



Tokenomics: Aligning Digital Labor, Token Costs, and ROI with ValueOps

What did AI cost us — and was it worth it?
Your Guide: Brian Nathanson, Broadcom

Sponsored by

ValueOps
by Broadcom

Clarity
by Broadcom

Rally
by Broadcom

ConnectALL
by Broadcom

Insights
by Broadcom

A quick show of hands

Who can tell me what your organization spent on AI last month?

Keep your hand up if you can tell me which projects/investments it was for.

The work moved. The tracking didn't.

YESTERDAY

A developer's hours on a funded build.

Recorded on a timesheet or a work item. Capitalized.
On the balance sheet.

TODAY

Some of that work is done by AI.

Untraceable by default — so it's expensed. That
depresses EBITDA.

Same work. A different line on the P&L.

Nobody decided that. It's what happens when the tool changes and the tracking doesn't.

Meet Atlas

\$24,000

AI spend, last month

A FUNDED PROJECT

Internal customer portal

Six developers

In flight, approved last quarter

Illustrative numbers — don't check my arithmetic.

We assume that an approved investment adds value

Two asks, one blocker

LEADERSHIP ASKS FOR

A forecast

AND FOR

Control

BOTH DEPEND ON

Saying what it bought

**What percentage of your AI spend
can you explain?**

Not what it cost. What you can explain — tie to a project, a team, a decision.

If your answer today is twenty percent, that's a fine place to start.

Why nobody can answer • Forecasting it • What it's worth getting right • Seeing and governing it • When it becomes your problem

Why nobody can answer

The money is real. The trail isn't.



Four concepts: Brief anatomy of an AI transaction

Token

A chunk of text. Roughly three quarters of a word. Everything is priced per million of them.

Input / Output

What you send, and what comes back.
Output costs several times more.

Cache

Repeated context, billed at a discount — if the tool is set up for it.

Assistive / Agentic

One answer, versus a loop that plans, calls tools, and retries.

An AI transaction has more cost layers than this — the compute, the data you feed it, the engineers who built the feature.

Every one of those already had a home in your portfolio. Tokens are the one that didn't.

Three things that surprise people

1

It's not how often you ask. It's how much you send.

Pulling project history and status notes into one request can cost fifty times a lean one — for the same answer.

2

What comes back costs more than what you send.

Output is priced several times higher than input. Long answers are the expensive ones.

3

Retries and failures still bill.

A call that errors out, or returns something unusable, is charged the same as one that works.

How you write the prompt is now a budget decision.

Three levers, all of them defaults nobody revisited: which model, how much context, how many turns.

It doesn't ramp. It jumps.

Assistive

Summarize this. Draft that.

A few cents. Predictable. Easy to forecast.

Agentic

Plans, calls tools, loops, retries.

Several dollars. Ten to a hundred times more per outcome — with much wider swings.

Each turn resends everything before it, so late turns cost far more than early ones.

Any forecast built on today's usage breaks the day you deploy your first agentic capability.

And the pricing itself is still moving — seat licenses turning metered, allowances changing, unit prices falling while consumption per task climbs.

Where the \$24,000 went

Coding assistant

\$14,500

Six developers, metered usage

Chat and drafting

\$6,100

Spread across the team, seat-based

Brokered by Clarity

\$3,400

The only slice that arrives already attributed

Three systems. Three owners. Each one sees a piece.

And the invoice can't disentangle it. Providers bill at the level of the key, the workspace, or the deployment — never the request.

So how many keys you created, and how you carved them up, already decided whether this question is answerable. That was settled in an afternoon by whoever set it up first.

Every other cost has a gate. This one doesn't.

Software license

Negotiated. Approved. On a purchase order.

Cloud infrastructure

Requested, provisioned, and tagged before it runs.

Contract labor

Statement of work, PO, invoice matched against it.

AI consumption

A developer decides. The money is spent.

