THE LINUX FOUNDATION | Research

MEASURING OSPO VALUE:
A FRAMEWORK FOR ROI, RESILIENCE, RISK FORESIGHT, AND STRATEGIC INFLUENCE

A Report for Executive Leaders,
OSPO Practitioners, and Technology Strategists

Ibrahim Haddad, *PhD*

Foreword by
Ana Jiménez Santamaría, *The Linux Foundation*

June 2026

In partnership with

TODO

MEASURING OSPO VALUE: A FRAMEWORK FOR ROI, RESILIENCE, RISK FORESIGHT, AND STRATEGIC INFLUENCE

Open source powers ***products***, ***infrastructures***, ***and operations*** across modern enterprises.

OSPO performance should be judged by ***business outcomes***, not activity counts or workflow volume.

OSPO value is best measured as a ***portfolio of indicators*** that capture different forms of enterprise value.

Weak OSPO measurement leads to ***underinvestment***, ***mis-scoped governance***, and strategic ***blind spots***.

A ***credible OSPO scorecard*** should balance economics, preparedness, foresight, and strategic influence.

OSPOs create value by improving ***governance***, ***visibility***, and ***decision quality*** across the enterprise.

Resilience means the organization is prepared before disruption forces ***improvisation and expensive reaction***.

Strong OSPO governance depends on ***clear decisions and communication***, not just documentation and tracking.

Organizations gain ***more influence*** when they ***actively engage*** in the open source ecosystems they depend on.

Organizations that measure OSPO value well make ***better open source decisions*** with less friction.

AI-generated code adds new ***governance complexity*** that standard compliance workflows were not designed to handle.

Risk foresight for OSPOs means ***detecting emerging problems*** while they are still governable.

TABLE OF CONTENTS

FOREWORD

We are witnessing a fundamental shift across the technology industry: infrastructure is no longer an invisible layer operating quietly in the background. It has become a strategic determinant of organizational value, resilience, and long-term competitiveness.

Nowhere is this shift more visible than in AI. Open protocols and interoperability standards such as the Model Context Protocol (MCP) and A2A (Agent2Agent) protocol, open source agentic tools such as Goose, and emerging open source agent platforms such as Hermes and OpenClaw are becoming foundational building blocks for how organizations develop, connect, and operate AI-powered systems.

In this context, open source has become a business capability, a risk surface, a governance challenge, and a strategic environment in which organizations participate intentionally.

This makes the role of the **Open Source Program Office (OSPO)** increasingly important. As these ecosystems mature, organizations need a function that can help them understand what they depend on, how those technologies are governed, where risks may emerge, and how to participate in ways that create long-term value rather than passive dependency. OSPOs and their open source specialists are well positioned to connect these realities.

Yet the value of OSPOs has often been difficult to communicate. Too often, it has been described through activity counts: tickets closed, repositories scanned, contribution requests reviewed, or training delivered. These signals can be useful, but they do not fully explain why an OSPO matters to the business.

That is why this report is timely. It connects open source management to the capabilities organizations need to operate: cost efficiency, operational resilience, risk anticipation, and strategic influence in the open ecosystems that shape their infrastructure.

In my work with open source developer communities and practitioners responsible for open source management inside organizations, I have seen this need become increasingly urgent. OSPOs need a framework that executives can act on: one that explains not only what the OSPO does, but what organizational *value* it enables.

This need is becoming more pressing as new agentic workflows, open protocols, interoperability standards, and open source AI tools will create new dependencies and new governance questions. For OSPOs, the challenge is not only to support adoption, but to show how their work creates measurable value for the organization through better decisions, stronger resilience, clearer risk management, and more strategic ecosystem participation. This report provides a practical foundation for making that shift visible, measurable, and actionable.

Ana Jiménez Santamaría,
Senior Project Manager,
The Linux Foundation

EXECUTIVE SUMMARY

This report proposes a structured method to measure the business value of an Open Source Program Office (OSPO) across four dimensions: cost efficiency, operational resilience, risk anticipation, and strategic influence. The reference model framework links OSPO activities to measurable business outcomes through explicit causal pathways. It distinguishes between activity metrics, operational outcomes, and financial or risk-adjusted value. For example, increasing SBOM coverage improves vulnerability detection speed. Faster detection reduces mean time to remediation, which lowers the probability and impact of production incidents. As a result, the organization avoids outage costs or reduces security exposure. The framework is designed for executive decision-making. It provides a baseline model that organizations can adapt to quantify OSPO contribution, compare investment against outcomes, and integrate OSPO metrics into broader engineering and risk dashboards.

INTRODUCTION

Open source software is embedded in the majority of modern software systems. Industry studies consistently show that most commercial codebases depend on externally maintained open source components. This dependency introduces both leverage and exposure. Recent security incidents have shown that vulnerabilities in widely used components can propagate rapidly across industries. That dynamic has elevated open source from an engineering concern to a board-level risk and strategy topic.

Organizations have established OSPOs to manage such complexities. However, most OSPOs are evaluated using activity-based metrics such as contribution counts or policy adoption. These indicators do not demonstrate business value. This report addresses that gap by defining a measurement framework that connects OSPO activities to business outcomes and quantifiable impact.

As enterprise dependency on open source software deepens, the need for coordinated governance and interactions with that ecosystem deepens as well. OSPOs emerged in response to that reality. In their most basic form, they provide policies, processes, training, tooling, and compliance support. At higher maturity levels, OSPO also shape contribution strategy, developer enablement, supplier expectations, ecosystem participation, and executive visibility into open source risks and opportunities.

However, OSPO value has been historically difficult to measure for at least four reasons:

- **Much of the OSPO value is preventive:** When an OSPO identifies a licensing issue before a software release, improves a dependency intake path, or detects a governance concern early, the result is often the absence of a crisis rather than a visible, escalated, and costly event.
- **The effects are distributed:** OSPO benefits accrue across product teams, engineering operations, legal, security, procurement, and strategy rather than inside a single budget line. Capturing these effects is not a straightforward exercise.
- **The impact is multi-horizon:** Some benefits, such as shorter review cycle times, may appear within a quarter; others, such as ecosystem influence or reduced lock-in, emerge over years.
- **The work is cross-functional:** OSPO outcomes are a mix of technical, legal, operational, cultural (community), and strategic.

The issue is often the lack of language that makes the value of OSPOs clear to decision-makers and executive sponsors. A credible measurement model or an OSPO value framework must satisfy two constraints: it must be rigorous enough to support executive review and prioritization, and it must be flexible to reflect the real character of open source governance.

Some of this measurement difficulty is also structural: the value of OSPOs is preventive, distributed, and cross-functional within the organization. Some of it is organizational: many enterprises have not yet built the data pathways, reporting discipline, or cross-functional instrumentation needed to capture that value reliably. The distinction is important because it reminds leaders that imperfect visibility is not the same thing as absent impact.

THE URGENCY OF MEASURING THE VALUE OF OSPOS

Measuring the OSPO value has become a more urgent question because the operating context for open source has significantly changed in the last few years. In the following subsections, we explore five major factors pressuring OSPOs into demonstrating value across all enterprises:

- Open source now sits closer to revenue-critical systems
- Security and supply chain expectations are rising
- Tightened regulations affecting open source software
- Evolving expectations from leadership
- Added governance complexity with AI-generated code

Open source now sits closer to revenue-critical systems

As products become more software-defined, open source dependencies increasingly sit in the direct path of customer experience, release confidence, security posture, and operational continuity. This environment changes the significance of OSPO performance. Open source now affects product delivery speed, remediation quality, and the enterprise's ability to move quickly without accumulating hidden exposure to technical debt.

Security and supply chain expectations are rising

The OSPO function now covers open source compliance, contribution management, dependency transparency, vulnerability response coordination, SBOM workflows, and broader supply chain readiness. That shift in scope and responsibilities is reinforced by public guidance around SBOMs and software supply chain governance, which increasingly treats transparency and traceability as normal elements of responsible software management rather than exceptional practices.

Rise of regulations

The **European Union's Cyber Resilience Act** (EU Regulation 2024/2847), which entered into force on 10 December 2024, is the most consequential regulatory development for open source governance in recent years. It establishes mandatory cybersecurity requirements for products with digital elements sold into the EU market, with direct implications for how organizations manage open source dependencies. Specifically, it creates obligations around vulnerability disclosure timelines, SBOM transparency, conformity assessment, and the treatment of commercial open source integrators who place products on the EU market.

