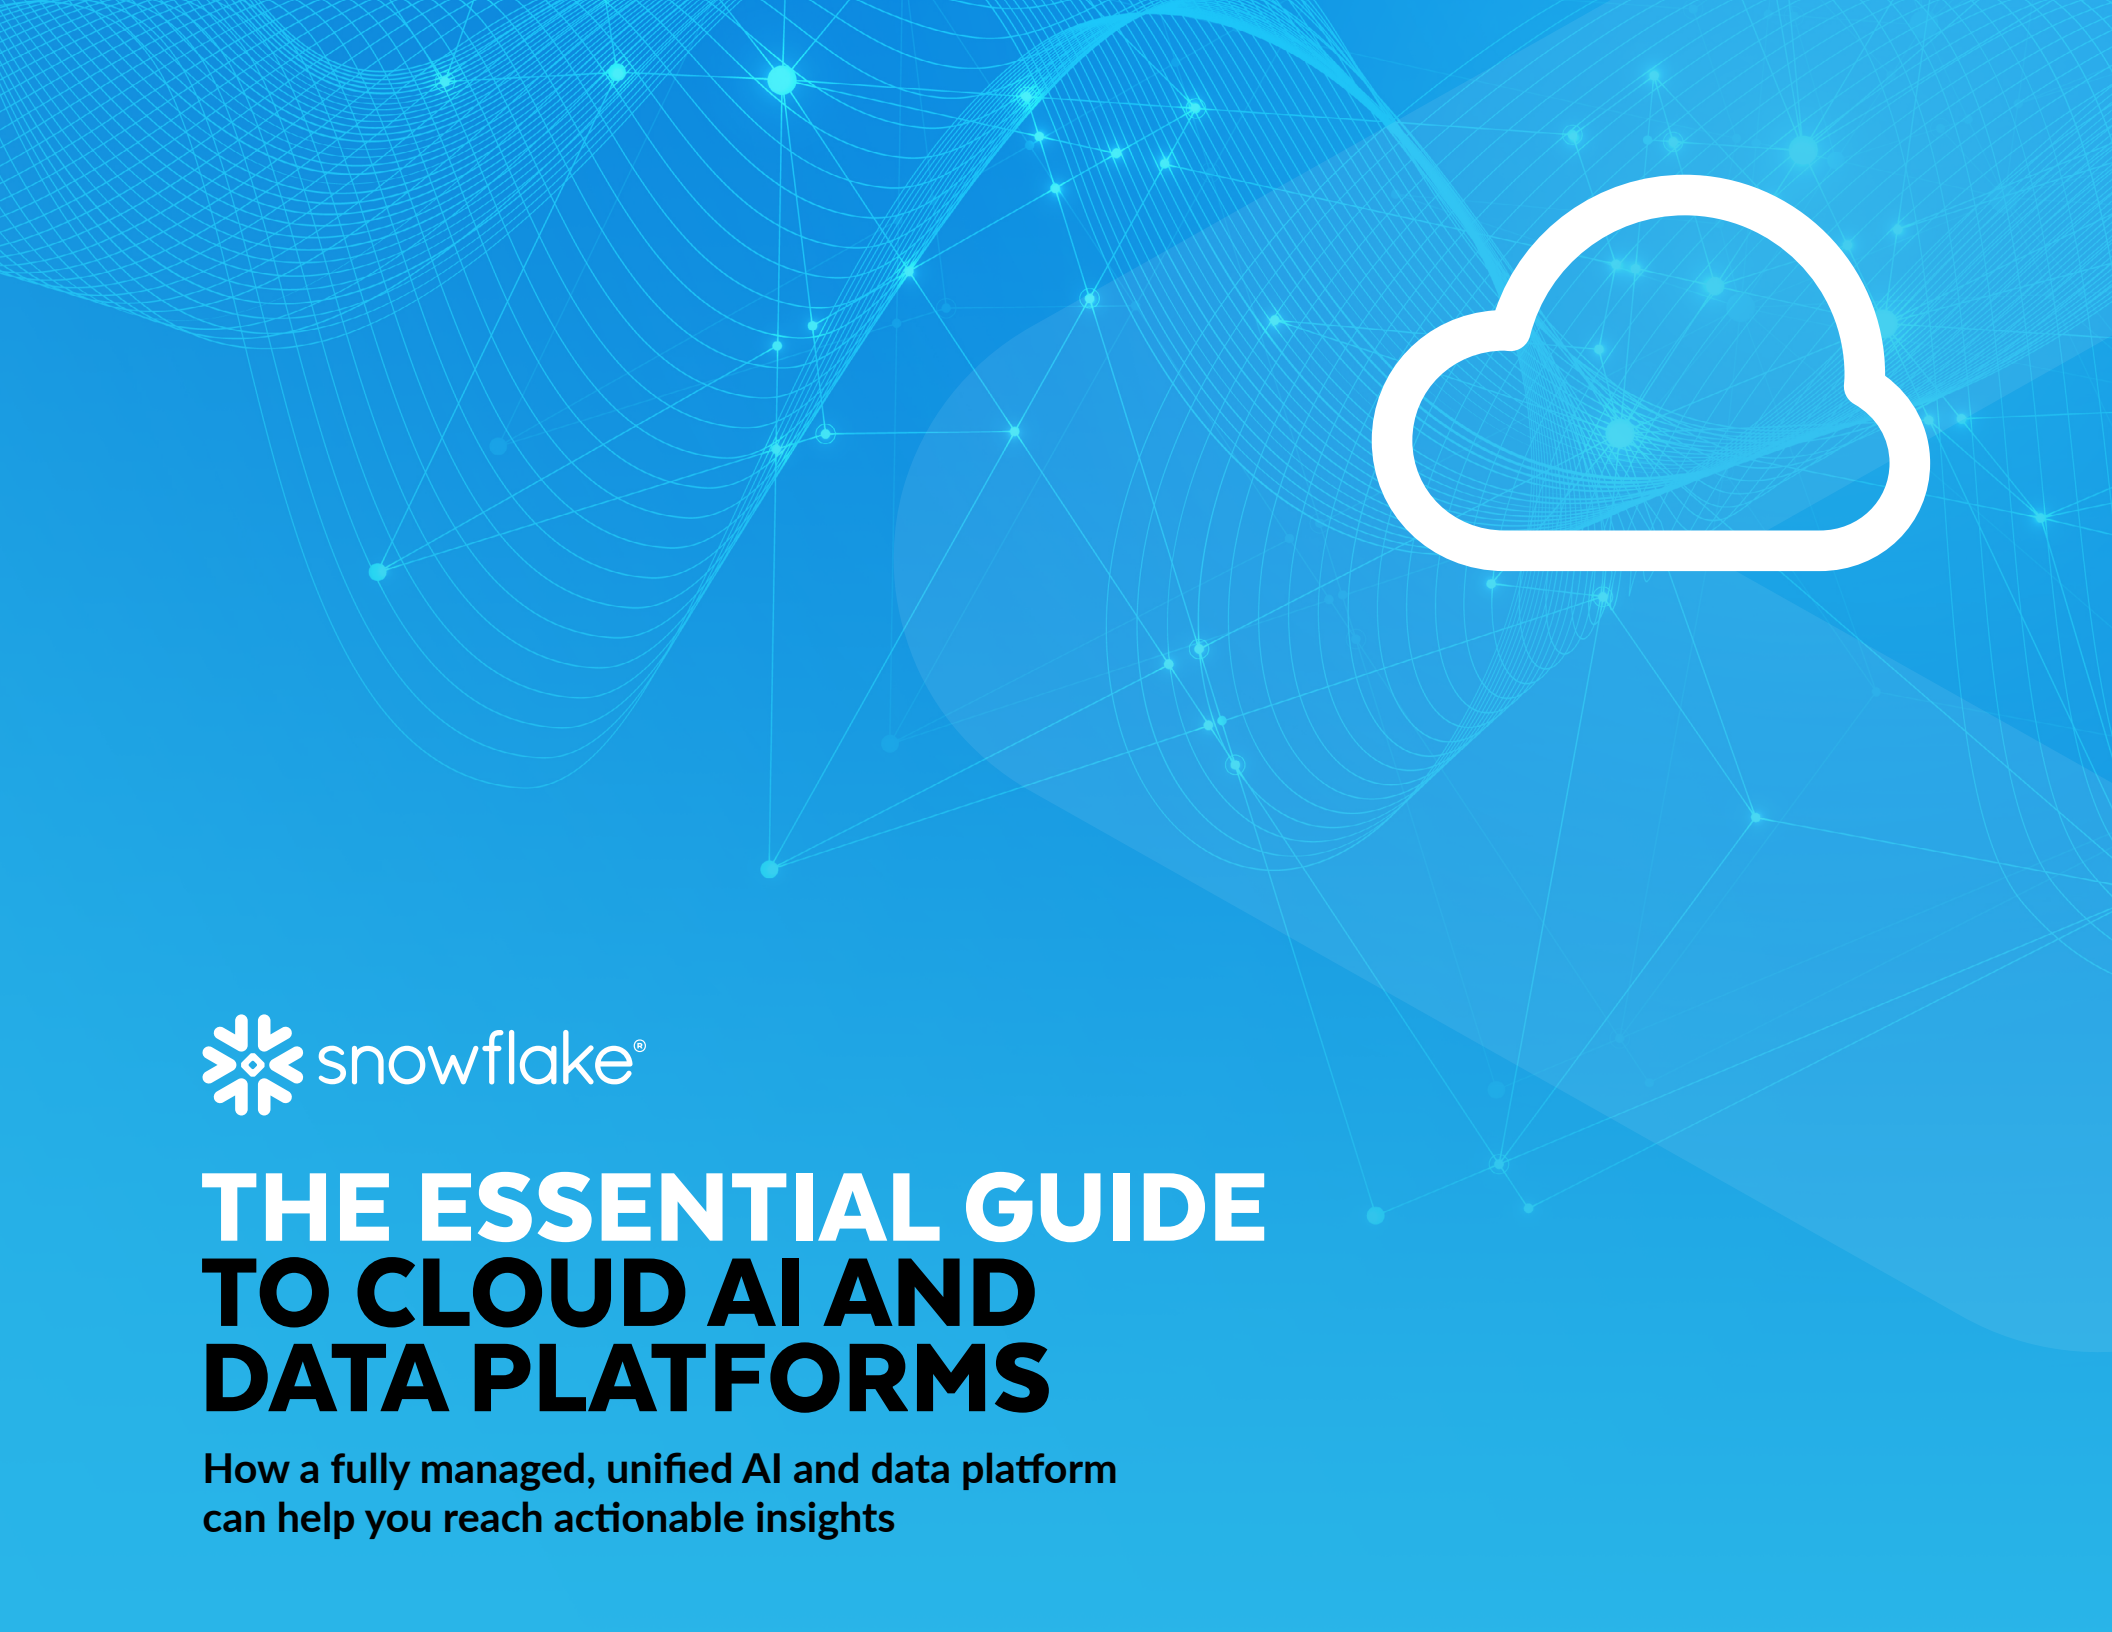

snowflake®

THE ESSENTIAL GUIDE TO CLOUD AI AND DATA PLATFORMS

How a fully managed, unified AI and data platform can help you reach actionable insights

TABLE OF CONTENTS

INTRODUCTION

Today's businesses are awash in data from enterprise applications, mobile phones, websites, point-of-sale terminals, and billions of interconnected devices that track the world around us. The total amount of data created, captured, copied and consumed globally will grow to more than 180 zettabytes by 2025.

Legacy on-premises and typical data platforms can't keep up with the relentless creation, storage, analysis and sharing of these diverse data sets — don't even get started on how ready they are for organizations to use this data for accelerating AI and building and distributing applications. Much of the data is semi-structured, which means it doesn't fit neatly into a traditional data warehouse architecture. Some data types, such as images and audio files, are also wholly unstructured. These platforms also don't offer the flexibility and interoperability to work with data on-premises or in open table formats, which removes lock-in and allows your organization to adapt to any architectural pattern.

