SOLUTIONS BY JEWEL · AI GOVERNANCE CONSULTING

AI ECOSYSTEM DEPENDENCY MAP

Mapping Capital, Infrastructure, Technology, Ownership, Dependencies, Decision Authority, and Evidence Across the AI Ecosystem

SBJ

SOLUTIONS BY JEWEL

SBJ ORIGINAL GOVERNANCE METHODOLOGY

LOCATE → ASSIGN → RECORD → SCALE

Purpose & Core Principle

The AI Ecosystem Dependency Map is a governance instrument, not a technology diagram. Its function is to force an organization to answer, in writing and with evidence, where its AI capability actually comes from, who controls each link in that chain, and what happens the day a link breaks.

Every AI system an enterprise deploys sits inside a chain of dependencies it did not build and does not fully control — capital, chips, cloud capacity, foundation models, training and fine-tuning data, and software platforms, each held by a small number of providers. Governance failures in this ecosystem are rarely caused by a single bad model. They are caused by unexamined concentration: multiple "different" vendors that all sit on the same chip, the same cloud region, or the same foundation model, so that one upstream event cascades through an organization's entire AI portfolio at once.

> This instrument is designed to make ten questions answerable, with evidence, at any time: what we depend on, who provides it, what infrastructure sits beneath it, where financial and ownership relationships exist, where concentration is forming, where the critical dependencies sit, who holds decision authority, what controls apply, what evidence supports each governance decision, and what the organization's actual exposure is if a provider becomes unavailable or changes materially.

Core Principle

The AI ecosystem is a network of interdependencies, not a linear supply chain. The reference architecture below moves from capital to end users, but real dependency relationships cut across layers — a cloud provider is also an investor in a model provider; a semiconductor company both supplies and invests in a hyperscaler; a data provider is also a deployment partner. The Map is built to capture cross-layer relationships explicitly, not to force them into a straight line.

Layer	Position in the ecosystem
1	Capital — investors, financing sources, capital commitments funding the ecosystem
2	AI companies — the organizations capital flows into; ownership and equity structures sit here
3	Infrastructure — semiconductors, servers, networking, power, data-center operators
4	Compute — aggregated processing capacity drawn from infrastructure
5	Models — foundation models and the providers that train and host them
6	Data — training, fine-tuning, and proprietary data sources feeding the models
7	Software / platforms — developer platforms, APIs, orchestration and MLOps layers
8	Applications — the products built on top of models and platforms
9	Enterprise deployment — internal systems, business functions, and end users consuming the application

Classification key used throughout this instrument: VERIFIED RELATIONSHIP · REPORTED / ORGANIZATIONAL CLAIM · ANALYTICAL INFERENCE · RISK ASSESSMENT · GOVERNANCE RECOMMENDATION

Executive Ecosystem Map

The nine-layer AI value chain, encircled by the five governance disciplines that must be applied at every layer. Relationship lines indicate the type of dependency crossing between layers — not every dependency runs strictly downward.

