SafePaaS

Board-Ready KPI One-Pager: Time-to-Coverage

> **Time-to-coverage** measures how long a critical application, integration, or non-human identity remains outside policy, review, and audit-ready evidence after it goes live.
>
> For boards, CROs, and CISOs, it is one of the clearest ways to judge whether identity governance is reducing exposure or simply trailing behind it.

Why This KPI Belongs In The Board Pack

Most identity programs report activity: connected systems, completed reviews, approvals processed, workflows deployed. Those numbers do not show how long new risk remains outside governance.

Time-to-coverage does.

If a finance SaaS platform, ERP module, or critical automation goes live this quarter but remains outside formal governance for the next two, the business is operating inside a known exposure window. That delay should be visible alongside cost, audit findings, and control effectiveness.

The Four Numbers To Put On One Page

1 Current Coverage Of Critical Applications

Report the percentage of SOX-in-scope and other business-critical applications that are currently under policy, monitoring, review, and evidence.

> **Board question:** *How much of the application estate that can move money or change critical data is actually governed today?*

2 Current Time To Onboard A New Critical Application

Show the average elapsed time from identifying a new in-scope system to bringing it under an auditable governance standard.

> **Board question:** *How long does a new critical system operate before governance catches up?*

3 Estimated Time To 50% Coverage

Show how long it will take, on the current path, to bring half of critical applications under governance.

> **Board question:** *When does governance move beyond the pilot phase and start affecting the broader risk surface?*

4 Estimated Time To 80% Coverage

Show how long it will take to extend governance into the long tail of SaaS, business applications, and integrations where centralized programs often stall.

> **Board question:** *Do we have a credible path to broad coverage, or will the program flatten after the first wave?*

What The Board Should See In The Same View

- Current coverage percentage
- Current onboarding time per new in-scope system
- Time to 50% coverage
- Time to 80% coverage
- Separate visibility for human and non-human identities where possible
- Trend direction quarter over quarter

If these numbers do not exist, that is already a finding. It usually means the organization is monitoring effort, not exposure.

What Good Looks Like

A useful board-ready KPI view should make three things obvious:

- Whether the coverage ceiling is moving
- Whether onboarding is getting faster as the program grows
- Whether non-human identities are part of the governance model or still sitting outside it

If onboarding system 20 costs as much time and effort as system 1, the architecture has not scaled. If system 50 still looks like a custom project, the program is not expanding coverage fast enough to match business change.

The Board-Level Comparison That Matters

Centralized IGA Model	Federated Factory Model
• First wave of systems onboarded with heavy effort • Each additional system requires custom connector, mapping, and workflow work • Central teams remain the bottleneck • Time-to-coverage stretches as complexity grows	• First few systems establish reusable onboarding patterns • New applications follow the same connect-transform-govern path • Local owners participate within central guardrails • Time-to-coverage drops as patterns are reused

This is the difference between a connector strategy and a coverage strategy.

Why SafePaaS Changes The Curve

SafePaaS improves coverage by using a federated governance model that separates policy, risk, and evidence from any single application's implementation. Instead of rebuilding connectors, mappings, and workflows for every new system, the platform follows the same three-step pattern each time.

✓ **Connect via DataProbe** to reach target systems through a standard connectivity layer, so new applications plug into the same entry point instead of needing bespoke integration work.

✓ **Transform via DataPaaS** to convert local structures into a reusable governance schema for identities, entitlements, and high-risk actions, so the data looks familiar to the control plane no matter where it came from.

✓ **Govern via SafePaaS** to apply policies, SoD rules, review patterns, and evidence requirements in one control plane that works across the estate.

This is what changes the curve. The next application is not treated as a fresh design problem. Each onboarding pattern becomes an asset the next team can reuse, not a project they have to rediscover.

Questions The Board Should Ask

- How long does it currently take to bring a new critical system under governance to an auditable standard?
- How long will it take to reach 50% and 80% coverage on the current path?
- Are non-human identities included in those numbers?
- Does the proposed architecture reduce onboarding effort as the program grows?
- Are we funding a path that changes the coverage curve, or just extending the current bottleneck?

Decision Signal

If the program cannot show current onboarding time, projected time to 50% and 80% coverage, and separate visibility into non-human identities, then the board does not yet have a reliable view of identity governance exposure.

Time-to-coverage should not sit in an implementation dashboard. It belongs in the board pack because it shows how long the business stays exposed before governance catches up.