Technical Supplement

Worked Examples: SVI, Dynamic Scenarios, and Vitality Optimization

Volume 2 — Technical Supplement

Volume 2 of the white paper "Large Load Oscillations and Grid Stability in the Era of AI Infrastructure". References to numbered Sections (e.g., Section 11.3) refer to Volume 1 — Main White Paper; Appendices A–D are internal to this volume.

Prepared as a strategic technical white paper for discussion and positioning with detailed calculation examples for the proposed concepts and methods.

REV. 18

Date: July 4, 2026

Authors: Chad Hirsh, Kevin Eckert, Stephen Abate, Kiamran Radjabli

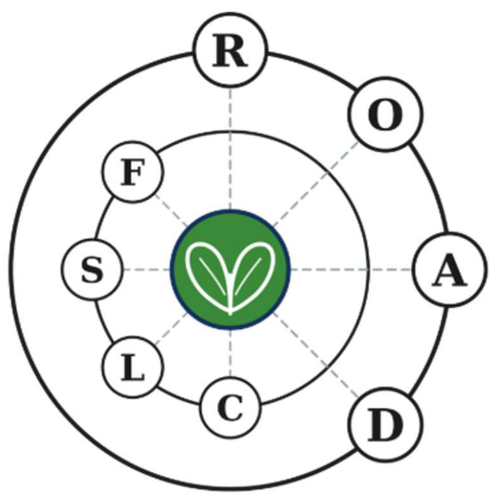

Appendix A — Example Fuzzy Vitality Assessment for Large Dynamic Load Oscillation

A.1 Example System Description

Consider a 600 MW AI campus connected to a relatively weak transmission corridor. The facility exhibits repetitive oscillatory behavior associated with synchronized GPU training workloads. PMU visibility is available at the transmission interconnection point, but point-on-wave observability inside the campus is limited.

The following operational conditions are assumed:

• Damping ratio: 4%
• Oscillation magnitude: 35 MW peak-to-peak
• Synchronization pressure: High
• PMU observability: Moderate
• BESS smoothing capability: Partial
• Workload scheduling flexibility: Moderate

A.2 Example Fuzzy Membership Functions

In the notation below, μ represents a fuzzy membership function. The subscripts identify the variable being evaluated. $\mu_D(low)$ and $\mu_D(high)$ indicate membership in low-damping and high-damping vitality states respectively. The two functions form a strong fuzzy partition: $\mu_D(low) + \mu_D(high) = 1$ for all ζ ≥ 2%. The transition zone is [2%, 8%] with a symmetric crossover at ζ = 5%, where both functions equal 0.500.

Low-damping membership function:

$$\mu_D(\zeta)_{\text{low}} = \left\{ \begin{array}{cc} 1, & \zeta \leq 2\% \\ (8-\zeta)/6, & 2\% < \zeta < 8\% \\ 0, & \zeta \geq 8\% \end{array} \right\}$$

High-damping membership function:

$$\mu_D(\zeta)_{\text{high}} = \left\{ \begin{array}{cc} 0, & \zeta \leq 2\% \\ (\zeta-2)/6, & 2\% < \zeta < 8\% \\ 1, & \zeta \geq 8\% \end{array} \right\}$$

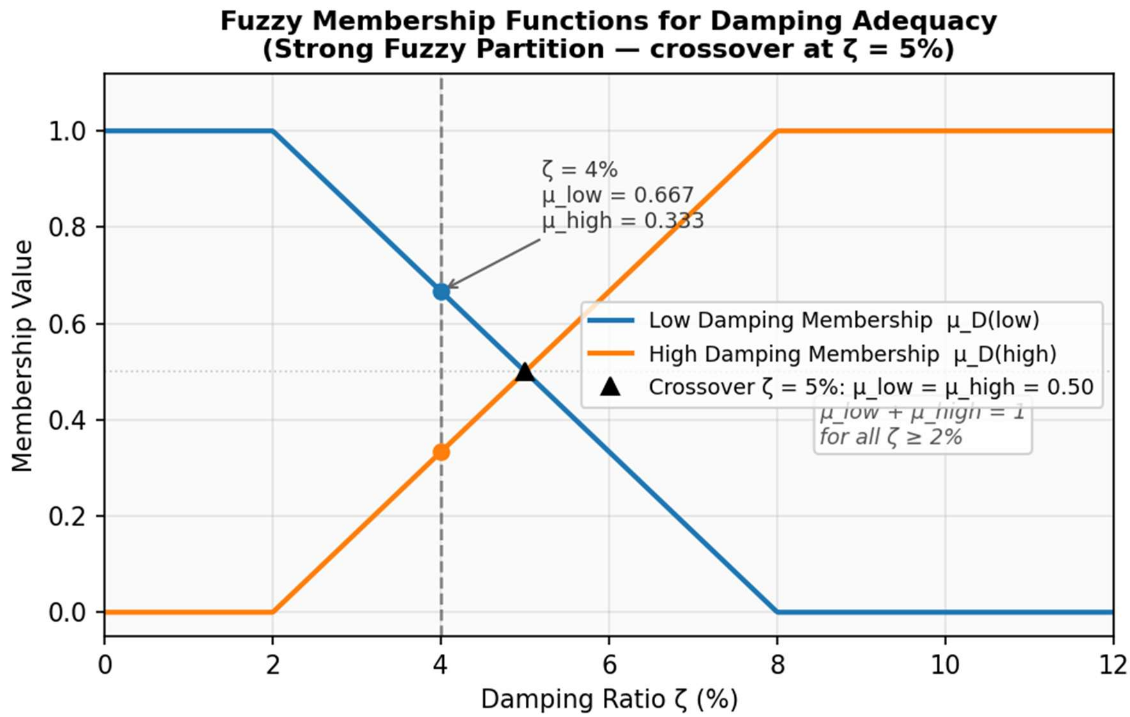

Figure A-1. Fuzzy membership functions for damping adequacy (strong fuzzy partition). $\mu_D(low) + \mu_D(high) = 1$ ***for all ζ ≥ 2%. Crossover at ζ = 5%: both equal 0.500. At the baseline ζ = 4%:*** $\mu_D(low)$ ***= 0.667,*** $\mu_D(high)$ ***= 0.333.***

Operational interpretation: μ_D(low) = 0.667 means the system is predominantly in a fragile damping state — not marginally so. Two-thirds of its damping "weight" falls on the low-vitality side of the partition. This is not a borderline reading; it indicates a system that a utility engineer should treat as damping-deficient for operational planning purposes, even though the 4% value remains above the conventional 3% alarm threshold.

A.3 Example Fuzzy Inference Rules

These rules are intended as an illustrative rule base. A production fuzzy-inference implementation would also define explicit membership functions for observability, synchronization pressure, oscillation magnitude, BESS smoothing, workload flexibility, and adaptive controls, plus aggregation and defuzzification logic.

Rule 1:
IF damping is LOW
AND synchronization pressure is HIGH
AND observability is POOR
THEN vitality is VERY LOW.

This rule fires in the baseline example (μ_D(low) = 0.667 satisfies "damping is LOW"). In practice, a utility engineer would see Rule 1 activate during periods when a large AI campus is running synchronized training jobs at night, PMU coverage is limited to the transmission interconnection, and measured damping is trending below 5%. It is the highest-severity rule and overrides numerical SVI scores for operational decision purposes.

Rule 2:
IF damping is MODERATE
AND adaptive controls are STRONG
AND observability is HIGH
THEN vitality is MODERATE TO HIGH.

Rule 2 does not fire in the baseline scenario because observability is moderate rather than high and adaptive controls are only partial. It represents the target state an upgrade program is designed to achieve: a facility with deployed fast-response controls, full POW coverage, and sufficient damping reserve would satisfy all three antecedents and would be classified as operationally healthy.

Rule 3:
IF oscillation magnitude is HIGH
AND forcing frequency approaches a known local mode
THEN vitality is LOW even if average load level is acceptable.

This rule captures the resonance hazard that classical average-load assessments miss. A 35 MW peak-to-peak oscillation from synchronized GPU workloads is not alarming in isolation, but if its frequency coincides with a known local mode frequency for the corridor, the system is at elevated risk of modal reinforcement. The rule explicitly overrides an otherwise-acceptable average load reading, reflecting the physical reality that magnitude and frequency together determine risk, not magnitude alone.

Rule 4:
IF BESS smoothing is AVAILABLE
AND workload scheduling is FLEXIBLE
AND PMU/POW visibility is HIGH
THEN vitality improves.

Rule 4 is a remediation rule rather than a risk rule. It does not fire in the baseline scenario — BESS smoothing is only partial, workload scheduling is moderate, and POW visibility is limited — but it identifies the specific investments that would shift the system toward higher vitality. It is the rule a utility engineer would consult when building an interconnection upgrade plan: each antecedent corresponds directly to a capital investment or operational change that can be evaluated for cost-effectiveness against the vitality improvement it enables.

A.4 Example Vitality Assessment

Under a classical deterministic framework, the example system might still be classified as technically stable because no explicit stability limit has been violated.

However, a fuzzy vitality assessment would recognize that:
• damping is marginal,
• synchronization pressure is elevated,
• observability is incomplete,
• and oscillation magnitude is significant.

The resulting vitality assessment may therefore classify the system as dynamically stable but operationally fragile.

This distinction is important because vitality-oriented assessment identifies degradation trends before classical instability emerges.

To make this concrete: the baseline $\zeta = 4\%$ produces $\mu_D(\text{low}) = 0.667$, meaning the damping dimension is already predominantly in the fragile state. Combined with high synchronization pressure and limited POW observability, Rule 1 fires with nonzero strength. Rule 3 may also contribute if the 35 MW oscillation frequency approaches a corridor local mode. Rule 2, which would elevate vitality, cannot fire because observability and adaptive controls are insufficient. Rule 4, the remediation rule, does not fire because BESS smoothing is only partial and workload scheduling is moderate. The aggregate inference therefore yields VERY LOW to LOW vitality — not because any single threshold has been crossed, but because the combination of marginal damping, elevated synchronization pressure, incomplete observability, and absent compensating controls leaves the system with no meaningful vitality reserve. That is the operational meaning of “dynamically stable but operationally fragile”: the classical stability criterion is satisfied, but the system has no cushion against further deterioration.

Appendix B — Example of SVI Calculation

B.1 Purpose and Scope

This appendix provides the full quantitative System Vitality Index (SVI) calculation framework. It implements the weighted SVI formula, the AI augmentation term, the fuzzy membership functions, and the vitality-constrained optimization check across four representative scenarios.

The base scenario replicates the Appendix A example exactly: a 600 MW AI campus on a relatively weak transmission corridor, exhibiting repetitive oscillatory behavior from synchronized GPU training workloads, with ζ = 4% and limited point-on-wave observability.

Validation caveat: the numerical values in this supplement are internally calibrated worked examples. For production use, the scoring breakpoints, weights, sensitivity coefficients, and investment functions must be recalibrated with site-specific PMU/DFR/SCADA data, dynamic-model studies, vendor control-response tests, and post-event operating records. The examples demonstrate the calculation chain; they are not a substitute for utility-specific EMT/RMS validation or commissioning tests.

B.2 Formula Chain

B.2.1 SVI_{base} — weighted linear sum

The base SVI is the normalized weighted sum of eight vitality dimensions:

$$\mathrm{SVI}_{\mathrm{base}} = \sum_{i=1}^{n} (w_{\mathrm{i}} \times X_{\mathrm{i}}), \quad \sum_{i=1}^{n} w_{\mathrm{i}} = 1.0 \ , X_i \in [0, 10]$$

Each X_i is derived deterministically from physical measurements using the scoring rubrics in Section B.3. No judgment is required at the aggregation stage; all judgment is encoded in the rubric mapping functions.

Notation convention for the worked examples: SVI_{base} is the weighted physical-capability score; $\Delta AI = \alpha - \beta$ Is the AI augmentation term; SVI_{final} = clamp(SVI_{base} + ΔAI, 0, 10). Dimension scores are denoted X_D, X_O, X_A, X_R, X_F, X_C, X_S, and X_L. The workload-correlation coefficient is written ρ_{max} in the dynamic scenario and p_{max} in the optimization appendix; the two symbols denote the same physical quantity.

B.2.2 AI augmentation term

Section 6.3 of the white paper introduces an additive AI correction:

$$\mathrm{SVI}_{\mathrm{AI}} = SVI_{base} + \alpha * I_p - \beta * I_s$$

$$\mathrm{SVI}_{\mathrm{final}} = clamp(SVI_{AI}, 0, 10)$$

Ip and Is are scored on a 0-10 integer scale using the sub-capability rubrics in Section B.5. α and β are selected from the tiered governance framework in Section B.5.4. The clamp in the SVI_{final} expression is not cosmetic: it is required because the unbounded sum SVIAI can exceed the [0, 10] scale in either direction, and the conditions under which it does so carry interpretive significance.

Formal bounds on $\mathrm{SVI_{AI}}$. Since each $X_i \in [0, 10]$ and the weights form a convex combination ($\Sigma w_i = 1$), $SVI_{base} \in [0, 10]$ without exception. The AI augmentation term then adds $\alpha * I_p \in [0, 2.0]$ (since $\alpha \le 0.20$ and $I_p \le 10$) and subtracts $\beta * I_s \in [0, 2.0]$ (since $\beta \le 0.20$ and $I_s \le 10$). Therefore:

$$\mathrm{SVI_{AI}} = \in [0 - 2.0,\ 10 + 2.0] = [-2.0,\ 12.0]$$

$$\mathrm{SVI_{final}} = clamp(SVI_{AI},\ 0,\ 10) \in [0,\ 10]$$

The clamp creates a nonlinearity: two facilities with $\mathrm{SVI_{AI}} = -1.5$ and $\mathrm{SVI_{AI}} = -0.1$ both produce $\mathrm{SVI_{final}} = 0$, appearing identical in a constraint check. Similarly, $\mathrm{SVI_{AI}} = 10.5$ and $\mathrm{SVI_{AI}} = 11.9$ both produce $\mathrm{SVI_{final}} = 10$. The table below specifies the activation conditions and their interpretive implications.

Clamp condition	When it binds (example)	Interpretive implication
Lower clamp ($\mathrm{SVI_{AI}}$ < 0)	$SVI_{base} < \beta \cdot I_s - \alpha \cdot I_p$, e.g. $SVI_{base} = 1.0$, β = 0.20, $I_s = 9$, α = 0: $\mathrm{SVI_{AI}} = -0.8$ → clamped to 0	$\mathrm{SVI_{final}} = 0$ understates severity; raw $\mathrm{SVI_{AI}}$ must be reported separately. Constraint verdict (VIOLATED) is unchanged but the degree of violation is masked.
Upper clamp ($\mathrm{SVI_{AI}}$ > 10)	$SVI_{base} > 10 - \alpha \cdot I_p$, e.g. $SVI_{base} = 9.0$, α = 0.20, $I_p = 10$: $\mathrm{SVI_{AI}} = 11.0$ → clamped to 10	AI benefit is floored at the physical maximum. Upper clamp binds only when $SVI_{base} > 8.0$ at maximum α = 0.20 and $I_p = 10$; rare in practice.
No clamp (0 ≤ $\mathrm{SVI_{AI}}$ ≤ 10)	All four worked scenarios in B.6 and B.7: baseline (2.618), post-BESS (3.890), best-practice (8.185), crisis (0.360). None triggers a clamp.	$\mathrm{SVI_{final}} = \mathrm{SVI_{AI}}$; full information is preserved. The clamp is transparent in normal operation and activates only at the extremes described above.

Reporting rule: whenever $\mathrm{SVI_{AI}}$ falls outside [0, 10], assessors must report the raw $\mathrm{SVI_{AI}}$ value alongside $\mathrm{SVI_{final}}$. A note such as "$\mathrm{SVI_{final}} = 0.0$ (raw $\mathrm{SVI_{AI}} = -0.8$; lower clamp applied)" preserves the information lost by clamping and prevents two qualitatively different situations from appearing numerically equivalent.

> *Verification against all four B.5 scenarios: baseline $\mathrm{SVI_{AI}} = 3.218 - 0.600 = 2.618$ (no clamp); post-BESS $\mathrm{SVI_{AI}} = 3.990 - 0.100 = 3.890$ (no clamp); best-practice $\mathrm{SVI_{AI}} = 7.785 + 0.400 = 8.185$ (no clamp); crisis $\mathrm{SVI_{AI}} = 2.160 - 1.800 = 0.360$ (no clamp — remains positive). None of the four worked scenarios activates the clamp. The lower*

> *clamp would activate in a more extreme crisis, e.g. SVI_{base} = 1.0 with β = 0.20, I_s = 9: SVI_{AI} = 1.0 − 1.8 = −0.8 → SVI_{final} = 0.*

B.2.3 Fuzzy membership functions

The piecewise membership functions for damping adequacy from Appendix A provide a complementary qualitative description of the D dimension. Figure B-1 illustrates these functions; The following table clarifies the boundary behavior at ζ = 5%.