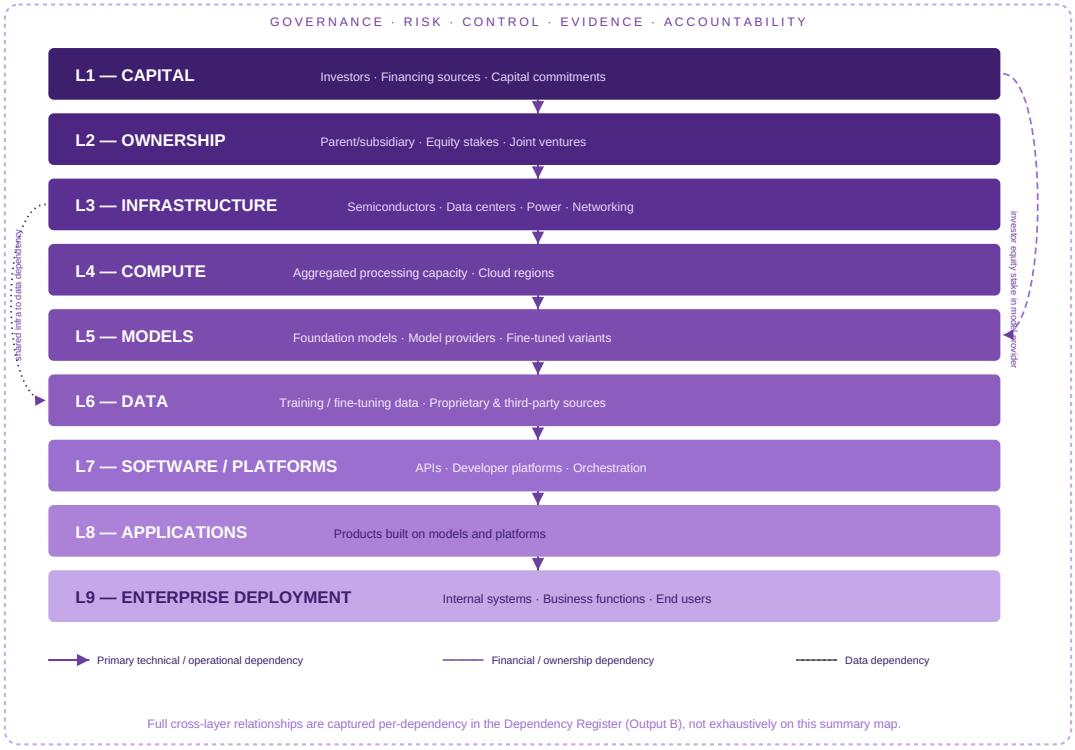

Full cross-layer relationships are captured per-dependency in the Dependency Register (Output B), not exhaustively on this summary map.

Dependency Register — Structure

The register is the system of record. One row equals one dependency. Every field below is mandatory for a dependency to be considered governed; a dependency with unresolved required fields is, by definition, ungoverned and should carry a Status of "Incomplete."

- Dependency ID
- Organization
- Business Function
- AI System / Use Case
- Dependency Type (A–H category)
- Dependency Description
- Provider
- Provider Tier
- Technology Layer (L1–L9)
- Criticality
- Concentration Level
- Financial Relationship
- Ownership Relationship
- Contractual Relationship
- Technical Dependency
- Data Dependency
- Operational Dependency
- Decision Authority
- Risk Owner
- Control Owner
- Primary Control
- Evidence Required
- Evidence Location
- Failure Impact
- Substitutability
- Switching Difficulty
- Recovery Time
- Geographic / Jurisdictional Exposure
- Regulatory Exposure
- Known Conflicts
- Monitoring Frequency
- Last Reviewed
- Next Review
- Status
- Risk Rating
- Remediation Owner
- Remediation Due Date
- Notes

Dependency Type Taxonomy (A–H)

Cat.	Category	Representative sub-types
A	Financial dependency	Investor, investment relationship, financing source, debt exposure, guarantee, capital commitment, revenue dependency, funding concentration
B	Ownership dependency	Parent company, subsidiary, acquisition, equity stake, joint venture, strategic ownership, cross-ownership
C	Infrastructure dependency	Semiconductor provider, server provider, networking provider, data-center operator, power provider, cloud provider, colocation provider
D	Technology dependency	Foundation model, AI model provider, model platform, AI framework, API, developer platform, software dependency, open-source dependency
E	Data dependency	Training data, fine-tuning data, proprietary data, third-party data, data provider, data storage, data processing
F	Operational dependency	Critical vendor, service provider, managed service, support provider, deployment partner, integration partner
G	Governance dependency	Decision authority, policy owner, model owner, system owner, risk owner, compliance owner, security owner, procurement authority, escalation authority
H	Evidence dependency	Decision record, contract, risk assessment, model documentation, audit report, security assessment, compliance documentation, testing evidence, monitoring record, incident record

Example Completed Register (Illustrative)

ILLUSTRATIVE All organization and provider names below are fictional and constructed solely to demonstrate register structure. No real financial, ownership, or contractual relationships are asserted. See the Case Study (Section 09) for the full narrative version of this same illustrative scenario.

