identara

WHITE PAPER · IDENTITY GOVERNANCE

Measuring *What Moves*

A Dynamic Metrics Framework for Identity Governance Responsiveness

AUTHOR

Vidyaa Ganesh

PUBLISHER

Identara

PUBLISHED

May 26, 2026

CANONICAL URL

identara.ca/papers/measuring-what-moves

SUGGESTED CITATION

Ganesh, V. (2026). "Measuring What Moves: A Dynamic Metrics Framework for Identity Governance Responsiveness." Identara. https://identara.ca/papers/measuring-what-moves/

Abstract

Identity Governance and Administration (IGA) programs are designed to ensure that appropriate access is maintained across organizational resources. However, the operational metrics used to evaluate these programs remain overwhelmingly static and activity-based: certification completion rates, provisioning speed, and policy violation counts. These metrics measure whether governance work was performed, not whether the governance program is keeping pace with the rate at which access states change. This paper examines the gap between the continuously evolving nature of access environments and the periodic, checkpoint-based governance models that most organizations rely on. Drawing on a review of current IGA literature, industry frameworks, vendor approaches, and standards-body publications, we identify the absence of a formalized, vendor-neutral measurement vocabulary for access dynamics at the governance program level. We propose four named metrics to address this gap: Entitlement Drift Rate (decomposed into gross, net, and unreviewed variants), Governance Lag, Justification Half-Life, and Trust Gradient. These metrics are designed to function as program-level indicators of governance responsiveness, analogous to the role that DORA metrics play in measuring software delivery performance. The framework is applied to both human and non-human identity (NHI) governance contexts, with analysis of how each metric behaves differently across identity types. We position this framework not as a replacement for existing IGA measurement approaches but as a complementary layer that captures the dynamic properties current metrics do not address.

Keywords: *Identity Governance, IGA, Access Dynamics, Non-Human Identity, Governance Metrics, Zero Trust, Continuous Identity, DORA Metrics, Entitlement Management, Service Identity Governance*

1. Introduction

The practice of Identity Governance and Administration has matured significantly over the past decade. Modern IGA platforms from vendors including SailPoint, Saviynt, Okta, Microsoft, and Omada provide comprehensive capabilities for identity lifecycle management, access certification, role-based provisioning, separation of duties enforcement, and compliance reporting.[1] Organizations can now automate provisioning workflows, execute certification campaigns at scale, and generate auditor-ready evidence of governance activity.

Yet a fundamental tension persists. Access environments are dynamic: entitlements are granted, roles change, projects spin up and wind down, service accounts proliferate, and organizational structures shift. These changes are continuous. The governance mechanisms designed to manage them, however, remain largely periodic. Certification campaigns run quarterly or semi-annually. Access reviews are triggered by compliance calendars. Alerting is based on static thresholds. The result is a structural mismatch between the pace of access change and the cadence of governance response.

This mismatch is not a new observation. Multiple industry analysts, vendors, and standards bodies have identified the limitations of periodic governance models. Gartner has called for IGA to become "continuous, orchestrated, and intelligent."[2] SGNL, which CrowdStrike announced a definitive agreement to ac-

quire in January 2026 (subject to customary closing conditions),[3] has built its product positioning around "Continuous Identity" and Zero Standing Privilege. The Continuous Access Evaluation Profile (CAEP), a specification under the OpenID Foundation's Shared Signals Framework, provides protocol-level mechanisms for real-time session evaluation and has been championed by major identity ecosystem participants.[4] The IDPro Body of Knowledge has published work on optimizing access recertifications, including event-driven review triggers.[5]

The industry consensus that governance should be more continuous is well established. What remains underdeveloped is a formalized, vendor-neutral measurement vocabulary for the dynamic properties that make continuous governance necessary. Individual risk signals, detection heuristics, and vendor-specific analytics capabilities exist across the market (discussed in Section 2). These have not, however, been consolidated into a program-level governance metrics framework comparable to the shared vocabulary that DORA metrics provide for software delivery performance.[6]

This paper proposes a framework of four named metrics designed to address that gap. The framework draws an explicit structural analogy to the DORA metrics (Deployment Frequency, Lead Time for Changes, Change Failure Rate, and Mean Time to Restore) developed by Google Cloud's DevOps Research and Assessment team.[7] We argue that the absence of a shared IGA dynamics vocabulary is not merely a conceptual gap but an operational one: without metrics that capture the dynamics of access change, governance programs cannot objectively assess their own responsiveness, cannot justify their cadence decisions with data, and cannot detect when their operating tempo has fallen out of alignment with their environment.