**Every other cost you manage has a control point upstream.
This one has attribution downstream — or nothing.**

What we can see — and what we can't

VISIBLE

- AI calls within Clarity
- Usage you send us — from a gateway, from other tools, or as a file

INVISIBLE

- Anything routed around the instrumentation
- Project data pasted into a consumer chatbot — the cost and the value both
- Work nobody thought to tag

Clarity is not trying to be your enterprise AI cost management system.

We want to tie your AI costs to your investment intent so you can determine value.

Forecasting it

The first thing your leadership asked for.



Three parts. Three ways to forecast.

RATE IS GIVEN

Seats and licenses

You don't forecast this. You read it off the contract.

The rate is handed to you and it's fixed. Multiply by headcount. A shrinking part, as vendors move seats to metered pricing.

RATE IS DERIVED

Steady-state usage

No rate exists. You derive one from your own history.

Observed cost per active user over the last several months, projected forward with a growth assumption. A run rate — but not a stable one.

NO RATE YET

New capabilities

No history to project from. Each one is a business case.

Agentic features don't ramp, they step. A discrete decision with a discrete cost — not a percentage on last quarter.

A blended growth rate hides the one part that breaks the forecast.

When the rate moves, ask why

The rate will move. Size isn't the test — cause is.

YOU CAN NAME A CAUSE

A capability was enabled. A model was swapped. A default changed.

That was a step, not drift. Reset your baseline from here — and if nobody logged the change, that's the thing to fix.

YOU CAN'T NAME A CAUSE

Nothing changed except how people are working.

Genuine drift. Re-derive the rate from the recent months and shorten how often you revisit it.

Declare, don't request. A change log, not a change request.

Watch the rate, not the total.

Cost per active user climbing while headcount is flat is the only early warning you get.

Forecast in units of work

Unit prices are falling.

Consumption per task is rising.

At different rates. So don't forecast the dollars directly.

Cost per developer, per month

Cost per ticket

Cost per report

Convert to dollars as a separate assumption you revisit quarterly.

Tokens aren't comparable across vendors — the same document counts differently depending on who you send it to. Units of work is the only measure that survives a model or vendor change.

What it's worth getting right

Where the money actually is.



The hinge

Whatever you can't attribute isn't capitalizable.

Only direct costs capitalize. A cost is direct only if it can be traced to a specific asset-creation effort.

Untraced spend is indirect by definition — and indirect costs are expensed.

Coverage isn't a maturity score. It's the ceiling on what's eligible.

Your contract decides this before your tooling does

Does the fee move with what you built? Then it can be direct. Paid regardless? Overhead.

Usage-based Pay-as-you-go, or a prepaid commitment that draws down against measured use

Direct if traceable

Capitalizable

Commitment + overage Where you expect to exceed the minimum

Minimum acts as a floor, not a fixed fee

Capitalizable

Flat fee for access Enterprise subscription, reserved capacity

Paid whether or not anything is built

Expensed

Procurement is making an accounting decision without knowing it.

What ASU 2025-06 changed

BEFORE

Which stage of the project are you in?

Infer the accounting treatment from what people were doing.

NOW

Has management authorized and committed funding — and is the project likely to finish and be used?

**That's an approval, an investment record, a stage gate.
It already lives in your portfolio — with an audit trail.**

Which costs qualify didn't change. Only when the clock starts.

The part nobody's telling you

Significant development uncertainty — unproven functionality, or requirements still being substantially revised — means the bar isn't met.

And that uncertainty only clears by building and testing. Exploratory AI-assisted work is often exactly that.

FASB itself expects capitalization could decrease for cloud-delivered software.

**The boundary got easier to evidence.
The bar got harder to clear for genuinely novel work.**

Atlas, revisited

\$15,100 After funding approval, on work likely to finish

Potentially capitalizable

\$5,200 Exploration, before the shape was settled

Expensed

\$3,700 Production support and defect work

Expensed

Provable today: none of it.

Seeing it and governing it

The second thing your leadership asked for.



