RiskPrism® Maturity

Why Cybersecurity Governance Fails When the Map Is Wrong

A White Paper by CubicPrism Risk Analytics

1. Executive Summary

Cybersecurity maturity assessments aligned to established frameworks are widely used to inform decision making, audit oversight, and risk governance. In this role, they function as maps: representations intended to help boards and executives understand where the organization is, where risk resides, and how to navigate improvement. Many assessments rely on aggregated maturity scores, a representation choice that favors simplicity and comparability. While useful for reporting, such representations may limit visibility into how maturity varies across complex organizational environments, particularly where conditions are uneven or perspectives differ.

Traditional maturity assessments often convert uneven operational reality into a single reassuring score. RiskPrism Maturity preserves the distribution of maturity, the perspective of the assessor, the confidence behind the assessment, and the evidence supporting the conclusion. This produces a governance-ready view of cybersecurity maturity that shows where risk is concentrated, where perspectives diverge, and where leadership confidence may exceed operational reality.

RiskPrism® Maturity (*RPM*), developed by CubicPrism Risk Analytics (CubicPrism), applies a patent-pending, distributional, multi-assessor maturity assessment methodology aligned with the NIST Cybersecurity Framework. Rather than representing maturity as a single consolidated value, *RPM* preserves how maturity is distributed across assessed components and retains assessor perspective and confidence as part of the assessment record. This methodology provides a map of variation, convergence, and divergence within the organization, supporting governance, audit, and oversight contexts in which understanding the terrain is as important as summarizing position. When used for governance, averaged maturity scores create systematic blind spots that can materially misstate cybersecurity risk.

RPM produces cybersecurity risk information in a form compatible with enterprise risk registers, risk detail records, scenario analysis, risk appetite/tolerance evaluation, and enterprise risk profile development as described in NIST IR 8286A Rev. 1 (December 2025), while also providing distributional visibility that supports compliance traceability to NIST SP 800-171 Rev. 3 controls.

2. Introduction: When the Map Is Wrong

Imagine relying on a map labeled Chicago, only to discover that it is a map of Detroit. In that situation, no amount of effort, optimism, or discipline can correct the outcome – the underlying model is wrong.[1]

Cyber maturity assessments serve the same purpose as a map. Boards, audit committees, and executives rely on them to understand where the organization is, where risk resides, and how to

[1] Paraphrased from Covey, Stephen R. The 7 Habits of Highly Effective People: Powerful Lessons in Personal Change. Free Press, 1989.

move forward. When the assessment model accurately reflects reality, disciplined execution leads to improvement. When it does not, even well-intentioned effort becomes misdirected.

Many cyber maturity tools simplify complex organizational realities into rolled-up scores and averages. These models are easy to read and easy to report, but they frequently fail to represent the realities of cybersecurity across large, heterogeneous environments. The result is not a failure of leadership or execution, but reliance on a map that does not match the terrain – "the map is not the territory[2]."

Early quantitative cyber risk platforms, including foundational work leading to systems such as US Patent 10,686,825 (issued 2020), advanced decision support by processing threat logs to deliver dynamic, probabilistic risk quantification and role-tailored presentations. Yet they operated downstream of maturity posture and did not systematically capture or version the organization's cybersecurity capability distribution across assets, processes, and teams. This upstream gap limits the fidelity of inputs to modern enterprise risk management (ERM) processes that require granular, defensible cybersecurity risk data for risk registers (CSRRs), risk detail records (RDRs), scenario construction, and risk appetite/tolerance alignment (NIST IR 8286A Rev. 1, December 2025). Closing this gap calls for a more faithful representation – one that moves beyond averaged scores to a distributional model capable of preserving the full variation in maturity across the organization's diverse assets, systems, and operating environments, as well as the important variations that arise from multiple independent assessors.

3. What "Maturity" Actually Represents

Before maturity can be meaningfully measured, it must be honestly defined. In *RPM*, maturity is not a claim of security and not a prediction of outcomes; it is an indicator for how well ***cybersecurity risk*** is understood, identified, managed, and governed.

