CONTINUITY HUB • QUANTITATIVE RISK PACK

Part 1: Annual Loss Expectancy (ALE) Calculator & 80/20 Prioritization Log

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

WORKSHEET 1: ANNUAL LOSS EXPECTANCY (ALE) & THE 80/20 PRIORITIZATION LAW

ALE converts qualitative risk subjectivity into a singular dollar baseline for recovery budgeting: `ALE = Annual Probability × Dollar Impact`. In business continuity portfolios, **20% of risks generate 80% of total ALE exposure**. Use this worksheet to calculate baseline ALE, rank scenarios by exposure, and establish the 80% capital allocation cutoff.

Worked Baseline Reference: Regional Power Outage (Published Guide Benchmark)

VERIFIED STANDARD

Threat Scenario:	Regional Power Outage • Primary Utility Grid Loss
Annual Probability (P):	**8% (0.08)** • Based on 12-year local substation outage data
Impact per Event (I):	**$2,500,000** • 14 hours lost revenue, idle payroll, spoiled batch
Calculated Baseline ALE:	**$200,000 / year**

Audit & Prioritization Rule:
A point estimate of $2.5M feels overwhelmingly large, while 8% feels comfortably rare. Multiplying them establishes the true annual carrying cost: **$200,000/yr**. Any mitigation (e.g. secondary grid tie or generator lease) costing <$200K/yr yields positive risk retirement.

Enterprise ALE Prioritization & 80/20 Ranking Log

FILLABLE WORKSHEET

Rank	Risk Scenario & Affected System	Threat Category	Annual P (%)	Impact / SLE ($)	Calculated ALE ($)	Cumul. ALE ($)	80/20 Action
#1	Regional Power Outage (Primary Facility)	Infrastructure	8.0%	$2,500,000	$200,000	$200,000	TOP 20%
#2							[ ] Top 20%
#3							[ ] Top 20%
#4							[ ] Top 20%
#5							[ ] Top 20%
#6							[ ] Defer
#7							[ ] Defer
#8							[ ] Defer
#9							[ ] Defer
#10							[ ] Defer
				Total Portfolio ALE (100%):	**$**	**$**	**80% Threshold**

Execution Guide: Sort risks descending by Calculated ALE ($). Tally Cumulative ALE. The risks comprising the first 80% of Cumulative ALE are designated **Top 20% Mandatory Spend**. Risks below the 80% boundary accept operational retention or low-cost transfer.

CONTINUITY HUB • QUANTITATIVE RISK PACK

Part 2: Return on Risk Investment (RORI) & Payback Period Worksheet

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

WORKSHEET 2: RORI, RODR & CAPITAL RECOVERY PAYBACK ANALYSIS

CFOs reject qualitative arguments ("This is critical infrastructure"). Quantitative justification requires proving that the annual cost of the DR solution is less than the expected annual loss avoided. Use **Gross RORI** (>1.0 is break-even), **Net RODR** (RORI − 1), and **Payback Months** to secure executive budget approvals.

Worked Baseline: Data Center Hot Standby vs. Cold Recovery (Published Guide Benchmark)

VERIFIED STANDARD

Baseline Outage Recovery Time:	**16 hours (Cold recovery from offsite tapes)**
Baseline Annual Loss Expectancy (ALE_0):	**$400,000 / year**
Proposed DR Mitigation Architecture:	**Hot Standby at secondary carrier facility**
Revised Recovery Time Target:	**30 minutes (Synchronous mirroring)**
Revised Annual Loss Expectancy (ALE_1):	**$80,000 / year**
Net Annual ALE Reduction (ΔALE):	**$320,000 / year ($400K – $80K)**
Annual Total DR Standby Cost:	**$150,000 / year (Hosting, links, licenses)**

`RORI = $320,000 / $150,000 = 2.13`
"For every $1 spent on DR, save $2.13 in avoided losses."

