



Past the Pilot: Turning AI Investment into Realized Value

A recap of the “From AI Strategy to Scalable Enterprise Impact” panel hosted by SAS, in partnership with Red Hat, First Citizens Bank, and Hylaine.

<p>SAS</p>  <p>Jay Upchurch CIO and GM, Hosting Business</p>	<p>FIRST CITIZENS BANK</p>  <p>Ash Kaduskar Head of AI and Advanced Analytics</p>	<p>RED HAT</p>  <p>Stefanie Chiras SVP, Innovation Hub</p>	<p>HYLAINE</p>  <p>Adam Boitnott Founder and CEO (Moderator)</p>
--	---	---	--

Author

Justin Goff

Director of Technical Delivery, Hylaine

Contents

- 01** Foreword
- 02** The Shiny Toy Problem
- 03** Past P&L: Intent and Outcome
- 04** Governance as an Accelerator
- 05** Data That Carries AI
- 06** The Business and Technology Truce
- 07** Build, Buy, or Customer Zero
- 08** The Regulatory Moving Target
- 09** From Insight to Execution
- 10** About the Panel

01 Foreword

Every technology cycle picks up its own superlative. AI has collected more than most. Bigger than the internet. Bigger than the printing press. Bigger than anything the industry has ever seen.

The phrases are useful for headlines. They are almost useless for deciding what to fund next quarter.

That gap is what brought four executives together on the SAS campus in Cary, North Carolina, for the “From AI Strategy to Scalable Enterprise Impact” panel discussion. The conversation was hosted by SAS, supported by Red Hat and Hylaine, and moderated by Adam Boitnott, Hylaine’s founder and CEO. The panelists were Jay Upchurch, CIO of SAS, Ash Kaduskar, Head of AI and Advanced Analytics at First Citizens Bank, and Stefanie Chiras, SVP of the Innovation Hub at Red Hat.

The organizers set the tone at the start. The session was not about the hype. It was not about tools. It was a leadership conversation about trusted data, governance, operating models, and aligning business with technology. Those are the ingredients that decide whether AI scales or stays stuck in pilot.

This paper is our recap. It is organized around the themes that surfaced during the panel, in roughly the order the panelists raised them, with their own framing carried through. At a few places, my team at Hylaine has added a short perspective on what it looks like to operationalize these ideas inside a client engagement. Those contributions are meant to complement the panel, not to relitigate it. The panelists are the stars of this paper. Our role is to help translate their insights into execution.

02 The Shiny Toy Problem

THE HYPE

- Bigger than the internet.
- Bigger than the printing press.
- Bigger than anything the industry has ever seen.
- Vendors overselling generative AI.
- Neighbors who say it can be done in two days.

THE REALITY

- Legacy systems. Skeptical users. Imperfect data.
- Risk tollgates, legal review, regulatory assessment.
- Governance speed bumps at the 11th hour.
- Demos that don't survive contact with the enterprise.
- Two-day prototypes that take three months to deploy.

The failure mode is not the technology. It is treating a technology proof like a business solution.

Stefanie Chiras opened by revisiting the first time AI genuinely surprised her. She had deployed Cursor for the first time, built a working website in under an hour, then walked her daughter through changing the color of the whole site with two words.

"It was amazing."

— Stefanie Chiras, Red Hat

That reaction is common. It is also, in her telling, the root cause of most failed pilots. Teams fall in love with the wow moment. They recreate it in isolation, on curated data, without legacy enterprise systems or skeptical end users. They mistake a technology proof for a business outcome.

Ash Kaduskar described the same phenomenon from inside the enterprise. First Citizens Bank is actively managing about 60 AI use cases today. Most are product-driven, because every vendor in his procurement pipeline is embedding AI into their platform. Some of those vendors, he noted, are overselling their use of generative AI, and that overselling has a direct cost. If a vendor claims generative AI, the tool has to pass through more stringent risk tollgates, even when the underlying product is not using generative AI at all. Adoption slows. Reviews multiply. The cost of the hype lands on the enterprise.

Adam offered the same gap from the consulting side. Hylaine had scoped a seven-to-eight-month AI engagement with a client. The client's neighbor told them it could be knocked out in two days. Joking or not, that is a tough perspective to land against.

Ash answered the anecdote directly.