The CRA regulates the act of placing products that contain open source into commercial supply chains. It shifts governance expectations from upstream communities to the enterprises that integrate and ship those components. Organizations that lack documented intake processes, dependency inventories, and vulnerability response workflows face commercial and operational exposure as CRA obligations phase in. Manufacturer obligations to report exploited vulnerabilities apply from 11 September 2026, and the full set of obligations applies from 11 December 2027.

The governance infrastructure an OSPO builds, dependency visibility, SBOM practices, community collaboration, upstream relationship management, and vulnerability response coordination, is increasingly the same infrastructure that regulators, enterprise customers, and insurance underwriters are beginning to treat as a baseline expectation.

Organizations seeking to demonstrate CRA readiness through their OSPO scorecard should pay particular attention to four metrics:

- SBOM coverage rate (Appendix A.3), which maps directly to CRA's transparency obligations
- Time to awareness for upstream issues (Appendix A.3), which relates to vulnerability notification timelines
- High-risk issues identified pre-release (Appendix A.4), which demonstrates active governance, and
- Compliance cycle-time reduction (Appendix A.2), which provides evidence of a functioning intake process.

These four metrics together constitute a minimum credible evidence set for a CRA-focused governance conversation with customers, auditors, or regulators.

Evolving expectations

Executive expectations are also shifting as OSPOs become more involved in security, AI infrastructure, community collaboration, and supplier-facing governance. Leadership no longer expects the OSPO to just guide the internal adoption of open source and process license compliance tickets. Executive leadership expects the OSPO to help the enterprise govern complexity. Any OSPO value measurement model that remains trapped in narrow activity counts is therefore already outdated.

These developments increase the importance of OSPOs and increase the cost of measuring their value and impact poorly, because executive misperception now translates more directly into underinvestment, mis-scoped governance, and blind spots.

AI-generated code: an added governance complexity

AI-assisted development tools are changing how code enters the enterprise, and the governance questions they raise do not fit neatly into existing intake workflows. The 2025 TODO Group and Linux Foundation survey reports that 79% of OSPOs now rate themselves effective at managing generative AI risks, up from 65% in 2024. That signals real progress, but readers should note these figures are self-reported and no external benchmark for OSPO AI-risk management has yet been published.

The core concerns with AI-generated code are provenance and license contamination. Large language models trained on public code repositories may reproduce patterns, idioms, or fragments whose license terms are unclear or disputed. The legal status of AI-generated code varies by jurisdiction and remains unsettled. For OSPOs, this landscape creates a new category of compliance exposure and most current SCA tooling does not reliably detect it.

The practical implication for measurement is that both the resilience and foresight dimensions should include indicators specific to AI code governance: whether the organization has a policy governing AI-generated code in products, whether developers are trained on that policy, and whether there is a review pathway for AI-assisted contributions before they enter production codebases.

DESIGN PRINCIPLE FOR MEASUREMENT

A useful OSPO measurement portfolio must span short-term economics, structural preparedness, anticipatory detection, and long-horizon ecosystem position rather than overconcentrating on any one of them. However, these dimensions should not be treated as standalone categories or as a fixed list of metrics. Following the Goal-Question-Metric approach used by the OSPO Metrics WG, OSPOs are recommended to first clarify the organizational goals they want to support, identify the executive and operational questions they need to answer, and then select metrics that provide meaningful evidence for decision-making.

The core purpose of OSPO measurement is to communicate impact: how open source work creates economic, operational, strategic, and ecosystem value for the organization. This includes communicating that impact across different teams, such as engineering, product, legal, security, procurement, community, and executive leadership.

In this sense, the four dimensions proposed in this report below can be understood as complementary lenses for communicating OSPO impact. This is conceptually similar to the way leaders evaluate cybersecurity, platform reliability, architecture health, or product quality. That design principle leads directly to the framework proposed in this report.

CHAOSS Practitioner Guide

The **CHAOSS Practitioner Guide Series** provides practical resources to help OSPOs maximize the value of their open source programs. The guides cover key areas including demonstrating organizational impact, assessing dependency health and sustainability, managing project archival and retirement, and measuring the outcomes of open source funding initiatives. Additional guidance supports the improvement of open source projects across contributor sustainability, community responsiveness, organizational engagement, security, and leadership diversity.

THE FOUR-DIMENSION OSPO VALUE FRAMEWORK

There isn't a single metric that can capture the impact of an OSPO. Effective OSPO value measurement requires a structured portfolio of evidence. The OSPO value framework proposed in Figure 1 and detailed in Table 1 organizes the OSPO value into four dimensions. Each dimension is defined through a causal chain linking OSPO activity to business outcomes.

- **ROI and cost avoidance:** OSPO standardization reduces duplication and rework, which improves engineering efficiency and lowers development and maintenance cost.
- **Resilience:** Dependency visibility and governance improve issue detection and remediation speed, which reduces incident frequency and duration and improves service continuity.
- **Risk foresight:** Continuous monitoring of dependencies enables earlier awareness of vulnerabilities, which reduces the exposure window and lowers the probability and impact of security incidents.
- **Strategic influence:** Active upstream participation increases influence over project direction, which enables faster adoption of required capabilities and reduces the need for internal development.

FIGURE 1
OSPO VALUE FRAMEWORK

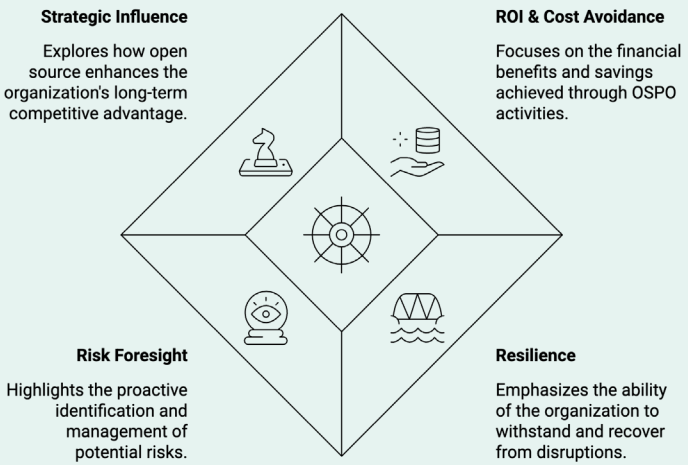

To avoid attribution debates, it helps to state the value logic explicitly.

OSPO mechanisms
(policy and decision rights, standardized processes, workflows and tooling, contribution governance, supplier expectations, escalation design, and internal enablement)

strengthen enterprise capabilities
(visibility, standardization, response speed, and cross-functional alignment),

which in turn drive outcomes
(reduced avoidable cost, improved continuity, earlier intervention, and greater ecosystem agency).

The scorecard should therefore be read as evidence that these capability pathways are strengthening, not as a claim that the OSPO alone owns all downstream results. A useful test of the pathway: if OSPO maturity rises but the dimension's outcome metrics stay flat for two consecutive periods, the pathway is not yet operating as designed and the OSPO should investigate whether artifacts and processes are actually being consumed in decisions.

TABLE 1

KEY CHARACTERISTICS OF THE DIMENSIONS IN THE OSPO VALUE FRAMEWORK

Value dimension	Core question the dimension addresses	Primary audience	Measurement horizon	Typical evidence	Appendix section
ROI & Cost Avoidance	What would we spend without the OSPO?	CFO, Finance, COO	Monthly / Quarterly	Savings estimates, cycle-time improvement, and avoided external spend	A.2
Resilience	How prepared are we to withstand supply chain and governance shocks?	CTO, CISO, Engineering Leadership	Quarterly / Annual	Coverage, dependency visibility, and readiness indicators	A.3
Risk Foresight	What material issues did we detect early enough to keep governable?	Legal, Security, Board, Risk Committees	Quarterly / Annual	Early warnings, pre-release interventions, and avoided escalations	A.4
Strategic Influence	How does open source help the organization achieve their product goals, and more broadly support the organization execute on its mission and vision? How does open source strengthen our long-term leverage and optionality?	CEO, Strategy, Product Leadership	Annual / Multi-year	Ecosystem role, standards participation, roadmap alignment, dependency leverage	A.5

Accountability across the OSPO value framework

Before applying this framework, organizations should be explicit about what the OSPO owns versus what it enables. Conflating the two is a common source of scorecard inflation and cross-functional friction.

The OSPO is accountable for the measurement system design, the governance mechanisms (policies, standard workflows, decision rights, escalation paths, and reusable artifacts), and the cross-functional coordination that enables outcomes. It is not accountable for every outcome that those mechanisms touch.

Legal is accountable for legal interpretation and legal risk acceptance decisions. The OSPO provides standard intake paths and reduces the volume of recurring matters requiring counsel; legal owns the decisions that remain.

