PRACTITIONER RESEARCH ARTICLE

From Monitoring to Validation

A Proposed Governance Assurance Model for Continuously Validated Operational Confidence

Ferdinand C. Guarin · Architect of iSOAF · 2026

ABSTRACT

Governance functions across sectors continue to rely on monitoring dashboards, periodic audits, and point-in-time assessments to answer a persistent operational question: are the controls that are supposed to be working actually working, right now, under current conditions? This article proposes the governance assurance gap as a three-component analytical model — a configuration gap, an evidence gap, and a governance gap — that together explain why visibility does not automatically establish reliance. It formalizes a continuous validation doctrine comprising eight governing principles and an eleven-stage assurance journey, and situates both within the continuous auditing literature and established IT governance theory. The model is illustrated with twelve months of operational observation from a live, continuously operating enterprise environment. The article is written for governance practitioners, internal and external auditors, risk professionals, and standards contributors evaluating how continuous validation disciplines might extend — rather than replace — existing governance, risk, and compliance practice.

1. Introduction: An Assumption Model Under Strain

Governance, risk, and compliance programs are built on a familiar cadence: controls are designed, documented, and periodically reviewed. Between review cycles, organizations largely operate on the assumption that what was confirmed still holds. This assumption model has served governance practice for decades, but it carries a structural weakness that becomes more consequential as operating environments grow more dynamic: the moment a review closes, the confidence it produced begins to age.

Audit findings routinely describe conditions that had been accumulating well before the audit took place. Customers and end users frequently notice service degradation before it appears on a dashboard. Strategic and operational decisions are made on the basis of what was true at the time of the last assessment, not necessarily what is true today. None of this reflects a failure of the professionals conducting the reviews; it reflects a structural limitation in a governance model organized around periodic confirmation rather than continuous validation.

This article proposes a governance assurance model designed to make that gap explicit, measurable, and — where evidence permits — closeable. The model does not propose replacing monitoring, audit, or established control frameworks. It proposes an additional evaluative layer that asks a narrower and more specific question than either: given the evidence available right now, can the resulting conclusion still support reliance?

2. The Governance Assurance Gap: A Three-Component Model

The three-component Governance Assurance Gap model presented in this article emerged through repeated operational analysis conducted during the design, implementation, and twelve-month observation of the iSOAF assurance discipline. Rather than introducing three independent theoretical constructs in isolation, the model formalizes recurring patterns observed while evaluating the relationship between operational evidence, governance accountability, and reliance decisions. It is presented as an analytical abstraction derived from operational practice and positioned for further academic evaluation.

The assurance gap is proposed here as a composite of three distinct, simultaneously occurring gaps, each of which requires a different form of evidence to close.

2.1 The Configuration Gap

A control that is deployed and documented is not the same as a control that is currently functioning as intended. Configuration confirmed at the point of implementation, or at the last assessment cycle, can degrade, drift, or silently fail without triggering any alert, because most monitoring is designed to detect known failure signatures rather than gradual departure from an expected operating state. The governing question is not whether a control exists, but whether it is still behaving as designed.

2.2 The Evidence Gap

Governance conclusions are frequently reported without evidence that remains current at the moment of reporting. A report, by construction, describes a historical state. Where the interval between the state being measured and the state being relied upon is unmanaged, the conclusion drawn from that report is describing the past while being used to inform a present decision. The governing question is whether the evidence underlying a conclusion is current enough to support the reliance being placed on it.

2.3 The Governance Gap

Accountability structures frequently exist in policy but cannot be demonstrated through operational evidence without manual reconstruction — a process that is often dependent on the presence and institutional memory of a specific individual. Where evidence of accountability is not self-sustaining, an organization's capacity to demonstrate governance is coupled to personnel continuity rather than to the governance system itself. The governing question is whether accountability can be evidenced independently of any one person's presence or explanation.

These three gaps are proposed as concurrent, not sequential: an organization may close one while the other two remain open, and a control may be simultaneously well-configured, poorly evidenced, and weakly governed. Treating the assurance gap as a single undifferentiated problem, rather than as three separable ones, is proposed here as a primary reason governance interventions frequently address symptoms without resolving the underlying condition.

3. Theoretical Grounding

3.1 Continuous Auditing Literature

The proposition that assurance should be continuous rather than periodic is not new. Vasarhelyi and Halper's foundational 1991 study of online systems auditing at AT&T argued that the increasing automation and continuous operation of business processes required audit approaches capable of operating on the same continuous basis as the systems being audited, rather than relying solely on retrospective, batch-oriented review. Subsequent work by Alles, Kogan, and Vasarhelyi through the 2000s extended this into a broader continuous assurance literature, examining how analytic monitoring, exception-based testing, and process-embedded controls could support assurance conclusions closer to real time, while also identifying the practical and professional obstacles — materiality judgment, independence, and audit evidence standards among them — that continuous approaches would need to address rather than bypass.

