databricks

EXECUTIVE BRIEF

How Telecommunications Leaders Are Putting AI to Work

Based on Hundreds of Databricks Deployments

Introduction

Telecommunications providers operate some of the most complex data environments in any industry, spanning network telemetry, subscriber activity, billing systems and call center interactions across fragmented OSS/BSS stacks. At the same time, rising infrastructure costs for 5G and fiber are colliding with stagnant revenue and increasing customer churn. Decisions that should take minutes take days. As a result, teams struggle to act in real time: network issues are addressed after they impact service, customer experience signals are buried in transcripts and tickets and critical decisions still depend on slow, manual analysis. Meanwhile, as fraud and cybersecurity threats become more sophisticated, yesterday's risk modeling is no longer sufficient to prevent the attacks of today.

The companies pulling ahead aren't simply deploying AI; they're deploying AI on a unified, governed data platform. That's the difference between a pilot that impresses and a system that scales.

Key takeaways

Multi-agent architectures, not single-purpose chatbots, are the production standard for leading companies.

At AT&T, fraud prevention is powered by a coordinated system of 100+ ML models. This reduces attacks by 80%, resulting in millions of dollars saved.

Data governance is the architectural prerequisite for AI at scale.

Unity Catalog's fine-grained access controls, column-level security and full lineage through agent workflows are what allow telecom providers to deploy AI across network, subscriber and operational data. All while maintaining strict control over sensitive customer information, meeting requirements like GDPR and CCPA, and ensuring every decision is fully auditable.

The most effective AI works across every data type — structured and unstructured — in a single coordinated system.

SQL and Genie agents query network telemetry, subscriber records and billing data, while RAG and Vector Search agents analyze call center transcripts, trouble tickets, field reports and equipment documentation. With Lakebase powering real-time read and write access to operational data, teams can act on those insights instantly. A network operator, field technician or customer service rep can ask a question and get an answer that pulls from both live systems and unstructured service records, without needing to know where the data lives.

Time to value is measured in weeks, not quarters.

Across Databricks deployments, companies are reaching production faster than expected: Digital Nasional Berhad (DNB) achieved their 5G coverage rollout goal a full year ahead of schedule.

The efficiency gains are real and attributable.

Frontier Communications reduced pricing calculation times for fiber installations from weeks to seconds, improving service quality and boosting sales. At NTT DOCOMO, automating LLM usage analysis cut monthly effort from 66 hours to just six, a 90% improvement that freed teams to focus on higher-value work.

Align your stakeholders before you build.

The most successful deployments involve CDOs, CTOs, Chief Customer Officers, heads of service delivery, sales leaders, IT security and finance from day one. This cross-functional collaboration is what separates pilots that scale from pilots that stall.

Millions	Cost Savings
90%	Productivity Increase
10x	ROI Yield

Use cases

Multi-Agent Supervisor Architectures

Most companies rely on dozens of systems spread across databases and document types. Multi-agent architectures solve this by assigning specialized agents to handle different tasks, all managed by a central supervisor that routes each question to the right place.

At AT&T, fraud prevention isn't handled by a single model or interface. It's coordinated across more than 100 ML models that ingest different data types and generate real-time alerts and recommendations, triggering actions across call centers, retail systems and digital channels. That shift from isolated models to a coordinated system is what enables proactive fraud prevention, reducing attacks by up to 80% while operating at the scale of millions of transactions per second.

DNB enables coordinated workflows (data ingestion, processing, analytics, and AI-driven assistants) where multiple specialized components operate together in real time. This approach allows teams to move beyond simple query-response systems toward intelligent, collaborative agents that deliver actionable insights, automate decisions and scale efficiently with the complexity of 5G operations.

Built on: *Unity Catalog (access controls and data lineage), Genie agents (structured data/SQL), RAG agents (unstructured documents), function-calling agents (vector search), Databricks Model Serving*

Natural Language Data Access and Analytics

What if anyone on your team could get answers from your company's data just by asking a question: no SQL, no analyst queue, no waiting? At NTT DOCOMO, that vision became reality with Genie. By introducing a chat-based interface for analytics, teams can explore LLM usage data, generate visualizations and uncover insights using natural language alone. This removed the need for technical expertise and dramatically sped up how quickly stakeholders could access answers.