Higher maturity reflects intentionality: responsibilities are defined, practices are documented and consistently executed, performance is monitored, and improvement is deliberate rather than reactive. Lower maturity reflects reliance on informal knowledge, uneven execution, or the absence of shared understanding.

This distinction matters for proper governance. An organization may deploy advanced security technologies while operating at low maturity if accountability is unclear or decision making is ad hoc. Conversely, an organization with modest technical controls may demonstrate higher maturity if risk is consistently identified, decisions are documented, and execution aligns with intent.

[2] Korzybski, A. (1995). Science and sanity: An introduction to non-Aristotelian systems and general semantics (5th ed.). Institute of General Semantics. (Original work published 1933)

RPM therefore treats maturity as a proxy for the cybersecurity risk lifecycle. Higher maturity does not imply immunity to incidents, but rather that, when incidents occur, the organization is more likely to understand why, respond deliberately, and improve systematically.

4. The Limits of Average-Based Cybersecurity Maturity

Many cyber maturity tools implicitly assume that maturity can be meaningfully summarized as a single value. Individual findings are normalized and rolled up to produce a category-level, function-level, and (ultimately) a framework-level score. That score is then used as an indicator of cybersecurity effectiveness.

The limitation of this approach is not a mathematical defect, but a conceptual constraint commonly described as the *flaw of averages*[3]: the tendency for conclusions based on average conditions to obscure how complex systems vary in practice. When applied to cybersecurity maturity, average-based representations may fail to fully reflect how risk and capability are distributed across an organization's operational environment. By compressing variation into a single value, averages reduce visibility into assessment distribution, areas of agreement, and areas of divergence that are often material for governance and oversight.

Consider two organizations that report the same overall cybersecurity maturity score. In the first, core business systems operate under defined, consistently executed practices, while a small number of peripheral systems rely on informal or ad hoc controls. In the second, strong practices exist in a limited set of centrally managed environments, while a substantial portion of operational systems – often legacy or decentralized – operate with minimal governance or undocumented procedures. Although both organizations present the same average maturity, their risk exposure is materially different. In the latter case, risk is concentrated where visibility and accountability are weakest, creating a governance blind spot that is not apparent in aggregated reporting. This contrast is illustrated in the table below:

[3] Savage, S. L. (2002). The flaw of averages. Harvard Business Review, 80(11), 20–27. https://hbr.org/2002/11/the-flaw-of-averages

Governance Information Available	Average-Based Assessment	RPM Distributional Assessment
Reported maturity score	**Yes**	**Yes**
Maturity distribution across environment	**No**	**Yes**
Areas operating at low or absent maturity	**No**	**Yes**
Whether strong areas mask weak areas	**No**	**Yes**
Concentration of risk in legacy systems	**No**	**Yes**
Concentration of risk in decentralized environments	**No**	**Yes**
Assessor role, proximity, and confidence	**No**	**Yes**
Supporting justification and evidence	**No**	**Yes**
Disagreement between assessors	**No**	**Yes**
Defined assessment basis	**No**	**Yes**
Clear basis for remediation priority	**No**	**Yes**
Usable input for risk registers and scenario analysis	Partial	**Yes**

Marc Andreessen, a prominent Silicon Valley entrepreneur, investor, and early web pioneer, has described the "big gray cloud": layers of management interpretation that can insulate leadership from operational reality. Traditional average-based cybersecurity maturity reporting can produce a similar governance effect. By converting uneven practices, assessor disagreement, weak ownership, low-maturity pockets, and evidence gaps into a single score, the average becomes easy to report but less faithful to the terrain leadership must govern.

In financial oversight, boards and audit committees have long understood that averages can conceal material issues. A statement such as "the organization is profitable on average" may be true while still masking persistent losses in specific business units or concentrated exposure in critical operations. Cyber maturity behaves the same way. Strong practices in some areas can offset weak or informal practices elsewhere, producing a reassuring score that does not reflect actual risk.