Three ways to capture it

At the gateway

A central proxy stamps each call with team, user, and feature on the way out.

Best data available. But some of you route only part of your traffic — and some of you turned gateways down over the delay they add.

At the tool

Coding assistants report their own usage, tagged with team and cost center, through a settings file admins push out.

No developer has to remember anything. It isn't in the request path, so it costs nothing in performance.

From the provider

Usage and cost reports, grouped by key, workspace, and model.

The coarsest of the three — but it's the total, and it's the only one that's complete.

Which one you land on is set by your architecture, not your intent. Most of you will use more than one.

Most of you will get a file for Clarity to process

Your platform or security team owns the provider credential. They are not handing it to a PMO application.

What they will do is run the pull themselves and hand over an extract. Broadcom does exactly this.

The file isn't the fallback. For most of you it's the normal path — and we're building it that way.

Which means no real-time enforcement. Leave that to the gateway and the provider's own budget controls.

Your close is monthly anyway. Same cadence, new cost type.

Coverage is the metric

Partial instrumentation is more dangerous than none. The numbers look complete, and they aren't.

What you can attribute

÷

The AI consumption invoice

Report the unexplained remainder as a number in its own right.

Answerable at any level of tooling. Improves incrementally. Requires buying nothing.

This is coverage of the consumption bill. The other cost layers you already track.

The ask is smaller than you think

Some of you chose to replace timesheets with alternatives: capitalizable work items in Jira or Rally, story-point allocation, role-based percentages.

Regardless, the plumbing for AI attribution is the same plumbing.

You're not being asked to build a new mechanism. You're being asked to extend one your controller may already depend on — with the audit precedent already established.

And it solves the harder problem for free: a ticket's type and state tells you feature versus defect versus support.

Four layers. Two naturally fall into Clarity's scope.

Prevent Spending caps, rate limits, per-user budgets, cheaper models for cheap work.

No

Detect Alerts on spend spikes, coverage drops, unexpected model mix.

Shared

Attribute Spend tied to funded work. Restrict the capitalizable pool to defined roles.

Yes

Decide Was this project's AI spend justified by what it delivered? Should it continue?

Yes

Tokens measure what you spent, not what you got.

A spending cap stops runaway usage. It can't tell you whether the usage was worth it.

ROI: you already have the other half

Coverage gives you AI cost per investment. Your benefit plans already say what that investment was meant to deliver.

IF YOU TRACK BENEFITS WELL

AI cost per unit of delivered value becomes computable.

Years of unglamorous benefits work suddenly pays off — you're the only ones who can answer the question everyone is about to ask.

IF YOU DON'T

You're about to get scrutiny on benefits realization that was never there before.

"We assumed an approved investment adds value" has held up fine for twenty years. It won't once there's a number next to it.

AI exaggerates what you already do — in both directions.

This isn't a new problem. It's new attention to whatever you've always had.

Demo

- Home
- Data Ingestion**
- Hybrid Workforce
- Team Productivity Dashboard
- AI Productivity Dashboard
- Digital AI Investments
- AI Capitalization
- AI ROI
- Digital Thread
- Administration

Data Ingestion 00000021

Vendor Dashboard Vendors Billing Data Internal Tagging



View Vendor Dashboard Save [Configure](#)


AI Budget16.4M

Spent to Date9.37M

Remaining7.03M

Investment		Contracts	
Name *	Amount	Spent to Date	Remaining
Total	16,400,000	9,370,000	7,030,000
Google	5,000,000	4,820,000	180,000
OpenAI	1,200,000	550,000	650,000
Anthropic	1,800,000	2,510,000	(710,000)
Microsoft	2,800,000	940,000	1,860,000
Amazon Web Services	2,400,000	280,000	2,120,000
Meta	3,200,000	270,000	2,930,000

1-6 of 6 < >



<< Collapse

- 🏠 Home
- 📊 Data Ingestion
- 👥 Hybrid Workforce
- 📈 Team Productivity Dashboard
- 📈 AI Productivity Dashboard
- 🚀 Digital AI Investments
- 💰 AI Capitalization
- 🔍 AI ROI
- 🎧 Digital Thread
- ⚙️ Administration

Data Ingestion 00000021

- Vendor Dashboard
- Vendors
- Billing Data
- Internal Tagging

Actions ▾

New from Template

🔍 Search...