Security is accountable for vulnerability response and security risk posture. The OSPO enables dependency transparency, upstream awareness, and consistent supply chain governance interfaces; security owns the response.

Engineering, product, and architecture are accountable for implementation choices and operational execution. The OSPO enables coherent options, escalation routes, and reusable artifacts; those teams own the build and release decisions.

Procurement and suppliers are accountable for commercial terms and supplier compliance where applicable. The OSPO helps define open source requirements and evidence expectations; procurement owns the contractual relationships.

Reading the four-dimension scorecard without this accountability map produces attribution errors in both directions: the OSPO either claims too much credit or absorbs blame for outcomes it could not have owned.

Each dimension captures a different mechanism by which OSPOs enable enterprise value:

- ROI and cost avoidance measures whether the OSPO reduces unnecessary spend, friction, and waste.
- Resilience measures whether the organization is structurally prepared to absorb open source and supply chain disruption.
- Risk foresight measures whether the organization can detect emerging issues early enough to act before they harden into incidents.
- Strategic influence measures whether the organization is building leverage, optionality, and a more intentional position in the ecosystems it materially depends on.

Resilience and risk foresight, for example, are related but not interchangeable. Resilience is about the state of preparedness: the extent to which the organization has visibility, controls, artifacts, and response paths in place. Risk foresight is about early detection and anticipatory intervention: the extent to which the organization can see emerging issues while they are still governable. One dimension asks whether the enterprise is built to withstand disruption; the other asks whether it can recognize disruption early enough to change the outcome.

Note on AI-generated code

As of 2025, risk foresight should include coverage of AI-generated code governance as a distinct signal category, separate from standard license and vulnerability detection. The detection methodology differs because standard SCA tooling was not designed for this class of exposure.

A worked illustration: reading all four dimensions from a single event

An OSPO that monitors upstream project health detects that the sole maintainer of a cryptographic library used in three revenue-critical products has not merged a pull request or responded to issues in for example 47 days, while new unresolved pull requests and new open issues are accumulating at a rate of several per day. The library has no co-maintainer listed and no succession documentation.

Reading this through the four dimensions:

- **Under risk foresight, this is an early warning.** The organization has detected a governance instability before it becomes a support or security incident. The foresight metric records one high-risk upstream issue identified with a 47-day detection lag from the first signal.
- **Under resilience, the situation exposes a gap.** The dependency has no assigned owner, no documented mitigation path, and no fork or replacement plan. The resilience metric reflects a critical dependency without a designated escalation path, which is exactly what the coverage metric should surface.
- **Under ROI, if the OSPO intervenes now, the cost is measurable.** The engineering hours required to identify an alternative library, engage the project community, or initiate a co-maintainership contribution. As an illustrative range, 40 to 80 hours of senior engineer time. The counterfactual is an emergency migration under time pressure after the library is abandoned or compromised, which is typically four to ten times more expensive across engineering, security review, and release re-qualification, plus delay costs if a release slips. The avoided-disruption value is the difference between those two scenarios, discounted by the probability that the library actually fails without intervention. The figures here are illustrative ranges, not industry benchmarks; organizations should anchor them to their own incident records or to upstream maintainer-health signals such as those tracked by CHAOSS or the OpenSSF Critical Score. Even with a probability of failure as low as 20 to 30 percent, the expected value of early intervention typically exceeds its cost by a multiple. This is the logic the OSPO should document, not just the action.
- **Under strategic influence, this is an opportunity.** If the organization has the engineering capacity to take a co-maintainer role in a library it already depends on, it converts passive exposure into dependency leverage. That is a strategic outcome with a clear pathway.

This single event does not dominate any one dimension's scorecard. But reading it across all four shows how the framework works as a diagnostic system rather than a checklist.

DIMENSION ONE: ROI AND COST AVOIDANCE

This section frames ROI and cost avoidance as the economic outcomes of better open source governance: reduced avoidable spend, lower external dependency, fewer late-stage disruptions, and efficiency gains that compound across teams. Because much of this value is preventive or distributed, it should be measured through disciplined estimation with explicit assumptions rather than false precision. In this dimension especially, credibility improves when reporting distinguishes the value the OSPO directly creates from the value it enables across engineering, legal, security, and procurement.

The importance of this dimension

For many executive audiences, the first question remains economic: if the organization invests in an OSPO, what does it get back? That question is legitimate, but it is often framed too narrowly. OSPOs rarely generate attributable revenue in a direct and isolated sense. Their financial contribution is more often visible through cost avoidance, efficiency, and reduced waste. This is one of the most credible ways to express value in enterprise settings. Finance leaders routinely evaluate programs based on avoided spend, reduced external dependency, improved process efficiency, and fewer delivery disruptions.

What belongs in this dimension?

ROI and cost avoidance should capture economic outcomes such as governed substitution of higher-cost alternatives, reduced reliance on external legal review for recurring open source matters, lower remediation costs because issues are found earlier, less release disruption caused by late-stage surprises, and less duplicated effort because teams follow standard guidance and approved paths.

It should also include efficiency gains that are small in isolation but meaningful in aggregate, such as faster intake workflows, clearer decision rights, and reusable artifacts that reduce repeated work across teams.

How to measure ROI and cost avoidance?

The most effective financial reporting for OSPOs relies on disciplined estimation rather than false precision. Executives will generally accept approximations when the assumptions are clear, the methodology is stable, and the claims are modest.

A practical way to keep estimates credible is to anchor them to a baseline and an explicit counterfactual. An organization can use historical remediation and release delay records, sampled products, repeat external counsel patterns, and time-spent analysis on recurring governance work to establish what they typically pay today. Then separate what governance is controllable from what is structural. This approach keeps ROI claims disciplined and reduces debates about attribution.

If Finance will not validate an avoided cost estimate, report it as an operational proxy with transparent assumptions, not as a financial claim. This approach keeps the scorecard credible while still supporting prioritization decisions.

Differentiating OSPO-created vs OSPO-enabled value

An important aspect to consider is creating a distinction between the value the OSPO directly creates and the value it enables across the organization. OSPO-created value is best reserved for outcomes the OSPO owns end to end, such as

a standardized intake path, a reusable compliance artifact pipeline, or a governance decision that replaces repeated ad hoc review. OSPO-enabled value describes outcomes delivered by engineering, legal, security, procurement, or product teams, where OSPO mechanisms made the outcome faster, safer, or cheaper.

For example, an OSPO may enable a release delay prevention by clarifying license options early, but the product team and release owners still execute the remediation. Similarly, an OSPO may enable reduced outside counsel spend by creating templates and escalation routes, while legal operations realize the spend reduction. Treating these as enabled value keeps attribution honest and tends to increase trust.

Common mistakes

Organizations tend to make four recurring errors in this dimension: double-counting savings already recognized elsewhere, treating all open source benefits as OSPO-generated benefits, presenting optimistic estimates without explicit assumptions, and ignoring hidden support or maintenance costs when making substitution arguments.

A related guardrail is to avoid metrics that improve simply by changing accounting treatment or shifting work to other teams. If savings are real, the assumptions, sources, and cost owners should remain visible and stable over time.

DIMENSION TWO: RESILIENCE

This section defines resilience as the organization's preparedness to absorb open source and software supply chain disruption without improvisation or late-stage disruption. It focuses on structural readiness, including dependency visibility, ownership, standardized artifacts, escalation paths, and decision rights, rather than on whether issues occur. Resilience is therefore distinct from risk foresight: resilience measures readiness posture, while foresight measures early warning and timely intervention.

From governance to continuity

If ROI explains why the OSPO matters to finance, resilience explains why it matters to engineering and security leadership.

Modern software supply chains are fragile in ways that are easy to underestimate:

- Critical packages may be maintained by small teams.
- Governance may be uneven across essential upstream projects.
- Transitive dependencies can hide risk.
- Build systems can accumulate invisible complexity.
- Supplier visibility may be incomplete.

Under those conditions, resilience becomes a strategic property. The OSPO contributes to resilience by helping the organization see, structure, and govern this complexity. It creates visibility into dependencies, strengthens artifact quality, supports consistent control paths, and clarifies what the enterprise will do when upstream or supply chain conditions deteriorate.

What belongs in this dimension?

Resilience should be understood as the organization's preparedness posture. It includes dependency visibility, SBOM coverage, ownership of critical dependency clusters, standard artifact generation, escalation paths for licensing or provenance concerns, and decision readiness around patching, contributing, replacing, forking, or escalating supplier issues.

SBOM coverage as a metric requires careful definition before it is useful. Three qualifications matter in practice.