When leadership relies on such scores, improvement efforts tend to focus on raising the maturity number rather than reducing exposure. Oversight becomes score-based rather than risk-based. Over time, confidence paradoxically rises as fidelity to reality declines.

5. Empirical Evidence: Perception Gaps and Systemic Limitations

Drawing on extensive hands-on experience conducting and participating in NIST Cybersecurity Framework maturity assessments across heterogeneous organizations, CubicPrism has directly observed the dynamics described above: managers advancing claims materially inconsistent with operational reality, technical experts being overridden by bureaucratic or political priorities, and cultural reluctance to surface findings that might 'rock the boat.'

Empirical research confirms these patterns are systemic rather than anecdotal. The Bitdefender 2025 Cybersecurity Assessment[4] found C-level executives more than twice as likely as operational and mid-level teams to report being 'very confident' in organizational readiness (45% versus 19%), a disconnect explicitly linked to executives' strategic lens versus practitioners' daily exposure to inconsistent controls, legacy environments, and informal practices. Complementary analyses of NIST CSF assessments, including BluOcean Cyber's 2025 review of maturity-model limitations[5], document that management-led or consensus-driven scoring routinely overstates maturity levels, with high-level averages masking asset-specific vulnerabilities and producing scores that 'do not pinpoint where biggest exposures lie.' Governance literature further identifies 'management override' as a recurring failure mode in which bureaucratic or operational priorities quietly sideline technical assessments, allowing exceptions to accumulate without visibility. When single opinions or forced consensus are then compressed into rolled-up averages, the result is not merely imprecise but actively misleading for oversight.

6. A Distributional View of Cybersecurity Maturity

RPM replaces rolled-up scoring with a distributional model aligned to the NIST Cybersecurity Framework. In practice, for each function, category, and subcategory, assessors explicitly allocate 100% of in-scope factors – real assets, systems, business processes, or control instances – across RPM's internal 0–5 maturity scale. The scale is aligned to NIST CSF outcomes and informed by the risk-governance concepts reflected in the NIST CSF Tiers, but it is not presented as an official NIST CSF Tier rating. This structure enables precise distributional allocation from 0 = Absent through 5 = Optimized while remaining aligned with the framework's intent. Every allocation is tagged with the assessor's role, operational proximity, confidence level, and supporting justification or evidence. The result is a preserved distribution, visible divergence across multiple independent assessors, and auditable inputs ready for enterprise risk registers and scenario analysis.

This approach models operational reality; organizations rarely function at a single maturity level across all domains, particularly across legacy and peripheral environments. *RPM* reveals this distribution rather than "averaging it away."

For each assessment, the assessor must identify the assessment basis before allocating maturity. The assessment basis may be a defined population of systems, applications, business processes, data flows, third-party services, control instances, or another explicitly stated

[4] Bitdefender. (2025). 2025 Cybersecurity Assessment Report. https://www.bitdefender.com/en-us/business/campaign/2025-cybersecurity-assessment

[5] BluOcean Cyber. (2025). Beyond NIST CSF: Limitations of Maturity Assessments and How to Fix Them. https://bluoceancyber.com/beyond-nist-csf-limitations-of-maturity-assessments-and-how-to-fix-them/

population. The 100% allocation applies to that defined assessment basis, not to an undefined or implied enterprise total.

This visibility is essential for governance. Risk is seldom evenly distributed; it accumulates where ownership is unclear, oversight is weak, or inadequate practices persist. A model that conceals this distribution does more than omit detail – it misleads the organization by presenting a simplified picture that does not reflect how risk is borne across the organization.

For a single NIST Cybersecurity Framework subcategory, *RPM* requires assessors to account for all in-scope systems, processes, and practices by explicitly describing how maturity is distributed across the organization's operating environment. In practice, this often reveals that a substantial portion of activity is governed by informal, case-by-case practices, while a material remainder operates without defined procedures, frequently concentrated in legacy or peripheral environments. Rather than collapsing these conditions into a representative maturity score, *RPM* preserves the full distribution and associates it with assessor perspective, proximity to operations, confidence, and supporting justification. This approach reveals where maturity is uneven and where cybersecurity risk is concentrated, enabling governance conclusions that focus on exposure and accountability rather than on an averaged indicator of progress.