2. Literature Review

2.1 The Continuous Governance Imperative

The critique of static, periodic governance is widespread across analyst literature and vendor positioning. Microsoft's Entra ID Governance documentation ties identity governance to Zero Trust execution, emphasizing automation, visibility, and ensuring the right people have the right access at the right time.[8] Palo Alto Networks positions IGA as "the continuous policy-enforcement point for all access" and a prerequisite for mature Zero Trust architecture.[9] Zluri has observed that Identity Security Posture Management "works in real time" as opposed to traditional IGA, which "works in cycles, like quarterly access reviews and annual certifications."[10]

Simon Moffatt, founder of The Cyber Hut, authored a three-part series on Continuous Identity for SGNL. His framework identifies three temporal layers at which access decisions occur: admin time (a manager decides whether someone should have access), login time (the identity provider evaluates conditions at authentication), and event time (changes in state trigger access re-evaluation via standards like CAEP).[3] This taxonomy is valuable for understanding runtime access control, but it operates at the authorization layer rather than the governance program layer.

CrowdStrike's 2024 Global Threat Report documented that intrusions involving identity abuse typically reach lateral movement within minutes, making the staleness of governance data an operational security concern.[11] SafePaaS has observed that "no matter how advanced your Zero Trust stack is, if the underlying entitlements are wrong, you are continuously verifying the wrong permissions."[12]

2.2 Non-Human Identity Governance

The governance challenge for non-human identities (NHIs), including service accounts, API credentials, workload identities, and machine-to-machine integrations, has become a central concern. The Cloud Security Alliance published "The Non-Human Identity Governance Vacuum" in May 2026.[13] Published NHI-to-human identity ratios vary by source and environment, with estimates ranging from dozens to more than one hundred NHIs per human identity; Entro Labs' H1 2025 research reports a ratio of 144:1 in cloud-native environments.[14] The OWASP NHI Top 10, released in 2025, formalizes security risks specific to non-human identities, including overprivileged NHIs, offboarding gaps, and auditing failures.[15]

Apono's analysis notes that NHIs "rarely go through onboarding or offboarding, rely on static API keys or long-lived tokens, and are frequently overprivileged."[16] These characteristics make NHIs particularly susceptible to the governance dynamics this paper examines: unchecked accumulation, unmeasured behavioral change, and unbounded trust decay.

2.3 Existing IGA Metrics and Adjacent Concepts

The current IGA metrics landscape is dominated by operational and compliance-oriented indicators. C1.ai's 2026 publication "10 IGA Metrics Every Security Team Should Use" advocates a shift from activity-based to outcome-based measurement.[17] Their proposed metrics include time to onboard/offboard, approval cycle time, certification auto-approval rate, high-risk identity reduction, and automation ROI. Omada's guidance similarly emphasizes ongoing measurement programs rather than single reports.[18] KuppingerCole's IGA Leadership Compass notes the importance of monitoring IGA processes "against defined key performance indicators."[19]

Several adjacent concepts in the market address aspects of the dynamics this paper examines. Cyberhaven uses the term "Permission velocity" as a detection signal, defined as "an entitlement count growing faster than peers in the same role."[20] Zluri has published quantitative analysis of mover-driven privilege creep operating in both horizontal and vertical dimensions.[21] SailPoint describes AI-driven governance capabilities including unusual access detection, peer-group recommendations, and risk-based certification prioritization. Microsoft Entra access reviews include smart recommendations based on inactivity and application activity signals. Okta positions governance recommendations around unified risk signals.

These capabilities represent important progress. However, they function as vendor-specific risk signals and detection heuristics rather than as a shared, vendor-neutral program-level measurement vocabulary. An organization using SailPoint cannot compare its governance responsiveness to an organization using

Saviynt using these signals, because there is no common framework. The metrics proposed in this paper are designed to provide that common layer.

2.4 The DORA Metrics Precedent

The DORA framework, formalized through the annual Accelerate State of DevOps Report and detailed in Forsgren, Humble, and Kim's *Accelerate* (2018), provides the structural precedent for this paper.[7] DORA introduced four specific, named metrics that gave software engineering teams a shared vocabulary for measuring delivery performance. Before DORA, DevOps maturity was discussed qualitatively; after DORA, it could be measured, benchmarked, and compared across organizations.[22]

The parallel is structural, not substantive. DORA measures software delivery velocity and stability. The proposed IGA metrics measure governance responsiveness to access dynamics. The analogy lies in the role each framework plays: providing a shared, named, quantitative vocabulary where previously only qualitative assessment existed. (Note: the acronym DORA also refers to the EU Digital Operational Resilience Act, Regulation 2022/2554, an unrelated financial services regulation.[23] This paper refers exclusively to the DevOps Research and Assessment framework.)