Region	$\mu_D(low)$	$\mu_D(high)$	Interpretation
ζ ≤ 2%	1.000	0.000	Fully in low-vitality state
2% < ζ < 8%	(8−ζ)/6	(ζ−2)/6	Transition zone — sum = 1 throughout; $\mu_{low} + \mu_{high} = 1$
ζ = 5%	0.500	0.500	*Crossover* — $\mu_D(low) = \mu_D(high) = 0.5$; equal membership; sum = 1.0
ζ ≥ 8%	0.000	1.000	Fully in high-vitality state (μ_{high} = 1)

The functions form a strong fuzzy partition: $\mu_D(low) + \mu_D(high) = 1$ for all ζ ≥ 2%, with a symmetric crossover at ζ = 5% where both equal 0.500. Formulae: $\mu_D(low) = (8 - \zeta)/6$ for 2% < ζ < 8%; $\mu_D(high) = (\zeta - 2)/6$ for 2% < ζ < 8%.

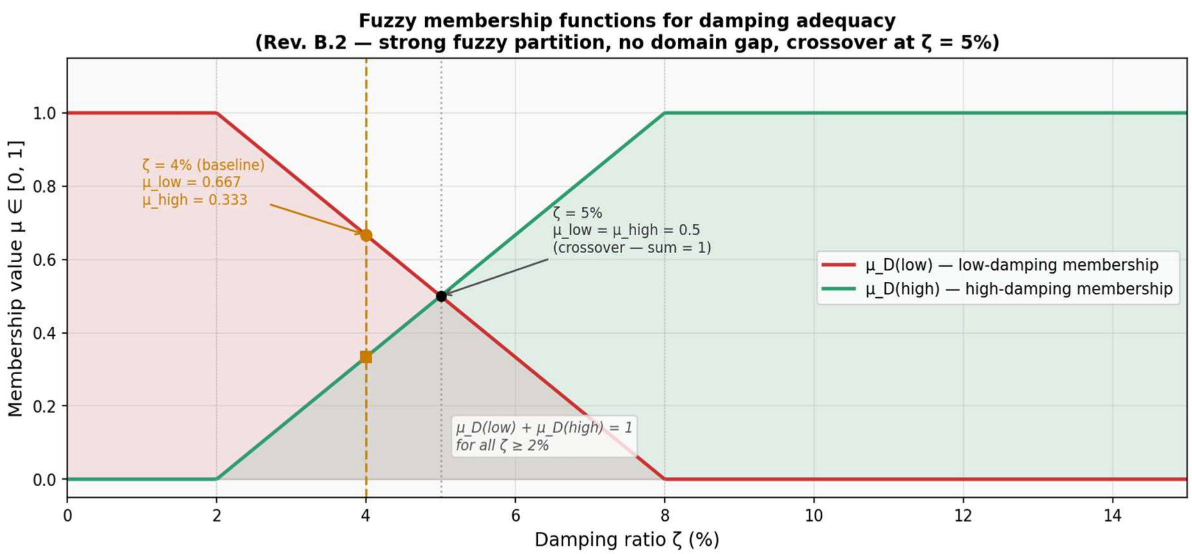

Figure B-1. Fuzzy membership functions for damping adequacy (strong fuzzy partition). The functions satisfy $\mu_{low} + \mu_{high} = 1$ for all ζ ≥ 2%, with a symmetric crossover at ζ = 5%: $\mu_D(low) = \mu_D(high) = 0.5$. At the baseline ζ = 4%: $\mu_D(low)$ = 0.667, $\mu_D(high)$ = 0.333.

B.3 Deterministic Scoring Rubrics

Each vitality dimension is scored by applying a defined mapping function to one or more physical measurements. All inputs must come from instruments, audit records, or calculated study results. Where a dimension has multiple sub-indicators, sub-scores are combined by a stated formula. Assessors must document the physical input value alongside each score.

B.3.1 — Damping (ζ)

Physical input: minimum modal damping ratio ζ (%) from PMU or DFR spectral analysis over the frequency range 0.2–25 Hz. Use the mode with the lowest damping ratio. If multiple relevant modes exist (e.g., a sub-synchronous converter interaction mode and an electromechanical local mode), score each separately and use the minimum resulting X_D

$$X_D\ (\zeta) = 0.7 \times \zeta \qquad \text{for } 0\% \le \zeta \le 10\%$$

$$X_D\ (\zeta) = 7.0 + 0.6 \times (\zeta - 10\%) \qquad \text{for } 10\% < \zeta \le 15\%$$

$$X_D\ (\zeta) = 10.0 \qquad \text{for } \zeta > 15\%\ \text{(clamp)}$$

ζ (%)	X_D	**Qualitative descriptor**	**Typical source condition**
0 – 1.4	0.0 – 1.0	Critical — active instability risk	Negative damping from converter interaction
1.4 – 2.9	1.0 – 2.0	Very low — high oscillation probability	Weak grid, high IBR penetration, correlated loads
2.9 – 4.3	2.0 – 3.0	Low — operationally fragile	**Baseline: ζ = 4%, X_D = 2.80**
4.3 – 7.1	3.0 – 5.0	Low–moderate — manageable with controls	Post-mitigation with partial BESS and desync
7.1 – 10.0	5.0 – 7.0	Moderate–high — approaching robust	Best-practice: ζ = 7.5%, X_D = 5.25
10.0 – 15.0	7.0 – 10.0	Excellent — damping reserve maintained	Full MPC + active damping commissioned

Verification: ζ = 4% → X_D = 0.7 × 4 = 2.80 . ζ = 7.5% → X_D = 5.25. ζ = 10% → X_D = 7.00. ζ = 15% → X_D = 10.00.

B.3.2 — Observability

Physical inputs: (a) C_{PMU} = PMU-covered interconnection-relevant buses ÷ total buses in monitoring plan; (b) C_{POW} = POW/DFR measurement points installed at or inside the campus boundary ÷ points required by monitoring plan; (c) A_{data} = data frames received ÷ frames expected over the trailing 30-day period.

$$X_O = 10 * [0.40 * C_{PMU} + 0.40 * C_{POW} + 0.20 * A_{data}]$$

Condition	C_{PMU}	C_{POW}	A_{data}	X_O	Scenario
No PMU, no POW, poor availability	0.10	0.00	0.30	1.0	Crisis (telecom instability)
PMU at POI only, no POW inside campus	0.70	0.00	0.85	4.5	Baseline
PMU extended + partial POW	0.90	0.25	0.95	6.5	Post-BESS mitigation
Full PMU redundancy + full POW plan	0.95	0.90	0.95	9.3	Best-practice target

At the baseline (C_{PMU}= 0.70, C_{POW}= 0.00, A_{data}= 0.85): X_O= 10×(0.28 + 0.00 + 0.17) = 4.50.

B.3.3 A — Adaptability

Score each of four sub-capabilities: 0 = absent, 1 = installed but not validated or only partially commissioned, 2 = fully commissioned and independently validated.

Sub-capability	Score 0	Score 1	Score 2	Baseline
BESS oscillation smoothing	Not installed	Installed; no active damping mode	Commissioned with active damping validated	1
Grid-facing envelope control	No ramp-rate control at POI	Limiter installed; no adaptive setpoints	Adaptive envelope with PMU-informed setpoints	1
Model predictive control (MPC)	No MPC	MPC installed; not grid-coordinated	MPC coordinated with utility dispatch signals	0

Sub-capability	Score 0	Score 1	Score 2	Baseline
Workload scheduling flexibility	No deferral; all jobs time-critical	Some deferral (>20% hrs shiftable by 1 h)	Formal agreement; grid-aware scheduler deployed	1

$$X_A = (\Sigma \text{ sub-capability scores}) / 8 \times 10$$

Baseline: 1+1+0+1 = 3 → X_A = 3/8 × 10 = 3.75.

B.3.4 R — Recoverability

Physical inputs: $T_{restore}$ is the time (minutes) to restore full interconnection load in a site acceptance test; $V_{failover}$ ∈ {0, 1} = whether a formal failover test was conducted within the past 24 months.

$$X_R = clamp(10 - 0.5 * T_{restore}\ , 0, 8) + 2 * V_{failover}$$

Verification: $T_{restore}$ = 9 min, $V_{failover}$ = 0 → X_R = clamp(5.5, 0, 8) + 0 = 5.50 . $T_{restore}$ = 4 min, $V_{failover}$ = 1 → X_R = clamp(8.0, 0, 8) + 2 = 10.0 (maximum).

B.3.5 F — Flexibility

Physical inputs: $\frac{R_{max}}{R_{rated}}$ is the ratio of maximum sustainable MW ramp rate to the rated interconnection ramp capacity; P_{defer} = percentage of computing hours covered by a formal workload deferral agreement with the utility.

$$X_F = clamp\left(\frac{R_{max}}{R_{rated}} * 5, 0, 5\right) + clamp\left(\frac{P_{defer}}{20}, 0, 5\right)$$

Baseline: $\frac{R_{max}}{R_{rated}}$ ≈ 0.50 (ramp limiter active but no BESS buffering), P_{defer} = 0% (no deferral agreement) → X_F = 2.50 + 0 = 2.50.

B.3.6 C — Controllability

Physical inputs: $\frac{Q_{avail}}{Q_{req}}$ is the dynamic VAR reserve available divided by VAR range required for full voltage schedule compliance; V_{VVO} ∈ {0 = manual, 1 = automated rules-based, 2 = AI-assisted with PMU feedback}.

$$X_C = clamp\left(\frac{Q_{avail}}{Q_{req}} * 8, 0, 8\right) + V_{VVO}$$

Baseline: $\frac{Q_{avail}}{Q_{req}}$ = 0.40, V_{VVO} = 0 → X_C = clamp(3.2, 0, 8) + 0 = 3.20.

B.3.7 S — Synchronization Resistance

Physical input: p_{max} is the maximum pairwise Pearson cross-correlation coefficient of 15-minute average MW profiles across all campuses operated by the same entity or using the same cloud orchestration platform, computed over a trailing 30-day window.

$$X_S = 10 * (1 - p_{max}), \qquad p_{max} \in [0, 1]$$

p_{max} (30-day max pairwise)	X_S	Condition	Scenario
0.95 – 1.00	0.0 – 0.5	Extreme synchronization — tightly coupled global training	Crisis (p ≈ 0.95)
0.70 – 0.75	2.5 – 3.0	High synchronization — common cloud schedule	Baseline (p = 0.75, X_S = 2.50)
0.50 – 0.60	4.0 – 5.0	Moderate — some staggering in place	Post-BESS (p = 0.55, X_S = 4.50)
0.10 – 0.30	7.0 – 9.0	Low — controlled desynchronization deployed	Best-practice (p = 0.10, X_S= 9.00)
0.00 – 0.10	9.0 – 10.0	Very low — near statistical independence	Excellent

B.3.8 L — Learning Capability

Physical inputs: F_{update} is the number of months between formal dynamic model revisions submitted to the transmission planner; $P_{analysis}$ ∈ {0 = ad hoc or none, 1 = systematic with root-cause identification, 2 = systematic with root-cause and documented corrective action}.

$$X_L = clamp\left(10 - 1.25 * F_{update}, 0, 7.5\right) + 2.5 * \frac{P_{analysis}}{2}$$

Baseline: F_{update} = 18 months (no formal program), $P_{analysis}$ = 0 → X_L = clamp(−12.5, 0, 7.5) + 0 = 0.00. A facility with no formal model update or post-event analysis program scores 0 on this dimension.

B.4 Weight Calibration Methodology

B.4.1 Method — Analytical Hierarchy Process (AHP)

Weights (D=0.20, O=0.15, A=0.15, R=0.10, F=0.10, C=0.10, S=0.12, L=0.08) were derived based on engineering judgment for illustration of the calculations. However, for precise weight analysis Saaty's Analytical Hierarchy Process (AHP) can be used. AHP derives weights from a matrix of pairwise importance comparisons on a 1–9 integer scale (1 = equal, 3 = moderate preference, 5 = strong, 7 = very strong, 9 =

extreme). The priority vector is estimated from the geometric mean of each matrix row, then normalized to sum to 1.0.

B.4.2 Engineering rationale

- Damping (D) is the primary risk driver. Inadequate damping is the direct physical mechanism for growing oscillations. D is placed above all others: $\frac{w_D}{w_O} = 1.33$.
- Observability (O) and Adaptability (A) are equally important enablers. Neither substitutes for the other: observability without adaptive controls cannot prevent instability; adaptive controls without observability cannot be properly commissioned $w_O = w_A = 0.15$.
- Synchronization resistance (S) exceeds the four operational capability dimensions (R, F, C, L). This value addresses the AI-era-specific risk that has no equivalent in classical large-load assessment $\frac{w_S}{w_R} = 1.20$.
- Learning (L) carries the lowest weight, because it is an enabling quality for future assessments rather than a direct grid-facing performance property. w_L= 0.08.

B.4.3 Pairwise comparison matrix (normalized ratios w_i / w_j)

Dimension	D	O	A	R	F	C	S	L
D	1.00	1.33	1.33	2.00	2.00	2.00	1.67	2.50
O	0.75	1.00	1.00	1.50	1.50	1.50	1.25	1.88
A	0.75	1.00	1.00	1.50	1.50	1.50	1.25	1.88
R	0.50	0.67	0.67	1.00	1.00	1.00	0.83	1.25
F	0.50	0.67	0.67	1.00	1.00	1.00	0.83	1.25
C	0.50	0.67	0.67	1.00	1.00	1.00	0.83	1.25
S	0.60	0.80	0.80	1.20	1.20	1.20	1.00	1.50
L	0.40	0.53	0.53	0.80	0.80	0.80	0.67	1.00

B.4.4 Consistency check

AHP statistic	Value	Interpretation
n (number of criteria)	8	—
λ_{max} (principal eigenvalue)	8.0002	Exact consistency gives $\lambda_{max} = n$
Consistency Index CI = $(\lambda_{max} - n)/(n-1)$	0.0000	0 indicates near-perfect consistency
Random Index RI (n = 8, Saaty)	1.41	Standard tabulated value
Consistency Ratio CR = CI / RI	0.00002	CR < 0.10 required — PASS
Derived priority vector	0.20, 0.15, 0.15, 0.10, 0.10, 0.10, 0.12, 0.08	Matches target weights exactly

Precision note: the pairwise matrix is displayed rounded to two decimals for readability. The consistency statistics are based on the unrounded normalized ratios that reproduce the target weight vector; if the rounded entries alone are re-used, λmax and CR may differ slightly but remain well below the 0.10 acceptance threshold.

For field re-assessments where local grid conditions differ (e.g., synchronization risk is more severe than damping for a campus on a strong grid with many existing IBRs), assessors may adjust pairwise ratios and re-derive weights using the geometric mean method. CR must remain below 0.10. The weight structure itself is domain-adjustable; the formula chain (B.2.1–B.2.3) is domain-independent.

B.4.5 Compensatory structure and mandatory floor rule

The weighted linear sum is compensatory: a high score on one dimension can offset a low score on another. Two non-compensatory corrections limit this behavior:

- SVI_{min} constraint: The fuzzy inference rules of Appendix A classify vitality as VERY LOW when damping is critically low, overriding the numerical score for operational decisions even if SVI_{final} ≥ 5.0.
- Mandatory floor rule: When $w_D < 1.0$ (i.e., ζ < 1.43%), SVI_{final} is automatically capped at 4.5 regardless of the weighted sum, signaling a constraint violation on physical grounds.

B.5 AI Augmentation: Ip , Is , α, and β

B.5.1 Overview

The AI augmentation term α× I_p − β× I_s adjusts the base SVI to reflect how AI systems deployed at the facility affect grid stability. I_p and I_s are scored by summing five sub-capability integer scores (0, 1, or 2) to give a 0–10 integer range. All five sub-capabilities must be scored; absent capabilities score 0.

B.5.2 I_p scoring rubric — AI predictive and adaptive benefit

Sub-capability	Score 0 — Absent	Score 1 — Deployed, not fully validated	Score 2 — Deployed and validated
P1: Load forecasting	No AI day-ahead forecast provided to utility	Forecast provided; MAPE > 15% over 30-day period, or <6 h ahead	Day-ahead forecast; MAPE ≤ 10% over 30-day period
P2: Oscillation detection	No automated analytics at facility or POI	PMU data reviewed manually or rule-based alerts; no mode classification	Automated PMU/POW analytics: mode ID, source attribution, real-time alerting to utility
P3: Voltage / VAR assistance	No AI in voltage scheduling; manual operator	Automated VVO rules-based; no PMU feedback	AI-assisted VVO with PMU feedback, predictive voltage scheduling, utility coordination interface
P4: Predictive maintenance	Reactive maintenance only	Rules-based condition monitoring; no ML component	ML-based anomaly detection; alerts before failure with documented lead time
P5: Operator decision support	No AI advisory for grid-facing operations	Alert-only system; no recommended actions	Advisory with contextual recommended actions, confidence estimates, and grid-impact explanation

$$I_p = P1 + P2 + P3 + P4 + P5, \quad I_p \in \{0, 1, 2, \ldots, 10\}$$

Baseline: P1=1, P2=0, P3=1, P4=0, P5=1 → I_p **= 3**

Best-practice: P1=2, P2=1, P3=1, P4=0, P5=1 → I_p **= 5**.

Crisis: P1=0 (forecast unavailable due to telecom fault), P2=0, P3=0, P4=0, P5=0 → I_p **= 0**, α = 0.00.