While any maturity model that requires full distributional visibility and multiple independent perspectives necessarily demands more thoughtful input than a single rolled-up score, *RPM* is engineered for rapid, low-friction adoption—even in mid-market and resource-constrained environments. Its clean, assessor-friendly interface presents the NIST CSF hierarchy, 100% allocation controls, proximity and confidence qualifiers, and supporting justification fields in a single intuitive workspace, enabling complete assessments without unnecessary wizard-driven workflows, excessive administrative overhead, or extensive training.

7. From Single Opinion to Collective Insight

Many maturity assessments assume that one assessor's judgment (or several assessors together in a meeting but forced to reach consensus) can represent organizational truth. *RPM* rejects this assumption.

The platform is designed to support multiple independent assessments of the same NIST CSF components. These assessments may – and should – come from different roles, teams, or perspectives and may occur repeatedly over time. Rather than forcing agreement or collapsing differences into a single value, *RPM* preserves convergence and divergence as meaningful signals; when perspectives align, confidence increases. But when they diverge, the disagreement itself provides insight, often pointing to communication gaps or misaligned assumptions.

As importantly, *RPM's* multiple assessor support exposes patterns that no single assessment could reliably identify. By incorporating multiple independent perspectives, *RPM* applies principles of collective intelligence (also known as "wisdom of crowds") in a way that reduces overreliance on individual judgment and limits the risk of institutional blind spots.

8. The Governance Value of Assessor Role

In audit and financial oversight, the source of an assessment bears directly on its interpretation. A control evaluation provided by management is interpreted differently from one provided by internal audit or an external reviewer. Each perspective has value, but each carries inherently different implications.

RPM applies this same discipline to cybersecurity assessment. Every assessment tracks the assessor's role, their proximity to operations, the nature of the assessment performed, and the assessor's confidence. This enables reviewers to distinguish between governance intent and practical execution. It also highlights where leadership perception diverges from operational reality. Risk, in this context, is not only about control strength, but about confidence in who is reporting on that strength.

9. Integrating *RPM* Outputs into Enterprise Risk Management

RPM is designed not only as a high-fidelity cybersecurity maturity assessment tool but also as a practical enabler of enterprise risk management (ERM) integration, consistent with NIST IR 8286A Rev. 1 (December 2025), "Identifying and Estimating Cybersecurity Risk for Enterprise Risk Management."

By preserving full distributional allocations (100% coverage across maturity levels 0–5 per NIST CSF subcategory), independent multi-assessor perspectives with metadata (proximity, confidence, type), justifications, evidence, and temporal versioning, *RPM* generates a structured precursor baseline cyber risk profile. This baseline reflects operational reality rather than an averaged abstraction and serves as credible, auditable input to cybersecurity risk registers (CSRRs), risk detail records (RDRs), scenario-based analysis, risk appetite/tolerance evaluation, and enterprise risk profile aggregation.

RPM is not a snapshot tool. It leverages system-versioned temporal tables to preserve every past assessment in an immutable version history that permits authorized superseding updates and corrections. This approach supports ongoing monitoring of risk evolution while still allowing yearly snapshots to be generated from the historical record for board reporting, insurance, and regulatory purposes – delivering a living, evolving view of risk posture rather than a static annual picture.

RPM assessments also support traceability to NIST SP 800-171 Rev. 3 security requirements by providing distributional visibility across applicable requirement families and mapped control areas. For organizations preparing for CMMC, RPM can also support mapping to CMMC assessment domains, evidence expectations, and self-assessment or third-party assessment preparation.

10. Transparency, Evidence, and Defensibility

RPM is designed with scrutiny in mind. Assessments can be supported by narrative justification, technical notes, references to supporting artifacts, and the attachment of artifacts themselves. When required substantiation is missing, assessments are explicitly marked as incomplete rather than silently accepted, preventing false assurance and post-hoc rationalization of control effectiveness. RPM does not treat insufficient evidence as proof of low maturity. "Unknown," "unsupported," and "absent" are distinct governance conditions. Where the assessor cannot support a maturity allocation with sufficient context, justification, or evidence, the assessment remains incomplete or unsupported rather than being converted into a speculative low-maturity score.

