W H I T E P A P E R

Risk Vector Analysis

A Mathematical Framework for Organizational Risk Posture Measurement, Trending, and Decision Support

Prepared by: **Cameron Kinsel**

Organization: **Ironclad Security LLC**

Date: **March 2026**

Version: **1.0**

Executive Summary

Organizations today face an overwhelming volume of security findings from vulnerability scanners, penetration tests, configuration audits, and threat intelligence feeds. The standard approach to synthesizing this data—likelihood-times-impact matrices, color-coded heatmaps, and subjective risk registers—fails to provide leadership with actionable, mathematically defensible answers to the question: are we getting better or worse?

The Challenge: Hundreds of Findings, Dozens of Sources, No Common Language

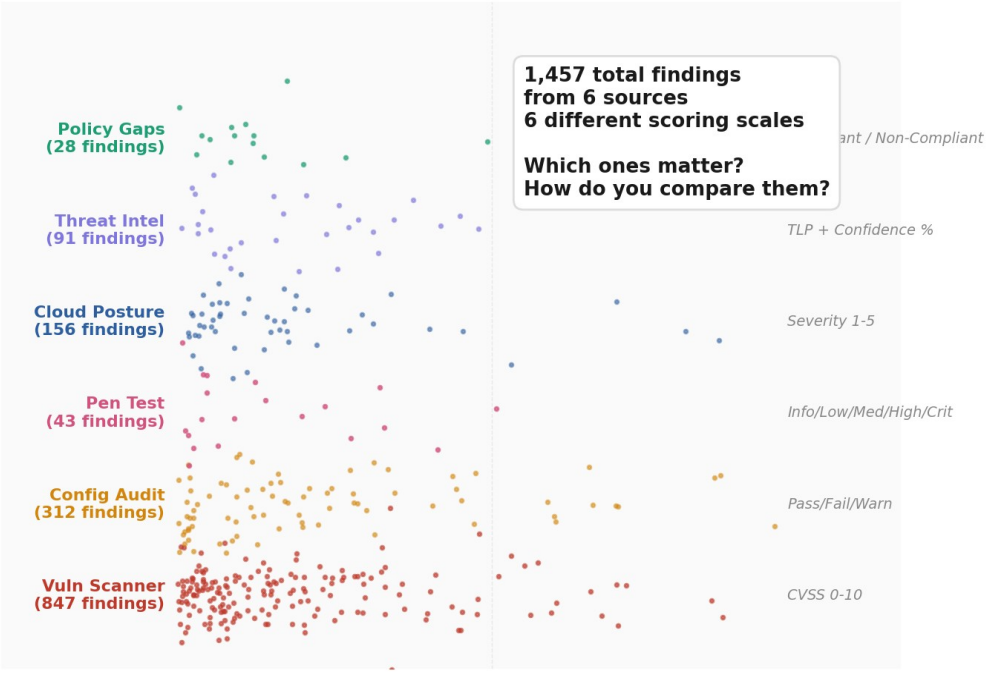

A typical organization faces hundreds of findings from diverse sources. Which ones matter most?

This white paper introduces Risk Vector Analysis (RVA), a quantitative framework that plots every organizational risk finding on a continuous two-dimensional coordinate system, calculates a single organizational risk score as a percentage, and—critically—measures the direction and magnitude of risk posture change over time using vector mathematics.

The result is a framework where a board member can understand organizational risk posture in thirty seconds, a CISO can make resource allocation decisions based on precise sensitivity analysis, and an analyst can prioritize remediation based on measurable score impact rather than subjective urgency labels.

The Problem with Current Risk Scoring

Discrete Matrices Destroy Information

The most common risk assessment tool in cybersecurity is the likelihood-by-impact matrix. A 3×3 grid produces nine possible outcomes. A 5×5 grid produces twenty-five. In either case, findings with materially different risk profiles collapse into identical cells. A vulnerability affecting 12% of assets and one affecting 48% of assets may both land in the same “Medium” likelihood column. The information that distinguishes them—the actual percentage—is discarded at the moment of classification.

This is not a minor inconvenience. It is a fundamental design flaw. When two findings with different remediation costs, different blast radii, and different exploit probabilities receive identical scores, the resulting prioritization is arbitrary. Resources are allocated based on classification artifacts rather than actual risk differentials.

Subjective Scoring Introduces Uncontrolled Variance

Traditional risk registers ask assessors to assign likelihood and impact scores on ordinal scales. Is this vulnerability a 3 or a 4? That single-point disagreement represents a 25% difference in risk score—a gap that could entirely determine whether a finding gets remediated this quarter or next. Reasonable practitioners will disagree, and that disagreement introduces variance that compounds across hundreds of findings. The final “risk score” reflects assessor psychology as much as organizational reality.

RVA eliminates subjective scoring entirely. Both axes use externally validated, normalized data sources: CVSS v3 base scores for severity (an industry standard maintained by NIST and FIRST) and empirical asset inventory data for spread. Neither axis requires human judgment to populate.

No Directional Trending

Perhaps the most significant limitation of existing approaches: they provide point-in-time snapshots with no mathematical framework for measuring change. If your risk score was 72 last quarter and 68 this quarter, all you know is that it decreased by four points. You do not know whether that improvement came from reducing severity concentration, narrowing asset exposure, or some combination. Different causes require different responses, and a scalar delta cannot distinguish between them.

Worse: what happens when the score stays flat? Your organization scores 62 this year and 62 last year, but your security budget increased 15%. Where did the money go? A flat scalar score contains zero information about whether the investment shifted risk from critical severity to distributed low-severity issues, whether new findings offset remediated ones, or whether the program is genuinely stagnant. Leadership cannot make funding decisions on a number that doesn’t move.

RVA solves this by treating year-over-year risk posture change as a vector with both magnitude (how much did we move?) and direction (which way did we move?). The direction, expressed as

an angle, maps to a four-quadrant interpretation framework that tells leadership not just that risk changed, but how it changed and what that implies for resource allocation. Even a flat score reveals its story: a CAP that moved from (30, 6.0) to (40, 4.2) traveled a measurable vector—the score may be similar, but the organization shifted from severity-concentrated to distributed risk, a directional change with specific remediation implications.

The Risk Vector Analysis Framework

Coordinate System

RVA plots every risk finding on a continuous XY coordinate system with the following axes:

- X-Axis: Percentage of organizational assets affected (0–100%). Derived from asset inventory correlation with finding scope. Continuous, not bucketed.
- Y-Axis: Normalized severity (0.0–10.0). The default mapping uses CVSS v3 base scores, but the axis is deliberately agnostic. Any severity input that can be encoded to a 0–10 scale integrates directly: Tenable's VPR, FIRST's EPSS probability scores (scaled), CVSS v4 environmental scores, or proprietary internal scoring. Organizations can further adjust individual finding scores based on contextual factors—asset criticality, production vs. non-production environment, microsegmentation status, ease of remediation, or business-specific exposure. This allows the same framework to reflect industry-standard severity alongside organization-specific risk context.

The origin (0, 0) represents zero risk—no assets affected at zero severity. The theoretical maximum (100, 10) represents total compromise at maximum severity. Every finding occupies a unique point in this continuous space, preserving the full granularity of both dimensions.

The Four Zones

The coordinate space divides naturally into four interpretive zones based on which quadrant a finding or cluster occupies relative to the center:

Zone	Region	Interpretation
Ideal Zone	Low X, Low Y	Low severity, narrow asset spread. Organizational risk is well-managed. Findings here represent acceptable residual risk.
Critical Corrections	Low X, High Y	High severity but limited spread. A small number of critical findings demand immediate attention. Quick wins are available: patching one system may eliminate the highest-severity risks.
Distributed Risk	High X, Low Y	Low severity but broad spread. Risk is distributed across many assets at low individual impact. This may reflect intentional prioritization of critical findings over widespread low-severity issues, or it may indicate systemic configuration drift requiring programmatic remediation.
Danger Zone	High X, High Y	High severity and wide spread. Organizational risk is critical. Leadership escalation is required. Findings in this zone represent existential risk to the organization and demand emergency resource allocation.

