CONTINUITY HUB • FIELD QUICK CARD

Annual Loss Expectancy (ALE) & Quantitative Risk Reading Guide

Current version — Oct 31, 2026
Update history: continuityhub.org/quantitative-risk-analysis-monte-carlo-loss-distribution-scenario-modeling/

CORE PRINCIPLE: REPLACE QUALITATIVE LABELS WITH ACTIONABLE CURRENCY

Qualitative labels ("High/Medium/Low") introduce political bias: IT calls cyber critical, operations calls infrastructure moderate, and finance underestimates business interruption. Quantitative analysis grounds decisions in **Annual Loss Expectancy (ALE)**, **Return on Risk Investment (RORI)**, and **Confidence Intervals** to justify recovery spend and prioritize RTO/RPO targets.

1. ALE: The Cornerstone Formula

FOUNDATIONAL

```
ALE = Probability (Annual) × Impact (Dollar Loss)
```

Single expected annual dollar loss for a specific threat scenario, enabling direct 80/20 risk ranking.

Worked Guide Example: Regional Power Outage

Annual Probability (P):	**8% (0.08)**
Impact (Lost Revenue):	**$2,500,000**
Annual Loss Expectancy (ALE):	**$200,000 / year**

The 80/20 Prioritization Law: Organizations consistently find that **20% of identified risks drive 80% of total ALE**. Allocate recovery investments strictly to the top quintile.

2. Return on Risk Investment (RORI)

CFO JUSTIFICATION

```
RORI = Annual ALE Reduction / Annual Recovery Cost
```

$$RODR = (\Delta ALE - \text{Annual Cost}_{DR}) / \text{Annual Cost}_{DR} = RORI - 1$$

Worked Example: Data Center Hot Standby

Current Baseline ALE (16h recovery):	**$400,000 / yr**
Revised ALE with Hot Standby (30m):	**$80,000 / yr**
Net Annual ALE Reduction:	**$320,000 / yr**
Annual DR Operating Cost:	**$150,000 / yr**
Gross RORI (Benefit-Cost):	**2.13 ($2.13 saved per $1)**
Payback Period:	**7 Months**

Payback = $150,000 / ($320,000 / 12) = 5.625 mos ≈ **7 months** (conservative). Net RODR = +1.13. Positive RODR proves the recovery solution pays for itself in avoided loss.

3. The 4 Fundamental Probability Distributions in Risk Modeling

DISTRIBUTION ENGINE

Distribution	Mathematical Profile	When to Deploy in Business Continuity	Guide Benchmark Example
Normal (Gaussian)	Symmetric bell curve; defined by Mean (μ) and Std Dev (σ). Clustering around central tendency.	Physical impact or recovery duration when variations are random and symmetrically bounded.	Recovery time averages **4 hours (σ=1h)**; 68% of recoveries complete between 3–5 hours.
Lognormal	Positively skewed, long right tail; cannot drop below zero. Captures catastrophic extreme events.	Financial losses, business interruption duration, repair costs where severe tail-risk dominates.	Power outages: typical **1–2 hours**, but rare severe outages extend to **24+ hours**.
Beta	Continuous, bounded strictly between 0 and 1; defined by shape parameters α and β.	Annual probability estimation under sparse empirical data; expert elicitation calibration.	Ransomware occurrence within 12 months elicited from experts as between **2% and 8%**.
Poisson	Discrete event count over fixed time interval; defined by single arrival rate parameter (λ).	Frequency modeling: number of distinct outages, outages, or cyber incidents per year.	Critical facility failures occur at rate of **λ = 1.2 failures per year**.

4. Frequency × Severity Compound Modeling

MONTE CARLO FOUNDATION

Why Separate Frequency & Severity?

- **Separate Controls:** Frequency is reduced by preventive controls (firewalls, UPS); severity is limited by containment and DR (backups, redundant sites).
- **Tail-Risk Visibility:** Captures low-frequency, high-severity catastrophic breaches that single point estimates erase.
- **Epistemic Uncertainty:** Widens confidence intervals honestly for rare events.

