Digibee

AI in Enterprise Integration

When it’s the smart solution, when it isn't, and what "AI-Native" integration actually means

EXECUTIVE SUMMARY

The short version

Coding agents like Claude Code are good at throwaway integration work: one-time migrations, well-documented systems, low-stakes scripts. They fall apart on enterprise integration necessities: scheduled runs, retry logic, audit trails, credential management, monitoring that catches silent failures. The gap isn't a maturity problem that better models will fix. It's structural. When a model hits its "lost in the middle" limit on a long context window, it falls back on training data and produces code that passes tests and fails quietly in production, sometimes six months later.

We present a decision framework for sorting backlog items by what they actually need, not by which tool is trendiest. Deterministic logic wins when failure has zero tolerance, when a regulator might ask for an audit trail, or when a competent developer could write the rules in an afternoon. Agentic AI earns its place when inputs are unpredictable or the task needs judgment a rule set can't encode. Most work falls in between: a "deterministic-plus" pattern where a governed pipeline does the heavy lifting and an LLM handles one bounded step, like flagging an invoice exception or triaging a helpdesk ticket. Most of the backlog still belongs on the deterministic end. Adding AI to something that isn't broken doesn't make it better; it makes it harder to govern.

In the last section, we explore what Digibee built to be "AI-native." This included building a new UI specifically for agents, and injecting integration knowledge into specialized agents. We also offer a sniff test: ask any integration vendor whether they rebuilt for agents or just added a chatbot.

THE LIMITS

Why coding agents can't carry enterprise integration—and why the gap is structural.

THE FRAMEWORK

Sort the backlog by need: deterministic, agentic, or deterministic-plus.

AI-NATIVE

What a genuine rebuild for agents takes, and the sniff test for vendors.

CONTENTS

What's inside

00 **Introduction** 04
Deciding where AI fits and where it doesn't

01 **Why coding agents can't handle enterprise integration** 06
Where generic agents fail, and why the gap is structural

What coding agents are actually good at 07
Why the gap is structural 08
What skills fix, and what they can't 13

02 **AI isn't all or nothing** 14
A framework for matching each task to the right solution

Start with outcomes, not solutions 15
The deterministic → agentic spectrum 18
Deterministic plus 21

03 **What "AI-native" actually means** 23
Rebuilding a platform from zero for a digital workforce

The engineering work behind the label 24
Human collaboration & higher altitude 26
The sniff test 29

→ **See the platform in action** 30

The real work isn't adopting AI. It's deciding where it fits.

Boards, CEOs, and your C-suite peers are pushing AI-first mandates, and stakeholders believe agents can solve any software problem. The flood of AI initiatives keeps expanding the integration backlog, and the integration team can't keep up.

Integration specialists have been handed generic AI tools designed for greenfield software development and told to move faster. In integration settings, those tools offer little help. In this guide, we examine the tension between integration and AI in three parts.

Demand keeps scaling

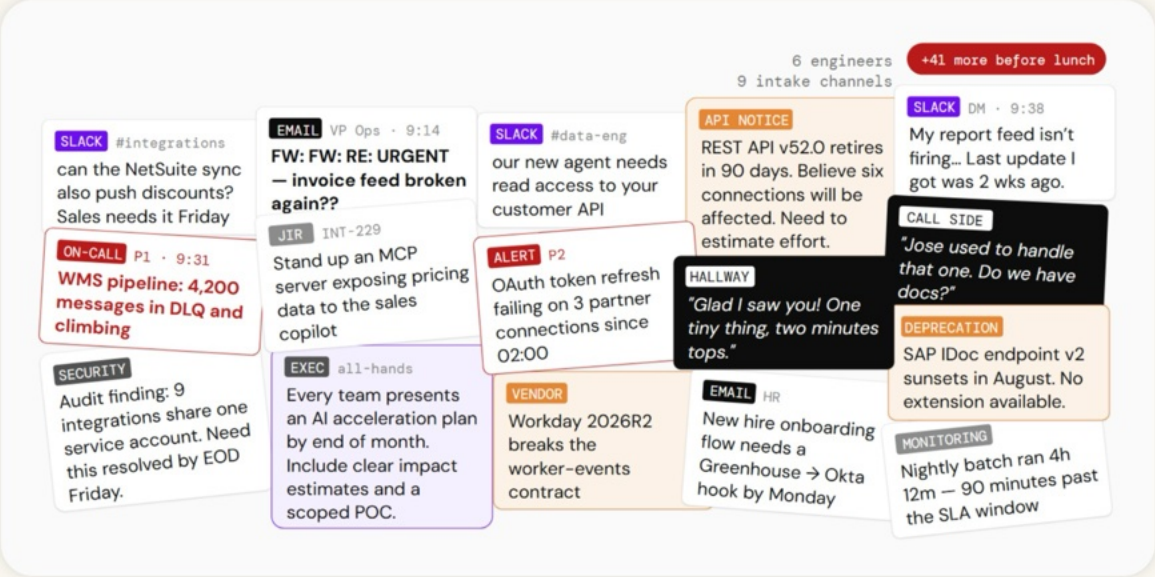

Generic coding agents fail

- RISK Training AI on patterns, best practices
- RISK Safe system and data access
- RISK Governed decisions
- RISK Ease of team collaboration
- RISK Efficiency through re-use
- RISK Standardizing organizational protocols
- RISK Support SLAs for integration use cases

THE DELIVERY GAP

Demand keeps rising. Manual capacity doesn't.