3. Gap Analysis

The literature review reveals a consistent pattern. The industry has converged on the diagnosis (periodic governance is insufficient) and on the directional prescription (governance should be continuous). Multiple vendors have built products that address continuous evaluation at the runtime access control layer. Individual risk signals and detection heuristics exist across the vendor landscape. What remains underdeveloped is a formalized, vendor-neutral program-level measurement vocabulary for the dynamics these signals individually address.

This gap can be characterized along four dimensions:

Rate of change. No standard program-level metric exists for measuring the rate at which unreviewed access accumulates across a governed population. Cyberhaven's "Permission velocity" addresses this at the individual user level as an anomaly detection signal.[20] No equivalent exists at the program level, answering: "How fast is unreviewed access accumulating across our entire governed population, and does our certification cadence match that rate?"

Detection latency. While Mean Time to Detect (MTTD) and Mean Time to Respond (MTTR) are established security operations metrics, no IGA-specific equivalent measures the time gap between an access state becoming inappropriate and the governance program identifying and remediating it. A recent industry analysis of identity in the SOC identifies "decision latency" as a critical vulnerability,[24] but this concept is applied to real-time threat response, not governance program cadence.

Justification relevance. IGA platforms capture business justification at the point of access request. No standard metric tracks the decay of that justification's relevance over time. The IDPro Body of Knowledge suggests triggering reviews based on changes to entitlements, user risk, or peer-group deviation,[5] but does not model the justification itself as a quantity whose relevance degrades at a measurable rate.

Trust currency. Zero Trust architecture emphasizes continuous verification, and CAEP provides a protocol for real-time session evaluation.[4] However, these mechanisms operate at the session and authentication layer. At the governance layer, the question of how much confidence should be placed in a standing access grant has no standardized measure. Trust in access grants is treated as binary (certified or not certified) rather than as a continuously evolving quantity.

4. Proposed Framework: Identity Governance Dynamics Metrics

We propose four named metrics designed to address the gaps identified above. Together, these metrics provide a measurement vocabulary for the dynamic properties of access governance, complementing existing operational IGA metrics in the same way that DORA metrics complemented existing software engineering process metrics. These metrics formalize concepts that exist individually and informally across vendor implementations and practitioner intuition; the contribution is their consolidation into a vendor-neutral, program-level governance responsiveness framework.

4.1 Entitlement Drift Rate (EDR)

Definition: The rate at which entitlements change across a governed population between governance checkpoints, decomposed into three complementary measures.

A single net drift measure, while useful, can obscure dangerous access volatility. If 1,000 entitlements are granted and 950 are removed within a governance cycle, the net drift is only 50, but the environment experienced significant access churn that may warrant investigation. EDR is therefore decomposed into three variants:

VARIANT	FORMULA	PURPOSE
Gross Entitlement Velocity (GEV)	GEV = G / (I × D)	Total access change activity; captures churn independent of direction
Net Entitlement Drift (NED)	NED = (G − R) / (I × D)	Directional accumulation; measures whether the environment is expanding or contracting
Unreviewed Entitlement Drift (UED)	UED = U / (I × D)	Governance debt; directly measures unreviewed accumulation

Where G = total new entitlement grants in the period, R = total entitlement removals in the period, U = new entitlement grants not yet reviewed or certified in the period, I = number of governed identities, and D = number of days in the period.

Operational definition of "unreviewed." For the purposes of UED calculation, an entitlement is classified as unreviewed if it was granted after the completion date of the most recent certification campaign for its applicable scope and has not been subject to any of the following governance actions prior to the next campaign: inclusion in a completed certification decision (approved or revoked), review through an event-triggered micro-certification (such as a role-change review), or evaluation through an automated policy-based access decision. Entitlements granted through approved request workflows are not considered "reviewed" by virtue of initial approval alone; the approval establishes the entitlement, but governance review refers to subsequent validation that the entitlement remains appropriate. This definition ensures that UED captures the accumulation of access that has been provisioned but not yet subjected to ongoing governance scrutiny.

