

Future-proofing retail banking:

How Spanner is enabling
next-generation
financial services

Author:

Szabolcs Rozsnyai

Contributor:

Mohsin Imam

White paper

October 2025




Contents

Introduction	3
Industry challenges	4
Core banking and transaction processing	5
Customer experience and digital banking	6
Fraud, security, and compliance	7
Modernization and Infrastructure	8
Spanner: a solution for modern banking challenges	9
Security, regulatory compliance and resilience	11
Availability, resilience and data sovereignty	12
Data protection: access controls and CMEK encryption with optional EKM through trusted partners	14
Concentration risk: mitigation through open standards	17
Full end-to-end audit trails with tamper-proof retention locking	18
Cost efficiency	20
Cost effective elasticity with constant latency	21
Tiered storage to optimize costs for hot and cold data	23
Cost-optimized incremental backups for DR	24
Managed service – low-ops	25
Evolving customer expectations for modern banking	26
Always-on availability for seamless customer experiences	27
Cross-channel consistency	28
Low-touch integration of data silos	30
AI-powered personalization	32
Use case example – digital online banking platform	34
Architecting a Spanner-powered online banking system of engagement	36
Multi-channel support	37
Operational apps	38
Core banking	41
Integration and distribution	42
Analytics and warehouse	43
Governance and tools	45
Real-time fraud and risk analytics	46
Use case example – CIAM (Customer Identity and Access Management)	47
The challenges	49
Architecture example	51

Introduction

The retail banking sector is navigating a period of transformation, driven by heightened digital competition, evolving customer expectations for seamless experiences, and increasing regulatory requirements. Traditional infrastructures often struggle to adapt, hampered by legacy systems, data silos, and the high costs associated with maintaining resilience and compliance. This whitepaper delves into these pressing industry challenges and presents Spanner as a transformative database solution uniquely positioned to address them. Readers will gain a comprehensive understanding of how Spanner's core capabilities – including global consistency, high availability, robust security, and cost-effective scalability – directly tackle critical pain points in core banking, customer experience, fraud prevention, and compliance.

By leveraging advanced capabilities like graph processing to understand complex customer relationships for hyper-personalization, integrated full-text search for intuitive information discovery, and built-in vector search to enable GenAI-driven experiences like semantic search and intelligent chatbots, Spanner provides the tools to move beyond foundational stability to power these next-generation interactions. Furthermore, this paper provides practical architectural blueprints and explores specific use cases, such as modern digital online banking platforms and Customer Identity and Access Management (CIAM), illustrating how Spanner can serve as the foundation for secure, resilient, and innovative retail banking systems.



Industry challenges

The retail banking industry today faces a unique combination of pressures: heightened competition, rising customer expectations, and increasingly stringent regulatory requirements.

Traditional banking institutions now find themselves competing not only with each other but also with agile, digital-native challenger banks from the FinTech sector for the same cohort of customers. The rapid growth of mobile and online banking, accelerated by the coronavirus pandemic (with a 72% rise in FinTech app usage in Europe alone¹), has fundamentally changed customer expectations. Customers demand fast, simple, and seamless digital experiences. This increased digital adoption and the increasing shift away from cash transactions present a significant opportunity for banks to grow and retain customers, and to leverage data for enhanced customer engagement and new revenue streams.

However, several challenges hinder financial institutions from capitalizing on these opportunities and effectively implementing key use cases:

Legacy systems and infrastructure:

The speed at which banks can adapt to changing market conditions and evolving consumer behavior is often slowed by legacy systems, costly data centers, and the expenses associated with maintaining a large geographic footprint. These systems can create data silos, hinder agility, and increase vulnerability to disruptions.

Security threats and data vulnerability:

The ever-present threat of cyberattacks and data breaches necessitates robust security measures to protect sensitive customer data across global operations. Failures in this area can have severe consequences, including financial penalties and reputational damage.