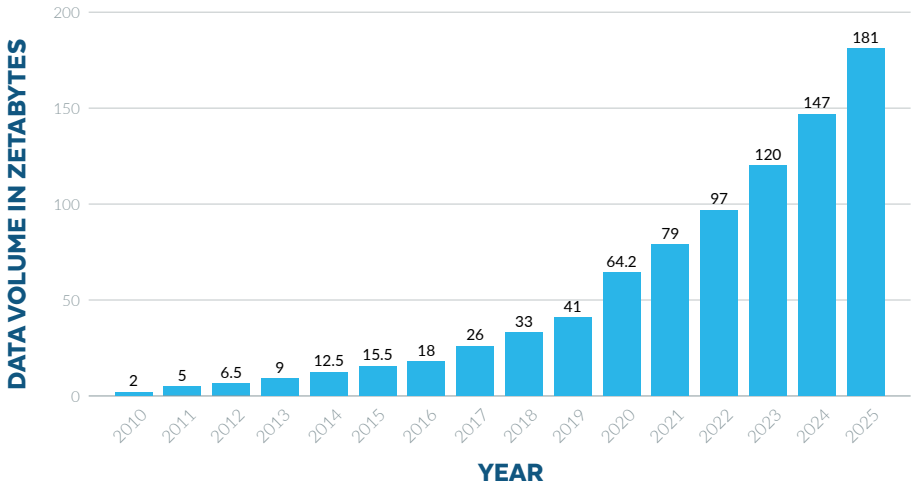

Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025.

In response, many organizations have established unique solutions for each type of data and each type of workload, such as a data warehouse for operational reporting of structured data, data marts for departmental reporting and analytics, and data lakes for ad hoc analytics on semi-structured and unstructured data. They also have implemented specialized extract, transform and load (ETL) tools to rationalize different types of data into common formats and set up data pipelines to orchestrate the exchange of data among databases and computing platforms. Each of these "point solutions" requires independent and specialized tools and services, with each connection point creating a new silo with painful implementation and maintenance issues.

Additionally, organizations are racing to implement generative AI (gen AI) tools and app development into their business. But app development and running AI/ML workloads can create silos by limiting the accessibility and usefulness of data, hindering collaboration and knowledge sharing, restricting resource utilization, and making it difficult to standardize processes across an organization.