“Doing a demo on your personal computer using Cursor, Claude Code, or Replit is a completely different ballgame as compared to deploying the same baby demo website into an enterprise.”

— Ash Kaduskar, First Citizens Bank

The developer may not see the full picture of enterprise governance, legal review, compliance, or regulatory assessment. That is the picture that turns two days into three months.

Across all three panelists the theme was consistent. The failure mode is not the technology. It is treating a technology proof like a business solution.

Put Into Practice

In my work leading delivery at Hylaine, the fastest way to deflate the shiny toy reflex is to give the builder the enterprise checklist up front. If a developer is excited about a demo, the fastest path to maturity is walking them through the fifty questions that privacy, compliance, legal, cyber, and third-party risk will ask in production. That transparency does not kill the enthusiasm. It redirects it toward problems worth solving.

03 Past P&L: Intent and Outcome

Adam Boitnott turned the panel toward the question underneath all of this. **How should leaders actually prioritize use cases so they deliver measurable value and not just impressive demos?**

Ash’s answer gave the conversation its sharpest framework. He started with the cost reality. AI is not cheap. The people you hire are not cheap. Infrastructure is not cheap. Usage is not cheap. He noted that developers in the Bay Area are now spending more on tokens than they earn in salary, and Bay Area developers are not paid poorly.

His conclusion followed naturally.

“So anything that you do, if it doesn’t impact P&L, just stop it.”

— Ash Kaduskar, First Citizens Bank

P&L impact is the table stake. Everything above it, Ash broke a six-box checklist into two dimensions, three intents, and three outcomes.

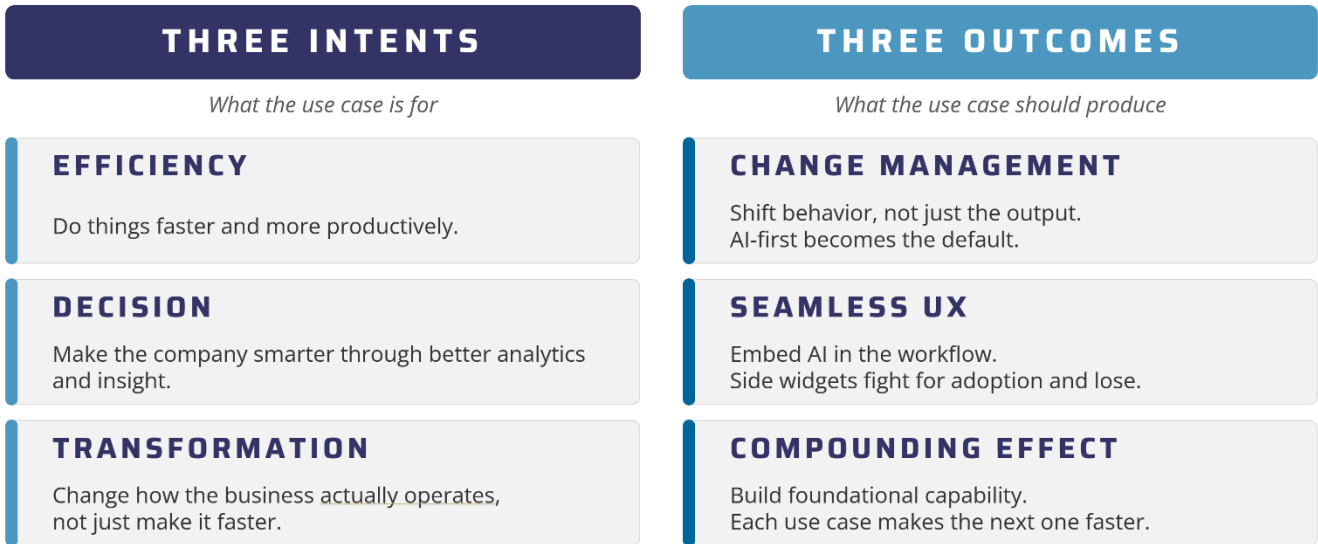


TABLE STAKE → P&L Impact *Without a P&L tie-out, stop the use case.*

Rule of thumb: a strong use case checks at least three of the six boxes, with P&L impact as the prerequisite.

The three intents

Every use case should align with at least one of three intents.

