10.48047/jocaaa.2026.35.06.27

Early Warning for Catastrophic Infrastructure Collapse: Critical Slowing Down Theory for Reliability Phase Transitions

Satya Sampreeth Boppudi†, Saketh Gowerneni‡, Pradeep Kandepaneni§, Praneeth Ganta¶, Murali M. Chittabathini*

†Independent Researcher, USA, sampreeth.boppudi@gmail.com
‡Independent Researcher, USA, gsaketh369@gmail.com
§Independent Researcher, USA, pradeepkandepaneni@gmail.com
¶Independent Researcher, USA, praneethganta@gmail.com
*Independent Researcher, USA, murali.mohan.rao@gmail.com

Abstract—Production infrastructure rarely fails all at once; it tips. Post-mortems repeatedly show that the signs were present hours before collapse, recovery from small perturbations slowing, latency variance widening, cross-service correlations tightening, yet reliability tooling reacts only once something has visibly broken. This paper introduces RPT-EWS, a framework that applies critical slowing down theory from dynamical systems to infrastructure reliability. Four early-warning signals that theory shows must precede any tipping point, recovery-rate decay, variance inflation, rising autocorrelation, and growing spatial coherence, are computed from telemetry teams already collect, combined into a composite order parameter, and used to raise an alarm and classify the approaching bifurcation so the right intervention is chosen. The cited mathematics is established and validated across ecology, finance, and climate science; the performance figures reported here are projected design targets illustrated on sample scenarios, not measured outcomes, and empirical validation on production telemetry is left to future work. The contribution is the mapping of critical-transition theory onto observable SRE metrics, the composite order parameter, the bifurcation classifier, and a comparative analysis against threshold and burn-rate alerting.

Index Terms—site reliability engineering, critical slowing down, tipping points, bifurcation, early warning signals, phase transition, observability, autocorrelation, order parameter, cascade failur

Introduction

Every reliability tool in current production use shares one assumption: something must visibly go wrong before the system can respond. An error rate crosses a threshold, a burn rate breaches its budget, an alert fires. At its best this gives thirty minutes of response time; in tightly coupled distributed systems it often means responding after the failure has already cascaded past graceful recovery.

The cost of this lag is not evenly distributed. For the large class of failures that are genuine tipping points, slow saturation, runaway retry amplification, cascading resource exhaustion, the gap between the first visible symptom and unrecoverable collapse can be minutes, far too short for a human to diagnose and intervene. The result is a war room convened to manage a collapse that, in retrospect, had been building for hours. The information needed to act earlier existed in the telemetry the whole time; what was missing was a way to read it as a signal of approaching transition rather than as ordinary noise.

Yet experienced operators recognise, in hindsight, that the collapse was foreshadowed. Recovery from minor perturbations took progressively longer. The variance of tail latency widened. Independent services began to move together. These are not vague impressions: they are the precise, theorem-backed signatures that dynamical systems theory proves must appear as a system approaches a tipping point. The reliability community has not used them because it has not known they exist.

Critical slowing down is the phenomenon that, as a dynamical system nears a bifurcation, its recovery from perturbation slows because the restoring force weakens. The same mathematics has provided early warning of regime shifts in ecology, of abrupt climate transitions, and of financial crises [1], [2], [3], [4]. Its signatures are computable from ordinary time series. Bringing this established theory to infrastructure reliability, where the telemetry needed is already collected, is the gap this paper fills.

The main contributions of this work are: 1) a mapping of the four canonical early-warning signals onto observable SRE metrics; 2) a composite order parameter that fuses them into a single tip-proximity score; 3) a bifurcation classifier that names the kind of transition approaching and selects a matched intervention; and 4) a comparative analysis against threshold and burn-rate alerting. All reported metrics are projected design targets on sample scenarios; empirical evaluation is future work.

10.48047/jocaaa.2026.35.06.27

The remainder of this paper is organized as follows. Section II reviews critical-transition theory and reliability alerting. Section III states the theoretical basis, Section IV the signal mapping, and Section V the order parameter and bifurcation classifier. Section VI presents the architecture, Section VII the comparative analysis, and Sections VIII through X cover limitations, future work, and conclusions.

II. Related Work

A. Early-warning signals for critical transitions

The theory of early-warning signals holds that a broad class of systems exhibits universal precursors before a tipping point, regardless of the system's substrate [1]. Rising autocorrelation and variance have been demonstrated ahead of abrupt shifts in lakes, climate records, and physiological collapse [2], [4], and methodological surveys catalogue how to estimate these indicators robustly from finite, noisy time series [3]. The mathematics rests on bifurcation theory and the slowing of the dominant eigenvalue of the linearised dynamics near a fold [5], [6].