A cloud AI and data platform is not a disparate set of tools. It's a unified, pervasive platform that houses all your data, apps and AI and puts it to work, securely and consistently. It enables advanced analytics with modern data warehouses and data lakes, easy-to-use and cost-effective ingestion and transformation for data engineering, the creation and use of models and services for AI and machine learning, the development of data and AI applications, and the secure sharing of data, apps and models for collaboration.

The platform's architecture should also empower organizations to easily scale their computing resources up or down as needed, eliminating infrastructure silos and promoting efficient resource utilization. A single cloud-based AI and data platform for data storage, processing and sharing actually breaks down barriers between different teams and fosters a collaborative environment that accelerates innovation and drives better decision-making.

This essential guide will illuminate the capabilities and benefits of a modern cloud AI and data platform, how to select and optimize the right platform, and how Snowflake can help your organization mobilize your data, apps and AI to reach actionable insights.

WHAT IS A CLOUD AI AND DATA PLATFORM?

A cloud AI and data platform can support various types of frameworks and models for enterprise AI — including integrating popular AI/ML tools and libraries to allow data scientists and engineers to work with their preferred frameworks. Storing and managing large volumes of structured and semi-structured data is crucial for training and deploying AI/ML models. A cloud AI and data platform can also run large language models (LLMs) within the platform for capabilities such as text analysis, sentiment analysis, entity recognition and language translation. The platform should also be able to integrate with other popular language models like GPT-3 through their APIs to provide organizations with a wide range of capabilities. This gives organizations valuable insights from unstructured data and enhances their AI applications.

Forward-looking organizations seek a single, powerful and extensible cloud AI and data platform that can handle legacy needs in conjunction with a raft of new and pressing requirements, such as contending with a broad array of data types and architecture patterns, supporting new types of analytics, enabling robust and automated data pipelines, securely sharing data, developing applications, and scaling effortlessly to handle the data demands of cutting-edge technologies like generative AI.

It's important that a cloud AI and data platform also gives organizations flexibility while using elastic computing resources that enable organizations to scale their compute capacity up or down based on the needs of their diverse workloads to ensure optimal performance and cost efficiency.

Organizations need a platform combining the best components of cloud data warehouses, modern data lakes, AI/ML, and app development platforms with secure ways to share and consume shared data, apps and AI from within a single, cohesive solution. They need an AI and data platform.

STAYING AHEAD OF THE TRENDS

A cloud AI and data platform takes advantage of five important technology trends:

1. **The rise of the cloud:** Traditional data center infrastructure is sized for a known set of data management tasks. The cloud offers unlimited capacity for storing and processing data, apps and AI. This opens the door to unprecedented concurrent high-performance workloads within a centralized platform.
2. **The explosion of data:** Data will continue to grow in both size and variety, driven, in part, by the proliferation of internet, mobile, social, IoT and AI technologies that produce immense quantities of raw but vital data. This growth happens every second of every day. And within the new data lies valuable insights for organizations with the technology, resources and commitment to tap its potential. And data is vital to training models for AI/ML use cases and for building applications.
3. **The democratization of analytics:** On some level, every business user is an analyst. When, where and how these users perform analytics is changing quickly, however, as data and analytics are democratized across the enterprise. As the appetite for data and AI continues to grow, analytics will become central to more and more business processes, from historical quarterly reporting to forward-looking predictive and prescriptive analytics. This will also lead to an increase in analytics use cases. Analytics encompasses several approaches to extract insights from data. Descriptive analytics, for example, focuses on understanding past events and trends by analyzing historical data, while diagnostic analytics goes a step further, identifying the root cause of issues or events by examining data patterns. Predictive analytics forecasts future outcomes based on historical data and statistical models, while prescriptive analytics recommends actions based on predictions to optimize future outcomes.
4. **Build and distribute applications:** Organizations need to scale their business with applications. It's a superior way to derive value from data both inside and outside a company. A cloud AI and data platform helps your organization with streamlined development, deployment and distribution of apps.
5. **Gen AI use cases:** Gen AI and ML capabilities go beyond analytics, focusing on creating intelligent systems that can learn, adapt and act in various situations. Specifically, they cover a range of techniques that help intelligent systems perform tasks and learn from data.

HOW TO SELECT A MODERN CLOUD AI AND DATA PLATFORM?

With a traditional platform, fixed compute and storage resources limit concurrency — the ability to perform many tasks simultaneously and/or to allow many users to access the same data and resources. Users may also have to stitch together multiple services, with multiple bills that makes it hard to understand total operating costs. Furthermore, these platforms often require manual expertise to manage and optimize.

Organizations should select a cloud AI and data platform with a multi-cluster, shared data architecture that separates compute and storage resources so these can be scaled independently to leverage all the resources of the cloud. This allows multiple users to query the same data without degrading performance, even as other workloads are taking place simultaneously, such as ingesting data via a data engineering pipeline or training a machine learning model. The platform should also be a fully managed service that efficiently supports users and workloads at scale while reducing the time, effort and expertise required to operate.

A fully managed cloud AI and data platform with a multi-cluster, shared data architecture should have five components that are logically integrated:

1. **Interoperable storage:** This component refers to the centralized, scalable and secure data storage provided by a cloud AI and data platform. It allows users to store and manage structured, semi-structured and unstructured data in a single location, making it easily accessible to various users and workloads. The interoperability means data can be seamlessly shared and used across the platform.
2. **Elastic compute:** This is the computing power provided by a cloud AI and data platform, which allows users to process and analyze large volumes of data without the need for additional hardware or infrastructure. The elasticity means users can scale their computing resources up or down as needed, ensuring optimal performance and cost efficiency.
3. **AI models and services:** This encompasses the AI and machine learning capabilities provided by a cloud AI and data platform, such as built-in models and tools. These features enable users to develop and deploy AI-driven applications and services.
4. **Cloud services:** This refers to the comprehensive security, governance and maintenance features provided by a cloud AI and data platform. It ensures the platform is secure, compliant and well-maintained without requiring significant effort from users. Key aspects include data encryption, access control, and maintenance and tuning.
5. **Global flexibility and high availability:** A cloud AI and data platform should support an organization's global operations across diverse workloads, providing users the flexibility to deploy and manage their data, apps and models across different regions and cloud providers. This includes multi-region, multi-cloud support for collaboration, governance and business continuity.

These components provide the foundations for a cloud AI and data platform that should be built on five key fundamentals:

1. A single platform that supports many different workloads
2. Secure, governed access to data, apps and AI
3. Instant and near-infinite performance and scalability
4. A zero-maintenance, cost-effective cloud service that's instantly available and extremely easy to use
5. A common technology layer across regions and clouds that lets you operate at a global scale

A SINGLE PLATFORM THAT SUPPORTS MANY DIFFERENT WORKLOADS

When properly architected, a fully managed cloud AI and data platform can offer virtually unlimited scale for every workload, making it possible for thousands of users to analyze and share data concurrently, with no bottlenecks. It further empowers organizations to harness the full potential of their data, enabling businesses to accelerate their AI initiatives and build apps more effectively. This unified platform allows a diverse range of users — from data scientists to business analysts — to interact with the same copy of data seamlessly, making the platform an invaluable asset for organizations looking to unlock the potential of their data.