`RODR = ($320,000 – $150,000) / $150,000 = +1.13` (+113% Net Return)

`Payback Period = $150,000 / ($320,000 / 12) = 5.6 mos ≈ 7 Months`

CFO Verdict: Immediate approval. Capital recouped in 7 months; produces ongoing net risk savings of $170,000 every year thereafter.

Enterprise Disaster Recovery Cost-Benefit & RORI Table

FILLABLE WORKSHEET

Workload / System	Proposed DR Solution	Baseline ALE ($)	Revised ALE ($)	ΔALE Saved ($)	Annual DR Cost	Gross RORI	Payback	Decision
Tier-0 Core Transaction DB	Hot Standby (30m)	$400,000	$80,000	$320,000	$150,000	**2.13**	7 Mos	APPROVE
								[ ] Yes / No
								[ ] Yes / No
								[ ] Yes / No
								[ ] Yes / No
								[ ] Yes / No
			Portfolio Totals:	**$**	**$**	**RORI:**	**Avg:**	**Summary**

CFO Investment Memorandum Paragraph Template (Ready for Submission)

EXECUTIVE COPY

"The proposed Disaster Recovery modernization for **[System / Workload Name]** replaces our current **[Baseline Hours]**-hour recovery posture with a **[Proposed Architecture: Hot / Warm / Cloud]** solution, reducing recovery time to **[Revised Minutes/Hours]**. This investment reduces our Annual Loss Expectancy (ALE) from **$[Baseline ALE]** to **$[Revised ALE]**, delivering an annual risk reduction benefit of **$[Net ΔALE]** against an annual operating cost of **$[Annual DR Cost]**. This represents a gross Return on Risk Investment (RORI) of **[RORI Ratio]** ($[RORI Value] saved in avoided loss per $1 invested), a positive Net RODR of **+[Net RODR %]**, and full capital payback achieved in **[Payback Months]** months."

CONTINUITY HUB • QUANTITATIVE RISK PACK

Part 3: Probability Distribution Quick-Reference Sheet & Calibration Guide

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

WORKSHEET 3: SELECTING AND CALIBRATING THE CORRECT PROBABILITY DISTRIBUTION

Single-point estimates conceal real-world operational variance. In quantitative risk analysis, every input variable is assigned a probability distribution. Choosing the incorrect shape (e.g. modeling long-tailed losses with a symmetric Normal curve) grossly underestimates tail risk. Use this quick-reference sheet to match data attributes to distribution shapes.

The 4 Essential Probability Distributions: Parameters, Use Cases & Guide Benchmarks

MATHEMATICAL GUIDE

Distribution	Mathematical Parameters	Data Profile & Bounds	Appropriate DR/BC Use Cases	Published Guide Example
Normal (Gaussian)	Mean (μ), Standard Deviation (σ)	Symmetric, unbounded (−∞ to +∞), bell curve. 68% within ±1σ, 95% within ±2σ.	Recovery execution times, labor hours, physical repair durations where outcomes cluster around average.	System recovery time averages **4 hours (σ=1h)**; 68% of recoveries complete between 3 and 5 hours.
Lognormal	Log-mean (μ_{log}), Log-std (σ_{log})	Positively skewed, strictly bounded at 0 (non-negative), heavy right tail for rare extremes.	Financial losses, business interruption duration, ransom demands, regulatory fines, reputational damage.	Most power outages last **1–2 hours**, but rare severe grid failures extend to **24+ hours**.
Beta	Shape Alpha (α), Shape Beta (β), [Min, Max]	Continuous, strictly bounded [0, 1] or between Min and Max. Highly flexible asymmetry.	Annual incident probability where historical data is sparse and calibrated via expert elicitation.	Expert elicitation establishes ransomware probability in next 12 months falls between **2% and 8%**.
Poisson	Rate parameter Lambda (λ)	Discrete event count (0, 1, 2, ... k) over a fixed annual time window. Variance equals mean.	Frequency modeling: counting annual occurrences of server failures, telecom drops, water leaks.	Critical facility failures occur at Poisson rate of **λ = 1.2 per year**.

