

NAVIANT & FEI – October 6, 2026

The CFO's Guide to the **Agentic Enterprise**: Why an **Enterprise Strategy** Is Your Biggest Financial Lever



Mitchell Suter

Chief Innovation Officer
masuter@naviant.com



Jason Bruner

Chief Financial Officer
jbruner@naviant.com



Chris Krause

SVP, Commercial Innovation
ckrause@naviant.com



Brett Hughes

Growth Executive, Commercial
bhughes@naviant.com

naviantTM

Before we begin

A Quick Show of Hands

01

BOARD MANDATE

Do you have a board mandate to drive AI adoption this year?

02

P&L KPI

Is a hard P&L KPI tied to dollar-for-dollar returns on that AI spend?

03

THE “AI TAX”

Are vendors adding a \$30–\$50 per-user AI fee while your software spend keeps rising?

04

LABOR & CAPACITY

Is your team at max capacity, but you can't hire fast enough to keep up?

AI Use Accelerates. EBIT Impact Stagnates.

88%

Use AI in a business function

Adoption is near-universal



80%

Report productivity gains

Individual gains are real



37%

See any EBIT impact

Flat year over year



THE PARADOX

Individual productivity is real, but enterprise-level financial returns lag behind.

Source: McKinsey, The State of AI in 2026: On the Road to ROI (August 2026), 1,719 respondents.

The Dual Squeeze: Margin Pressure Meets Labor Scarcity

MARGIN COMPRESSION

Costs rise faster than prices

- + **Sticky cost structures:** Wage growth and rising SaaS vendor fees compress gross margins.
- + **Pricing power limits:** Customers resist price increases; savings must come from efficiency.
- + **The SaaS tax:** Vendor AI add-ons raise IT spend without expanding operational throughput.

WORKFORCE & CAPACITY

Teams at capacity, hiring can't keep up

- + **Demographic realities:** Accounting and operational talent is scarce and expensive to hire.
- + **The 110% capacity trap:** Knowledge workers spend 80% of their day on manual data gathering.
- + **Non-linear growth:** Output must grow without scaling headcount linearly.

THE IMPERATIVE

AI transformation isn't about replacing your workforce. It's about scaling capacity without expanding overhead.

The Unspoken Truth About AI Transformation

Everyone is selling you agents.

Almost no one is telling you the truth.

Transformation means understanding three things—**People**, **Flow**, and **Context**—and rebuilding your organization around them.

PEOPLE

FLOW


CONTEXT



PEOPLE

Adapting software to people

Designing digital experiences around how teams actually collaborate and solve problems.



FLOW

How work actually moves

Orchestrating end-to-end processes so value moves across disconnected systems.




CONTEXT

Connecting data to business intent

Fusing operational data with business rules so every action drives the right decision.

Employees as “Human Middleware”



Hard Truth of AI

A faster engine doesn't shorten the trip when you hit a red light every mile.

The 30-year design failure

Specialized software built in silos forces skilled workers to toggle across 11 different applications just to complete routine tasks.

The real operational overhead

Employees spend up to 80% of their time simply gathering and assembling information across those tools. That context assembly is the core bottleneck to solve.

Three Pillars of Agentic Success

01

PILLAR 01 | PLAN

Strategic Blueprint

- + Align AI spend to P&L outcomes before writing code
- + Prevent “random acts of AI” through continuous blueprinting

02

PILLAR 02 | BUILD

Enterprise Architecture

- + Unify fragmented data into an authoritative knowledge layer infused with business intent
- + Orchestrate many flows to keep knowledge authoritative and power autonomous action.

03

PILLAR 03 | RUN

Continuous Innovation

- + Anchor change models to evolving business needs, not SaaS constraints
- + Implement systematic performance tracking and optimization

Protecting Financial Systems & Enterprise Data

01

PILLAR 01 | PRIVACY

Data Privacy & Residency

- + Rented models run sandboxed
- + Corporate data is never stored, exposed or used to train public models

02

PILLAR 02 | CONTROL

Programmatic Hard Gates

- + Agents validate, reconcile and stage transactions for review
- + Humans retain exclusive General Ledger and credit posting authority

03


PILLAR 03 | AUDIT

Audit Provenance

- + Immutable decision logs for every agent action
- + Every recommendation links back to source-system evidence

AI agents calculate and stage; financial executives retain exclusive authority over execution.