1. **Format:** SBOMs generated in SPDX and CycloneDX are not interchangeable across all downstream consumers, and organizations should specify which formats they generate, for which consumers, and whether those consumers can actually ingest the output.
2. **Freshness:** an SBOM generated at release is accurate at that moment. If the underlying dependencies change during a product's operational life and the SBOM is not updated, high coverage scores can coexist with operational blindness.
3. **Consumption:** an SBOM that exists but is not referenced in procurement decisions, vulnerability response workflows, or supplier conversations is an artifact, not a capability.

Coverage metrics should therefore be paired with a short operational test: in the last quarter, how many times was an SBOM actually used to make or accelerate a decision?

To make this test operational, organizations should designate a single tracking point for SBOM-referenced decisions, whether that is a field in the release checklist, a tag in the vulnerability response ticket, or a line in the procurement sign-off record. The

question requires that someone is responsible for noting, each time an SBOM is consulted, what decision it informed. Twelve such records per year is credible evidence of operational use. Zero records are evidence of an artifact practice, regardless of coverage score.

Resilience is therefore the presence of a structure that makes risk more governable when stress arrives. Resilience metrics are usually leading indicators rather than lagging financial outcomes. Useful measures include the percentage of products with current SBOMs, the percentage of critical products whose key dependencies are identified and assigned owners, the time it takes responsible stakeholders to become aware of relevant upstream issues, the percentage of releases covered by standard compliance artifacts, and the share of higher-risk dependencies with explicit mitigation paths.

How to measure resilience?

To prevent false confidence, we should interpret these indicators with a critical systems-first lens: prioritize the most safety, revenue, and release critical products, and be explicit about scope. High coverage across the long tail can coexist with serious gaps in the systems that matter most, so executive reporting should make scope visible.

As a guardrail, we should avoid treating higher coverage numbers as proof of readiness unless the artifacts are current, used in decision-making, and connected to ownership and escalation paths. Otherwise, the metric can become a compliance artifact rather than an operational capability indicator.

Resilience is often undervalued

A resilient open source posture does not eliminate dependency risk. It reduces the likelihood that the organization is surprised by it, paralyzed by it, or forced into expensive late-stage decisions.

Resilience is easy to underinvest in because it shows up only when it fails. When readiness work is done well, nothing dramatic happens. Releases proceed as planned, and issues are handled early. That success looks like normal operations. This is why resilience metrics matter. They give leaders evidence of preparedness before a disruption forces an expensive test.

DIMENSION THREE: RISK FORESIGHT

This section explains risk foresight as the organization's ability to detect emerging open source issues early enough to change outcomes, before they become incidents, release disruptions, licensing issues, or escalations. Unlike resilience, which measures preparedness posture, foresight measures early warning and timely intervention. It is best evidenced through a combination of leading indicators and short, disciplined near-miss narratives that show what was detected, what decision changed, and what was likely avoided.

Some OSPO value is not captured fully by resilience alone. The OSPO frequently acts as an early-warning mechanism: identifying a license incompatibility before release, recognizing when an upstream project's governance is weakening, surfacing a policy issue before it escalates, or detecting patterns of process drift while corrective action is still practical. This function is better described as risk foresight than simply risk management. It is about both handling known risk categories and recognizing weak signals early enough to change outcomes.

What belongs in this dimension?

Risk foresight includes pre-release detection of problematic licenses or unapproved terms, early identification of regulatory or policy shifts relevant to open source use, recognition of upstream instability that may create future support or security issues, detection of teams bypassing approved pathways, and escalation of provenance or supply chain concerns before they become incidents.

In practice, foresight signals typically come from multiple sources, such as upstream security advisories and vulnerability feeds, SBOM or dependency-diff monitoring, maintainer and project-health signals, policy and regulatory tracking, supplier disclosures, and internal patterns. Mature foresight is therefore partly about coverage of signal channels, and not just the number of issues found. In this context, it is worth flagging the "**CHAOSS Practitioner Guide: Assessing Viability**", developed primarily by Gary White (Principal Engineer, Verizon OSPO), as part of the CHAOSS OSPO WG. The guide is intended primarily for OSPOs and other teams within organizations that need to understand the viability and risks associated with the open source software that they are consuming.

How to measure risk foresight?

Preventive value is inherently difficult to measure. It is rarely possible to prove with certainty what would have happened had the OSPO not intervened. But this does not make measurement impossible. It means the evidence model must combine counts with judgment and a short narrative.

Useful indicators include the number of high-risk issues identified pre-release, the severity mix of those issues, the trend in late-stage discoveries, the number of issues redirected before formal escalation, the number of policy advisories issued in response to external change, and the time between issue detection and decision.

A common interpretation error is to treat more issues found as worse performance. In foresight, a short-term increase can indicate better detection and healthier disclosure. The more meaningful executive signals are severity mix, late-stage discovery trend, and time to decision.

A near-miss is "validated" only when all three of the following conditions are met. First, the issue would have caused a measurable downstream impact had it not been intercepted: a release block, a security exposure, a license violation, or a regulatory finding. Second, the interception is documented

with a named owner, a date, and the decision taken. Third, the counterfactual impact is estimated with a stated assumption set rather than asserted. Counts that fail any of the three are useful operational telemetry but should not be reported as validated near-misses on the executive scorecard.

Foresight vs. resilience

The distinction between resilience and foresight is worth citing. Resilience asks whether the enterprise is structurally prepared for open source shocks. Risk foresight asks whether it can see emerging problems soon enough to intervene before those shocks mature. One measures preparedness; the other measures anticipatory detection and action.

If resilience is the organization's capacity to withstand disruption, risk foresight is its capacity to detect disruption while it is still governable.

DIMENSION FOUR: STRATEGIC INFLUENCE

This section defines strategic influence as the long-horizon value created when an organization moves from passive dependence on open source to deliberate ecosystem engagement that increases agency, leverage, and technology optionality. Because this value is often directional rather than transactional, it is best evaluated through a mix of concrete engagement evidence (representation, alignment, and collaboration) and concise executive narrative over annual and multi-year timeframes.

The most mature OSPOs do more than reduce risk and support compliance. They shape how the organization participates in the ecosystems it depends on. They influence standards, contribution priorities, dependency strategy, supplier expectations, and the internal posture toward emerging technologies. They help move the enterprise from passive consumption to intentional participation.

This matters because open source is a sourcing model and a strategic environment in which visibility, contribution, and governance participation affect long-term leverage. OSPOs become strategically valuable when they help their enterprise shape the ecosystems that affects its future, not just consume them.

Three strategic mechanisms

Strategic influence should be grounded in concrete business mechanisms rather than abstract claims of thought leadership. Three mechanisms are especially important.

1. **Dependency leverage:** The organization has more agency when it is meaningfully engaged in the projects, communities, and governance bodies behind critical dependencies rather than merely exposed to them.
2. **Technology optionality:** A stronger open source posture can improve the organization's ability to adopt, shape, substitute, or exit technologies with less lock-in and greater confidence.
3. **Institutional positioning:** Participation in relevant foundations, standards bodies, and working groups gives the enterprise a seat where future technical and governance decisions are made.

The strategic cost of passive dependence is reduced agency: the organization becomes more exposed to technical and governance decisions made elsewhere without sufficient influence over the conditions that shape them.

Strategic influence in practice

In practical terms, strategic influence may include: representation in relevant foundations or steering groups, contribution strategies tied to product roadmaps, better access to ecosystem information and collaboration channels, reduced vendor lock-in through stronger internal capability, improved employer credibility among senior engineers, and faster paths to adopting technologies that mature in open ecosystems.

Strategic influence also depends on internal alignment among the OSPO, product strategy, architecture, procurement, and executive leadership. External participation alone rarely creates durable leverage.

Why is this dimension harder to quantify?

Strategic influence is usually directional rather than transactional. It is better evaluated over longer horizons and through mixed evidence: representation, roadmap alignment, visible ecosystem participation, dependency leverage, and concise leadership narrative explaining why a given position mattered. The right standard here is disciplined modesty. Be

concrete about the influence that can be described. Avoid inflated claims. Connect ecosystem engagement to specific organizational outcomes.

To keep this dimension evidence-based, the organization should apply a three-tier test before claiming influence rather than participation.

1. **Tier one is presence:** the organization holds a membership, attends meetings, or makes occasional contributions that are not tied to any internal priority.
2. **Tier two is engagement:** the organization has named contributors actively working in the project, their work is linked to a product or platform roadmap, and there is internal sponsorship with a defined accountability owner.
3. **Tier three is influence:** the organization can point to a specific outcome where its participation changed a technical decision, governance outcome, or standard that affected its downstream cost or risk.