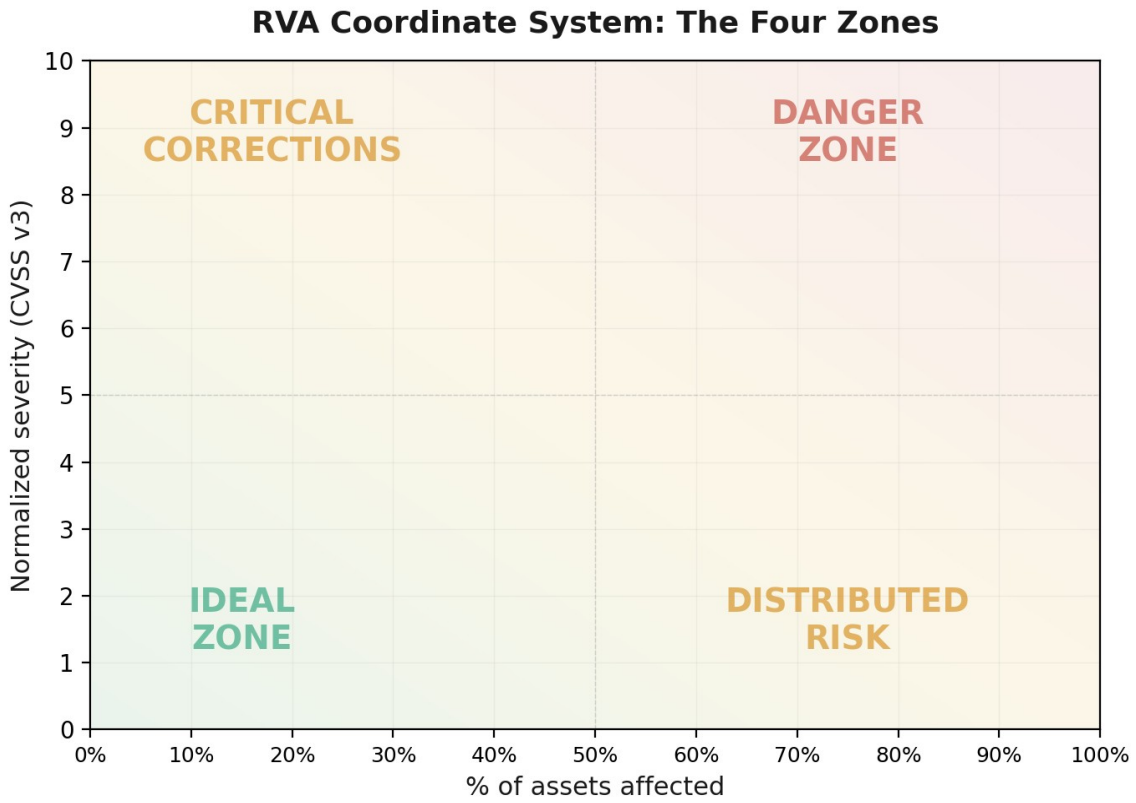

Figure 1: The RVA coordinate system with four interpretive zones

Distance from Origin: The Single-Integer Risk Score

Each finding's risk contribution is quantified as its Euclidean distance from the origin, normalized to account for the different scales of each axis:

$$d = \sqrt{((x/100)^2 + (y/10)^2)}$$

Dividing x by 100 and y by 10 normalizes both dimensions to a 0–1 scale, ensuring equal weighting. The maximum possible distance is √2 (the point 100, 10). To produce an intuitive organizational score, we calculate the Calculated Average Point (CAP)—the mean x and mean y of all findings—and express its distance as a percentage of maximum:

$$\text{Risk Score (\%)} = (d(\text{CAP}) / \sqrt{2}) \times 100$$

A score of 0% represents perfect security posture. A score of 100% represents total compromise at maximum severity across all assets. In practice, organizational scores typically range between 15% and 55%. Scores above 50% indicate critical posture requiring immediate executive attention.

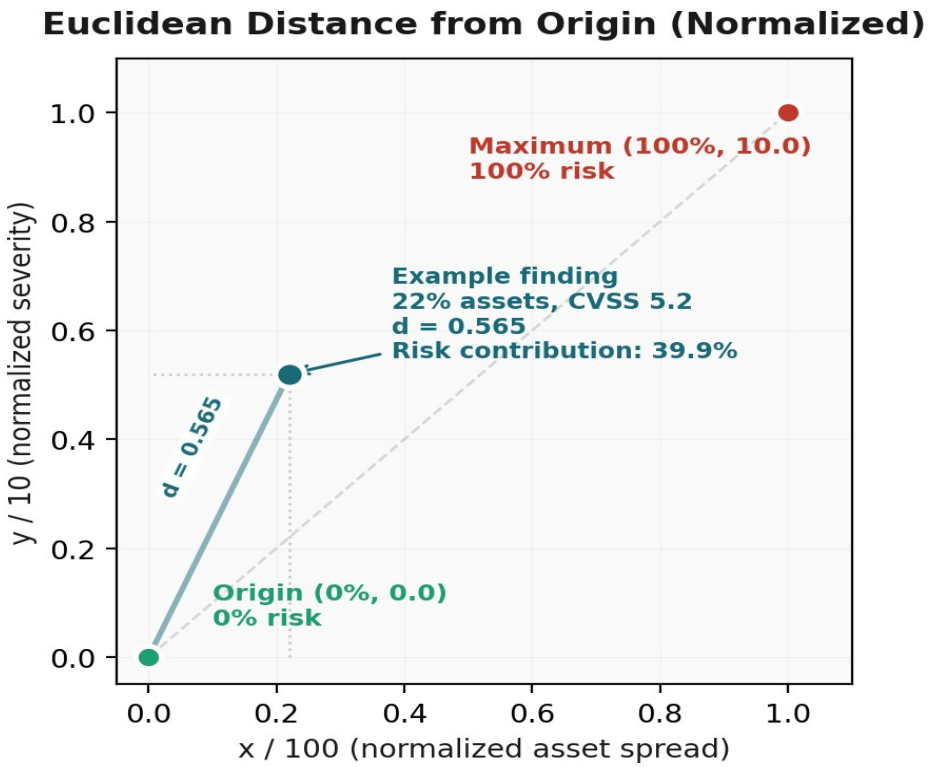

Figure 2: Euclidean distance calculation in normalized coordinate space

The Calculated Average Point (CAP)

Definition and Computation

The CAP is the arithmetic mean of all finding coordinates:

$$CAP = (\Sigma x_i/n , \Sigma y_i/n)$$

Where n is the total number of tracked findings, x_i is the asset percentage for finding i, and y_i is the normalized severity for finding i.

The CAP serves as the single point that represents the organization's aggregate risk posture. Its position on the chart tells leadership where the organization "lives" in risk space. Its distance from the origin is the organizational risk score. Its position relative to the Business Appetite Target (see below) indicates whether the organization is operating within defined tolerance.

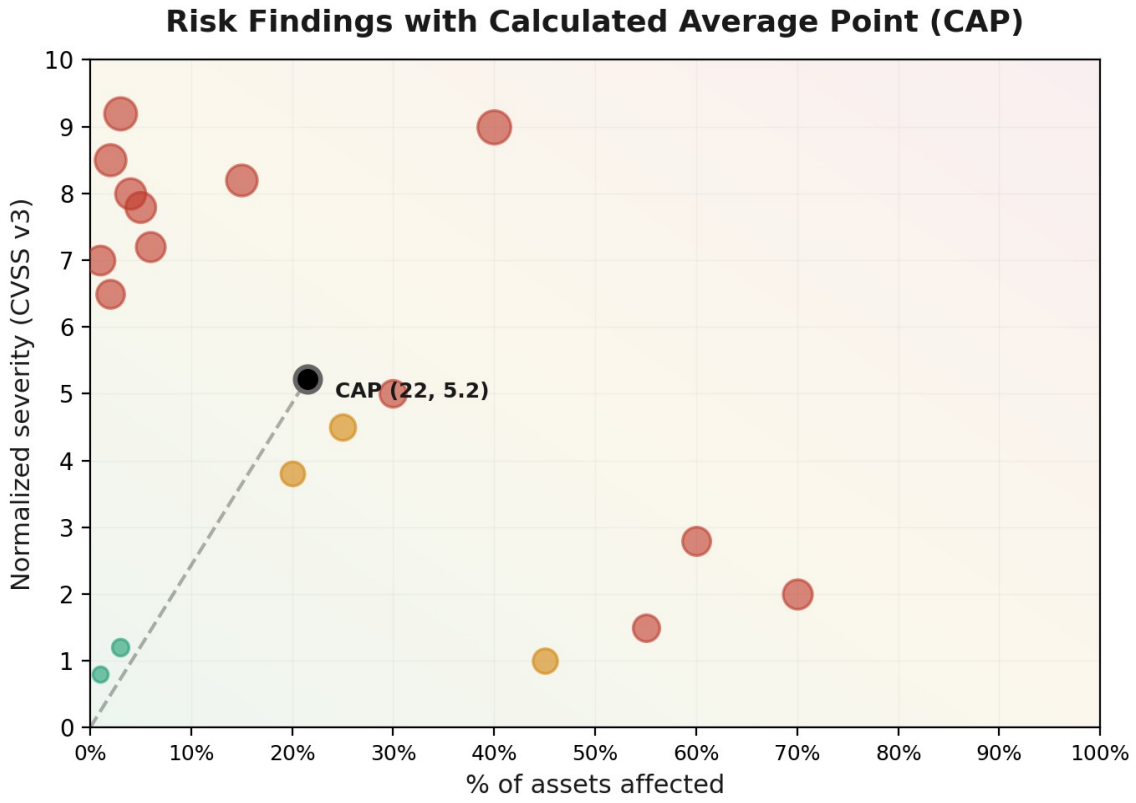

Figure 3: Risk findings plotted with CAP and distance vector

Year-Over-Year Tracking

The CAP is computed at defined intervals—annually, quarterly, or per assessment cycle. Each historical CAP is preserved as a data point, creating a trajectory through risk space that visualizes organizational risk evolution over time.