Whether your data is stored in the cloud, an AI and data platform itself, or in an external repository, such as an object store from one of the cloud infrastructure providers, all users should have a single interface enabling them to view and manage that data. A cloud AI and data platform should allow you to access data in these external tables just as easily as you can from the main platform — and with exceptional performance. It should be flexible enough to let your organization work with data on-premises or in open table formats. This removes lock-in and allows you to adapt to any architectural pattern.

Traditional data platforms often require a lot of effort to integrate with AI models and services because they may not have built-in support or they may require complex data pipelines to be developed. This can lead to increased development time, costs and errors in the data processing. A cloud AI and data platform should offer seamless integration with a wide range of AI models and services. This allows you to accelerate your organization's AI initiatives and build apps more effectively on the same platform that already houses your data. This simplifies the process of working with AI models and services, reducing the time and resources required for integration and improving overall efficiency.

Secure, governed access to data, apps and AI

A cloud AI and data platform should focus on simplifying how you protect and govern your content across data, apps and models to thwart breaches while complying with all industry and regional regulations. Security and governance depend on three essential factors:

1. **Confidentiality:** Preventing unauthorized access to content
2. **Integrity:** Ensuring content is not modified or corrupted and is properly governed
3. **Collaboration:** Providing ways for disparate teams to collaborate across shared sets of curated content

A cloud AI and data platform should also have provisions for securing your content while enforcing comprehensive governance. This ensures that no matter what type of workload, application or procedure is invoked, the content is seen only by authorized users, and that all users obtain consistent results.

Content is an asset. To maximize its use, you must be able to securely make it available — across and outside your organization — for a number of benefits. You must be able to share that content, so it doesn't have to be copied or moved.

INSTANT AND NEAR-INFINITE PERFORMANCE AND SCALABILITY

Traditional data platforms are derived from legacy environments that work on closed networks with local data. As data-driven workloads gain popularity, these legacy platforms slow down under the weight of too many concurrent users and burgeoning volumes of data. These platforms typically rely on a fixed infrastructure, meaning that resources are allocated for peak demand, even if they are not utilized consistently, which can lead to inefficiencies and higher costs. On the other hand, a cloud AI and data platform has elastic compute, which allows you to scale your organization's computing resources up or down as needed, ensuring you only pay for what you use. The platform's architecture also separates compute and storage, allowing your organization to scale both independently, further optimizing resource utilization and cost efficiency. This flexibility is a key differentiator between traditional data platforms and a cloud AI and data platform's elastic compute.

A ZERO-MAINTENANCE, COST-EFFECTIVE CLOUD SERVICE THAT'S INSTANTLY AVAILABLE AND EXTREMELY EASY TO USE

While easy access and shareability for data, apps and AI are top priorities for today's organizations, most want to avoid the tedious database maintenance, system management and IT administration needed for legacy data management.

A cloud AI and data platform should enable you to run the workloads you need without having to provision infrastructure or manage a complex software environment. These should be taken care of as part of a zero-maintenance platform that automates security and manages each workload to guarantee top performance. Zero maintenance means no infrastructure to manage and no knobs to turn, so you can focus on the data and not on managing the platform.

Your platform should also offer per-second billing, which enables each user and workgroup to pay only for the precise storage and compute resources used in per-second increments, so you never have to pay for idle capacity.

A COMMON TECHNOLOGY LAYER ACROSS REGIONS AND CLOUDS THAT LETS YOU OPERATE AT A GLOBAL SCALE

A cloud AI and data platform should have a common metadata layer across clouds and regions to let you operate at a global scale. This enables organizations to collaborate on the same content — regardless of the cloud and region each team member is using — to govern and enforce policies consistently everywhere from a single place, and to maintain business continuity and flexibility across regions and clouds.

CROSS-REGION, CROSS-CLOUD COLLABORATION

You should be able to instantly discover, access and share data, services and apps across clouds and regions — without ETL — on your cloud AI and data platform. This enables collaboration on the same data regardless of the cloud and region each team member is using through a single, consistent user experience.

CROSS-REGION, CROSS-CLOUD GOVERNANCE

When working with multiple clouds across many regions, you need to ensure the same governance and security configurations apply. Policies, tags and lineage should follow the data for consistent enforcement across users, workloads, clouds and regions. Your platform should offer a single governance pane across regions and clouds to minimize risk and unlock more value from sensitive or regulated data.

CROSS-REGION, CROSS-CLOUD BUSINESS CONTINUITY

A cloud AI and data platform should be able to replicate databases, accounts, pipelines and more across regions and clouds for resiliency, durability and failover by choice or in a stressed event. This safeguards mission-critical accounts and datasets so if an outage takes place, you can instantly resume activities without incurring downtime.

Your cloud AI and data platform should allow you to easily move workloads across clouds and regions so you can lift and shift operations flexibly as needed based on changes in business strategy. Having this type of deployment flexibility assists with geographic expansion and mergers and acquisitions. Furthermore, each department and division within your organization may have unique requirements and preferences. Rather than demand all business units use the same provider, a multi-cloud strategy allows each unit to use the cloud that works best for that unit. This is a strategic advantage for global companies because not all cloud providers operate in all regions.

THE BENEFITS OF A MODERN CLOUD AI AND DATA PLATFORM?

Today's organizations want an easier way to cost-effectively load, transform, integrate and analyze work on unlimited amounts of structured, semi-structured and unstructured data — whether that data is housed on-premises, in the cloud or in open table formats — in a versatile AI and data platform. They want to simplify and democratize the exploration of that data, automate routine data management activities, and support a broad range of data and analytics workloads while incorporating generative AI into the business. And they want to collect, store and analyze their data in one place, so they can easily obtain all types of insights from all their data.

A properly architected cloud AI and data platform brings this all together with a common set of services that streamline how data, apps and AI is used. The platform also enables you to consolidate diverse analytic activities, orchestrate the secure sharing and exchange of data, derive insights with AI and ML, and create modern applications.

Here are some of the most important reasons to consider a cloud AI and data platform:

FASTER TIME-TO-INSIGHTS

Organizations are increasingly transitioning from on-premises legacy systems to the cloud for several compelling reasons, foremost is time-to-value because deploying a traditional platform in your data center can take a year or more. During this time, business requirements may change, stakeholders may question the project's purpose, and the project may be exposed to risks such as staff attrition, economic downturns, revenue shortfalls and more. By comparison, a modern cloud AI and data platform can be up and running much faster, allowing your organization to start generating insights and value almost immediately. This rapid deployment helps your businesses respond more quickly to changing market conditions and customer needs, ultimately leading to a competitive advantage. A modern, cloud-based AI and data platform also provide faster insights by leveraging the power of scalable computing resources and advanced analytics tools, enabling your organization to process and analyze large volumes of data in near real time. This agility and speed on getting insights from your data can lead to better decision-making and improved business outcomes.

LOWER TOTAL COST OF OWNERSHIP

A cloud AI and data platform can help your organization avoid the high, fixed, upfront costs associated with on-premises solutions, which typically include hardware and infrastructure costs from computer servers, storage devices, data center space, high-speed networks, redundant power supplies and more. These platforms operate on a pay-as-you-go model, allowing your organization to scale its usage up or down as needed without incurring unnecessary costs. This flexibility can lead to substantial savings compared to the fixed costs of on-premises solutions. Storage and compute resources can scale independently, elastically and infinitely to meet the shifting demands of each workload, as well as to accommodate peak usage periods without reducing the performance of many concurrent workloads.