Distribution Selection Decision Logic

SELECTION DIAGNOSTIC

Follow this 4-question decision gate:

1. **Are you counting discrete events per year?**
   → *Yes:* Use **Poisson**. Estimate average occurrences per year (λ).
2. **Are you estimating an uncertain probability bounded between 0% and 100%?**
   → *Yes:* Use **Beta / PERT**. Elicit Min, Most Likely, and Max bounds.
3. **Are you estimating financial loss or outage duration with tail risk?**
   → *Yes:* Use **Lognormal**. Skew prevents negative losses and models catastrophic tails.
4. **Is the variable symmetric, clustering tightly around an established mean?**
   → *Yes:* Use **Normal**. Truncate at zero if physical negative is impossible.

Expert Elicitation Calibration (3-Point PERT Conversion)

FIELD ELICITATION

When hard telemetry is absent, elicit 3 points from 3–5 Subject Matter Experts (SMEs):

PERT Mean: μ = (Min + 4×Most Likely + Max) / 6
PERT Std Dev: σ = (Max − Min) / 6

Interview Questions for SMEs:

- *"In the best 10% of cases, what is the fastest recovery?"* (Min)
- *"What is the typical outcome under normal operations?"* (Mode)
- *"If compounding failures occur, what is the 90th percentile worst-case?"* (Max)

Operational Variable Distribution Mapping Log

FILLABLE WORKSHEET

Variable Description	Assigned Distribution	Parameter 1 (Mean / λ / Min)	Parameter 2 (Std / Mode)	Parameter 3 (Max / Skew)	Data Source
Recovery Duration (Hrs)	Normal	Mean = 4.0 hrs	Std Dev = 1.0 hr	Min bound = 1.0 hr	Historical Logs
Incident Financial Loss ($)	Lognormal	Median = $180,000	Typical = $50K–$200K	Tail Max > $5.0M	Claims Data

WORKSHEET 4: THE 4-STEP MONTE CARLO IMPLEMENTATION AUDIT CHECKLIST

Monte Carlo simulation is mandatory when risks cascade or exhibit complex interdependencies that simple point formulas cannot capture. Iterating 10,000 to 100,000 random trials compiles a continuous loss distribution revealing true tail exposure. Complete this 4-step checklist prior to presenting model results to executive committees.

Monte Carlo 4-Step Implementation Checklist

AUDIT CHECKLIST

Stage	Mandatory Implementation Requirement	Verification Criteria & Target Metric	Sign-Off	Status
Step 1: Model the System	**1.1 Variable Inventory:** Define all critical variables across failure probability, recovery duration, and financial revenue impact per hour.	Variable list documented with explicit units of measure.	Risk Lead	[ ] Pass
	1.2 Distribution Assignment: Assign appropriate probability distributions based on empirical historical logs or calibrated SME elicitation.	No subjective point-estimates remaining in model core.	Lead Modeler	[ ] Pass
	1.3 Cascading Relationships: Map cause-and-effect dependencies and positive correlation links (e.g. ransomware × recovery difficulty).	Interdependency matrix / logic tree finalized.	IT / DR Lead	[ ] Pass
Step 2: Run Simulations	**2.1 Random Value Sampling:** Generate random draws from each assigned distribution using certified pseudo-random number generator.	Fixed random seed recorded for reproducibility.	Modeler	[ ] Pass
	2.2 Scenario Output Calculation: Calculate full annual financial loss outcome for every single simulated iteration.	Model calculation verified against manual benchmark trial.	Reviewer	[ ] Pass
	2.3 Sample Size Target: Execute **10,000 to 100,000 iterations** to ensure convergence of extreme tail percentiles.	Standard error of the mean < 1.0% of expected value.	Modeler	[ ] Pass
Step 3: Analyze Results	**3.1 Distribution Compilation:** Generate frequency histogram of annual loss outcomes and plot cumulative probability curve.	Visual distribution graph produced and inspected.	Analyst	[ ] Pass
	3.2 Percentile Extraction: Calculate and record 10th percentile (P10), 50th percentile (Median), Mean (Expected Value), and 90th percentile (P90).	P10, P50, P90, and Mean tabulated in standard executive format.	Lead Modeler	[ ] Pass
	3.3 Tail-Risk Quantification: Explicitly calculate probability of losses exceeding catastrophic risk threshold (e.g., *Loss* > *$5M*).	Tail probability expressed as percentage (e.g. 2.4%).	Risk Committee	[ ] Pass
Step 4: Sensitivity Analysis	**4.1 Tornado Chart Ranking:** Vary input assumptions ±20% to identify which variables contribute greatest variance to total expected loss.	Tornado diagram generated with top 3 levers highlighted.	Analyst	[ ] Pass
	4.2 Mitigation Targeting: Channel risk mitigation capital and data collection effort strictly into the highest-sensitivity variables.	Mitigation roadmap aligned with sensitivity ranking.	CFO / VP DR	[ ] Pass