4.1.1 Sample Calculation

Consider a business unit with 5,000 governed identities operating on a quarterly (90-day) certification cadence. During the quarter, 12,000 new entitlement grants are issued and 4,000 entitlements are removed. At the start of the next certification cycle, 8,000 of the 12,000 grants have not yet been reviewed.

```
GEV = 12,000 / (5,000 × 90) = 0.027 grants per identity per day
```

```
NED = (12,000 − 4,000) / (5,000 × 90) = 0.018 net accumulation per identity per
day
```

```
UED = 8,000 / (5,000 × 90) = 0.018 unreviewed grants per identity per day
```

Over the 90-day cycle, this translates to approximately 1.6 unreviewed entitlements per identity entering each certification campaign. If the organization shortened its cycle to 60 days, the projected UED backlog per identity would decrease proportionally, providing a quantitative basis for cadence adjustment.

4.1.2 Relationship to Existing Concepts

EDR is related to Cyberhaven's "Permission velocity" concept,[20] which measures entitlement growth rate at the individual user level to flag anomalous accumulation. EDR differs in scope and purpose: it is a program-level metric designed to evaluate whether governance cadence matches organizational access change rate. An organization could have no individual users flagged by permission velocity (all users accumulating at roughly the same rate) while still having a dangerously high UED (all users accumulating rapidly, but uniformly). The concepts are complementary: permission velocity is a detection signal; EDR is a governance responsiveness indicator.

4.1.3 Limitations

EDR treats all entitlements equally. A low-risk read-only entitlement and a privileged administrative credential contribute identically to the metric. Risk-weighted variants (discussed in Section 8) would address this. Additionally, EDR requires reliable entitlement grant and revocation timestamps, which may not be consistently available across all connected applications in an IGA environment.

4.2 Governance Lag (GL)

Definition: The elapsed time between the moment an access state becomes inappropriate and the moment the governance program detects and initiates remediation.

Formula:

```
GL = T(detection) − T(inappropriateness)
```

Where T(detection) is the timestamp of the governance action (certification decision, automated flag, or manual review) and T(inappropriateness) is the timestamp of the business event that rendered the access no longer warranted.

Rationale: Existing operational metrics measure the speed of governance actions once initiated (time to provision, time to revoke). No standard metric captures the latency before governance action begins. This latency is the true risk window. An organization that revokes access within two hours of detection but takes four months to detect the access was inappropriate has a Governance Lag of four months, regardless of its revocation speed.

4.2.1 Event Anchors

The primary measurement challenge for Governance Lag is determining T(inappropriateness). The following table defines recommended event anchors by scenario:

SCENARIO	EVENT ANCHOR	DATA SOURCE
Employee termination	Termination effective date	HR system
Role change (mover)	Role change effective date	HR system
Project-based access	Project end date or phase completion	Project management system
Temporary elevation	Approved expiration date	IGA platform / ticketing system
NHI / service account	Last validated owner date, integration retirement date, or last usage date	IGA platform, CMDB, or SIEM
Vendor/contractor access	Contract end date	Procurement / vendor management system

4.2.2 Sample Calculation

During a quarterly certification campaign, 340 entitlements are revoked across 5,000 governed identities. For each revocation, the team identifies the business event that rendered the access inappropriate. Analysis shows: 120 revocations were for access that became inappropriate due to role changes (median 67 days before detection), 80 were for project-based access where the project ended (median 104 days before detection), 90 were for terminated contractor access (median 38 days before detection), and 50 were for NHI access tied to decommissioned integrations (median 211 days before detection).

```
Program-level median GL = weighted median across all categories = 74 days
```

```
NHI-specific GL = 211 days
```

This reveals that the governance program carries, on average, a 74-day risk window for human identities and a 211-day risk window for non-human identities. The NHI-specific GL of 211 days is a direct measure of the governance vacuum for machine identities.

4.2.3 Limitations

Governance Lag depends on the ability to identify when access became inappropriate, which requires correlation between IGA data and HR, project management, or CMDB data. Not all scenarios produce a clean event anchor. Gradual organizational restructuring, for instance, does not generate a single date at which access becomes inappropriate. In such cases, GL may need to be approximated using the date of the organizational change announcement or the date of the restructured reporting structure taking effect.

4.3 Justification Half-Life (JHL)

Definition: The estimated duration before the business justification for an access grant of a given type loses half its original relevance, derived from a composite of governance and usage signals.