<

Per-Period Metrics - 4 Periods

>

🔍

📊

☰

📄

📅

▾

View

Tokenomics (Total Tokens) ▾

Save

🔗

⚙️

Top Agents (M Tokens)

AI Software Engineering Agent

1.04M

Customer FAQ Assistant

942.11K

Customer Complaint Response Agent

655.28K

Select all

Deselect all

📄 Group by

+

	AI Agent	AI Agent Type	User ID ↑	2026-Q1	2026-Q2	2026-Q3	2026-Q4		Totals
				M Tokens	M Tokens	M Tokens	M Tokens		M Tokens
	Total			3,831,673	2,033,734	0	0		5,865,406
☐	Customer FAQ Assistant	🤖: Agentic	amydonegal	315,374	221,870	0	0		537,244
☐	Marketing Content Review Agent	🤖: Agentic	amydonegal	262,812	184,892	0	0		447,704
☐	Card Dispute Intake Agent	🤖: Agentic	carolynMurphy	245,291	172,566	0	0		417,857
☐	Loan Operations Copilot	🤖: Agentic	cherylAmos	7,819	13,624	0	0		21,443
☐	Customer FAQ Assistant	🤖: Agentic	christopheLefevre	6,842	11,921	0	0		18,763
☐	Marketing Content Review Agent	🤖: Agentic	christopheLefevre	5,864	10,218	0	0		16,082
☐	Employee Service Desk Agent	🤖: Agentic	danaLewis	6,353	11,069	0	0		17,423
☐	Loan Operations Copilot	🤖: Agentic	danaLewis	8,308	14,475	0	0		22,783
☐	Employee Service Desk Agent	🤖: Agentic	dianeStewart	25,986	0	0	0		25,986
☐	Regulatory Policy Review Agent	🤖: Agentic	dianeStewart	36,380	0	0	0		36,380

1-49 of 49

<<

<

>

>>

<< Collapse

Columns

Details

- Home
- Data Ingestion
- Hybrid Workforce
- Team Productivity Dashboard
- AI Productivity Dashboard
- Digital AI Investments
- AI Capitalization
- AI ROI
- Digital Thread
- Administration

Hybrid Workforce 00000022

AI Usage Dashboard Org Chart Hybrid Workforce Token Budget AI Usage Data



View
AI Usage Dashboard

Save

Configure

AI Budget

16.4M

Spent to Date

9.37M


Remaining

7.03M


Ungoverned C...

2.43M

Actual Token Cost per User




Actual Tokens per User




Ungoverned Cost per Department




Department AI Budget vs Actuals



Actuals per Role



Actuals per CC



- 🏠 Home
- 🗄️ Data Ingestion
- 👤 Hybrid Workforce
- 📊 Team Productivity Dashboard
- 📈 AI Productivity Dashboard
- 🚀 Digital AI Investments
- 💰 AI Capitalization
- 🔍 AI ROI
- 🎵 Digital Thread
- ⚙️ Administration
- ⏪ Collapse