Monte Carlo Tooling Architecture (Guide Software Review)

TOOL EVALUATION

Commercial Enterprise Add-ins:
- **@Risk (Lumivero):** Industry-standard Excel add-in. Deep integration with finance spreadsheets and BC databases. Ideal for corporate risk teams.
- **Crystal Ball (Oracle - retired legacy):** Legacy Excel add-in common in financial services and insurance underwriting portfolios.
- **Analytica (Lumina Decision Systems):** Visual influence diagrams, non-spreadsheet native engine for modeling complex interdependent enterprise networks.

Open-Source & Scripted Modeling:
- **Python (`numpy` + `scipy.stats`):** Zero licensing cost, unlimited simulation scale, full API integration with telemetry pipelines.
- **R (`stats` package):** Advanced statistical modeling and automated reporting.
- **Spreadsheet RAND():** Suitable only for quick 2-state scenario proofs; avoid for production tail-risk analysis.

CONTINUITY HUB • QUANTITATIVE RISK PACK

Part 5: Frequency × Severity Modeling & Cascading Failure Worksheets

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

WORKSHEET 5: DECOMPOSING RISK INTO FREQUENCY AND SEVERITY LEVERS

Treating risk as a single compound probability blinds leadership to effective controls. **Frequency** (how often events strike) is controlled by preventive hardening; **Severity** (the financial toll when an event strikes) is controlled by recovery architecture and containment. Model them as distinct distributions to capture true tail dynamics.

Cybersecurity Frequency × Severity Compound Benchmark (Published Guide Case)

VERIFIED STANDARD

Frequency Distribution:	**Poisson (λ = 2.5 incidents / year)**
Historical Evidence:	Internal telemetry & sector threat intelligence
Severity Distribution:	**Lognormal** (pos-skewed, heavy tail)
Typical Event Severity:	**$50,000 – $200,000** per incident
Extreme Tail Breach:	**> $5,000,000** (rare major database compromise)
Compound Simulation:	**Draw N ~ Poisson(2.5), then sum N × Lognormal draws**

Mitigation Strategy Insights:
• *Endpoint Protection:* Reduces Poisson rate λ from 2.5 → 1.0/yr.
• *Immutable Air-Gapped Backups:* Compresses Lognormal severity, cutting maximum breach loss from $5M+ down to $500K.
Modeling both separately proves to the board why backups are essential even when firewalls are strong.

