

CLOUDERA

EBOOK

Unleashing the Power of Generative AI in Your Business



Table of Contents

Introduction	2	Establishing Trust in Your Generative AI	12
Understanding the Basics of Generative AI	3	Security	13
Generative AI	4	Governance	13
GenAI vs AI	4	Accelerating Generative AI Innovation with Cloudera	15
GenAI In Action	4	Accelerator for ML Projects (AMPs) in Action	16
The Challenges of Implementing Enterprise Generative AI	5	AMPs built with Cloudera engineering	17
Data	6	AMPs from HuggingFace Spaces	17
Infrastructure Flexibility	7	AMPs built by community contributors	17
Security, Governance, and Transparency	7	Customer Success Story: OCBC Bank	18
Future-Proofing Your Generative AI Investment	8	Business Impact	20
Differentiating with Your Own Data	10	Conclusion	21
Hallucinations	11	Glossary	23
Factual but Out-of-Context Answers	11	About Cloudera	24

Introduction: The Future is Hybrid


If you're reading this, chances are you've heard about the incredible potential of Generative AI (GenAI) to transform everyday life and reshape businesses across industries. From creating engaging content to optimizing complex processes, this technology is rapidly opening up a world of seemingly endless possibilities.

As the technology continues to evolve, success in the broader business landscape will be directly linked to how well organizations can leverage AI and, in this case, GenAI to redefine both internal processes and customer experiences—whether that's with a customer-facing chatbot, or adoption of automation that streamlines IT infrastructure management.


But that promise isn't without risk and complication. GenAI, while incredibly powerful and full of potential, is not as simple as a drag-and-drop solution to every problem. To get the most out of GenAI, you need to understand it, and more specifically, you need to understand how it should be implemented in your organization, what data should be leveraged with it, and what intended outcomes you're looking to accomplish. It's a lot to adapt to, but it's an endeavor that is becoming widely seen as a must for nearly every enterprise.

So, what exactly is Generative AI, and how can your business harness its power?

The following guide will serve as your trusty companion on a journey to understanding and implementing GenAI, at scale, within your enterprise. We'll break down the basics, tackle common challenges, and show you how Cloudera's Open Data Lakehouse provides a comprehensive solution to unlock the full potential of AI. Get ready to dive into the world of Enterprise GenAI.



According to Forrester's 2023 Artificial Intelligence Pulse Survey, 89% of AI decision-makers reported that they were expanding, experimenting with, or exploring the use of Generative AI within their own organizations.



Understanding the Basics of Generative AI

First things first, let's get to know GenAI. By now, you're likely familiar with the explosion in popularity of GenAI apps like ChatGPT, Claude, or Gemini; the various image and video-generating tools like DALL-E and Sora respectively; or the rising use of chatbots among various major brands to help streamline and enhance customer experiences. But with that said, what exactly is GenAI? What's behind the content it can produce? Let's take a closer look at some of the key elements of this burgeoning technology.

Generative AI

Whether it's unstructured data like text, images, audio, or video, GenAI, a subset of artificial intelligence, is focused on generating new content based on learned patterns and rules.

All those outputs and creations don't just magically materialize, though. The secret behind what enables GenAI to put out those impressive results? It's all about data. Before a GenAI model can start powering that new chatbot or generate an image that matches the prompt you feed it, it needs to be trained on a

set of data to help inform its output. And, in short, the better the data, the better the result. We'll cover more on the data needed to train GenAI models in an upcoming chapter.

GenAI Vs. AI

With those definitions in mind, it's worth examining what makes GenAI different from AI itself. Ultimately, it comes down to a couple of key points: regular AI is like a smart detective that can identify and classify things. A practical application of this might be using AI to seek out and identify areas where processes can be optimized or streamlined to generate cost savings. GenAI, on the other hand, is more like an artist, creating original content based on a given prompt from the user, built upon the data that it is provided.

GenAI In Action

Now that we know what GenAI is, what are the practical applications in a business setting? And how can a tool like this generate value for your business? There are a lot of ways—and it can seem like the possibilities are essentially endless. Imagine generating personalized marketing campaigns, designing unique products, or even creating virtual experiences for your customers. It's like having a creative genie on your team, working 24/7 to bring your ideas to life.

Within organizations, the impact of GenAI is expected to grow significantly as it relates to application development. A Gartner report indicated that over 80% of enterprises will adopt GenAI APIs and models or deploy GenAI-enabled applications in their production environments by 2026—a figure that is up from approximately 5% in early 2023.