At Frontier Communications, natural language querying enables teams across the business to explore unified, real-time data and generate insights instantly, accelerating decisions across sales, pricing and operations without relying on technical specialists. This supports initiatives like analyzing over 40,000 contracts with GenAI and helping teams quickly surface revenue opportunities and operational improvements. Frontier's B2B sales team leverages AI-powered lead prioritization, driving a 20% increase in new logo sales and reducing ramp-up time for account executives.

Built on: *Databricks Model Serving, open-source LLM support (Llama, Mistral), Genie text-to-SQL, Mosaic AI critique agent framework*

Autonomous Network Operations and Predictive Fault Detection

Unplanned outages and network degradation are among the most expensive and customer-impacting challenges in telecommunications. By using real-time network telemetry, equipment logs and traffic patterns, AI can detect anomalies and predict failures before they affect service. Operators can resolve issues proactively instead of reacting after disruptions occur.

T-Mobile's predictive models built on unified data from 60+ sources guide cell tower placement, forecast demand and optimize supply chains. This helps drive a >10x increase in weekly site construction starts while reducing data latency by 75%, enabling faster, more reliable service delivery at scale.

Built on: *Databricks streaming pipelines, Delta Lake, real-time ML model serving, Mosaic AI*

Data Governance as Enabler

Governance is what makes enterprise-wide AI deployment possible by giving providers the confidence to scale across highly sensitive subscriber and network data. Unity Catalog's access controls, data masking and lineage tracking are built into the platform, allowing AI agents to operate across datasets like call detail records and billing information, while meeting strict regulatory requirements such as GDPR, CCPA and CPNI, and ensuring every action is fully auditable.

Unity Catalog provides Lumen Technologies centralized, secure data governance across cloud and on-prem systems, enabling GenAI applications to scale confidently on fragmented, sensitive datasets. This unified control layer ensures consistent access, security and data management as AI is embedded directly into technician workflows.

Built on: *Unity Catalog (column-level security, dynamic data masking, data lineage tracking through agent workflows, audit logging), fine-grained access controls for multi-agent systems*

Strategic Recommendations

Build Your Foundation

Before your first agent deployment, implement Unity Catalog. Its built-in access controls, data lineage and audit logging are what allow you to expand AI beyond your first team without creating governance debt. Start by auditing high-volume telecom workflows like network incident triage and fraud detection, and identify where teams rely on manual queries or delayed reporting. Then, run a focused two-week pilot on a real use case (e.g., surfacing churn risk signals, answering network performance questions with natural language or flagging anomalous SIM activity) to prove value quickly and establish a repeatable foundation for scale.

Scale What Works

Once a workflow is proven in one business unit, treat it as a template for the next. Databricks enables reusable pipelines, shared notebooks and governed data assets that let you take a text-to-SQL workflow or multi-agent architecture and expand it across an organization without rearchitecting. Use MLflow to track model performance and Unity Catalog to manage access as new teams and data domains come online. The goal is to build a library of reusable agent components instead of a collection of isolated deployments.

Skip What Doesn't

Not every investment is worth making. Three to avoid:

- **DIY MLOps infrastructure.** Databricks and major cloud providers have already solved this; building your own means spending engineering time on maintaining infrastructure instead of solving business problems.
- **Single-purpose AI point tools.** Every additional vendor increases the total cost of ownership and adds integration debt.
- **Premature model optimization.** LLMs improve every quarter. Build on an open architecture that lets you swap in better models as they become available without rebuilding your workflows.

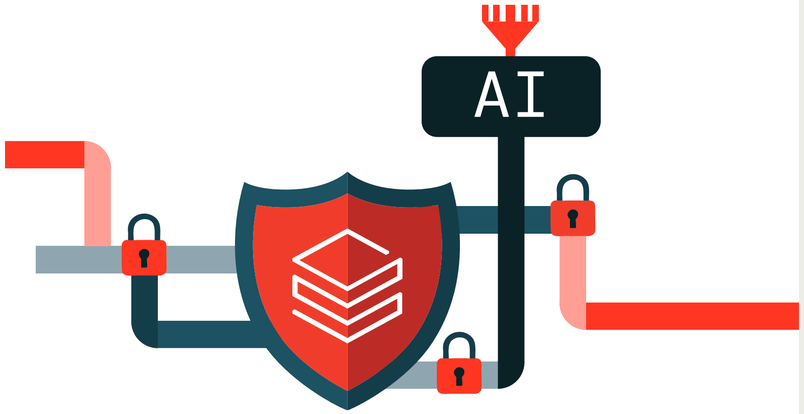

Results from the Field

Three patterns emerge across these deployments:

Efficiency gains are concentrated in data operations: processing, querying and enrichment. This repetitive, high-volume work is exactly what AI handles best.