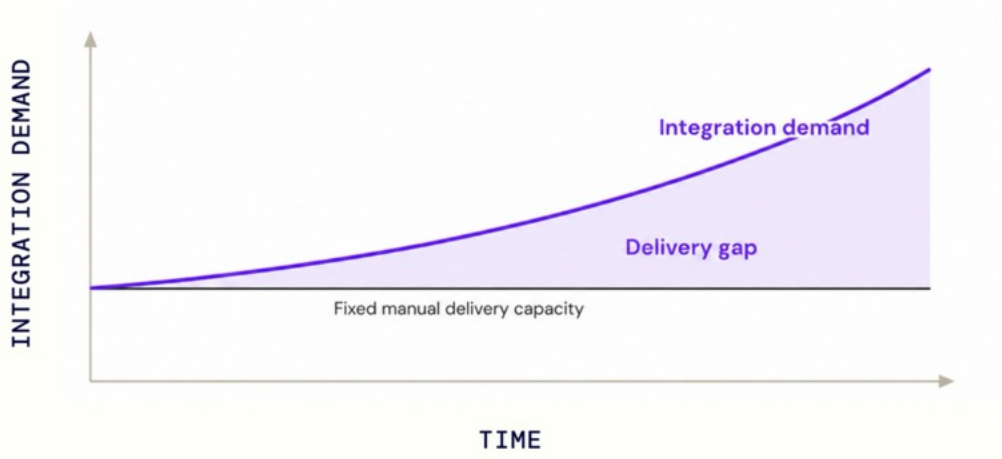

CH. 01

Where agents fail

Where generic coding agents break down on enterprise integration, and why the gap is structural, not temporary.

CH. 02

A decision framework

How to match each task to the right solution: deterministic, agentic, or "deterministic-plus."

CH. 03

Building AI-native

Our journey building an AI-native platform, and the challenges we had to solve to get there.

Most organizations are treating AI adoption as binary.
The real work is deciding where AI fits and where it doesn't.

CHAPTER 01

Why coding agents can't handle enterprise integration

Every team at Digibee uses Claude Code. We want to start there, because what follows is not a dismissal from people unfamiliar with the tools.

ENTERPRISE INTEGRATION IS A DIFFERENT CATEGORY OF WORK

Coding agents have changed how our engineers work, how fast we prototype, and how much cognitive load we save on routine tasks. If your teams aren't using them, you're leaving real productivity on the table. But their success in greenfield development created an assumption that they can be pointed at any software problem. Enterprise integration is a different category of work, with different failure modes that coding agents weren't designed to handle.

What coding agents are actually good at

Generic coding agents handle integrations well under specific conditions:

- The systems involved are well-documented.
- The task is one-time rather than recurring.
- Failure is low-stakes and recoverable.
- Nothing in production is at risk during testing.

A one-time migration from a REST API to a CSV? Go for it. A quick script to pull data from a public endpoint? Perfect. A throwaway ETL task with no downstream dependencies? Good use of the tool. The problems start the moment the integration faces more complex conditions.

Coding agents don't produce any of this. They produce code—code that has to live somewhere, be secured somehow, and be maintained by someone. Falling short on this can create expensive process or compliance failures.

THE FIT CHECK

Integrations: coding agent capabilities vs. enterprise needs.

CODING AGENTS WORK WELL WHEN

- The systems involved are well documented
- The task is one-time, not recurring
- Failure is low-stakes and recoverable
- Nothing in production is at risk during testing

One-time, low-stakes work: go for it.

ENTERPRISES NEED

- Reliable runs on a schedule or real-time trigger
- Retry logic and failure recovery
- Audit trail: what moved, when, and why
- Credentials managed in one place, not a dozen configs
- Monitoring that catches silent failures
- Maintainable by someone other than the builder

Generated code gives you none of these.

Why the gap is structural

It's tempting to frame this as a maturity argument: coding agents are new and will improve. But the limitations are structural, not incidental. Enterprise integration depends on accumulated knowledge encoded in pre-built connectors that account for how NetSuite sequences operations, how SAP handles idempotency, or how Salesforce structures its data model. A coding agent reasons about each of those nuances from scratch, every time. That's where subtle, production-only bugs are born.

Coding agents figure out every system from scratch

Enterprise integration targets are often poorly documented. That's the reality of legacy systems, accumulated data contracts, and decades of organizational complexity. Most of the behavior that matters only surfaces under load, and none of it is in the spec. Even when documentation exists, an agent may ignore it. Researchers have found that LLMs can lose track of content in the middle of a long context window—an effect sometimes called "lost in the middle." When that happens, the model falls back on its training data. The more obscure the package or API, the more likely the agent is to generate code that fails against it.