The Challenges of Implementing Enterprise Generative AI




The Challenges of Implementing Enterprise Generative AI

Now that you're excited about the potential of GenAI, let's talk about the elephant in the room: implementing it in your enterprise. It's not always a smooth ride, but don't worry, Cloudera has you covered.

Data: GenAI needs to be fed data to work its magic. But that doesn't mean you can simply throw a mountain of data at it and expect an impactful result. As a tool that relies so heavily on data, the quality of the content produced is directly linked to the quality of the data that it was trained on. Imagine trying to bake a cake with rotten eggs and expired milk. Not a good idea, right? The same goes for GenAI.

You must ensure that your data is properly integrated, cleaned, and ready for consumption. Getting to that point of trustworthiness starts by looking at how your organization handles enterprise-wide data management and storage. Organizations need to ensure their data architecture is capable of supporting data in all its forms—be it structured, semi-structured, or unstructured.

The ideal scenario is to adopt a hybrid data management approach, leveraging both on-premises and public cloud infrastructure and data strategy. Cloudera's survey found that while only one-third of its respondents currently deploy multi-cloud or hybrid data architectures, 93% stated they agreed that multi-cloud and hybrid capabilities for data and analytics were a key piece of successfully adapting to change. With a hybrid approach, organizations get the advantages of both cloud and on-premises infrastructure. With the cloud, leading-edge software can be quickly deployed without the expense and long lead times of local installations. On-premises infrastructure allows enterprises to more easily keep data secure, and follow regulations and industry standards.



A Cloudera survey of IT decision-makers found that a quarter (26%) of respondents identified contextual (Trusted) data was a challenge to implementing AI within their organizations.

Structured data is data that can be stored in relational databases, and unstructured data is unorganized data that does not have a predefined format, making it difficult to store or manage.

Infrastructure Flexibility: The AI landscape is like a rapidly changing fashion trend. What's new and innovative today might be outdated tomorrow. You don't want to be stuck with a rigid infrastructure that can't adapt to new models and frameworks. Flexibility is key. However, building that flexibility into existing data architecture can seem like a daunting task.

As noted earlier, the pace of change with emerging technologies is rapid, and meeting those changes head-on means having an infrastructure in place that can implement new tools and technologies at an enterprise scale while ensuring the most optimal business outcomes are achieved.

Security, Governance, and

Transparency: You wouldn't leave your house unlocked with all your valuables inside, would you? The same principle applies to your AI initiatives. You need to ensure that they are secure, compliant with regulations, and transparent to build trust both internally and externally. As all the new and innovative GenAI tools come into use, they all depend on vast swaths of data to be effective. That data still needs to be monitored and used in compliance with both internal guidelines as well as external regulations. Those regulations, whether it's something like the General Data Protection Regulation (GDPR) or a policy like the Health Insurance Portability and Accountability Act (HIPAA), are constantly evolving and failing to keep security and governance standards up to date could spell disaster for businesses.

As all the new and innovative GenAI tools come into use, they all depend on vast swaths of data

to be effective. That data still needs to be monitored and used in compliance with both internal guidelines as well as external regulations. Those regulations, whether it's something like the General Data Protection Regulation (GDPR) or a policy like the Health Insurance Portability and Accountability Act (HIPAA), are constantly evolving and failing to keep security and governance standards up to date could spell disaster for businesses.

But don't worry, you don't have to tackle these challenges alone. The key is to adopt a comprehensive approach that addresses data management, infrastructure flexibility, and governance holistically – and choosing the right vendor to help you do so.

The global average cost of a data breach in 2023 was USD 4.45 million, a 15% increase over 3 years.

Adopting a modern data architecture is crucial to building flexibility.

Among Cloudera survey respondents, 38% of IT decision-makers cited flexibility as a key benefit of this kind of approach.

A survey of enterprise leaders, commissioned by Cloudera and 451 Research, found that 72% of respondents either completely or mostly agreed that data governance was an enabler of business value.

Future-Proofing Your Generative AI Investment



Future-Proofing Your Generative AI Investment

Investing in AI is about making sure it's useful for your enterprise for years to come. That's where futureproofing comes in. You need a platform that can adapt to the ever-changing AI landscape, so you don't end up with an outdated tool. So, what should you look for when it comes to implementing a future-proof platform?

