

# The public sector guide to delivering value from data and AI

5 steps for maximizing AI  
success in your agency



# AI Success Relies on Data

95%

foresee an increase in AI in their agency over the next 1-2 years<sup>1</sup>



48%

plan to use AI to support predictive analytics for decision making, up from 35% in 2023<sup>1</sup>



53%

expressed concerns regarding budget/cost of AI initiatives<sup>1</sup>



67%

believe AI will improve their agency's overall efficiency and productivity<sup>2</sup>



69%

are investing in training agency staff to close the AI skills gap<sup>2</sup>



Across the board, public sector agencies are eager to embrace gen AI, but they're facing challenges in scaling from pilot projects to full production deployments. What's holding them back?

## Their data.

As a data leader, you are at the forefront of fueling your agency's AI transformation. Gen AI thrives on comprehensive, accurate, and high-quality organizational data. Unfortunately, many agencies struggle with fragmented data silos, poor data quality, lack of proper governance, and other challenges.

To unlock the true potential of gen AI for transformation, agencies need to rebuild their data strategies—focusing on building modern data platforms that serve as a robust foundation for gen AI initiatives.

Based on our work with public sector agencies on their data and AI platform strategy, this guide shows you five steps to building a solid foundation for AI success.

# 5 steps for building a data foundation for AI success



01

Create an AI-first data strategy



02

Connect all your data to AI to drive insights



03

Boost productivity with agents in every data task



04

Ensure data security and governance



05

Improve efficiency to reduce data cost and scale



# Create an AI-first data strategy

# 01

# Data is everybody's business, and it's the fuel for AI.

Combining a strong data strategy with a thriving data culture ensures that every team member understands, uses, and contributes to the value of your data—and leads to better utilization of AI.



# How do you build a strong data strategy?

One that aligns with your organizational culture, fosters collaboration across teams, and creates a more accessible data ecosystem. And most importantly, a data strategy that's tailored to meet the demands of the AI era.

Here's how. An AI-first data strategy should:



## Treat data as a strategic asset.

**Shift the agency's mindset** to recognize data's intrinsic value and align people, processes, and technology to meet the agency's core mission.

**Ensure that your data and AI strategy supports your agency mission** – to help improve government services, enhance decision-making, and ultimately create a more efficient and effective public sector.

**Weigh the implications of inaction** in the context of today's landscape.

**Implement strong data governance policies** and invest in modern cloud data management tools that address the data to AI lifecycle.

**Consider whether your proposed data strategy** is able to support your future AI strategy at scale—or whether it will be a lift to organize your data for each model.

**Provide employees with the skills,** tools, and incentives to utilize data effectively in their decision-making processes.

**Create data products** and strategies that unlock new operational efficiencies. Streamline your data outputs with embedded analytics and AI infusion.



## Create product owners and define roadmap ownership.

**Consider leveraging a data mesh strategy.** Appoint a central data team that establishes governance requirements, supports the data platform, and creates the data strategy framework. This team should periodically collect feedback from stakeholders to understand if the data stack is working well and continually meets emerging agency needs.

**Organize the rest of the team,** including stakeholders and data practitioners, by data products and data domains—resulting in stewardship, continuous management, and improvement in their respective domains. Stakeholders should be responsible for their data products from sourcing to offering insights.