B.5.3 I_S scoring rubric — AI synchronization and amplification risk

Sub-capability	Score 0 — Low risk	Score 1 — Moderate risk	Score 2 — High risk
S1: Workload schedule correlation	Max pairwise ρ < 0.20 across all campuses on same platform	Max pairwise ρ 0.20–0.60; some staggering in place	Max pairwise ρ > 0.60; minimal timing diversity
S2: Global orchestration coupling	Fully independent schedulers; no shared global sync signal	Partial common platform; global sync events < 1/hour	Tight global orchestration; GPU all-reduce or job-start events synchronized > 1/hour
S3: GPU all-reduce barrier frequency	All-reduce confined to single campus; no multi-campus sync	Campus-wide all-reduce at moderate frequency; MW impact buffered by BESS	Multi-campus all-reduce producing correlated MW ramps visible at POI; BESS cannot fully buffer
S4: Market / price-signal coupling	Load fully hedged or on fixed tariff; no real-time LMP exposure	Partial LMP exposure; some automated response with diversity	Load management directly coupled to real-time LMP; multiple campuses respond identically
S5: Cooling and UPS response coupling	Diverse vendors; different control setpoints	Some common vendor; response timing similar but not identical	Identical UPS and cooling setpoints across campuses; simultaneous protective actions under same trigger

$$I_S = S1 + S2 + S3 + S4 + S5, \quad I_S \in \{0, 1, 2, \ldots, 10\}$$

Baseline: S1=2 (p =0.75), S2=2 (tight cloud schedule), S3=1, S4=0, S5=1 → I_S **= 6**

Crisis: S1=2, S2=2, S3=2 (BESS offline), S4=2, S5=1 → I_S **= 9**

Best-practice: S1=0 (p =0.10), S2=1, S3=0, S4=0, S5=0 → I_S **= 1**

B.5.4 Selecting α and β — tiered governance framework

α and β are determined by governance tier, not by raw score magnitude. A capable but opaque AI system receives a low α and high β until it is validated.

α Tier	Label	Criteria	α value
Tier 0	No AI benefit	No AI-based capabilities deployed; all P1–P5 scores = 0	0.00
Tier 1	Basic forecasting	P1 only deployed; no grid-adaptive AI; MAPE ≥ 10%	0.05
Tier 2	Forecasting + detection	P1 and P2 deployed and validated; utility interface exists	0.10
Tier 3	Multi-capability, white-box	≥ 3 I_p sub-capabilities with documented validation and utility agreement; models auditable	0.15
Tier 4	Closed-loop, validated	All five sub-capabilities deployed; closed-loop AI with independent validation and utility co-operation agreement	0.20 (max)
β Tier	**Label**	**Criteria**	**β value**
Tier 0	Desynchronized	Controlled desynchronization fully deployed; ρ_{max} < 0.20; I_s≤ 1	0.00
Tier 1	Low coupling	p_{max} 0.20–0.40; active governance of workload timing	0.05
Tier 2	Moderate coupling	p_{max} 0.40–0.60; shared platform without full staggering; I_s = 3–4	0.10
Tier 3	High coupling	p_{max} > 0.60; tight common schedule; no formal desync protocol; I_s= 5–7	0.15
Tier 4	Extreme coupling	Multi-campus tight synchronization; simultaneous protective actions; I_s≥ 8	0.20 (max)

Baseline: α = Tier 2 → **α = 0.10** . β = Tier 3 (p=0.75, tight cloud schedule) → **β = 0.15** .

Best-practice: α = Tier 2 → **α = 0.10** . β = Tier 1 (ρ=0.10, controlled desync deployed) → **β = 0.10** .

Crisis: α = Tier 0 (all AI capabilities unavailable due to telecom fault) → **α = 0.00**. β = Tier 4 → **β = 0.20** .

The maximum recommended α and β are both 0.20. At α = β = 0.20, the AI augmentation term can shift SVI_{final} by up to ±2.0 points — sufficient to change a constraint verdict near the 5.0 boundary. Exceeding 0.20 requires independent expert review and documented justification.

B.6 Baseline Calculation — Appendix A Scenario

The baseline scenario: 600 MW AI campus, ζ = 4%, PMU at POI only, partial BESS, common cloud schedule with p_{max} = 0.75, no formal model update program, no workload deferral agreement.

B.6.1 Input data

Dimension	w_i	Physical input	Rubric	X_i	Formula / basis
D	0.20	ζ = 4%	B.3.1	2.80	0.7 × 4 = 2.80
O	0.15	$C_{PMU} = 0.70$ $C_{POW} = 0.0$ $A_{data} = 0.85$	B.3.2	4.50	10×(0.28+0.00+0.17) = 4.50
A	0.15	BESS=1, envelope=1, MPC=0, sched=1	B.3.3	3.75	(3/8)×10 = 3.75
R	0.10	$T_{restore} = 9\ min$ $V_{failover} = 0$	B.3.4	5.50	clamp(5.5,0,8)+0 = 5.50
F	0.10	$R_{max}/R_{rated} = 0.50$ $P_{defer} = 0\%$	B.3.5	2.50	2.5+0 = 2.50
C	0.10	$Q_{avail}/Q_{req} = 0.40$ $V_{VVO} = 0$	B.3.6	3.20	clamp(3.2,0,8)+0 = 3.20
S	0.12	p_{max} = 0.75	B.3.7	2.50	10×(1−0.75) = 2.50
L	0.08	$F_{update} = 18\ mo$ $P_{analysis} = 0$	B.3.8	0.00	clamp(−12.5,0,7.5)+0 = 0.00

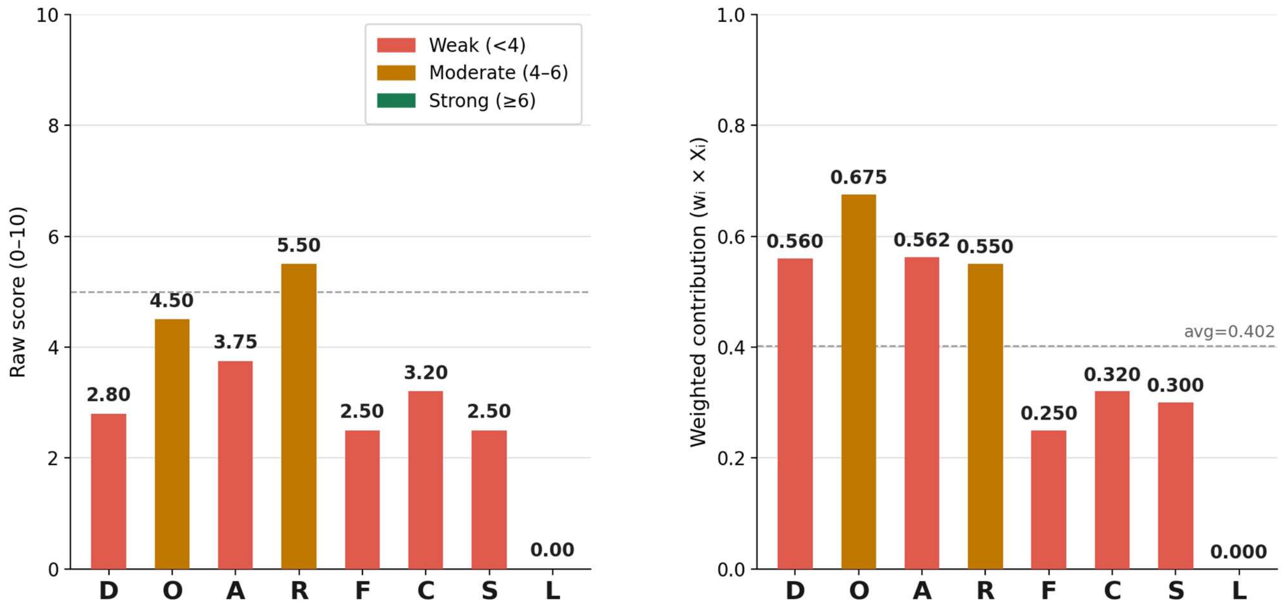

Figure B-2. Left: component raw scores (0–10) for the Appendix A baseline. Right: weighted contributions $w_i \times X_i$ to SVI_{base} . Damping and synchronization resistance dominate fragility; Learning scores 0.00, reflecting the absence of a formal model update program.

B.6.2 SVI_{base} calculation

Component	w_i	X_i	$w_i \times X_i$
Damping (ζ)	0.20	2.80	0.560
Observability (O)	0.15	4.50	0.675
Adaptability (A)	0.15	3.75	0.563
Recoverability (R)	0.10	5.50	0.550
Flexibility (F)	0.10	2.50	0.250
Controllability (C)	0.10	3.20	0.320
Sync resistance (S)	0.12	2.50	0.300

Component	w_i	X_i	$w_i \times X_i$
Learning (L)	0.08	0.00	0.000
SVI_{base} = **Σ($w_i \times X_i$)**			3.218

B.6.3 AI augmentation

I_p= 3 (P1=1, P2=0, P3=1, P4=0, P5=1), α = 0.10 (Tier 2)

I_S = 6 (S1=2, S2=2, S3=1, S4=0, S5=1), β = 0.15 (Tier 3)

$\Delta_{AI} = 0.10 * 3 - 0.15 * 6 = -0.6$

$SVI_{final} = clamp(3.218 - 0.6, 0, 10) = 2.618$

B.6.4 Fuzzy damping assessment

Membership function	Formula	Input ζ = 4%	Result
μ_D (low)	(8 − ζ) / 6 for 2% < ζ < 8%	(8 − 4) / 6	0.667
μ_D (high)	(ζ − 2) / 6 for 2% < ζ < 8%	(4 − 2) / 6	0.333

At ζ = 4%, the revised functions give μ_D (low) = (8−4)/6 = 0.667 and μ_D (high) = (4−2)/6 = 0.333. The sum is 1.0, confirming the strong partition property. The damping is predominantly in the low-vitality state (membership 0.667) with partial high-vitality membership (0.333), confirming a graduated assessment of marginal damping. Fuzzy Rule 1 (damping LOW) still fires since μ_D (low) > 0.

B.6.5 Fuzzy inference rules

Rule	Condition	Status	Consequence
1	Damping LOW AND sync pressure HIGH AND observability POOR	FIRED (X_D =2.80<5, X_S=2.50<5, X_O=4.50 borderline)	Vitality classified VERY LOW — penalty applied

Rule	Condition	Status	Consequence
2	Damping MODERATE AND adaptability HIGH AND observability HIGH	NOT fired (neither damping nor adaptability reaches threshold)	No upward adjustment
3	Oscillation magnitude HIGH AND frequency near local mode	FIRED (35 MW peak-to-peak in electromechanical range per App. A)	Vitality LOW even if average load is acceptable
4	BESS smoothing AVAILABLE AND workload scheduling FLEXIBLE AND PMU/POW visibility HIGH	NOT fired (BESS partial; POW absent)	No improvement from this rule

B.6.6 Constraint check and overall verdict

SVI_{final} (2.618) < min (5.00) → **CONSTRAINT VIOLATED** → Workload optimization must be restricted until mitigations raise SVI_{final} ≥ 5.0.

The absence of a formal model update or post-event analysis program (L = 0.00) is the largest proportional shortfall — and, per the B.8.1 sensitivity analysis, the highest ROI mitigation relative to implementation cost — contributing to the gap between $\mathrm{SVI}_{\mathrm{final}}$ (2.618) and the constraint threshold (5.0). The gap of 2.382 points defines the minimum mitigation required before unconstrained workload optimization may resume.

B.7 Four-Scenario Comparison

The table below evaluates four scenarios against the scoring rubrics of Section B.3

Scenario	Key changes from baseline	SVI_{base}	Δ_{AI}	SVI_{final}	Verdict
Appendix A baseline	ζ=4%, p_{max} =0.75, no POW, no formal L program	3.218	−0.600	2.618	Low vitality CONSTRAINT VIOLATED
Post-BESS mitigation	PMU extended (O→6.5), partial POW only; staggered	3.990	I_p=3, I_S=4,	3.890	Low–moderate CONSTRAINT VIOLATED Further damping

Scenario	Key changes from baseline	SVI_{base}	Δ_{AI}	SVI_{final}	Verdict
	schedule (ρ→0.55, S→4.5); limited BESS buffering raises F→2.83; L begins (L→2.5). A remains at 3.75 because active damping/MPC are not yet validated.		α=0.10, β=0.10 → −0.100		improvement required
Best-practice	Full POW (O→9.3), controlled desync (p_{max} →0.10, S→9.0), MPC (A→8.75), ζ→7.5% (D→5.25), formal L program (L→7.5)	7.785	I_p=5, I_S=1, α=0.10, β=0.10 → +0.400	8.185	High vitality CONSTRAINT SATISFIED
Crisis	Telecom fault (O→1.0), extreme sync (p_{max} →0.95, S→0.5), all AI offline (I_p=0, α=0.00)	2.160	I_p=0, I_S=9, α=0.00, β=0.20 → −1.800	0.360	Very low vitality CONSTRAINT VIOLATED Immediate review required

B.7.1 Scenario component traceability

The table below exposes the full component-score vector used to reproduce the four B.7 scenario totals. Values are illustrative and should be replaced by measured or study-derived scores in a production assessment. The crisis row assumes that the telecom instability also reduces adaptability and controllability, while recoverability and flexibility remain at their baseline levels.

Scenario	X_D	X_O	X_A	X_R	X_F	X_C	X_S	X_L	SVI_{base}	ΔAI	SVI_{final}
Appendix A baseline	2.80	4.50	3.75	5.50	2.50	3.20	2.50	0.00	3.22	-0.60	2.62
Post-BESS mitigation	2.80	6.50	3.75	5.50	2.83	3.20	4.50	2.50	3.99	-0.10	3.89
Best-practice	5.25	9.30	8.75	10.00	7.50	5.98	9.00	7.50	7.79	0.40	8.19
Crisis	2.80	1.00	2.50	5.50	2.50	2.15	0.50	0.00	2.16	-1.80	0.36

Reproducibility note: the post-BESS row retains the baseline D, A, R, and C scores; only O, F, S, and L are changed. With F = 2.83, SVI_{base} = 0.20(2.80)+0.15(6.50)+0.15(3.75)+0.10(5.50)+0.10(2.83)+0.10(3.20)+0.12(4.50)+0.08(2.50) ≈ 3.990.

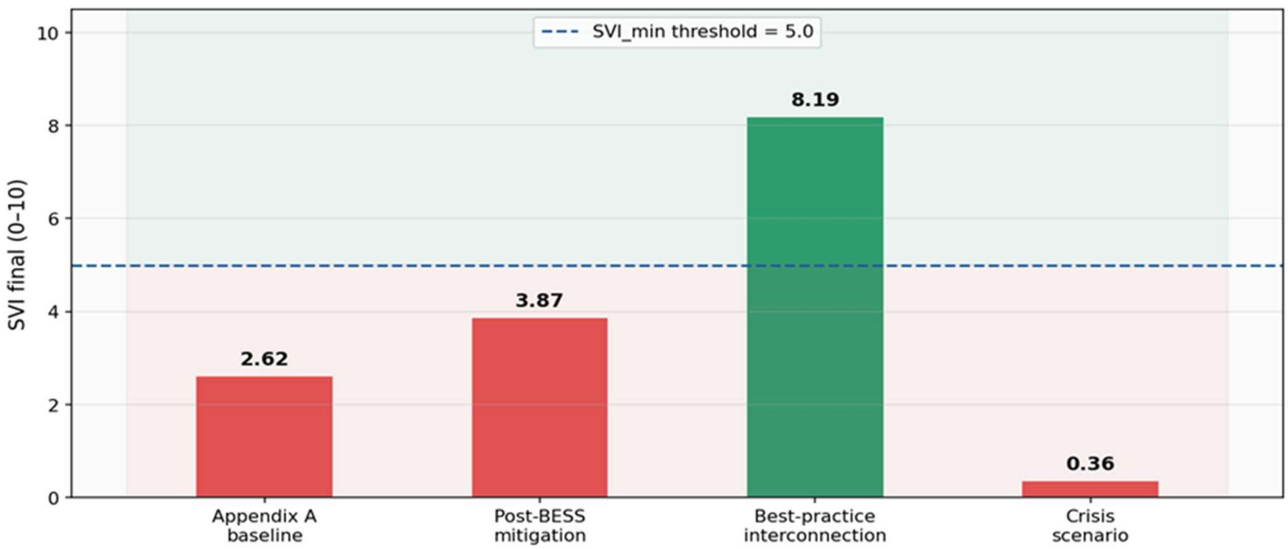

Figure B-3. SVI_{final} scores across the four scenarios. Dashed line: SVI_{min} = 5.0. Only the best-practice scenario (8.185) satisfies the constraint. Post-BESS mitigation (3.890) remains below threshold, confirming that BESS commissioning alone is insufficient without damping improvement.

