Christopher J Waller: Operationalizing AI at the Federal Reserve

Speech (via webcast) by Mr Christopher J Waller, Member of the Board of Governors of the Federal Reserve System, at the Federal Reserve Bank of Boston 2026 Technology-Enabled Disruption Conference "Shaping the future of finance and payments", Boston, Massachusetts, 24 February 2026.

* * *

Thank you, President Collins, and thank you for the opportunity to speak to you today.[1] Artificial intelligence is a technological phenomenon that is taking the world by storm. We read about it every day and many of us have used it in some way or another. In my lifetime, I have never seen a technological revolution like this-and I have seen the birth of space exploration, the rise of the personal computer, the explosion of the internet and then smartphones. While all transformative, none match the potential that AI has in changing our lives and doing so at breathtaking speed. Firms, households, and every government are all trying to incorporate it into the way they function and operate. And I am here to tell you that the Federal Reserve is no different.

So, given the theme of this conference and the audience makeup, I felt this was a good time to discuss how the Federal Reserve System is using artificial intelligence to build and optimize the systems that support our work, as well as embedding it into other internal applications.

Now most people associate the Federal Reserve with monetary policy-interest rates, inflation, and the decisions that make headlines when Fed officials meet eight times a year. But the bulk of our daily activity is doing operational work such as payments, financial management, human resources, and providing financial services to the U.S. Treasury. A critical element of this operational work is technology. AI is the latest technology that we are building into our daily work to achieve operational efficiencies.

The Fed was designed in 1913 as a system of regional banks, and historically many technology decisions were made Bank by Bank. That made sense in an earlier era, when each Reserve Bank effectively operated as a separate organization. But as the Fed's work has become more digital and interconnected-alongside the evolution of the banking system and the broader economy-that approach increasingly creates duplication, inefficiency, and operational risk.

When systems are interconnected, decisions made in isolation create coordination problems-especially given the scale and criticality of the systems the Fed operates. Maintaining central-bank standards of resilience and security requires a more coordinated model.

That's why we're moving toward a Federal Reserve System-first approach-with shared standards and infrastructure, while preserving decentralization where it matters most, particularly for monetary policy and economic research.[2]

And the urgency of making this transition is growing. The volume and velocity of technological change continue to increase. As a public institution with an important role

in the U.S. and global financial systems, the Fed must keep pace to deliver effective, reliable services alongside the private sector.

That velocity of change was evident at the recent Payments Innovation Conference I hosted, where a range of participants and attendees talked about the convergence of AI, stablecoins, tokenization, and payments-a theme I also hear regularly in my conversations with industry and one I suspect will surface throughout today's panels.

At that pace of change, a Bank-by-Bank approach simply won't work, especially for frontier technologies such as tokenization, quantum computing, and generative AI. These are bold challenges-and opportunities-that cut across the Federal Reserve System.

Meeting the moment requires Systemwide coordination, a bias for action, and disciplined execution at scale. That's why we deliberately built a System-centric innovation practice. Rather than duplicating effort-doing the same thing a hundred times across the System-we can instead do a hundred different things. A System approach allows us to move ideas from experimentation into execution more quickly and efficiently.

AI is a case study of what this approach looks like in practice. I don't have to tell this audience how quickly AI is moving-reshaping how work gets done and how organizations operate. As builders and leaders, you experience it firsthand. And the Federal Reserve is no exception. It's imperative that we keep pace. Yes, we're a central bank; "break things and ask forgiveness" won't work here. With great power comes great responsibility. AI systems can amplify errors as quickly as they amplify efficiency. They can hallucinate. They can introduce real risks around data protection, model risk, bias, and operational resilience. We cannot approach AI casually. As a central bank, we hold ourselves to a high standard. That means clear guardrails on how and where it's used, strong information-security controls, rigorous model validation, human accountability for decisions, and ongoing evaluation as the technology evolves. Innovation and risk management are not competing priorities here - they reinforce each other.

We need to respect these principles, but passivity isn't an option either. We can't afford to be late or to examine transformational change in a fragmented way. So, we're doing this differently. In implementing AI across the Fed, we're moving as one System, with shared direction and alignment. We have developed a common internal general-purpose AI platform for all Reserve bank employees to use. Our approach is intentionally business-led and AI-enabled. We start with the problem to be solved and the business need, then apply the right capability from across the AI stack. That discipline helps us deliver real business value while avoiding unnecessary complexity and cost.