> **The cloud, by its very nature, gives you flexible capacity, ensuring you'll always have what you need when you need it.**

Not all cloud AI and data platforms are the same, however, so you'll want to make sure you have the correct architecture to take full advantage of what the cloud offers. Cloud AI and data platforms also often have lower software licensing fees and may not charge per named user, making them more cost-effective for your organization if you have a large number of users or frequent changes in user count. This is different from legacy solutions and some cloud-based platforms, which may require additional fees for each named user or impose limitations on user counts.

Typical platforms also require organizations to retain skilled IT personnel to configure, deploy and maintain systems. These skilled and often hard-to-find specialists will likely spend their time on low-value activities, such as tuning, rather than creating new applications that add value to the business. This can create potential bottlenecks when issues arise, because you may not have enough staff to handle development, maintenance and troubleshooting. A cloud AI and data platform should automate maintenance and performance tuning for you to reduce overhead costs.

Furthermore, many platforms require tedious and manual configurations for managing sensitive content, which involves constant attention to firewall protection, security posture management, encryption, user roles and access privileges, as well as responding in real time to cybersecurity threats. Effective governance and security are complex and costly to implement, especially in terms of human resources. Poorly implemented governance and security measures expose you to even more costs if your platform is breached. Before your organization can accelerate with gen AI and LLMs — or onboard any workload — you need to ensure the right people have access to the right content across data, apps and models while also protecting sensitive and personally identifiable information (PII) and adhering to a growing list of regulatory requirements.

BUILT-IN GOVERNANCE AND DISCOVERY

At the same time, your organization needs to be able to empower employees to quickly find and collaborate on relevant content to derive insights faster. To do this, you need to be able to effectively govern and discover with a unified set of capabilities. A cloud AI and data platform with built-in governance and security can help your organization quickly uncover and resolve cross-cloud security risks, universally enforce access controls across clouds, and easily apply out-of-the-box, proven governance protections to sensitive content. You can also quickly search, discover, access and share these governed data, apps and models from across your ecosystem to boost privacy-preserving collaboration. These capabilities streamline governance and security monitoring, lower TCO even more and prevent escalation of risks.

HOW DO YOU USE CLOUD AI AND DATA PLATFORMS?

A cloud AI and data platform is about maximizing the value of your data. This is done by bringing together modern technologies for storing, sharing and analyzing that data; for ingesting new types of data; for building new applications; and for delivering cutting-edge AI/ML initiatives. A modern cloud AI and data platform can enable, automate and improve these important workloads.

RUNNING ADVANCED ANALYTICS

A cloud AI and data platform should enable a diverse set of analytics use cases such as BI and reporting and geospatial analytics against a single copy of data. The platform should power a full spectrum of data formats and architectural patterns including data warehouses, data lakes, data lakehouses and data mesh while supporting multiple languages including SQL, Python, Java and Scala.

STREAMLINING DATA ENGINEERING

Traditional data pipelines require complex ETL procedures to extract, transform and load data. Furthermore, these legacy solutions can't handle and process all types of data, such as structured data from enterprise applications and machine-generated data from IoT systems, streaming data from social media feeds, JSON event data, and weblog data from internet and mobile apps. A modern cloud AI and data platform allows you to easily and efficiently ingest all these data types while supporting a range of popular data ingestion styles, including batch integration and streaming integration.

COLLABORATING ON CONTENT IN A PRIVACY-PRESERVING WAY

Collaboration shouldn't involve copying and moving content. A cloud AI and data platform should enable organizations to easily share content across data, apps and models in a secure and governed way, without requiring ETL. Rather than physically transferring content to internal or external teams and systems, the platform should enable seamless access to live, updated data, apps and AI products in a governed, privacy-compliant way. This type of advanced content sharing encourages collaboration by making it easier to share across teams, business units, business partners and customers. It also provides the opportunity to monetize content to create new revenue-generating services.

The cloud AI and data platform you choose should also extend collaboration to a much bigger scale in the form of a marketplace where organizations can offer or acquire third-party products, combine it with their own data estate, and gain new insights. Because data isn't copied or moved in these scenarios, you eliminate the cost, headache and delays associated with traditional data sharing, exchanges and marketplaces, which deliver only subsets or "slices" of data that must be continually refreshed through manual intervention.

BUILDING AND DISTRIBUTING APPLICATIONS

Software developers need cloud-native development tools to accelerate the process of developing and deploying new applications. A cloud AI and data platform provides unlimited compute and storage resources for development, iteration, testing and quality assurance (QA) activities. It eliminates the need to build infrastructure and automatically handles provisioning, availability, tuning, data protection and other operations across multiple clouds. You should also be able to easily distribute and procure apps securely on your platform.

ACCELERATING AI

A modern cloud AI and data platform should satisfy the entire data lifecycle of generative and predictive application development. For gen AI, the platform should provide support for models and services that make it easy to develop and deploy conversational interfaces with data – whether that data is a document like a sales contract or a database of tables with sales records. For ML, the platform should not only equip data scientists with a complete end-to-end ML and ML Ops platform to develop, productionize and consume ML models using Python libraries but also provide easy ways for analysts, data engineers and other SQL users who don't have Python or ML expertise to generate forecasts or find anomalies.

HOW CLOUD AI AND DATA PLATFORMS DRIVE SUCCESS IN INDUSTRIES

Cloud AI and data platforms can help organizations across industries gain insights quickly, drive efficiency, improve decision-making, personalize customer, patient and citizen experiences, and more. Organizations across industries can use a cloud AI and data platform as the foundation for key capabilities like enterprise analytics, applications, gen AI and data collaboration to power a wide range of valuable use cases.

INDUSTRY	USE CASES
Healthcare and life sciences	• Automate administrative and clinical tasks • Improve patient care management • Augment clinical decision-making • Personalize patient experiences
Financial services	• Make informed investment decisions and mitigate risk • Predict company performance • Personalize content to improve marketing ROI
Manufacturing	• Analyze cycle time to improve efficiency • Optimize yield to increase efficiency • Measure overall equipment effectiveness (OEE) • Employ predictive maintenance
Advertising, media and entertainment	• Improve subscriber experiences and grow lifetime value • Leverage applications to curate content, enhance personalization and monetization • Accelerate ad revenue with data clean rooms
Public Sector	• Securely modernize and share data • Collaborate with mission-critical data • Enrich data with third-party data for greater insights

HOW GEN AI IMPROVES PATIENT CARE AND ORGANIZATIONAL EFFICIENCY IN HEALTHCARE AND LIFE SCIENCES

Healthcare and life sciences organizations are increasingly embracing AI, but its use has been largely limited to data science specialists and a small subset of business analysts. Gen AI provides virtually everyone in an organization access to powerful AI tools, reducing the requirements of technical knowledge to obtain data-driven insights. Here are four ways gen AI can transform healthcare organizations:

1. **Improve care management:** Patient care management is complex and dynamic. Gen AI can help address its many factors and nuances by analyzing vast amounts of data in near real time and suggesting the next best actions, including analyzing patient and social determinants of health data to generate personalized care plans and assisting in managing chronic diseases by continuously monitoring patients' health status and adherence to treatment regimens.
2. **Personalize patient/member experiences:** Delivering personalized, effective care is increasingly important as more healthcare organizations adopt value-based care models. By analyzing vast data sets, gen AI allows healthcare payers and providers to quickly determine patient or member preferences, behaviors, sentiments and health trends to develop customized care plans and communications. This in-depth analysis allows them to create highly targeted and relevant plans and content, and refine them throughout the patient's care journey.
3. **Accelerate drug discovery and development:** Research and development in life sciences is an expensive and lengthy process that often takes more than 10 years. By analyzing vast amounts of biomedical data, including genetic information and clinical trial data, gen AI can predict interactions, identify novel drug targets, and optimize drug efficacy and safety profiles, speeding up drug discovery and development. It can also expedite personalized medicine by tailoring patient treatments based on in-depth clinical data.
4. **Assist clinical content generation and submission:** There is no shortage of content to be created and submitted in the life sciences industry. Through advanced automation and data-driven insights, gen AI can streamline and optimize how life sciences organizations generate and submit clinical content. Gen AI can be used for drafting, formatting and translating documents; checking for regulatory compliance; automating documents for submissions in alignment with journals and conferences; and curating content for researchers to stay current on developments in their field.

HOW FINANCIAL SERVICES COMPANIES LEVERAGE THIRD-PARTY DATA TO POWER DECISION-MAKING AND GAIN INSIGHTS

With customers' increasing demands for seamless experiences and a heightened focus on sophisticated modeling to power both traditional and predictive analytics, financial services companies have ever-expanding data needs for a wide range of purposes. Third-party data enables financial services organizations to optimize the performance of data models, improve decision-making based on more accurate forecasting, and deliver personalized experiences to customers. External data empowers teams across the business to make better data-driven decisions, especially when it's integrated with first-party data.

Below are three ways third-party data can power investment decisions, uncover insights that help predict company performance, and personalize content and offers when it's combined with first-party data sets.

1. **Make informed investment decisions and mitigate risk:** Financial services firms can leverage third-party data to enrich their portfolio analysis and help determine if investments will be profitable, risky or ethically questionable. For example, they can pull in an external data set to join with their own profit and loss data. Investors might also choose to integrate multiple data sources into their machine learning models to help with real-time decision-making for when to buy and sell.
2. **Leverage alternative data sources for critical decision-making:** When a public company issues reports about its financial situation, investors use this information to decide whether to buy or sell shares in that company. Institutional investors also have access to private markets with information about company fundamentals that aren't publicly available. Whether public or private, smart investors recognize that a balance sheet provides only a snapshot in time. That's why hedge funds and other financial services organizations use alternative data gathered from nontraditional sources. With the combination of social media streams, employment data, satellite imagery and more, alternative data helps these organizations predict whether a listed company will meet its earnings forecast and whether to buy or sell shares.
3. **Personalize content to improve marketing ROI:** By leveraging third-party demographic and life-event data — which might include information on household size, number of children in the household, net worth, home ownership status and property type — banks and other financial institutions can enrich the customer information that's already in their databases and obtain a true customer 360. From there, they can refine their segmentation and targeting strategies to increase the likelihood of customers receiving content and offers that are relevant to them.

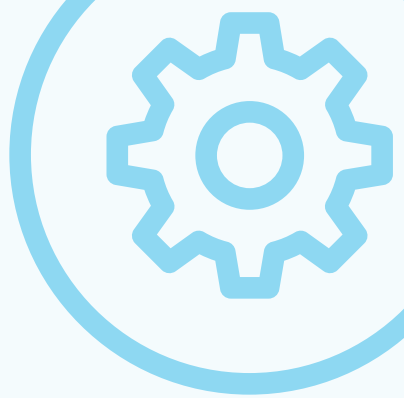

HOW MANUFACTURERS LEVERAGE DATA-DRIVEN ENTERPRISE ANALYTICS TO OPTIMIZE EFFICIENCY AND PERFORMANCE

Manufacturing companies need to be able to collect and analyze vast amounts of data, enabling more advanced analytics use cases and accelerated decision-making. To generate data-driven insights, businesses need a foundational cloud platform with data pipelines that can ingest both IT and OT data at scale. Here are four ways manufacturers are leveraging a single cloud AI and data platform to gain data-driven insights:

1. **Analyze cycle time to increase efficiency:** Cycle time is the time it takes for a machine or process to complete a single production cycle. Reducing cycle time can improve efficiency and productivity because more units can be produced in a given period. Data from machines coupled with other data sets can be used to identify bottlenecks and inefficiencies in the production process, and AI/ML algorithms can analyze this data and suggest process improvements. And near real-time monitoring of cycle time allows for real-time optimization.
2. **Improve supply chain performance:** The supply chain presents significant challenges for manufacturers as they navigate supply and demand volatility, lack of visibility into the supply chain, and the need for accurate and complex forecasting. A cloud AI and data platform provides manufacturers with real-time visibility and traceability across the value chain to internal and external stakeholders as well as customers. Event-driven architectures allow for risk assessment and preventive action by helping to predict, identify and respond to disruptions and risks. And with robust forecasting capabilities, companies can plan, forecast and manage inventory amid fluctuations in supply and demand.
3. **Measure yield and overall equipment effectiveness (OEE):** Yield is the proportion of product that is successfully produced by the total amount of raw materials used in the production process. It's an important measure of efficiency in manufacturing, as it can impact the overall cost of production. Improving yield can involve identifying and addressing bottlenecks or inefficiencies in the production process, which can potentially lead to quality improvements. OEE measures how well a production process is used and considers factors such as availability, performance and quality to determine efficiency by measuring factors such as the percentage of time a production line is available for production and the speed at which the production line operates.
4. **Employ predictive maintenance:** Predictive maintenance is a strategy that uses data and analytics to predict when equipment or machines are likely to fail or require maintenance, so maintenance can be scheduled in advance. Scheduling maintenance in advance can minimize the risk of unplanned downtime and improve the overall reliability and efficiency of the production process.

HOW THE ADVERTISING, MEDIA AND ENTERTAINMENT INDUSTRY USES APPS TO POWER COMPETITIVE ADVANTAGES

Data cloud platforms play a pivotal role in the advertising, media and entertainment industry. Platforms that can power capabilities such as AI, apps and data clean rooms provide flexible, innovative solutions to enhance consumer engagement, deliver personalized experiences and accelerate advertising revenue — all while protecting consumer privacy. Here are three ways industry leaders can use cloud AI and data platforms to gain competitive advantages:

1. **Improve subscriber experiences and grow lifetime value:** With a cloud AI and data platform, companies can gain a customer 360 and leverage data at every touchpoint. They can create a custom personalization engine, connect behavior across properties and enrich profiles with third-party data to tailor recommendations to subscribers. A cloud AI and data platform also helps companies analyze marketing and advertising campaigns to help improve efficiency and performance. And by bridging subscriber behaviors and marketing campaigns across analytics environments and providers, they can gain deep insights into campaign performance and identify opportunities for optimization. They can also understand churn and retention propensity across properties.
2. **Leverage applications to curate content, enhance personalization and monetize:** With apps, companies can curate content based on user interests, creating customized playlists, news feeds or streaming queues tailored to individual tastes. Apps offer interactive experiences to engage users and gather feedback for content optimization by analyzing user behavior and preferences. They can recommend personalized content, integrate with social media, and use push notifications and game-like elements to enhance user engagement. Companies also can monetize apps with in-app ads, in-app purchases and sponsored content placements. Subscription models can offer access to premium content, ad-free experiences and early access to generate revenue, further diversifying revenue sources.
3. **Accelerate ad revenue with data clean rooms:** Data clean rooms, powered by a data cloud platform, can serve as the critical foundation for secure data sharing that preserves consumer privacy. Data clean rooms provide a secure environment for multiple companies — or divisions of a single company — to collaborate with and across data sets for identity, enrichment, joint analysis, targeting or media measurement in a privacy-preserving way. This enables data interoperability between companies and their partners and customers while maintaining data security and governance. Advertisers can analyze and enrich the data to increase campaign effectiveness, drive media optimization and define the media mix — ultimately driving advertising revenues. For publishers, the end result can be improved audience experiences, increased subscription revenue and differentiated ad products.