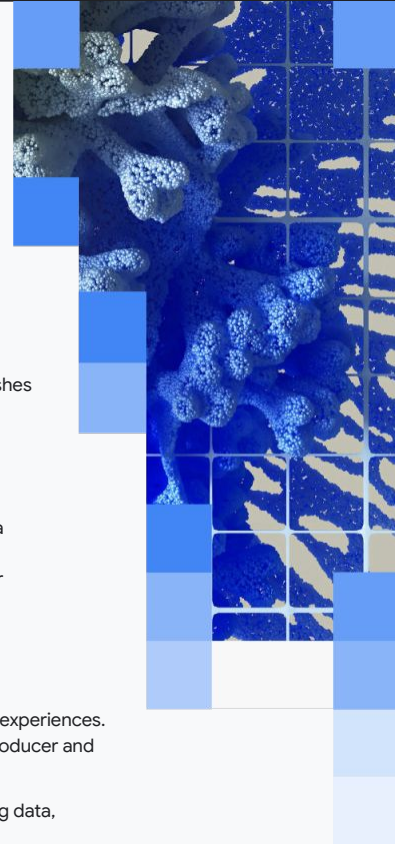


## Build a data ecosystem.

**Empower data producers and consumers** with streamlined, unified, and productive experiences. Also, ensure your data ecosystem is able to seamlessly support the blurring of data producer and consumer roles as models are trained and retrained.

**Establish principles** and practices for publishing, discovering, building on, and trusting data, enabling efficient collaboration and innovation.

**Provide a roadmap** to address inefficiencies in the “usability” of data. This roadmap helps you realize the true value of data by maximizing its discoverability, accessibility, and quality.



## An effective data culture creates a shared understanding of using, owning, and creating value from data.

This shared set of values, expectations, and capabilities is the foundation for successfully implementing a data strategy that can help serve and accelerate mission impact.

What will yours look like? Your agency is unique, and so your data culture and strategy will be, too. But know that the more clearly you establish and internally communicate your data culture and strategy, the easier it is to drive internal adoption and develop the required skills internally to unlock the value of AI.

Establishing and sharing an AI-first data vision is crucial to removing the central bottlenecks and silos in your data.



**Connect all your  
data to AI to  
drive insights**



**02**

# Agencies thrive on real-time insights.

## And they're powered by today's AI technologies.

But in order for gen AI to deliver, your data needs to be available as input. All of it. Not sitting in silos.

The need to unify your data has not only become more urgent in the AI era, it's also become more challenging. Simply putting all your data in one place is often not practical or cost effective. That's why it's important to create a cohesive cost optimized data foundation that breaks down silos as well as works across clouds—establishing a single source of truth that acts as the bedrock for AI initiatives.

Here's how to do it.



# An open data approach to fuel your lakehouse for AI

Multimodal data often becomes fragmented across different platforms, and potentially even stored in diverse table formats, making it difficult to unify your data and incorporate unified security and governance. To take full advantage of AI, it's imperative to create a single source of truth with a unified platform across all forms of data, including open formats.

In addition, a highly scalable architecture is crucial, enabling the unification of transactional and analytical systems without impacting performance. This allows for seamless data access between these systems, even with demanding workloads.



Here are best practices when developing an open data approach:



## Map capabilities you need in the data platform.

When selecting a platform for your data foundation, prioritize scalability, flexibility, and integration across data management, processing, governance, ML, AI, and BI capabilities. Map all your use cases and workflows to the capabilities of your data platform to identify any gaps.



## Connect structured, semi-structured, and unstructured data.

Your data platform should accommodate all types of data, including structured data (e.g., relational databases), semi-structured data (e.g., key-value, JSON, XML), and unstructured data (e.g., text, images, video). An increasing amount of data—including constituent information, employee handbooks, and more—is raw, unstructured, and diverse. In fact, unstructured data currently makes up 90% of data, and so it's important to make sure your platform of choice is suited for extracting insights from a variety of formats.



## Define your strategy for multimodal data for enhanced search and AI applications.

Build a single-copy architecture for all your structured and unstructured data across clouds and open file formats. Accomplish this by converting multimodal data into numerical representations, or vectors, for AI models including LLMs to generate more contextually relevant and meaningful responses when interacting with users or processing information.



## Simplify your data platform with a single pane of glass across open formats.

This includes a single UX and collaborative workflows. Ensure compatibility with popular open data formats like Apache Iceberg, Apache Hudi, and Delta Lake, each offering unique benefits for data management and analytics. Unification also alludes to interoperability across engines (i.e., SQL, Spark, Python).