At least one of the following must be documentable: a design proposal that was accepted and reduced a future maintenance burden, a governance role that altered the conditions under which a critical dependency operates, standards language that the organization shaped and that applies to its products, or a project roadmap change that served the organization's platform interests.

The reporting test is simple: if you remove the organization's participation entirely, would the ecosystem outcome have been materially the same? If yes, then you would report it as presence. If the honest answer is probably not, then you would report it as influence with the specific outcome attached or described. Strategic influence is the dimension most vulnerable to inflation, and a single overclaim here damages the credibility of the entire scorecard.

For tier-three claims, the claim should be corroborated by at least one of the following: a named maintainer or TSC member outside the organization who can confirm the contribution's significance, a documented project record (accepted proposal, merged design document, published standard text) that is publicly verifiable, or a formal acknowledgment by the relevant foundation or working group. Internal narratives alone are insufficient for tier-three claims and should be labeled as tier-two engagement until external evidence is available.

As a guardrail, it is highly recommended to avoid treating counts (commits, memberships, meetings attended, conference presentations) as evidence of strategic influence unless you can also articulate the outcome pathway and why it mattered to the business.

PRINCIPLES FOR BUILDING A CREDIBLE OSPO VALUE MEASUREMENT SYSTEM

Before defining metrics, organizations should agree on a few design principles.

1. **Measure outcomes, not just activity:** Reviews completed, training delivered, repositories scanned, and policies published may all be useful operational indicators. But they are not evidence of enterprise value. They should be treated as supporting measures rather than headline indicators.
2. **Prefer a small number of stable indicators:** A publishable scorecard should be legible. Two to four metrics per dimension are usually sufficient for executive review. The goal is decision support, not metric abundance.
3. **Combine quantitative and narrative evidence:** OSPO value becomes clearer when indicators are paired with short interpretive text. In domains such as risk foresight and strategic influence, narrative is not a weakness; it is part of the evidence model.
4. **Be explicit about assumptions:** Where cost avoidance or avoided disruption is estimated, the assumptions should be documented. Transparent estimation builds trust; hidden estimation erodes it.
5. **Distinguish enabled value from owned value:** A credible scorecard should avoid implying that the OSPO directly "creates" all the value associated with open source. In most organizations, the OSPO enables value by improving how the enterprise governs, adopts, contributes to, and de-risks open source. Measurement should reflect that enabling role. When an outcome metric improves (MTTR drops, audit prep time falls, late-stage discoveries decline), name at least one alternative cause that could have produced the same movement (concurrent tooling change, security-team initiative, vendor SLA shift) and state how the OSPO contribution is being isolated from those alternatives.
6. **Avoid metrics that punish disclosure:** If teams believe surfacing issues will worsen their performance metrics, the system will produce blindness rather than improvement. OSPO measures should reward earlier visibility, not create incentives to hide problems.
7. **Design for maturity:** Organizations should start with the data they can collect reliably and improve over time. A good measurement program is iterative.
8. **State the limits of the framework:** This framework does not claim that we can monetize all OSPO values. It does not claim perfect attribution. It does not replace detailed engineering, security, or legal reporting. And it does not eliminate executive judgment. Its purpose is narrower and more practical: to provide a structured way of assessing how open source governance affects enterprise outcomes.
9. **Protect metric continuity:** Definitions should change deliberately and rarely. If a metric's logic, denominator, or source basis changes, that change should be documented and versioned visibly. Otherwise, the dashboard risks appearing to improve simply because the measurement system moved. When the OSPO's organizational scope changes materially, whether through restructuring, expansion, or merger with another function, the scorecard should include a versioning note explaining what changed and how much historical comparability is affected. Continuity of measurement should be a design principle for organizational decisions, not only a reporting principle.

EXECUTIVE REPORTING MODEL

The most effective OSPO reporting model is a balanced executive scorecard delivered on a regular cadence. The scorecard should include a concise one-page summary of the four dimensions, a small set of trend indicators for each, short commentary explaining movement and implications, and selected case notes where a near-miss or intervention materially clarifies the numbers.

For executive usefulness, each indicator should be paired with a simple decision loop: what the metric means, what actions are available if it improves or deteriorates, who owns the action, and what time horizon is expected for effect. This approach keeps the scorecard from becoming descriptive reporting and turns it into a governance instrument.

Audience-specific tailoring

Different executive audiences need different emphases even when the underlying data is the same.

- **CFO / COO:** cost avoidance, cycle-time improvement, reduction of external spend, and release economics
- **CTO / Engineering Leadership:** dependency transparency, release readiness, contribution leverage, and operational continuity
- **CISO / Risk Leadership:** pre-release issue discovery, time to awareness, control coverage, and governance quality
- **CEO / Strategy / Board:** ecosystem leverage, regulatory readiness, strategic contribution alignment, and long-term capability building

Audience-specific tailoring should change emphasis, not the underlying truth conditions of the scorecard; the objective is coordinated cross-functional interpretation, not parallel narratives for different stakeholders.

- **OSPO:** accountable for the measurement system design, cross-functional coordination, and the governance mechanisms that enable enterprise outcomes.
- **Legal:** accountable for legal interpretation and legal risk acceptance decisions (for example, licensing positions and exceptions), with the OSPO providing standard paths and intake discipline.
- **Security:** accountable for security risk posture and vulnerability response, with the OSPO enabling dependency transparency, upstream awareness, and consistent supply chain governance interfaces.
- **Engineering, product, and architecture:** accountable for implementation choices (patch, replace, fork, contribute) and operational execution, with the OSPO enabling coherent options, escalation routes, and reusable artifacts.
- **Procurement and suppliers:** accountable for commercial terms and supplier compliance where applicable, with the OSPO helping define open source requirements and evidence expectations.

Sample scorecard

Table 2 below presents a sample OSPO quarterly scorecard populated with illustrative values. Status thresholds are at the dimension level. All values are illustrative. Organizations should calibrate targets to their size, industry, and OSPO maturity stage before adopting this format. They also should establish baseline values and set targets at the start of the measurement year rather than adopting the illustrative figures here.

There are two details in the table worth flagging. SBOM coverage and SBOM freshness appear as separate rows because they measure different things: an organization can have SBOMs for nearly all production dependencies while still letting most of them go stale. Both matter for EU CRA readiness, and collapsing them into a single metric obscures which problem needs fixing. Similarly, the validated near-miss count looks low at three. That is intentional. Only near-misses meeting all three conditions are counted; inflating the number by including informal reports that don't meet the threshold defeats the purpose of the metric.

TABLE 2
SAMPLE OSPO SCORECARD

Dimension	Metric	Current	Target	Status	Trend	Owner
ROI & Cost Avoidance	Internal reuse rate	34%	45%		▲	Engineering Platforms
	License remediation spend	$250,000	$150,000		►	OSPO / Legal
	Audit preparation time (days)	11	5		▼	OSPO
	Approved component adoption rate	61%	80%		▲	Engineering
Resilience	SBOM coverage (production deps)	78%	95%		▲	Security / OSPO
	SBOM freshness (updated <90 days)	64%	90%		►	Security
	Mean time to patch critical CVEs	9 days	5 days		▼	Security
	Critical single-maintainer deps	14	<5		►	OSPO

Dimension	Metric	Current	Target	Status	Trend	Owner
Risk Foresight	Validated near-misses (quarter)	3	4+		▲	OSPO
	CVEs identified pre-disclosure	2	4+		▲	Security
	AI code policy coverage (teams)	55%	100%		▲	Engineering
	CRA readiness score (self-assessed)	62/100	80/100		▲	Legal / OSPO
Strategic Influence	Tier-1 presence (strategic projects)	8	10		▼	OSPO
	Tier-2 governance roles held	4	6		▲	OSPO
	Tier-3 external corroboration events	2	2+		▶	OSPO
	Dependency influence ratio	0.31	0.40		▲	Engineering

Legend:

▲ Increasing ▶ Flat ▼ Decreasing

performance at or above target

performance within 20% of target

performance more than 20% below target or no meaningful progress over the quarter

Note on CRA readiness:

The score in this row is self-assessed and should be read as a preparedness indicator, not as audit-grade evidence. When an external rubric becomes available (for example, future OpenChain guidance), align this score to that rubric and remove "self-assessed" from the row.

Note on Tier-1 presence:

This is a count metric and is reported as context, not as evidence of strategic influence. Influence claims should reference the Tier-2 and Tier-3 rows.

MEASUREMENTS EVOLVE AS THE OSPO EVOLVES

Organizations should not attempt an advanced scorecard before they have the foundations to support it. Measurement maturity generally progresses across four stages, and each stage has a different reporting objective and a different risk of overreach (see Table 3 and Figure 2).