Formula (estimated):

```
JHL = T at which P(revocation or disuse) ≥ 0.5 for a given access category
```

Where P(revocation or disuse) is the cumulative probability that an entitlement in the given category will be revoked during certification or will become dormant (unused for a defined threshold period) by elapsed time T after initial approval.

Rationale: IGA platforms capture business justification at the point of access request. Once approved, that justification is treated as indefinitely valid until the next review cycle. In practice, the relevance of a justification decays at different rates. Access to a long-term ERP module decays slowly. Access to a project-specific development environment decays rapidly once the project concludes. The half-life metaphor,

drawn from physics and applied in other domains such as the "half-life of knowledge" concept in scientometrics,[25] provides a quantitative model for a phenomenon governance practitioners recognize intuitively but do not currently measure.

4.3.1 Multi-Signal Estimation

Relying solely on certification revocation data to estimate JHL is vulnerable to the rubber-stamping problem identified in the IDPro Body of Knowledge.[5] If reviewers routinely approve all entitlements, revocation rates will understate how quickly justification actually decays. JHL should therefore be estimated from multiple signals:

SIGNAL	RELEVANCE TO JHL ESTIMATION
Certification revocation rate	Direct reviewer decisions to remove access, segmented by access type and time since approval
Entitlement usage decay	Decline in usage frequency over time; access that is no longer being exercised is access whose justification may have expired
Role/job change frequency	Rate of organizational mobility; high mover rates shorten JHL for role-specific access
Project or contract end dates	Known expiration of the business context that justified the access
Peer-group deviation	Growing divergence from the entitlement profile of current-role peers indicates the access may no longer be role-appropriate

By incorporating usage and contextual signals alongside certification outcomes, JHL becomes robust to the failure mode of unreliable reviewer decisions. The weighting of these signals should be calibrated against historical outcomes in each organization's environment.

4.3.2 Application to Governance Decisions

JHL enables risk-proportionate certification cadence. Rather than applying a uniform review frequency across all access types, organizations can use JHL data to design tiered certification strategies: shorter cycles for access categories with short half-lives (project-based access, temporary elevations, developer environment permissions) and longer cycles for access categories with long half-lives (core job function entitlements, stable role-based access). This directly addresses the rubber-stamping problem by reducing the volume of genuinely stable access that reviewers must evaluate, concentrating reviewer attention on the access most likely to have become inappropriate.

4.3.3 Limitations

JHL is an estimated metric, not a directly observed one. Its accuracy depends on the quality and availability of the underlying signals, particularly entitlement usage data (which requires SIEM or application-level telemetry integration) and project lifecycle data (which requires integration with project management sys-

tems). In organizations where these data sources are unavailable, JHL estimation will necessarily rely more heavily on certification revocation patterns, with the acknowledged risk of rubber-stamping distortion.

4.4 Trust Gradient (TG)

Definition: A composite, continuously updated confidence score estimating the degree to which a standing access grant is believed to remain appropriate, based on governance-layer signals including time since last validation, usage recency, organizational context stability, and peer-group alignment.

Conceptual formula:

> $TG = f(w_1 \times TimeSinceValidation, w_2 \times UsageRecency, w_3 \times ContextStability, w_4 \times PeerAlignment)$

Where w_1 through w_4 are organization-specific weights calibrated against historical certification outcomes, TimeSinceValidation is a decaying function of elapsed time since last certification, UsageRecency reflects the frequency and recency of entitlement exercise, ContextStability reflects the stability of the identity's role and organizational position, and PeerAlignment reflects the degree to which the entitlement matches the current profile of role peers.

Rationale: Current IGA systems treat access appropriateness as binary: certified (appropriate) or flagged for review. The certification event refreshes trust to full confidence; between events, no measurement of trust currency exists. This binary model ignores the reality that confidence in an access grant degrades continuously as time passes, organizational context shifts, and absence of usage signals accumulates.

4.4.1 Distinction from Vendor-Specific Analytics Capabilities

Trust Gradient must be clearly distinguished from the risk scoring, identity analytics, peer comparison, and access recommendation capabilities offered by vendors including SailPoint, Microsoft, and Okta. Some of these capabilities operate at the runtime access control layer (adaptive authentication, session risk scoring), while others operate within governance workflows (SailPoint's AI-driven certification recommendations based on peer-group analysis, Microsoft Entra's access review recommendations based on inactivity and application activity signals). Trust Gradient differs from both categories not in the layer it addresses but in its design as a vendor-neutral, portable governance confidence score. Existing vendor capabilities are product-specific: an organization using SailPoint cannot compare its governance confidence posture to one using Saviynt using these tools. Trust Gradient proposes a standardized, composite measure of standing entitlement appropriateness, calibrated against historical outcomes and designed to prioritize review attention across any IGA platform. It is intended as a calibrated decision-support score, not as an objective truth determination.