Historical CAPs are displayed with progressive visual fading: the most recent CAP is fully opaque, and older CAPs become increasingly transparent. This creates an intuitive “trail” that shows where the organization has been and where it is headed.

Rolling Averages for Spike Smoothing

Individual assessment cycles may include anomalous findings (a newly disclosed zero-day, a one-time misconfiguration) that spike the CAP temporarily. To prevent single-cycle anomalies from distorting the trend line, RVA supports optional rolling average smoothing. A three-period rolling average dampens spikes while preserving genuine directional trends. Organizations should select a smoothing window appropriate to their assessment cadence.

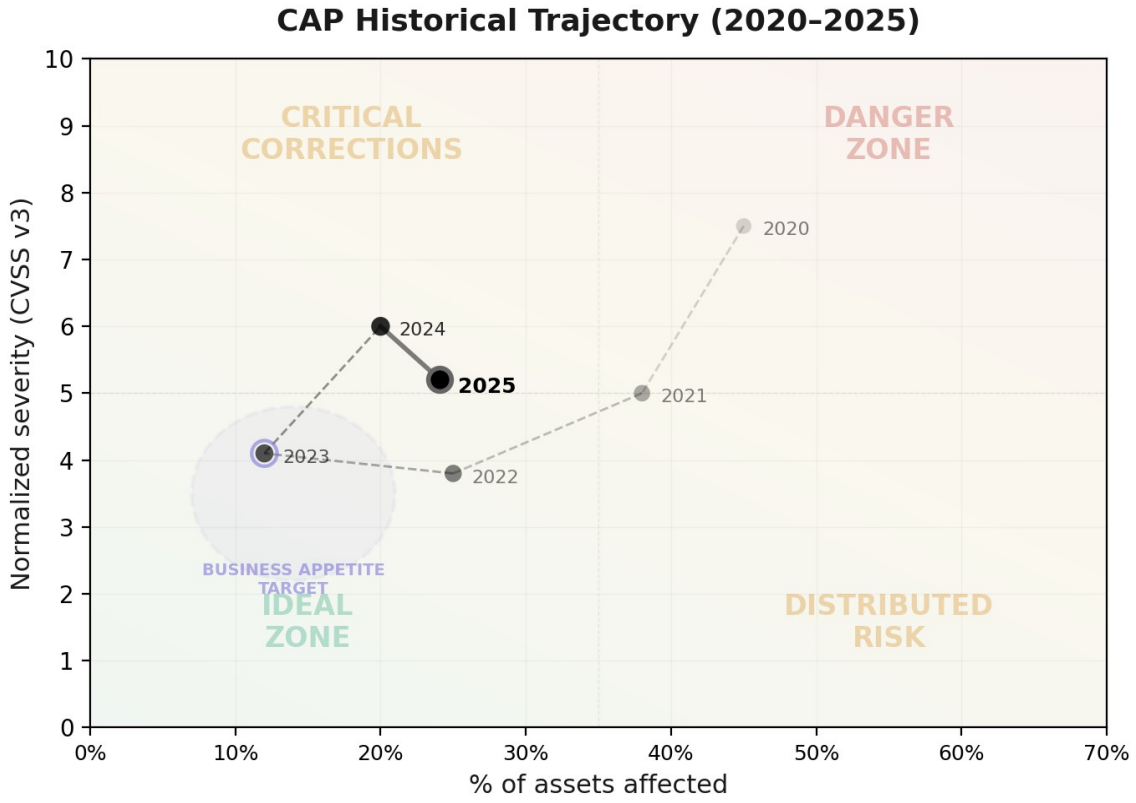

Figure 4: Six-year CAP trajectory showing organizational risk evolution

Vector Analysis: Direction of Change

This is the core differentiator of the RVA framework. While traditional risk scoring can tell you that a score changed, vector analysis tells you how it changed—and the how determines the appropriate organizational response.

Computing the Trend Angle

The vector between two consecutive CAPs defines the direction of risk posture change. The angle of this vector, measured in standard mathematical notation (0° = east/right, counterclockwise), maps to a four-quadrant interpretation framework:

$$\theta = \text{atan2}((CAP_2.y - CAP_1.y)/10 , (CAP_2.x - CAP_1.x)/100)$$

The division by 100 and 10 normalizes the axes before angle computation, ensuring that the angle reflects proportional change rather than raw coordinate difference.

The Four Trend Quadrants

Angle Range	Direction	Meaning	Recommended Response
0° – 90°	↗️ Risk Increasing	Both severity and asset spread increased. The organization is measurably worse off than the prior period. Overall risk grew in both dimensions.	Executive escalation. Emergency resource allocation. Board notification if sustained across multiple periods.
90° – 180°	↖️ Critical Corrections	Asset spread narrowed but severity increased. Fewer systems are affected, but the remaining findings are more critical. The overall distance from origin may be similar, indicating concentration risk.	Targeted remediation of high-severity findings. Investigate root cause of severity increase. Validate that spread reduction reflects genuine remediation, not scope reduction.
180° – 270°	↙️ Ideal Trend	Both severity and spread decreased. The organization is moving toward the Ideal Zone. The security program is producing measurable results across both dimensions.	Continue current program investment. Monitor for diminishing returns. If CAP enters appetite target, evaluate whether security spending can be reallocated to innovation.
270° – 360°	↘️ Distributed Risk	Severity decreased but more assets are accumulating lower-severity findings. Critical risks are being managed while distributed risk grows. This may be a deliberate strategic choice: prioritizing critical remediation while accepting broader low-impact exposure.	If intentional: validate that critical reduction is tracking to plan. If unintentional: evaluate whether increased spread reflects improved detection coverage or genuine degradation. Consider automated remediation for high-volume, low-severity classes when ready.

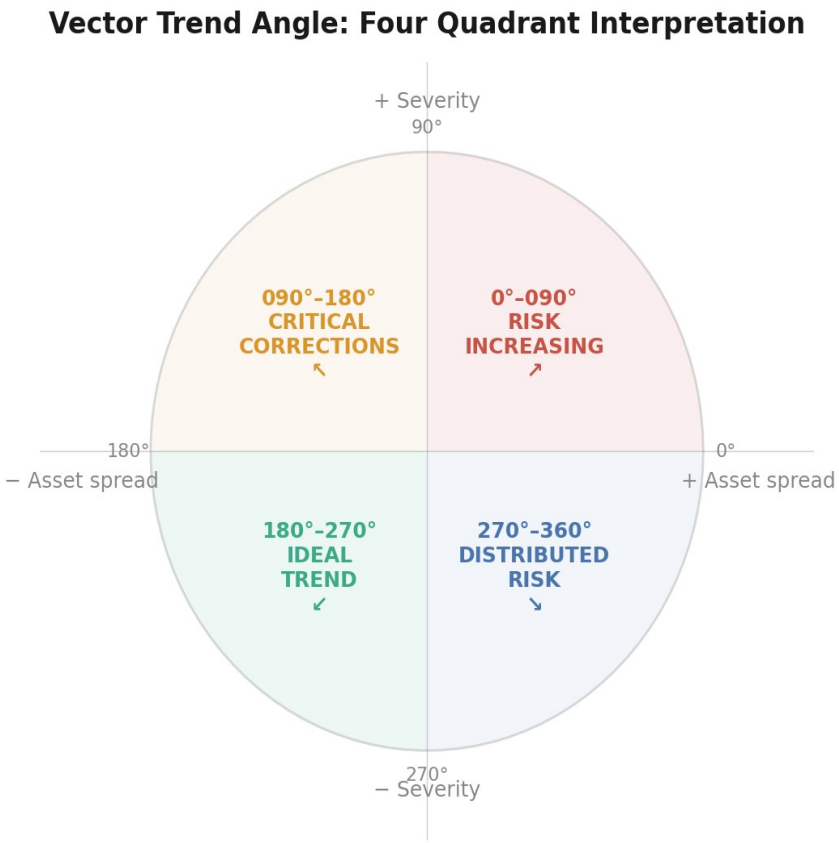

Figure 5: Vector trend angle quadrant interpretation

Magnitude: How Urgently Must We Act?