Being business-led also means building AI into how the Fed operates-not treating it as a collection of random experiments or zombie projects.

The goal isn't novelty. It's utility.

So, let's get into the AI use cases in action. We're deploying AI in three focused, complementary ways that reflect how work gets done across a large, complex institution-broad access for all employees, specialized tools for builders, and embedded capabilities within enterprise workflows-which together make it part of how the Federal Reserve operates day to day.

Let me start with general-purpose AI for all employees, because that's where the biggest day-to-day impact shows up first.

This is about making AI a baseline capability of everyday work, not a niche tool. Every employee has access to Fed approved AI solutions they can use throughout the day-to draft, summarize, analyze information, and get unstuck faster. For many, it functions as a pervasive digital assistant-a sounding board they can return to as they work through problems and complete daily tasks. The goal isn't to turn everyone into a technologist; it's to reduce friction in routine work so people can spend more time on judgment, problem solving, and higher-value activities.

In many ways, this mirrors how people already use AI in their personal lives. In my own household, AI has become an everyday tool-much like a smartphone or even a microwave you rely on without thinking about it. My wife uses it as a daily companion: planning trips, helping our kids think through career options, comparison-shopping, or turning small, annoying tasks into something more manageable.

It's not exotic. It's a tool.

And that's exactly how we should be thinking about AI at work as well.

Let me make this real.

Fed staff often receive substantial non-FOMC background materials to help prepare for various meetings. To help synthesize the information, they sometimes use the Fed's internal general-purpose AI tool to generate key themes quickly. Of course, it does not replace preparation or judgment. It compresses the mechanical work so more time can be spent on substance and the questions that matter.

Another example is when a colleague returned from a well-deserved vacation-a true digital detox, meaning she did not have access to work device or emails-and she came back to a full inbox and document queue. Rather than spending days sorting through everything, she used the Fed's internal AI tool to summarize and triage what had accumulated. That allowed her to go straight to what required her expertise.

In both cases, the tool handles the volume and the first pass. The human makes the decisions.

The second area where we're seeing real impact is with developers and hands-on builders-the people turning ideas into implementation.

Coding assistants are helping developers optimize work across the software development life cycle-from documentation and refactoring to writing code and unit

testing. This helps teams work through backlogs faster, improve quality and reliability, modernize systems, and ship more value and innovation.

But this isn't just about speed.

AI is taking on some of the most time-consuming and least satisfying parts of software development, so developers can focus on security and quality. That matters for an institution like the Federal Reserve, where reliability and resilience in production systems are critical.

Take unit testing. It's essential for quality and resilience, but it's not the part of the job developers get excited about. In several teams, tasks that used to take days are now completed in hours with AI assistance. One developer told me directly, "What used to take me two days now takes two hours." This frees up time for higher-value work like strengthening security and building new capabilities. As these tools mature, the benefits compound.

There's a broader point about capacity. When coding assistants lower the cost of producing software while improving quality, they expand what's possible. We can write more code, build more capabilities, and deliver more business value. As scarcity eases, capacity rises-allowing us to tackle backlogs and technical debt that accumulate over time.

A helpful analogy I've heard is the iPhone and photography. Putting a camera in everyone's pocket didn't eliminate professional photography. It lowered the cost of production, increased volume, and expanded the market. More photos were taken, and demand for high-quality work actually grew. I think coding assistants will work the same way for software.

At the Fed, we're already seeing strong early uptake-with hundreds of developers adopting these tools quickly-which tells us this capability is meeting a real need.

Let me give you a different kind of example-one that's less about code and more about listening to the communities we serve.

Across the Federal Reserve, we gather an enormous amount of qualitative information-conversations with businesses, community leaders, and market participants. Historically, synthesizing that information across regions and time periods has been labor-intensive.

Using AI tools, analysts can now pull targeted themes from large volumes of interview notes, compare patterns across cycles, and surface shifts in sentiment much more quickly. That doesn't replace human judgment-it accelerates the first pass so economists can spend more time interpreting what matters.

The third way we're putting AI into production is by embedding it directly into the workflows people already use.