B. Reliability alerting

Modern SRE practice centres on service-level objectives and error budgets, with multi-window burn-rate alerts extending the response horizon by extrapolating current degradation [7]. Machine-learning anomaly detection flags deviations from a historical baseline. Both are valuable, but neither distinguishes a system that is merely anomalous from one that is approaching a qualitative state change, and burn-rate extrapolation is linear where the dynamics near a tipping point are not [8]. This paper is complementary: it adds the missing nonlinear precursor signal rather than replacing budget-based alerting.

A separate strand of work mines logs and traces to localise faults once they manifest, and a growing AIOps literature applies sequence models to predict imminent failures from learned patterns. These methods can be powerful, but they are statistical pattern-matchers trained on past failures; they do not carry a guarantee that a particular precursor must appear. The appeal of critical-transition theory is exactly that guarantee: the four signals are consequences of the mathematics of bifurcations, not correlations discovered in a training set, so they transfer to failure geometries the system has never exhibited before.

More recent AIOps work applies sequence models and LLM-based agents to failure triage and repair in microservices [14]; like the methods above, these act on or after a failure manifests rather than detecting the dynamical precursors of one, which is the gap this paper addresses.

III. Theoretical Basis

The foundation is exact rather than empirical. In a system governed by a potential landscape, the recovery rate from a small perturbation is proportional to the curvature of the potential around the stable equilibrium. As a control parameter drives the system toward a tipping point, that curvature flattens and the recovery rate falls toward zero, the defining meaning of critical slowing down [1], [5].

Three measurable consequences follow. Variance grows because the flattened potential lets the state wander further before returning. Lag-one autocorrelation rises toward one because each state increasingly resembles the previous one. And across coupled subsystems, spatial coherence increases because the independent stabilising forces weaken and shared coupling comes to dominate. These four signals, slowing recovery, inflating variance, rising autocorrelation, and growing coherence, are theorems about any system near a fold, not heuristics tuned to one domain [2], [6].

Concretely, linearise the dynamics about the stable equilibrium and let the dominant eigenvalue of the Jacobian be lambda, which is negative for a stable state. The characteristic return time after a perturbation scales as the reciprocal of the magnitude of lambda, so as the system approaches a fold and lambda rises toward zero the return time diverges. This single fact, established by Wissel and used throughout the critical-transitions literature [13], is the mathematical root of all four signals: a return time that grows without bound is observed, in sampled data, as slower recovery, larger variance, autocorrelation approaching one, and tighter coupling among components that the weakening restoring force can no longer hold apart.

Two properties make these indicators attractive for an operational setting. They are model-free: estimating them requires no knowledge of the system's governing equations, only its time series, which is essential when the 'equations' are a sprawling microservice deployment. And they are leading rather than lagging: they rise while the system is still nominally healthy, before any threshold is crossed, which is precisely the window in which a daytime, planned intervention is still possible.

A complementary, stochastic reading sharpens why variance and autocorrelation are reliable. Treat the

10.48047/jocaaa.2026.35.06.27

system as fluctuating under continual small perturbations around its equilibrium, an Ornstein-Uhlenbeck picture. The stationary variance of such a process is inversely proportional to the magnitude of the restoring rate, and its lag-one autocorrelation approaches one as that rate approaches zero. So the same flattening of the potential that slows recovery simultaneously inflates the equilibrium variance and lifts the autocorrelation, tying all the indicators to a single underlying parameter. This is why the signals tend to move together as a tipping point nears, and why their coherent co-movement, rather than any one crossing a line, is the trustworthy alarm.

IV. Mapping Signals to SRE Metrics

RPT-EWS instantiates each theoretical signal with telemetry that reliability teams already collect, so no new instrumentation is required. Recovery-rate decay is estimated from the trend in mean-time-to-recover across recent minor perturbations. Variance inflation is the rolling variance of p99 latency. Rising autocorrelation is the lag-one (AR(1)) coefficient of the error-rate series. Spatial coherence is the leading correlation of the cross-service health matrix.

Normalisation against each signal's own recent baseline is what makes the approach portable. A latency-variance level that is alarming for a tightly toleranced payments path is unremarkable for a batch analytics service; by reacting to the relative rise of each indicator rather than its absolute value, the framework applies the same logic to both without per-service threshold tuning. This sidesteps the central operational burden of threshold alerting, the endless hand-calibration of limits, and replaces it with a single calibrated mapping from coherent signal movement to tip proximity.