ID	Function	Type	Provider	Layer	Crit.	Concentration	Risk Owner	Decision Auth.	Substitut.	Status
DEP-001	Customer support automation	D — Technology	Northgate Frontier Models, Inc.	L5 Models	5	Single-provider	Head of AI Risk	CAIO	Low	Active
DEP-002	Customer support automation	C — Infrastructure	Meridian Cloud Services	L3 Infra / L4 Compute	5	Single-provider	Infrastructure Owner	CTO	Low	Active
DEP-003	Underwriting model	E — Data	Halyard Data Partners	L6 Data	4	Dual-provider	Data Governance Lead	Chief Risk Officer	Medium	Active
DEP-004	Underwriting model	A — Financial	Northgate Frontier Models, Inc. (investor: Ashcombe Capital)	L1 Capital / L5 Models	3	Shared financial dependency	CFO Office	CFO	N/A	Monitoring
DEP-005	Enterprise search	D — Technology	Vantpoint Platform (built on Northgate models)	L7 Platform	3	Hidden concentration — converges on DEP-001	Head of AI Risk	CAIO	Medium	Flagged

RISK ASSESSMENT DEP-001 and DEP-005 present as separate vendor relationships but converge on the same upstream foundation-model provider. A material change at Northgate Frontier Models (outage, pricing, licensing, ownership change) would concurrently impact customer support automation and enterprise search. This is a concentration finding, not a hypothetical — see Concentration Analysis, Section 06.

Blank Reusable Template — Row Structure

The blank template mirrors the field structure in Section 03. Distribute as a spreadsheet or database table with one row per dependency; do not attempt to capture multiple dependencies in a single row even when they share a provider.

ID	Org.	Function	Use Case	Type	Provider	Layer	Crit.	Concentration	Risk Owner	Decision Auth.	Status

Dependency Risk Scoring Methodology

Each dependency is scored 1–5 across nine dimensions. A score of 1 represents minimal exposure; 5 represents maximum exposure. Scores are assigned by the risk owner and control owner jointly and must cite the evidence used to justify the score — a score without cited evidence is itself a governance gap (see Section 07).

Dimension	1 (Low)	3 (Moderate)	5 (Severe)	Scoring question
Criticality	Nonessential function	Important, workaround exists	Mission-critical, no workaround	Does the business function stop without this dependency?
Concentration	Many interchangeable providers	Handful of providers	Single provider, no alternative	How many providers can deliver this capability today?
Substitutability	Swap in days	Swap in a quarter	Effectively non-substitutable	Could this be replaced without a material outage?
Switching difficulty	Contractual only	Contractual + moderate rework	Deep technical/data lock-in	What does switching actually require operationally?
Financial exposure	Immaterial spend	Material but bounded spend	Spend concentration or revenue dependency	Does the relationship create material financial exposure?
Operational exposure	Contained blast radius	Cross-team impact	Enterprise-wide outage risk	How far does a failure propagate operationally?
Governance exposure	Clear, documented authority	Authority exists, partially documented	No clear owner or authority	Is there a named, accountable decision owner today?
Regulatory exposure	Unregulated function	Indirectly regulated	Directly regulated / high-risk classification	What regulatory regime attaches to this dependency?
Evidence quality	Full evidence on file	Partial evidence	Claimed only, no evidence	Is the governance claim actually supported by evidence?

Base Score

Base Score = simple average of the nine dimension scores, expressed on a 1–5 scale, then normalized to 100 for comparability across dependencies: **Base Score (100-pt) = (Σ of 9 scores ÷ 45) × 100**.

Weighted Score

A flat average treats a governance-exposure gap the same as a minor regulatory nuance. SBJ recommends weighting only where a dimension has asymmetric downside — i.e., a high score in that dimension alone can cause enterprise-level harm even if every other dimension scores low. This is a recommendation, not a mandatory feature; unweighted base scores are valid where an organization prefers a simpler model.

