DataArt

TRAVEL & HOSPITALITY
A POINT OF VIEW
JUNE 2026

The AI Differentiation Dividend

The advantage won't go to whoever has the most AI but to whoever knows which work to give it, and can trust what it gives back.

The argument in brief

As I head towards Q3 2026, my view is simple: **advantage won't go to whoever deploys the most AI, but to whoever has the best system to turn it into a real increment: predictable output they can trust.**

This isn't a tooling race. Almost everyone now has AI; few get a predictable, repeatable return. The problem usually isn't the AI, it's the system around it: the process, data and context you build and run it with.

DORA's 2026 research on AI ROI confirms it: **AI is a mirror and a multiplier. It amplifies a strong system of work and magnifies the cracks in a weak one. That remains true.**

The payoff is consistent, predictable delivery, and the freedom to invest where you differentiate.

Trust runs through it all: you earn it by measuring output quality, the way you earned uptime. That's the discipline we help businesses build.

Because more AI doesn't make a great business. Good AI, AI you can trust is far more valuable.

DataArt

More spend, uneven return.

Almost every team now uses AI, and token spend has climbed with it. That part is certain. But most reach for it the way they already use chat tools, a quick fix bolted onto today's work. That builds a faster horse, not a new way to move, so the increment stays uneven and rarely company-wide.

TOKEN SPEND

Rising

INCREMENT

Uneven

Adoption and spend are up for certain; by 2026, a third more changes merge with no review at all (Faros). The increment stays uneven and rarely company-wide, until the system changes.

That's the trap. The token spend is guaranteed; the increment isn't. AI doesn't make the return predictable; the system around it does. What turns scattered, team-by-team wins into a reliable, company-wide increment isn't more AI. It's the system.

The variable isn't the model. It's the way you work.

Sources: DORA, 2026; Faros telemetry, 2026; DataArt Data & AI Reality Report, 2025.

DataArt

What decides whether AI pays back

DORA's research on AI ROI, the work I lean on most here, calls **AI a mirror and a multiplier: strong systems pull ahead, weak ones see their cracks magnified.** The return comes from the system of work, not the tool.

Five parts of that system decide whether it pays back:

Point AI at a real problem	**User-centricity:** DORA's top prerequisite for AI to pay back.
A platform to build on	**A high-quality internal platform is what unlocks AI's value.**
Safety nets for speed	**AI lifts throughput but strains stability;** testing and fast feedback hold the line.
Healthy teams	**Well-being and psychological safety still predict** who pulls ahead.
Measure what matters	**SPACE and DORA's team profiles:** people and flow, not lines of code.

For me that's the whole point: performance is a property of the system, not the tool, which is exactly why AI amplifies rather than fixes.

Sources: DORA 2026, ROI of AI-Assisted Software Development; SPACE; Team Topologies.

03 · THE THESIS

Automate the ordinary. Invest in the extraordinary.

The way I think about it, every hour of work sits on one side of a line: undifferentiated grind no traveller thanks you for, or the differentiating judgment and relationships that earn loyalty.

UNDIFFERENTIATED

- Profile & data syncing
- Rate & fare auditing
- Confirmation chasing
- Content normalising

No traveller ever thanks you for it.

DIFFERENTIATING

- Revenue optimisation: right offer, right moment
- True personalisation
- Relationships & service
- Loyalty that earns the next trip

What actually sets you apart.

THE DIFFERENTIATION LINE

I'd hand the ordinary to agents, but keep people in the loop on both sides, judging at the right moments. The scaffolding is commodity now, you can vibe-code it; spend your attention on what truly sets you apart.

3-10×

The upside is real, but earned: the teams reaching it rewired the work, not just added the tool, and only on structured tasks.

DataArt

Trust is the quality of the output. And quality is all context.

That context has to be gathered twice: enough to write the right code when you build AI, and enough to read what the user meant when you run it. The same discipline on both sides, and both are easy to get wrong.