## Implement data cataloging and metadata management.

Create a comprehensive data catalog to enable easy access, and understanding of your assets through AI-enabled data discovery, profiling, quality, lineage, and lifecycle management. Use a unified metastore and governance model to ensure consistency, flexibility, risk mitigation, compliance, and trust across all assets.

# Harness all data types and formats.

Once you've integrated disparate datasets, you empower your AI models to leverage the full spectrum of organizational information—resulting in more accurate, actionable insights. To effectively harness the power of your data and AI in this complex landscape, a robust and adaptable strategy for analytics is essential.

Here are some tips:



## Create a unified storage engine.



Build a unified interface for analytics and AI engines to query data in a secure, governed, and performant manner. Maintain a single copy of data and make it uniformly accessible across engines, while centrally managing security policies in one place, and have it consistently enforced across the engines.



## Enable cross-cloud data access and analysis.

Implement solutions that allow you to access, process, and analyze data stored in different cloud environments seamlessly. This approach reduces risk, minimizes data duplication thus reducing cost, and improves security by centralizing data governance and access controls.



## Support diverse analytics engines.

Your architecture should support different types of analytics engines. For example, Spark for large-scale data processing, graph processing to find hidden connections in your data, a BI engine for interactive dashboards, and TensorFlow for ML model development.

# Activate your data for real-time insights.

In order to support your mission, public sector agencies need to make decisions and respond to events.

Real-time data streaming has emerged as a critical technology enabling agencies to process and analyze data as it's generated. Incorporating streaming analytics into your data and AI strategy requires careful consideration of several architectural aspects to ensure efficient and effective processing of real-time data.

Real-time data processing requires a platform that offers:



## High-throughput, low-latency data ingestion.

Your platform must efficiently handle diverse, high-volume data sources with minimal delay. It should use message queues (e.g., Kafka, Pub/Sub) and ingestion frameworks (e.g., Kafka, Kinesis, Pub/Sub) to support various data formats (e.g., JSON, Avro) while ensuring data consistency through a schema registry.



## Scalable and fault-tolerant processing.

Horizontal scaling handles increasing data volumes and maintains availability in the event of failures. This requires distributed processing frameworks (e.g., Spark Streaming, Flink, Dataflow), microservices architecture, and containerization (e.g., Docker, Kubernetes). Stateful processing is essential for computations that span multiple events.



## Flexible data transformation and enrichment.

Real-time data transformation, enrichment, and aggregation using stream processing APIs (e.g., Beam, Kafka Streams), integrated with external data sources and machine learning models is crucial for advanced analytics. Data quality checks should be implemented to ensure data integrity.



## Real-time analytics and visualization.

Integrating with the appropriate tools enables the creation of real-time dashboards, alerts, and ad-hoc querying. Deeper analysis can be achieved through connectivity with business intelligence tools.



## Operational monitoring and management.

Robust logging, monitoring, and alerting mechanisms ensure operational visibility and proactive issue resolution. Auto-scaling and load balancing are essential for optimal performance under varying workloads.



# Central Texas Regional Mobility Authority: Building the Road Ahead with Google

Faced with rapid population growth and increasing traffic, the [Central Texas Regional Mobility Authority](#) (CTRMA) initiated a plan to implement an advanced road management system. With \$800 million already invested in road construction, CTRMA recognized the need for a system capable of handling complex data from various sources, including road sensors, toll transactions, and license plate reading cameras. The agency sought a custom-built, centralized data platform to support its high-volume datasets and standardize input from other roadway toll systems.

CTRMA partnered with [Google Public Sector](#) and [Deloitte](#) to design and build this platform, leveraging Google Cloud's [BigQuery](#) for data management, [Looker](#) for analytics and visualization, and [Apigee](#) for third-party integrations. This design prioritized scalability and flexibility. To ensure a smooth transition from the legacy system, the new and old systems ran in parallel for six months. Since its launch in August 2023, the platform processes 750,000 toll transactions daily and has generated over \$225 million in revenue.