Dimension	Recommended weight	Rationale for the weight
Criticality	1.5×	Directly determines business continuity impact; understating it understates everything downstream
Concentration	1.5×	Systemic risk research consistently identifies concentration, not any single vendor's quality, as the primary source of cascading failure
Governance exposure	1.3×	An undefined decision authority means every other control can silently lapse without detection
Evidence quality	1.3×	A claimed-but-unevidenced control provides no actual risk reduction; weighting protects against false assurance
All other dimensions	1.0×	No identified asymmetric downside beyond ordinary linear risk

Weighted Score (100-pt) = (Σ of weighted scores ÷ Σ of max possible weighted points) × 100. Max possible weighted points = (5 × 1.5) + (5 × 1.5) + (5 × 1.3) + (5 × 1.3) + (5 × 5 × 1.0) = 39.5, scaled the same way as the base score.

Overall Dependency Risk Score — Thresholds

80–100 Severe	60–79 High	40–59 Moderate	20–39 Low	0–19 Minimal

Thresholds are an SBJ-recommended starting point, not a regulatory standard. Organizations with lower risk tolerance (e.g., regulated financial services) should compress the bands, moving "High" to 45+.

Concentration Analysis

Concentration analysis exists to surface dependencies that look distributed on the surface but are not underneath. The most consequential finding this instrument can produce is not "we depend heavily on Vendor X" — it is "we believed we had three vendors and actually have one."

Concentration Classifications

Classification	Definition
Single-provider dependency	Exactly one provider is capable of delivering this capability to the organization today
Dual-provider dependency	Two providers are viable; concentration risk reduced but not eliminated
Highly concentrated layer	A single ecosystem layer (e.g., L4 Compute) where the organization's entire portfolio routes through one or two upstream providers
Shared infrastructure dependency	Multiple, seemingly distinct organizational dependencies terminate at the same infrastructure provider (chip, data center, cloud region)
Shared model dependency	Multiple applications or use cases are built on the same foundation model, directly or via a downstream platform
Shared cloud dependency	Multiple dependencies route through the same cloud provider or the same specific region/availability zone
Shared semiconductor dependency	Multiple compute dependencies ultimately rely on the same chip architecture or fabrication source
Shared data dependency	Multiple models or use cases are trained or fine-tuned on overlapping data sourced from the same provider
Shared financial dependency	Multiple providers in the organization's portfolio share a common investor, parent company, or financing source

Hidden Convergence — Worked Illustration

ILLUSTRATIVE, FICTIONAL Three business units each independently select what they believe is a distinct AI vendor:

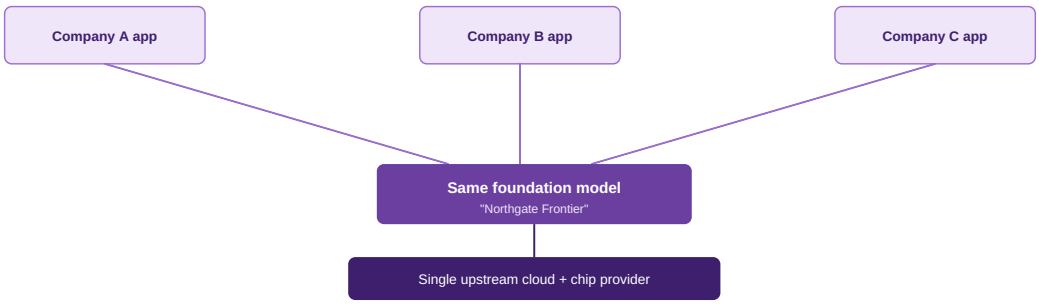

RISK ASSESSMENT An outage, price change, licensing change, or ownership change at "Northgate Frontier" or its underlying cloud/chip provider affects Companies A, B, and C simultaneously, despite each having procured through a different application vendor. This is the pattern the Concentration Map exists to detect. Classification: Shared model dependency + shared infrastructure dependency.

Decision Authority & Governance View

Every material dependency requires seven questions answered by name, not by title placeholder: who owns the relationship, who approves its use, who accepts the risk, who monitors it, who can terminate it, who can approve an exception, and who receives escalation. A dependency without a named answer to all seven is, by definition, ungoverned.