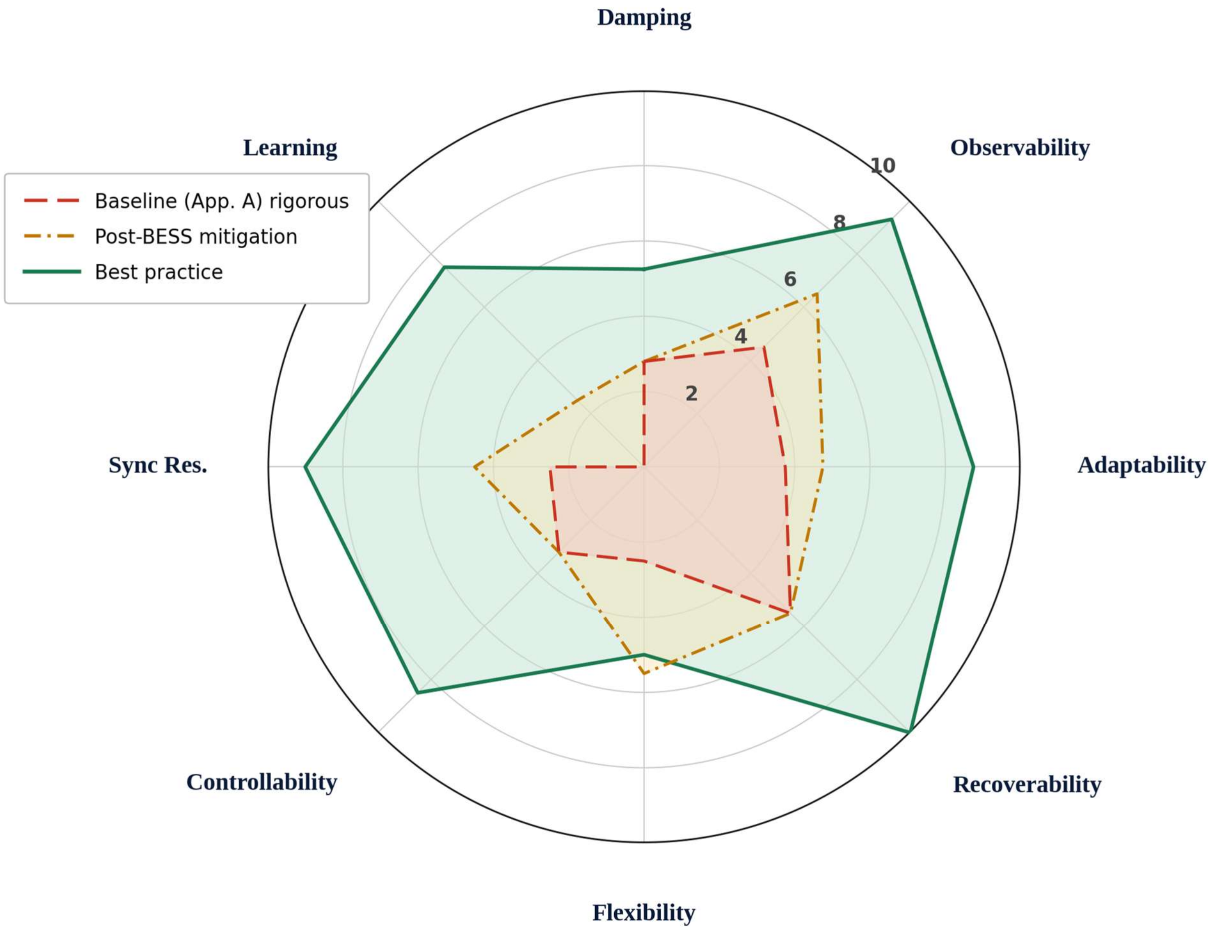

Figure B-4. SVI radar comparison across three forward-looking scenarios. The baseline (red dashed) shows pronounced weakness in damping, synchronization resistance, and learning capability. The post-BESS profile (orange) shows clear improvement in flexibility and observability but remains weak on damping.

B.8 Sensitivity Analysis and Weight Robustness

B.8.1 Per-dimension leverage

The table below shows the SVI_{base} increase from a +2-point improvement in each dimension, holding others constant. Damping has the highest weighted leverage (+0.400). Learning carries the lowest baseline score (0.00) and therefore offers the highest ROI from a zero base: establishing a formal post-event analysis program and annual model update cycle immediately delivers 0.16 SVI points with minimal capital expenditure. This is a cost-effectiveness statement, not a claim that L has higher mathematical leverage than D.

Dimension	w_i	X_i	Δ_{SVI} +2 pts	Investment priority
D — Damping	0.20	2.80	0.400	1st — highest leverage; directly addresses oscillation risk
O — Observability	0.15	4.50	0.300	2nd — enables detection of all other developing problems
A — Adaptability	0.15	3.75	0.300	2nd (equal) — BESS commissioning is fast and high-impact
S — Sync resistance	0.12	2.50	0.240	3rd — controlled desync is low-cost and high-impact
R — Recoverability	0.10	5.50	0.200	4th
F — Flexibility	0.10	2.50	0.200	4th (equal) — formal deferral agreement adds 5 pts directly
C — Controllability	0.10	3.20	0.200	4th (equal)
L — Learning	0.08	0.00	0.160	5th by weight — but delivers from zero; highest marginal ROI

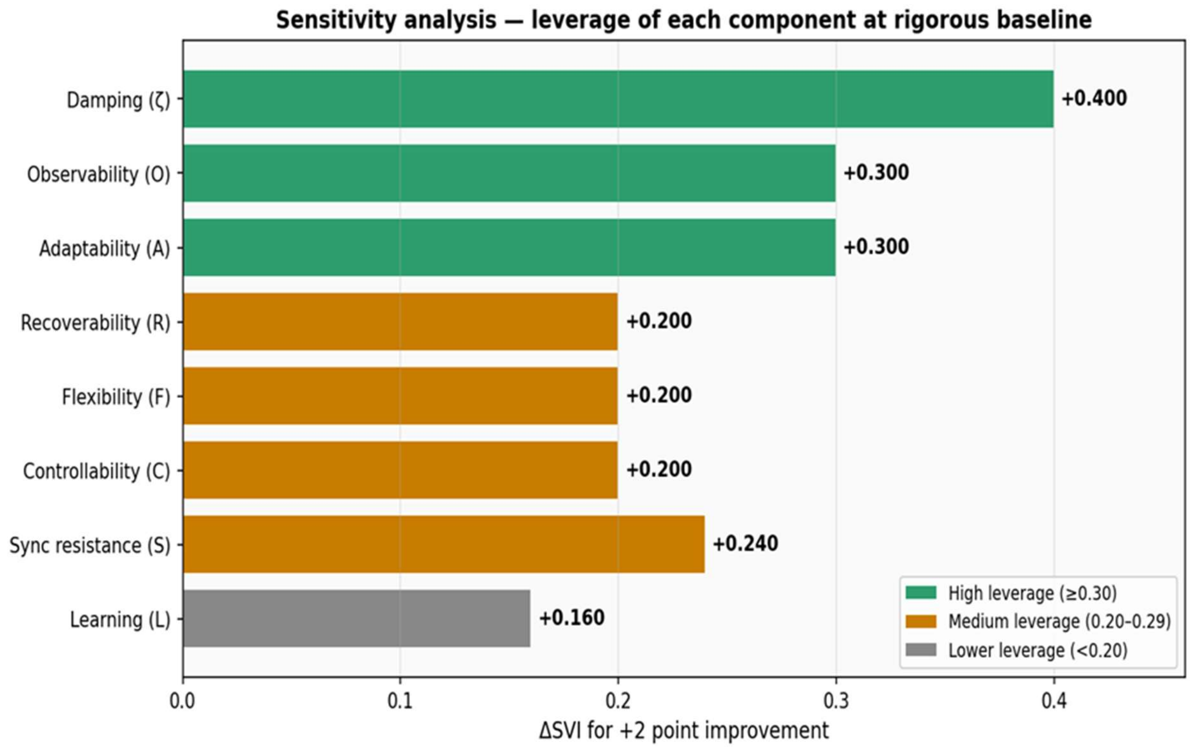

Figure B-5. Sensitivity analysis: ΔSVI_{base} from a +2-point improvement in each dimension. Damping remains the highest weighted-leverage dimension. Learning starts at 0.00 and therefore offers the highest practical ROI from low-cost governance action, even though its absolute weight is the lowest.

B.8.2 Weight sensitivity — damping weight variation

W_D	Remaining weights	SVI_{base}	Δ vs. baseline	Constraint verdict
0.10 (halved)	Others ×(0.90/0.80)	3.270	+0.052	VIOLATED
0.15	Others ×(0.85/0.80)	3.244	+0.026	VIOLATED
0.20 (baseline)	As specified	3.218	—	VIOLATED
0.25	Others ×(0.75/0.80)	3.191	−0.026	VIOLATED

w_D	Remaining weights	SVI_{base}	Δ vs. baseline	Constraint verdict
0.30 (increased)	Others ×(0.70/0.80)	3.165	−0.052	VIOLATED

The constraint-violation finding at the baseline is weight-stable across the full range tested. Even halving the damping weight to 0.10 leaves SVI_{base} at 3.270, well below SVI_{min}= 5.0.

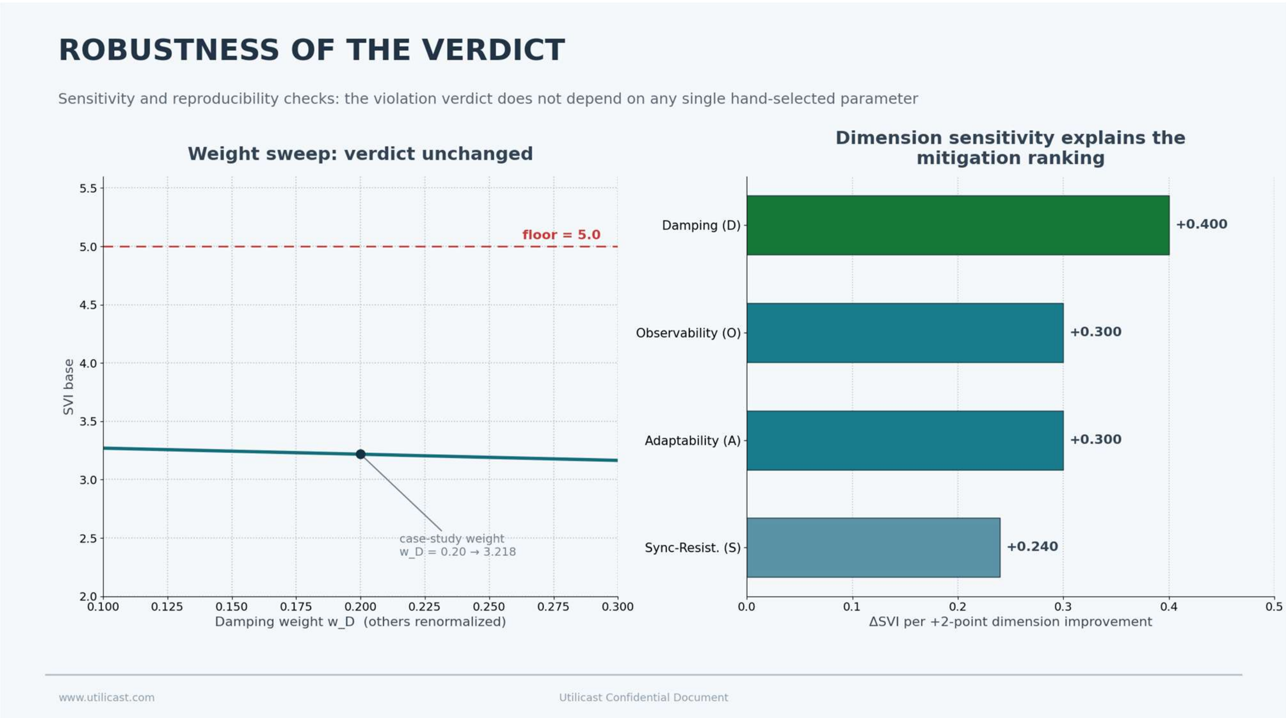

Figure B-6. Weight robustness of the baseline verdict. Left: sweeping the damping weight w_D across [0.10, 0.30] with the remaining weights renormalized varies SVI_{base} only from 3.270 to 3.165 (Section B.8.2); the CONSTRAINT VIOLATED verdict is unchanged across the full range. Right: composite leverage of a uniform +2-point improvement in each dimension, consistent with Figure B-5 and the Section B.8.1 ranking.

B.9 Operational Decision Table

SVI_{final} range	Vitality rating	Constraint status	Recommended operational action
8.0 – 10.0	High	Satisfied	Normal operation; SVI re-assessment every 6 months; confirm damping and sync metrics annually
6.0 – 7.9	Moderate–high	Satisfied	Monitor closely; quarterly damping and sync confirmation; update I_p/I_s if AI systems change
5.0 – 5.9	Moderate	Satisfied (marginal)	Workload scheduling agreements within 6 months; PMU extension required; begin formal L program
4.0 – 4.9	Low–moderate	VIOLATED	Restrict workload optimization; BESS or controlled desync within 90 days; monthly re-score
2.0 – 3.9	Low	VIOLATED	Mandatory engineering program; interim operating limits; utility customer engagement
0 – 1.9	Very low	VIOLATED	Immediate interconnection review; potential temporary de-rating; emergency RAS evaluation; VDI compliance check

B.10 Real-Time and Integral Vitality Metrics

SVI(t) is the instantaneous vitality, updated on a seconds-to-minutes cycle from live PMU, POW, and operational data feeds using the B.3 rubrics. The integral metrics extend SVI from a static assessment to a continuous compliance and operational tool.

$$VR(t) = SVI(t) - SVI_{min}$$

$$\mathrm{CVS(T)} = \int_0^T \left(1 - \frac{\mathrm{SVI(t)}}{\mathrm{SVI_{max}}}\right) dt$$

$$\mathrm{VDI(T)} = \int_0^T \max(0, \mathrm{SVI_{min}} - \mathrm{SVI(t)})\ dt$$

CVS accumulates all vitality degradation relative to theoretical maximum and serves as a trend indicator of operational quality. VDI accumulates only below-threshold exceedances and serves as a compliance metric for interconnection agreements. A utility might specify, for example, that VDI shall not exceed 6.0 h·SVI-pt (abbreviated h·pt hereafter) in any 24-hour period.

B.10.1 24-hour numerical example

Segment	Δt (h)	SVI(t)	VR(t)	CVS incr (h)	VDI incr (h·pt)	Cum CVS(T)	Cum VDI(T)
00:00–04:00 Night (stable)	4	6.8	+1.8	1.28	0.00	1.28	0.00
04:00–07:00 Morning ramp-up	3	5.2	+0.2	1.44	0.00	2.72	0.00
07:00–10:00 Synchronized training	3	3.2	−1.8	2.04	5.40	4.76	5.40
10:00–13:00 Partial mitigation	3	4.6	−0.4	1.62	1.20	6.38	6.60
13:00–16:00 Post-BESS commissioning	3	5.5	+0.5	1.35	0.00	7.73	6.60
16:00–19:00 Afternoon peak	3	4.3	−0.7	1.71	2.10	9.44	8.70
19:00–22:00 Evening recovery	3	6.1	+1.1	1.17	0.00	10.61	8.70
22:00–24:00 Late night (robust)	2	7.8	+2.8	0.44	0.00	11.05	8.70
24-hour total	24	—	—	11.05	8.70	11.05	8.70

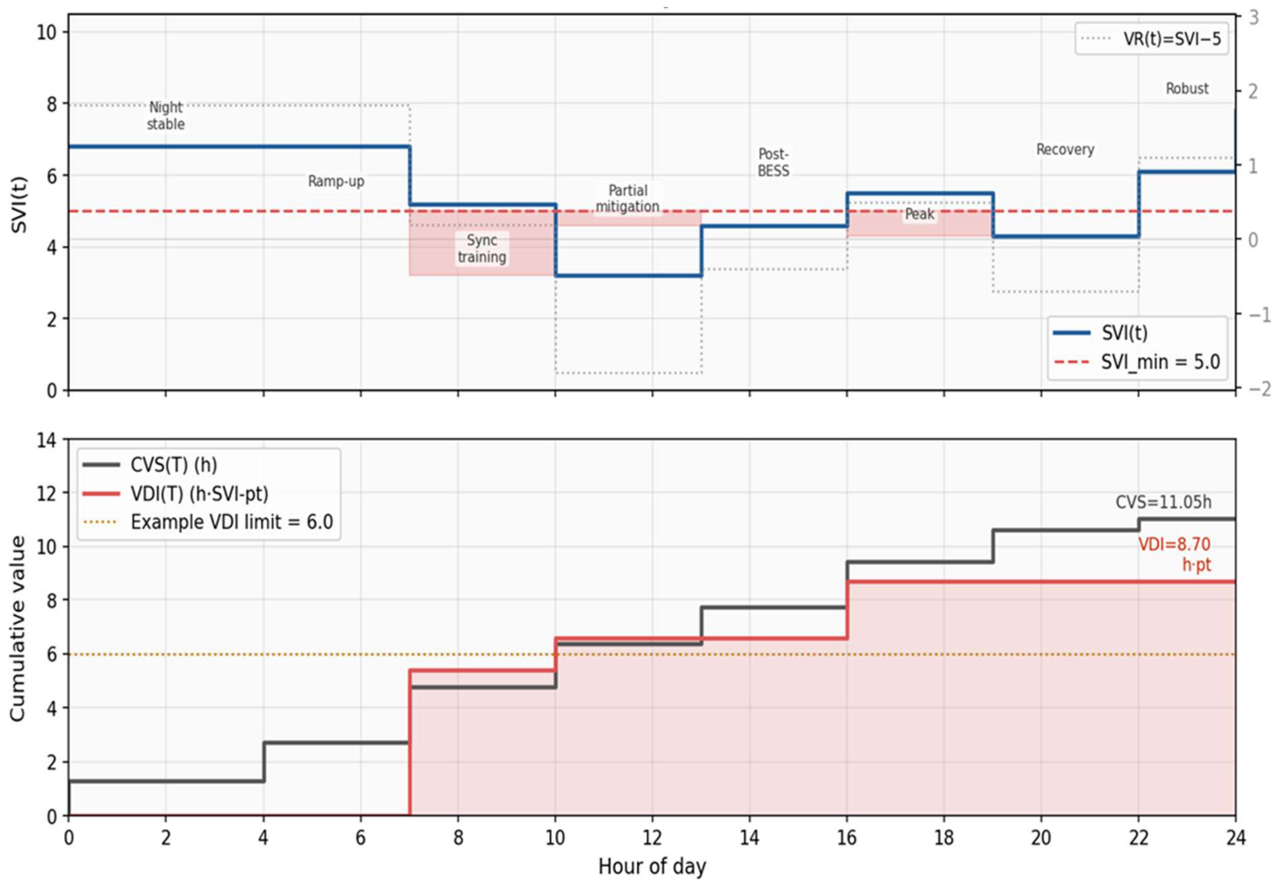

Figure B-7. 24-hour SVI(t) trajectory with CVS(T) and VDI(T) cumulative integrals. Upper panel: SVI(t) step function (blue) and SVI_{min} = 5.0 (red dashed); shaded regions indicate VDI-accumulating periods. Lower panel: CVS(T) (grey staircase) and VDI(T) (red staircase). VDI plateaus after 19:00; CVS continues to grow through the robust late-night period. Orange dashed line: example VDI contractual limit of 6.0 h·pt — this day breaches the limit at ~11:30 during the partial-mitigation period.