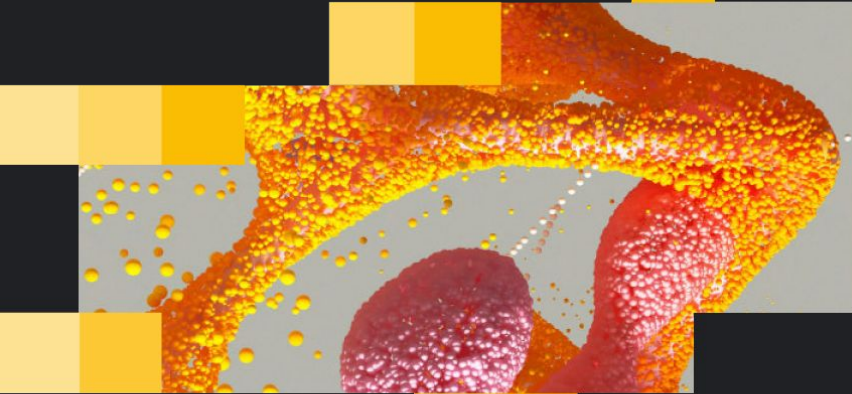
The new platform has enabled CTRMA to own and visualize data in real-time, providing insights into trends and improving operational efficiency. Call center agents can quickly access transaction details and images, reducing the time spent on inquiries. Future plans include integrating real-time traffic data from [Google Maps](#) and using [Gemini](#) to allow staff to analyze data using natural language. CTRMA also intends to expand the platform to include data from other regional transit systems and parking garages, with a continued focus on improving the motorist experience.

[Read the full story](#)

Google Cloud



**Boost productivity  
with agents in  
every data task**



**03**

# It's a virtuous cycle.

Unified data enables you to harness the power of AI and gain real-time insights. And it's through the use of AI that you can unlock the full potential of your data.

Since AI is a suite of tools that can be used at many different levels and for different purposes, the potential for positive gains amplifies with each integration.

Here's how to get started and accomplish your larger AI initiatives.



# Accelerate your data to AI journey.

Want it to be simple for users with various levels of technical abilities to build, deploy, and manage AI models and applications? Have a look at your platform. By selecting one that brings AI closer to your data and integrates AI agent experiences into every user journey, you can improve productivity and outcomes. **Here are the best practices to bring AI agents to your data:**

## Drive productivity with AI agents.

Gen AI is more than just content creation—it can transform how users interact with your data platform and accelerate their work. Integrate gen AI capabilities into your data platform to provide always-on, AI-powered assistants for your users. AI agents can assist with writing queries, migrating and managing data, creating data transformations, and building data visualizations. They can also simplify analysis by interpreting natural language questions and translating them into appropriate queries. Additionally, AI agents can provide intelligent recommendations to optimize budget-conscious resources, while continuing to ensure security, governance, and compliance.

## Improve enterprise knowledge discovery.

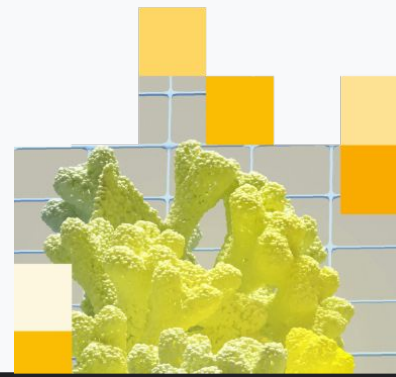
Integrate model development, tuning, and inference capabilities directly into your data platform. This eliminates the need to move data to external AI management tools, which can be time-consuming, costly, and introduce risks. Built-in gen AI capabilities such as vector embeddings and vector search allow for a deeper, more meaningful understanding of your data directly within your data platform, whether the data resides in an operational database or a data warehouse.