Decouple Intelligence from Context



Financial safeguard: Rented models run sandboxed; your data is never used to train public models.

What Is an Enterprise Knowledge Layer?

SYSTEMS OF RECORD



+ Defined above systems

Core business entities (Customer, Employee, Invoice, Asset) exist **outside any SaaS application**.

+ The enterprise ontology

A unified, software-agnostic map of **how your business actually operates**.

+ Permanent IP


Applications can be swapped; your knowledge layer is your **permanent enterprise asset**.

“An ontology is simply your enterprise dictionary: **one authoritative definition above legacy silos.**”

The Customer Model: Separating Business Rules from Live Facts

THE ENTERPRISE DICTIONARY & BUSINESS RULES

Shared business logic above software



- + Defines what a “**Customer**” means across every system.
- + **System-agnostic**: works the same whether data comes from Salesforce, SAP or ServiceNow.

LIVE OPERATIONAL FACTS & TRANSACTIONS

Real-time system truth

Customer = Acme Corp

\$100K
Active deal

\$50K
Overdue AR

3
Open tickets

- + Holds **live records and state changes** from across your software stack.
- + **Execution layer**: background agents act directly on these live facts.

Business rules define how your enterprise operates. **Live facts tell your agents what to do right now.**

Why Background Agents Require Authoritative Context

WITHOUT A KNOWLEDGE LAYER: DANGEROUS



✗ WHAT IT MISSES

Fails to see the **\$50K overdue AR balance** in the Subledger → drives a bad financial decision

WITH AN AUTHORITATIVE KNOWLEDGE LAYER: SAFE



✓ WHAT IT DOES

Evaluates complete enterprise context → reconciles risk → **executes a safe financial workflow**

1 Headless background execution

Agents run on business events without waiting for a human prompt.

2 Context over prompting

An agent is only as smart as the complete enterprise context it receives.

3 Ambient curation

As agents execute work, your knowledge layer updates automatically as a byproduct.

“An agent with 20% of the data makes flawed decisions. **Authoritative context ensures safe execution.**”

9 Economic Levers

DESKTOP TRACK 4 levers

- 1 Universal access with a value plan**
Shift focus from seat counts to targeted role capacity
- 2 Build the skills catalog**
Capture role wins as version-controlled enterprise skills
- 3 Connect to enterprise knowledge**
Ground desktop tools in authoritative company context
- 4 Token literacy**
Build cost culture; route routine work to lighter models

OPERATIONAL TRACK 5 levers

- 5 Escape the “SaaS AI toll booth”**
One AI capability above your stack; skip 5 vendor taxes
- 6 Dynamic model routing**
Route each task to the lowest-cost correct engine
- 7 System hunting to one-click**
Eliminate human middleware; pre-assemble context & math
- 8 Compounding incrementalism**
Build #2 costs 50% less, Build #3 costs 75% less
- 9 Tech stack rationalization**
Decommission replicable SaaS and retire legacy UI seats

TRACK 01

Desktop Productivity

Universal assistants, governed skills, enterprise knowledge and token literacy



01

The Omnipresent Assistant

THE DESKTOP IMPERATIVE

1

Universal access with a value plan

Buying seats is easy

Every employee gets desktop AI. Shift focus from license counts to targeted role capacity gains.

2

Build the skills catalog

Discovered through role enablement

Capture role wins as governed, versioned skills that scale across teams.


3

Connect to enterprise knowledge

Own the context

Ground desktop tools in authoritative enterprise context so they deliver real business output, not generic text.


REFRAMING THE VALUE MODEL



None of this works without role-based change management

Shift the conversation from license count to role impact, one role at a time.

How Desktop Assistants Leverage Authoritative Customer Context



Build Token Literacy

Use the right model for the work.