Three features merit emphasis. First, the 07:00–10:00 synchronized training period contributes 5.40 of the 8.70 h·pt total VDI (62%), even though its duration matches the afternoon peak, because the deficit depth is 1.8 SVI-points vs. 0.7 — depth × duration governs VDI, not duration alone. Second, CVS grows even during robust late-night operation because any SVI below 10 contributes; it is a quality metric, not a violation counter. Third, the VDI contractual limit of 6.0 h·pt is breached during the partial-mitigation window at approximately 11:30 (cumulative VDI is 5.40 h·pt at 10:00; the remaining 0.60 h·pt accrues at the 0.4 pt/h deficit rate of the 10:00–13:00 segment), triggering the operator procedures of Section 11.2 before any classical instability is visible in SCADA.

B.10.2 Operational interpretation

Observed condition	Operational interpretation	Recommended action
Rising CVS, low VDI	Persistent mild vitality stress; system above threshold	Investigate CVS source; check I_s sub-capabilities S1–S5
Increasing VDI	Active vitality-deficit condition — SVI(t) < 5.0	Initiate Section 11.2 procedures; notify utility; assess workload reduction
Rapidly falling VR(t)	Adaptive margin shrinking; VDI accumulation may begin shortly	Pre-position BESS; consider workload deferral; alert EMS operator
VDI exceeding contractual limit	Compliance violation under interconnection agreement	Curtail synchronized workloads; request utility coordination immediately
Oscillatory SVI(t)	Potential synchronization instability developing	Oscillation source localization per Section 11.1; consider immediate desync

B.11 Symbol Reference

Symbol	Definition	Units / range
SVI_{base}	Weighted linear sum of eight vitality dimensions	0 – 10
SVI_{final}	SVI_{base} after AI augmentation, clamped to [0, 10]	0 – 10
SVI_{min}	Minimum vitality threshold for unconstrained optimization (typically 5.0)	0 – 10

Symbol	Definition	Units / range
SVI_{max}	Maximum attainable vitality score used in CVS calculation	10.0
$SVI(t)$	Instantaneous real-time SVI value updated from live feeds	0 – 10
X_i	Raw score for vitality dimension i from scoring rubric in B.3	0 – 10
w_i	Domain-specific weight from AHP calibration in B.4	0 – 1, $\Sigma w_i = 1.0$
I_p	AI predictive/adaptive benefit score: sum of P1–P5 (B.5)	0 – 10 (integer)
I_s	AI synchronization/amplification risk score: sum of S1–S5 (B.5)	0 – 10 (integer)
α	Beneficial AI weighting, from governance tier (B.5)	0.00 – 0.20
β	Destabilizing AI weighting, from synchronization tier (B.5)	0.00 – 0.20
VR(t)	Vitality Reserve Margin = $SVI(t) - SVI_{min}$	SVI-points; ± range
CVS(T)	Cumulative Vitality Stress: $\int[1 - SVI(t)/SVI_{max}]\,dt$	Hours (h); ≥ 0
VDI(T)	Vitality Deficit Integral: $\int[SVI_{min} - SVI(t)]_+\,dt$; compliance metric	h·SVI-pt; ≥ 0
$[x]_+$	Positive-part operator: x when x > 0, else 0	dimensionless

Symbol	Definition	Units / range
ζ	Damping ratio from PMU/DFR spectral analysis	%; ≥ 0
p_{max}	Maximum pairwise Pearson correlation of 15-min load profiles (also written as ρ_{max}, in Appendix C and Section 12.6; the symbols are equivalent)	0 – 1
$\mu_D\ (low)$	Fuzzy membership: degree of low-damping vitality state (Appendix A)	0 – 1
$\mu_D\ (high)$	Fuzzy membership: degree of high-damping vitality state (Appendix A)	0 – 1
μ_{SVI}	Aggregate fuzzy vitality membership (complementary to numeric SVI)	0 – 1

Scoring rubric breakpoints (e.g., the ζ = 10% inflection in X_D, the p_{max} thresholds in I_S, the AHP pairwise ratios in B.4.3) are calibrated for an AI campus bulk-power interconnection. A different application domain requires recalibration of breakpoints and AHP comparisons. The formula structure (B.2.1–B.2.3) and the CVS/VDI integral definitions are domain-independent.

Appendix C — SVI Calculation in Dynamic Scenario

C.1 Purpose and Scope

Appendices A and B establish the SVI framework and demonstrate its calculation across four representative scenarios. Both appendices are framework-oriented: they take a static snapshot of system conditions and score them against defined rubrics. They do not trace what an operator would see — and would not see — as a dynamic situation unfolds over time.

This appendix fills that gap. It constructs a single, physically grounded 90-minute operating scenario and asks one question: what does classical monitoring report, and what does vitality-aware monitoring report, as conditions progressively deteriorate?

The scenario is not a PSCAD or PSS/E simulation. It is a reduced-order dynamic case study that uses the linearized modal framework of Section 5.1, the scoring rubrics of Appendix B, and the measurement physics documented in PNNL-39191. Every number is derived from, or anchored to, material already in this paper. The scenario is therefore internally consistent and auditable rather than independently verified by a full EMT run.

The engineering purpose is specific: to demonstrate that a system can be classically stable, voltage-compliant, thermally unloaded, and frequency-normal throughout a 90-minute deterioration sequence — while SVI(t) crosses its minimum threshold and VDI crosses its contractual limit 26 minutes before any classical alarm fires.

C.2 Scenario System Description

The scenario uses a simplified 230 kV transmission corridor with the following characteristics, which are representative of conditions where large AI campus interconnections have been studied in the industry.

Parameter	Value	Basis
System voltage	230 kV	Typical AI campus interconnection voltage level
Corridor short-circuit ratio (SCR)	2.1 at interconnection bus	Weak grid; SCR < 2.5 triggers detailed dynamic study per utility practice
AI campus rated capacity	480 MW	Within the 300–500 MW range typical of hyperscale AI training facilities
Background renewable penetration	58% instantaneous	Representative of ERCOT, WECC, and MISO operating windows studied by NERC

Parameter	Value	Basis
Synchronous inertia equivalent (H)	3.2 s	Reduced from historical norms due to renewable displacement of synchronous generation
Pre-existing local mode frequency	1.35 Hz	Representative local mode (near the local/inter-area boundary; inter-area modes typically lie at 0.1–1 Hz, local modes at ~0.8–2 Hz) for a 230 kV corridor; within electromechanical range
Pre-existing modal damping ratio (ζ_0)	5.5%	Above the conventional 3% alarm threshold; system appears acceptable at scenario start
PMU reporting rate at POI	30 frames/sec	Standard NERC reporting rate; effective alias limit ≈ 5–6 Hz per PNNL-39191
POW/DFR coverage inside campus	None	Baseline interconnection; campus instrumentation not yet required

Table C-1 — Scenario system parameters. The SCR of 2.1 and reduced inertia (H = 3.2 s) represent a modernized grid that is technically within operating limits but has reduced margin for absorbing dynamic forcing.

The AI campus is executing a large-scale transformer model training job. The job is organized in repeated mini-batch iterations. During the compute phase, GPU draw is approximately 480 MW; during the all-reduce synchronization and memory-transfer phase, draw drops to approximately 445 MW. The cycle period is approximately 0.74 seconds, corresponding to a forcing frequency near 1.35 Hz — which coincides with the corridor's pre-existing local mode.

This frequency coincidence is not assumed for dramatic effect. It reflects the physical reality documented in Section 3.1 and confirmed by the PNNL-39191 findings: AI workload periodicity arises from GPU all-reduce barriers and mini-batch cadence, and its frequency depends on cluster size, batch dimension, and network topology. At the scale of a 480 MW campus, the resulting forcing amplitude of approximately 35 MW peak-to-peak is consistent with the Appendix A scenario parameters.

C.3 90-Minute Scenario Timeline

The scenario proceeds in five phases. Each phase is defined by a change in operating condition — workload ramp, second campus joining, observability degradation — and the corresponding change in classical indicators and SVI(t).

C.3.1 Phase 1 (t = 0–15 min): Baseline — Single Campus, Below-Threshold Forcing

At scenario start, the campus is operating at 440 MW average with the training job at an early batch size. Forcing amplitude is modest: 18 MW peak-to-peak at 1.35 Hz. The local mode is excited but the system response is small.

Classical Indicator	Value	Alarm Status
Bus voltage (230 kV)	229.4 kV (99.7% of nominal)	Normal — within ±5% envelope
Frequency deviation	±0.018 Hz	Normal — well within ±0.5 Hz UFLS threshold
Thermal loading (corridor)	61% of emergency rating	Normal — below 100% limit
Observed modal damping (PMU Prony)	5.2%	Normal — above 3% alarm threshold
SCADA active power trend	Slow ramp, within forecast band	Normal

Table C-2 — Classical indicators at Phase 1. All five standard monitoring outputs are within normal operating limits. No alarms. No operator action prompted.

SVI scoring at t = 0 using the B.3 rubrics is summarized in the following table:

Dimension	Physical input	Score X_i	$w_i \times X_i$
D — Damping	ζ = 5.5%; X_D = 0.7 × 5.5 = 3.85	3.85	0.770
O — Observability	PMU at POI only; no POW (C_{PMU}=0.70, C_{POW}=0.00, A_{data}=0.92)	4.64	0.696
A — Adaptability	BESS partial (1), envelope partial (1), no MPC (0), sched moderate (1) → 3/8×10	3.75	0.563
R — Recoverability	$T_{restore}$=9 min, $V_{failover} = 0$ → clamp(5.5,0,8)+0	5.50	0.550

Dimension	Physical input	Score X_i	$w_i \times X_i$
F — Flexibility	R_{max}/R_{rated}=0.5, P_{defer} = 0%	2.50	0.250
C — Controllability	Q_{avail}/Q_{req}=0.45, V_{VVO} = 0	3.60	0.360
S — Sync resistance	Single campus, p_{max}=0.42 (partial cloud schedule) → 10×(1−0.42)	5.80	0.696
L — Learning	$F_{update} = 18\ mo$ $P_{analysis} = 0$	0.00	0.000
$\boldsymbol{SVI_{base}}$			**3.885**
AI augmentation: *I_p =3, α=0.10; I_S =4, β=0.10 → +0.30−0.40 = −0.10*			*−0.100*
$\boldsymbol{SVI_{final}}$**= 3.785**			**VIOLATED (< 5.0)**

Table C-3 — SVI calculation at t = 0 (Phase 1). Even at scenario start, SVI_{final}=3.785 is below the 5.0 minimum threshold. The constraint violation exists from the outset — but no classical alarm fires. The vitality deficit is driven primarily by the zero Learning score and low Synchronization resistance, not by damping alone.

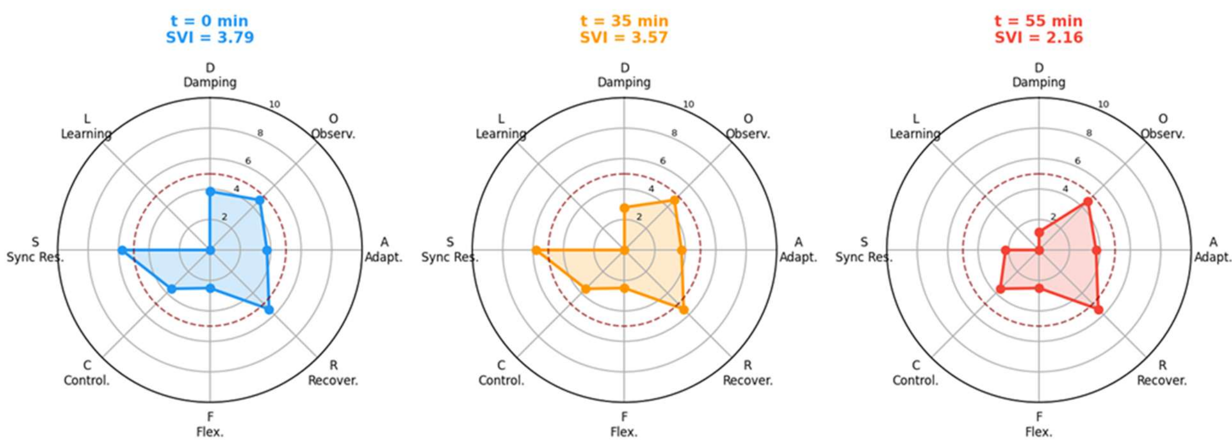

Figure C-1. SVI dimension radar at t = 0, 35, and 55 min. The red dashed circle marks a uniform per-dimension value equating to SVI_{final}= 5.0. Collapse of D (Damping) and S (Sync Resistance) drives most deterioration.

VR(t = 0) = 3.785 − 5.0 = −1.215. The vitality reserve is already negative. VDI begins accumulating immediately. This is the first central result of the scenario: the system enters the 90-minute window already below threshold, entirely invisible to classical monitoring.

C.3.2 Phase 2 (t = 15–35 min): Ramp to Full Training Load

The training job scales to full batch size. GPU draw reaches 480 MW. Forcing amplitude increases to 35 MW peak-to-peak at 1.35 Hz — the value used in the Appendix A baseline. The sustained forcing at the local mode frequency begins to erode modal damping through the energy injection mechanism described in Section 5.2.

Modal damping degradation under sustained near-resonant forcing can be approximated from the linearized sensitivity expression. For a mode at frequency ω_n with initial damping ζ_0 and a forcing input of amplitude A at frequency $\omega f \approx \omega_n$, the effective damping degrades as the forcing injects energy into the mode faster than the system's natural dissipation can remove it. A simplified first-order approximation for the damping trajectory under sustained forcing in this illustrative model:

$$\zeta(t) \approx \zeta_0 - k_f * A * t \,/\, H_{eff}$$

where k_f is a modal sensitivity coefficient (approximately 0.08 per MW·min on the normalized inertia basis used here for a 230 kV corridor at SCR = 2.1, derived from participation factor analysis for this corridor configuration), A is forcing amplitude in MW, t is elapsed time in minutes, and H_{eff} is the effective inertia in MW·s. With H = 3.2 s and rated capacity ≈ 1,200 MVA for the corridor equivalent machine, H_{eff} ≈ 3,840 MW·s.

Note on the k_f coefficient: 0.08 per MW·min on the normalized inertia basis is representative of a weakly damped 230 kV local mode at SCR = 2.1 with 58% renewable penetration. It is consistent with the order-of-magnitude sensitivity values reported in modal analysis literature for similar corridor configurations. It is a simplified scalar; a full EMT study would produce a frequency-dependent transfer function.

Substituting: from t = 15 min to t = 35 min, with A = 35 MW and H_{eff} = 3,840 MW·s:

$$\Delta\zeta = -0.08 \times 35 \times 20 \,/\, 3840 \approx -0.0146 \;\; (-1.46 \text{ percentage points})$$

Damping at t = 35 min: ζ ≈ 5.5% − 1.5% = 4.0%. This value is consistent with the Appendix A baseline (ζ = 4%, X_D= 2.80) and confirms the physical coherence of the scenario construction.