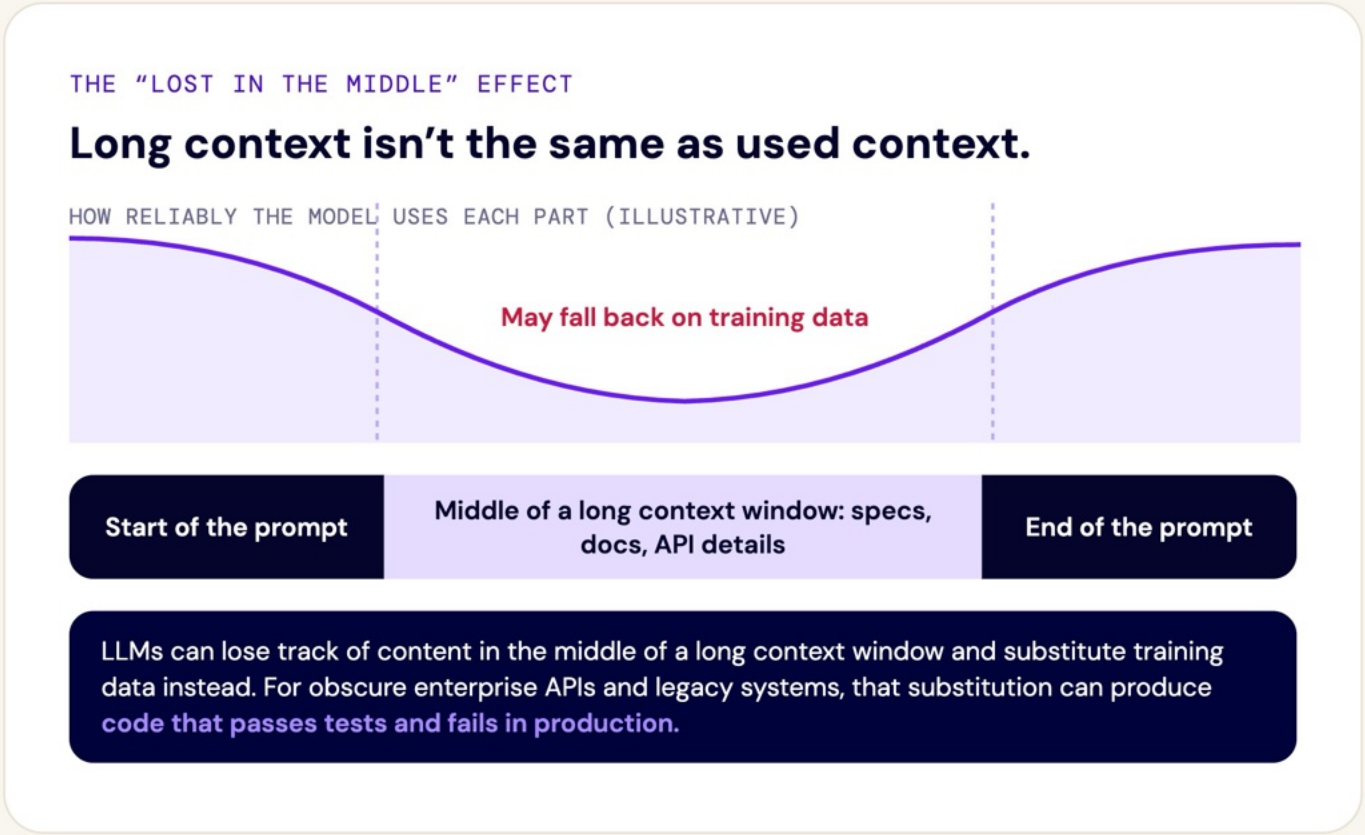

Agents don't improve on this automatically. They don't accumulate understanding across sessions. Agents parse complex sequencing logic and non-obvious idempotency rules anew each time. The code they produce is bespoke, hard to read, and difficult to govern consistently across an enterprise estate.

The worst outcome: bugs that don’t break loudly on deployment. They silently corrupt records or double-write transactions in edge cases that only surface under real load. The CIO learns about them six months later.

Coding agents build code. Integrations require more

Someone has to provision infrastructure, manage uptime, handle scaling, implement retry logic, build failure recovery, and create visibility into what happens in production.

You could prompt your way around each gap individually. But now you’re building an integration plus everything required to run it, each with its own maintenance burden.

Coding agents are also optimized to iterate fast, not to fail safely. When an integration fails in production, payments aren't processed, freight bids aren't made, or orders don't ship. An agent that generates business logic without logging, alerting, or circuit breakers gives you half a solution —the easy half.

WHAT'S MISSING

Code is the tip of the iceberg.

THE EASY PART

Generated code, the part you can see

- Infrastructure
- Scaling
- Retry logic
- Failure recovery
- Logging and alerting
- Circuit breakers

Someone still has to build, run, and maintain all of it.

APIs change. Tokens expire. Systems update. An agent has no relationship with what it built and no memory of why decisions were made. The integration lives in the head of whoever prompted it. When that person leaves (even for vacation) so does the rationale behind every design choice.

THE COMPOUNDING COST

"One coding-agent integration is manageable. A hundred of them—each with its own retry logic, credential handling, and error assumptions—is a hundred codebases to govern."

Coding agents don't own what they build

Technical debt is visible. Governance debt isn't—until it's a breach, an audit finding, or someone's last day.

No intermediate artifact, no way to audit

Generic coding agents generate code directly from a prompt. There's no intermediate artifact—no specification, no mapping document, no structured record of requirements, edge cases, or error behavior. That's a control gap, and it affects any team that needs to understand, audit, or change an integration after the fact.

Any integration that moves data between systems needs a reliable record of what moved, when, and why. Generated code doesn't produce it.

Credentials with no owner

API keys, OAuth tokens, and service credentials need to be stored, scoped, and rotated correctly. Generated code has no opinion on this and no ownership over it. Every custom integration becomes a new artifact to secure, a new attack surface, and a new dependency with no ownership model.

The cost compounds

A hundred integrations, each with its own retry semantics, its own error handling, its own credential assumptions, are a hundred codebases to secure, update, and maintain.