Context is fragile. As the work scales, that understanding degrades, and the output starts to look right while being wrong. Four ways it slips:

Context compresses

Without sub-agents to divide the work, an agent crams everything into one window, and as it fills it drops detail and quietly gets less capable.

Meaning gets misread

Without the right context, "Naples" could be Italy or Florida. One wrong read and a confident answer points to the wrong place.

Retrieval misses

Pull the wrong document and the agent answers fluently from the wrong facts. Retrieval quality is the ceiling on answer quality.

It can't judge quality

Trained on your Code repo, good code and bad, an agent copies what's common, not what's right. Telling craft from debt is the human's job.

For the business: confident, wrong output erodes customer trust and inflates cost. Humans in the loop are the safeguard.

DataArt

You can't measure human trust - but you can measure AI trust.

We don't treat AI quality as a launch checkpoint. For us it's a standing question. We build observability and governance in from the start and keep measuring. Three questions we never stop asking:

01 **Is it right today, not just running?**

02 **When it's wrong, where and why?**

03 **Did the last change make it better, or just look better?**

IN PRACTICE

Accuracy: low-50s → high-80s. Cost per answer: down 33%

On a live travel AI assistant. Continuous, attributed measurement got us there, then let us route routine work to cheaper, smaller models, so reliability and efficiency moved together.

You can't improve what you won't measure.

Measurement isn't a tax on shipping. It's what turns "is this safe to deploy?" into "how fast can we scale?"

DataArt

What it's worth to the business

Predictable, trustworthy output that shows up in the metrics the business runs on.

Better ways of working and measured trust don't just speed up engineering; they move the numbers the business is run on.

Ship safely, faster	**More tested and shipped with confidence**: faster revenue and RevPAR experiments.
Grounded answers	**Fewer wrong outputs**: lower service cost, higher CSAT/NPS, fewer refunds.
Right-sized models	**Run routine work on cheaper, smaller models** you've proven safe. Token and inference cost fall.
Observable & governable	**Measured, attributed and auditable**: compliance and duty of care for enterprise buyers.