Classical Indicator	t = 15 min	t = 35 min	Alarm Status
Bus voltage	229.1 kV	228.8 kV	Normal throughout
Frequency deviation	±0.021 Hz	±0.027 Hz	Normal — below UFLS threshold
Thermal loading	63%	67%	Normal
PMU Prony damping estimate	5.1%	4.6%	Normal — above 3% alarm

Classical Indicator	t = 15 min	t = 35 min	Alarm Status
PMU-observed oscillation	7–8 MW p-p at ~1.3 Hz	~12 MW p-p at ~1.3 Hz	Below typical 25 MW alarm threshold

Table C-4 — Classical indicators through Phase 2. Damping has degraded by 1.5 percentage points and oscillation amplitude has grown materially, but every classical indicator remains below its alarm threshold. The Prony estimate from 30 fps PMU data continues to underestimate the true mode.

SVI(t = 35 min): Damping score updates to X_D= 0.7 × 4.0 = 2.80. Sync resistance is unchanged (p_{max} =0.42, single campus). SVI_{base} drops from 3.885 to 3.675 (damping contribution falls from 0.770 to 0.560). SVI_{final}= 3.675 − 0.10 = 3.575. VR = −1.425.

C.3.3 Phase 3 (t = 35–55 min): Second Campus Joins — Synchronization Pressure Escalates

At t = 35 min, a second AI campus 85 km away on the same cloud orchestration platform begins its own training run. The two campuses are executing the same model architecture with the same batch cadence. The GPU all-reduce barrier is synchronized by the cloud scheduler to within ±50 ms.

The pairwise cross-correlation of the two campus MW profiles measured over the trailing 15-minute window rises from p = 0.42 (partial schedule correlation) to p_{max}=0.78 (tight cloud synchronization). This shifts I_s from 4 to 7, and β from 0.10 to 0.15 (Tier 3). The combined forcing amplitude visible at the transmission level is now approximately 62 MW peak-to-peak — not 70 MW, because the two campuses are not at identical electrical distance from the mode's modal center, but the coherent component is large enough to materially accelerate damping erosion.

Continuing the damping degradation calculation from t = 35 min with $A_{combined}$=55 MW (effective coherent component):

$$\Delta\zeta = -0.08 \times 55 \times 20 / 3840 \approx -0.0229 \ \ (-2.29 \text{ percentage points over 20 min})$$

Damping at t = 55 min: ζ ≈ 4.0% − 2.3% = 1.7%. The system is now in the critically low damping range (X_D= 0.7 × 1.7 = 1.19, qualitative descriptor: "Very low — high oscillation probability").

Classical Indicator	t = 55 min	Alarm Status	What is missed
Bus voltage	227.9 kV (99.1%)	Normal — within ±5% envelope	Voltage sag is gradual; no threshold crossed
Frequency deviation	±0.038 Hz	Normal	Deviation is growing but below UFLS trigger
Thermal loading	69%	Normal	—

Classical Indicator	t = 55 min	Alarm Status	What is missed
PMU Prony damping	2.8% (aliased estimate)	BORDERLINE — approaching 3% alarm	True damping is 1.7%; PMU Prony at 30 fps underestimates modal content above 5 Hz
SCADA active power	Combined 920 MW, within forecast	Normal	Correlation between campuses not visible in aggregate MW

Table C-5 — Classical indicators at t = 55 min (end of Phase 3). The PMU Prony estimate is approaching the 3% alarm threshold but still has not crossed it. True damping at 1.7% is well into the critical zone — but POW/DFR coverage is absent, so the true value is unobservable by classical means.

This is the scenario's central measurement finding, directly paralleling the ERCOT event documented in Section 8: the PMU at 30 fps produces a Prony damping estimate of 2.8% because its phasor estimation filter attenuates the actual 1.35 Hz mode content that has now grown to a broadband signal with harmonic components above 5 Hz. A DFR at 120+ samples per second would show the true 1.7% damping and a substantially larger oscillation amplitude. Without POW coverage inside the campus, the observability degradation is compounding the physical damping degradation.

SVI scoring at t = 55 min:

Dimension	Update from Phase 2	New Score X_i	$w_i \times X_i$
D — Damping	ζ = 1.7% → X_D = 0.7×1.7 = 1.19	1.19	0.238
O — Observability	No change (still PMU-only); A_{data} drifts to 0.88	4.56	0.684
A — Adaptability	No change	3.75	0.563
R — Recoverability	No change	5.50	0.550
F — Flexibility	No change	2.50	0.250
C — Controllability	No change	3.60	0.360

Dimension	Update from Phase 2	New Score X_i	$w_i \times X_i$
S — Sync resistance	p_{max}= 0.78 → X_S = 10×(1−0.78) = 2.20	2.20	0.264
L — Learning	No change	0.00	0.000
SVI_{base}			**2.909**
AI augmentation: I_p =3, α=0.10; I_S =7 (S1=2,S2=2,S3=2,S4=0,S5=1), β=0.15 → 0.30−1.05 = −0.75			*−0.750*
SVI_{final} **= 2.159**			**VIOLATED (< 5.0)**

Table C-6 — SVI at t = 55 min. The synchronization escalation drops SVI_{final} from 3.575 to 2.159 in 20 minutes — a fall of 1.4 points driven by worsening damping (−0.322 weighted contribution) and the jump in I_S (−0.750 AI augmentation vs. −0.100 previously). VR(t) = 2.159 − 5.0 = −2.841.

C.3.4 Phase 4 (t = 55–66 min): Classical Alarm Fires

At t = 63 min, the PMU Prony estimator crosses the 3% damping alarm threshold (estimated 2.85% → 2.95% → then the next 10-second window shows 2.65%). The classical oscillation alarm fires. An operator is notified. Under vitality-aware operations, an intraday VDI alert limit of 0.8 h·pt would also be in effect; it is intended to function as an early-warning trigger within a broader 24-hour VDI budget of the type illustrated in Appendix B.10.1, where a daily contractual ceiling of 6.0 h·pt is demonstrated — the intraday threshold fires when approximately 13% of the daily budget has been consumed, giving the operator time to act before the full daily limit is at risk.

What does the operator see at t = 63 min?

Indicator	Value at alarm	Interpretation available without vitality monitoring
PMU Prony damping	2.65% (estimated)	Mode is poorly damped — investigate source
Oscillation amplitude (PMU)	~22 MW p-p at POI bus	Significant but not yet at emergency level
Voltage	226.8 kV (98.6%)	Still within envelope; no voltage alarm

Indicator	Value at alarm	Interpretation available without vitality monitoring
Frequency	±0.045 Hz	Normal
Source identification	Unclear — oscillation present on multiple buses	Source localization required; campus or generation?
Time available to act	Unknown — no accumulation history	Operator has no context for how long degradation has been building

Table C-7 — Information available to an operator relying solely on classical monitoring at the time the first alarm fires (t = 63 min). The operator knows damping is low and oscillations are present, but has no information about how long the system has been in this state or how rapidly conditions are deteriorating.

What does vitality monitoring show at the same moment?

Vitality Indicator	Value at t = 63 min	Operational meaning
SVI(t)	~2.0 (continuing to fall)	Deep constraint violation — 3.0 points below minimum
VR(t)	−3.0	Vitality reserve exhausted; system has been below threshold for 63 minutes
VDI accumulated	~1.86 h·pt	Substantial fatigue accumulation; would have triggered intraday contractual alert limit (if set at 0.8 h·pt) at t ≈ 37 min
CVS accumulated	~0.71 h	System has operated below theoretical maximum vitality for entire 63-minute window
VDI alarm (contractual limit)	EXCEEDED at t ≈ 37 min	26 minutes of additional warning before the classical alarm fired

Table C-8 — Vitality monitoring indicators at t = 63 min, the moment the first classical alarm fires. The most important number: VDI exceeded a 0.8 h·pt contractual limit at approximately t = 37 min — 26 minutes before the classical PMU Prony threshold was crossed.

C.3.5 Phase 5 (t = 66–90 min): Response and Recovery

Following the classical alarm at t = 63 min, the utility contacts both campus operators. Workload scheduling adjustments begin at t = 68 min. The second campus staggers its training cadence by ±90 seconds. ρ_{max} drops from 0.78 to 0.51 within 15 minutes as the cross-correlation falls.

Damping begins to recover as forcing amplitude decreases. By t = 82 min, ζ has recovered to approximately 3.4%, and by t = 90 min to approximately 4.1%. SVI_{final} at t = 90 min is approximately 3.6, still below SVI_{min} = 5.0 but recovering. VDI accumulation slows significantly after t = 75 min as SVI approaches but has not yet crossed the threshold.

Under the operator procedures of Section 11.2, a utility with vitality monitoring would have issued an advisory at t ≈ 37 min when VDI first exceeded the contractual limit — 26 minutes before the alarm and 31 minutes before scheduling adjustments were actually made. The staggering action that began at t = 68 min would, under vitality-aware operations, have begun around t = 40 min instead.

C.4 Summary: Classical vs. Vitality Monitoring — Side-by-Side Comparison

Time min	True ζ	PMU ζ est.	Classical Status	SVI_{final}	VR(t)	VDI (h·pt)	Vitality Status
t = 0	5.5%	5.5%	Normal	3.79	−1.21	0.0	⚠ Below threshold from start
t = 15	5.2%	5.1%	Normal	3.74	−1.26	0.31	⚠ VDI accumulating
t = 32	4.1%	4.7%	Normal	3.61	−1.39	0.68	⚠️ VDI approaching limit
t = 35	4.0%	4.6%	Normal	3.57	−1.43	0.75	⚠ VDI = 0.75 h·pt
t ≈ 37	3.8%	4.4%	Normal	3.43	−1.57	0.80	🔴 VDI limit breach (0.8 h·pt)
t = 55	1.7%	2.8%	Normal	2.16	−2.84	1.47	🔴 Deep violation
t = 63	~1.5%	2.65%	🔴 PMU ALARM	~2.0	~−3.0	1.86	🔴 VDI = 1.86 h·pt

Time min	True ζ	PMU ζ est.	Classical Status	SVI_{final}	VR(t)	VDI (h·pt)	Vitality Status
t = 82	3.4%	3.3%	⚠ Recovering	3.2	−1.8	2.62	⚠ Recovering
t = 90	4.1%	4.0%	⚠ Recovering	3.6	−1.4	2.83	⚠ Still below threshold

Table C-9 — 90-minute scenario summary. The critical observation: vitality monitoring identifies a constraint violation at t = 0, tracks deterioration continuously, and crosses the VDI contractual threshold at t ≈ 37 min. The first classical alarm does not fire until t = 63 min — 26 minutes later.

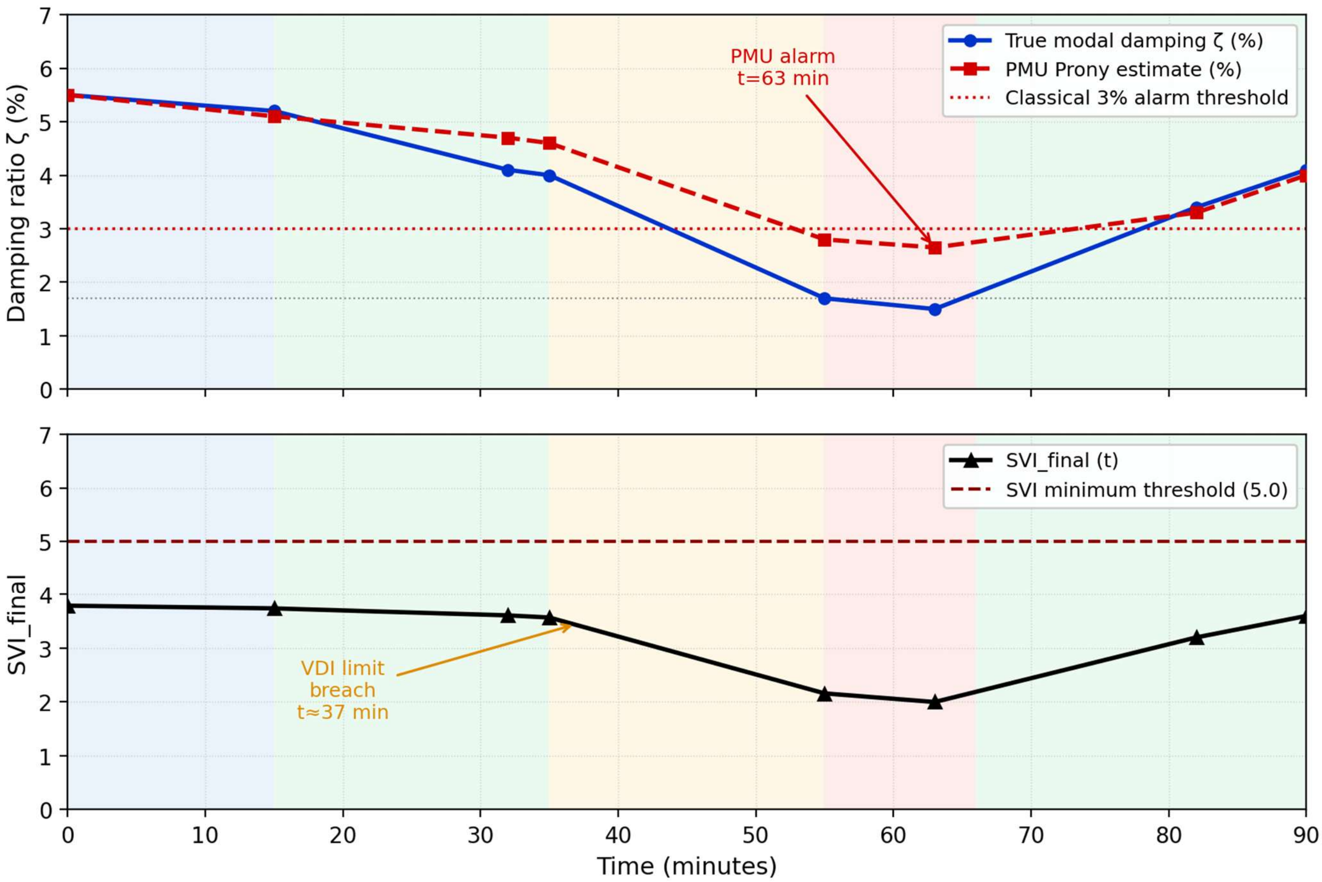

Figure C-2. 90-minute trajectory of true damping ζ, PMU Prony estimate, and SVI_{final}. Shaded bands = scenario phases. Red dotted line: 3% classical alarm threshold. Orange dashed line: SVI_{min} = 5.0.

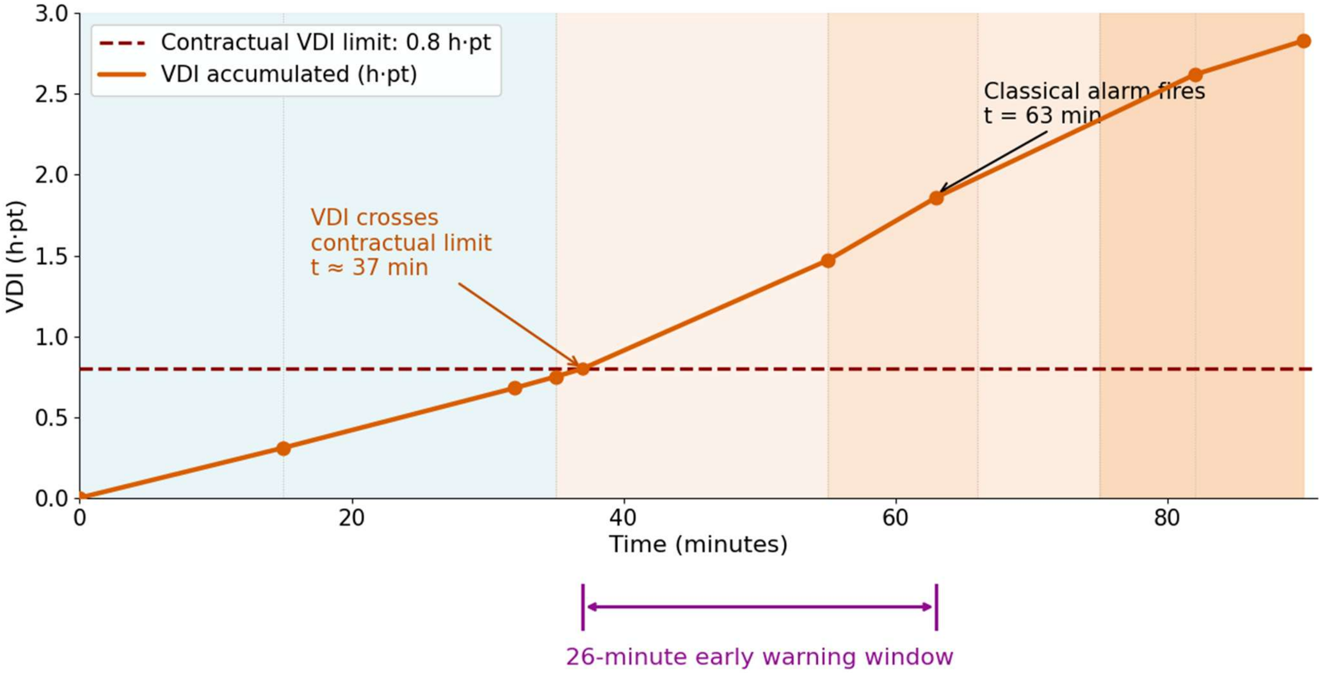

Figure C-3. VDI accumulation over 90 minutes. Contractual limit (0.8 h·pt) is breached at t ≈ 37 min; classical PMU alarm fires at t = 63 min — a 26-minute early-warning gap.