Each signal is normalised against its own recent baseline so that the framework reacts to relative change rather than absolute level, which makes it portable across services of very different scale. The four normalised signals are the inputs to the order parameter defined next.

The choice of estimator for each signal matters in practice. Recovery rate is read from the decay of the autocorrelation function or, operationally, from the trend in restore time across recent minor incidents; variance and AR(1) are computed over a sliding window whose length trades sensitivity against false alarms; and coherence is the leading eigenvalue of the windowed cross-service correlation matrix, which summarises how much of the fleet's variation has collapsed onto a single shared mode. Each estimator is cheap, incremental, and runs on the one-minute data a monitoring stack already produces.

V. Order Parameter and Bifurcation Classification

The four normalised signals are fused into a composite order parameter Phi in the interval [0,1], a single tip-proximity score whose weights are calibrated on historical incidents. A two-level policy turns the score into action: at Phi >= 0.65 the framework raises a medium alert recommending investigation, and at Phi >= 0.85 it escalates to an emergency. The evaluation runs in well under a second on one-minute data, so it adds negligible overhead to an existing monitoring pipeline.

The weights that fuse the four signals are not arbitrary. They are fit on a corpus of historical incidents by maximising separation between the pre-collapse window and ordinary operation, so a signal that is noisy or uninformative for a given fleet is down-weighted automatically. Because the order parameter is a scalar, it also gives operators something a four-panel dashboard does not: a single trend line whose slope conveys not just that the system is near a tipping point but how fast it is approaching one, which is the quantity that distinguishes a watch from a page.

Proximity to a tipping point is not enough; the kind of tipping point determines the right response. The framework classifies the approaching bifurcation into one of three canonical types and selects a matched intervention: a fold, which produces an abrupt cascade, calls for load shedding; a Hopf bifurcation, which produces growing oscillation in latency, calls for oscillation damping; and a transcritical transition, a graceful exchange of stability, calls for a managed migration. Figure 2 sketches these regimes. The projected classification accuracy and the prevention gains it enables are design targets, stated as such, and are reported alongside the comparative analysis.

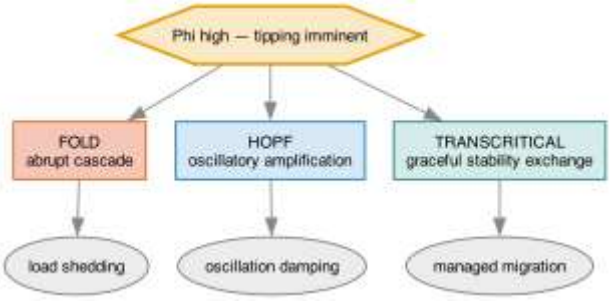

Fig. 2. Bifurcation regimes and matched interventions (illustrative). A fold drives an abrupt cascade (load shedding), a Hopf bifurcation drives oscillatory amplification (damping), and a transcritical transition is a graceful stability exchange (managed migration).

10.48047/jocaaa.2026.35.06.27

VI. System Architecture

Figure 1 shows the pipeline. Telemetry already gathered through metrics, tracing, and incident tooling feeds four signal estimators, one per early-warning indicator. Their normalised outputs are fused into the order parameter Phi, which drives the two-level alarm and, when proximity is high, the bifurcation classifier that selects the intervention. The framework emits a tip-proximity score, a bifurcation label, and a recommended action into existing on-call tooling, running as a lightweight background consumer of the monitoring stream.

Integration is deliberately undemanding. The framework subscribes to the existing metric and trace streams, computes the four estimators incrementally, and writes the order parameter back as just another metric, so it can be graphed, alerted on, and recorded with the same tooling teams already operate. The bifurcation label and recommended intervention are attached to the alert payload, giving the on-call engineer a proximate diagnosis and a starting action rather than a bare anomaly. Nothing in the design requires changing application code, adding agents, or instrumenting new signals.

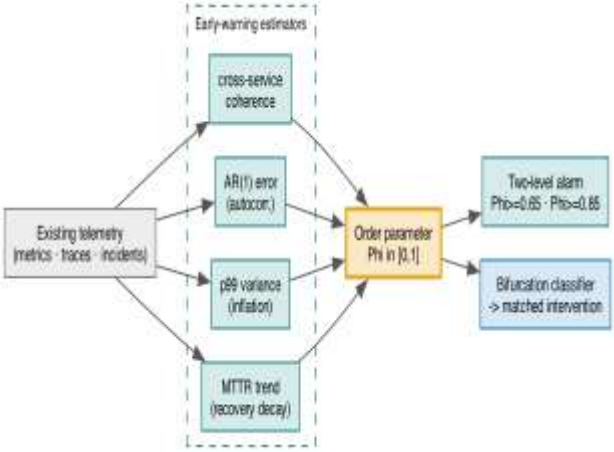

Fig. 1. RPT-EWS architecture. Four early-warning estimators draw on existing telemetry, fuse into the order parameter Phi, and drive a two-level alarm plus a bifurcation classifier that selects a matched intervention.