Supply Chain Cascading Failure Model (Published Guide Benchmark)

CASCADING PATHWAYS

Context: Single tier-1 supplier provides 40% of critical manufacturing sub-assemblies. Pathways are interdependent:

Failure Pathway	Annual Probability	Cascading Mechanism	Simulated Impact on Operations
Supplier Facility Failure	**1.2% / year**	Factory fire, physical damage, utility outage	3-week factory halt; buffer stock buffers 10 days
Supplier Bankruptcy / Distress	**3.5% / year**	Financial insolvency, asset freeze by creditors	Permanent tooling loss; 90-day requalification delay
Geopolitical Disruption	**5.0% / year**	Tariff embargo, regional conflict, export closure	Port shutdown; alternate sourcing required
Transportation Interruption	**4.0% / year**	Carrier strike, port congestion, maritime route loss	2-4 week logistics transit delay

Enterprise Frequency × Severity Threat Decomposition Log

FILLABLE WORKSHEET

Threat Scenario	Frequency Model	Frequency Param (λ / P)	Severity Model	Severity Range (P50 - P90)	Annual ALE / Mean	Lead Control
Ransomware Encryption	Poisson	λ = 0.03 / yr	Lognormal	$2.0M – $8.0M	$240,000 / yr	Air-gap DR

CONTINUITY HUB • QUANTITATIVE RISK PACK

Part 6: Loss Distribution Interpreter & DR Expected Value Comparison

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

WORKSHEET 6: PERCENTILE INTERPRETATION, CONFIDENCE STATEMENTS & DR TOPOLOGY EV

The final output of Monte Carlo simulation is a continuous loss distribution curve. Never present a single average without its confidence spread. This worksheet provides the reading guide for **P10/P50/P90/Mean** metrics, the fillable 80% confidence executive script, and the complete **Hot vs Warm vs Cold DR Site Expected Value** decision model.

Loss Distribution Percentile Reading Guide & Tail Skew Analysis

VERIFIED STANDARD

Metric	Guide Value	Statistical & Practical Meaning
P10	**$50,000**	**Optimistic:** Only a 10% chance losses will fall below this floor.
P50	**$180,000**	**Median (Best Guess):** 50% probability above, 50% below.
Mean	**$250,000**	**Expected Value:** Pulled $70K above median by catastrophic tail events.
P90	**$600,000**	**Pessimistic:** 90% confidence ceiling; only 10% chance of exceeding.

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

80% Confidence Executive Template:
*"Annual loss for **[Risk Scenario]** is expected at **$[Mean Value]**, with **80% confidence** that annual losses will fall between **$[P10 Value]** and **$[P90 Value]**."*

Discrete Scenario Expected Value Model: Hot vs. Warm vs. Cold DR Topology

DECISION MODEL

DR Architecture	Scenario 1: No Incident (92%) Annual Standby Cost Only	Scenario 2: Major Failure (6%) Recovery Cost + $100K/hr Downtime	Scenario 3: Extended Event (2%) Multi-Day Facility Unavailability	Expected Value (EV) Annual Expected Cost
Hot Standby	92% × $350K = $322,000	6% × ($500K direct + 1h × $100K = $500K) = $30,000	2% × $1.0M (fully mirrored) = $20,000	**$372,000 / year**
Warm Site	92% × $300K = $276,000	6% × ($250K direct + 6h × $100K = $850K) = $51,000	2% × ($200K + $1.1M impact) = $26,000	**$353,000 / year**
Cold Site	92% × $100K = $92,000	6% × ($100K direct + 18h × $100K = $1.9M) = $114,000	2% × ($100K direct + $5.0M loss) = $102,000	**$308,000 / year***

***Strategic Analysis:** Cold site shows the lowest mathematical expected annual cost ($308K) due to minimal standby overhead ($100K). However, in Scenario 3, Cold site incurs an uncontained **$5.1M catastrophic loss** that may threaten corporate solvency. Hot Standby ($372K EV) represents a $64K annual premium that strictly eliminates catastrophic downtime risk.

Blank Enterprise DR Expected Value Calculator

FILLABLE DECISION TABLE

Evaluated DR Option	Standby Cost ($)	Scen 1 EV (P_1 = %)	Scen 2 EV (P_2 = %)	Scen 3 EV (P_3 = %)	Total Expected Annual Cost ($)
Option A:	$	$	$	$	$
Option B:	$	$	$	$	$
Option C:	$	$	$	$	$