The angle tells you what to work on. The magnitude tells you how urgently to prioritize protection efforts over growth and innovation work. A 220° angle with a 5% year-over-year improvement may represent the maximum rate of change the business can reasonably absorb given current staffing and change windows. The same angle with a 25% improvement needed represents a structural gap that demands dedicated funding and potentially deferred projects.

This is fundamentally a rate-of-change-over-time problem. If an organization can sustain 8–10% CAP improvement per year at a given budget, the magnitude tells leadership how many periods of sustained effort separate the current posture from the target. A magnitude equivalent to 30% improvement at a 10%-per-year rate means three years of consistent execution—or a budget increase to accelerate the rate. Together, angle and magnitude form a complete description of organizational risk velocity in polar coordinates: direction defines the strategy, distance defines the investment.

The Business Appetite Target

Definition

The Business Appetite Target is a business-defined zone on the risk chart, represented as an ellipse centered on the organization's ideal CAP position. It answers the question: where should our CAP be if we are operating within acceptable risk tolerance?

The center coordinates and radii of the appetite ellipse are set by organizational leadership, not by the framework. This is critical: the framework provides the measurement; the business defines the threshold. Different organizations in different industries with different regulatory requirements will define different appetite zones.

Three Posture States

Posture	Implication
CAP Inside Appetite	The organization is operating within defined risk tolerance. No additional resource allocation required beyond maintenance. The security program is effective and appropriately funded. Leadership can confidently redirect incremental budget toward innovation or other priorities.
CAP Above Appetite	The organization exceeds its risk tolerance. Resources should be diverted toward remediation. The vector angle indicates which dimension is driving the excess: severity (invest in patching and vulnerability management), spread (invest in asset hardening and configuration management), or both (escalate to leadership for budget increase).
CAP Below Appetite	The organization is consistently more secure than its defined tolerance requires. While this may seem desirable, it often indicates over-investment in security controls that create operational friction without proportional risk reduction. Security spending may be reducible without meaningful posture impact.

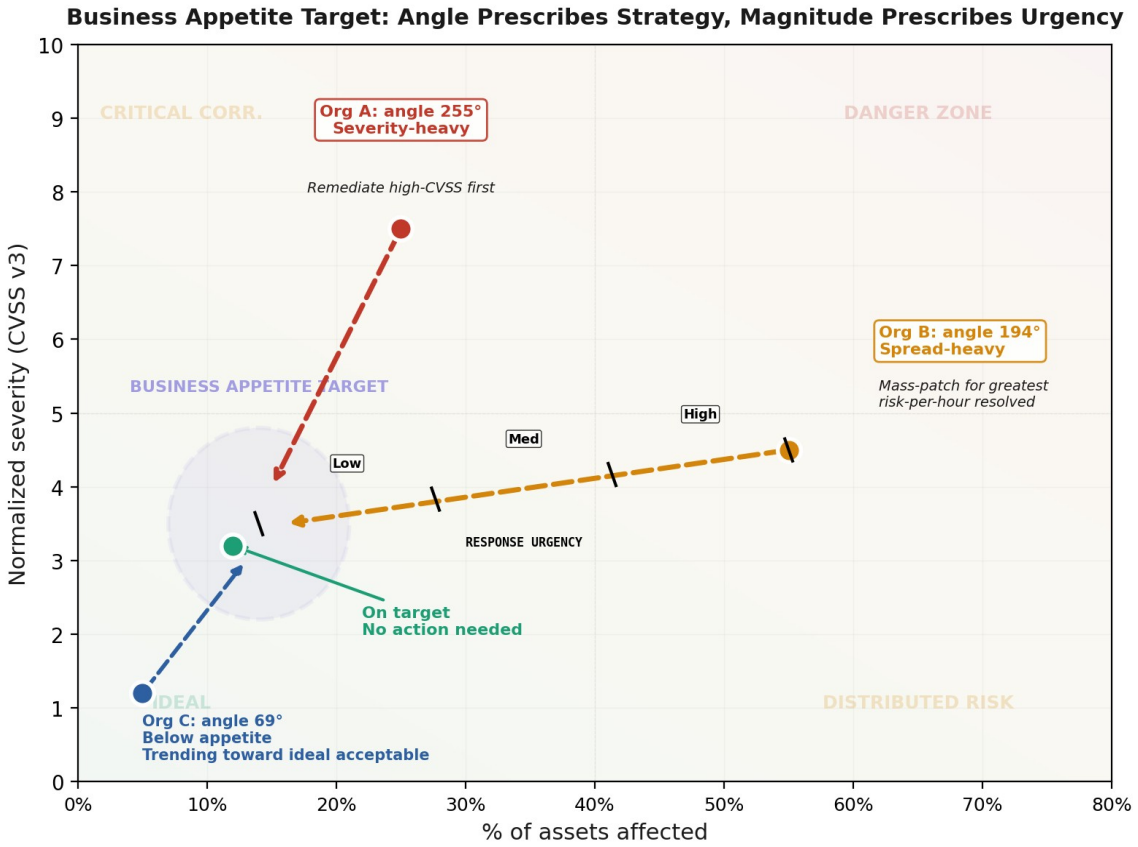

Figure 6: Business appetite target — angle prescribes strategy, magnitude prescribes urgency

The Angle of Attack: Prescriptive Remediation Strategy

The vector from the current CAP to the appetite center is not merely a measurement—it is a prescription. The angle of this vector describes what category of remediation will most efficiently return the organization to its target zone, and the magnitude describes how urgently.

Consider two organizations, both outside their business appetite target. Organization A's CAP sits at (25, 7.5)—high severity, moderate spread. The angle from their CAP to the appetite center is 255°, meaning the ideal vector points steeply downward. The prescription is clear: remediate high-CVSS findings first, because severity is the dominant contributor to the gap between current posture and target. Every point of CVSS remediated moves the CAP further toward the appetite zone than the same effort spent on reducing asset spread.

Organization B's CAP sits at (55, 4.5)—moderate severity but extreme spread. Their angle to the appetite center is 194°, pointing sharply left with minimal vertical component. The prescription inverts: mass-patching and configuration hardening across broad asset populations will produce the greatest risk-per-hour-spent reduction. Chasing the highest-severity finding first would barely move their CAP, because spread is what's keeping them outside the zone.

An organization already inside the appetite zone but trending outward (below-appetite, for example) sees a different prescription entirely: an angle pointing up-right toward the appetite

center means that trending toward higher risk is acceptable, because the organization is currently over-secured relative to its defined tolerance.

The magnitude of the vector—the distance from CAP to the business appetite target center—scales the urgency. Two organizations may have identical angles (identical remediation strategies) but different magnitudes. A magnitude representing a 5% improvement needed suggests a minor course correction achievable within normal operational cycles. A magnitude representing a 30% improvement needed suggests a structural gap requiring dedicated budget allocation, executive sponsorship, and potentially a multi-quarter remediation program.

When an organization sustains a consistent rate of CAP improvement—say 10% per year at a given budget level—the framework enables calculus on the rate of risk change over time. That rate can be mapped against funding, headcount, or any unit the business values: risk reduction per dollar spent, per FTE allocated, per change window consumed. This transforms risk management from a cost center narrative into an efficiency equation that leadership can optimize alongside other business functions.

CAP Total: Measuring Aggregate Risk on the Board

The Dilution Problem

The CAP location tells you where risk concentrates. It does not tell you how much total risk exists. This creates a vulnerability in the framework: an organization that discovers 50 new low-severity findings affecting zero critical assets will see its CAP drift toward the Ideal Zone—not because risk decreased, but because the denominator grew. The average moved; the total didn't.

Similarly, if new findings happen to land near the existing CAP coordinates, the position doesn't change at all—but there are now more risks on the board. A flat CAP with a growing finding count is not stability. It is accumulation masked by averaging.

