CASCADE:
Cross-Agent Supply Chain Audit and Defense Ecosystem

A Novel Framework for Systematic Risk Assessment in Autonomous AI Dependency Networks

Owen Sakawa
Elloe AI Research Lab · February 2026

Abstract

CASCADE-AI is a whitepaper on dependency risk in autonomous systems. It proposes a graph-based representation of agents, models, tools, data sources, and policy controls so that organizations can reason more clearly about where upstream failures propagate downstream.

The central argument is that model-level testing does not fully capture systemic exposure. Teams need a way to represent relationships, simulate cascade effects, and compare how risk moves through a live AI environment.

At a glance

- **Model:** CASCADE-AI treats an AI system as a dependency graph made up of models, agents, tools, data sources, and control points.
- **Question:** The whitepaper is trying to answer a practical governance problem: how do teams measure systemic risk when failures propagate across linked components rather than appearing in one model alone?
- **Use:** It is most useful as a framing tool for risk reviews, simulation exercises, and discussions about where operational controls belong in a more complex AI stack.

Why dependency risk needs its own model

Autonomous systems are rarely single artifacts. They are stacks of models, tools, retrieval systems, memory layers, policies, and external services. When something fails upstream, the downstream consequences can be difficult to see without a system view.

CASCADE-AI is useful because it shifts the frame from isolated model risk to connected system risk. That makes it more relevant for governance teams dealing with orchestration, tool use, and multi-component workflows.

How CASCADE represents systems

The whitepaper models systems as directed graphs. Nodes represent models, agents, tools, data sources, infrastructure, or policy controls. Edges represent dependency relationships with weights attached to them.

That structure supports a simple but practical move: treating propagation as measurable. Instead of asking whether a component is risky in the abstract, teams can ask how much exposure it creates once the rest of the system depends on it.

- Dependency graphs for connected AI systems
- Weighted edges to reflect coupling strength
- Risk scoring and simulation for scenario analysis

How it helps governance and assurance teams

CASCADE-AI is best understood as a governance aid rather than a replacement for evaluation. It helps teams identify where controls should sit, which upstream nodes deserve more scrutiny, and how scenario exercises can reveal systemic exposure before an incident forces the issue.

That makes it especially useful for regulated or operationally sensitive deployments, where the real risk often comes from interaction effects rather than a single isolated failure.

Source note: This curated PDF is derived from the public summary prepared for owensakawa.com and is intended as a clean download companion to the indexable research page.