GOVERNANCE DEBT, ITEMIZED

- No specification or mapping document
- No owned model for credential rotation
- No structured record of edge cases
- No reliable record of what moved, and why

What skills fix, and what they can't

Teams using Claude Code often rely on "skills"—embedded example libraries and documented patterns that improve an agent's output on well-defined tasks. Applying skills to integration would mean researching how each target system behaves beyond its documentation: edge cases, rate limits, version-specific quirks. That's not one-time work; enterprise systems change, and any skill would need to evolve alongside them.

Even well-crafted skills don't eliminate the risk of buggy code. LLMs behave probabilistically. They might produce perfect code for one prompt; for the next, their output might pass tests and fail silently in production. Idempotency violations don't throw errors. Double-writes accumulate until reconciliation breaks several days later.

A skill can document how to respond to a failure, but it can't notice one. The agent doesn't run alongside the integration. A platform does: monitoring for silent failures, alerting before reconciliation breaks, and keeping the operational layer running between builds.

Great tools, but not for this job

Coding agents fall short on recurring, high-stakes, enterprise-scale integration. Not because the technology isn't impressive, but because the problem doesn't fit the tool. Knowing that, the question becomes: when does AI belong in integration at all, and when doesn't it?

Next: a framework for every item in your backlog →

CHAPTER 02

AI isn't all or nothing

How to anchor to outcomes—not solutions—when solving enterprise integration challenges. You have a year's worth of work in the backlog. How do you match each problem to the right solution?

Start with outcomes, not solutions

You have a year's worth of work in the backlog, you're racing to validate where AI fits, and you face real system, team, and data constraints.

How do you match each problem to the right solution?

Work through the backlog item by item. Some tasks fit deterministic workflows, where AI would only introduce risk, cost, and delay. Others present genuinely ambiguous problems where agentic AI unlocks automation that wasn't previously possible. A surprising number fall in between. It's natural to bias toward a preferred platform, language, or the latest technique. Ask first what success requires for the process.

SIGNALS TO BUILD DETERMINISTIC

Zero tolerance for failure

At 99% accuracy, a process that runs 100,000 times produces 1,000 errors. For some tasks, that cost is unacceptable.

Decision explainability

If a regulator or compliance body might ask why your system made an automated decision, stick with deterministic logic. It produces an audit trail. LLMs produce probabilistic outcomes that sometimes violate instructions. When they do, it's difficult or impossible to understand why.

Simple A-to-B logic

If a skilled developer could write the rules in an afternoon, write the rules. Leave simple things simple and save AI for harder problems.

THE MATH OF "99% ACCURATE"

At 99% accuracy, a process that runs 100,000 times is wrong 1,000 times. Imagine the consequences of 1,000 failed paychecks.

SIGNALS THAT AGENTIC AI ADDS VALUE

Unpredictable inputs

If the project calls for ingesting documents in varying formats, natural language requests, or other unstructured data sources, LLMs are typically the only reasonable solution.

Contextual, run-time decision-making

An LLM can reason across ambiguous inputs in ways rule sets cannot. If the system needs to employ judgment, it needs AI.

Considerations for unclear cases

Throughput and latency

High-volume processes with tight timing should lean deterministic. AI inference adds latency and cost at scale.

Cost predictability

Agentic workflows carry variable operational costs at any scale. Model the costs carefully before committing.

Total cost of ownership

Code pipelines cost more in employee time upfront; a quickly-built agent can cost more over time. Is a pipeline that takes hours to build but fails 2% of the time cheaper than one that takes a week and never fails? It depends.

THE RUNAWAY-COST WARNING

The $47,000, 11-day agent spiral was an extreme case. But agentic workflows carry variable operational costs at any scale. If budget predictability matters, model it before committing. (Source)

THE RUNAWAY-COST LOOP

Two agents kept asking each other for more.

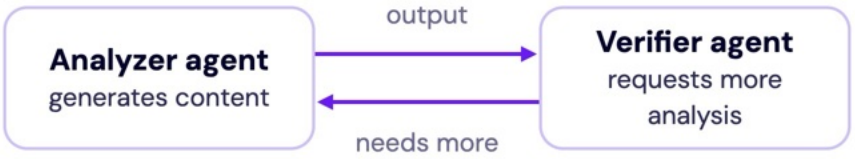

No step limit

No cost ceiling

No alert acted on

THE BILL

$47,000

in 11 days, before anyone noticed

The three guardrails that were missing

Cap the steps

Set a hard limit on how many rounds agents can loop.

Set a cost ceiling

Stop the run when spend crosses a budget you chose.

Act on alerts

Route cost alerts to someone who can pause the workflow.

The spectrum

Modern integration work sits on a spectrum from deterministic to agentic. In between lies a wide middle ground where deterministic foundations are enhanced by targeted agentic steps. This framework applies at the individual workflow level and across an entire project portfolio.

Matching the task to the solution

	Deterministic	Deterministic Plus	Agentic
USE WHEN	Requirements are stable, inputs are structured, governance matters, failure isn't an option.	A governed workflow could deliver more value with a targeted agentic step; edge cases or outputs benefit from AI judgment.	The problem requires judgment, synthesis, or creativity; some output variability is acceptable.
COMMON EXAMPLES	Password recovery workflows Compliance audit log exports Bank fraud flag escalations	IT helpdesk routing with AI-assisted triage notes Invoice processing with LLM exception flagging Release notes from structured commit data	Contract summarization & risk flagging RFP response drafting Generating social content from a prompt
WATCH OUT FOR	Underestimating AI's value on edge cases and unstructured inputs.	Letting the agentic step expand beyond its bounded role.	Variable costs at scale; audit requirements; silent failure modes.