Definition

CAP Total solves this by tracking the sum of all individual Euclidean distances from origin across every finding:

$$\text{CAP Total} = \Sigma \sqrt{\left(x_i/100\right)^2 + \left(y_i/10\right)^2} \quad \text{for all findings } i$$

This is not an average. It is a sum. Every new finding increases the total. Every remediated finding decreases it. CAP Total represents the aggregate risk burden the organization is carrying, regardless of where that risk is distributed on the chart.

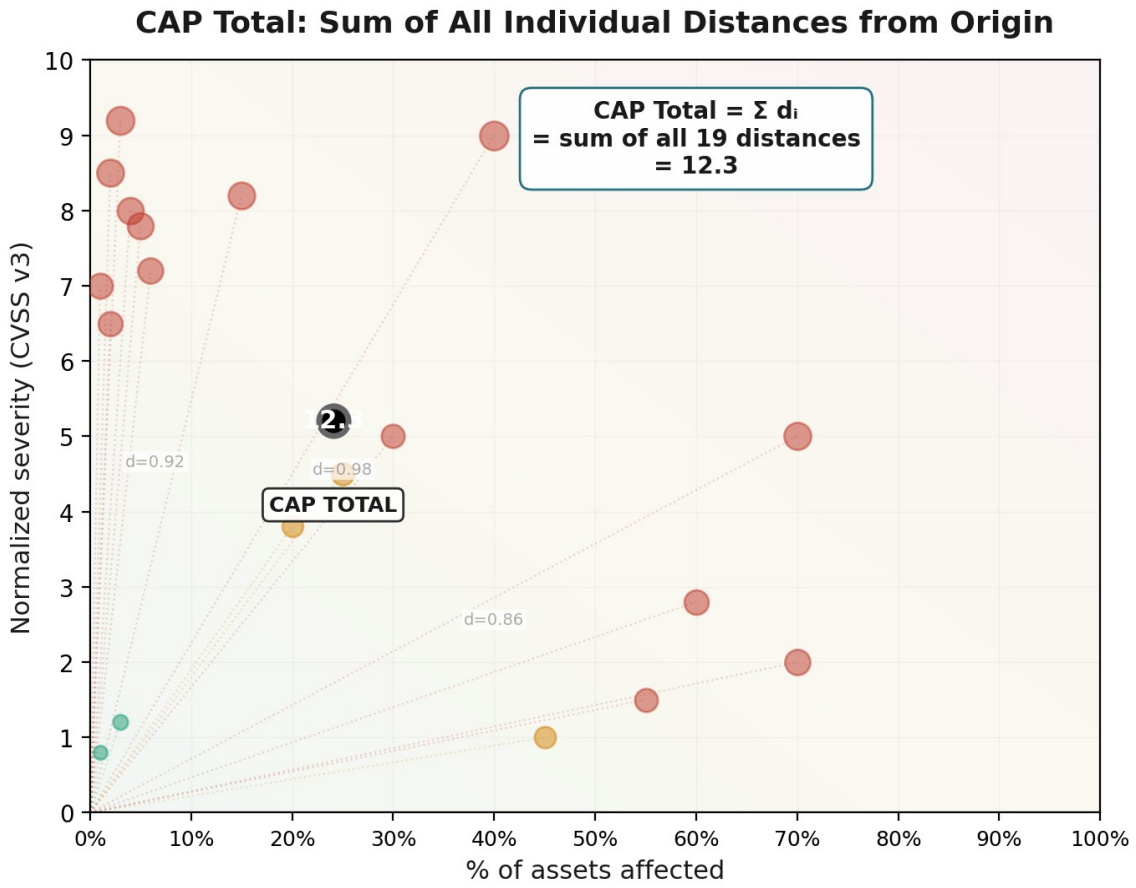

Figure 7: CAP Total calculation — the sum of all individual distances from origin, displayed as a single aggregate number

CAP Location and CAP Total are now tracked as a pair. Three scenarios become distinguishable: CAP moves toward ideal and Total decreases (genuine improvement), CAP moves toward ideal but Total increases (dilution by low-severity discovery—cosmetic improvement, not real), and CAP stays flat but Total increases (accumulation masked by averaging).

Bollinger Bands for Risk Trending

Applying Bollinger Bands—a technique from financial technical analysis—to the CAP Total time series creates an expected-range channel for organizational risk. A rolling mean of CAP Total with standard deviation bands establishes what “normal” risk accumulation and reduction looks like for the organization over time.

When CAP Total breaks above the upper band, something unexpected happened: a new vulnerability scan discovered an entire class of findings, a zero-day expanded scope across the environment, or a merger introduced unassessed assets. The band break is a signal, not a score—it triggers investigation and potentially emergency response.

When CAP Total breaks below the lower band, the security program is producing better-than-expected results. This is the mathematical signal that justifies reallocating security budget to

innovation—the data shows the program is outperforming its historical trend, and continued investment at the current level may yield diminishing returns.

Candlestick overlays on the same chart show period-over-period dynamics: each quarter's open (Total at period start), close (Total at period end), high (peak during period), and low (minimum during period). A series of candlesticks with rising closes means risk is accumulating faster than remediation. Falling closes means the program is winning. The visual language is immediately familiar to any executive who reads financial reports.

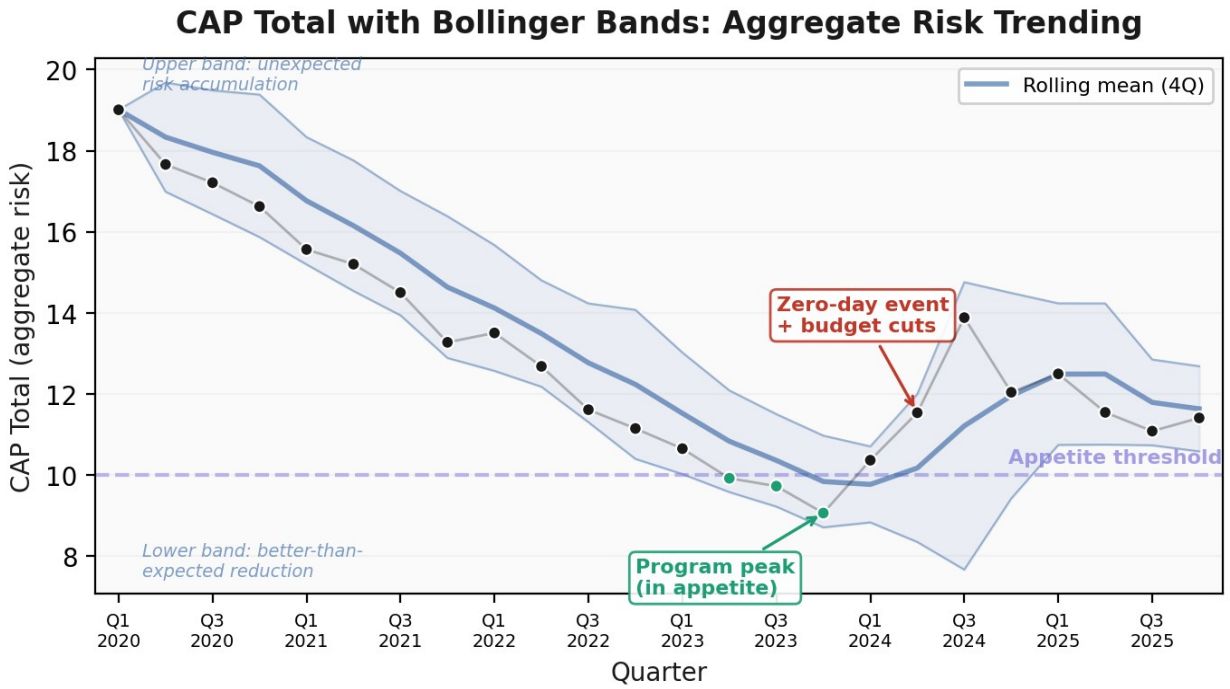

Figure 8: CAP Total with Bollinger Bands — aggregate risk trending with expected-range channels

Three-Dimensional Risk Space

When CAP Location (x, y) and CAP Total (z) are combined, the organizational risk posture occupies a point in three-dimensional space. The business appetite target is no longer an ellipse on a flat chart—it becomes a volume: an ellipsoid in 3D space that defines acceptable ranges for where risk concentrates, how much total risk exists, and where the two intersect. A CAP that sits within the 2D appetite ellipse but carries excessive total risk is outside the 3D volume—the dilution problem becomes geometrically visible.