One of the biggest components to future-proofing your GenAI investment is adopting a platform with multi-cloud capabilities. Managing and analyzing data across any form factors, public and private, and at the edge is crucial to supporting AI implementations.

Building a data architecture with greater flexibility is so important to future-proofing efforts. We've talked a lot about just how quickly the AI space can change, and organizations that find themselves with highly rigid data management practices could be leaders one day and left in the dust the next day.

As AI models evolve, the volume of data that needs to be trained grows as well. With the amount of data constantly growing, it's critical

to make sure that data silos don't proliferate across the enterprise. Cloudera's platform ensures AI models have access to the right, high-quality, data and delivers a flexible approach designed to help you lower TCO costs even as data volumes grow.

Cloudera's future-proof platform seamlessly integrates with the latest and greatest AI models and frameworks, so you never have to worry about compatibility issues or being left behind. Leveraging the Cloudera platform, you can deploy storage, compute, and access, all with the freedom offered by the cloud, avoiding vendor lock-in and taking advantage of best-of-breed solutions.

But it's not just about staying current—it's also about getting to the finish line faster. Data practitioners need every tool at their disposal to help keep up with the rapid changes that come with GenAI. With Cloudera's Accelerators for ML Projects (AMPs), you can hit the ground running. These pre-built reference examples are like a cheat sheet for your GenAI projects. They're tailored to specific use cases, so you can quickly adapt and customize them to your unique needs. But more on that later in this eBook.



The Cloudera platform gives you the flexibility to choose, and easily switch between, open-source and closed-source models, making it easier to adapt to evolving needs.

Differentiating with Your Own Data



Differentiating with Your Own Data

Now, here's where things get exciting. One of the biggest advantages of GenAI is the ability to use your own unique data to create solutions that are tailored specifically to your business—making it all your own, and most useful for your needs.

By integrating your enterprise-specific data into GenAI models, you can develop large language models (LLMs)—a form of GenAI that is trained on vast amounts of text data—that speak the language of your business. Unlike broader GenAI, LLMs are specifically trained on text-based data—think tools like OpenAI's ChatGPT—and are, in turn, effective at creating text-based outputs.

Imagine having an AI assistant that understands your industry jargon, and your company's unique challenges, and can provide insights that are relevant to your context.

All that said, there are some limitations to overcome when leveraging an LLM in your business. One of the key limitations is that LLMs lack contextual understanding of the enterprise-specific questions one might have for the system. Some of the more well-known LLMs available are trained on huge volumes of publicly available text from a variety of sources on the Internet. But without access to enterprise knowledge bases or proprietary data sources,

the results that they generate are far less likely to create a result that is accurate to what the business user is looking to achieve. In these cases, there are two common responses that an LLM might exhibit:

Hallucinations


As you may be able to guess from the name, this is the tendency of LLMs to generate fictional information, in a manner that can seem highly realistic. These hallucinations can be much harder to spot than you might think as they tend to be more of a blend of factual information with erroneous, made-up information.

Factual but Out-of-Context Answers

Another variation that could come from an LLM where the model is unsure about the specific answer it needs to generate. In this case, the LLM will come up with an answer that, while technically true, is very generic and not tailored to the context with which the original query was posed. So, what can you do to ensure your LLMs have the contextual data they need to be effective? The Cloudera platform makes it easy to create these specialized LLMs by allowing you to train your models on your proprietary data. The result? Unparalleled accuracy, relevance, and a significant competitive edge. Together with Cloudera, you can implement LLMs with confidence and unlock the true potential of your data to stand out from the crowd.

Tip: With Cloudera, you can easily implement a hybrid data architecture that enables you to quickly prototype AI applications in a public cloud environment and then deploy those applications in a private cloud to lower the total cost of ownership (TCO).

Establishing Trust in Your Generative AI



Establishing Trust in Your Generative AI

We've talked a lot about data to this point—ensuring that GenAI models have high-quality and trustworthy data, to be exact. But what does that really mean? What makes data trustworthy, or untrustworthy? And what makes trust so important to good GenAI?

Trust is imperative to keep your GenAI initiatives together. Without it, your efforts will fall apart faster than a house of cards. Any time you're pulling huge swaths of data into new projects, there's always the risk that doing so may put you at odds with either internal guidelines or broader regulations. Questions around data access and visibility could result in a slipup along the way that could derail not just a GenAI project, but the entire organization at large. This is why Cloudera is committed to providing the right security and governance capabilities to support your AI projects.

