

U101 UNDERSTAND · FOUNDATION

Technology Portfolio Intelligence

Building an evidence-based model of your applications, infrastructure and dependencies that stays current.

FORMAT


Knowledge pack · 13 slides

FOR

Architects and portfolio managers

LEVEL

Foundation



THE PREMISE

A portfolio model maintained continuously from evidence

A survey-based inventory begins to go out of date while it is still being compiled. Portfolio intelligence is the capability to answer questions about the estate from **current evidence**, whenever they are asked.

Current

Refreshed from the systems themselves, by scan and event.

Evidenced

Every attribute carries its source and its evidence grade.

Connected

Applications linked to capabilities, data, infrastructure and cost.

PURPOSE

Eight questions it must answer on demand

1

What do we run?

Every application, service and platform, including unregistered SaaS and AI agents.

2

Who owns it?

An accountable business owner and a technical owner.

3

What does it do?

The capabilities and processes it supports.

4

What does it depend on?

Upstream and downstream systems, data and identities.

5

What does it cost?

Fully loaded run cost, not license fees alone.

6

What data does it hold?

Data classes, sensitivity and retention duties.

7

How exposed is it?

Vulnerabilities, end-of-support components, privileged access.

8

What breaks if it changes?

The blast radius of a change, migration or outage.


EVIDENCE

What each evidence source sees and misses

SOURCE	SEES WELL	MISSES
CMDB and service catalog	Declared ownership and service relationships	Drift and unregistered systems
Discovery and cloud APIs	What runs, where, and how it is configured	Business purpose, ownership and SaaS
Network flows	Actual communication between systems	Owning application, payloads, jobs outside the window
Source code and builds	Frameworks, libraries and internal dependencies	Deployed versions, configured endpoints
Identity and access logs	Who signs in to what, and how often	Anything outside SSO, and activity after sign-in
Finance and contracts	Cost, vendors and renewal dates	Deployment and actual use

METHOD

From raw evidence to a trusted record



CONFIDENCE

Grade the evidence behind every attribute

Each decision should state the grade it requires. Retiring an application needs observed usage data over a full business cycle.

Unknown

No evidence yet.
Recorded as a gap to close.

1

Inferred

Derived from names,
patterns or adjacent
records.

2

Declared

Stated by an owner or
entered manually in the
CMDB.

3

Observed

Seen directly by
discovery, telemetry,
network flows or access
logs.

4

Confirmed

Two independent sources
agree, at least one of
them observed.

5

STRUCTURE

Model the estate in connected layers

Capabilities

What the organization does, independent of systems. It changes least often.

Applications

What delivers each capability, with owners and lifecycle stage.

Components and data

Code, libraries, databases, interfaces and data classes.

Infrastructure

Hosts, containers, platforms, networks and cloud accounts.

Identities

People, service accounts, AI agents and the privileges that link them to systems.

WHY LAYERS

Most questions cross layers. Finding which capabilities rely on an end-of-support database means following the links.

DEPENDENCIES

Six kinds of dependency to map

Observation must cover a full business cycle. Quarter-end and year-end processes stay invisible to a short sample.

Runtime calls

APIs, queues and service-to-service traffic.

Shared data

Common databases, file transfers and extract jobs.

Identity

Service accounts, single sign-on and shared credentials.

Infrastructure

Shared hosts, clusters, DNS, certificates and networks.

Schedules

Batch jobs that run monthly, quarterly or once a year.

People

Knowledge held by only a few individuals.

USE

Decisions the model supports

DECISION	THE QUESTION IT ANSWERS
Rationalization	Where capabilities overlap, and what is unused or redundant
Modernization	What depends on a system, and how hard it is to change
Security	What is exposed, and what an attacker could reach from it
Cost	What each capability costs to run, and who consumes it
Resilience	Which paths are critical, and where single points of failure sit
Mergers and acquisitions	Where two IT environments overlap, and what a carve-out must separate

AI AND AGENTS

What AI and agents change

ROLE OF AI	WHAT CHANGES	WHAT TO REQUIRE
Part of the estate	Models, agents and AI features in SaaS products need records, as any application does	Owner, purpose, data accessed, tools called and the identity each agent uses
Evidence reader	Language models extract owners, terms and purpose from contracts, tickets and runbooks	Grade the output Inferred until evidence or the owner confirms it
Entity matcher	Models match one application across sources that name it differently	Keep each match with its reason, and review a sample before relying on it
Consumer of the model	Agents that assess change impact or triage incidents query the model, often through MCP servers	Return each fact with its grade and last-verified date; limit what each agent can read

MEASUREMENT

Measures of record quality and use

Coverage

Share of applications seen directly by at least one automated source.

Freshness

Median and 90th-percentile time since key attributes were last verified.

Confidence

Share of key attributes graded observed or confirmed.

Ownership

Share of applications whose owner confirmed the record in the last six months.

Dependency coverage

Share of critical applications whose upstream and downstream links are observed.

Time to answer

Time to answer the eight questions for one application.

PITFALLS

Why inventories fail

✗ Starting with surveys

✓ Start with evidence, then ask owners to confirm

✗ Waiting for completeness

✓ Make it useful for one decision first

✗ Running it as a project

✓ Run it as a standing capability with an owner

✗ Treating the CMDB as truth

✓ Treat it as one source among several

✗ Sampling a short window

✓ Observe a full business cycle

✗ Collecting without using

✓ Tie every attribute to a decision that needs it

SUMMARY

Key takeaways

- 01 Refresh the record continuously from evidence; use surveys only to confirm it.
- 02 Reconcile many sources and grade every attribute by the evidence behind it.
- 03 Anchor the model on business capabilities and follow dependencies across layers.
- 04 Observe a full business cycle before declaring anything unused.
- 05 Measure coverage, freshness, confidence and the time it takes to answer a question.

NEXT COURSE · U201

AI-Ready Data

Understand · Practitioner