TABLE 3
EVOLUTION OF MEASUREMENTS AS OSPOS EVOLVE

Stage	Description	Reporting	Objectives	Avoid
OSPO Stage 1: Foundational	The OSPO is building policy baselines, intake processes, review pathways, and basic education	Reporting is mostly activity-based.	Visibility and consistency	Avoid demanding fully formed ROI proof before the OSPO function has stable processes or data
OSPO Stage 2: Operational	The OSPO can now report cost-avoidance proxies, coverage levels, cycle times, and basic risk indicators.	Data ownership is clearer, and some dashboards emerge leading to improved reporting.	Repeatability	Avoid adding too many unstable metrics too early
OSPO Stage 3: Strategic	The OSPO begins connecting its work to product roadmaps, dependency criticality, contribution priorities, and resilience outcomes.	Reporting is more outcome-focused	Business relevance	Avoid confusing ecosystem activity with strategic influence
OSPO Stage 4: Transformational	The OSPO operates as a recognized enterprise capability with measurable influence on ecosystem position, governance posture, and strategic foresight.	Reporting integrates financial, operational, and strategic lenses.	Enterprise integration	Avoid rhetoric that outruns evidence

This model describes measurement maturity, which does not always track governance maturity at the same pace or in the same order. An organization can have sophisticated open source governance, well-designed policies and processes, consistent contribution practices, and strong upstream relationships, while its measurement infrastructure remains at Stage 1 because it has not invested in data collection, reporting discipline, or cross-functional instrumentation. The reverse is also possible: a modest OSPO with a well-instrumented dashboard can appear more mature than it is because the numbers are clean.

FIGURE 2
CORE MEASUREMENTS PER OSPO STAGE

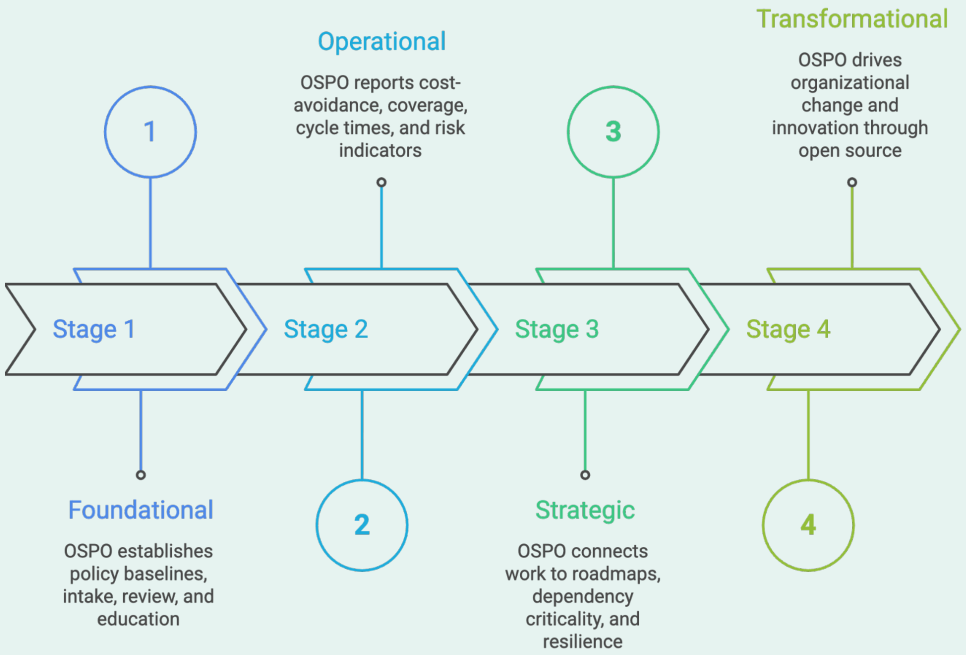

Leaders should therefore use the model diagnostically across two axes. The first is governance capability: how mature are the underlying practices? The second is measurement capability: how reliably can the organization observe and report those practices? Mismatches between the two are common and worth naming explicitly, because the remediation strategies are different.

In practice, maturity is not always uniform across all dimensions of the framework. An organization may be relatively advanced in compliance visibility and resilience indicators while remaining early-stage in strategic influence or foresight reporting. The purpose of the model is, therefore, diagnostic rather than strictly sequential. The goal is not to rush to Stage 4. It is to build a reporting model that reflects actual organizational maturity rather than aspirational language.

PRACTICAL ROADMAP FOR IMPLEMENTATION

A practical implementation sequence is straightforward (see Figure 3):

1. **Align with the executive leadership:** Before selecting metrics, it is important to have an understanding of the type of insights the leadership is looking for and engage with them on the metrics you plan on adopting leading to the desired insights. This process is typically an interactive one with an ongoing feedback loop.
2. **Select two to three metrics per dimension:** Start small. Choose indicators that are material, understandable, stable, and collectible.
3. **Define ownership and calculation rules:** Every metric should have an owner, a calculation note, a cadence, a documented caveat, and an identified source system or source workflow. Where manual collection or judgment is involved, that should be made visible so that leaders understand the confidence and limitations of the measure.
4. **Pilot the dashboard for one or two quarters:** Treat the initial reporting period as calibration. Use it to identify weak definitions, poor data quality, or metrics that create unhelpful incentives.
5. **Add narrative case notes:** During the pilot, treat any sustained metric movement as a prompt for action, not just interpretation. If a measure deteriorates for two periods, require an owner, a root-cause analysis, and a dated corrective action. If a measure improves, capture what changed so the organization can repeat it. Especially for foresight and strategic influence, a short case note often communicates more value than another numeric field.
6. **Review the system annually:** As the OSPO matures, the scorecard should evolve. Updates include retiring measures that no longer matter, and adding measures that better reflect current executive concerns.

FIGURE 3

STEPS TO FOLLOW WHEN IMPLEMENTING THE OSPO VALUE FRAMEWORK

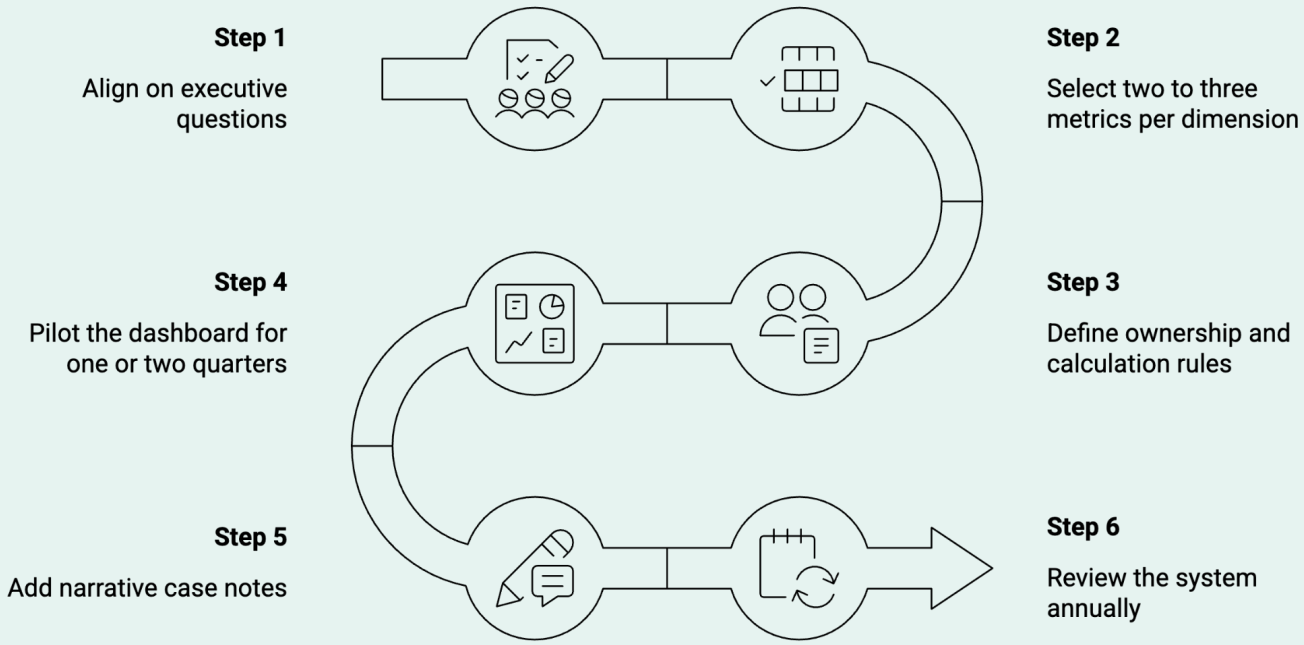

In a focused implementation with executive sponsorship and an OSPO team of three or more people, Steps 1 through 3 can typically be completed in four to six weeks in organizations with prior cross-functional governance alignment. Where those working relationships are still forming, add four to eight weeks for the stakeholder alignment in Step 3, which is usually the critical path. The first pilot quarter (Step 4) requires three months by definition. Steps 5 and 6 are ongoing. A credible first executive scorecard, at Stage 1 to Stage 2 maturity, should therefore be achievable within five to seven months of starting. Reaching Stage 3 reporting (outcome-focused, multi-dimensional) typically takes longer and depends on cross-functional data instrumentation. In larger organizations with fragmented data ownership, double the timeline. These figures reflect the author's experience advising organizations through this work, not benchmark data from a published study. The most common failure mode is not technical difficulty; it is waiting for perfect data before publishing anything. Publish a clearly labeled first-draft scorecard at the end of the pilot quarter, even if several metrics carry explicit uncertainty ranges. An imperfect dashboard that ships builds more organizational credibility than a perfect one that does not.