This transparency supports defensible governance. Results can be explained, challenged, and relied upon with a clear understanding of their basis and limitations, which is particularly important when assessments are used in audit, insurance, regulatory, or legal contexts.

11. WHO *RPM* IS – *AND* IS *NOT* – FOR

RPM is built for organizations that value clarity over comfort. It is for boards, audit committees, and senior leaders who understand that oversight requires more than a reassuring score. It is for security and risk leaders who are willing to describe reality honestly, even when that reality is uncomfortable or inconvenient.

It is also for internal audit functions, risk consultants, cybersecurity advisors, and cyber insurance reviewers who must defend their conclusions. The platform preserves perspective, confidence, evidence, and assessment basis so that findings can withstand scrutiny in governance, audit, advisory, underwriting, and risk-transfer contexts.

This distinction is deliberate. RPM is designed for organizations that want the assessment map to reflect the actual terrain, including the uncomfortable parts. It is built for leaders who value defensible visibility, assessor independence, and evidence-backed governance over simplified reporting.

RPM is *NOT* for organizations seeking a quick answer or a single number to report and move on. It is *NOT* for organizations in which disagreement must be smoothed away or in which assessments exist primarily to satisfy compliance requirements. It is *NOT* for organizations that want a promise of certainty, or the elimination of risk.

12. CONCLUSION: CHOOSING THE RIGHT MAP

> *"If the map is wrong, the more you rely on it, the more lost you become."*
> *--Stephen R. Covey*

Cyber maturity is uneven, contextual, and dependent on perspective. Treating it as a single averaged score creates a map that is easy to read but fundamentally wrong. When governance relies on average-score representations, the result is not merely imprecision but systematic

decision error, as attention and resources are guided by a map that does not accurately reflect where risk resides.

RPM offers a different map – one that preserves distribution, exposes convergence and divergence, highlights inconsistencies, and reflects how organizations *really* operate. For leaders responsible for governance, audit, and accountability, this difference matters.

When the map is wrong, progress is illusory. When the map preserves distributional reality, temporal fidelity, and compliance traceability to frameworks such as NIST SP 800-171 Rev. 3, cybersecurity risk becomes meaningfully actionable. RPM's governance value is not that it produces a more complicated score. The value is that it preserves the facts usually lost before leadership ever sees the assessment: distribution, evidence, confidence, proximity, disagreement, and change over time.

Organizations whose current maturity reporting remains average-based are encouraged to pilot RPM on 6–10 high-value subcategories across three assessor roles, compare the resulting distribution and divergence against their existing rolled-up score, and review the differences with their governance committee.

13. About the Author

Frederick Doyle is a cybersecurity, risk, and governance practitioner with experience assessing cybersecurity maturity, designing defensible control frameworks, and translating technical findings into governance-ready risk information. His work focuses on improving the fidelity of cybersecurity assessment methods so that boards, executives, audit functions, and risk leaders can make decisions based on operational reality rather than simplified scoring abstractions.

14. About CubicPrism Risk Analytics

CubicPrism Risk Analytics is a cybersecurity risk consultancy and the creator of RiskPrism Maturity, a high-fidelity assessment platform for cybersecurity, risk, and governance, applied in advisory, audit, and oversight contexts where defensibility and accountability matter. Our approach is grounded in Richard Feynman's principle of fidelity to reality over simplified abstractions, articulated during the Challenger investigation: "reality must take precedence over representation, because nature cannot be fooled." Accordingly, *RPM* is designed to resist comforting abstractions and to reflect operational conditions as they exist.

15. References

NIST IR 8286A Rev. 1, "Identifying and Estimating Cybersecurity Risk for Enterprise Risk Management," National Institute of Standards and Technology, December 2025.

NIST SP 800-171 Rev. 3, "Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations," National Institute of Standards and Technology, May 2024.