

FIRST-YEAR CYBER BUDGET DECISION · ILLUSTRATIVE WORKED EXAMPLE

First-year cyber budget allocation after acquisition close

Example Manufacturing Co. (fictional post-close portfolio company) · annual gross loss (no insurance recovery modeled)

DECISION OWNER**PE operating partner (fictional)****DECISION DUE****Day 45 of the post-close 100-day plan (fictional)****FIRST-YEAR BUDGET****\$100,000****OBJECTIVE****Minimize modeled P95 annual loss within budget**

Recommendation

Fund **Recovery tabletop exercise (sponsor-required in this example) + Recovery capability investment** for **\$90,000**, inside the \$100,000 first-year budget. This is the in-budget option that minimizes modeled P95 annual loss (\$3,512,449 baseline → \$1,908,058). The mandatory tabletop exercise (\$15,000) is included in every option below, including this one.

Alternative considered

Recovery tabletop exercise (sponsor-required in this example) + Access hardening for \$70,000 is cheaper and has a *lower modeled mean* annual loss (\$402,566 vs. \$513,439 for the recommended option), but a *higher* modeled P95 (\$2,092,283 vs. \$1,908,058). Mean and P95 do not agree on a winner here — see page 3 for why, and for the sensitivity case that flips this ranking.

Evidence-held: not funded this cycle

Network segmentation (evidence-held) is excluded from the comparison rather than costed at \$0 or assumed to have zero effect — its cost and effectiveness are simply not yet known. Owner: **Portfolio-company security lead (fictional)**. See page 4 for the specific evidence requested and the condition for revisiting it.

What these dollars are. Every cost above is a hypothetical, all-in year-one assumption invented for this illustrative example — not an observed vendor quote, and not modeled beyond year one. The \$15,000 tabletop exercise is a fictional sponsor policy for this worked example, not a legal or insurer requirement: Fictional sponsor policy for this worked example, not a legal or insurer mandate. No loss reduction is claimed or modeled for this action; its cost is included because it is required in this scenario.

What changes, option by option

MODELED ANNUAL LOSS BY OPTION (FIXTURE: EXAMPLE-MANUFACTURING-CO-CYBER-DECISION-V1, AS OF 2026-09-05)

OPTION	COST	MEAN	P50	P95	ZERO-LOSS YEARS	STATUS
BASELINE (MANDATORY ONLY)	\$15,000	\$791,066	\$124,681	\$3,512,449	45.3%	Deferred: not P95-minimizing within budget
M + ACCESS HARDENING	\$70,000	\$402,566	\$0	\$2,092,283	67.4%	Deferred: not P95-minimizing within budget
M + RECOVERY CAPABILITY	\$90,000	\$513,439	\$275,501	\$1,908,058	45.3%	Proposed
M + BOTH ACTIONS	\$145,000	\$257,921	\$0	\$1,285,222	67.4%	Deferred: budget

How to read these numbers

Each option's annual loss is a **sum of simulated incident severities over one year**, including years with zero incidents (a Poisson-distributed event count, with independently drawn lognormal severities summed across events — the FAIR 2e / HTMA-Cyber compound loss model). **Mean** is the average across simulated years, pulled upward by rare large-loss years. **P50 (median)** is the typical simulated year, and is not equal to the mean when the loss distribution is skewed — it usually is not. **P95** is an *empirical percentile* of the simulated annual-loss distribution: the loss level exceeded in only about 5% of simulated years. It is not a confidence interval, not a worst case, and not an insurance return period.

M + Both actions is *not* the sum of the other two options' individual improvements. Recomputed jointly through the engine at the access-hardening frequency and recovery-capability severity assumptions — not the sum of the two actions independent quantile deltas.

No return-on-investment multiplier, valuation, or EBITDA-benefit figure is claimed anywhere in this pack. The only claims made are about modeled annual loss under invented input assumptions.

The recommended option lowers mean and P95, but raises P50

Versus baseline, Recovery tabletop exercise (sponsor-required in this example) + Recovery capability investment moves the modeled median (P50) annual loss **up**, from \$124,681 to \$275,501, even though it moves both the mean (\$791,066 → \$513,439) and the P95 (\$3,512,449 → \$1,908,058) down. This is a real, expected consequence of *which* assumptions this action changes — not a sign of a modeling error, and not a measurement of the control's real-world effectiveness.