Today, most organizations we work with find the majority of their backlog belongs at or near the **deterministic end.** That's appropriate; these are proven, solved integration goals. Agentic workflows open up new automation opportunities, and our customers are increasingly finding ways to add agentic value to deterministic pipelines without compromising the core.

DECISION AID

Run each backlog item through four questions.

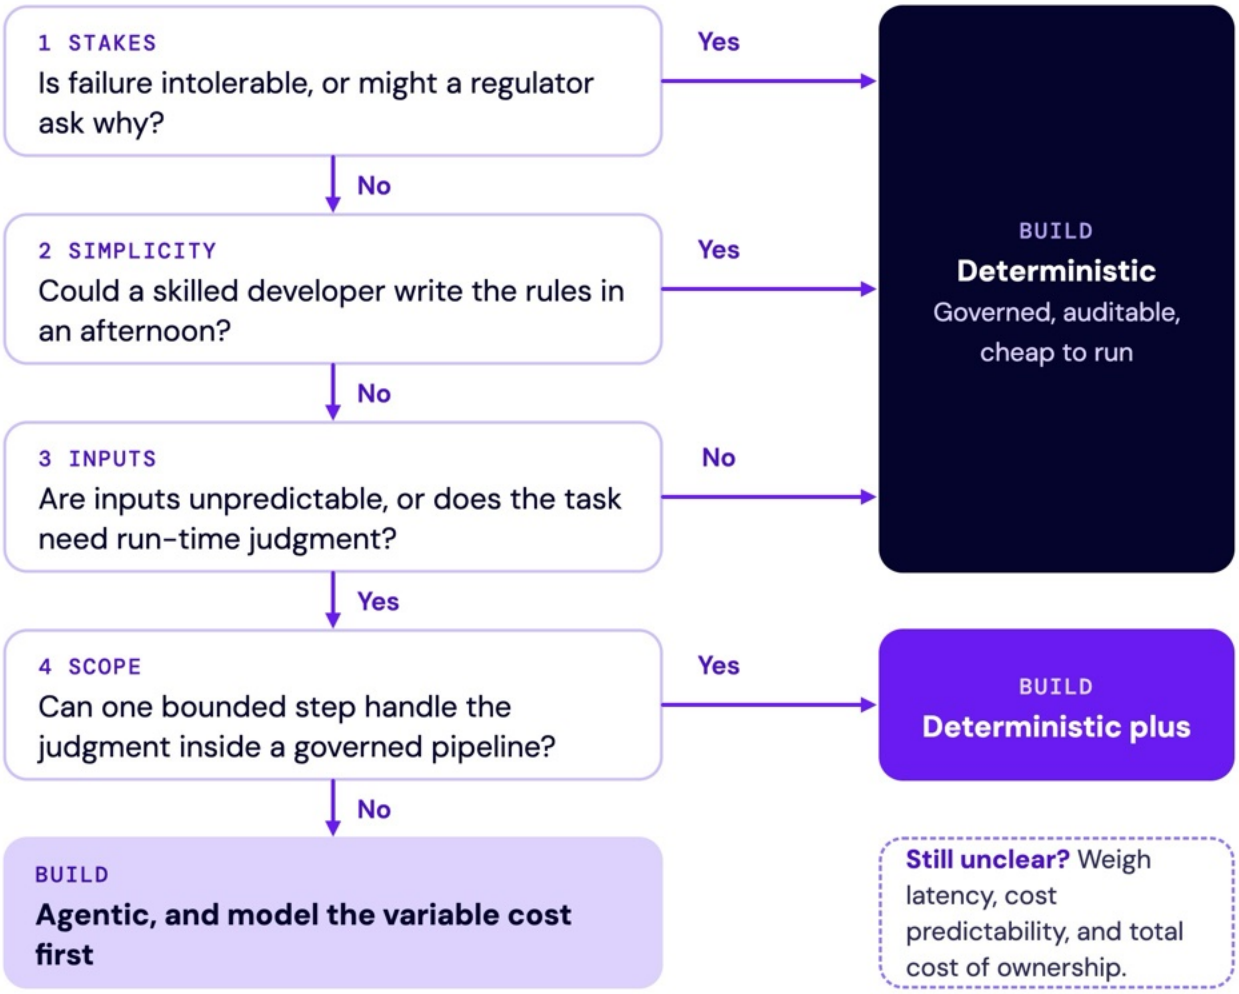

Deterministic workflows and agentic workflows

Before combining them, it's worth being precise about what each end of the spectrum does well—and where each one costs you.

Deterministic workflows

The backbone of enterprise integration. They deliver reliable, auditable, repeatable execution at low cost. When requirements are stable and inputs are structured, code pipelines are almost always the right call, and they're often undervalued in the current AI climate.

- Password recovery workflows
- Compliance audit log exports
- Bank fraud flag escalations

Agentic workflows

Handle what code can't: ambiguity, variable inputs, synthesis, and reasoning. They unlock automation for processes that previously required humans (often expensive specialists with limited availability) to make repetitive and tedious judgments. The tradeoff: output is inherently variable, runs cost more, and the work is harder to audit and monitor.

- Contract summarization & risk flagging
- RFP response drafting
- Generating social posts

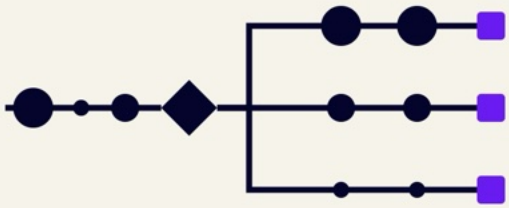

Deterministic

Fixed rules route each input. The same input always takes the same path.