RACI — Standard Governance View

Activity	Business Owner	Risk Owner	Control Owner (e.g. CAIO/CISO)	Procurement	Governance Committee
Relationship ownership	A	C	C	R	I
Approve use	R	A	C	I	I
Accept residual risk	C	A	R	I	I
Ongoing monitoring	I	R	A	I	I
Terminate relationship	C	C	C	R	A
Approve exception	I	C	C	I	A
Receive escalation	I	I	R	I	A

R = Responsible · A = Accountable · C = Consulted · I = Informed. Role titles above are illustrative placeholders — substitute the organization's actual named roles before use.

Governance Escalation Path

Escalation triggers: dependency risk score enters "High" or "Severe" band · substitutability changes · provider ownership or financial event · evidence gap identified

Evidence Layer & Evidence Register

A governance decision without evidence is an opinion. This layer connects every decision to a rationale, an evidence artifact, a control, and a review cycle — and flags the gap explicitly when that chain is broken.

1. **Decision** — what was decided about this dependency (accept, mitigate, remediate, terminate)
2. **Decision owner** — the named individual accountable for the decision
3. **Decision date** — when the decision was made
4. **Decision rationale** — the reasoning, in writing
5. **Supporting evidence** — the artifact(s) that justify the rationale
6. **Control** — the control implemented as a result
7. **Review** — the cadence and next review date
8. **Outcome** — what the review found; feeds back into the register

Required Evidence Types by Dependency Category

Category	Required evidence
A — Financial	Contract / financing terms, investment disclosure, revenue concentration analysis
B — Ownership	Corporate structure documentation, equity/ownership disclosure, M&A notices
C — Infrastructure	SLA, uptime history, security assessment, business continuity attestation
D — Technology	Model documentation / model card, API terms, version change log
E — Data	Data provenance documentation, data processing agreement, data quality assessment
F — Operational	Vendor risk assessment, service contract, incident history
G — Governance	Decision record, RACI assignment, escalation log
H — Evidence	Audit trail confirming the above artifacts exist, are current, and are retrievable on demand

> **GOVERNANCE FLAG** Any dependency where a decision record exists but the cited "Supporting Evidence" field is blank, expired, or unretrievable should be flagged **Evidence Gap** in Status, regardless of how the dependency otherwise scores. Evidence Gap findings route directly to the Governance Committee.

Locate → Assign → Record → Scale

The Dependency Map is operated through four repeatable steps. Each step maps to a specific artifact within this instrument, so the operating model is not aspirational language — it is a literal instruction for which section to use and when.

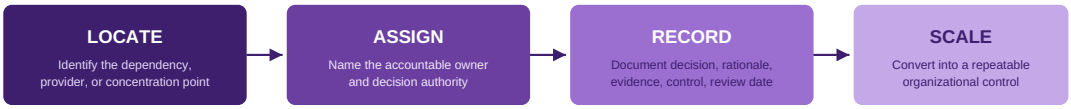

Step	Instrument artifact used	Output produced
Locate	Executive Ecosystem Map + Concentration Map	A logged dependency, tagged by layer and concentration classification
Assign	Decision Authority Matrix (RACI)	A named business owner, risk owner, control owner, and escalation path
Record	Dependency Register + Evidence Register	A completed register row with decision rationale and cited evidence
Scale	Governance workflow + Signal alignment	The dependency type becomes a standing control applied to all future instances of that pattern

A dependency is not considered governed until it has passed through all four steps. Partial completion (e.g., located and assigned, but not recorded with evidence) should carry Status = "In Progress," not "Active."

SIGNAL Framework Alignment

The Dependency Map is the primary data source feeding SBJ's Signal methodology. The five Signal phases are not a separate exercise — they are how the register, once populated, gets scored and matured over time.