- **Efficiency:** doing things faster and more productively.
- **Decision:** making the company smarter through better analytics and insight. Ash flagged this as particularly underserved, noting that generative AI augmenting predictive AI and ML in the analytics stack is a space primed for value.
- **Transformation:** changing how the business actually operates, not just making the existing workflow faster.

The three outcomes

Beyond intent, he evaluates what a use case is actually producing.

- **Change management:** the shift in behavior, not just the output. Summarizing an RFP response is an outcome. Changing how the organization thinks about RFPs, so that everyone assumes an AI-first approach, is a behavior change.
- **Seamless user experience:** embedding AI into the workflow rather than bolting on a side widget. Side widgets fight for adoption and usually lose.
- **Compounding effect:** the one closest to his heart. If a use case builds a foundational capability that makes the next five use cases faster, it has a flywheel built in. If it does not, it is a one-off with a ceiling.

His test is simple. A strong use case checks at least three of those six boxes, with P&L impact as the prerequisite.

Jay's enterprise decision rubric

01 Business Intent	02 Alignment	03 Readiness	04 Success Criteria	05 Governance
What is the problem? Is the team passionate about solving it?	Are technology, business, legal, and cyber teams in sync?	Is the organization ready? The data? The development team?	Is there P&L impact? Is the vision defined?	Who raises the child after go-live? Use cases live on.

Jay added a complementary rubric his team uses for larger AI investments. Business intent and a clear problem definition. Alignment across technology, business, legal, and cybersecurity. Readiness of the organization, the data, and the development team. Success criteria that include P&L impact and a stated vision. Governance plans for ongoing management after go-live.

"Who's going to raise the child at some point? Because these things live on."

— Jay Upchurch, SAS

Use cases outlast their go-live dates. Someone has to manage them.

Stefanie closed the theme with the line that stuck. When she advises startups, she tells them not to fall in love with the technology. Fall in love with the problem.

Put Into Practice

The opportunity mapping work my team does at Hylaine sits directly in this space, and our framing is deliberately industry-tuned. In banking and financial services, the highest-leverage opportunities tend to cluster in compliance, risk, fraud, and customer operations, where the compounding effect is real and P&L tie-outs are clean. In insurance, claims triage, underwriting support, and broker enablement dominate. In healthcare and life sciences, prior authorization, clinical documentation, and regulatory submissions are the first places to score a use case against these six boxes. The framework Ash laid out is remarkably close to how we evaluate engagements in every one of those verticals.

04 Governance as an Accelerator

THE OLD WAY

1. Stand up an oversight committee.
2. Write an acceptable use policy.
3. Tell teams to go fast.
4. Hit a governance speed bump.
5. Risk team appears at go-live and asks for 30 days.

THE NEW WAY

1. Build risk and compliance checks into the process.
2. Bias, drift, & lineage controls are part of the build.
3. Data ethics, legal, cyber, & tech engage from day one.
4. Governance accelerates delivery, not gating it.
5. Whole company participates, not just pilot teams.

Adam cited a statistic that landed hard. A report from Grant Thornton on AI Impact surveyed 950 business leaders across 10 industries from Feb. 13 to March 18, 2026, and found 78 percent lacked confidence they could pass an AI governance audit within 90 days.

Jay reframed the question before he answered it. The assumption built into the statistic, that speed and governance are at odds, is not always true.

He walked through the arc he had seen at SAS. In the early days of generative AI, enterprises led with governance. Everyone wrote an acceptable use policy. Everyone stood up a cross-divisional oversight committee. Then the organization was told to go fast. The result was predictable. Teams hit a governance speed bump. Compliance showed up at the end of the project, where the risk team asked for thirty days right before go-live and slowed the rollout.

The alternative, he argued, is to build speed and governance into the same process. Risk and compliance checks move with the work, not after it. Bias detection, drift monitoring, and transparency controls are part of the build, not an audit on top of it. When the organization sees that governance actually accelerates delivery, curiosity and creativity grow instead of shrinking, and the whole company participates rather than a few pilot teams.

“If you can put speed and governance together in the process that you build, it actually becomes an accelerator.”

— Jay Upchurch, SAS

Ash's three baby steps

1

AGREE ON PATTERNS

Classify risk into simple, medium, and high. Align the full risk team on the shapes before you address individual cases.