Hybrid Workforce 00000022

AI Usage Dashboard

Org Chart

Hybrid Workforce

Token Budget

AI Usage Data

📊

▼

View Hybrid Workforce

Save

🔗

⚙️

Resource Availability

5

Restore

1-46 of 46

⏪ < > ⏩

 Columns

Full Name	Person Type	Skills	2026-Q1	2026-Q2	2026-Q3	2026-Q4
AI Software Engineering Agent	Agent	Code Review; Secure Coding	6,238 / 6,328	33,916 / 33,916	42,192 / 42,192	50,469 / 50,469
AML Case Triage Agent	Agent	AML Case Management; Suspicious A	811,046 / 811,046	587,757 / 587,757	699,401 / 699,401	699,401 / 699,401
Amos, Cheryl	Employee	AI Ethics and Bias Mitigation; AI/ML M	424 / 512	560 / 520	268 / 528	264 / 521
API Documentation Agent	Agent	API Documentation; Technical Writin	34,422 / 34,422	62,441 / 62,441	70,846 / 70,846	79,252 / 79,252
Bauer, Joyce	Employee	AI Ethics and Bias Mitigation; AI/ML M	512 / 512	520 / 520	528 / 528	528 / 528
Branch Advisor Copilot	Agent	Customer Advisory; Retail Banking Pr	61,342 / 61,342			
Card Dispute Intake Agent	Agent	Customer Case Intake; Dispute Mana	5,925 / 5,745			
Chapin, Tom	Employee	Content Creation; Market Research	8 / 512	0 / 520	0 / 528	0 / 528
Cloud Runbook Remediation Agent	Agent	Cloud Operations; Runbook Automat	112,602 / 112,602	72,151 / 72,151	92,377 / 92,377	92,377 / 92,377
Cohen, Rita	Employee	AI/ML Model Training; HTML; JavaScr	478 / 512	696 / 520	634 / 528	438 / 528
Conner, Ian	Employee	Data Preprocessing; Deep Learning; I	512 / 512	520 / 520	528 / 528	528 / 528
Customer Complaint Response Agent	Agent	Complaint Handling; Customer Comr	107,610 / 107,610	210,135 / 210,135	240,892 / 240,892	271,649 / 271,649

- 🏠 Home
- 🗄️ Data Ingestion
- 👤 Hybrid Workforce
- 📊 Team Productivity Dashboard
- 📈 AI Productivity Dashboard
- 🚀 Digital AI Investments
- 💰 AI Capitalization
- 🔍 AI ROI
- 🎵 Digital Thread
- ⚙️ Administration

⏪ Collapse

Hybrid Workforce 00000022

AI Usage Dashboard

Org Chart

Hybrid Workforce

Token Budget

AI Usage Data

⬇️

🔍 Search...

< Per-Period Metrics - 2 Periods >

🔍 2

📄

📊

▼

View Tokenomics - Token Budget

Save

🔗

⚙️

Select all

Deselect all

Group by										
	1↓	Full Name 2↑	Primary Role	Person Type	Skills	Totals				
						Token Budget ...	Token Budget (M Uni...	Actual Token Cost	Actual Tokens (...	Budget Var...
	Tot...					9,020,000	9,655,821	7,790,000	6,373,979	1,230,000
☐	★	Amos, Cheryl	Engineer	Employee	AI Ethics and	131,147	491,803	37,440	140,401	93,707
☐	★	Bauer, Joyce	ML Engineer	Employee	AI Ethics and	419,516	521,989	145,841	180,272	273,675
☐	★	Chapin, Tom	Solutions Mark...	Employee	Content Crea	13,721	114,055	48,422	402,504	(34,701)
☐	★	Cohen, Rita	UX	Employee	AI/ML Model	104,421	10,442	26,275	2,627	78,147
☐	★	Conner, Ian	Engineer	Employee	Data Preproc	232,483	63,819	412,134	113,135	(179,650)
☐	★	Donegal, Amy	Analyst Relations	Employee	Market Rese.	11,796	105,479	37,549	335,757	(25,753)
☐	★	Evans, Nick	UX	Employee	AWS Certified	99,314	320,367	8,056	25,986	91,258
☐	★	Gillian, Erin	Engineer	Employee	AI Ethics and	142,932	111,666	80,451	62,852	62,481
☐	★	Gupta, Minesh	Engineer	Employee	AutoCAD	378,528	118,570	937,171	293,380	(558,643)
☐	★	Jones, Mike	QA	Employee	Adobe Creati	107,552	66,580	30,309	18,763	77,243
☐	★	Lefevre, Christophe	Product Market...	Employee	Market Rese.	13,131	110,846	42,276	356,863	(29,144)
☐	★	Lewis, Dana	Project Manager	Employee	AI Ethics and	937,045	940,402	119,026	246,474	818,019
☐	★	Lewis, Nicole	QA	Contractor	Code Review	112,571	11,257	48,826	4,883	63,745
☐	★	Maxwell, Faye	Engineer	Employee	AI/ML Model	215,423	59,315	348,267	93,927	(132,845)
☐	★	McCarthy, John	Product Owner	Employee	Agile Methoc	201,400	342,441	339,389	288,075	(137,989)
☐	★	McKay, Erin	Architect	Contractor	AI Ethics and	920,398	1,539,018	141,230	803,117	779,168