4.4.2 Application to Governance Decisions

Trust Gradient enables continuous prioritization of governance attention. Rather than reviewing all access at the same cadence, organizations can surface entitlements with rapidly declining TG scores for early or targeted review, allocating reviewer attention to the access most likely to be inappropriate. This converts

certification from a bulk periodic exercise into a continuously prioritized governance activity.

4.4.3 Limitations

Trust Gradient is the most complex metric in the framework and carries the highest implementation burden. Composite scoring can become opaque if the weighting model is not transparent, and poorly calibrated weights can produce misleading confidence scores. Organizations implementing TG should validate the model against historical certification outcomes (does a low TG score predict revocation?) and should treat TG as a prioritization tool, not an automated decision mechanism. The weights will need periodic recalibration as organizational dynamics change.

5. Application to Non-Human Identity Governance

Each metric exhibits distinct behavior when applied to non-human identities. Because NHIs lack the natural lifecycle events and oversight structures that partially constrain human identity governance, the dynamic properties these metrics measure tend to be more extreme for NHIs.

Entitlement Drift Rate. For human identities, EDR is bounded by certification cycles. For NHIs frequently excluded from certification campaigns, EDR may be unbounded. Published NHI-to-human ratios ranging from dozens to over one hundred per human identity[14] suggests that total unreviewed NHI access volume in many organizations significantly exceeds total unreviewed human access volume, despite governance investment being overwhelmingly directed at human identities.

Governance Lag. For human identities, GL is bounded by certification cadence and partially mitigated by lifecycle events (role changes, terminations). For NHIs, GL can extend to years. The integration a service account was created for may be decommissioned, the project team disbanded, and the vendor contract expired, with no corresponding event in the IGA platform. The GL for many NHIs equals the age of the service account itself.

Justification Half-Life. NHI access justifications tend to have shorter half-lives than human access justifications in dynamic environments. Service accounts are frequently created for specific integration projects with finite timelines. The business justification becomes invalid at project completion, but unlike human project-based access, no organizational process signals that the justification has expired.

Trust Gradient. Usage patterns for NHIs may be highly regular (automated scheduled processes) or completely absent (orphaned accounts). Peer-group comparison is less applicable because NHIs are typically unique in function. The net effect is that TG for NHIs tends to decay faster and with less visibility than for human identities, reinforcing the need for explicit governance mechanisms such as time-based expiration and automated dormancy detection.

6. Relationship to Existing Frameworks

Relationship to Continuous Identity and CAEP. SGNL's Continuous Identity positioning and the CAEP specification address continuous evaluation at the runtime layer: whether a session should be allowed, a token revoked, or an access decision re-evaluated.[3,4] The proposed metrics operate at the governance program layer: whether the entitlements themselves should still exist. An organization could have a mature CAEP implementation (excellent runtime evaluation) while still having poor Governance Lag (long delays in identifying inappropriate standing access). Both layers are necessary.

Relationship to existing IGA metrics. Operational metrics like time to provision, time to revoke, and certification completion rate measure the efficiency and completeness of governance execution. The proposed metrics measure whether that execution is appropriately timed and targeted. An organization with 100% certification completion but a median GL of six months is executing governance work efficiently but not responsively. Both measurement layers are needed.

Relationship to Zero Trust. Zero Trust architecture, as defined in NIST SP 800-207, establishes that trust should not be assumed and must be continuously verified.[26] JHL and TG directly quantify the decay of trust in standing access grants, operationalizing "never trust, always verify" at the entitlement governance level rather than only at the session layer.

7. Discussion

7.1 Practical Implementation

The data required to compute these metrics exists in most mature IGA platforms. Entitlement grant and revocation timestamps, certification campaign results, user role change dates, and entitlement usage data (where SIEM integration exists) are standard data elements. Implementation complexity varies by metric. We recommend a phased adoption approach:

PHASE	METRIC	DATA REQUIREMENTS
Phase 1	Entitlement Drift Rate (all three variants)	IGA entitlement grant/revocation logs and certification cycle dates
Phase 2	Governance Lag	IGA data correlated with HR system, project management, and CMDB data
Phase 3	Justification Half-Life	Historical certification outcomes, entitlement usage telemetry, project lifecycle data
Phase 4	Trust Gradient	Real-time composite scoring from multiple data sources with calibration against outcomes

7.1.1 Minimum Data Model

Practical implementation requires specific data elements from across the identity and IT ecosystem. The following table identifies the minimum data fields required to compute each metric and the typical source system for each field.