Rather than asking teams to adopt entirely new tools or build bespoke solutions, we're activating AI capabilities within the platforms that already support day-to-day work across areas like legal, risk, procurement, operations, and other enterprise functions.

Adoption follows workflow. When AI is embedded-not bolted on-people don't have to change how they work to get value from it.

If you're a frequent online shopper or traveler, you've probably experienced this shift already. When something goes wrong-a delayed flight, a missed connection, a damaged package-it's often faster now to resolve the issue through a simple chat or text, or even that often-dreaded call. In many cases, AI is working behind the scenes to summarize context, route the issue, or resolve it outright-and when a person is involved, they're better equipped to help. The experience is simpler, faster, and often better.

That same dynamic applies at the Fed. By embedding AI into existing enterprise systems, we can improve speed, consistency, and quality of service without creating fragmented solutions. It's also fiscally responsible. Given how quickly the technology is evolving, consuming AI through vendor platforms allows us to benefit from ongoing improvements, rather than building and maintaining tools that can become costly or stale.

Taken together, these examples show how we're moving AI from exploration to execution through a coordinated, system-first approach.

The result is increased capacity across the organization, enabling teams to tackle complex problems and deliver more value, while improving productivity and cost efficiency through responsible innovation.

While this has been a solid start, we're not interested in being one-hit wonders. What matters is durability and that demands a focus on adoption, accountability, and leadership.

This is where many AI efforts succeed or fail-when the baton shifts from early adopters to day-to-day operators. The technology isn't the hard part anymore; change management is. It comes down to how quickly people adopt the tools, how deeply they embed them into daily workflows, and whether that adoption translates into results.

That's why we've taken an adoption-first approach. We treat AI as a human-capital investment, not side-of-desk experimentation. Training and upskilling happen on paid time, not nights and weekends.

And that training isn't one-off or theoretical. It's evergreen, hands-on, and role relevant. Employees are learning by using AI in real workflows, through applied workshops, hands-on education sessions, and prompt jams. This "hands-on-keys" approach matters, because comfort and competence come from use, not from slides.

We've also been explicit about expectations. AI use is not optional. Baseline literacy and application are being built into employee performance goals across the System. What gets measured gets done.

I've seen this firsthand. When I was Research Director in St. Louis, we crafted a strategic plan that largely sat on a shelf. What changed behavior was putting those priorities directly into employee goals. Once people knew what mattered and how they'd be measured, execution followed. That experience has shaped how I think about making change stick.

Leadership plays a critical role here. Setting expectations and making investments is necessary, but not sufficient. Teams need to see leaders walk the talk and communicate that they are committed-and that they are students of this technology themselves. That leadership signal is what turns early momentum into sustained behavior change.

This is how we move from early wins to lasting capability by pairing technology with training, accountability, and leadership, so AI becomes a durable part of how the Federal Reserve operates.

Conferences like this one focus on how technology is reshaping the future. What I've tried to show today is how we're approaching that challenge at the Federal Reserve-delivering our mission through technology-enabled innovation, with a clear focus on execution and efficiency. And in doing so, we demonstrate how a public institution can adopt AI responsibly and in a way that strengths public trust.

As we see technologies like tokenization and agentic AI coming into view, it's worth remembering that this isn't the first time our industry has navigated shifts like this. When ATMs were first introduced, they didn't eliminate bank tellers. Instead, they changed how banking worked. Routine transactions became cheaper, faster, and more accessible, while human effort shifted toward higher-value activities. The real impact wasn't automation alone-it was how institutions reorganized around technology.

AI is similar. The biggest gains won't come from simply adding AI to existing processes. They'll come from rethinking workflows, roles, and systems to take advantage of what this technology makes possible.

What we don't know-and can't know-is exactly when these technologies will hit their full inflection point. We won't get a clear signal when AI moves from rapid progress to truly systemic impact. But waiting for perfect clarity isn't a strategy. If we want to be ready when that moment comes, the work has to start now.

AI is one clear example of how the Federal Reserve can execute boldly, at scale, when we take a System-first approach.

---

[1] The views expressed here are my own and not necessarily those of my colleagues on the Federal Reserve Board.

[2] In this speech, as the Oversight Governor for the Reserve Banks, when I say 'System' I am referring to the 12 Reserve banks, not the Board of Governors.