## Simplify model development and deployment.

Building, deploying, and managing ML models can be complex and require specialized expertise. A modern data platform should simplify this process by providing easy-to-use tools and frameworks for model development and deployment. These tools can abstract away the complexities of coding languages and infrastructure management, allowing data scientists and analysts to focus on building and deploying effective models without needing deep knowledge of underlying technologies.

To empower users further, consider seeking out a data platform that offers:

- **Pre-trained models** for a variety of tasks such as speech-to-text that users can adapt to their specific needs instead of building them from scratch.
- **Tools for fine-tuning models** to specific needs and data characteristics, resulting in more accurate and relevant outcomes. Not all models are one-size-fits-all.
- **Capabilities to analyze and connect insights** from unstructured data with structured data, which opens up opportunities for semantic search and model fine-tuning. Learn more about [multimodal AI](#).







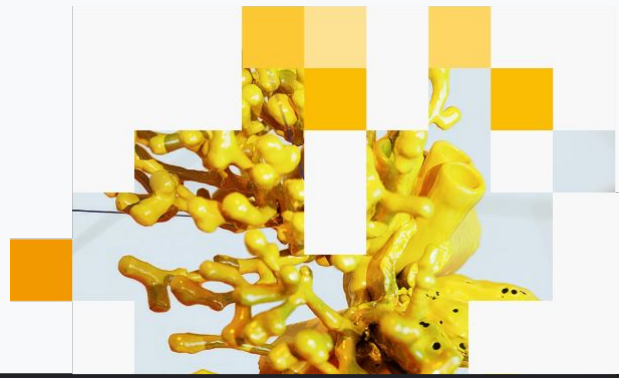
# Enable conversational analytics

Conversational analytics delivers a new approach to data analysis that combines natural language processing with gen AI, enabling users to ask questions of their data in plain language and receive clear, concise answers. Fueled by semantic modeling and modern knowledge graph capabilities, analytics empower agencies to unlock deeper insights, make better decisions, and drive innovation. Conversational analytics eliminates the need for users to understand complex data structures or query languages, making data analysis accessible to everyone.

A data platform built for AI needs a semantic model built on shared definitions and figures, enabling anyone in the agency to ask detailed questions and gain accurate insights. Multi-turn capabilities allow for follow-up questions and greater specificity.

Conversational analytics offer several benefits:

-  **Increased data accessibility** empowers everyone to access and analyze data, regardless of technical expertise, leading to better decision-making.
-  **Reduced time to insights** streamlines the data analysis process, allowing users to quickly get the information they need.
-  **More agile and accurate decision-making** provides real-time access to data, enabling users to quickly identify trends and anomalies.
-  **Reduced reliance on centralized IT and data analysts** frees up their time to focus on other strategic initiatives by enabling users to self-serve their data needs.





**Ensure data  
security and  
governance**

**04**

# You have more data than ever before.

And keeping it secure and compliant—while extracting valuable insights—is an increasingly complex task.

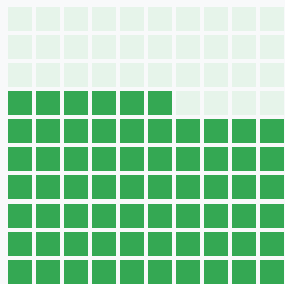
The more data you collect and analyze, the more security and governance matter.

We'll cover some best practices for governing and securing your data.



# Manage and govern data, end-to-end.

Data governance has always been a necessity, but the exponential growth of data in the AI era has made achieving it more complicated. It's how you can connect high-quality data to successful AI initiatives, ensuring compliance, empowering data-driven decision-making, and maximizing efficiencies by addressing data-related challenges like dark, duplicate, and dirty data. The goal is to establish end-to-end, centralized governance and organization-wide search. Start by creating a data to AI catalog across distributed data and implementing AI-powered data profiling for end-to-end data lineage.