DATA ELEMENT	REQUIRED FOR	SOURCE SYSTEM	NOTES
Identity ID	All metrics	IGA platform	Unique identifier for each governed identity (human and NHI)
Account ID	EDR, GL, TG	IGA platform	Links identity to target system accounts
Entitlement ID	All metrics	IGA platform	Unique identifier for each entitlement or access grant
Entitlement grant timestamp	EDR, JHL, TG	IGA platform	Date and time the entitlement was provisioned
Entitlement revoke timestamp	EDR, GL, JHL	IGA platform	Date and time the entitlement was removed
Certification decision timestamp	GL, JHL, TG	IGA platform	Date of reviewer approve/revoke decision per entitlement
Certification campaign completion date	EDR (UED)	IGA platform	Defines the boundary for unreviewed entitlement classification
User termination date	GL	HRIS	Event anchor for terminated-employee GL calculation
Job/department/manager change date	GL, JHL, TG	HRIS	Event anchor for mover GL; context signal for JHL and TG
Project or contract end date	GL, JHL	PM system / procurement	Event anchor for project-based and contractor GL
Entitlement last usage timestamp	JHL, TG	Application logs / SIEM	Indicates whether the entitlement is actively exercised
Entitlement risk tier	Risk-weighted variants	IGA platform / GRC	Enables future risk-weighted EDR and GL (see Section 8)
NHI owner	GL, TG	IGA platform / CMDB	Identifies accountable party for non-human identity governance
NHI last validated date	GL, TG	IGA platform / CMDB	Event anchor for NHI GL; decay input for NHI TG
Identity type classification	All metrics (segmentation)	IGA platform	Human vs. NHI classification; enables separate metric reporting

Not all data elements are required for initial adoption. Phase 1 (EDR) requires only the first seven fields, all of which are typically available within the IGA platform itself. Subsequent phases introduce cross-system data dependencies that increase implementation complexity but also increase metric precision.

7.2 Implications for Governance Program Design

Certification cadence becomes a data-driven decision rather than a compliance-calendar default. Organizations with high Entitlement Drift Rates need shorter cycles or more targeted continuous review. Organizations with low, stable EDR can justify longer cycles with data. Review prioritization becomes risk-proportionate through Trust Gradient scores. NHI governance becomes measurable for the first time, providing a quantitative basis for investment decisions in an area the industry increasingly recognizes as critical.

7.3 General Limitations

Beyond the per-metric limitations discussed in Section 4, several framework-level limitations should be acknowledged. First, the metrics are proposed based on conceptual analysis and literature review rather than empirical validation in production IGA environments. Empirical validation across diverse organizational contexts is a necessary next step. Second, the framework does not yet include benchmark thresholds. DORA's impact was amplified by its classification of teams into performance tiers (Elite, High, Medium, Low). Equivalent tiers for governance dynamics metrics would require industry-wide data collection beyond the scope of this paper. Third, all four metrics treat entitlements equally by default. Practical implementations would benefit from risk-weighting (discussed in Section 8).

8. Future Research Directions

Several directions for future work emerge. Empirical validation through pilot implementations in production IGA environments would provide calibration data and enable the development of performance benchmarks. Risk-weighted variants of each metric (Risk-Weighted EDR, Risk-Weighted GL) that account for entitlement sensitivity would increase operational usefulness by ensuring that privileged access accumulation is weighted more heavily than low-risk entitlement growth. The extension of the framework to agentic AI identities, which request and exercise access autonomously at machine speed,[12] presents a particularly urgent research question. The development of industry benchmark tiers, analogous to DORA's performance classifications, would enable cross-organizational comparison. Finally, integration pathways with existing IGA platform analytics capabilities warrant investigation, including whether current architectures can support Trust Gradient computation in near-real-time.

9. Conclusion

Modern IGA has metrics for execution, but not enough metrics for tempo. The industry has reached consensus that governance must become more continuous, more dynamic, and more responsive. Multiple vendors and standards bodies have contributed important capabilities at the runtime access control layer. What has remained underdeveloped is a vendor-neutral governance program measurement vocabulary for the dynamics that make continuous governance necessary.