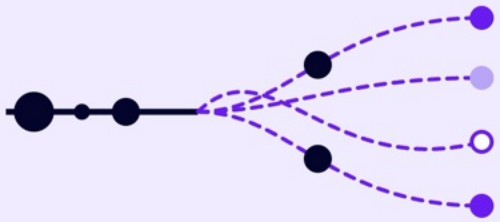

Agentic

Judgment instead of cutoffs. The same input can create different outputs.

Deterministic plus

Most integration workflows should start as deterministic pipelines. “Deterministic plus” describes what happens when you enhance a proven, governed workflow by adding one or more agentic steps that deliver clear, bounded value.

This isn’t a 50/50 split. The deterministic pipeline remains the backbone; the agentic step adds targeted value.

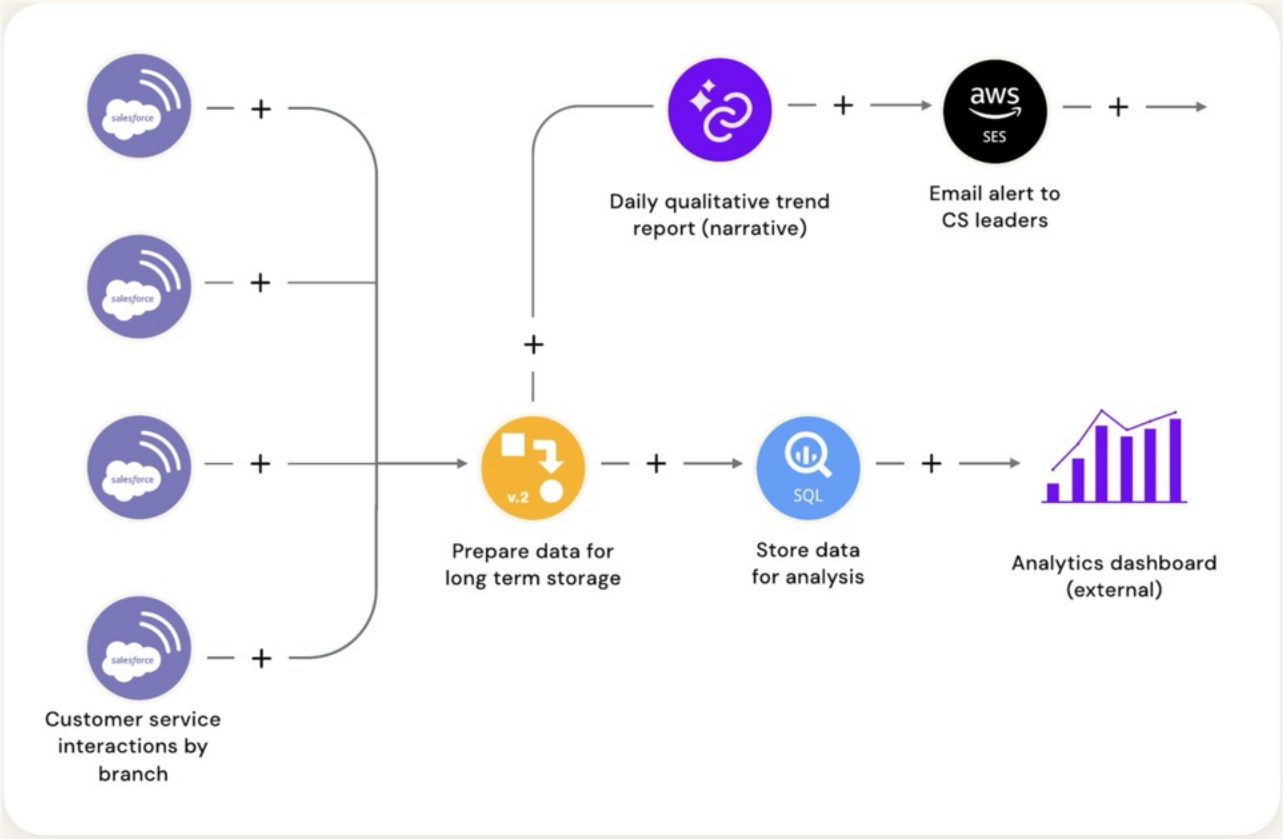

ONE GOVERNED PIPELINE, ONE BOUNDED AI STEP

The AI touches one step. Everything else stays predictable, auditable, and cheap to run.

A complementary variant: separate workflows

A fully deterministic pipeline ingests or moves a batch of data, then triggers an agentic workflow to analyze the batch for insights. One of our customers is experimenting with a centralized evaluation workflow to assess the performance of others across its integration estate. The integration logic stays clean; AI operates only where variability is acceptable.

Deterministic plus lets organizations capture AI value without exposing critical infrastructure to agentic failure modes. It's also how most integration estates will actually evolve—incrementally, deliberately, and with governance intact. **This pattern is what Digibee's architecture is built to support.**

THE BEST STRATEGY IS A DELIBERATE ONE

The most capable integration leaders aren't doing the most with AI. They're the ones who understand the tradeoffs. If a deterministic workflow isn't broken, don't fix it. The backlog doesn't need to be an AI backlog. It needs to be a solved backlog.

CHAPTER 03

What "AI-native" actually means

A year ago, we asked ourselves a simple question: how would we build our platform if we started from zero today? Here's what we built and the reasoning behind it.

