Alright — here's the *real* bridge: cascades are not just "like" phase transitions; many cascade models **are** phase transitions in disguise. And once you see that, renormalization group (RG) becomes the natural language for "lock-in," "recoverability collapse," and "interaction-order geometry."

I'll give you (1) the canonical cascade→phase-transition correspondences, (2) the RG picture, (3) the "distance to tipping" observables you can actually compute, and (4) how your five-node variables slot into that formally.

---

1) Cascades as phase transitions: the core equivalence

The universal structure

Most cascade systems can be written as dynamics on a graph with a control parameter that tunes propagation:

- **control**: how easily activation spreads (coupling, load, infection rate, adoption rate…)
- **state**: fraction of nodes activated/failed/committed
- **threshold**: a critical point where propagation becomes self-sustaining

A minimal form:

A minimal form:

$$x_{t+1} = F\left(x_t ; \theta\right)$$

where x_t is "activated fraction" and θ is the propagation control. A phase transition happens when $x=0$ changes stability.

The condition is the same across models:

$$\left.\frac{\partial F}{\partial x}\right|_{x=0} = 1$$

That's the critical boundary: below it perturbations die; above it they grow.

2) Canonical examples that show "cascade = criticality"

(A) Branching / reproduction number cascades

Early cascade growth is a branching process. Let m be expected "children per activation."

- $m<1$: subcritical → dies out

- $m = 1$: critical → power-law cascade sizes
- $m > 1$: supercritical → macroscopic cascade probability ¿ 0

That's a phase transition, and it has RG structure (below).
Observable: estimate an effective reproduction number m_{eff}.
Distance to criticality: $\delta = 1 - m_{\text{eff}}$.

(B) Percolation (connectivity-driven cascades)

If cascades require paths through a vulnerable subgraph, you get percolation.
Control: effective occupation probability p (or vulnerable fraction).
Order parameter: giant connected vulnerable cluster size.
Critical point p_c where a giant component appears. That's literally a statistical physics phase transition.
Observable: size of largest vulnerable component, susceptibility (cluster-size variance), correlation length proxy.

(C) Threshold cascades

Nodes flip if enough neighbors are active.

$$s_i(t+1) = 1\left(\sum_j A_{ij} s_j(t) \geq \phi_i\right)$$

This has a global cascade condition tied to network degree distribution + threshold distribution. It exhibits critical boundaries and universality classes depending on heterogeneity.

Observable: spectral measures + local threshold distribution → predicts cascade boundary.

(D) Sandpile / self-organized criticality (SOC)

Load accumulates until local thresholds trigger avalanches. Many "pressure-buffer" systems live here. SOC naturally produces near-critical dynamics; *buffer erosion* pushes you from stable to avalanche-dominated regimes.

Observable: avalanche size distribution, power-law tail emergence, finite-size scaling.

3) The RG picture: why "lock-in" is a fixed point

Renormalization group is: **coarse-grain the system and see how parameters flow**.

Coarse-grain a cascade

Take a graph. Group nodes into blocks of size b. Define a block state "failed if $\geq k$ nodes failed" (or similar).

The microscopic cascade rule induces an effective cascade rule at the block scale. That defines an RG map:

$$\theta' = R_b(\theta)$$

- θ could be infection rate, coupling strength, effective reproduction number, etc.
- θ' is the "same" parameter after coarse-graining.

Fixed points

A fixed point satisfies:

$$\theta^* = R_b(\theta^*)$$

Interpretation:

- stable fixed point: the system flows toward it under coarse-graining → robust regime
- unstable fixed point: separatrix / critical point → boundary between regimes
- runaway flow: system flows toward "fully cascading" attractor →lock-in

So lock-in is literally "RG flow into the cascading basin of attraction."

Why phase transitions feel sudden

Near the unstable fixed point, small changes get amplified across scales:

$$|\theta' - \theta^*| \approx b^y |\theta - \theta^*|$$

where $y > 0$ is a relevant exponent. That's the formal reason tiny parameter shifts can cause macroscopic regime change.

4) What becomes measurable: distance-to-criticality indicators

Once you treat cascades as critical phenomena, you inherit a powerful toolkit.

(A) Susceptibility

Response to a small perturbation:

$$\chi = \frac{\partial \langle x \rangle}{\partial h}$$

where h is "seed pressure" (external shock). Near criticality, χ blows up.

In practice: a tiny perturbation causes disproportionately large response.

(B) Correlation length

Critical systems develop long-range correlations:

$$\xi \sim |\theta - \theta_c|^{-\nu}$$

In operational telemetry, ξ shows up as: disturbances propagating farther before damping out.

(C) Critical slowing down

Recovery time diverges near the transition:
Linearize near equilibrium:

$$x_{t+1} - x^* \approx \lambda\left(x_t - x^*\right)$$

As $\lambda \to 1$, return time explodes. Observable early warnings:

- rising autocorrelation
- rising variance
- slower decay after shocks

This directly matches your "recoverability narrowing."
(D) Power-law cascades

At criticality, cascade size distribution often follows:

$$P(S) \sim S^{-\tau}$$

If your datasets show a shift from exponential tails to heavy tails, that's a classic signature of approaching criticality.