Historical CAPs trace a trajectory through this 3D space. The spiral pattern visible in the 2D trajectory gains a vertical dimension: years where total risk decreased show the trajectory descending, while spikes in total risk (zero-day events, merger integrations) push the trajectory upward. The ideal path spirals downward and inward toward the center of the appetite volume.

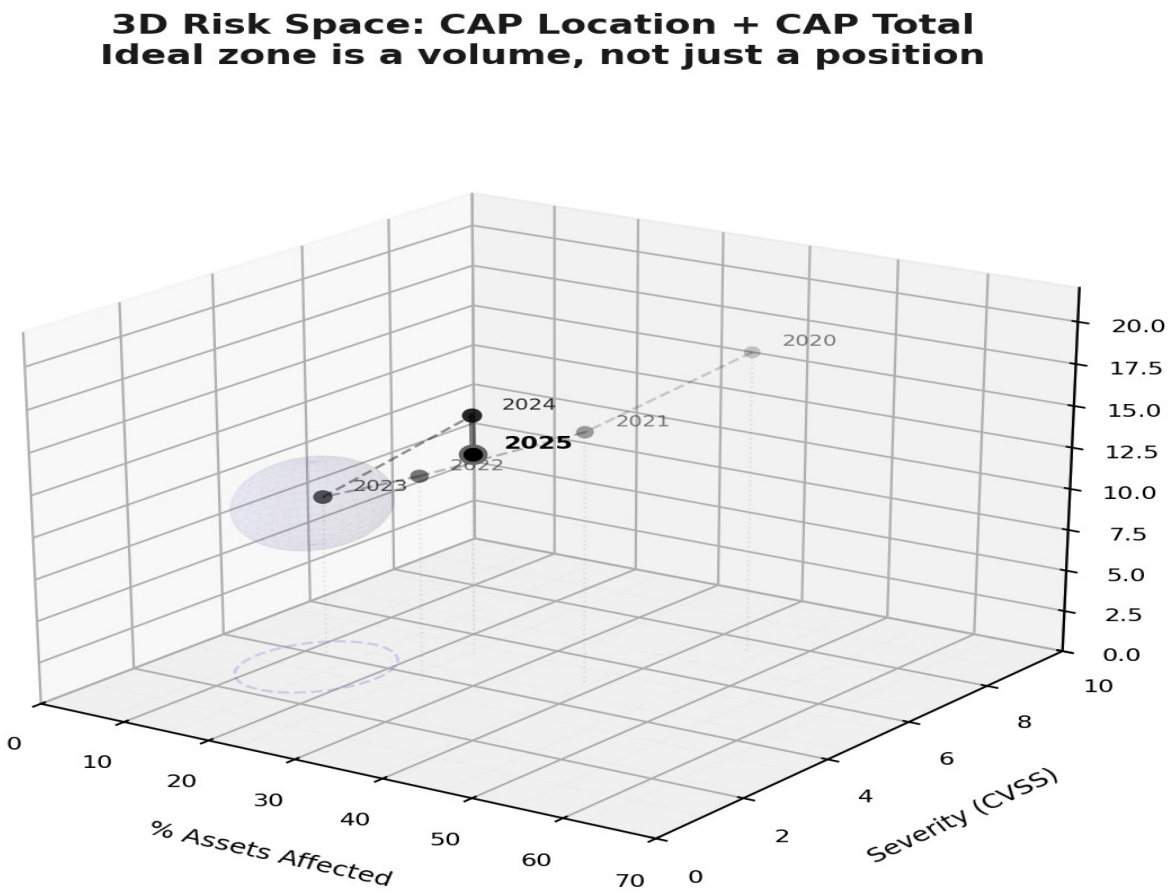

Figure 9: Three-dimensional risk space — CAP trajectory with ideal zone as a volume, not just a position

Sensitivity Analysis: Elimination Impact

Concept

One of the most powerful applications of the RVA framework is answering the question every CISO faces: if I can only fix three things this quarter, which three produce the greatest risk reduction?

Because the organizational risk score is derived from the CAP, and the CAP is the mean of all finding coordinates, we can compute the exact score impact of eliminating any individual finding by recalculating the CAP without that finding:

$$CAP' = \left((\Sigma x - x_i)/(n-1) \, , \, (\Sigma y - y_i)/(n-1) \right)$$

$$Impact_i = Score(CAP) - Score(CAP')$$

This produces a precise, integer-point score reduction for each finding. Findings are then ranked by elimination impact, giving leadership a mathematically defensible prioritization that accounts for both severity and spread simultaneously.

Practical Application

In a typical assessment with 18 tracked findings, the top three remediation actions by Euclidean distance might account for 60–70% of the achievable score reduction. This concentration effect is common: a small number of high-severity, moderate-spread findings (or moderate-severity, high-spread findings) disproportionately drive the organizational score. Identifying these findings analytically—rather than by subjective triage—ensures that limited remediation resources produce maximum posture improvement.

The elimination impact can also be expressed as a percentage of current score, enabling leadership to understand remediation ROI in relative terms. Stating that eliminating a specific finding would reduce organizational risk by 12% is more actionable than stating it would reduce the score by 4 points.

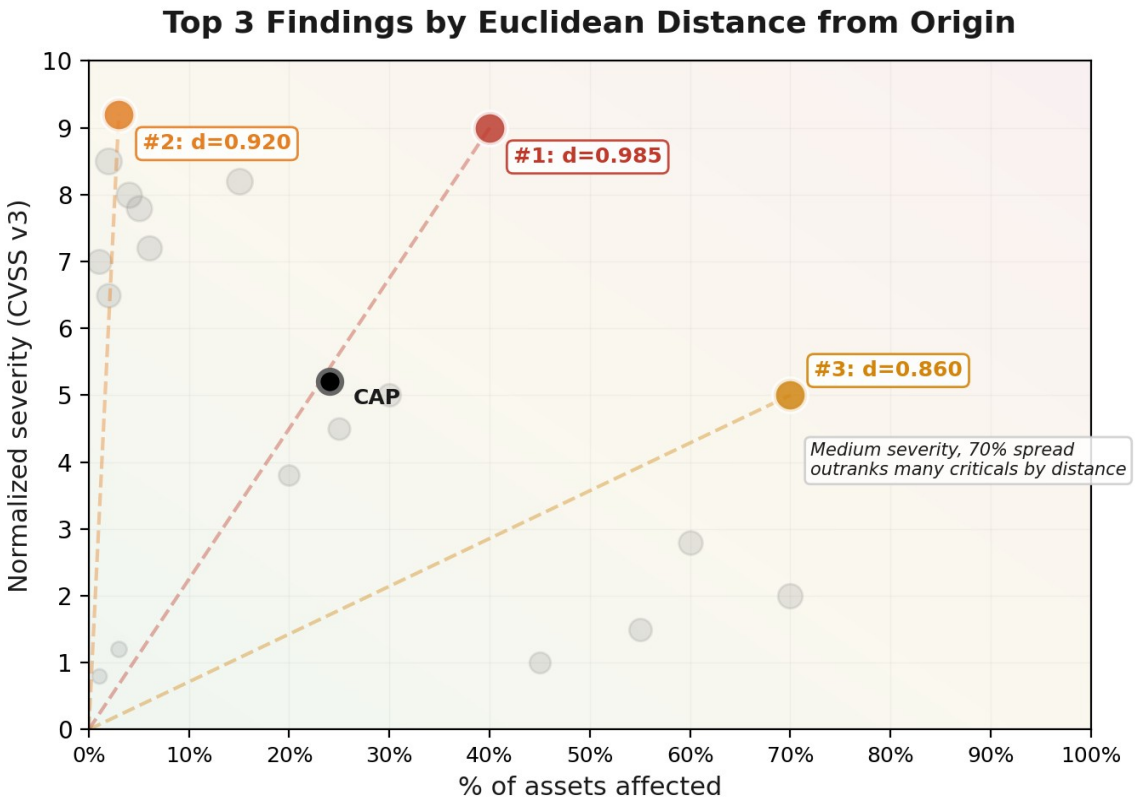

Figure 10: Top three findings ranked by Euclidean distance from origin

Executive Consumption: From Math to Meaning

Board-Ready in Thirty Seconds

The mathematical rigor of RVA means nothing if the output is not consumable by non-technical decision-makers. The framework is designed so that hovering over any CAP on an interactive dashboard produces an instant executive summary: the current risk score as a percentage, the prior-year comparison, the year-over-year delta in both points and percentage, the trend angle with a plain-language directional label, whether the organization is inside or outside its appetite target, a historical sparkline showing score trajectory across all periods, and the top three remediation actions by Euclidean distance with exact score reduction values.