LEVER

4

naviant™

Cost literacy in practice

Three habits that keep AI spend in line.

1

Know the price tiers

Every request has a cost, from \$ to \$\$\$\$.

2

Start with the everyday model

Escalate to frontier reasoning only when the task needs it.

3

Review spend by role

Leaders see where reasoning budget goes and coach to it.

Match the model to the work

Use more intelligence only when the task requires it.
work



Complex reasoning

\$\$\$\$

Strategic analysis, difficult decisions,
application building



Standard work

\$\$

Drafting, analysis, research




Everyday tasks

\$

Formatting, summaries, routine admin

More
intelligence
(higher cost)

Less
intelligence
(lower cost)



Don't play "Token Cops." Build AI cost literacy into your workforce culture.



KEY TAKEAWAY

Freedom within guardrails: role-based access, approved models, usage limits, and the judgment to use the **lowest-cost model**

Token Routing Across Customer Data Lookups and Complex Analytics

SIMPLE FACT RETRIEVAL

\$

Simple Customer lookup

TASK

"What is Acme Corp's billing address and tax ID?"

NATURE Static fact lookup across live databases

MODEL Lightweight model (e.g., Haiku / Flash)

COST \$0.0001

DEEP MULTI-VARIABLE REASONING

\$\$\$\$

Deep Customer reasoning

TASK

"Evaluate Acme Corp's payment delays against support ticket volume to calculate churn risk."

NATURE Multi-variable reasoning across complex operational facts

MODEL Frontier reasoning model (e.g., Opus / GPT-4o)

COST High reasoning budget



KEY TAKEAWAY

Don't burn Tier 1 reasoning budgets on **simple Customer fact lookups**.

TRACK 02

Operational Workflows

Core business nouns, model routing, and system abstraction



02

Escaping the “SaaS AI Toll Booth”

THE SAAS TOLL TRAP

5 AI taxes for 5 partial views


Salesforce	AI add-on +\$30–50/user/mo	own data only
SAP	AI add-on +\$30–50/user/mo	own data only
ServiceNow	AI add-on +\$30–50/user/mo	own data only
Workday	AI add-on +\$30–50/user/mo	own data only
Microsoft 365	AI add-on +\$30–50/user/mo	own data only

Result: 5 add-ons, 5 partial views → workers still act as manual human middleware



THE ALTERNATIVE

One enterprise AI capability above your stack




The twin failures of SaaS AI add-ons

- + **Expensive and fragmented:** You pay \$30 to \$50 per user across multiple vendor portals for isolated point tools that cannot communicate.
- + **Perpetuates “human middleware”:** Application-bound AI cannot reach across your stack, leaving your workforce to manually gather context and bridge the gaps between systems.

Don't pay the AI tax five times for five disconnected point solutions.

Why Vendor AI Add-Ons Fail the Customer Model

THREE VENDOR BOTS, THREE PARTIAL VIEWS OF ONE CUSTOMER



THE PROBLEM


You pay **\$150/user/mo**, but **no vendor bot** sees the complete Customer truth.

THE SOLUTION

One **Knowledge Layer** sits above the portals and unifies all live Customer facts.

Dynamic Model Routing: Continuous Optimization & Unit Economics

INCOMING TASKS



Models are fungible utilities

Hosted models have strict shelf lives and shifting API pricing. Centralized configuration lets you swap models without breaking agent logic.

Lowest-cost correct execution

Operational workflows don't need brand-name models. Route tasks to the cheapest model that delivers a 100% verified result.

Knowledge is the true asset

Building an agent is trivial when your enterprise context exists. Focus engineering on your enterprise data map, not prompt tuning.

Operational Execution Tiers for Customer Workflows

TIER 3

Local / open-source

\$

Formatting phone numbers and addresses during routine **Customer CRM data syncs**.

TIER 2

Purpose-built SLM

\$\$

Extracting line items and verifying signatures on incoming **Customer PDF purchase orders**.

TIER 1

Frontier reasoning engine


\$\$\$

A nightly background sweep evaluating **overdue AR balances** against business credit rules.