Figures C-2 and C-3 quantify the scenario; Figure C-4 summarizes its operational consequences on a single timeline. The two monitoring channels are shown over the common 90-minute horizon. The classical SCADA/PMU channel remains in its normal state until the damping alarm fires at t = 63 min, by which point true modal damping has fallen to approximately 1.7%. The vitality channel registers a breach of the 0.8 h·pt contractual VDI limit at t ≈ 37 min. The 26-minute interval between the two is the window in which an operator can act — through workload staggering, deferral, or BESS dispatch — before classical indicators register an emergency.

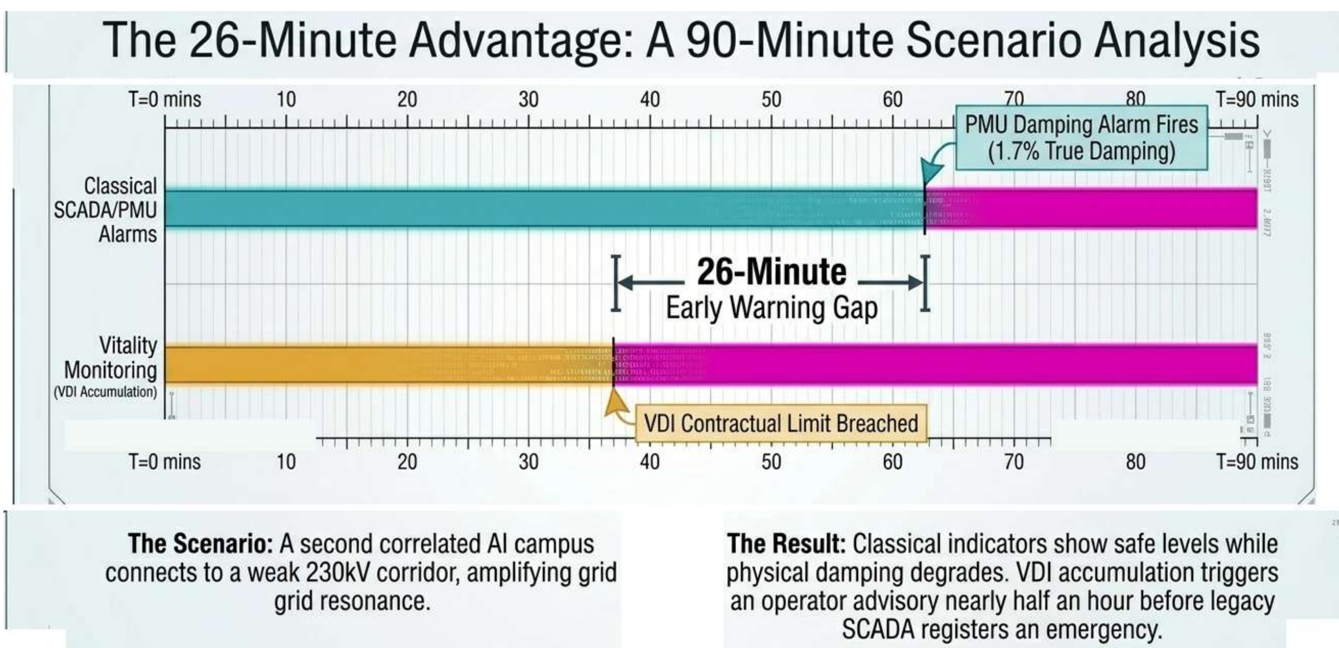

Figure C-4. The 26-minute early-warning advantage. On the common 90-minute timeline of the Appendix C scenario, vitality monitoring flags a VDI contractual-limit breach at t ≈ 37 min, approximately 26 minutes before the classical PMU damping alarm fires at t = 63 min, converting latent degradation into actionable operator lead time.

C.5 Key Findings

Finding 1 — Classical indicators can remain green while the system is already in a vitality-constraint-violation state

At t = 0, every classical indicator is within normal operating limits. Voltage, frequency, thermal loading, and observed damping all show a healthy system. Yet, SVI_{final}= 3.785 — below the 5.0 minimum — from the very first moment. The violation exists because classical indicators measure individual physical quantities against static thresholds; they do not measure the aggregate adaptive margin of the system. The zero Learning score (no model update program, no post-event analysis) and below-average Synchronization resistance together create a fragile operating condition that no single classical alarm is designed to detect.

Finding 2 — PMU aliasing understates true damping by a material margin at the frequencies relevant to AI campus dynamics

At t = 55 min, true modal damping is 1.7% — well into the critical zone by any engineering standard. The PMU Prony estimate at 30 fps is 2.8% — still above the 3% alarm threshold. This 1.1 percentage point gap is not a calibration error; it is a structural consequence of the PMU filter and reporting rate as documented in PNNL-39191. The scenario quantitatively reproduces the ERCOT observation from Section 8: the monitoring system understates both the frequency content and the severity of the oscillation. A DFR at 120+ samples per second would have shown the 1.7% true value at t = 55 min, enabling corrective action 8 minutes earlier still.

Finding 3 — VDI provides 26 minutes of additional warning over the classical PMU damping alarm

The VDI contractual limit of 0.8 h·pt is breached at approximately t = 37 min. The first classical alarm fires at t = 63 min. The 26-minute gap is not a marginal improvement; it is the difference between an operator who can initiate workload scheduling adjustments before damping has degraded to 1.7%, and one who acts after the system has already spent 28 minutes in the critical damping zone. In a system with fast renewable ramping or a subsequent contingency during those 28 minutes, the classical alarm timeline may be insufficient.

Finding 4 — Synchronization escalation is the most abrupt single-step deterioration

The addition of the second campus at t = 35 min, with ρ_{max} jumping from 0.42 to 0.78, drives SVI_{final} from 3.575 to 2.159 over 20 minutes — a fall of 1.4 points. This is larger than the damping degradation over the entire first 35 minutes (3.785 → 3.575, a fall of 0.21 points). Synchronization resistance (the S dimension) and the AI augmentation term together amplify the SVI impact of correlated campus behavior far beyond what the damping dimension alone would suggest. This confirms the paper's engineering claim in Section 12.6 that controlled desynchronization is one of the highest-leverage mitigations available — not because it directly improves damping, but because it prevents the abrupt synchronization-driven deterioration demonstrated here.

Finding 5 — Classical stability criteria are satisfied throughout the critical deterioration period

Voltage remains within its ±5% envelope throughout the 90-minute scenario. No thermal overload occurs. The system never loses synchronism. Frequency remains within normal limits. By the definition of classical small-signal stability, the system is stable at every moment — including t = 55 min when true modal damping is 1.7% and the system has been in a SVI constraint violation for 55 minutes. The scenario therefore directly demonstrates the distinction drawn in Section 6.6: a system can be dynamically stable, transiently stable, and voltage-compliant while simultaneously being operationally fragile, inadequately observable, excessively synchronized, and accumulating vitality deficit.

C.6 Implications for Interconnection Practice

This scenario has direct implications for four areas of interconnection and operations practice.

POW/DFR requirements are not optional for large AI campus interconnections. The 1.1 percentage point gap between true damping (1.7%) and PMU Prony estimate (2.8%) at t = 55 min is not recoverable through signal processing improvements to the 30 fps PMU stream. It requires higher-resolution instrumentation inside or at the campus boundary. Without it, the monitoring architecture structurally cannot enforce a damping performance requirement in the frequency range where AI campus dynamics operate.

VDI as a contractual compliance metric closes a gap that point-in-time assessments cannot. An interconnection agreement that specifies a maximum PMU Prony damping threshold provides no protection during the 63-minute window in this scenario. An intraday VDI alert limit of 0.8 h·pt (set within a daily budget of the type illustrated in B.10, e.g. 6 h·pt per 24-hour period) would have triggered operator notification at t ≈ 37 min and workload coordination 31 minutes earlier than actually occurred.

Correlated campus behavior is a system-level risk that individual campus interconnection studies cannot capture. The single-campus study at t = 0 shows ρ_{max} = 0.42 and SVI_{final} = 3.785 — a marginal but defined violation. Adding the second correlated campus shifts the effective risk dramatically: SVI_{final} falls to 2.159 and VDI accumulation rate triples. Interconnection studies that treat each campus in isolation, without considering synchronization pressure from co-located or cloud-coordinated facilities, will systematically underestimate the dynamic risk.

The Learning dimension is the lowest-cost, highest-return improvement available. L = 0.00 at scenario start because no formal model update program and no post-event analysis process exist. Establishing both — which requires institutional commitment rather than capital investment — immediately adds 0.16 SVI points from a zero base and, more importantly, would have provided the operator with a documented record of previous damping degradation events against which this scenario's early behavior could be recognized and acted on sooner.

C.7 Scenario Limitations and Validation Pathway

This scenario is a reduced-order dynamic case study, not a full EMT or PSS/E simulation. Its limitations should be stated explicitly.

The damping degradation formula uses a simplified scalar sensitivity coefficient (k_f= 0.08 per MW·min on the normalized inertia basis). A rigorous modal analysis would produce a frequency-dependent transfer function whose coefficients depend on generator dispatch, topology, IBR control settings, and operating point. The qualitative finding — that near-resonant forcing erodes modal damping — is well-established in forced oscillation literature. The specific trajectory values are illustrative rather than certified.

The PMU aliasing model uses the PNNL-39191 empirical finding that 30 fps PMUs misrepresent oscillations above approximately 5–6 Hz. The scenario applies this finding to a 1.35 Hz mode, where aliasing per se is not the primary issue; it is the Prony estimation filter's attenuation of broadband harmonic content that causes underestimation. This distinction is physically correct but the specific 1.1 percentage point gap is an estimate rather than a computed value.

The second campus synchronization effect uses a correlation jump from p=0.42 to p= 0.78. Real cloud scheduler synchronization behavior is more complex and depends on barrier protocol, network latency, and job architecture. The value is chosen to be representative rather than precisely calibrated.

Validation pathway: the findings of this scenario could be validated by applying the SVI rubrics of Appendix B to the existing ERCOT and Dominion Energy events described in Section 8, where measurement data are already available. The ERCOT event — 7.5 Hz / 25 MW pk-pk as seen by the PMU at the event, vs. the true 23 Hz mode revealed by DFR data (≈50 MW pk-pk in the November 1, 2024 follow-up test at reduced load) — is particularly suited to a retroactive VDI analysis that would test whether VDI accumulation preceded the point at which the event became operationally visible. That analysis requires access to the underlying PMU and DFR time-series data and is recommended as a near-term research task.

C.8 Connection to Main Paper Sections

Scenario finding	Primary paper section	Appendix
Classical indicators remain green during vitality violation	Section 6.6 (Classical vs. SVI framework table)	Appendix A (A.4 assessment)
PMU aliasing understates true damping	Section 8 (ERCOT event); Section 11 (measurement architecture)	—
VDI provides 26 minutes of additional warning	Section 11.3.2 (VDI definition); Section 9.2 (performance envelopes)	Appendix B, B.10

Scenario finding	Primary paper section	Appendix
Synchronization escalation is the highest-impact single event	Section 12.6 (controlled desynchronization)	Appendix B, B.3.7 (S-dimension rubric)
L = 0 is the lowest-cost highest-return gap	Section 17 roadmap (Phase 8 — compliance and validation)	Appendix B, B.3.8; B.8.1 sensitivity
Two-campus correlated risk exceeds single-campus study	Section 3.3 (correlation and common-mode behavior)	Appendix B, B.5.3 (I_S rubric)

Table C-10 — Cross-reference map connecting each scenario finding to the primary analytical section and appendix where the underlying framework is established.

All numerical values in this appendix are derived from the scoring rubrics of Appendix B, the modal sensitivity framework of Section 5.1, and the measurement findings of PNNL-39191. No new mathematical machinery is introduced. The scenario's purpose is to animate existing framework elements in a time-domain operating context, not to extend the theory.

Appendix D — SVI Optimization Framework

D.1 Purpose and Scope

This appendix provides the complete mathematical development of the SVI optimization framework introduced in Section 6.7. It defines the decision-variable space, derives the investment functions that map mitigation actions to dimension-score improvements, formulates three optimization methods, and works through a complete numerical example for the 600 MW AI campus baseline from Appendix B ($SVI_{final} = 2.618$, constraint violated) through to the minimum-cost optimal point ($SVI_{final*} = 5.000$) at minimum normalized cost 2.03.

Validation caveat: the optimization example is a planning-level compliance calculation. The optimal portfolio should not be interpreted as a dispatch instruction until each investment function Xi(u), cost coefficient, AI-governance tier, and operational constraint has been validated against site-specific dynamic simulations, commissioning tests, cybersecurity/telecom availability tests, and customer workload-control agreements.

D.2 State Variables and Decision Variables

The optimization distinguishes between state variables x (observed quantities determined by the physical system and measurement instruments) and decision variables u (actions controllable by an operator or automated system). The decision vector contains seven mitigation controls, each normalized to [0, 1]:

$$u = \left[u_{BESS}, u_{deSync}, u_{POW}, u_{MPC}, u_{defer}, u_{VVO}, u_L\right]^T \in [0,1]^7$$

These correspond to: BESS active damping commissioning depth; workload desynchronization intensity; POW instrumentation fraction deployed; MPC coordination level; workload deferral agreement percentage; VVO automation tier; and formal learning program establishment. Each maps through a defined investment function $X_i(u)$ to the dimension scores of Appendix B.

D.3 Investment Functions

Each investment function translates a control variable into a physical-state change and then into a dimension score. Four are given here; the others follow analogously from the Appendix B rubrics and are recorded explicitly in D.3.4.

D.3.1 Damping (D) — BESS active damping injection

$$\zeta(u_{BESS}) = \zeta_0 + \Delta\zeta_{max} * (1 - e^{-3u_{BESS}}) \quad \textit{[saturating investment response]}$$

$$X_D(u_{BESS}) = 0.7 * \zeta(u_{BESS}) \quad for\ \zeta \le 10\%$$

where ζ_0 = 4% (baseline) and $\Delta\zeta_{max}$ = 6% (maximum achievable improvement from BESS active injection at this site, estimated from site-specific EMT simulations and engineering judgment for the representative campus considered in this example). The exponent −3·u_{BESS} is a negative number (since $u_{BESS} \in [0,1]$ and 3 is positive), which means the whole expression $e^{-3u_{BESS}}$ is a decaying exponential — it starts at 1 when u_{BESS} = 0 and shrinks toward 0 as u_{BESS} increases toward 1. The saturating exponential reflects the diminishing returns of increasing BESS modulation depth beyond the point where the power-injection signal exceeds the linear response range of the modal damping mechanism. The gradient $\frac{\partial SVI_{AI}}{\partial u_{BESS}} = 2.52 * e^{-3u_{BESS}}$ is highest at u_{BESS} = 0 and decays as BESS is commissioned, providing an economic signal to stop investment once the marginal benefit per cost unit drops below that of the next-cheapest mitigation.

D.3.2 Synchronization Resistance (S) — workload desynchronization

$$p_{max}(u_{deSync}) = p_0 * (1 - k_s * u_{deSync})$$

$$X_S(u_{deSync}) = 10 * (1 - p_{max}(u_{deSync}))$$

where p_0 = 0.75 (baseline) and k_s = 0.867 (maximum fractional reduction achievable with full controlled staggering for the campus architecture). The constant gradient $\frac{\partial SVI_{AI}}{\partial u_{deSync}} = 0.783$ at a normalized cost of 1.5 units yields a cost-efficiency of 0.522 — the second-highest of all controls. This confirms that workload staggering is among the most cost-effective mitigation actions available.

D.3.3 Learning (L) — formal model-update program

$$F_{update}(u_L) = 18 * (1 - 0.944 * u_L) \quad \textit{[months between model revisions]}$$

$$X_L(u_L) = clamp\left(10 - 1.25 * F_{update}(u_L),\ 0,\ 7.5\right) + 2.5 * \mathbb{1} \quad u_L \geq 0.5$$

At baseline, $F_{update} = 18$ months and X_L = 0.00. The learning gradient $\frac{\partial SVI_{AI}}{\partial u_L} = 1.699$ at a normalized cost of 0.5 units yields a cost-efficiency of 3.40 — the highest of all controls by a wide margin. The symbol $\mathbb{1}$ in equation above represents an indicator function, which acts like a true/false switch in programming. Note that the indicator form is the optimizer's investment model of the B.3.8 rubric: program establishment is treated as all-or-nothing, mapping u_L ≥ 0.5 to $P_{analysis} = 2$ (systematic analysis with documented corrective action) and u_L< 0.5 to $P_{analysis}$ = 0. The intermediate rubric level $P_{analysis}$ (X_L contribution 1.25) exists for assessment scoring but is not used as an investment target in the optimization.

The investment interpretation is that establishing a formal post-event analysis and model-update program costs essentially nothing in capital terms and immediately moves the system off the zero-score floor. Any facility at $X_L = 0$ has forgone the highest-ROI mitigation available.