The engineering work behind the label

We founded Digibee to take full advantage of cloud technology when it was changing what was possible for B2B software. Most legacy iPaaS platforms marketed themselves as cloud-native while still running lift-and-shift workloads. Our customers have seen and felt that difference ever since. We won't accept anything less in the AI era.

We reassessed Digibee from first principles. Our goal has always been for integration professionals to build enterprise-grade workflows (complete with credential management, security, and auto-scaling cloud infrastructure) as fast as available technology allows. With agentic AI, work shifts from execution to delegation and strategy. Making that real meant solving three challenges:

01 **Make the platform UX just as intuitive for agents as it is for people.**

02 **Specialize and harness AI agents to operate as seasoned integration engineers.**

03 **Evolve how people work with this new, digital workforce.**

Agent access

Humans need visual UIs: menus, buttons, form fields. AI needs atomized, self-explained interaction points. Given multiple ambiguous endpoint names, an LLM will confidently pick one and proceed. So we created APIs and MCP tools documented with clear, self-contained language, written for a reader that has no surrounding context to fall back on.

Context delivery

When users build in Digibee's canvas, the platform surfaces the error message for the connector that failed. We could supply the entire execution log, but chose not to. Users don't need it, and the full log would slow their search for relevant information.

LLMs benefit from processing the entire execution log. An agent doesn't necessarily understand what each node before the failure was supposed to do, and the additional context helps. Human users accumulate context through many vectors—the visual representation of the workflow, their experience with similar integrations. Agents gather context only from what they're given at run time. We built our agent-facing services to deliver deep context structured for LLM-friendly consumption.

CONTEXT DELIVERY

Same failure. Two kinds of reader.

A PERSON IN THE CANVAS

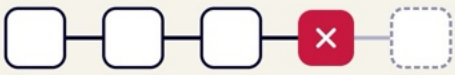

Sees the error message for the failed connector

Fills in the rest from the visual workflow and past experience.

AN AGENT, VIA AGENT-FACING SERVICES

NODE 1	purpose	input	output
NODE 2	purpose	input	output
NODE 3	purpose	input	output
NODE 4	purpose	input	error

Gets the full execution log, structured for LLMs. Knows only what it's given at run time.

Integration expertise

Coding agents excel at work within a codebase. Integration work exists at the level of organizational knowledge. The systems integrations touch are frequently undocumented, and LLM training data includes very few examples of how to wire together obscure legacy systems. Digibee encodes integration logic into a "flowspec" document, and we'd be surprised if any major LLM's training data has seen more than a handful of these.

Building a workflow requires many divergent sub-tasks. Forcing one agent to handle all of them undermines results—the context grows too long, and the agent loses coherence. We distilled our knowledge into discrete prompt templates and distributed them to specialized sub-agents that deliver expert-grade output. The user sees none of this. They interact with one primary digital worker, who synthesizes the sub-agents' output into a cohesive end-to-end project.

Human collaboration

Our approach uses AI to draft projects and workflows, but in collaboration with the user. Users define and guide the project, then verify the output. The integration specialist's job doesn't go away; it changes. Achieving this required interfaces for two tasks: defining the integration specification, and reviewing what the digital worker built.

Checking output was straightforward; we already have the Digibee canvas. For collaborative planning, we built a new interface pairing a chat window with a multi-tab living document. Users begin by describing their high-level goal, then collaborate with the digital worker to craft a detailed project specification, complete with edge cases, trigger conditions, and error handling.