CROSSED

2

STANDARDIZE QUESTIONS

One canonical question set per pattern. Stop assessors from asking a different 50 questions depending on what they read last night.

IN PROGRESS

3

AUTOMATE

Generative AI runs the first pass. Human reviews the output. Proven in a single risk area. Full rollout projected within ~2 months once Step 2 lands.

PROJECTED

Standardization before automation. Patterns before questions. Humans in the loop on the outcomes.

Ash took the framing into practice with a three-step approach at First Citizens Bank.

Step one: agree on patterns. Risk comes in shapes. Some use cases are simple, some are medium, some are high. Identifying and agreeing on those patterns with the full risk team is the first baby step.

Step two: standardize the questions. Ask most banks for the complete set of questions their risk team uses to assess AI, and no one has it. Assessors rotate by availability, not expertise, and each one asks a slightly different set based on whatever article or video they saw the night before. That inconsistency is not a minor inefficiency. It is a quality problem in its own right.

Step three: automate. Once patterns and questions are stable, the actual risk assessment is largely a pattern-matching exercise. Generative AI can run the first pass. A human in the loop reviews the output. Ash reported that his team has already proven the concept against a single risk area. Once the standardized question set is finalized, he projects that the bank's full AI risk assessment can be largely automated within roughly two months.

The stakeholders that need to be at the table

Jay named the groups that need to be part of governance from the beginning, rather than showing up at the end. SAS's data ethics practice chairs the AI oversight committee. Legal is active throughout. Cybersecurity tracks data lineage so the company can respond to regulators when they ask. Technology keeps the patterns repeatable so the company can scale without accumulating one-off decisions.

Shadow AI is already here

400+

AI AGENTS FOUND

Discovered running in a single enterprise environment via a Microsoft Agent 365 scan.

1 in 3

EMPLOYEES USING THEM

No one asked for a literacy program. No one waited for a policy. The agents were already there.

78%

CANNOT PASS AN AUDIT

Of 950 C-suite executives surveyed, most lack confidence they could pass an AI governance audit today.

You cannot govern what you cannot see. The inventory work has to come first.

Jay then shared a number that landed as hard as the 78 percent. SAS ran an internal scan using Microsoft's Agent 365 tooling. They found over 400 individual AI agents already running in the environment. That is roughly one agent for every three employees. Nobody had requested a literacy program. Nobody had waited for a policy. The agents were already there.

The next question was not whether to allow them. It was whether they were safe, whether they were governed, and which ones had enterprise applicability worth hardening and releasing to the broader organization.

Put Into Practice

Shadow AI at this scale is the governance challenge of the next twenty-four months. The patterns we recommend at Hylaine track closely to what Ash and Jay described. Start with tiered risk patterns. Standardize the assessment question set before you automate anything. Build a lightweight inventory of what is already running in your environment before you try to govern what you think is there. The 78 percent figure is not a warning. It is a map of where the audit work has to happen first.

05 Data That Carries AI

STRUCTURED DATA	UNSTRUCTURED DATA
<p data-bbox="334 443 610 468">SLOWER TO UNLOCK</p> <ul data-bbox="175 506 768 785" style="list-style-type: none"> • Historically the focus of roughly 99% of Chief Data Officers. • Still flows through multi-year governance and modernization programs. • On-premises to cloud migrations of 8–12 months are common. • Sequencing matters: align data migration with AI use case priority. 	<p data-bbox="997 443 1287 468">READY FOR USE NOW</p> <ul data-bbox="846 516 1422 785" style="list-style-type: none"> • Documents, images, videos, audio, social content. • Previously outside the remit of most CDO organizations. • Generative AI goes leaps and bounds on unstructured data. • High-value domains: legal, procurement, call center, compliance, fraud, marketing.

The data question came from the audience. **AI is moving faster than the data foundations underneath it. How should leaders handle that mismatch?**

Stefanie framed it around trust. Is the data decision-grade? Is it governed? Is someone accountable for it? Is it safe, secure, and transparent in its lineage? All of that sits under the AI conversation, and it does not get easier as platforms proliferate. The platform sprawl inside data today, she noted, looks a lot like where AI is now.