# 66%

of organizations have at least half of their data dark, posing significant risk<sup>6</sup>

Here's a detailed look at how to build a robust data and AI governance framework:

## ⌘ Prioritize data quality to drive efficiency.

To drive greater efficiency through AI initiatives, public sector agencies must prioritize data quality by ensuring data is fully accurate, regularly updated, and complete. Implementing data observability tools can help monitor data lineage, detect anomalies, and address data quality issues proactively. High-quality data builds trust, empowers decision-making, and leads to improved user experiences and operational efficiency.

## ▮ Rise above architecture for compliance and efficiency.

As requirements evolve, agencies need a proactive, end-to-end data governance strategy that transcends specific architectures. This involves implementing policies that ensure compliance across jurisdictions and technological environments.

## 🕒 Banish dark, duplicate, and dirty data to ensure accuracy and consistency.

Data that has been collected but isn't in use; data that appears more than once; and data that is inaccurate, incomplete, or inconsistent must be addressed. Tools and processes are needed to identify and classify dark data, including unstructured data from various sources. Data deduplication techniques are essential for ensuring data accuracy and consistency. Implementing data quality management processes helps identify and correct errors, ensuring AI systems are trained on reliable data.

## ✦ Govern AI models in production.

In a changing landscape, and as new regulations for AI are being established, it is crucial to maintain transparency and accountability in your AI initiatives. By tracking the training data and retraining schedules, and building model explainability frameworks, you can help ensure the model's reliability, fairness, and compliance. This can help foster trust and allows for continuous improvement and responsible AI deployment.

[Learn more about data governance →](#)

# Incorporate multi-layer security.

Mandiant, in their [2025 M-Trends Report](#), warns that global attackers are now seizing every opportunity to further their objectives. One way they are doing this is through the use of infostealer malware, which is increasingly being used to enable intrusions using stolen credentials. Another growing trend is the targeting of unsecured data repositories, brought on by the lack of basic security hygiene.

Additionally, attackers are exploiting the gaps and risks introduced as organizations continue their migrations to the cloud.

Here are some first steps you can take to secure your data:



## Implement strong access controls.

Establish a robust identity and access management system with granular control over data access. Utilize multi-factor authentication with multiple verification methods, integrate with third-party identity providers, and define governance rules for resource usage and configurations.



## Protect the perimeter of your data platform.

Create a protective perimeter around your valuable data and define precisely which users or services can reach the data within your protected zone. Deploy a scalable cloud-native firewall with threat detection and prevention capabilities. Implement a web application firewall (WAF) to protect against DDoS, XSS, and SQLi attacks.



## Protect sensitive data.

Scan and classify data to understand its composition and associated risks. Employ data obfuscation, de-identification, redaction, masking, tokenization, and transformation techniques. Utilize a secrets management system for secure credential storage and management. Consider managed encryption keys for greater control over data encryption.



## Enforce continuous monitoring and auditing.

Maintain data integrity and availability through continuous security monitoring, anomaly detection systems, and data backup and recovery mechanisms. Implement data validation and reconciliation processes. Collect and analyze logs in a SIEM system for real-time insights into system activities and data access patterns. Ensure compliance with data privacy regulations through appropriate security controls and regular audits. Utilize tools that simplify compliance and security management for sensitive data and workloads in the cloud.



# Connecting residents with high-quality, affordable healthcare at unprecedented speed

Covered California is the state's health insurance marketplace, connecting uninsured Californians with affordable, high-quality health plan options. However, the eligibility determination process, which involves manual paperwork processing, can be slow and inefficient, causing delays in approval and active coverage. To address this, Covered California sought a more streamlined, automated solution for verifying documents.

Covered California partnered with Google Cloud and Deloitte to develop a proof of concept using Google Cloud's Document AI. This generative AI-powered solution automates resident information verification by using machine learning to improve data extraction from documents. Key considerations for Covered California included verification accuracy, the volume and variety of documents, and the solution's security and compliance capabilities, given the sensitive data involved.

