

## Governance Stress-Test Checklist

### Use this as a working list with your team.

The answers will tell you whether you can squeeze more out of your centralized model or whether it's time to move to a federated one.

### 1. Coverage: where do we actually govern?

- ☐ Percentage of SOX-in-scope and business-critical applications with defined identity policies, lifecycle controls, and periodic access reviews—not just in an inventory.
- ☐ Percentage of nonhuman identities (service accounts, integration users, bots, agents) in those systems with clear owners, purpose, and review patterns.
- ☐ Number of critical applications where approvals still live mainly in spreadsheets, email threads, or chat logs rather than a control plane.

*If your "governed" list is short and your shadow-governed list is long, your coverage problem is structural, not temporary.*

### 2. Time-to-coverage: how fast do we bring new risk under control?

*For your last three in-scope additions outside the original core (for example, a finance SaaS app, an integration platform, a regional workflow tool):*

- ☐ Days or weeks from "identified as in scope" to "under standard lifecycle, review, and SoD controls."
- ☐ Number of teams required: IAM, security, integration, app owners, audit, external consultants.
- ☐ Amount of bespoke work: custom connectors, scripts, schema mappings, and manual workflow designs.

*If the story reads "months, many, and a lot," your time-to-coverage won't improve without changing the operating model.*

### 3. Certification fatigue: are reviews still doing real work?

*Look at your last large access review:*

- ☐ Total number of systems and entitlements reviewed, and total number of reviewers involved.
- ☐ Percentage of approvals vs. revocations, and how many items included comments or follow-up.
- ☐ Volume and tone of feedback from reviewers—are they flagging risk, or complaining about volume and relevance?

*High completion, low change, and high frustration is the signature pattern of a centralized campaign model under stress.*

### 4. Role and SoD health: are our models helping or hurting?

- ☐ Total roles in your IGA platform and how fast that number is growing.
- ☐ Proportion of roles with only a single user or marked "temporary" with no clear retirement path.
- ☐ Number of active SoD rules, number of open exceptions, and how many of those exceptions have a named owner and next review date.

*If your roles and SoD exceptions are growing faster than your governed estate, they are carrying stress instead of managing it.*

### 5. Operating-model pressure points: who owns what?

- ☐ For a new in-scope application today, who is responsible for connectivity, data mapping, role modeling, policy assignment, and certification design—and how many of those steps lead back to the same central team or small set of experts?
- ☐ For a change in a global SoD or access policy, how do you push that change into each relevant system, and how do you confirm it actually landed?

*If the answers involve manual coordination, custom work, and hope, your current model has already failed the stress test. You're just living with the consequences.*

The point of this stress-test is not to declare your IGA a failure. It's to be explicit about what it can and cannot do under load. If your answers point to the same pattern—limited coverage, slow time-to-coverage, stressed certifications, unwieldy roles, brittle SoD—then you've hit the natural limit of a centralized model.

**At that point, you have a choice. You can keep pushing more connectors and campaigns through a structure that's already buckling. Or you can redesign the operating model around a federated control plane, local ownership, and a standard onboarding fabric that changes the economics of coverage entirely.**