The recommended action changes **both** the severity distribution's mean (\$1,000,000 → \$650,000 per event) and its log-sigma / shape (1.2 → 0.6), tightening the distribution around a lower center. A tighter distribution with more of its mass at a somewhat higher typical (median) value, but a much shorter tail, is not a uniformly lower outcome in every simulated year — it trades a higher typical year for far smaller upper-tail exposure. That tradeoff is a property of the assumed loss-distribution shape, not a measurement of the control's effectiveness in the field.

Sensitivity: does the choice hold up?

MODELED P95 UNDER A WEAKER ASSUMPTION FOR THE RECOVERY-CAPABILITY ACTION

CASE	SEVERITY MEAN	LOG- SIGMA	P95	WINNER BY P95
AS MODELED ABOVE	\$650,000	0.6	\$1,908,058	M + Recovery capability
IF RECOVERY CAPABILITY ONLY REDUCES MEAN SEVERITY TO \$800,000 (VS. \$650,000 ASSUMED ABOVE)	\$800,000	0.6	\$2,348,379	M + Access hardening

Under this weaker assumption, the P95-minimizing choice flips to **M + Access hardening**. The recommendation on page 1 depends on the recovery-capability severity assumption holding roughly as modeled — it is not robust to every plausible weaker assumption.

FOR REFERENCE: A WEAKER ASSUMPTION FOR THE ACCESS-HARDENING ACTION

CASE	FREQUENCY (Λ)	P95	CHANGES RANKING?
IF ACCESS HARDENING ONLY REDUCES FREQUENCY TO 0.6/YR (VS. THE 0.4/YR ASSUMED ABOVE)	0.6/yr	\$2,832,883	No

Held for evidence, not funded this cycle

Network segmentation (evidence-held)

Cost and effectiveness are unknown for this option — it is held out of the comparison entirely, not assumed to cost \$0 or to have zero effect. Cost and effectiveness assumptions are not yet available for this option. It is excluded from the comparison below rather than modeled at zero cost or zero effect.

OWNER

Portfolio-company security lead (fictional)

EVIDENCE REQUESTED

- A scoped cost quote for the segmentation project
- Tested (not just designed) segmentation coverage across the relevant network boundary
- An agreed mapping from that tested coverage to this decision's annual-loss scenario, so a frequency or severity assumption can actually be set

REVISIT WHEN

Return to this funding decision once a cost quote and a tested-coverage report exist; until then this option is held, not costed at \$0 or assumed to have zero effect.

How this was computed

Annual aggregate gross loss $S = \sum_{k=1..N} X_k$, $N \sim \text{Poisson}(\lambda)$ (event count), $X_k \sim \text{Lognormal}(\mu_{\text{LogUsd}}, \sigma_{\text{Log}})$ iid per-event severity (FAIR 2e / HTMA-Cyber compound collective-risk model). No insurance recovery modeled (gross loss only).

SIMULATION	50,000 trials, seed 20260905, as of 2026-09-05
ENGINE SOURCE	Vendored, byte-identical from product SHA 5c6ffc07c8af773a2944fe7ed22c4c8a2035d485, except:
ONLY ADAPTATION	Only external (non-relative) import in the vendored engine files; replaced with a local adapter that throws if reached. Every engine call in this generator passes explicit evtTail/pairedCrn/reprojectionMultiplier/degradationFactor options, so the adapter is never invoked on any path this generator exercises.
PAIRED RANDOM STREAMS	Enabled — improves numerical comparability between scenarios; it is not evidence of a causal effect.
ASSUMPTIONS	Frequency, severity, cost, and budget inputs are invented for this demonstration, not calibrated to observed rates. Only the loss metrics above are simulated — cost and budget figures are simple assumptions.
INDEPENDENCE	Event count is independent of severity. Per-event severities are independent and identically distributed (iid).
FINITE-SAMPLE VARIABILITY	All figures carry ordinary Monte Carlo sampling variability at this trial count; they are not exact closed-form values.
NOT VALIDATED BY	No actual customer outcome. Reproducible computation (same seed, same code, same output) is not the same thing as independent human validation of the modeling choices.

Accessible HTML equivalent

A stable, accessible HTML equivalent of this worked example is published at <https://valty.ai/proof/sample>, with a live methodology reference at <https://valty.ai/methodology>. This is a synthetic example of decision structure, not paid customer evidence; this fixed PDF does not recompute when those pages change.

Illustrative worked example. Company, decision, and all input assumptions are fictional. Not a customer outcome. This pack demonstrates a decision artifact's structure, not a production capability claim: no general-availability, deployment, or actuarial-validation claim is made anywhere in it.