Now let's dive a bit deeper into each of those elements—security and governance.

Security

When it comes to protecting your data, you need robust security measures and models to prevent unauthorized access and manipulation. Cloudera's robust data platform has you covered with features like encryption, access controls, and monitoring—making your AI initiatives secure.

Governance

Security alone isn't enough. A successful AI project also requires highly effective governance, which is, essentially, a set of rules that everyone in an organization agrees to follow. With so much data, much of it involving sensitive information, floating around, visibility into where data lives and who can access it is critical to good governance.

Governance ensures transparency and control over your AI-driven decisions. With Cloudera's comprehensive governance framework, you can establish clear policies, processes, and audit trails. No more making decisions in the dark or crossing ethical boundaries. You'll be able to ensure the responsible and ethical use of GenAI.

In the end, everything comes down to trust. There is already so much at stake with GenAI when it comes to generating the right outcomes from the technology, that a strong foundation of trust is critical. By prioritizing security and governance, you'll build trust with the most important stakeholders, both internal and external, and instill confidence in your AI initiatives, knowing that you have their best interests in mind. Plus, when it comes to regulatory compliance, you'll be able to sleep soundly at night knowing you're not only adhering to the current guidelines but are adequately prepared to adjust in the face of new policies as they arise.

Cloudera's comprehensive approach to data architecture ensures that you can trust your business' data and AI models while maintaining the highest standards of security and transparency.

Accelerating Generative AI Innovation with Cloudera



Accelerating Generative AI Innovation with Cloudera

You may be familiar with the concept of data lakes, data warehouses, or even data lakehouses. But Cloudera's Open Data Lakehouse delivers functionality that is one-of-a-kind and ideally suited to address the challenges that come with accelerating your GenAI journey?

First and foremost, the Open Data Lakehouse is like a big, friendly giant. It welcomes all kinds of AI models and vector databases with open arms. No more feeling trapped or limited by compatibility issues. That's part of what makes the Open Data Lakehouse 'open'. The interoperability and compatibility of the Open Data Lakehouse allow you to choose the tools that best fit your needs. That simplicity makes it much faster for data practitioners to get straight to the insights they need.

You're free to explore and innovate!

An Open Data Lakehouse has a host of benefits that can positively impact GenAI efforts. This architecture taps into the flexibility that you gain from a data lake, integrating it with the structured querying capabilities of a data warehouse—this makes it easier to handle both raw and processed data of any number

of types and formats. This unified data environment eliminates the need for maintaining separate data silos and facilitates seamless access to data for AI and analytics applications.

But what about when your data grows bigger and bigger? It's a very real concern, particularly as GenAI, and AI in general continues to demand more data to train more advanced models. Cloudera's cost-effective hybrid architecture allows you to scale effortlessly. Whether you're dealing with a little bit of data or a massive amount, the Open Data Lakehouse can handle it all. The capabilities within Cloudera's Open Data Lakehouse let you seamlessly scale infrastructure horizontally. That approach ensures the data architecture is prepared to handle increased data ingestion, processing, and storage demands. Ultimately, scalability is what enables the open data lakehouse to remain so adaptable, even when the complexity of data and AI models accelerate over time.

Data lakes provide a protective ring around the data stored in a cloud object store, including authentication, authorization, and governance support.

Data warehouses bring together disparate data into a single repository for use in data analytics.

Open Data Lakehouses serve as a unified data platform—an element critical to effective data management. An overwhelming 90% of respondents in Cloudera's AI survey agreed unifying the data lifecycle on a single platform is critical for analytics and AI.

Accelerators for ML Projects (AMPs) in Action



Accelerators for ML Projects (AMPs) in Action

Now, let's take a closer look at Cloudera's secret weapon: Accelerators for ML Projects (AMPs). The world of ML projects is a hectic one. Only 20% of ML models in the enterprise actually make it to production, according to over three thousand surveyed C-level executives. With IT teams increasingly pressed for time, they need the right tools to help support the development of ML projects and ensure that they successfully make it from production to completion.