Speed improvements are transformational, not incremental. Comcast reducing deployment times from weeks to minutes and NTT DOCOMO accelerating monthly analysis efforts by 90% fundamentally changes how fast teams can act on information.

Scale compounds as each AI workflow absorbs more throughput without adding operational cost. AT&T isn't analyzing 10 million transactions per second and supporting over 100 ML models with proportionally more people; the platform handles volume that would be impossible to manage manually.

In every case, the underlying enabler is the same: Databricks Platform that connects raw operational data to the AI layer without requiring teams to stitch together point solutions to bridge the gap.

T-Mobile	75%	**Data Latency Reduction** 6 hour refreshes → 1.5
AT&T	80%	**Decrease** In fraud attacks
COMCAST	10x	**Savings** Reduction in overall compute costs
FRONTIER	20%	**Growth** New logo sales

Realistic timeline

When it comes to AI, business decision-makers often wonder: How long before this technology delivers something real?

Based on Databricks deployments across media and entertainment companies:

9 MONTHS

Data Center Migration

AT&T moved their workloads (100 million petabytes across the network) from a costly on-premises Hadoop data lake to Azure Databricks — achieving a 300% ROI

1 YEAR

Accelerated Rollout Timeline

DNB had the ambitious goal of delivering 80% 5G coverage across Malaysia by 2024; they achieved this milestone in 2023, a year ahead of schedule, making it one of the fastest 5G rollouts globally

MINUTES

Faster Data Processing Unlocks New Revenue Streams

With Databricks Serverless Compute, CKDelta boosted data processing speeds by 93% - from 4-5 hours to 20 minutes - opening up new data monetization opportunities across industries in need of real-time insights, like tourism, transportation, retail, and advertising

Real, measurable value can be delivered in weeks, not the multi-year transformation cycles most enterprises are used to. The companies that move fastest have two things in place before they start: a governed data foundation, so agents have clean, accessible data to work with, and a clear business owner who defines what 'working' looks like. Both are achievable before your first sprint begins.

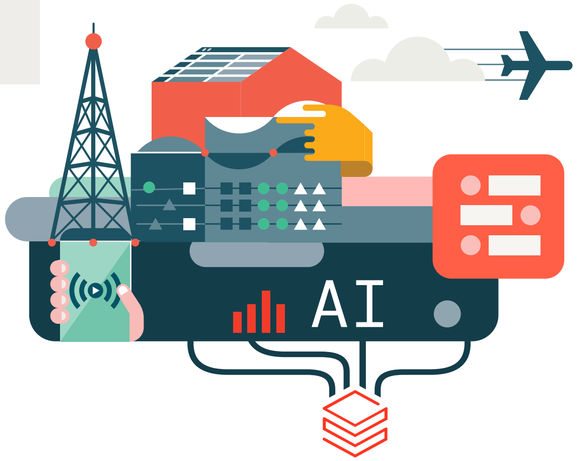

Conclusion

Telecommunications AI is no longer an experiment; it's an operational capability that leading companies are deploying at scale right now. The companies pulling ahead aren't doing it with a collection of point solutions or a single chatbot. They're doing it with a unified platform that brings data engineering, governed access, multi-agent orchestration and model serving together in one environment.

Databricks is the defensible choice for telecommunications AI at scale, and it comes down to three advantages:

Open architecture:
Native support for ChatGPT, Claude, Llama and other open-source models means no model lock-in, and you never have to rebuild when a better model emerges.

Governance at the platform layer:
Unity Catalog makes it possible to deploy AI broadly across an organization without compromising data protection or compliance.

The compounding advantage of a unified platform:
Every new use case built on Databricks shares the same data assets, pipelines and governance framework, so the tenth deployment takes a fraction of the effort the first one did.

The question for telecom leaders isn't whether to deploy AI. It's whether to spend the next two years stitching together tools that don't scale, or build on a unified foundation where every deployment makes the next one faster, smarter and more efficient.

databricks