

L101 LEAD · FOUNDATION

The AI-Native IT Organization

How the IT function changes when software agents carry a growing share of the work, and how to lead that change.

FORMAT

Knowledge pack · 14 slides

FOR

IT leadership teams

LEVEL

Foundation



THE PREMISE

An AI-native IT function redesigns work, decisions and controls

Adding a copilot to each workflow can speed up individual tasks. Redesigning the workflow around what agents can now do reliably, and around what still requires **human judgment**, changes what the IT function can deliver.

Work

Tasks agents execute are separated from decisions people own.

Knowledge

Agents work from a current model of the estate, maintained from live evidence.

Control

Policy is checked inside the workflow, before an agent acts.

WHAT THE EVIDENCE SAYS

AI use is widespread, while measured gains vary

DORA finds that AI acts as an amplifier, magnifying an organization's existing strengths and weaknesses.

90%

of technology professionals surveyed by DORA reported using AI at work, and more than 80% believed it had increased their productivity.

DORA, State of AI-assisted Software Development, 2025

19%

longer to complete tasks where AI was allowed, for 16 experienced open-source developers. Afterwards they estimated AI had cut their time by 20%.

METR randomized controlled trial, July 2025

30%

of DORA respondents reported little or no trust in AI-generated code. DORA linked AI adoption to higher delivery throughput but lower stability.

DORA, State of AI-assisted Software Development, 2025

In February 2026, METR called its newer data “an unreliable signal”, as developers unwilling to work without AI opted out, but believed developers are likely more sped up than in early 2025. Measure gains in your own workflows.

THE SHIFT

Six shifts define the AI-native IT function

FROM

Projects staffed by people

Documentation written after the fact

Periodic assessments and audits

Tickets routed to queues

Controls applied at review gates

Headcount as the limit on capacity



TO

Outcomes delivered by teams of people and agents

Context maintained continuously from code, telemetry and configuration

Continuous evaluation that supplies audit evidence


Issues resolved where they are detected, escalated by exception

Controls enforced in the workflow, with a complete audit trail

Review capacity, context quality and trust as the limits

DELEGATION FRAMEWORK

Decide what to delegate by consequence and reversibility



ARCHITECTURE OF THE WORK

Work moves across four layers

Intent

Outcomes, priorities and risk appetite, set by accountable leaders.

Decisions

Architecture, portfolio and risk acceptance, owned by people and informed by agents.

Execution

Changes, remediation, migrations and routine operations, increasingly performed by agents under policy.

Context

A current, evidence-based model of applications, infrastructure, code, data flows, identities and cost.

THE ASSET

Every workflow delegated later reuses this layer. DORA names AI-accessible internal data as one of seven capabilities that increase the benefits of AI.

PEOPLE

Roles shift toward specification, review and exceptions

New responsibilities appear in every team: owning each agent, evaluating it, managing its identity and curating the context it uses.

Engineers

Specify, review and test agent output, and own the design decisions that require judgment.

Architects

Maintain architecture decision records and the guardrails within which agents operate.

Operations

Turn runbooks into executable policy and handle the exceptions agents escalate.

Leaders

Own outcomes across human and agent work, and decide what is delegated.

DECISION RIGHTS

Write decision rights down before agents act

DECISION	AGENT MAY	APPROVAL	EVIDENCE RETAINED
Classify and route an incident	Act	None; outcomes sampled	Classification, confidence, outcome
Change an agent’s model, prompt or tools	Propose only	Agent owner, after evaluation passes	Versions and evaluation results
Patch a production system	Act after approval	Service owner or standard change	Change record, test results, rollback plan
Retire an application	Prepare evidence and the runbook	Business owner and architecture	Usage, dependencies, retention decision
Grant or elevate access	Provision approved requests only; never approve	Identity governance	Request and approval chain

CONTROL PLANE

One control plane for every agent

Decision rights hold only when they are enforced the same way for every agent, whichever team built it.

- ✓ **Register**
Each agent listed with its purpose, accountable owner, model, tools and decision rights.
- ✓ **Policy**
Decision rights checked at every tool call, outside the model.
- ✓ **Trace and cost**
Every action logged with its requester, agent, tool, result and cost.
- ✓ **Identity**
Its own non-human identity, with short-lived credentials scoped to the task.
- ✓ **Evaluation**
A versioned test set, rerun before release and after any model, prompt or tool change.
- ✓ **Stop**
A tested way to halt any agent and revoke its access at once.

MATURITY

Five stages from assisted to native

Moving to Delegated requires current context, controls enforced in the workflow and shadow results that meet thresholds.

Assisted

Individuals use copilots. Gains are personal and unmeasured.

1

Augmented

Teams standardize tools and review practices, and measure each workflow.

2

Delegated

Agents complete bounded tasks end to end under policy. People approve by exception.

3

Orchestrated

Agents coordinate across workflows through shared context and one control plane.

4

Native

Operating model, funding and metrics assume teams of people and agents.

5

MEASUREMENT

Six measures for delegated work

Cycle time

From request or detection to a verified change in production.

Change fail rate

Share of deployments that need immediate intervention, such as a rollback or hotfix.

Autonomous completion

Share of work agents finish without human intervention, with sampled quality checks.

Correction rate

Share of agent output that people correct or revert.

Cost per outcome

Fully loaded, including model, compute and review effort.

Context coverage

Share of the estate whose records were verified against live data within a set period.

AVOID Licenses activated, prompts per user and lines of code generated count activity; report them only beside outcomes.

PITFALLS

Six common failure patterns in IT's AI programs

✗ Buying tools before redesigning the work

✓ Redesign three workflows end to end, then choose tools

✗ Reporting adoption as success

✓ Report cycle time, quality and cost per outcome

✗ Granting agents broad standing access

✓ Scoped, short-lived permissions for each task

✗ Letting every team build its own stack

✓ One control plane for identity, logging, evaluation and policy

✗ Ignoring the quality of context

✓ Invest first in a current model of the estate

✗ Delegating without a shadow period

✓ Compare agent and human results before removing approvals

GETTING STARTED

A 180-day plan

DAYS 1–30

Choose and baseline

- ◆ Select three high-volume, reversible, measurable workflows
- ◆ Baseline cycle time, quality and cost
- ◆ Write decision rights and thresholds for each workflow

DAYS 31–90

Build the foundations

- ◆ Stand up agent identities, logs, evaluation and cost tracking
- ◆ Build current context for the systems in scope
- ◆ Run agents in shadow mode against human results

DAYS 91–180

Delegate and scale

- ◆ Approve by exception where shadow results meet thresholds
- ◆ Publish results against the baseline
- ◆ Extend to the next workflows; adjust roles and funding

SUMMARY

Key takeaways

- 01 AI-native IT redesigns work around what agents do reliably and what still needs human judgment.
- 02 Delegate by consequence and reversibility, and write decision rights down before agents act.
- 03 Current, evidence-based context and enforced controls make delegation safer.
- 04 Measure cycle time, stability, quality and cost per outcome; treat usage as context only.
- 05 Start with three workflows, one control plane and a published baseline.

NEXT COURSE · L201

IT Economics: Cost, Value and AI Spend

Lead · Practitioner