The Document AI pilot showed promising results, achieving an initial automated verification rate of 80-96%, a significant improvement over the previous 28-30% rate. Covered California is focused on further improving this rate and integrating new document types. By also using Google Cloud's Assured Workloads and Google Security Operations, Covered California is enhancing security and compliance. The implementation of these Google Cloud solutions aims to reduce costs, increase access to healthcare for Californians, and improve the overall customer and employee experience.

[Read the story](#)

Google Cloud





**Improve efficiency  
to reduce data cost  
and scale**

**05**

# Your data is both an asset and a liability.

While your data can power tremendous insights and streamlined output, it also is increasingly costly to store.

Refining the greatest efficiency from your data requires capitalizing on its value while simultaneously keeping pace with its growing demands.



# Minimize costs with thoughtful data platform design.

Public sector agencies are collecting and processing unprecedented volumes of data, driving the need for scalable and powerful analytics data platforms. However, the exponential growth of data also brings the challenge of escalating costs. Unchecked, these costs can strain budgets, hinder innovation, and limit the potential of AI initiatives.

Here are some best practices of optimizing data and analytics costs, so your AI initiatives deliver a higher ROI. And freeing your resources up for more innovation, too.



## Automate workloads with managed services

**Migrate your data and AI workloads** to a cloud environment to eliminate the need to monitor, troubleshoot, update, tune, and plan your underlying infrastructure as you scale.

**Minimize overhead and management costs** with serverless architectures, which eliminate patching and maintenance—and are optimized for cost.

**Scale up or down** as needed to meet changing demands with cloud architectures, enabling you to quickly act on new opportunities without the need to plan configuration requirements, pause databases, or spin up dedicated warehouses. This eliminates or reduces the time spent on database administration, ETL management, and new schema modification.



## Data ingestion and extraction optimization

**Consider the right format** and compression techniques for your data. Formats like Parquet and ORC are ideal choices as they offer efficient compression and are optimized for analytical queries by storing data in columns instead of rows.

**Leverage native connectors** to simplify the data ingestion process and reduce maintenance efforts.

**Load raw data into the warehouse** and perform transformations within the warehouse itself to maximize flexibility and cost-effectiveness.

**Use real-time streaming** only when necessary for your specific use case, as batch processing can often be a more cost-effective solution.



## Storage optimization

**Achieve cost-effective storage** by implementing a tiered storage system that uses a combination of hot and cold storage options.

**Delete or archive unnecessary data** according to established data retention policies to further reduce storage costs.

**Use data compression techniques** to reduce your overall storage footprint.

**Partition tables** into smaller segments and organize data within partitions using clustering techniques to improve query performance.



## Compute optimization

**Analyze query patterns** to understand your workloads. This will help you identify areas for improvement.

**Right-size resources** by selecting appropriate compute resources for different workloads. Also, consider autoscaling to enhance performance.

**Use caching** (which stores frequently accessed results in memory) to improve query response times and reduce compute resource usage.

**Pre-compute and store results** of complex queries with materialized views for faster response times.

**Adhere to best practices** for writing efficient SQL queries to significantly contribute to overall optimization.



## FinOps

**Implement cost monitoring and alerts** to track spending and identify unexpected increases, allowing for proactive cost control.

**Use chargeback and showback mechanisms** to enable the allocation of data and analytics costs to departments based on their usage, promoting accountability and cost awareness.

**Visualize costs through a dashboard** to gain insights into usage patterns, helping to identify costly queries and facilitate optimization efforts.

# How the World Bank provides crucial information to global policymakers in seconds

As one of the world's largest sources of funding for developing countries and a leading research institution in development economics, the World Bank operates in 189 countries. Manually curating research evidence often takes weeks if not months, making it challenging to obtain accurate, data-validated information when it is needed. This delay can hinder the ability to make timely, well-informed policy decisions on how to optimally invest global development resources.