HOW THE PUBLIC SECTOR IMPROVES DATA-DRIVEN DECISION-MAKING WITH SECURE DATA SHARING AND COLLABORATION

Government agencies need the agility and flexibility of a dependable, single-data platform that allows them to break down silos and access essential data instantly. With that foundation, agencies can securely share and collaborate with data within and across organizations and agencies to achieve strategic objectives and serve the public good. Here are three ways government and education organizations can benefit from data sharing and collaboration:

1. **Securely modernize and share data:** The public sector faces unique data challenges when it comes to data modernization because data is often siloed in legacy systems in different agencies' self-selected cloud applications. Sensitive data is also difficult to share securely when it needs to be copied and moved into another system. With faster data sharing, government agencies can benefit from timely delivery of service and insights. A modern cloud AI and data platform provides agencies with the ability to securely share data across teams, between departments and with essential partners in near-real time, without moving or copying data. They also can share data across different cloud providers.
2. **Collaborate with mission-critical data:** Governmental leaders want to collaborate on efforts such as improving emergency services, tackling cybersecurity and providing responsive citizen services. A cloud AI and data platform provides access to mission-critical data to enable the cross-agency flow of information for better-informed decisions and improved efficiency. With data collaboration, government organizations can scale elastically in times of need to support access to critical data for all users who need it.
3. **Enrich data with third-party data sets:** Government data analytics can be enhanced with real-time access to third-party data from trusted data partners. With third-party data sets, agencies can access data to enhance or fill data gaps in their own data while keeping personally identifiable information (PII) private. This provides leaders with access to all the available information in near-real time, empowering them with data-driven decision-making to help protect and serve the public and achieve mission success.

THE SNOWFLAKE AI DATA CLOUD: A MODERN CLOUD AI AND DATA PLATFORM THAT DELIVERS

Snowflake's vision has always been to open the door for everyone across an organization to work with data to drive innovation, increase productivity and sharpen their competitive edge.

The Snowflake AI Data Cloud provides a single, unified data foundation that puts the power of data and AI into the hands of your enterprise. Your organization can access data on your terms, including data that's unstructured, in open formats, and from third parties, simply by the AI Data Cloud eliminating silos and removing the complexity of managing legacy data infrastructure.

Snowflake's unique architecture unites all data types, workloads, languages and runtimes on a common data foundation, with elastic compute that can scale up and out as needed — allowing you to only pay for what you use.

The AI Data Cloud is also easy to use because it comes with built-in simplicity, the ability to scale easily and unified governance, offering a platform that is a fully-managed service, ultimately delivering massive TCO for you — making you the star of your organization.

The Snowflake AI Data Cloud also makes it easy to get value from AI and ML securely against your enterprise data, with efficiency, ease and trust. This same foundation can be used to easily build both internal and commercial apps, and bring AI and analytics to life without taxing your operation. Builders will love you because they can simply build, deploy, and even distribute and monetize their apps globally to other organizations in the AI Data Cloud.

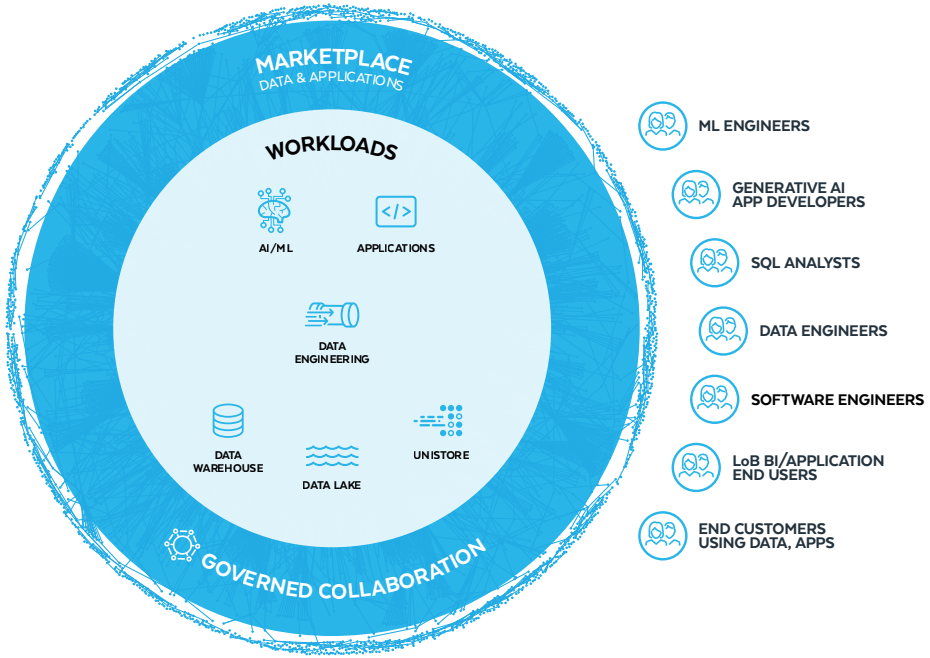

SNOWFLAKE'S PLATFORM ARCHITECTURE

- **Interoperable storage:** Snowflake's interoperable storage is the platform's ability to seamlessly work with various data storage systems and formats at any volume. With all this data together, Snowflake also provides unique automations and optimizations such as built-in encryption by default, storage compression and micro partitioning to keep your data fast, efficient and secure at scale. With Snowflake, you get one platform that can support a wide range of architectural patterns, including a data lake, data warehouse, lakehouse or data mesh. Interoperability to connect with data on-premises or stored in open table formats, like Apache Iceberg, further extend the value of the AI Data Cloud to your entire data estate.
- **Elastic compute:** Snowflake separates storage from compute and has one single engine that can support every workload — from running large, complex pipelines and machine learning models to interactive applications and analytics dashboards and beyond. With leading performance and concurrency out of the box, Snowflake elastically scales instantly to support practically any number of users, jobs, or data; and multi-cluster resource isolation means reliable performance every time. Whether users prefer to work in SQL, Python or Java or if they want to work in native development interfaces or even bring their own IDE or third-party tools, Snowflake provides flexibility and programmability.