AMPs are ML projects that can be deployed with one click directly from Cloudera Machine Learning. These accelerators enable data scientists to quickly go from an idea to a fully working ML use case in half the time. AMPs also provide an end-to-end framework for building, deploying, and monitoring business-ready ML applications instantly. Every AMP captures industry-leading practices for overcoming some of the most complex ML challenges that data practitioners face. Within Cloudera's AMP catalog there are three different types to choose from:

1. AMPs built with Cloudera engineering – These are built and supported by research teams that focus on the latest in AI and ML. They are also put through a rigorous testing and review process to guarantee the highest quality reference projects.

2. AMPs from HuggingFace Spaces – Like AMPs, Spaces are ML demo applications that are self-contained and instantly ready to deliver value upon deployment.

3. AMPs built by community contributors – The community AMP catalog is where anyone can contribute best-in-class solutions to an open-source repository of meaningful projects.

These pre-built reference examples are like a toolkit for your GenAI projects. They're packed with best practices, architectures, and code snippets that you can use to kickstart your initiatives and achieve success faster than you can say "AI"!

But what exactly can AMPs do for you? Let's explore a couple of examples. First, imagine you want to create a chatbot that can understand and respond to your customers' inquiries. With the LLM ChatBot AMP, you can build an intelligent conversational agent that leverages your enterprise-specific data to provide personalized and context-aware responses. It's like having a customer service

representative who never sleeps and knows your business inside out!

Next, let's say you want to adapt a pre-trained foundation model for your specific use case. Foundation models are like Swiss Army knives of the AI world — they're versatile and can be used for a wide range of tasks. With the fine-tuning AMP, you can take these models and tailor them to your unique requirements, getting the most bang for your buck.

By leveraging Cloudera's AMPs, you can accelerate your Generative AI journey and drive tangible results faster. It's like having a team of AI experts at your fingertips, guiding you towards success.

Tip: Every AMP is fully open source, meaning even though they are easiest to deploy in Cloudera Machine Learning, each project provides a README with instructions on how to deploy in any environment.

Customer Success Story: OCBC Bank



Customer Success Story: OCBC Bank

Enough of the theoretical, let's take a look at an example of GenAI in action. Meet OCBC Bank, a leading financial institution that has embraced Cloudera's Open Data Lakehouse to revolutionize its banking experiences with trusted AI. OCBC Bank realized it needed to introduce a more resilient infrastructure - a data platform that could manage projects with increasing volume, variety, and velocity of data, while also enabling real-time analytics.

OCBC Bank's journey with Cloudera began with a vision to harness the power of AI and data analytics at scale. Specifically, the bank was looking to enhance its data platform to leverage the private cloud and provide a secure, controlled, and customizable experience for storing and processing confidential data. OCBC Bank also wanted to use AI and ML technologies to make more data-driven decisions and improve customer experiences while mitigating risks

By adopting the Cloudera platform, OCBC Bank was able to transform and personalize

customer banking experiences. The bank's teams built Next Best Conversation, a centralized platform that uses ML to analyze real-time contextual data from customer conversations related to sales, service, and other variables to deliver unique insights and opportunities to improve operations.

Critically, Cloudera helped OCBC Bank break down the data silos that existed across the enterprise, democratize access to insights, and empower their employees to make data-driven decisions. It was like giving their team superpowers!

But the real magic happened when they started leveraging Generative AI to enhance their customer experiences. By training AI models on their vast repository of customer data, OCBC Bank was able to develop personalized recommendations, automate processes, and create engaging digital experiences. The result? Happy customers, increased efficiency, and a competitive edge in the market.

Leveraging powerful ML models, OCBC Bank could now send over 100 different personalized nudges on its mobile banking app, notifying customers about financial opportunities.

Business Impact

The success of OCBC Bank is a testament to the transformative power of Cloudera's Open Data Lakehouse. By providing a secure, scalable, and flexible platform for Generative AI, Cloudera enabled OCBC Bank to innovate at scale and drive meaningful business outcomes. It's a success story that will inspire you to embark on your own Generative AI journey.



OCBC customers now enjoy faster transactions with the aid of the bank's chatbots, which handle 10% of interactions on the bank's website.

Leveraging the Cloudera platform, OCBC Bank's Next Best Conversation platform helped earn \$100 SGD (\$75 USD) million annually, using data to curate personalized experiences for customers.

Conclusion



Conclusion

With that, we've reached the end of our overview of the ins and outs of GenAI. By now, you should have a solid understanding of Enterprise Generative AI, the challenges involved, and how Cloudera's Open Data Lakehouse can help you unlock its full potential.