No training is required to interpret this summary. A score of 37% with “−12% YoY” and “Ideal trend ↙️” communicates posture improvement to anyone in the room. A score of 54% with “+18%” and “Risk increasing ↗️” communicates a problem. The math runs underneath; the language runs on top.

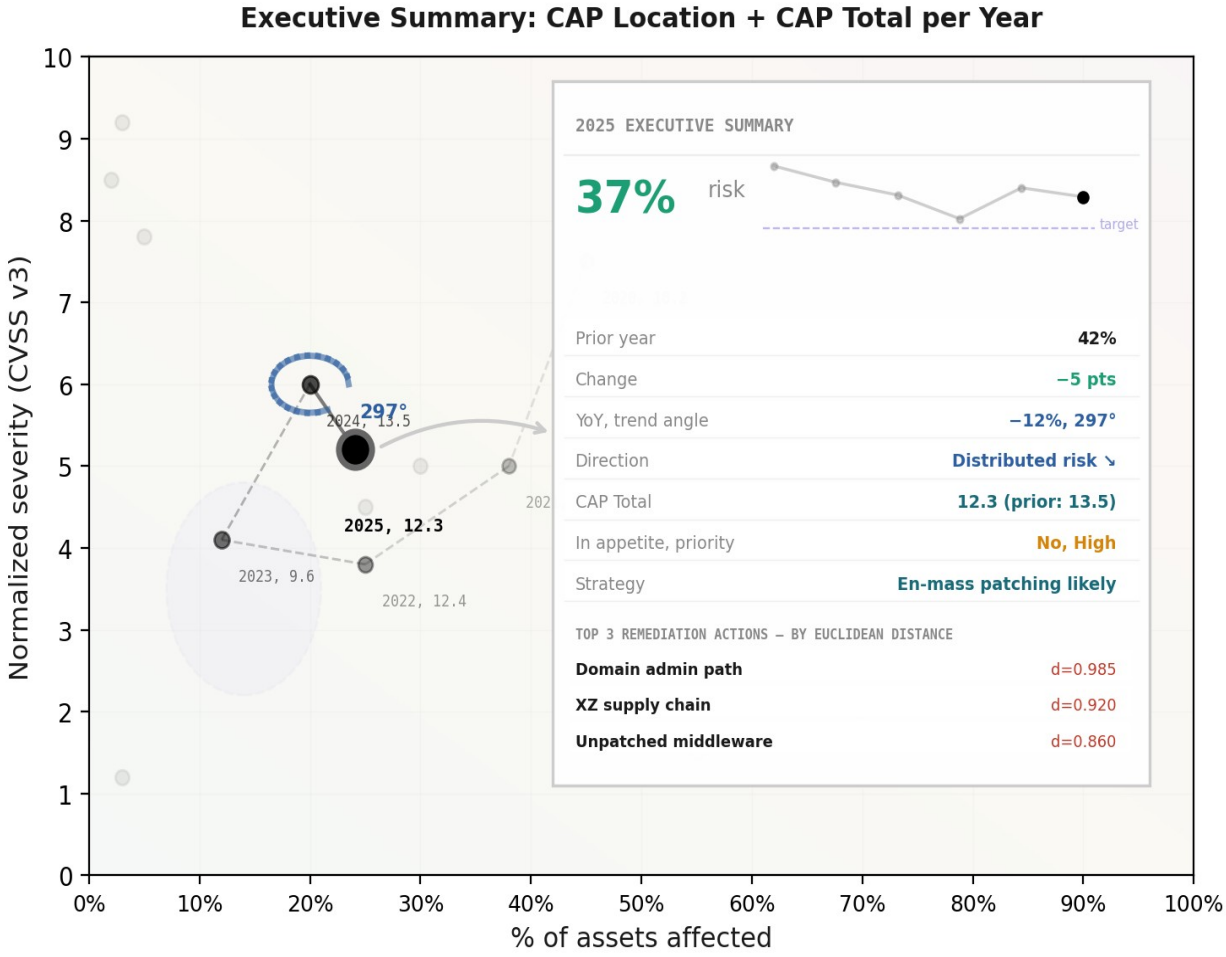

Figure 11: Executive summary panel — hover any CAP for an instant board-ready briefing

AI-Enabled Risk Inquiry

The RVA framework's structured data model is inherently compatible with AI-powered IT service management platforms. A board member could ask an AI assistant: "What are our top three risks right now?" or "What would happen to our score if we funded the EDR expansion?" and receive precise, mathematically grounded answers drawn directly from the live dataset. The elimination impact calculation already answers "what if we fix this?" with a single subtraction—an AI-enabled ITSM layer simply makes that calculation conversational.

Leadership gains direct, on-demand access to risk posture information without requiring a scheduled briefing, a dedicated analyst, or the ability to read a scatter plot. The data remains rigorous; the interface becomes natural language.

The Zero-Day Emergency Threshold

Purpose

Not all risk findings follow the standard remediation workflow. When a critical zero-day vulnerability is disclosed—particularly one with active exploitation in the wild—organizations need an automatic escalation mechanism that bypasses normal prioritization.

The Zero-Day Emergency Threshold is a curve plotted on the RVA chart above which any finding automatically triggers the Emergency Change Advisory Board (ECAB) process. Rather than relying on “patch now” directives from industry news or internal subjective opinion, newly identified risks are mapped onto the same chart as all existing organizational findings. The threshold calculates a defensible, math-driven strategy specific to your organization based on all known threats in play.

This distinction matters: nobody in the news knows which of your systems are non-production, microsegmented, public-facing, or regulated. External guidance is generic. Only your organization knows the context that determines whether a zero-day is a four-alarm emergency or a scheduled patch. RVA’s threshold provides the mathematical framework to make that determination objectively and defend it to auditors, leadership, or regulators.

Curve Definition

The threshold is defined as a logarithmic or linear decay function from a maximum severity at zero spread to a floor severity at full spread:

$$y_{thjesh} = \max(\text{floor}, \text{ceiling} - \text{decay} \times x)$$

Typical parameters: ceiling of 8.8 (a CVSS 8.8+ finding affecting even a single asset triggers emergency), floor of 6.0 (no finding below CVSS 6.0 triggers emergency regardless of spread), and a decay rate that connects them. These parameters are business-configurable and may be tied to the organization’s risk appetite—a more risk-tolerant organization may set a higher floor, while a regulated environment may lower the ceiling.

The threshold curve is displayed on the risk chart as a dashed line. Any finding plotted above this curve immediately routes to ECAB. This provides a visual, auditable, and mathematically defined escalation policy that removes subjective judgment from emergency classification.

Emerging Threat Detection

Not all findings appear above the threshold instantaneously. An emerging threat—a newly disclosed vulnerability with increasing exploitation activity, or a misconfiguration whose scope is expanding as new assets are deployed—may approach the threshold from below. Tracking the velocity of individual findings on the chart enables early warning: a finding trending upward toward the threshold line can trigger proactive change management before it breaches the emergency zone.

In organizations with mature CMDB integration, approaching-threshold findings can be cross-referenced against asset criticality, personnel overlap, and change windows. This transforms

the threshold from a reactive trigger into a predictive change-risk-management tool: the CMDB identifies which systems and teams would be affected by emergency remediation, enabling pre-staged response plans before the finding formally crosses the line.

