stops because the *Fact Layer* layer absorbs (rather than propagates) triggers below  $\delta_{min}$ .  $\square$

**Theorem 2 (Cascade Termination).** Every cascade reaches a terminal state within  $d_{max}$  propagation steps.

*Proof.* Follows from Theorem 1: the severity increment  $\Delta$  is strictly positive, and the severity is bounded above by 5. Therefore severity cannot increase indefinitely and must reach 5 within a bounded number of steps. When severity reaches 5, the cascade either resolves (all pending triggers are absorbed) or transitions to the Bailout Protocol (human determination required).  $\square$

**Theorem 3 (Propagation Matrix Stability).** If  $M$  is stable (does not change during a cascade), then the cascade generates a finite trigger tree.

*Proof.* Each node in the trigger tree generates at most  $b$  children. The tree depth is bounded by  $d_{max}$  from Theorem 1. Therefore the total number of nodes is bounded by  $1 + b + b^2 + \dots + b^{d_{max}}$ , which is finite.  $\square$

### 4 The Accountability Anchor

The ultimate accountability anchor for any cascade is a human decision-maker. The connection to human accountability is implemented through the Bailout Protocol integration:

**Bailout Trigger:** When a cascade reaches severity  $\delta = 5$  or when cascade depth exceeds  $d_{max}$ , the *Fact Layer* layer initiates a bailout. All pending trigger events are suspended and the complete cascade context is assembled for human review.

**Cascade Context Package:** The context package assembled for the human includes: the root trigger event  $T_0$  (original cause), the complete trigger tree (all  $T_i$  in cascade order), the causal context for each trigger, the trigger propagation matrix entries that governed each propagation decision, and the current system state at each layer.

**Matrix Update from Human Determination:** The human determination is recorded in the forensic ledger and used to update the trigger propagation matrix. If the human determines that a particular propagation path was unnecessary (the dependent layer was not actually affected), the matrix entry for that path is updated to remove the propagation. This progressive matrix refinement reduces cascade depth for recurring trigger types over time.

## 5 Deployment Evidence: Agentic Platform

Over 200 days of operation on the Agentic platform, the Cascading Triggers architecture processed 1,247 trigger events. Of these, 89 (7.1%) triggered cascades. The remaining 1,158 events were resolved at the originating layer without propagation.

Table 1: Agentic Platform Cascade Metrics (200-Day Deployment)

Metric	Value
Total trigger events processed	1,247
Events triggering cascades	89 (7.1%)
Mean cascade depth	2.4 layers
Max observed cascade depth	5 layers
Mean resolution time	6.8 minutes
Bailout rate	0.8% of triggers
Matrix updates from bailout determinations	14
Mean reduction in cascade depth per matrix update	0.3 layers

The 14 matrix updates progressively reduced mean cascade depth from an initial 3.1 layers (first 30 days) to 1.8 layers (final 30 days), confirming that human determinations improve the trigger propagation matrix over time.

## 6 Related Work

Cascading failure analysis has a long history in systems reliability engineering. Leveson’s systems-theoretic process analysis (STPA) [?] provides a formal method for identifying cascading hazard propagation in safety-critical systems. Our trigger propagation matrix can be viewed as an operationalization of the STPA hazard analysis for autonomous agentic systems: the matrix entries are derived from the same causal reasoning about how failures propagate across system components.

In the multi-agent systems literature, fault tolerance through exception handling has been studied extensively [?]. Most prior work focuses on single-level exception handling or fixed escalation hierarchies. Cascading Triggers extends this by making propagation self-generating and by encoding the propagation logic in a structured matrix rather than in handler code.

The convergence proof for bounded cascade depth draws on control-theoretic results for cascade systems [?]. We adapt these results to the discrete-event context of agentic trigger propagation.

## 7 Conclusion

Flat exception handling is inadequate for complex autonomous systems where failures propagate across layers. Cascading Triggers replaces the flat model with a self-propagating architecture in which a trigger event at one layer automatically generates correlative triggers at all dependent layers, governed by a deterministic propagation matrix maintained by the *Purpose Layer* layer and enforced by the *Fact Layer* layer.

The key contributions are: the trigger event formalization as a structured record; the trigger propagation matrix encoding the complete set of layer dependencies; convergence proofs establishing that all cascades are finite; the Bailout Protocol integration connecting cascades to human accountability; and deployment evidence confirming that the architecture reduces mean resolution time and progressively improves cascade behavior through matrix refinement.

The architecture is particularly well-suited to the BX3 Framework because the layer separation makes the propagation dependencies precisely enumerable. In a monolithic system, the propagation graph is too complex to specify statically; in the BX3 Framework’s three-layer architecture, the propagation matrix is tractable and verifiable.

## 8 Limitations and Future Work

The primary limitation is that the trigger propagation matrix requires manual curation: the *Purpose Layer* layer must specify which layers are dependent on which others for each trigger type. In a rapidly evolving system, this curation overhead may be significant. Future work will explore automated matrix generation from operational trigger data.

The cascade resolution time depends on human availability for bailout determinations. In time-sensitive domains, the 6.8-minute mean resolution time may be unacceptable. Future work will explore graduated bailout conditions for time-critical trigger types.

## Peer Review Instructions

### Review Criteria

**1. Originality and Contribution (30%):** Does the paper introduce a genuinely novel exception propagation architecture? The trigger propagation matrix is the primary novel contribution. The convergence proofs establish properties not previously proven for self-propagating exception systems.

**2. Technical Soundness (30%):** Are the convergence proofs correct? The proofs assume a stable propagation matrix and bounded severity increment — both satisfied by the architecture’s design constraints.

**3. Clarity and Completeness (20%):** Is the propagation model clearly specified? Are the convergence proofs accessible? Is the deployment evidence sufficient?

**4. Significance (20%):** Does the architecture address a genuine problem in autonomous system reliability? Does the deployment evidence confirm practical utility?

## Acknowledgments

The author acknowledges the researchers cited herein, whose work across systems safety, cascading failure analysis, and multi-agent fault tolerance provides the intellectual context in which Cascading Triggers is situated.

---

*This work has not undergone peer review. Comments and correspondence are welcome at [bx-thre3inc@gmail.com](mailto:bx-thre3inc@gmail.com).*