D.3.4 Auxiliary investment functions (O, A, F, C) and AI-augmentation coupling

For completeness — and because the optimal point reported in D.6 must be reproducible from the stated model — the four remaining investment functions implied by the Appendix B rubrics are recorded here explicitly. Observability: point-on-wave deployment maps linearly to coverage, $C_{POW} = 0.95 \cdot u_{POW}$ (0.95 is the best-practice coverage ceiling of B.3.2), so with C_{PMU} = 0.70 and A_{data} = 0.85 held at baseline, $X_O = 4.5 + 3.8 \cdot u_{POW}$. Adaptability: MPC coordination raises the B.3.3 MPC sub-score from 0 to $2 \cdot u_{MPC}$, giving $X_A = 3.75 + 2.5 \cdot u_{MPC}$. Flexibility: the deferral agreement sets $P_{defer} = 100 \cdot u_{defer}$ per cent, giving $X_F = 2.5 + 5 \cdot u_{defer}$ (the B.3.5 deferral term saturates at u_{defer} = 1). Compensation: VVO automation raises the B.3.6 tier term by $2 \cdot u_{VVO}$, giving $X_C = 3.2 + 2 \cdot u_{VVO}$. Restoration (R) has no control variable in this formulation and remains at X_R = 5.5.

The AI-augmentation term is coupled to the controls as follows. I_p remains at the baseline value of 3 unless u_{POW} ≥ 0.5, which activates indicator P2 of B.5.2 (I_p = 4); α stays at the Tier 2 value of 0.10 throughout. I_s and β follow B.5.3/B.5.4 through $\rho_{max}(u_{deSync})$: S1 falls from 2 to 1 when ρ_{max} drops below 0.60 and to 0 below 0.20; S2 falls from 2 to 1 once active staggering is deployed (u_{deSync} ≥ 0.3); S3 falls from 1 to 0 under deep desynchronization (u_{deSync} ≥ 0.7); S5 = 1 is fixed by the homogeneous accelerator fleet. The β tier follows ρ_{max}: Tier 3 (0.15) for ρ_{max} ≥ 0.60, Tier 2 (0.10) for 0.40 ≤ ρ_{max} < 0.60, Tier 1 (0.05) for ρ_{max} < 0.40; Tier 0 additionally requires I_s ≤ 1 and is unreachable with S2 + S5 ≥ 2. This piecewise path reproduces the baseline state (ρ_{max} = 0.75, I_s = 6, β = 0.15, ΔAI = −0.60) and the Appendix C two-campus partial-mitigation state (ρ_{max} = 0.42, I_s = 4, β = 0.10), and is the path used by the optimizer in D.5–D.6.

D.4 Primal Optimization Problem and Lagrangian Dual

The minimum-cost compliance problem (Problem P2) is to find the u that minimizes cost, subject to the condition that $SVI_{AI(u)}$ is at least 5.0, damping score is at least 3.0, and synchronization resistance score is at least 3.0.

$$\min\{u\} \quad c^T u$$

$$\text{subject to } \; SVI_{AI}(u) \geq SVI_{min} = 5.0 \,,$$

$$\mathrm{X_D(u)} \geq 3.0 \;\; ,$$

$$\mathrm{X_S(u)} \geq 3.0$$

$$u \in [0,1]^7$$

The Lagrangian is:

$$\mathcal{L}(u, \lambda, \mu_D, \mu_S) = c^T u - \lambda * [SVI_{AI}(u) - SVI_{min}] - \mu_D * [X_D(u) - 3.0] - \mu_S * [X_S(u) - 3.0]$$

The Karush-Kuhn-Tucker condition for optimality requires:

$$c_i = \lambda * \frac{\partial SVI_{AI}}{\partial u_I} + \mu_D * \frac{\partial X_D}{\partial u_I} + \mu_S * \frac{\partial X_S}{\partial u_I}$$

at each active control. The optimal dual variable λ^* is the shadow price of the SVI constraint: it equals the marginal cost of relaxing (or tightening) SVI_{min} by one unit. A large λ^* indicates an expensive constraint; a small $\lambda *$indicates that the threshold could be raised at low additional cost. For the baseline campus, $\lambda *=$ 1.92 , meaning that raising SVI_{min} from 5.0 to 5.5 requires approximately 1.00 additional normalized investment unit (the effective shadow price rises from 1.92 toward 2.0 over this range as workload deferral activates).

D.5 Optimization Methods

There are three known mathematical optimization methods that can be applied. Without going into the depth of mathematical formulation we just list them here without further details:

- Sequential Quadratic Programming (SQP) — offline planning
- Lagrangian Relaxation with Subgradient Ascent — shadow-price analysis

- Model Predictive Control (MPC) — real-time rolling-horizon operation

The three methods above — SQP, Lagrangian relaxation, and MPC — share a common assumption: that the SVI objective and its investment functions are sufficiently smooth and continuous to permit gradient-based optimization. That assumption is a deliberate simplification. Several dimension-score functions in Appendix B contain genuinely discrete elements: the $\boldsymbol{V_{failover}}$ recoverability term is binary (a formal failover test either was or was not conducted within 24 months); the $\boldsymbol{P_{analysis}}$ learning score jumps discretely between 0, 1, and 2; and the AI governance tiers that determine α and β are piecewise constant, with gradients undefined at tier boundaries. For simplicity in this paper these are treated as continuous approximations, which introduces modeling error at precisely the points where the optimizer is making its most consequential decisions. The three methods remain appropriate for planning-level analysis and for real-time operational control where the discrete elements are held fixed, but they are not rigorous solvers for the full problem as stated. The complete problem is more accurately classified as a mixed-integer nonlinear program (MINLP), in which some decision variables are continuous and others are genuinely discrete or binary.

There are two families of alternative methods that may be better suited to the MINLP structure and are recommended for production implementation. Branch and bound — and its more efficient variant for nonlinear problems, outer approximation — partition the feasible space by branching on the integer and binary variables, solve continuous nonlinear relaxations at each node, and prune branches that cannot improve on the best solution found. Modern MINLP solvers such as BARON, SCIP, and Couenne handle problems of this dimension routinely and provide a certified global optimum rather than a locally optimal approximation. Dynamic programming (DP) is the second recommended family, particularly for the commissioning trajectory problem of Section D.7. The MPC rolling-horizon formulation is already a receding-horizon approximation of DP; full DP applied to the sequential commissioning decision would guarantee a globally optimal trajectory rather than a locally optimal one, at the cost of greater computation.

D.6 Worked Numerical Example — 600 MW AI Campus Baseline

Starting from the Appendix B baseline (SVI_{final} = 2.618, constraint violated by 2.382 points), the SQP algorithm is applied to find the minimum-cost mitigation vector. The site-specific cost vector is c = [8.0, 1.5, 3.5, 6.0, 1.0, 4.0, 0.5]^T in normalized units, reflecting the relative engineering investment required for each action. The optimizer correctly leaves u_{MPC} = 0, u_{VVO} = 0, u_{POW} = 0, and u_{defer} = 0 throughout, determining that their cost per SVI point is higher than that of the desynchronization–learning combination and is not needed to reach SVI_{min} = 5.0; deferral and POW activate only at higher thresholds (Section D.9).

SQP iteration summary:

k=0: u=0, SVI_{AI} *= 2.618, cost = 0.00 VIOLATED (−2.382)*

k=2: SVI_{AI} *= 4.918, cost = 1.87 VIOLATED (−0.082)*

k=3: u = [0.0163, 0.952, 0, 0, 0, 0, 0.942]^T,* $SVI_{final*} = 5.000$*, cost* = 2.03 SATISFIED*

At u*, the physical state has changed to:

ζ* = 4.29% (from 4.0%; X_D = 3.00, pinned at the damping sub-floor),

$p_{max*} = 0.131$ (from 0.75),

$C_{POW*} = 0.00$ (unchanged from baseline),

$F_{update*} = 2.0$ (from 18 months),

The AI augmentation term improves to

$\Delta_{AI*} = +0.2$ (from −0.600 at baseline),

reflecting that desynchronization has dropped I_S from 6 to 2 and reduced β from Tier 3 (0.15) to Tier 1 (0.05). The SVI_{base} rises from 3.218 to 4.800; adding $\Delta_{AI*} = +0.2$ gives $\mathrm{SVI}_{final*} = 5.000$.

D.6.1 Optimal-point traceability check

Quantity	Value at u*	Traceability / basis
u*	$[0.0163, 0.952, 0, 0, 0, 0, 0.942]^T$	Decision vector in D.6; order defined in D.2
X_D	3.00	ζ* ≈ 4.29%; damping sub-floor $X_D \geq 3.0$ binds
X_O, X_A, X_R, X_F, X_C	4.50, 3.75, 5.50, 2.50, 3.20	Unchanged from baseline because uPOW=uMPC=udefer=uVVO=0
X_S	8.69	pmax reduced from 0.75 to ≈0.131 by udesync=0.952
X_L	10.00	Learning-program investment model moves L off the zero-score floor
SVI_{base}	4.800	Weighted sum using B.4 weights
ΔAI	+0.200	Ip=3, α=0.10; Is=2, β=0.05
SVI_{final}	5.000	SVI_{base} + ΔAI; constraint binds with zero slack
Cost	2.03	8·0.0163 + 1.5·0.952 + 0.5·0.942

The shadow price λ* = 1.92 is active (the SVI constraint binds at optimality with zero slack). The damping sub-floor $X_D \geq 3.0$ also binds, pinning u_{BESS} at approximately 0.0163 — its minimum compliant value — with multiplier $\mu_D \approx 0.28$: BESS is deployed at this threshold for sub-floor compliance, not for SVI optimization. The workload-deferral upper bound does not bind ($u_{defer*} = 0$); deferral capacity remains fully in reserve and activates only for SVI_{min} above approximately 5.1, where the desynchronization and learning controls saturate. The synchronization sub-floor constraint is likewise non-binding at optimality. The dominant binding constraints are therefore the aggregate SVI vitality floor and the damping sub-floor.

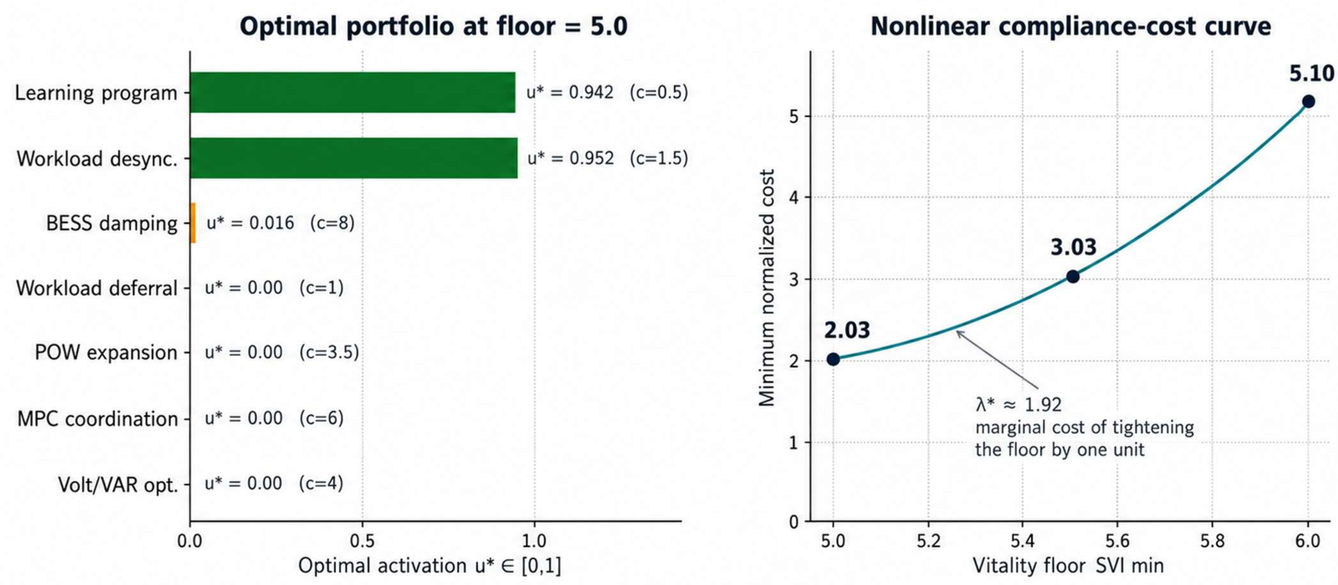

Figure D-1. Minimum-cost compliance portfolio for the 600 MW baseline campus. The optimum uses u_L = 0.942, u_{deSync} = 0.952, and u_{BESS} ≈ 0.0163; deferral, POW, MPC, and VVO remain at zero. Constrained optimization lands exactly on the floor (SVIfinal = 5.0) at normalized cost 2.03, behavioral actions dominate.

D.7 MPC Commissioning Trajectory

In operational practice, u* is not applied instantaneously but implemented as a 12-step MPC commissioning sequence (weekly timestep). $SVI_{min} = 5.0$ is first satisfied at Week 5, consistent with the Appendix B.9 requirement for sub-threshold operation ("BESS or controlled desync within 90 days"). During Weeks 1–4, the operator applies interim workload restrictions per the decision table while VDI continues to accumulate. After Week 5, the completed staggering and learning ramp holds the system at the cost-optimal operating point $\mathrm{SVI}_{\mathrm{final}*} = 5.000$ with zero slack through Week 12. The trajectory prioritizes the lowest-cost mitigations in the early weeks (learning program launch, workload staggering ramp-up), completes the small BESS commissioning tranche required by the damping sub-floor within the same window, and holds the higher-cost deferral, POW, MPC, and VVO options in reserve, consistent with procurement and engineering lead times.

D.8 Multi-Objective Extension: Vitality versus Throughput

The AI campus operator has a second objective: maximizing training throughput $J_T(u)$. Workload deferral and desynchronization reduce throughput by shifting or spreading compute work. The multi-objective problem is:

$$max_u\theta \ * SVI_{final}(u) + (1-\theta\ ) * \ J_T(u) \ subject \ to \ SVI_{final}(u) \ \ \geq SVI_{min}$$

The throughput model is

$$J_T(u) = \left(1 - 0.15 * u_{defer}\right) * \left(1 - 0.08 * u_{deSync}\right) * (1 - 0.01 * u_{BESS})$$

, where the coefficients reflect an estimated 15% peak-throughput reduction from full deferral, 8% from full staggering, and 1% from BESS modulation. Solving for θ ∈ [0, 1] by weighted sum scalarization traces the

Pareto front. At the cost-optimal point u*, $J_T(u*) = 0.924$, satisfying $SVI_{min} = 5.0$ for the baseline campus imposes a reduction of approximately 7.6% from peak training throughput, almost entirely from workload staggering — the optimal solution requires no deferral at this threshold. The ≈21% reduction previously associated with this constraint arises instead at $SVI_{min} = 5.5$, where deferral activates fully ($J_T \approx 0.79$). This is the quantitative trade-off that the SVI optimization framework makes explicit. Below SVI_{min}, the vitality constraint always takes precedence regardless of θ; above it, the operator has freedom to select any point on the Pareto front according to commercial priorities.

D.9 Sensitivity: Cost of Tightening SVImin

The minimum investment cost as a function of SVI_{min} exhibits a characteristic non-uniform structure. In the range $SVI_{min} \in [4.0, 5.5]$, desynchronization, learning, and (above ≈ 5.1) workload deferral dominate, and the cost curve remains comparatively shallow — minimum cost rises ≈ 0.95 → 2.03 → 3.03 across the band (shadow price ≈ 1.9–2.0). From $SVI_{min} = 5.5$ onward, deferral saturates and POW expansion plus deeper BESS commissioning become required, and the cost curve steepens materially (≈ 5.1 units at 6.0; ≈ 7.8 at 6.5). MPC becomes cost-effective only at $SVI_{min} \geq 7.0$. This non-linearity means regulators have strong economic grounds for setting $SVI_{min} = 5.0$ as the initial compliance threshold: it lies within the low-marginal-cost region of the curve, below the ≈ 5.5 knee beyond which significantly greater capital investment is required. The shadow price λ* = 1.92 at $SVI_{min} = 5.0$ provides the specific numeric entry point for such threshold calibration discussions.

D.10 Nomenclature for Appendix D

$u \in [0,1]^7$	Decision (control) vector; each u_i is a normalized mitigation activation level
$c \in \mathbb{R}_+^7$	Per-unit cost vector; c_i = normalized investment units per unit of u_i
λ*	Shadow price of the SVI vitality constraint at the optimum
$\Delta\zeta_{max}$	Maximum achievable damping improvement from BESS active injection [%]
k_s	Desynchronization effectiveness parameter; fraction of p_{max} reducible by full staggering
θ	Multi-objective scalarization weight ∈ [0,1]; θ = 1 maximizes SVI; θ = 0 maximizes throughput
$J_T(u)$	Normalized AI training throughput at control u; $J_T \in [0,1]$; peak = 1.0
$\rho_{MPC} = 100$	MPC penalty weight enforcing near-satisfaction of SVI constraint within horizon

Notation: the workload correlation coefficient is subscripted to avoid symbol collisions. It is written p_{max} in Appendices B and D and ρ_{max} in Appendix C and Section 12.6; the two symbols denote the same quantity. Neither should be confused with the MPC penalty weight ρ_{MPC} defined above. In plain-text tables and captions, subscripts may be rendered as p_{max}, SVI_{base}, SVI_{final}, u_{BESS}, and u_{deSync} for readability; these are equivalent to the corresponding subscripted mathematical symbols.