This paper proposes four metrics designed to fill that gap. Entitlement Drift Rate measures how fast unreviewed access accumulates. Governance Lag measures how long inappropriate access persists before detection. Justification Half-Life estimates how quickly business justifications lose relevance. Trust Gradient provides a continuous confidence score for standing entitlements. Together, these metrics capture the dynamic dimension of access governance that existing operational metrics do not address.

The framework's application to non-human identities highlights how each property is amplified for identity types that lack natural lifecycle governance mechanisms, providing a quantitative basis for the NHI governance investment decisions organizations are increasingly confronting.

Just as DORA metrics gave DevOps teams a shared language for measuring delivery performance, a governance dynamics measurement framework can give IGA programs the quantitative foundation to move from compliance-driven cadence decisions to evidence-based governance design. This paper offers an initial formalization of that vocabulary. Empirical validation and community refinement will determine its ultimate utility.

References

[1] "Best Identity Governance and Administration Reviews 2026." Gartner Peer Insights, 2026.

[2] "Guidance for Identity Governance and Administration." Gartner, October 14, 2025. Directs identity architects to design and operate IGA for continuous, orchestrated, and intelligent access controls.

[3] Moffatt, S. "No More Snapshots: Architecting Identity for a Real-Time World." SGNL / The Cyber Hut, November 2025. Three-part series on Continuous Identity.

[4] Continuous Access Evaluation Profile (CAEP) 1.0. OpenID Foundation, Shared Signals Framework. https://openid.net/specs/openid-caep-1_0.html

[5] Gupta, V. "Optimizing Access Recertifications." IDPro Body of Knowledge 1(16), April 2025. doi: 10.55621/idpro.119

[6] Forsgren, N., Humble, J., and Kim, G. *Accelerate: The Science of Lean Software and DevOps: Building and Scaling High Performing Technology Organizations.* IT Revolution Press, 2018. ISBN: 9781942788331.

[7] "DORA: DevOps Research and Assessment." Google Cloud. https://dora.dev/

[8] "Microsoft Entra ID Governance Documentation." Microsoft, 2025-2026. https://learn.microsoft.com/en-us/entra/id-governance/

[9] "Identity Governance and Administration (IGA)." Palo Alto Networks Cyberpedia, April 2026.

[10] "How Privilege Creep Compounds in Two Directions: The Mover's Journey." Zluri, January 2026.

[11] CrowdStrike 2024 Global Threat Report. CrowdStrike, 2024.

[12] "The Role of Identity Governance and Administration (IGA) in Zero Trust Security." SafePaaS / Security Boulevard, February 2026.

[13] "The Non-Human Identity Governance Vacuum." Cloud Security Alliance (CSA), May 2026.

[14] Entro Labs H1 2025 Research (144:1 NHI-to-human ratio in cloud-native environments). Additional NHI ratio estimates from Rubrik Zero Labs and others vary by source and environment definition. Referenced in CSA and The Hacker News, 2025-2026.

[15] OWASP Non-Human Identity (NHI) Top 10. OWASP Foundation, 2025.

[16] "Identity and Access Governance: Definition and Differentiation." Apono, October 2025.

[17] "10 IGA Metrics Every Security Team Should Use to Measure Success." C1.ai, January 2026.

[18] "The IAM Metrics That Really Matter." Omada Identity, January 2026.

[19] "Leadership Compass: Identity Governance and Administration (IGA)." KuppingerCole, June 2024.

[20] "What Is Entitlement Creep? Risks and Prevention." Cyberhaven, April 2026.

[21] "How Privilege Creep Compounds in Two Directions." Zluri, January 2026.

[22] "Accelerate State of DevOps Report." Google Cloud DORA Program, annual publication. https://dora.dev/research/

[23] Digital Operational Resilience Act (DORA). Regulation (EU) 2022/2554. European Union, 2022.

[24] "Identity in the SOC: From Decision Latency to Decisive Action." CIO.com, April 2026.

[25] Arbesman, S. *The Half-Life of Facts: Why Everything We Know Has an Expiration Date.* Penguin, 2012. ISBN: 9781591846512.

[26] Rose, S., Borchert, O., Mitchell, S., and Connelly, S. "Zero Trust Architecture." NIST Special Publication 800-207. National Institute of Standards and Technology, August 2020. doi: 10.6028/NIST.SP.800-207