The World Bank has developed an open-source, queryable database of causal research from randomized controlled trials focused on policy interventions, utilizing generative AI to extract insights from development research literature. Accessible via an intuitive chat interface, this AI tool, called "ImpactAI," simplifies policymakers' access to quality information by helping them understand the impact of various interventions and identify those best suited to achieve their objectives. ImpactAI enables informed decision-making within seconds, speeding up the process by 90% compared to previous manual methods.

The World Bank's AI tool, ImpactAI, initially incorporating 5,000 development impact studies and growing by 5,000 studies annually, has an anticipated user base of over 100,000 annually. This tool is expected to dramatically reduce literature review time by more than 90%. With such broad reach and efficiency, it has the potential to significantly accelerate progress towards critical development goals globally.

[Watch the video](#)



# Inside our unique approach



## Unparalleled infrastructure, data, and AI integration

Extract maximum value from your data with unparalleled speed and efficiency. With Google's Data Cloud, data teams can take advantage of end-to-end integration from infrastructure (TPUs and GPUs) to data management (AlloyDB, Spanner and BigQuery) to AI (Gemini and Vertex AI). And with AI capabilities and assistive agents integrated into our platform, you can deliver new use cases faster—from building and deploying models using simple SQL to analyzing complex data using natural language.



## A unified data foundation

Google Cloud's unified data foundation is built on BigQuery and database solutions, integrating your structured and unstructured data to deliver insights and power AI-driven applications. This unified data foundation enables you to manage your entire data lifecycle—making data management, analysis and AI development easier for different types of users across your agency. It integrates end-to-end data governance with complete lineage tracking. That includes everything from data ingestion and storage to model training and deployment, so that you can have a single, governed, source of truth for your data, across transaction systems, analytics and AI.



## The most open data platform for modernization

Google Cloud is committed to being the most open cloud provider, letting you build modern, data-driven applications wherever your workloads are. We support open source and open standards, and offer managed database services, such as Cloud SQL for MySQL, PostgreSQL and SQL Server, that are fully compatible with popular open-source engines and AI models. With AlloyDB Omni and BigQuery Omni, you can utilize data and modernize your applications across Google Cloud, AWS, Azure, and Google Distributed Cloud, without incurring the costs, security risks, and governance concerns associated with data migration.



## Enterprise-grade AI ecosystem

Google Data Cloud is an industry leader in efficiency, security, and scale—enabling you to use advanced techniques for grounding your models in your agency data. Build, deploy, customize, and manage models at scale with retrieval augmented generation (RAG), automated prompt engineering, and fine-tuning. Query agency data using natural language with greater accuracy and nuance with semantic layer integration across services. Power your applications with capabilities like semantic search and anomaly detection with highly efficient vector search technology that Google users for services such as Google Search and YouTube Search.

# Make your data cloud AI-ready

Now you know that having unified data is the foundation to successful AI initiatives.

1 Google for Government (2025), [Scaling AI's Impact in State and Local Government](#)

2 Google for Government (2025), [Scaling AI's Impact in the Federal Government](#)

3 MIT Management (2021), [Tapping the power of unstructured data](#)

4 Recent Google Cloud survey of IT Decision Makers and Influencers from organizations 500+ with AI workload use or interest

5 Google Cloud (2024) Customer Intelligence Data & AI Trends Research.

6 Forbes (2023), [Five Factors For Planning A Data Governance Strategy](#)

## Are you ready to get started?

01

Looker is the complete AI for BI platform from Google Cloud that empowers everyone in your organization to easily access and explore data. Tap into the full value of your data by requesting a free trial of Looker.

[Request a free trial](#) →

02

Stay up to date with our latest innovations and updates by signing up for the Google Public Sector newsletter.

[Sign up](#) →

03

Learn more about how Google's Data and AI solutions can empower your agency and see examples of how we are helping accelerate mission impact with AI.

[Learn more](#) →

Google Cloud