VII. Worked Example: A Saturating Connection Pool

Consider a payments service backed by a fixed-size database connection pool. Over an afternoon, a slow query regression raises mean hold time, and the pool begins to approach saturation, a classic fold: below a critical utilisation the pool absorbs bursts and recovers; above it, requests queue, hold times rise, and the system tips into a self-reinforcing collapse.

Hours before saturation, the four indicators move in the way theory predicts. Recovery from minor load bursts slows: the MTTR-trend estimator rises as the pool takes longer to drain after each spike. The rolling variance of p99 latency inflates as the flattening stability margin lets latency excursions run further. The AR(1) coefficient of the error-rate series climbs toward one as each interval increasingly resembles the last. And the cross-service coherence grows as upstream callers, blocked on the same pool, begin to move together. The composite order parameter rises through the medium-alert band well before any latency or error threshold would fire.

Crucially, the bifurcation classifier labels this regime a fold and recommends load shedding, the intervention matched to an abrupt cascade, rather than oscillation damping or a managed migration. An operator paged at this point is acting on a tip-proximity signal with hours of head room, not reacting to a saturation that has already begun. The same example also bounds the method: if the regression instead arrived as an instantaneous bad deploy that exhausted the pool in seconds, there would be no slow precursor to detect, the limitation Section VIII makes explicit.

The example also illustrates why a single composite score is operationally preferable to four separate alerts. During the afternoon, each individual indicator crosses and re-crosses its own noise band several times; alerting on any one of them in isolation would produce a stream of false starts. The order parameter, by requiring coherent movement across all four signals, stays quiet through the incoherent fluctuations and rises only when the indicators align, which is the regime the theory associates with a genuine approach to a tipping point.

VIII. Comparative Analysis

Because no production tool models critical-transition precursors, we compare RPT-EWS against the dominant practices it complements: SLO burn-rate alerting and machine-learning anomaly detection. Table I summarises the comparison. The values in the proposed column are projected design targets illustrated on sample collapse scenarios, not measured results; they express the behaviour the framework is designed to achieve and frame the evaluation that future work will perform on production telemetry.

The structural distinction is the predictive horizon. Burn-rate alerting extrapolates linearly from current degradation and so cannot see a nonlinear tipping point until degradation is already underway; anomaly

10.48047/jocaaa.2026.35.06.27

detection flags deviation but cannot distinguish anomalous-but-stable from anomalous-and-tipping. The order parameter and bifurcation label are signals neither baseline can produce, which is why several cells in the comparison are not merely lower but unavailable.

It is worth stating what the comparison is not claiming. RPT-EWS does not supersede burn-rate alerting or anomaly detection; a mature reliability practice will run all three. Budget burn remains the right instrument for slow, linear exhaustion of an error budget, and anomaly detection remains useful for point failures with no dynamical precursor. The contribution is orthogonal: a precursor signal for the specific and costly class of failures that are genuine phase transitions, which the other two instruments are structurally unable to anticipate.

TABLE I — PROJECTED CAPABILITY COMPARISON (DESIGN TARGETS, NOT MEASURED)

Capability	Proposed (projected)	Best baseline
Early-warning lead time	design target, hours	~0.5-2 hr - SLO burn rate
Collapse prevention rate	design target ~83%	~56% - ML anomaly
True-positive rate	design target ~91%	~78% - ML anomaly
False-positive rate	design target ~9%	~19% - burn rate
Bifurcation classification	design target ~87%	none - no baseline

IX. Limitations and Threats to Validity

The quantitative figures here are projected design targets on sample scenarios and carry no empirical validity; they describe intended behaviour and must be confirmed on production data before any operational claim is made. This is the central threat to validity, stated plainly.

Three substantive limitations remain. First, the early-warning theory assumes the collapse is a genuine bifurcation; a failure driven by a sudden external shock, a bad deploy that instantly breaks a service, has no slow precursor and will not be caught. Second, the signals require a baseline of normal fluctuation to estimate against, so newly launched services or rare failure modes offer too little history for a stable order parameter. Third, weight calibration on past incidents risks overfitting to the failure modes already seen; a previously unobserved collapse geometry may be scored incorrectly until the calibration set grows.