Jay acknowledged the tension directly. He worries about companies that hold back their AI work because they are waiting for pristine data. Most organizations do not have that luxury of time.

Ash cut to the resolution.

“Data is never going to be ready. It is a perennial problem.”

— Ash Kaduskar, First Citizens Bank

No institution in the world can honestly say its data is clean, complete, and AI-ready. Waiting for that state is waiting forever.

The practical shift, he argued, is to recognize that generative AI changed the data conversation entirely. Until about three years ago, roughly 99 percent of Chief Data Officers focused only on structured data. Documents, images, videos, audio, and social content were largely outside the remit. Generative AI inverted that. Structured data still moves slowly through governance and modernization programs. Unstructured data is ready for use now, and generative AI goes leaps and

bounds on unstructured data. Legal, procurement, call center, compliance, fraud detection, and marketing are all domains where unstructured data can deliver massive value without waiting for a structured-data program to finish.

Ash also described how First Citizens Bank is sequencing its on-premises to cloud migration. The bank had already prioritized which data sets would move first. His team is now cross-aligning that sequence with the AI use cases and processes the bank wants to transform. The sequencing is how that work moves from theoretical to applied.

Stefanie added a forward-looking note. Frontier models are commoditizing. The next decade of differentiation will come from small models built on well-curated proprietary data. She described a Boston-area startup she recently met with that collects marine data along the Northeast coast. The company does not sell the data. It sells small models built on that data, customized for construction, fishing, or environmental monitoring clients. The data asset became the product. That pattern, she believes, is coming to the enterprise.

06 The Business and Technology Truce

Adam asked about the historic push and pull between business and technology. **Has AI made that dynamic better or worse?**

Stefanie's answer was clear. AI has made it harder, not easier, because it has lowered the cost of building technology for non-technologists. A domain expert with an AI tool can now iterate in the line of business without waiting for IT. That dependency reduction is a good thing for velocity. It is a bad thing for outcomes when the collaboration between business, IT, legal, and risk is skipped entirely. The output is weaker than it could have been, and the gaps show up late.

Her more hopeful framing was that AI is an excellent interpreter of expertise across domains. When legal comes to the table early and says, here is what we need you to consider, when the line of business shows up with the P&L outcome, and IT shows up with the infrastructure constraint, the solution arrives faster and stronger than it would if each team waited to hand work off in sequence. Concurrent collaboration replaces serial handoffs.

THE THROUGH-LINE

AI scales when the discipline around it is real.

It stalls when the discipline is ceremonial.

Jay extended the point from inside SAS. AI literacy has progressed in phases. Early apprehension, then the governance speed bump, then the wave of pressure to deliver, and now the flood of demand that has outstripped capacity. SAS is in the prioritization phase now, sifting through which ideas are worth enterprise investment and which should stay inside the sandbox where they came from. That progression, he noted, is roughly the shape most enterprises follow, though not at the same pace.

07 Build, Buy, or Customer Zero

SAS

Customer Zero

- IT takes software straight out of CI/CD from R&D.
- First test case for every internal AI tool.
- Tight feedback loop between builders and operators.
- Learning compounds. Friction exists, but outcomes come faster.

FIRST CITIZENS BANK

Build common, buy niche

- BUILD common patterns in house: knowledge mgmt, decision engines.
- BUY niche, domain-deep tools that need continuous specialized care.
- Avoid generic widgets. A startup will absorb them in 12 months.
- Buy where the domain requires continuous specialized investment.

An audience question raised the tension between internal operations and product lines when both are using AI. **When does an organization use its own product, and when does it buy?**

Jay described SAS’s Customer Zero approach. His team takes software straight out of CI/CD from the R&D organization. That carries advantages. They become the first test case. Some of those trials go well. Some do not. Either way, the loop between product R&D and internal IT sharpens both sides. Learning happens faster when the builder and the buyer are talking every week.

He named the friction honestly. There is always going to be a question of whether the internal option or an outside alternative is the right fit, and sometimes the internal team is pointed toward its own tools for reasons other than pure product fit. The trade-off is worth it, in his view, because the shared learning between R&D and IT delivers faster outcomes than either group could alone.

Ash offered First Citizens Bank’s philosophy in a clean one-liner. Build common patterns in house. Buy niche, domain-specific tools that require continuous specialized investment his team cannot sustain on its own.