Phase	Signal activity	How the Dependency Map feeds it
SCAN	Identify the ecosystem and dependencies	Executive Ecosystem Map and initial register population (Locate step)
CALIBRATE	Determine materiality, criticality, concentration, exposure	Dependency Risk Scoring (Section 05) and Concentration Analysis (Section 06)
ASSESS	Evaluate risk, substitutability, controls, conflicts, evidence	Base/weighted scores, Substitutability and Switching Difficulty fields, Evidence Register
ALIGN	Assign authority, controls, accountability, remediation	Decision Authority Matrix, RACI, Remediation Owner / Due Date fields
EVOLVE	Monitor changes and reassess continuously	Monitoring Frequency, Last/Next Reviewed fields; 90-day decay applies to certification claims built on this register

GOVERNANCE RECOMMENDATION This instrument is designed to be SBJ's structural companion to Signal: Signal scores an organization's AI governance program; the Dependency Map is the underlying evidence base that makes several Signal criteria (concentration awareness, vendor governance, evidence quality) verifiable rather than claimed.

Critical Dependency Dashboard — Reporting View

The dashboard is the executive-facing rollup of the register, refreshed on the Monitoring Frequency defined per dependency.

Metric	Definition
Total dependencies logged	Count of register rows with Status ≠ blank
Critical dependencies (score ≥ 60)	Count and % of total in the High or Severe risk band
Single-provider concentration count	Count of dependencies classified Single-provider or Highly concentrated layer
Evidence Gap count	Count of dependencies flagged Evidence Gap (Section 08)
Overdue reviews	Count where Next Review date has passed
Unassigned governance	Count of dependencies missing a named Risk Owner or Decision Authority

Case Study (Hypothetical)

FICTIONAL All organizations below are invented for illustration. No relationship, financial position, or event described is real or attributed to any actual company.

Cast: Ashcombe Capital (investor) · Halyon Semiconductor (chip provider) · Meridian Cloud Services (hyperscaler) · Northgate Frontier Models (foundation model provider) · Vantpoint Platform (AI application platform) · Baytree Data Centers (colocation operator) · Corvale Financial Group (enterprise deploying AI).

Scenario. Corvale Financial Group deploys two AI-enabled products: a fraud-detection tool procured from Vantpoint Platform, and a customer-service assistant built directly on Northgate Frontier's API. Two procurement teams ran two separate vendor reviews and concluded Corvale had diversified its AI vendor exposure.

What the Dependency Map found. Vantpoint Platform is built entirely on Northgate Frontier's foundation models under a reseller agreement (technology dependency, D). Northgate Frontier's training and inference infrastructure runs exclusively on Meridian Cloud Services (infrastructure dependency, C), hosted in Baytree Data Centers' facilities. Meridian's AI-optimized capacity runs on Halyon Semiconductor chips under a multi-year supply agreement. Separately, Ashcombe Capital holds an equity stake in both Northgate Frontier and Halyon Semiconductor (financial and ownership dependency, A/B) — a relationship visible in Northgate's public financing disclosures.

RISK ASSESSMENT Corvale's fraud-detection and customer-service products — procured through two different vendors, reviewed by two different teams — both terminate at the same foundation model, the same cloud provider, the same data-center operator, and the same chip supply chain. A material event at any one of Northgate Frontier, Meridian Cloud Services, Baytree Data Centers, or Halyon Semiconductor affects both products simultaneously. The common investor relationship (Ashcombe Capital) additionally creates a shared financial dependency across two nominally separate technology dependencies.

Register Outcome

ID	Use case	Direct provider	True upstream convergence	Concentration classification
DEP-101	Fraud detection	Vantpoint Platform	Northgate → Meridian → Baytree → Halyon	Shared model + shared infrastructure
DEP-102	Customer service assistant	Northgate Frontier (direct)	Meridian → Baytree → Halyon	Shared infrastructure

Governance recommendation issued. RECOMMENDATION Corvale's Governance Committee should treat DEP-101 and DEP-102 as a single concentration cluster for continuity planning purposes, not as two independent vendor relationships, and require a documented failover or degradation plan covering both products jointly.

Executive Interpretation

For a governance committee or board audience, the register and scoring detail should distill to four questions, answerable in a single sitting from the Critical Dependency Dashboard:

- **Where are we exposed?** — the count and identity of Severe and High-band dependencies, and which business functions they support.
- **Where are we more concentrated than we think?** — the output of the Concentration Analysis, specifically any Hidden Convergence findings.
- **Is anyone actually accountable?** — the count of dependencies with an unassigned Risk Owner or Decision Authority.
- **Is our governance claim backed by evidence?** — the count of Evidence Gap findings, since an unevidenced control is not a control.