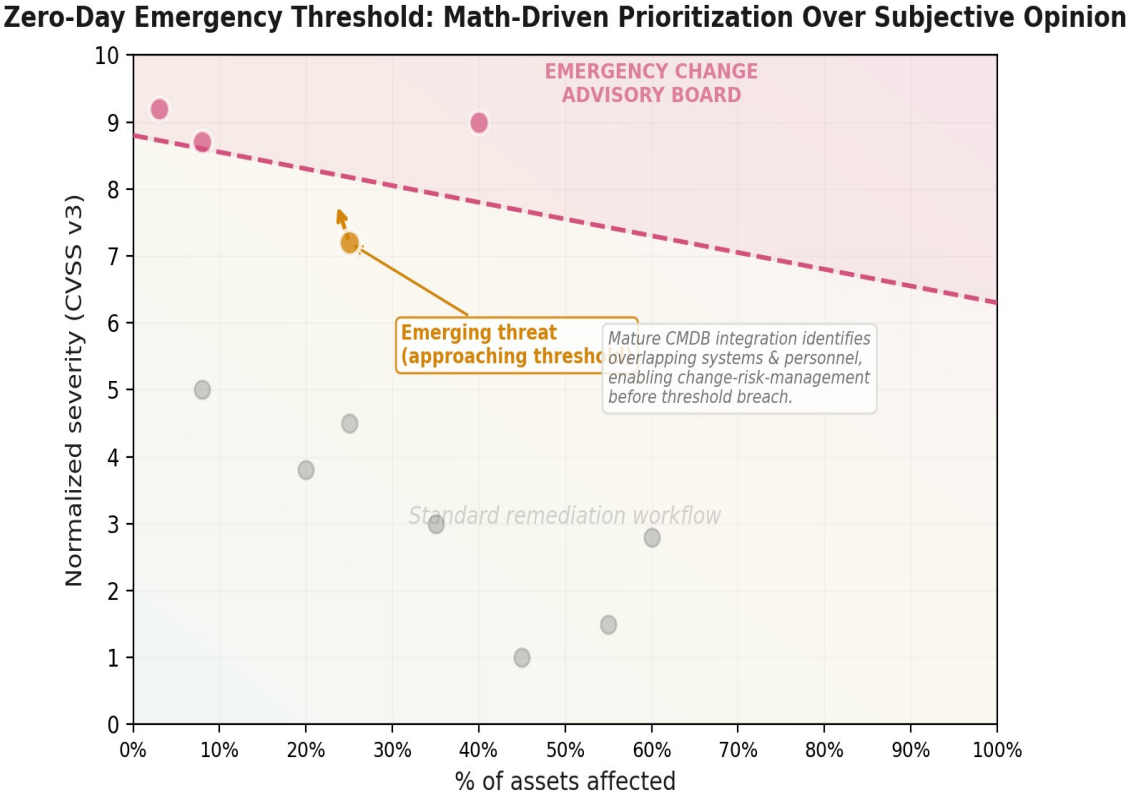

Figure 12: Zero-day emergency threshold — math-driven prioritization over subjective opinion

Data Normalization: Making Everything Comparable

The Universal Encoding Principle

The power of RVA lies in its ability to plot heterogeneous risk data on a single chart. Vulnerability scan results, penetration test findings, configuration audit gaps, compliance deficiencies, and threat intelligence indicators can all be visualized together—provided they can be encoded to the two-axis schema.

X-Axis Normalization (Asset Spread)

The x-axis requires expressing each finding's scope as a percentage of total organizational assets. For vulnerability scan results, this is typically straightforward: if 14 of 200 servers are affected, x = 7%. For findings that affect logical rather than physical assets (a policy gap, an architectural weakness), the encoding requires defining "assets" at an appropriate level of abstraction. A missing multi-factor authentication policy affecting all user accounts might be encoded as 100% of user-facing systems. The key requirement is consistency: the same encoding rubric must be applied across assessment periods to ensure CAP comparisons are valid.

Y-Axis Normalization (Severity)

For findings with CVSS v3 base scores, the y-axis is populated directly. For non-CVSS findings, a normalization rubric maps organizational severity to the 0–10 scale. This rubric should be documented, version-controlled, and applied consistently. Example mappings include: expired certificates (2.0–4.0 depending on scope and visibility), missing security headers (0.5–2.0), policy non-compliance (variable based on regulatory exposure), and architectural single-points-of-failure (4.0–7.0 based on blast radius).

The normalization rubric is itself a deliverable of the RVA engagement. It becomes organizational IP that ensures consistent, defensible encoding across assessment teams and reporting periods.

Organizations can also bake asset criticality weighting directly into the severity axis. A finding on a critical production system receives a multiplier—for example, 4× on a Low/Medium/High/Critical scale, or 3× on Low/Medium/High. This weighting is guaranteed to be applied equally across all findings because it derives from the asset inventory, not from analyst judgment. A CVSS 6.0 on a critical payment processing server becomes a weighted 8.0 (capped at 10.0), while the same CVSS 6.0 on a development sandbox remains 6.0. The framework surfaces the business impact that raw CVSS scores miss.

Implementation Architecture

Deliberate Technology Choice: Excel

RVA is implemented entirely in Microsoft Excel. This is a deliberate design decision, not a limitation. Excel provides several critical advantages for GRC tooling in the small and mid-size business market:

- Zero infrastructure cost. No servers, no SaaS subscriptions, no vendor lock-in.
- Universal availability. Every organization already has Excel.
- Data sovereignty. Your risk data stays on your systems, in your control. No third-party cloud dependency.
- Modifiability. Organizations can inspect, modify, and extend every formula. The tool is transparent, not a black box.
- Auditability. Every calculation is visible. An auditor can trace any output to its source data and verify the math.

Workbook Structure

The RVA Excel workbook contains the following sheets: a Findings Register where all risk data is entered with x and y coordinates, a CAP Calculator that computes the current-period CAP and organizational score, a Historical CAP sheet that stores prior-period CAPs for trending, a Vector Analysis sheet that computes year-over-year angles and magnitudes, an Elimination Impact sheet that ranks findings by score reduction potential, a Dashboard sheet with the XY scatter plot and all visual overlays, and a Configuration sheet where the organization defines its appetite target center, radii, and zero-day threshold parameters.

Known Limitations and Mitigations

Mean Sensitivity and CAP Behavior

The CAP is a mean, and means are sensitive to extreme values. A CVSS 10.0 finding affecting 100% of assets will pull the CAP significantly—and it should. That finding represents existential risk, and its weight in the average reflects its genuine impact on organizational posture. Any scoring system that mutes catastrophic findings in favor of a “balanced” average would be dangerous.

The practical nuance is one of communication, not calculation. When a catastrophic finding dominates the score, leadership may question whether the security program is making progress on the other 17 findings. The elimination impact analysis answers this directly: it shows the dominant finding’s specific contribution while simultaneously revealing that the remaining finding set has been improving. Both truths are visible in the same view.

Equal Axis Weighting

The current model normalizes both axes to 0–1 and weights them equally in the distance calculation. An organization might legitimately argue that a CVSS 9.2 finding affecting 5% of assets represents more risk than a CVSS 2.0 finding affecting 60% of assets, even though the latter has a larger distance from origin. Mitigation: introduce configurable weighting coefficients on each axis. The distance formula becomes $d = \sqrt{((w_1 \times x/100)^2 + (w_2 \times y/10)^2)}$ where w_1 and w_2 are business-defined weights. This becomes another parameter in the Configuration sheet.

Garbage In, Garbage Out

RVA operates on known findings. It cannot account for risks that have not been identified. An organization with excellent scanning coverage and a high risk score may be in better actual posture than one with poor scanning and a low score. The framework should always be presented with the caveat that it measures known risk, and that the completeness of the input data is a prerequisite for meaningful output.

Periodic and Continuous Modes

RVA supports both periodic assessment and continuous monitoring. In periodic mode, the CAP is computed at defined intervals—quarterly, annually, or per assessment cycle. In continuous mode, the framework accepts API feeds from vulnerability scanners, configuration management databases, and ITSM platforms, recomputing the CAP in real time as findings are discovered and remediated. The normalized coordinate system is inherently compatible with automated data ingestion: any source that can output a severity score and an asset count can feed the framework without manual intervention.

Conclusion

Risk Vector Analysis provides something the GRC industry has lacked: a mathematically rigorous, visually intuitive, and operationally actionable framework for measuring organizational risk posture over time. By preserving the full dimensionality of risk data through continuous coordinates rather than discrete matrices, computing organizational scores from Euclidean geometry rather than subjective weighting, and introducing vector mathematics to describe the direction and magnitude of change, RVA transforms risk reporting from a compliance exercise into a decision-support system.

The framework is intentionally transparent. Every formula is documented. Every calculation is auditable. The tool is Excel, not a proprietary platform. Organizations retain full ownership and understanding of their risk measurement process.

The question is no longer “are we compliant?” It is: “what angle are we trending, how fast are we moving, and does our trajectory land us inside the zone our business defined as acceptable?”

That question has a mathematical answer. This framework provides it.

About Ironclad Security LLC

Ironclad Security LLC provides practitioner-built governance, risk, and compliance frameworks for organizations that need enterprise-grade security posture management without enterprise-grade budgets. Our tools are built in Excel, delivered as one-time purchases, and designed to be operated independently. We consult ourselves out of a job—and that’s the point.

Contact: **ckinsel@pm.me**

Web: **ironcladsecurity.us**