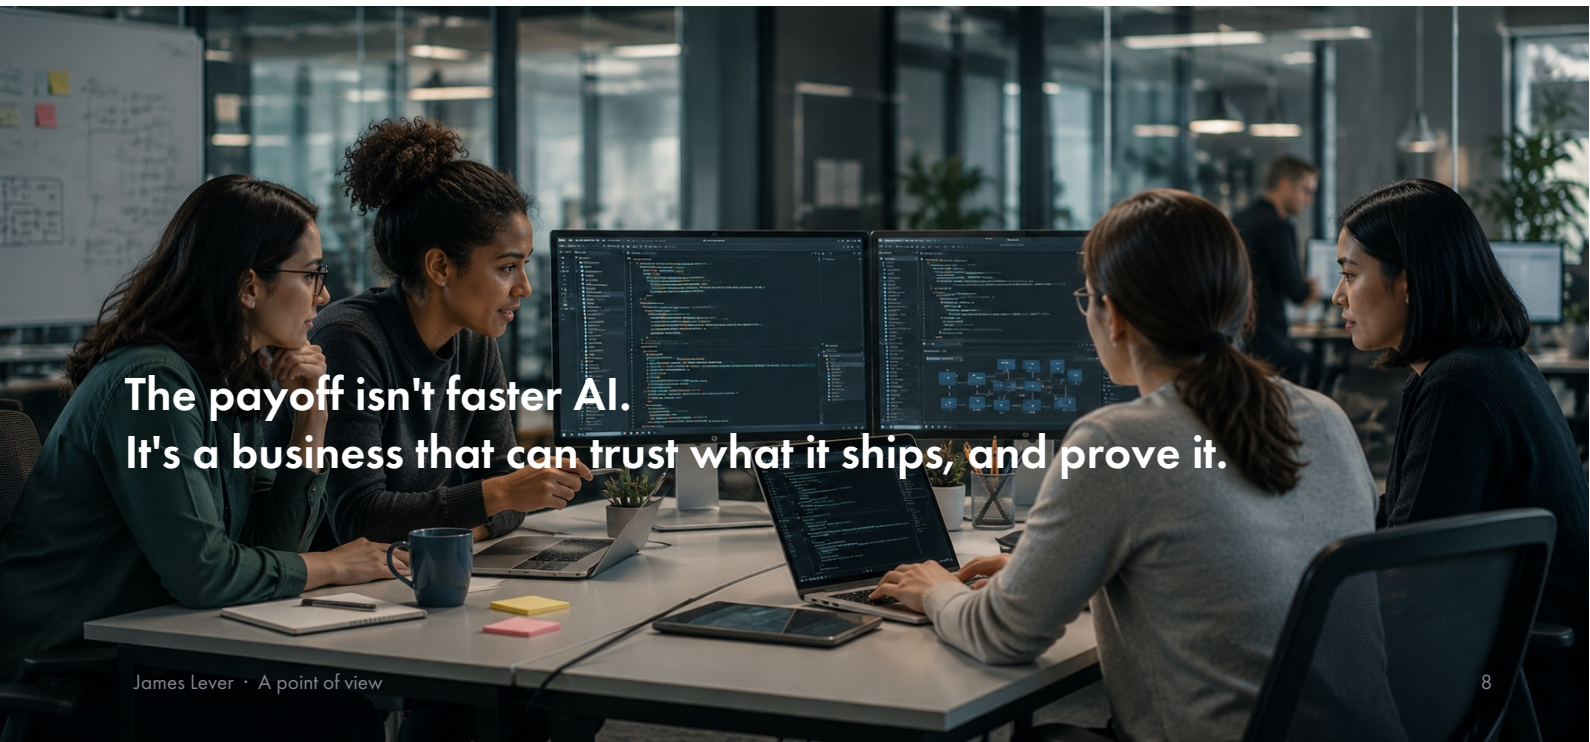

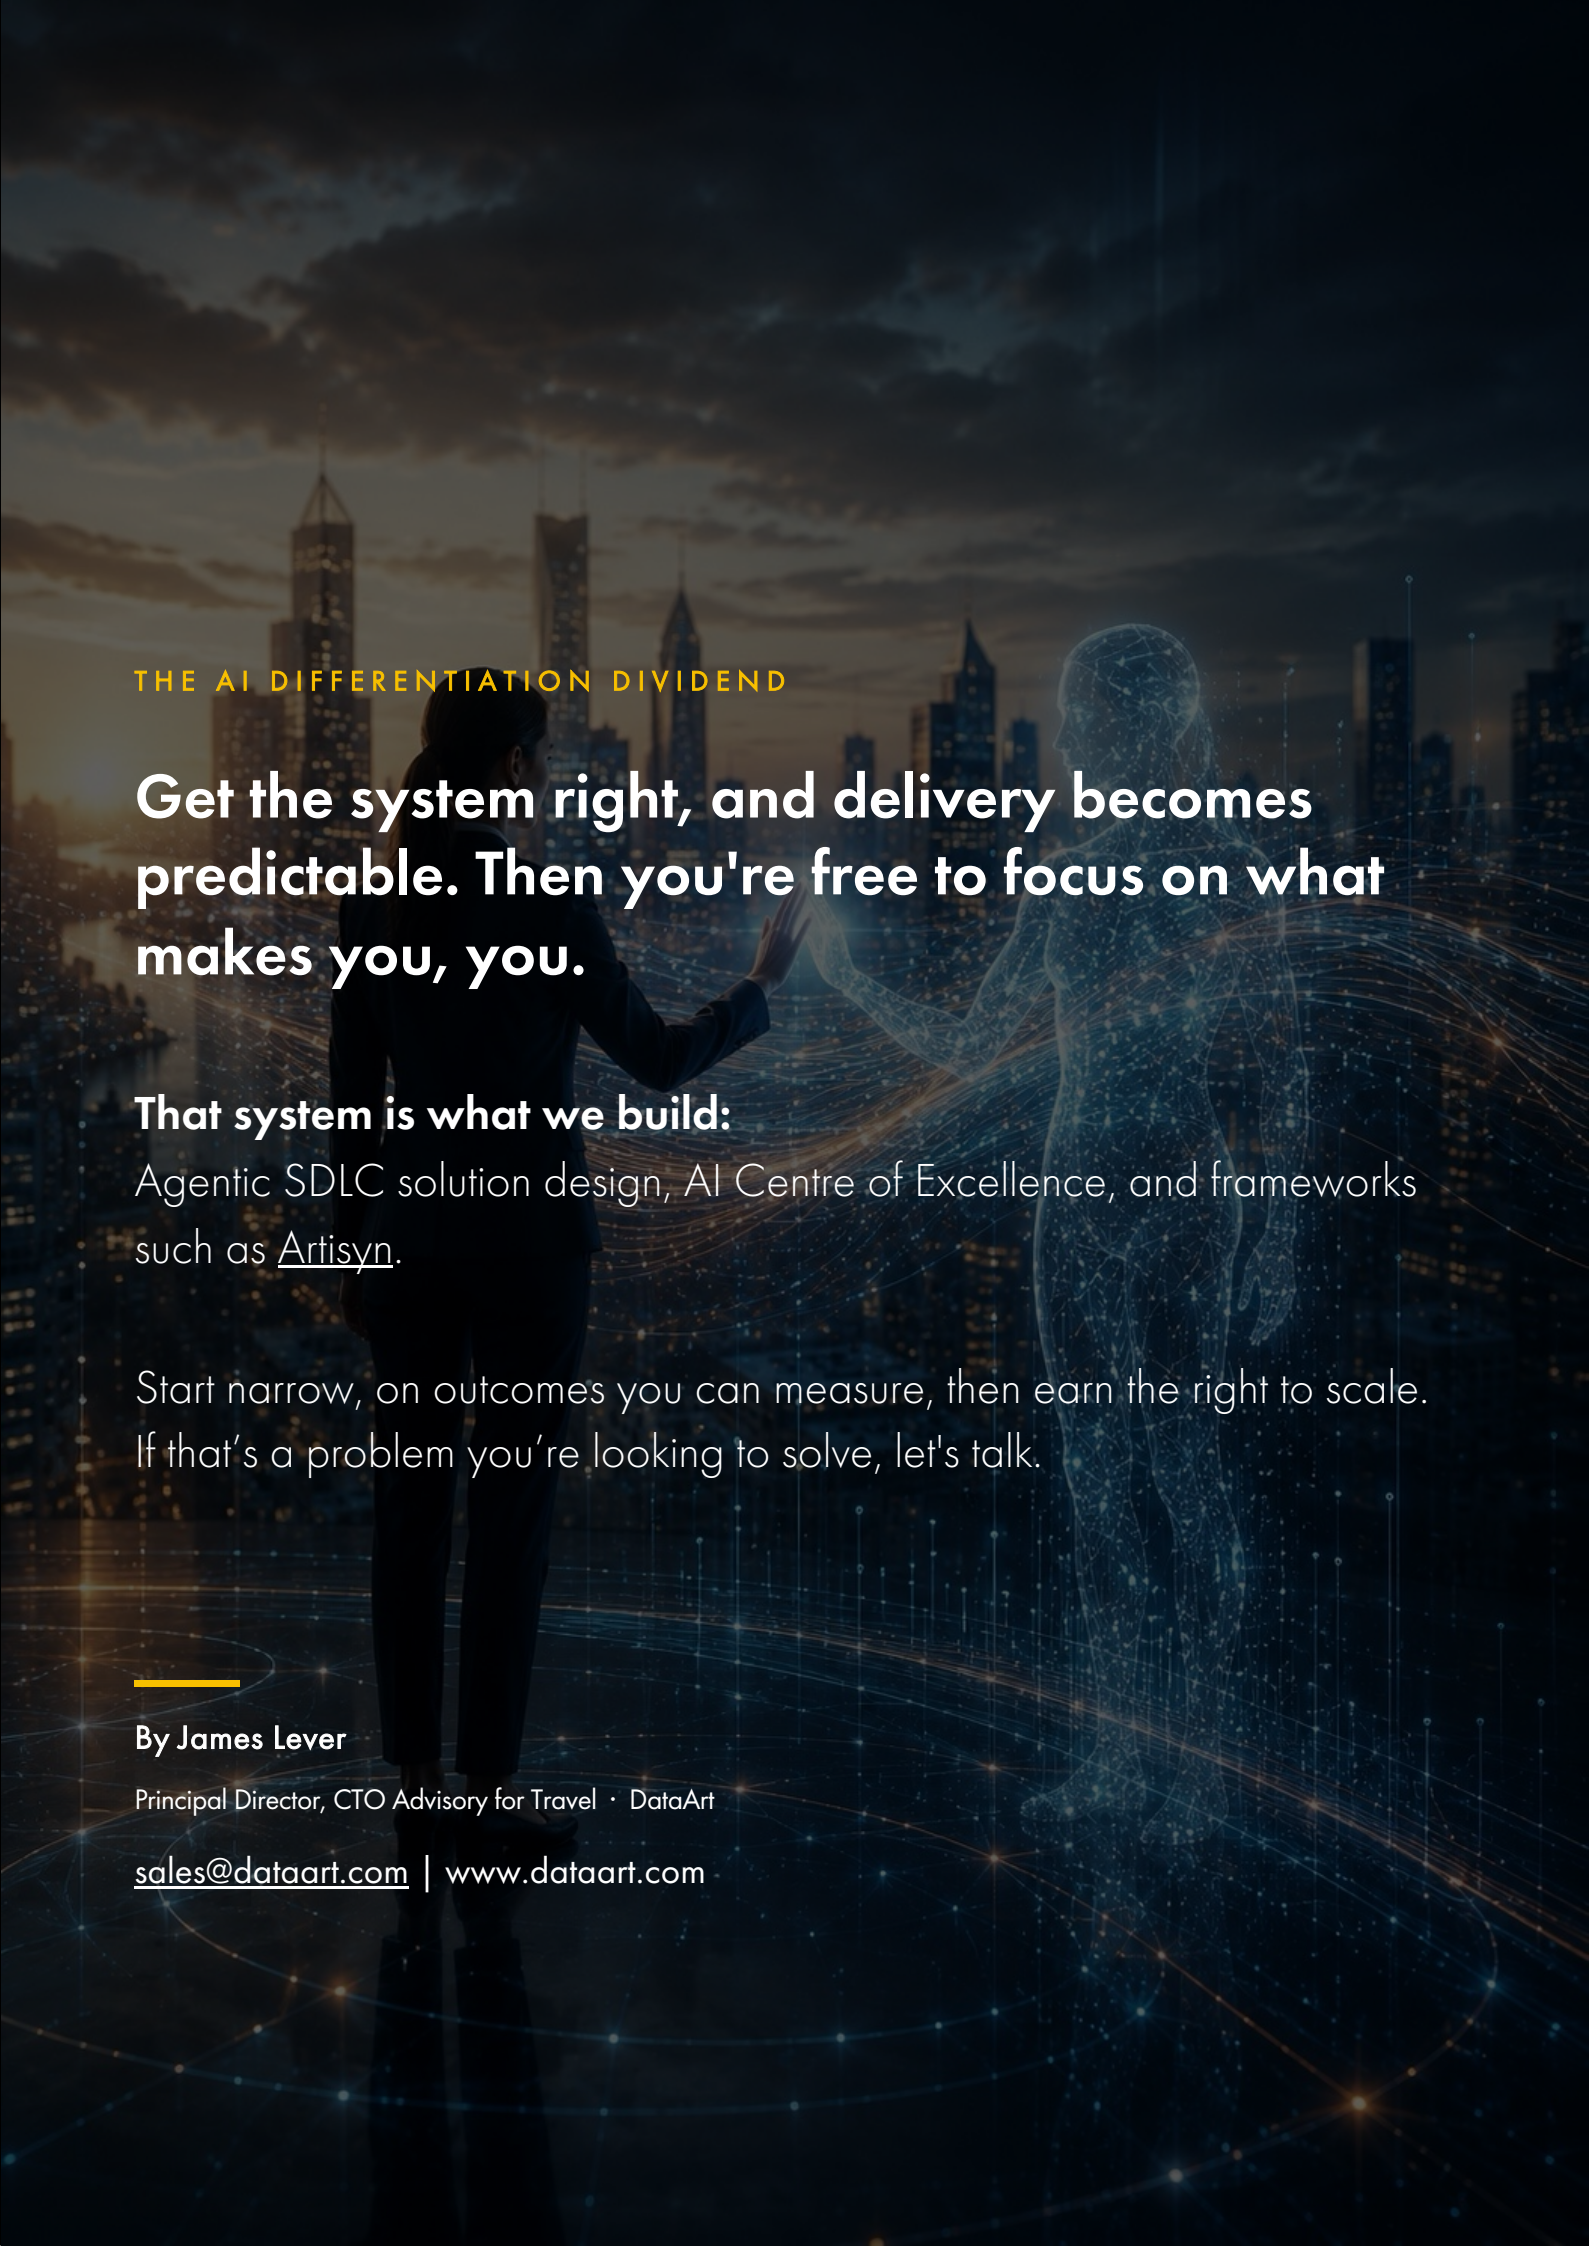

THE AI DIFFERENTIATION DIVIDEND

Get the system right, and delivery becomes predictable. Then you're free to focus on what makes you, you.

That system is what we build:
Agentic SDLC solution design, AI Centre of Excellence, and frameworks such as Artisyn.

Start narrow, on outcomes you can measure, then earn the right to scale. If that's a problem you're looking to solve, let's talk.

By James Lever

Principal Director, CTO Advisory for Travel · DataArt