> A committee that can answer all four questions with named individuals, current dates, and cited evidence has an operating governance program. A committee that answers with role titles, general assurances, or "we believe so" has a governance intention, not a governance program — and the Dependency Map exists specifically to make that distinction visible.

Reading the Risk Score in Committee

A single dependency's Overall Dependency Risk Score should never be read in isolation from its Concentration Classification and Evidence status. A Moderate-scoring dependency with an Evidence Gap and Single-provider concentration is a more urgent finding than a High-scoring dependency with full evidence and Dual-provider concentration — the score is a starting point for discussion, not a final verdict.

Implementation Instructions

Phase 1 — Locate (Weeks 1–3)

- Inventory every AI system, model, and platform currently in production or pilot, by business function
- For each, trace the provider chain upward through infrastructure, model, and data layers using vendor documentation, contracts, and technical architecture diagrams
- Open a register row (Dependency Register, Section 03) for every distinct dependency identified

Phase 2 — Assign (Weeks 2–4, overlapping)

- Name a Risk Owner and Control Owner for each dependency; do not proceed with a role title as a placeholder past this phase
- Apply the RACI structure (Section 07) and confirm escalation authority
- Route any dependency with no willing or obvious owner to the Governance Committee for assignment

Phase 3 — Record (Weeks 3–6, overlapping)

- Score each dependency across the nine dimensions (Section 05); require cited evidence for every score
- Complete the Evidence Register fields; flag Evidence Gaps immediately rather than backfilling scores to avoid the flag
- Run the Concentration Analysis across the full register to surface hidden convergence

Phase 4 — Scale (Week 6 onward, continuous)

- Set Monitoring Frequency and Next Review dates per dependency, weighted toward shorter cycles for Severe/High-band items
- Stand up the Critical Dependency Dashboard as a standing governance committee agenda item
- Convert recurring findings into standing procurement and architecture controls (e.g., a policy against new Single-provider dependencies without an explicit exception)

> **GOVERNANCE RECOMMENDATION** Treat the first full pass through Locate → Assign → Record as a fixed-scope project with an end date. Treat Scale as permanent operating cadence, not a project phase — concentration and dependency exposure change continuously as providers raise capital, get acquired, or change infrastructure.

SECTION 14

Classification Standard & Methodology Notice

Every statement in a completed Dependency Map must carry one of the following classifications. Analysts populating the register should tag each field accordingly; unclassified assertions should be treated as unverified.

Classification	Definition
FACT / VERIFIED RELATIONSHIP	Confirmed via primary source (contract, SEC filing, technical architecture record, direct vendor confirmation)
REPORTED / ORGANIZATIONAL CLAIM	Stated by the provider or a secondary source but not independently verified by the organization
ANALYTICAL INFERENCE	A reasoned conclusion drawn from available evidence, explicitly not a confirmed fact
RISK ASSESSMENT	A judgment about exposure or consequence, distinct from a factual claim about the relationship itself
GOVERNANCE RECOMMENDATION	SBJ or internal guidance on how to act; explicitly not a regulatory requirement

STANDING REQUIREMENT Analytical inferences must never be recorded or presented as verified relationships. Financial relationships, ownership relationships, dependencies, and regulatory requirements must never be fabricated to complete a register row; an incomplete row with a documented "unable to verify" note is more useful, and more defensible, than a fabricated one.

SBJ Original Methodology Notice

The AI Ecosystem Dependency Map, its field structure, scoring methodology, concentration taxonomy, and Locate → Assign → Record → Scale operating model are original governance methodology developed by Solutions by Jewel. This instrument is a proprietary advisory framework and is not an official regulatory standard, government requirement, or third-party certification. Alignment references to the Signal methodology refer to SBJ's own framework, also proprietary to Solutions by Jewel.

This instrument is designed for use during AI risk assessments, vendor assessments, AI governance audits, enterprise AI implementations, third-party risk reviews, procurement reviews, AI infrastructure planning, executive risk reviews, and governance committee meetings.