5) Now map this onto Clarus five-node geometry (formally)

Your five-node variables are almost a direct decomposition of the control parameters + dynamical responses that determine distance to criticality.

Clarus → criticality correspondences

Clarus variable	Physics / cascade analogue	Role in transition
pressure	external field / driving / seed rate h	pushes system toward activation
buffer	local thresholds / dissipation / margin	stabilizes, sets avalanche containment
lag	relaxation time / feedback delay	governs critical slowing & overshoot
coupling tightness	interaction strength / branching factor / effective transmission β	sets propagation amplification
cascade state	order parameter x / phase label	stable vs cascading regime

This is the key: phase transition is controlled not by one knob but by a composite effective parameter.
A minimal composite could be:

$$\theta_{\text{eff}} = \underbrace{C}_{\text{coupling}} \cdot \underbrace{P}_{\text{pressure}} \cdot \underbrace{\underbrace{g(L)}_{\text{lag gain}}/B}_{\text{buffer}}$$

Then a transition boundary is:

$$\theta_{eff} \approx \theta_c$$

This is not yet "the truth," but it's the right structural form: amplification divided by dissipation.

This is the key: phase transition is controlled not by one knob but by a composite effective parameter.
A minimal composite could be:

$$\theta_{\text{eff}} = \underbrace{C}_{\text{coupling}} \cdot \underbrace{P}_{\text{pressure}} \cdot \underbrace{\underbrace{g(L)}_{\text{lag gain}}/B}_{\text{buffer}}$$

Then a transition boundary is:

$$\theta_{eff} \approx \theta_c$$

This is not yet "the truth," but it's the right structural form: amplification divided by dissipation.

6) “Lock-in” = hysteresis = first-order transition structure

Many real cascades exhibit hysteresis:

- increasing pressure triggers cascade at $\theta_c^{+¿¿}$
- decreasing pressure does not immediately restore stability; recovery happens at θ_c^{-}

That gap is the "irreversibility window."
This is the formal backbone of your "recoverability collapse" language.
Hysteresis emerges when:

- feedback is delayed (lag)
- coupling is strong
- buffer restoration is slow
- redundancy/recording occurs (your copying analogy)

So "lock-in timing" becomes: the moment the system crosses into the basin where the reverse path is unavailable at current buffer/lag.

2. Where κ and ε can become mathematically real

If you want K/ε to survive contact with serious math, anchor them to standard objects.
κ as "coherence" ↔ correlation structure / recoverability
In cascades, a natural κ analogue is inverse correlation length or spectral gap.

- Let J be the Jacobian (linearized dynamics) or an effective propagation operator.
- The largest eigenvalue $\lambda_{\max}$ controls whether perturbations grow.

Define:

$$\kappa \equiv 1 - \lambda_{max}$$

- $\kappa > 0$: stable (returns)
- $\kappa \to 0^{+¿¿}$: critical (slow recovery)
- $\kappa < 0$: unstable (propagation)

That's a clean, defensible coherence/stability indicator.
ε as "uncertainty distribution" ↔ where variance lives
A defensible ε analogue is the distribution of variance / influence across nodes/links. For example:

- use per-node contributions to risk (Shapley / influence functions / gradient norms)
- or distribution of local branching factors m_i

A "spread" of ε corresponds to risk becoming distributed (systemic), rather than localized (containable).

8) What this buys you (and why it's a big deal)

If you adopt this framing, you can legitimately claim:

- Cascade boundaries are **critical manifolds** in parameter space.
- "Lock-in" corresponds to **hysteretic transitions** / basin capture.
- "Early warning drift" corresponds to **critical slowing down** and **rising susceptibility**.
- Interaction order corresponds to **dimensionality of the relevant RG flow** (how many coupled knobs are actually driving you).

And crucially:

You can compute real indicators from telemetry:

- largest eigenvalue (distance to instability)
- return-time / autocorrelation
- susceptibility to injected shocks

- heavy-tail emergence in event sizes
- finite-size scaling across subsystems

That's "mathematically real" in the sense physicists and complex-systems people recognize.

---

9) A crisp way to say it in one paragraph

Cascades are phase transitions on networks. Renormalization explains why: when you coarse-grain a propagation process, you get parameter flow toward stable or cascading fixed points. The "lock-in" you care about is basin capture or hysteresis — the regime where reversing the inputs no longer reverses the state. Your five-node variables decompose the effective control parameter and the dynamical response terms that govern distance to the critical manifold. That makes "recoverability" a measurable quantity via spectral gaps, susceptibility, correlation length proxies, and critical slowing down signals.

If you want, I can do the next step: **choose one of your dataset families** (e.g., privilege escalation, multi-agent coordination, deployment drift) and write down a concrete "minimum model" whose control parameters map exactly onto your five nodes, and show what the critical boundary and hysteresis loop look like. That turns this from theory into a directly testable mathematical scaffold.