The model proposed in this article sits within that lineage but narrows its scope. Where much of the continuous auditing literature is concerned with automating the audit function itself, this model is concerned with a related but distinct question: whether the evidence an organization already produces — from its own operational, security, and governance systems — is sufficient to support a specific reliance decision at a specific moment. It is offered as a complement to continuous auditing scholarship rather than a restatement of it.

3.2 Governance Theory: Decision Rights and Accountability

Weill and Ross's 2004 framework for IT governance defines governance as the specification of decision rights and an accountability framework intended to encourage desirable behavior in the use of information technology — explicitly distinguishing governance, which determines who has authority to decide, from management, which makes the decisions themselves. That separation is structurally significant for the model proposed here. A continuous validation capability that produces evidence and surfaces conclusions is, in Weill and Ross's terms, informing a decision domain — it is not occupying one. The model is therefore designed to preserve the decision-rights boundary that governance theory establishes: it is positioned to strengthen the evidentiary basis on which an accountable decision-maker acts, not to relocate the decision itself.

This grounding matters for how the model should be evaluated by governance practitioners. A continuous validation capability that quietly assumed decision authority — closing findings, authorizing remediation, or accepting risk without a named accountable party — would represent a governance design failure regardless of its technical sophistication. The model discussed in this article is deliberately bounded to avoid that outcome, a constraint discussed further in Section 5.

4. A Continuous Validation Doctrine

The eight principles described in this section should not be interpreted as newly independent governance requirements. They formalize the operational doctrine consistently implemented throughout the iSOAF assurance discipline and provide a structured vocabulary for concepts

previously expressed through the framework's operational architecture and evidence-first governance model.

Closing the three-component gap requires more than additional tooling; it requires a set of operating principles that determine how evidence is treated, how long a conclusion may be relied upon, and where governance authority is preserved. The model proposes eight such principles, summarized below in the sequence they are typically encountered in practice.

1. Evidence Before Reliance — an operational conclusion should be supported by validated evidence before it is treated as something an executive, operator, auditor, or regulator can rely upon.
2. Continuous Validation — assurance is not a permanent state; as evidence, controls, and operating conditions change, conclusions must be re-evaluated rather than assumed to persist.
3. Governance Before Decision — operational health and governance authorization are evaluated independently; a system can be operationally stable while remaining under active governance constraint.
4. Human Governance — operational decisions remain under accountable human authority; a validation capability may interpret evidence, but it does not replace organizational decision-making.
5. Evidence Lineage — assurance conclusions should remain traceable to the evidence that produced them, enabling independent review rather than requiring the original evaluator's explanation.
6. Continuous Audit Readiness — readiness for review should be a standing operational condition rather than a periodic preparation exercise undertaken before an audit.
7. Operational Drift Detection — validation should identify meaningful departure from an expected state before that departure materially affects governance confidence, rather than after.
8. Re-validation Before Closure — a governance finding or condition should not be treated as closed until independent confirmation demonstrates that the intended state has actually been restored.

Individually, several of these principles restate positions already established in audit and governance practice. Their proposed contribution is architectural rather than novel in isolation: applied together, as a closed operating discipline rather than as separate best practices invoked situationally, they are intended to prevent the common failure mode in which an organization satisfies most of these principles most of the time, while the gaps between them accumulate into the configuration, evidence, and governance gaps described in Section 2.

5. The Eleven-Stage Assurance Journey

The proposed eleven-stage assurance journey is not intended as a project lifecycle or organizational maturity model. Instead, it represents progressively increasing levels of evidentiary scrutiny applied

to an operational conclusion. The ordering reflects the observed progression during the development of the framework—from evidence quality and analytical consistency, through organizational reliance, and ultimately toward independent external scrutiny.

The model further proposes that assurance confidence should be understood as a graduated sequence rather than a binary state. As scrutiny of a conclusion increases — moving from an internal operational check toward external audit, regulatory review, or public challenge — the evidentiary support required to sustain that conclusion should increase correspondingly. The eleven-stage sequence proposed is: evidence, validation, alignment, coherence, confidence, stability, reliance, credibility, reliability, defensibility, and challenge.

Each stage is intended to represent an increase in the rigor of scrutiny a conclusion must withstand, not a chronological project timeline. A conclusion may reach 'reliance' quickly within a narrow internal scope while requiring substantially more evidentiary work to reach 'defensibility' under external audit, or 'challenge' under adversarial or regulatory questioning. Framing assurance this way is intended to give governance practitioners a shared vocabulary for a question that is otherwise often left implicit: not simply whether a conclusion is supported, but how much scrutiny it has been built to withstand.

6. Operational Observations: A Twelve-Month Field Record

The operational observations presented in this section demonstrate the practical implementation of the proposed model within one enterprise environment. They should not be interpreted as empirical proof of universal applicability or superiority over alternative governance approaches, but as documented operational evidence supporting the feasibility of the proposed assurance discipline.