A fourth, subtler limitation concerns the sampling rate relative to the speed of the transition. Early-warning indicators are reliable only when the telemetry is sampled fast enough to resolve the system's return time; a transition that unfolds faster than the monitoring cadence will not present a clean rising-autocorrelation signature. For most infrastructure tipping points, which develop over minutes to hours, one-minute telemetry is ample, but very fast electrical or kernel-level failures fall outside the method's reach and are better served by conventional threshold protection.

X. Future Work

The immediate next step is empirical evaluation on labelled collapse events across production clusters, measuring lead time, prevention rate, and bifurcation-classification accuracy against the projected targets stated here. Beyond validation, three directions follow: online, non-parametric estimation of the early-warning indicators to reduce dependence on a fixed baseline [3]; joint inference of the cross-service coupling structure when topology is uncertain; and closing the loop so that a high order parameter can trigger a pre-approved automated intervention, load shedding or migration, without waiting for human escalation.

A further direction is to exploit the bifurcation label beyond intervention selection. Knowing whether an approaching transition is a fold, a Hopf, or transcritical carries information about the expected post-transition state and the reversibility of the change, which could inform whether to fight the transition or to ride it to a degraded-but-stable mode. Connecting the early-warning signal to capacity planning is also natural: a service whose order parameter trends upward across deploys, even without ever tipping, is revealing a shrinking stability margin that warrants provisioning attention before any incident occurs.

XI. Conclusion

Infrastructure collapse is, in many cases, a phase transition, and dynamical systems theory proves that phase transitions announce themselves. Recovery-rate decay, variance inflation, rising autocorrelation, and growing coherence are computable from telemetry teams already keep, and fused into an order parameter they give a tip-proximity signal that linear burn-rate alerting cannot. The contribution here is the mapping, the order parameter, the bifurcation classifier, and the comparative analysis; the reported numbers are projected design

10.48047/jocaaa.2026.35.06.27

targets, and validating them on production telemetry is the work that follows.

XII. References

[1] M. Scheffer et al., "Early-warning signals for critical transitions," Nature, vol. 461, no. 7260, pp. 53–59, 2009.

[2] V. Dakos et al., "Slowing down as an early warning signal for abrupt climate change," Proc. Natl. Acad. Sci. USA, vol. 105, no. 38, pp. 14308–14312, 2008.

[3] V. Dakos et al., "Methods for detecting early warnings of critical transitions in time series illustrated using simulated ecological data," PLoS ONE, vol. 7, no. 7, e41010, 2012.

[4] R. M. May, S. A. Levin, and G. Sugihara, "Ecology for bankers," Nature, vol. 451, no. 7181, pp. 893–895, 2008.

[5] S. H. Strogatz, Nonlinear Dynamics and Chaos, 2nd ed. Boulder, CO: Westview Press, 2015.

[6] C. Kuehn, "A mathematical framework for critical transitions: Bifurcations, fast-slow systems and stochastic dynamics," Physica D, vol. 240, no. 12, pp. 1020–1035, 2011.

[7] B. Beyer, C. Jones, J. Petoff, and N. R. Murphy, Site Reliability Engineering: How Google Runs Production Systems. Sebastopol, CA: O'Reilly Media, 2016.

[8] B. Beyer, N. R. Murphy, D. K. Rensin, K. Kawahara, and S. Thorne, The Site Reliability Workbook. Sebastopol, CA: O'Reilly Media, 2018.

[9] T. M. Lenton et al., "Tipping elements in the Earth's climate system," Proc. Natl. Acad. Sci. USA, vol. 105, no. 6, pp. 1786–1793, 2008.

[10] S. R. Carpenter and W. A. Brock, "Rising variance: A leading indicator of ecological transition," Ecology Letters, vol. 9, no. 3, pp. 311–318, 2006.

[11] J. M. T. Thompson and J. Sieber, "Predicting climate tipping as a noisy bifurcation: A review," Int. J. Bifurcation and Chaos, vol. 21, no. 2, pp. 399–423, 2011.

[12] A. J. Veraart et al., "Recovery rates reflect distance to a tipping point in a living system," Nature, vol. 481, no. 7381, pp. 357–359, 2012.

[13] C. Wissel, "A universal law of the characteristic return time near thresholds," Oecologia, vol. 65, no. 1, pp. 101–107, 1984.

[14] J. Wang, X. Liao, and W. Wu, "STAR: A stage-attributed triage and repair framework for RCA agents in microservices," arXiv:2605.15581, 2026.