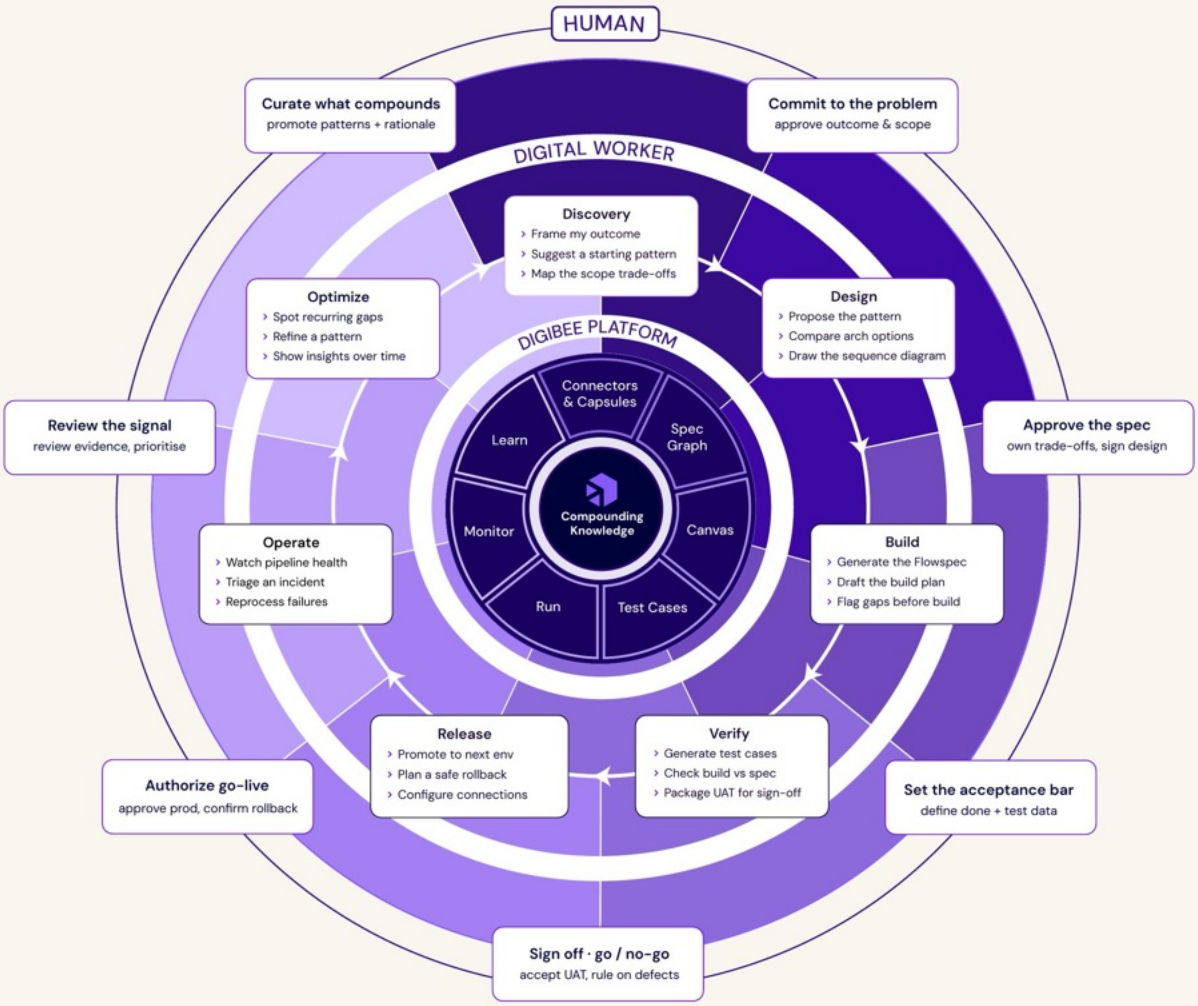

Integration at a higher altitude

For Digibee's users, AI-native means working at a higher altitude. While they review and adjust the digital worker's output in the canvas, they no longer drag-and-drop every component. Clearing the backlog faster frees integration specialists to think strategically: what processes can we centralize into reusable capsules? Where haven't we looked for business impact yet? These are questions integration teams have historically never had time for.

UP THE VALUE HIERARCHY

Every loop returns to real delivery–the altitude rises.

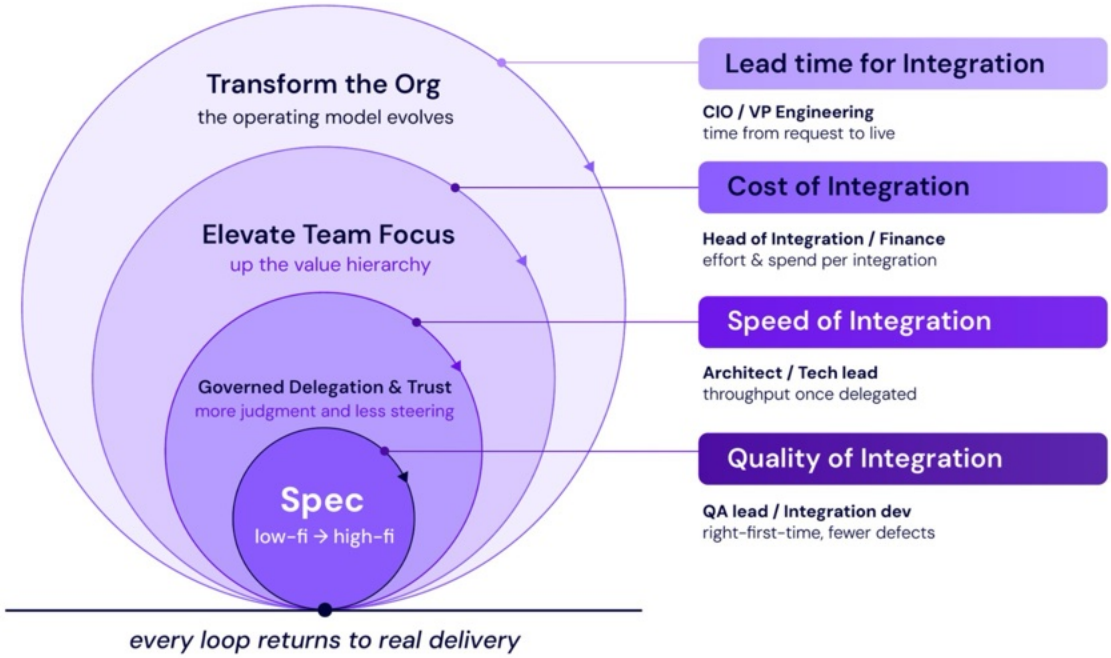

20x
FASTER

In early tests with design partners, users completed workflows up to 20x faster.

The sniff test

“Lift-and-shift” was the tell for cloud-washing. It moved old applications to the cloud without changing how they worked and didn't take advantage of what the cloud could do differently.

The AI-native equivalent is simpler: ask any vendor how they redesigned their platform for agents.

ASK ANY VENDOR THAT CLAIMS TO BE AI-NATIVE

How have you redesigned your platform to meet the needs of agentic users?

Adding AI onto existing products creates the same problems lift-and-shift did. It locks you out of what a genuine rebuild makes possible and keeps you in an intermediate state longer than you’d expect.

Ask the same question of your own workflows. You may find that the biggest gains don’t come from accelerating existing workflows. They come from asking which of those workflows should exist in their current form at all.

See the AI-native integration platform in action

Digibee is the first AI-native enterprise integration platform. If you're responsible for integration infrastructure and want to see what this looks like in practice, we'd like to show you.

Watch a demo → **Get early access →**