CONCLUSION: REFRAMING THE OSPO NARRATIVE

Three years after implementing the framework in this report, two types of organizations emerge. The first organization can walk a board through a populated scorecard, link a compliance near-miss to a structural change, and quantify what upstream influence saved them in avoided rework. The second organization is still relying on activity counts and hoping that the headcount justifies itself. The difference is whether the OSPO built measurement infrastructure before it needed to defend itself.

The four dimensions in this framework are not independent. For example, ROI without resilience produces an organization that looks efficient until a critical dependency goes unmaintained. Resilience without risk foresight produces one who responds well but never anticipates. Risk foresight without strategic influence produces an organization who sees threats coming but lacks the ecosystem standing to do anything about them upstream.

The framework's contribution is in the distinctions it draws. Resilience and risk foresight are not the same thing, and conflating them produces metrics that look similar but measure entirely different organizational capabilities. The near-miss standard exists precisely because lagging indicators (actual incidents, disclosed vulnerabilities, failed audits) arrive too late to change the behavior that caused them. The tier-three influence test exists because self-assessment at the ecosystem level is not credible.

What becomes possible with this infrastructure is a feedback loop between measurement and decision-making that most organizations currently lack. When the scorecard is live, the OSPO can trace a procurement decision to an upstream relationship, a staffing argument to a resilience gap, a compliance investment to a quantified risk reduction. That is a different conversation than the one most OSPO leaders are having today.

The proposed OSPO value framework is intended as a baseline reference. Organizations vary significantly in their strategic priorities, operating models, and maturity in open source engagement. As such, this framework should be used as a foundation to inform the design of a tailored OSPO value dashboard.

Practitioners are encouraged to adopt, adapt, and extend these metrics and dimensions based on what is most relevant to their context. In practice, effective OSPO value measurement emerges through iteration: refining indicators over time to better reflect organizational goals, stakeholder expectations, and evolving open source strategies. Organizations that apply this framework can build a value measurement approach faster than building from scratch, and the result will be more meaningful and actionable.

Organizations that build this framework now will be better positioned to measure OSPO value and make faster decisions with fewer organizational barriers around open source.

APPENDIX A. OSPO METRIC CATALOG

The following catalog provides sample metric definitions suitable for a quarterly executive scorecard. These are starting points rather than universal prescriptions and should be adapted to the organization's governance model, product portfolio, and data availability.

A.1 Metric design notes

Each metric below includes:

- **Definition:** what is being measured
- **Primary audience:** who uses it most directly
- **Cadence:** how often it should be reviewed
- **Caveats:** where interpretation can go wrong
- **Suggested owner:** who should maintain and track the measure

In addition, each metric should specify its baseline, denominator where relevant, and preferred presentation format (for example, count, percentage, ratio, or trend) so that interpretation remains consistent over time. Where relevant, metrics should also state scope explicitly (for example, limited to the most critical products or release trains) so that improved scores cannot be achieved by measuring only the easiest parts of the portfolio.

The following measures may be operationally useful but should not serve as primary evidence of enterprise value on their own:

- Number of tickets handled
- Number of trainings delivered
- Number of policies published
- Number of contribution requests processed
- Number of repositories scanned

These measures can help explain workload or process maturity, but none is sufficient as a headline measure of enterprise value.

A.2 ROI and cost avoidance metrics

TABLE 4
METRICS FOR ROI AND COST AVOIDANCE

Metric	Definition	Audience	Cadence	Caveats	Owner
Estimated commercial replacement value	Estimated annualized cost of commercial tools, platforms, or components that would need to be purchased if selected open source assets were unavailable.	CFO, CTO	Quarterly or annual	Avoid hypothetical replacements with no realistic procurement path. Document assumptions and confidence level.	OSPO + Finance business partner
External legal review spend avoided	Estimated spend avoided through standardized internal handling of recurring open source review classes that would otherwise require outside counsel or ad hoc expert review.	CFO, Legal leadership	Quarterly	Distinguish between avoided external spend and deferred work.	OSPO + Legal operations
Release delays prevented	Number of releases in which early OSPO intervention resolved an open source issue before it created a formal release gate or launch delay.	COO, CTO, Product leadership	Quarterly	Requires a documented intervention note. Do not count speculative possible delays.	OSPO + Release management
Compliance cycle-time reduction	Median time to complete common open source intake or review workflows compared with an agreed baseline.	Engineering leadership, COO	Monthly or quarterly	Reduced cycle time should not come at the expense of review quality.	OSPO operations lead

A.3 Resilience metrics

TABLE 5
RESILIENCY METRICS

Metric	Definition	Audience	Cadence	Caveats	Owner
SBOM coverage rate	Percentage of production-bound products or releases for which current, machine-readable SBOMs are available.	CTO, CISO	Monthly or quarterly	Report coverage scope explicitly: critical production systems only, or the full portfolio, since the two can produce very different numbers from the same governance program.	OSPO + Build or release the tooling team
Critical dependency visibility	Percentage of critical products whose key dependencies are identified, classified, and assigned owners or escalation paths.	CTO, CISO	Quarterly	Requires a documented definition of critical dependency.	OSPO + Architecture or supply chain governance team
Time to awareness for upstream issues	Median elapsed time between external publication or internal discovery of a relevant upstream issue and organizational awareness by responsible stakeholders.	CISO, Engineering leadership	Quarterly	Notification data may be incomplete if alerting paths are inconsistent.	OSPO + Security operations
Standard artifact coverage	Percentage of releases using standardized compliance artifacts, notices, or approved generation paths.	Legal, Engineering operations	Quarterly	Should be paired with quality checks rather than treated as a pure automation metric.	OSPO + Release engineering

Metric	Definition	Audience	Cadence	Caveats	Owner
AI tool provenance coverage (provenance assessed)	Percentage of AI coding tools approved for organizational use that have documented provenance assessments covering training data, license implications, and output review requirements. Distinct from the AI-generated code policy adoption metric in Appendix A.4, which measures the policy environment around the code AI tools produce.	CTO, Legal, CISO	Quarterly	Assessments should carry a review date and be revisited as tool versions change. Track assessment outcomes separately from coverage: report how many assessments resulted in conditional approval, restrictions, or rejection, so that coverage does not obscure governance decisions.	OSPO + Legal and procurement

A.4 Risk foresight metrics

TABLE 6
RISK FORESIGHT METRICS

Metric	Definition	Audience	Cadence	Caveats	Owner
High-risk issues identified pre-release	Count of high-risk licensing, provenance, policy, or security-related issues identified before formal product release.	Legal, Security, Risk committees	Quarterly	Volume alone is not inherently good or bad; interpretation depends on discovery quality and product mix.	OSPO + Security or Legal intake leads
Late-stage discovery trend	Trend in issues first discovered after formal release gate entry or other late development stages.	CTO, COO, Risk leadership	Quarterly	Late stage must be defined consistently.	OSPO program lead
Policy advisory issuance rate	This metric measures the OSPO's responsiveness to external change, not its workload. Number of formal advisories, guidance notes, or policy clarifications issued in response to changing regulatory, ecosystem, or governance conditions.	Legal, Security, Board risk committees	Quarterly or annual	A low rate in a stable regulatory and ecosystem environment may be appropriate. A low rate during a period of significant regulatory change (such as the CRA implementation period through 2027) or major ecosystem disruption is a warning sign of insufficient external monitoring. Interpret this metric in relation to the external event context, not as an absolute score. Each advisory issued should reference the specific trigger that prompted it.	Head of OSPO