“If you build a widget, tomorrow some startup comes in and builds a product, and your widget is already stale. So we don’t want to get into that game.”

— Ash Kaduskar, First Citizens Bank

His reasoning is practical. If a vendor is deep in a specific domain, delivering continuous investment to a niche problem, that is the right place to buy. Generic widgets that a product company will absorb into its platform inside twelve months are the wrong place to build.

08 The Regulatory Moving Target

Adam pulled the conversation toward the regulatory horizon. **Are startups addressing the governance and audit problem directly?**

Stefanie's answer was careful. Startups are focused on visibility and observability. Observability is one of the most active spaces in AI today, precisely because the regulatory environment is still a moving target. Nobody is confident enough in what the regulations will require to pitch compliance as a feature. Everyone is confident enough in the need for visibility to build the instrumentation layer first. It is a pragmatic bet. You cannot govern what you cannot see.

Jay agreed. Lobbying efforts are already underway, both for and against AI regulation. The requirement on an IT shop to track what is running in the dark corners of the business is already hard, and it is getting harder. Third-party SaaS, multi-cloud environments, and AI embedded in off-the-shelf platforms compound the challenge. The regulatory work ahead is going to be substantial. The visibility work has to come first.

09 From Insight to Execution

If there was a single through-line across the evening, it was this. **AI scales when the discipline around it is real. It stalls when the discipline is ceremonial.**

The panelists described that discipline in different vocabularies. Ash talked about P&L as table stakes and his six-box test. Jay talked about integrating speed and governance as one process rather than two. Stefanie talked about falling in love with the problem rather than the technology. The language is different. The underlying idea is the same. Enterprise AI value is not created by a better model. It is created by a better operating model.

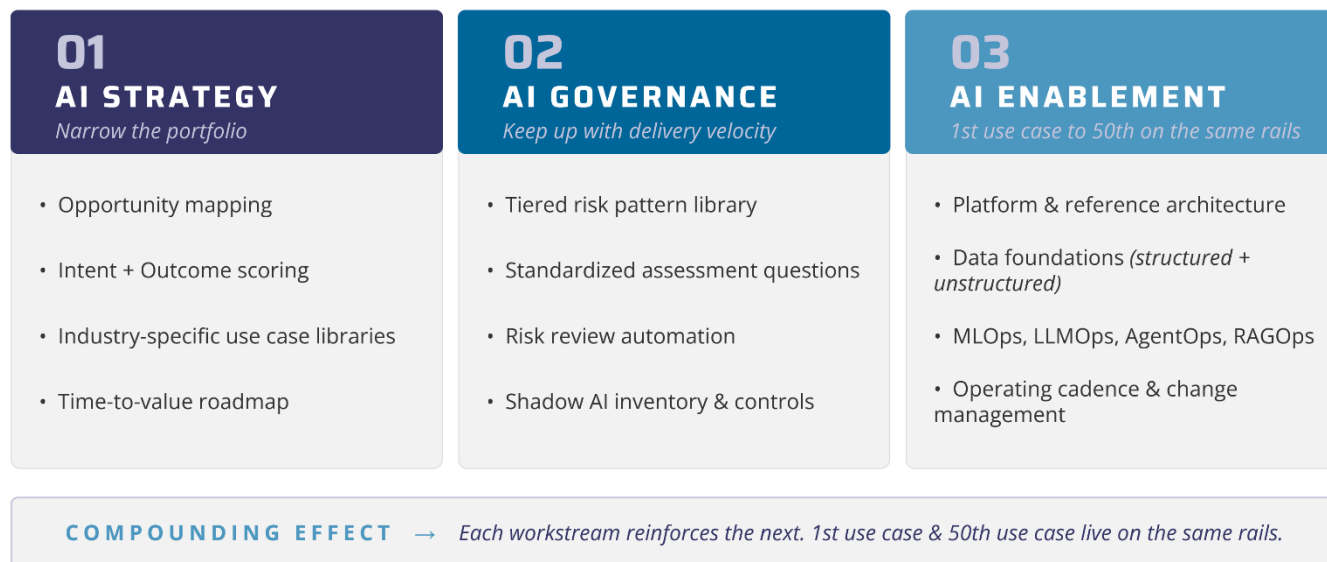
My view, and Hylaine's, is that the operating model is where most enterprises lose the most ground. The frameworks are available. The panelists shared several strong ones in a single hour. The gap is in operationalizing them across industries that each carry their own regulatory weight, data shape, and workflow reality.