The model was developed through, and subsequently tested against, twelve months of continuous operation within a live, 24/7 manufacturing enterprise environment managed under a lean internal IT function. In the interest of the framework's own evidentiary standards, the figures below are reported as historical validation results from that documented period, not as a claim about current, ongoing, or universally reproducible performance.

Over the documented period, the environment recorded 99.998% infrastructure availability; zero occasions on which a monitored condition reached a critical operating state, with degradation events consistently identified and routed for governance review at an earlier, amber-level threshold; 341 of 365 days recorded in a fully assured operating state, with the remaining days reflecting active governance events under review rather than undetected failure; and a governance record covering the full twelve-month period that required no manual reconstruction to support external review.

These results are offered as a single field record, not as a generalized performance claim. A central limitation of the model at this stage is precisely that it has been demonstrated in one operational environment, under one set of regulatory and organizational conditions. Independent replication in other sectors, control environments, and regulatory regimes is identified as a necessary next step, and is discussed further in the closing section.

7. Implications for Governance Practitioners, Auditors, and Risk Professionals

For internal and external auditors, the three-component gap model offers a diagnostic vocabulary for a category of finding that is often described informally — 'the control was configured but not effective,' or 'the evidence didn't hold up' — without a structured way to classify which of the three underlying gaps was actually responsible. Distinguishing a configuration failure from an evidence failure from a governance failure is proposed as a precondition for recommending the correct remedy, since the three call for different interventions.

For risk professionals, the eleven-stage assurance journey offers a way to communicate risk posture that distinguishes between evidence sufficient for internal reliance and evidence sufficient to withstand external challenge — a distinction that risk registers scored on a single static value often collapse into one number.

For governance and standards professionals, the model is offered as a complement to, not a substitute for, established frameworks such as ISO/IEC 27001, NIST CSF, COBIT, and COSO ERM. Those frameworks define what controls and governance structures should exist. The model proposed here is concerned with a narrower, adjacent question: whether the evidence available today still supports reliance on the conclusion that those structures are working. Alignment work connecting the model's terminology to the vocabulary of established frameworks is identified as ongoing.

8. Limitations and Future Research

In addition to the limitations described below, the Three-Component Governance Assurance Gap model and the Eleven-Stage Assurance Journey remain proposed conceptual constructs requiring broader academic scrutiny, comparative evaluation, and independent operational replication before broader theoretical generalization can be justified.

This article presents a proposed model, not a validated theory. Its principal limitations are the single-environment origin of the operational observations reported in Section 6, the absence to date of independent academic peer review, and the early-stage status of formal mapping between the model's eight principles and the control language of established standards. Planned next steps include a structured conference submission addressing the validation–authorization separation principle in greater technical depth, and a policy-oriented treatment for national and sector-level standards bodies evaluating continuous governance assurance as a policy instrument. Both are intended to subject the model to a wider range of professional and academic scrutiny than a single practitioner article can provide.

9. Conclusion

Existing governance programs answer a valuable but limited question: what was reported, what was detected, what was confirmed at the time of review. This article has proposed a narrower, adjacent governance assurance model built to answer a different question — can the conclusion drawn from that reporting still support reliance, right now, under current conditions — and has grounded that

proposal in established continuous auditing and IT governance theory rather than presenting it as a departure from either. The model does not claim to replace monitoring, audit, or existing control frameworks, nor does it claim decision authority over the accountable humans those frameworks are designed to serve. It is offered as a starting point for practitioner and academic scrutiny of a single, structural question: whether governance confidence, as currently practiced, ages faster than most governance programs are built to notice.

REFERENCES

Alles, M., Kogan, A., & Vasarhelyi, M. (2004). Principles of analytic monitoring for continuous assurance. Journal of Emerging Technologies in Accounting, 1(1), 1–24.

Alles, M., Kogan, A., & Vasarhelyi, M. (2008). Putting continuous auditing theory into practice: Lessons from two pilot implementations. Journal of Information Systems, 22(2), 195–214.

Alles, M., Brennan, G., Kogan, A., & Vasarhelyi, M. (2006). Continuous monitoring of business process controls: A pilot implementation of a continuous auditing system at Siemens. International Journal of Accounting Information Systems, 7(2), 137–161.

Chan, D. Y., & Vasarhelyi, M. A. (2011). Innovation and practice of continuous auditing. International Journal of Accounting Information Systems, 12(2), 152–160.

Vasarhelyi, M. A., & Halper, F. B. (1991). The continuous audit of online systems. Auditing: A Journal of Practice & Theory, 10(1), 110–125.

Weill, P., & Ross, J. W. (2004). IT governance: How top performers manage IT decision rights for superior results. Harvard Business School Press.

International Organization for Standardization. (2022). ISO/IEC 27001:2022 — Information security, cybersecurity and privacy protection.

National Institute of Standards and Technology. (2024). NIST Cybersecurity Framework 2.0.

This article proposes a model developed from operational practice; it does not disclose implementation architecture, scoring logic, or proprietary assurance methods. Figures reported are historical validation results and should not be read as current runtime state or as a claim of independent peer review, regulatory approval, or broad market adoption.