Metric	Definition	Audience	Cadence	Caveats	Owner
Mean time from issue detection to decision	Median or mean elapsed time between issue identification and a formal decision on next action (approve, mitigate, replace, escalate, fork, contribute, or reject).	Engineering leadership, Legal, Security	Monthly or quarterly	Complex issues may appropriately take longer; trend matters more than isolated values.	OSPO Operations Lead
AI-generated code policy adoption (engineering teams)	Percentage of engineering teams with a documented, communicated policy governing the use of AI coding tools and the review of AI-generated code before it enters production codebases. Distinct from the AI tool approval coverage metric in Appendix A.3, which measures the assessment status of the tools themselves.	Legal, CTO, CISO	Quarterly	Policy existence does not equal enforcement; pair with a sampled review of recent AI-assisted contributions to verify compliance. Engineering teams must be defined and agreed upon before this metric is first published. The definition should remain stable across reporting periods.	OSPO + Legal + Engineering leadership

A.5 Strategic influence metrics

TABLE 7
METRICS FOR STRATEGIC INFLUENCE

Metric	Definition	Audience	Cadence	Caveats	Owner
Strategic contribution alignment	Percentage of material outbound contributions or engagement efforts explicitly linked to product, platform, or roadmap priorities.	CTO, Product strategy, CEO	Quarterly or annual	Requires a simple rubric for what counts as aligned.	OSPO + Product or Architecture leadership
Ecosystem representation coverage	Number or percentage of priority ecosystems in which the organization has formal representation, maintained engagement, or designated sponsorship.	CTO, Strategy leadership	Annual	Presence is not the same as influence; pair with narrative on significance.	Head of OSPO or Standards lead
Dependency influence ratio	Share of critical open source projects for which the organization meets at least one of the following qualifying conditions: • At least one employee holds a maintainer, committer, or TSC role and has made substantive technical contributions in the past six months; • The organization has sponsored a working group or foundation program directly tied to the project's governance and has a named internal owner accountable for that relationship; • The organization has had a technical contribution accepted that reduced a measurable downstream maintenance or security burden, with documentation of that outcome (or third-party corroboration from a maintainer or foundation if internal documentation is not yet established, to avoid penalizing organizations that contribute substantively but document inconsistently). Passive membership, conference sponsorship, and attendance without contribution do not qualify.	CTO, Strategy, Risk leadership	Annual	The denominator (business-critical open source projects) must be defined and agreed upon with architecture leadership before the metric is first published. The definition should be reviewed annually.	OSPO + Architecture or Ecosystem Lead

RESOURCES

- Linux Foundation Research and TODO Group. **The 2025 State of OSPOs and Open Source Management**. 2025.
- Linux Foundation Research and TODO Group. **The 2024 State of OSPOs and Open Source Management**. 2024.
- Linux Foundation Research and TODO Group. **The 2023 State of OSPOs and OSS Initiatives.** 2023.
- Hoffmann, M., Nagle, F., and Zhou, Y. **The Value of Open Source Software**. Harvard Business School Working Paper 24-038, January 2024.
- Nagle, F., Powell, K., Zitomer, R., and Wheeler, D. A. **Census III of Free and Open Source Software: Application Libraries**. Linux Foundation, OpenSSF, and Laboratory for Innovation Science at Harvard, December 2024.
- Linux Foundation and TODO Group. **The Lifecycle of an Open Source Program Office: From Inception to Strategic Pivoting**. 2005
- Linux Foundation. **Measuring the Economic Value of Open Source**. 2023.
- Linux Foundation. **The Evolution of the Open Source Program Office (OSPO)**. 2023.
- Linux Foundation Research and OpenSSF. **Addressing Cybersecurity Challenges in Open Source Software**. June 2022.
- Linux Foundation and CNCF. **2024 Cloud Native Security Report**. September 2024.
- GitHub, Linux Foundation, and Harvard Laboratory for Innovation Science. 2024 **Open Source Software Funding Report**. December 2024.
- GitHub. **Octoverse 2024**. November 2024.
- Sonatype. **10th Annual State of the Software Supply Chain Report**. October 2024.
- Synopsys. **2024 Open Source Security and Risk Analysis Report**. February 2024.
- Black Duck Software. **2025 Open Source Security and Risk Analysis Report**. February 2025.
- McKinsey & Company. **Open Source in the Age of AI**. 2024.
- CHAOSS Project. **Community Health Analytics in Open Source Software**.
- CHAOSS Project. **CHAOSS Practitioner Guide: Assessing Viability**.
- Basili, V. R., Caldiera, G., and Rombach, H. D. **The Goal Question Metric Approach. In Encyclopedia of Software Engineering**, John Wiley and Sons, 1994.
- ISO/IEC 5230:2020. Information technology, **OpenChain Specification. International Organization for Standardization**, 2020. (See also the **OpenChain adoption checklist**.)
- European Union. **Regulation (EU) 2024/2847 of the European Parliament and of the Council of 23 October 2024 on horizontal cybersecurity requirements for products with digital elements** (Cyber Resilience Act). Official Journal of the European Union, 20 November 2024. Entered into force 10 December 2024.

FEEDBACK

Despite careful reviews, typographical errors or inaccuracies may remain. If you identify an error or have a suggested correction or improvement to the framework, please contact the author directly.

ACKNOWLEDGMENTS

The author thanks **Hilary Carter** and Anna Hermansen of LF Research, **Ana Jiménez Santamaría** (Open Source Project Manager, Linux Foundation), and **Dawn Foster** (Open Source Strategy Consultant) for their review and valuable input on this report. The author also recognizes the **TODO Group** for its foundational work and continued leadership in advancing the OSPO field.

DISCLAIMER

The views expressed in this report are those of the author alone. They do not reflect the views of any organization or employer with which the author is currently or has previously been affiliated.

ABOUT THE AUTHOR

Ibrahim Haddad, Ph.D., is Head of Infotainment Engineering at Volvo Cars, where he leads engineering for the company's next-generation in-vehicle infotainment software.

Prior to Volvo Cars, he served as Vice President of AI Strategic Programs at the Linux Foundation, leading LF AI & Data, and concurrently as Founding Executive Director of the PyTorch Foundation. He scaled LF AI & Data from 9 to 77 member organizations across 70 projects, building the neutral institution where developers and organizations could code, govern, and scale open source AI work. He co-created the Model Openness Framework, the reference standard for AI model openness, and launched the Generative AI Commons dedicated to fostering the democratization, advancement and adoption of efficient, secure, reliable, and ethical Generative AI open source innovations.

Earlier at Samsung Research, he served as VP of R&D, and founded the Open Source Group and scaled it to 100+ engineers. He also co-founded the Open Connectivity Foundation and served as its elected Vice President.

His broader career spans Ericsson Research, Motorola, Palm, Hewlett-Packard, and the Linux Foundation, with experience scaling engineering organizations, leading small focused teams, and serving as an individual contributor inside highly matrixed environments.

He has championed open source development driven by the belief that collaborative development is a faster, better, and cheaper path to innovation. He is a recognized expert in the field of open source license compliance, with extensive experience designing, founding and operating enterprise compliance programs. He is also a long-standing advocate for OSPOs as the structural backbone organizations need for serious external R&D engagement. He is a prolific author of 7 books, 18 e-books, and 150+ technical reports on open source strategy, AI governance, and engineering leadership.

Haddad earned a Ph.D. with honors in Computer Science from Concordia University where he was awarded both the J. W. McConnell Memorial Graduate Fellowship and the Concordia University 25th Anniversary Fellowship for academic excellence.

LinkedIn: **https://www.linkedin.com/in/ibrahimhaddad/**
Website: **https://ibrahimatlinux.com/**

THE LINUX FOUNDATION | Research

Founded in 2021, **Linux Foundation Research** explores the growing scale of open source collaboration, providing insight into emerging technology trends, best practices, and the global impact of open source projects. Through leveraging project databases and networks, and a commitment to best practices in quantitative and qualitative methodologies, Linux Foundation Research is creating the go-to library for open source insights for the benefit of organizations the world over.

x.com/linuxfoundation

facebook.com/TheLinuxFoundation

linkedin.com/company/the-linux-foundation

youtube.com/user/TheLinuxFoundation

github.com/LF-Engineering

To reference this work, please cite as follows: Ibrahim Haddad, "Measuring OSPO Value: A Framework for ROI, Resilience, Risk Foresight, and Strategic Influence", foreword by Ana Jiménez Santamaría, Linux Foundation, June 2026.