- Home
- Data Ingestion
- Hybrid Workforce
- Team Productivity Dashboard
- AI Productivity Dashboard
- Digital AI Investments
- AI Capitalization
- AI ROI
- Digital Thread
- Administration

Team Productivity Dashboard 00000025

Core Team Review Investment Hierarchy



View
Core Team - Next Gen Monthly Rev... Save [Link](#) [Configure](#)



Status Narrative

Core Team Delivery Update

The Core Team completed three NextGen Fraud Alert System features on June 5, 25 days ahead of plan:

- Account Lockouts
- High-Priority Fraud Alert Routing
- Suspicious Account Lockout

- 🏠 Home

📊 Data Ingestion

👤 Hybrid Workforce

📈 Team Productivity Dashboard

📈 AI Productivity Dashboard

🚀 Digital AI Investments

💰 AI Capitalization

🔍 AI ROI

🎧 Digital Thread

👤 Administration

AI Productivity

Unit of measure

FTE

Period

Dec 28, 2025 - Jan 2, 2027

Filter

2026

Save

Revert


AI Consumption VS Velocity

Feature Stats by Team


Planned VS Actuals by Epic

AI Efficiency by Team

Token Consumption




Token Cost



Team Velocity



Retired Scope



Tokens per FTE

Name ↑ ★	KPI	Per...	May 3, ...	May 10...	May 17...	May 24...	Ma
AI Team	Token Units per FTE	Actual	0.00	0.00	0.00	0.00	
API Team	Token Units per FTE	Actual	8,444.73	7,788.38	7,454.05	8,118.57	
Avengers Team	Token Units per FTE	Actual	1,023.58	1,108.48	1,083.94	1,052.21	
Business Operations	Token Units per FTE	Actual	0.00	0.00	0.00	0.00	
Core Team	Token Units per FTE	Actual	7,149.72	7,124.84	7,098.29	8,087.80	
Crushers Team	Token Units per FTE	Actual	1,574.05	1,579.25	1,437.92	1,427.48	
Global Operations	Token Units per FTE	Actual	122.48	136.33	137.33	130.74	

Velocity Uplift %

Name ↑ ★	KPI	Per...	May 3, ...	May 10...	May 17...	May 24...	Ma
AI Team	Velocity Uplift %	Actual	0.00	2.22	3.70	5.00	
API Team	Velocity Uplift %	Actual	0.00	3.85	7.69	11.54	
Avengers Team	Velocity Uplift %	Actual	0.00	1.43	2.86	4.29	
Business Operations	Velocity Uplift %	Actual	0.00	0.53	0.26	1.05	
Core Team	Velocity Uplift %	Actual	0.00	3.68	7.35	11.03	
Crushers Team	Velocity Uplift %	Actual	0.00	2.88	5.77	8.65	
Global Operations	Velocity Uplift %	Actual	0.00	0.63	1.56	2.50	

- 🏠 Home
- 🗄️ Data Ingestion
- 👤 Hybrid Workforce
- 📊 Team Productivity Dashboard
- 📈 AI Productivity Dashboard
- 🚀 Digital AI Investments
- 💰 AI Capitalization
- 🔍 AI ROI
- 🎧 Digital Thread
- ⚙️ Administration

AI Capitalization 00000024

AI Capitalization Dashboard

Properties

Investment Hierarchy

Token Usage Attestation

"Superhuman Contractor" Times...