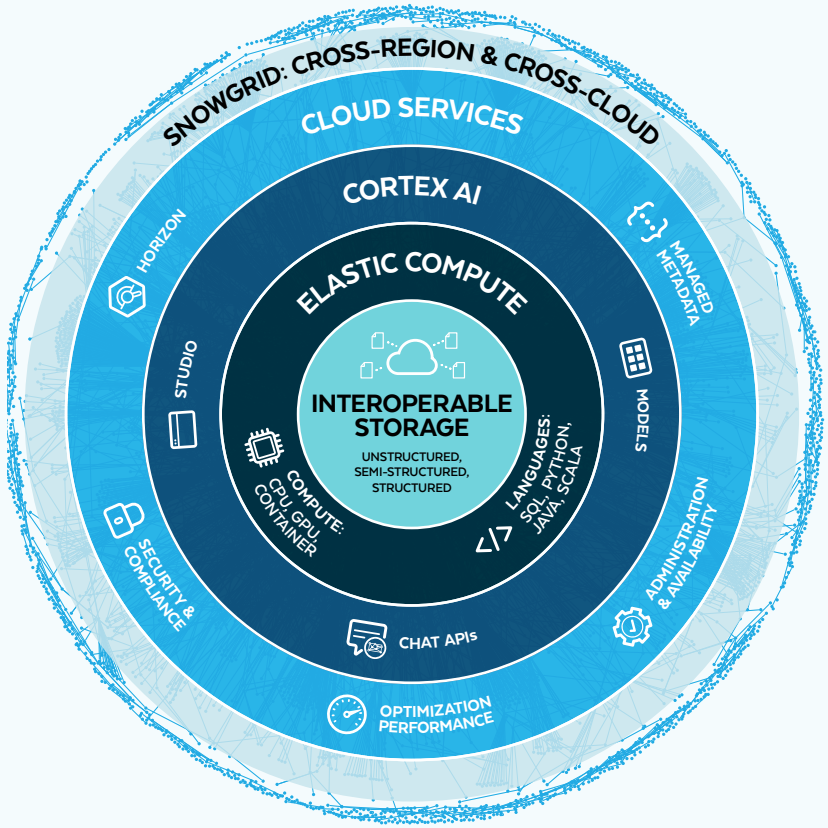

- **Cortex AI:** Cortex AI provides your organization the efficiency, ease of use and trusted benefits of the full platform that also apply to gen AI. The platform features a fully managed set of models, including Snowflake's Arctic and other leading models, which are ideal for tasks like text processing, sentiment analysis and custom summaries. To make data accessible to everyone in your organization, Snowflake offers two fully managed services for structured and unstructured data, such as documents. Accessing these features is easy with multiple code interfaces like SQL and Python, and Snowflake Studio provides a no-code development option. Snowflake trusts the reliability of Cortex AI and has built products on top of it — like Snowflake Copilot and Document AI — to further enhance the platform's capabilities.
- **Cloud Services:** Snowflake Cloud Services offers a fully managed service that automates costly and complex operations to reduce overhead and improve efficiency. With near-zero maintenance, there are no upgrades or downtime to worry about. Snowflake Horizon provides unified compliance, security, privacy, interoperability and access capabilities without additional configurations or protocols. Snowflake also offers optimized performance by default, with transparent performance enhancements automatically applied through releases.
- **Snowgrid:** Snowgrid is a cross-cloud technology layer to connect businesses' ecosystems across regions and clouds, enabling global collaboration and data sharing. It offers universally enforced governance, business continuity through data replication and synchronization, and it complies with changing regulations and regional controls. Snowgrid simplifies governance at scale, minimizes risks and unlocks more value from sensitive or regulated data. It allows for seamless cross-region and cross-cloud operations, ensuring business continuity and delivering leading experiences to customers.

WHAT YOU NEED FOR A STRONG DATA FOUNDATION THAT ACCELERATES ENTERPRISE AI, AND BUILDING AND DISTRIBUTING APPS

There are three requirements for your organization to achieve a solid data foundation. It requires going from:

1. **Endless silos to unified data.** Silos lead to complexity and moving data to get value. Switch to a single platform to access all data, including data that's unstructured, in open formats and from third parties.
2. **Hidden costs to superior TCO.** When you have multiple services, costs can add up and it can become hard to understand total operating costs — especially if the underlying infrastructure is billed separately or if they require manual expertise to manage and optimize. Switch to a fully managed service to efficiently support users and workloads at scale, while reducing the time, effort and expertise required to operate.
3. **Piecemeal policies to universal governance.** Inconsistent governance policies across systems and users can introduce security risk to your data. Switch to a single governance model with comprehensive compliance, security and privacy controls that are universally enforced.

Once your organization has built a strong data foundation, you're ready to succeed with AI. Snowflake can help your organization accelerate enterprise AI by:

- **Speeding up long cycles to make AI easily accessible:** Working with AI often takes long cycles and it often requires data scientists with expertise to unlock its true value. Snowflake, however, can help you get more users in your organization to benefit from the value of AI by making it easy for them to use industry-leading models in seconds.
- **Taking care of infrastructure management to help you quickly build AI apps:** Managing GPUs and other infrastructure required for AI can be costly and time-consuming — preventing your team from dedicating resources to innovation. With Snowflake, you can build custom AI-powered apps in minutes by accessing fully managed infrastructure and robust primitives.
- **Bringing AI to your secure, enterprise data so you can move beyond experimentation:** Building and deploying models often means needing to move data, which limits auditability and introduces risk, leaving most development stuck in experimentation mode. Snowflake brings AI to your enterprise data so you can safeguard your IP for end-to-end development and get full observability of assets without ever leaving the security boundary of your data foundation.

This same data foundation also can help your organization scale your business with applications to help you drive value from data inside and outside the company. The AI Data Cloud can help your organization successfully build and distribute applications by going from:

- **Architectural complexity to building with ease:** Most apps never get built because of the complexity and effort needed to make them. The AI Data Cloud lets you easily build any app — both internal and commercial — with capabilities ranging from LLMs to embedded analytics.
- **Unpredictable admin burden to efficient scaling:** App usage can often be unpredictable, and fluctuations can impact operations as well as end-user experiences. A strong data foundation helps you efficiently handle changing user demand with automatic scaling and fully managed infrastructure.
- **Limited distribution to secure global deployment:** Security concerns can prevent the distribution and procurement of apps. With a secure data foundation, you can share apps globally with teams inside your ecosystem or to the broader AI Data Cloud network. Consumers can also run apps where the data already is.

The Snowflake AI Data Cloud is the AI data and cloud platform that can unify all your data and workloads to eliminate complexity while delivering superior total cost of ownership and built-in governance to accelerate your enterprise AI and the building and distribution of applications.

Are you ready to learn more about how the AI Data Cloud can enable your organization to mobilize data to reach actionable insights?
Start your free 30-day trial of Snowflake today.

ABOUT SNOWFLAKE

Snowflake enables every organization to mobilize their data with Snowflake's Data Cloud. Customers use the Data Cloud to unite siloed data, discover and securely share data, and execute diverse artificial intelligence (AI) / machine learning (ML) and analytic workloads. Wherever data or users live, Snowflake delivers a single data experience that spans multiple clouds and geographies. Thousands of customers across many industries, including 647 of the 2023 Forbes Global 2000 (G2K) as of October 31, 2023, use the Snowflake Data Cloud to power their businesses.

Learn more at **snowflake.com**