<p>BANKING & FINANCIAL SERVICES</p> <p>HIGH-LEVERAGE USE CASES</p> <ul style="list-style-type: none"> • Compliance automation • Risk and fraud • Customer operations • Credit decisioning <p>Regulatory context: OCC, Federal Reserve, CFPB</p>	<p>INSURANCE</p> <p>HIGH-LEVERAGE USE CASES</p> <ul style="list-style-type: none"> • Claims triage • Underwriting support • Broker enablement • Policy servicing <p>Regulatory context: State-by-state DOI oversight</p>
<p>HEALTHCARE</p> <p>HIGH-LEVERAGE USE CASES</p> <ul style="list-style-type: none"> • Prior authorization • Clinical documentation • Care coordination • Revenue cycle <p>Regulatory context: CMS, HIPAA, state Medicaid</p>	<p>LIFE SCIENCES</p> <p>HIGH-LEVERAGE USE CASES</p> <ul style="list-style-type: none"> • Regulatory submissions • Clinical trial ops • Medical affairs • Pharmacovigilance <p>Regulatory context: FDA, EMA, GxP</p>

The playbook is the same. The execution is not. Each vertical carries its own regulatory weight, data shape, & workflow reality.

That gap is where our work lives. In banking and financial services, the operating model has to absorb OCC and Federal Reserve expectations alongside agentic automation that reaches into credit, fraud, and customer operations. Ash's approach at First Citizens Bank, building the risk assessment engine itself with generative AI, is exactly the kind of compounding move he described, applied to the governance function. In insurance, the operating model has to carry claims, underwriting, and broker enablement through state-by-state regulatory variability. In healthcare

and life sciences, the same operating model has to carry clinical documentation, prior authorization, and regulatory submissions under CMS, FDA, and HIPAA expectations that make even benign use cases non-trivial to deploy. The playbook is the same. The execution is not.



We organize that execution across three connected workstreams. AI Strategy narrows the portfolio to the use cases that actually clear the six-box test. AI Governance builds the patterns, the standardized question sets, and the automation that lets risk assessment keep up with delivery velocity. AI Enablement stands up the platform, the data foundations, and the operating cadence so the first use case and the fiftieth use case live on the same rails. Each workstream honors the principles the panel described, and each is designed to compound over time rather than solve for a single engagement.

Where this paper ends, and the next one begins

Adam closed the panel with an observation worth restating. This kind of conversation could happen every week, and the themes would continue to hold. The details will keep moving. The fundamentals will not.

Jay added the line that stayed with me. If you can show your organization that governance speeds up delivery rather than slowing it down, you pump fuel into the curiosity and creativity of the whole company. The shiny toys stop being the point. The business outcomes become the point. The whole organization participates.

That is the version of enterprise AI worth building toward. It is not about choosing between speed and discipline. It is about treating them as the same discipline. That is where my conviction sits. Insight is the easy part. Execution is the work.

CLOSING



Insight is the easy part. Execution is the work.

- Frameworks are available. The panelists shared several strong ones in a single hour.
- The gap is in operationalizing them across industries with their own regulatory weight, data shape, and workflow reality.
- That gap is where Hylaine's work lives.

Let's build the version of enterprise AI worth building toward.

10 About the Panel

The panel was hosted at the SAS campus in Cary, North Carolina, in partnership with Red Hat and Hylaine.

Panelists

- **Jay Upchurch**, Chief Information Officer and GM of the Hosting Business, SAS
- **Ash Kaduskar**, Head of AI and Advanced Analytics, First Citizens Bank
- **Stefanie Chiras**, Senior Vice President, Innovation Hub, Red Hat
- **Adam Boitnott**, Founder and CEO, Hylaine (Moderator)

About Hylaine

Hylaine is a values-first technology consulting firm. We partner with enterprises in financial services, insurance, healthcare, and life sciences to turn AI strategy into operating reality. Our work spans three connected workstreams: AI Strategy, AI Governance, and AI Enablement. To learn more, visit hylaine.com.