Actual Tokens

Token Forecast

Capital Approval

My Timesheet

Review & Approve

Timesheets

EM

Erin McKay

Architect

I affirm this timesheet's work and token allocation are accurate.

Select Resource

Close

Time Period

< Jun 1 2026 - Jun 7 2026 >

Status

Open

> Collapse all

Add Tasks

Remove Tasks

Name	Token %	Mon. 1 8.00	Tue. 2 8.00	Wed. 3 8.00	Thu. 4 8.00	Fri. 5 8.00	Sat. 6	Sun. 7	Totals
High-Priority Fraud Alert Routing (FEA00000081)									
New Feature Development (NFD1)	40%	3.20	3.20	3.20	3.20	3.20			16.00
Research (RES1)	10%	0.80	0.80	0.80	0.80	0.80			4.00
Suspicious Account Lockout Detection (FEA00000083)									
New Feature Development (NFD1)	40%	3.20	3.20	3.20	3.20	3.20			16.00
Research (RES1)	10%	0.80	0.80	0.80	0.80	0.80			4.00

Digital Thread 00000027

Enterprise Value Flow


Enterprise Value Hierarchy

AI Spend & Workforce Governance

Delivery & Capitalization

Portfolio AI ROI & Scenario Planni...

Period 2026-Q1



- 🏠 Home
- 🗄️ Data Ingestion
- 👥 Hybrid Workforce
- 📊 Team Productivity Dashboard
- 📈 AI Productivity Dashboard
- 🚀 Digital AI Investments
- 💰 AI Capitalization
- 🔍 AI ROI
- Digital Thread
- ⚙️ Administration

Digital Thread 00000027

Enterprise Value Flow Enterprise Value Hierarchy AI Spend & Workforce Governance **Delivery & Capitalization** Portfolio AI ROI & Scenario Planni...

Period Multiple selections ▼

Enterprise AI Budget	Ungoverned Spend	Hybrid Scope Spend	Invoiced AI CapEx	Invoiced AI OpEx	Freed Capacity	Product Budget Envelope
\$16.40M	\$2.43M	\$9.27M	\$4.27M	\$6.87M	19.74K	\$12.50M

NextGen Fraud Alert Planned Costs Planned Labor CapEx \$635.71K Planned Labor OpEx \$434.09K Planned AI CapEx \$3.01M Planned AI OpEx \$5.22M Actual Costs Actual Labor CapEx \$348.30K Actual Labor OpEX \$316.38K Actual AI CapEx \$1.14M Actual AI OpEx \$2.26M Total Hybrid Actual \$4.07M Features Completed 10 Freed Capacity 5.23M Labor Cost Saved \$405.12K	Digital Banking UX Planned Costs Planned Labor CapEx \$378.36K Planned Labor OpEx \$750.58K Planned AI CapEx \$2.33M Planned AI OpEx \$2.65M Actual Costs Actual Labor CapEx \$92.89K Actual Labor OpEX \$252.67K Actual AI CapEx \$1.08M Actual AI OpEx \$1.01M Total Hybrid Actual \$2.44M Features Completed 4 Freed Capacity 3.67M Labor Cost Saved \$783.37K	Loan Calculator Tool Planned Costs Planned Labor CapEx \$900.45K Planned Labor OpEx \$555.15K Planned AI CapEx \$1.43M Planned AI OpEx \$0.00 Actual Costs Actual Labor CapEx \$101.13K Actual Labor OpEX \$0.00 Actual AI CapEx \$626.35K Actual AI OpEx \$0.00 Total Hybrid Actual \$727.48K Features Completed 2 Freed Capacity 2.16M Labor Cost Saved \$1.35M	Commercial Credit Planned Costs Planned Labor CapEx \$3.00K Planned Labor OpEx \$823.99K Planned AI CapEx \$428.11K Planned AI OpEx \$0.00 Actual Costs Actual Labor CapEx \$16.87K Actual Labor OpEX \$0.00 Actual AI CapEx \$196.61K Actual AI OpEx \$0.00 Total Hybrid Actual \$213.48K Features Completed 1 Freed Capacity 1.04M Labor Cost Saved \$810.12K	Loan Assist eTransfer Planned Costs Planned Labor CapEx \$4.00K Planned Labor OpEx \$0.00 Planned AI CapEx \$1.04M Planned AI OpEx \$0.00 Actual Costs Actual Labor CapEx \$52.67K Actual Labor OpEX \$0.00 Actual AI CapEx \$471.48K Actual AI OpEx \$0.00 Total Hybrid Actual \$524.15K Features Completed 2 Freed Capacity 523.42K Labor Cost Saved (\$48.67K)
---	--	--	--	--