We've covered a lot of ground, so let's recap some of the most important takeaways:

- 01** Generative AI is a game-changer for businesses across industries, enabling the creation of novel and contextually relevant content.
- 02** Implementing Generative AI in an enterprise setting comes with challenges related to data management, infrastructure flexibility, and governance.
- 03** Future-proofing your Generative AI investment requires a platform that can adapt to the evolving AI landscape and accelerate your time-to-market.
- 04** Differentiating with your own data is key to creating highly contextualized and accurate Generative AI solutions.

05 Establishing trust through security, governance, and transparency is essential for the success of your AI initiatives.

06 Cloudera's Open Data Lakehouse provides an open, flexible, and scalable platform for accelerating Generative AI innovation.

07 Accelerators for ML Projects (AMPs) enable you to jumpstart your Generative AI journey and achieve rapid success.

08 Real-world success stories, like OCBC Bank, demonstrate the transformative potential of Cloudera's Open Data Lakehouse for Generative AI.

As you embark on your own GenAI journey, remember that the key to success lies in embracing a holistic approach that prioritizes trust, security, and scalability. With Cloudera by your side, you'll have a trusted partner to guide you every step of the way.

The future of Enterprise Generative AI is bright, and the possibilities are endless. By leveraging Cloudera's Open Data Lakehouse, you can unlock the transformative potential of AI for your business, drive innovation, and stay ahead of the competition. So, what are you waiting for? Let's get started on this exciting adventure together!

Glossary

Glossary

- **Generative AI (GenAI):** A subset of artificial intelligence that focuses on creating new content, such as text, images, audio, and video, based on learned patterns and rules.
- **Machine Learning (ML):** A subset of artificial intelligence that leverages data and algorithms to mimic human behavior when dealing with complex tasks.
- **Large Language Models (LLMs):** AI models that are trained on vast amounts of text data to understand and generate human-like language.
- **Open Data Lakehouse:** A platform that combines the best of data lakes and data warehouses, enabling organizations to store, process, and analyze structured and unstructured data at scale.
- **Accelerators for ML Projects (AMPs):** Pre-built reference examples that encapsulate best practices, architectures, and code snippets for common Generative AI use cases.
- **Foundation Models:** Large-scale AI models that are pre-trained on diverse datasets and can be fine-tuned for specific tasks or domains.

Resources

- **Cloudera Enterprise AI:** <https://www.cloudera.com/why-cloudera/enterprise-ai.html>
- **Cloudera Machine Learning:** <https://www.cloudera.com/products/machine-learning.html>
- **Cloudera Blogs:** <https://blog.cloudera.com/>
- **Get Your AI to Production Faster: Accelerators for ML Projects**
- **Navigating the Enterprise Generative AI Journey: Cloudera's Three Pillars for Success**
- **Data Architecture and Strategy in the AI Era**
- **Generative AI for the Enterprise**
- **CIO Whitepaper: Data architecture and strategy in the AI era**
- **Cloudera and 451 Research Report {Not on Cloudera Site Currently}**

Try Cloudera on Public Cloud

5-Day Free Trial

- Unlock actionable insights in minutes.
- Experience Cloudera through common use cases that also introduce you to the platform's fundamentals and key capabilities.



trycloudera.com



Generative AI

Powered by Cloudera Machine Learning.

Explore powerful generative AI applications within Cloudera Machine Learning using Accelerators for Machine Learning Projects.

About Cloudera

Cloudera is the only true hybrid platform for data, analytics, and AI. With 100x more data under management than other cloud-only vendors, Cloudera empowers global enterprises to transform data of all types, on any public or private cloud, into valuable, trusted insights. Our open data lakehouse delivers scalable and secure data management with portable cloud-native analytics, enabling customers to bring GenAI models to their data while maintaining privacy and ensuring responsible, reliable AI deployments. The world's largest brands in financial services, insurance, media, manufacturing, and government rely on Cloudera to be able to use their data to solve the impossible—today and in the future.

To learn more, visit cloudera.com and follow us on [LinkedIn](#) and [X](#). Cloudera and associated marks are trademarks or registered trademarks of Cloudera, Inc. All other company and product names may be trademarks of their respective owners.

Cloudera, Inc. 5470 Great America Pkwy, Santa Clara, CA 95054 USA

cloudera.com

© 2024 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice.

6413-001 September 4, 2024

CLOUdera