Same Customer data, three price points: **route each step to the cheapest model that gets it right.**

From System Hunting to One-Click Approval

OLD WORLD: HUMAN MIDDLEWARE



AI-NATIVE WORLD: CONTEXT-ASSEMBLED EXECUTION



Eliminate human middleware.

Agents gather every relevant fact across systems and present the recommended action. Humans simply review, approve, and move work forward.

TIME TO DECISION

45_{min}



30_{sec}

Same decision, with full context and recommendation pre-assembled.

How Agents Pre-Assemble Live Customer Facts for Humans

WORKFLOW Customer requests a \$25,000 credit limit increase

AGENT WORKS IN THE BACKGROUND

1

Event triggers agent

No human prompt required.

2

Agent queries business rules

Requires order history, AR status and a risk score.

3

Agent assembles live facts

Order volume **\$200K**

Overdue AR **\$0**

Payment score **98/100**

4

Agent recommends

"Approve \$25K credit line extension."

5

HUMAN DECIDES

Human approves

Verifies complete evidence and approves in 30 seconds with one click.

Compounding ROI Case Study

DAY 1

Build #1: First workflow

- › Map the core business entities
- › Deploy the first agent
- › Labor savings fund the platform

DAY 60

Build #2: Second workflow


- › Reuse the existing entities
- › Add only new data connections
- › Ships in 50% of the time at 50% of the cost

DAY 120

Build #3: Third workflow

- › Reuse entities and connections
- › Add only the new workflow logic
- › Ships in 25% of the time at 25% of the cost


Relative build cost, indexed to Build #1




Build #2 costs **half as much**, and Build #3 costs **75% less**, because the context already exists.

How Reusing the Customer Model Cuts Build Costs by 75%

RELATIVE BUILD COST



Tech Stack Rationalization: Decommissioning Replicable SaaS



Legacy systems keep the records. People stop working inside them.

- + **Build vs. buy shift:** Organizations build or abstract capabilities internally rather than buying expensive software add-ons.
- + **Headless back-ends:** As workers operate out of a unified task queue, legacy portals turn into raw data stores.
- + **Decommissioning dividend:** Systematically cancel redundant point solutions and retire high-cost UI seat licenses.


THE BUILD VS. BUY SHIFT

32%

of enterprise organizations have decided against purchasing software products or features because they can build or abstract them internally.

Decommission the software. Reclaim the balance sheet.

Why Pay UI Seat Licenses When You Own the Customer Model?



Legacy systems keep the records. People stop working inside them.

LEVER

9

naviant™

THE CFO CONCLUSION

When your Customer model lives in your Knowledge Layer and employees work out of one task queue, legacy UI seat licenses become obsolete.

Decommission the portals. Reclaim your balance sheet.

Strategic Execution: Solve Today, Build for Tomorrow

01

Enterprise strategy and architecture

Treat AI as a P&L strategy, not software purchases

- + Stop buying disconnected vendor add-ons. Own your enterprise knowledge layer above legacy silos.
- + Real ROI comes from pulling the 9 economic levers across desktop and operational workflows.

02

Solve today, fund tomorrow

Pragmatic, compounding execution

- + Don't try to boil the ocean. Target one high-friction, high-ROI operational workflow right now.
- + Immediate labor and process savings fund your platform foundation, so subsequent builds cost 50% less.

03

Continuous innovation and change

Change management as a core capability

- + Commercial models have short shelf lives. Complete model fungibility protects you from technical debt.
- + Invest in the organizational muscle, token literacy, and governed skill catalogs needed to absorb rapid change.

Don't boil the ocean. Solve one high-value problem today to build the compounding brain for tomorrow.

Thank you.



Mitchell Suter

Chief Innovation Officer
masuter@naviant.com



Jason Bruner

Chief Financial Officer
jbruner@naviant.com



Chris Krause

SVP, Commercial Innovation
ckrause@naviant.com



Brett Hughes

Growth Executive, Commercial
bhughes@naviant.com