Cybersecurity Compound Example (from Guide):
• **Frequency Model:** Poisson (λ = 2.5 incidents/year)
• **Severity Model:** Lognormal (most incidents $50K–$200K, rare breaches >$5M)
• **Compound Engine:** Monte Carlo draws N from Poisson, then draws N losses from Lognormal to sum total annual impact.

CONTINUITY HUB • FIELD QUICK CARD

Loss Distribution Interpretation & Scenario Modeling

Current version — Oct 31, 2026
Update history: continuityhub.org/quantitative-risk-analysis-monte-carlo-loss-distribution-scenario-modeling/

5. Loss Distribution Interpretation (P10 / P50 / P90 / Mean)

EXECUTIVE METRIC TRANSLATION

Metric	Guide Value	Operational Meaning
P10	**$50,000**	**Optimistic:** Only a 10% chance losses are lower (or 90% chance loss meets/exceeds this floor).
P50 (Median)	**$180,000**	**Most Likely Outcome:** True 50/50 midpoint — half of simulated years are below, half above.
Mean (Expected)	**$250,000**	**Mathematical Average:** Pulled $70K above median by severe long-tail outcomes.
P90	**$600,000**	**Pessimistic:** Only a 10% chance annual loss exceeds this threshold (90% confidence ceiling).

Uncertainty Spread Analysis:
Spread = P90 − P10 = $600,000 − $50,000 = **$550,000**.
A wide spread indicates high epistemic uncertainty, demanding structured expert elicitation or sensitivity reduction rather than artificial point precision.

Executive Statement Template (Ready for Board / CFO):
*"Annual loss for this risk is expected at **$250K** (mean), with **80% confidence** that annual losses will fall between **$50K** and **$600K**."*

6. The 4-Step Monte Carlo Implementation Protocol

SIMULATION WORKFLOW

Step 1: Model System
Define critical variables (outage P, recovery time, lost revenue/hr). Assign distributions. Map cause-effect & cascading failure trees.

Step 2: Run Iterations
Draw independent/correlated random values. Calculate total financial outcome per trial. Repeat **10,000 to 100,000 times**.

Step 3: Analyze Results
Plot frequency histogram. Extract P10, P50, P90, and Mean. Quantify tail risk: *"Probability of annual loss exceeding $5M?"*

Step 4: Sensitivity
Tornado analysis: identify which input variables drive variance in outcome. Focus data collection and controls on high-sensitivity levers.

7. Discrete Scenario Expected Value: Hot vs. Warm vs. Cold DR Site

DECISION MATRIX

DR Topology	Scenario 1: No Incident (92%) Annual Maintenance Only	Scenario 2: Major Failure (6%) Recovery Cost + $100K/hr Loss	Scenario 3: Extended Event (2%) Multi-Day Recovery & Unavailability	Expected Value (EV) Annual Expected Cost
Hot Standby	92% × $350K = $322.0K	6% × ($500K direct + 1h × $100K = $500K) = $30.0K	2% × $1,000K (contained) = $20.0K	**$372,000 / yr**
Warm Site	92% × $300K = $276.0K	6% × ($250K direct + 6h × $100K = $850K) = $51.0K	2% × ($200K + $1.1M impact) = $26.0K	**$353,000 / yr**
Cold Site	92% × $100K = $92.0K	6% × ($100K direct + 18h × $100K = $1.9M) = $114.0K	2% × ($100K direct + $5.0M loss) = $102.0K	**$308,000 / yr***

***Trade-Off Rule:** Cold has the lowest EV ($308K) due to low standby maintenance ($100K), but carries severe tail risk ($5.1M in Scenario 3). Hot ($372K EV) costs $64K more annually but caps maximum operational downtime at 1 hour.

8. Communicating Quantitative Risk Across 3 Audience Levels

STAKEHOLDER PROTOCOL

Level 1: Board / C-Suite
Lead with one number: Mean Expected Loss. Present 3-point range (Best $0, Most Likely $0, Worst $4M). Frame mitigation as avoiding $2.8M worst-case loss.

Level 2: Finance / Audit
Present full loss distribution table (P10, P50, P90, Mean). Show sensitivity tornado charts. Quantify confidence intervals (±30%) transparently.

Level 3: Technical / BC Team
Provide full mathematical documentation: distribution types, elicitation notes, correlation coefficients, iterations (100k), convergence checks, and code.