sales@dataart.com | www.dataart.com

Sources & references

METR (2025) • randomized trial; some experienced developers slower despite feeling faster, an early counterpoint to 2025's measured throughput gains.

DORA / Google (2026) • ROI of AI-Assisted Software Development. AI as an amplifier; the return comes from the system, not the tool.

Stack Overflow Developer Survey (2025) • high tool adoption alongside widespread distrust of AI output.

Gartner (June 2025) • 40%+ of agentic AI projects expected to be cancelled by 2027 (survey of 3,400+ organisations).

Gartner, via HyScaler • roughly 80% of AI projects never reach production.

Anthropic, via CIO.com • 70–90% of its own code AI-written, with mature engineering practice.

IOActive • 31.6% of AI-generated code samples fully exploitable in security testing.

JetBrains (2026) • 20–25% of AI code errors catchable by static analysis before review.

DX Research (2026) • ~91% developer AI adoption; ~22% of merged code AI-authored.

Faros (2026 telemetry) • across ~22,000 developers; throughput up, but PR review time +441% and 31% more PRs merged with no review.

Augment Code • elite engineering teams ~40% AI adoption versus ~29% for struggling teams.

McKinsey / Slack • median ~6.4 hours saved per developer per week.

Team Topologies · Project Aristotle · SPACE • stream-aligned teams, psychological safety, and multidimensional productivity.

Cursor • enterprise transformation guide · dual-lens velocity and quality measurement.

DataArt • Data & AI Reality Report (2025); AIDLC adoption case study, 100+ person travel platform.