When this becomes your problem

Probably not this quarter. And that's fine.



Cheap now. Expensive later.

1

How you buy it

Ask procurement what structure your AI agreements use. A flat access fee forecloses capitalization no matter how well you instrument.

2

Key and deployment structure

The reporting ceiling is set here, once. Re-keying a live estate later is genuinely unpleasant.

3

Tagging settings

Centrally managed, or left to individuals. One file now versus a hundred developer machines later.

4

Where usage lands

On the work item, or on the project. This decides what you can ever say about it.

No budget. No workstream. No approval.

What to watch for

Your first agentic capability goes to production

AI spend crosses the threshold where your finance partner starts asking

Your next AI contract renewal or true-up comes around

A business unit asks why they're paying for someone else's usage

Until one of these fires, this is genuinely tomorrow's problem.

If you change one thing

Admins

Make the tagging settings centrally managed.

One file you push out, rather than a hundred developer machines later.

PMO leads

Measure coverage before attempting anything else.

Showback long before chargeback. And forecast in units of work, not dollars.

Sponsors

Find out who owns the AI invoice — and how it's contracted.

A flat access fee forecloses capitalization no matter what anyone downstream does.

None of this is the project. It's the thing that makes the project possible later.

Atlas, completed

\$24,000

Spent

\$12,080

(of the \$15,100)

Substantiated, not guessed

80%

Coverage

And we know what the other twenty percent is.

Nothing here required a new discipline.

Rate. Driver. Forecast. Actual. Variance. And now coverage.

**This probably isn't your most urgent problem...yet.
Just start the record now. You can't build it backward.**

Questions?



Survey



Please take a few moments to fill out the class survey.

Your feedback is extremely important for future events.



Elevate Your Professional Expertise with Rego University Certifications

Rego is excited to continue our **certification programs**, designed to enhance your expertise in Clarity administration, leadership, and technical skills. These certifications provide hands-on experience and knowledge to excel in your career.



Certification Requirements:

✓ **Completion:** 12 units per certification track

✓ **Eligibility:** Open to all Rego University attendees



Important Reminder:

To have your certification **credits tracked**, ensure you **complete the class surveys in the app** after each session. This step is critical for certification progress.

Thank You For Attending Rego University

Instructions for PMI credits

- Access your account at pmi.org
- Click on **Certifications**
- Click on **Maintain My Certification**
- Click on **Visit CCR's** button under the **Report PDU's**
- Click on **Report PDU's**
- Click on **Course or Training**
- Class Provider = **Rego Consulting**
- Class Name = **regoUniversity**
- Course **Description**
- Date Started = **Today's Date**
- Date Completed = **Today's Date**
- Hours Completed = **1 PDU per hour of class time**
- Training classes = **Technical**
- Click on **I agree** and **Submit**



Let us know how we can improve!
Don't forget to fill out the class survey.



Phone

888.813.0444



Email

info@regoconsulting.com



Website

www.regouniversity.com