Cost inefficiency: Financial institutions are under constant pressure to optimize costs, improve ROI, and remain competitive. Legacy systems, outdated infrastructure, and inefficient processes often lead to high operational expenses, hindering profitability and the ability to invest in innovation.

Regulatory and compliance requirements:

The financial services sector is guided by a dynamic regulatory landscape that includes standards like DORA, KYC/AML, and GDPR. Organizations demonstrate adaptability by aligning their operations with these evolving requirements, which involves strategic investment in compliance functions. This regulatory environment shapes sophisticated approaches to data management and auditability, reinforces the adoption of strong security protocols, and encourages the development of resilient operational capabilities.

Need for operational resilience:

Regulations like DORA in the EU, frameworks mandated by the Reserve Bank of India (RBI) or FED, FFIEC, etc., emphasize the need for heightened IT resilience to withstand disruptions. Financial institutions require systems with rapid recovery capabilities and minimal service interruptions to ensure business continuity.

[1] Forbes, Coronavirus Drives 72% Rise In Use Of Fintech Apps

Core banking and transaction processing

A bank's fundamental operations rely on its core banking system and its methods for processing transactions and payments. Core banking involves managing customer accounts and essential financial services. Transaction and payment processing focuses on ensuring that money moves accurately and swiftly between accounts. These behind-the-scenes functions are critical for day-to-day banking. However, keeping these essential systems running effectively presents several major difficulties.

Banking function	Challenge and impact
Global data consistency and integrity	Inconsistencies lead to financial errors, regulatory violations, and loss of customer trust. Maintaining a single, accurate, and consistent view of financial data across all systems and locations is paramount, yet disparate systems and data silos often prevent a unified, reliable source of truth. This can result in conflicting information across different banking channels, making it difficult to reconcile accounts, generate accurate reports, and make informed business decisions. Impact: Financial losses, regulatory penalties, reputational damage, operational inefficiencies, and customer dissatisfaction.
Real-time operations and high availability	Downtime or delays are unacceptable. Customers and businesses demand instant access to information and services, expecting immediate transaction processing, up-to-the-minute balance inquiries, and continuous access to online banking platforms. Legacy systems, however, often struggle to provide the 24/7 availability and responsiveness that modern customers require, leading to frustration and lost opportunities. Impact: Lost revenue, customer churn, damaged reputation, and competitive disadvantage.
Scalability and performance under load	Traditional systems often hit scaling limits. Financial systems must handle massive and fluctuating transaction volumes, from daily peaks to unexpected surges during market events. When transaction volumes exceed the capacity of legacy systems, performance degrades, leading to slow processing times, declined transactions, and system outages, creating a poor customer experience and hindering business operations. Impact: Service outages, slow processing times, declined transactions, customer frustration, and lost business opportunities.
Transaction management complexity	Legacy systems struggle to manage this complexity efficiently. Building banking features like rollbacks, reversals, reconciliation, and settlements requires complex steps, especially in distributed systems, where transactions span multiple systems and databases. This complexity increases the risk of errors, makes it difficult to implement new features, and adds to the operational burden of maintaining these systems. Impact: Potential for financial errors, increased operational complexity, and difficulty in maintaining compliance.

Customer experience and digital banking

Banks are working to improve how customers interact with them using digital tools. This means offering easy, real-time account management and a good digital banking experience online and through apps. They also aim to provide personalized loyalty programs and offers across different channels, support simple digital wallets and prepaid accounts, and use connections with other banks (multi-bank and open banking) to help customers manage their finances better. These goals are all about making banking more customer-focused, connected, and digitally accessible. However, successfully adding these features brings significant challenges.

Banking function	Challenge and impact
Delivering a seamless, omnichannel experience	Legacy systems often operate in silos, hindering a unified customer view. Customers expect a consistent and personalized banking experience across all channels – mobile, online, and in-branch. However, when systems are siloed, customer data is fragmented, making it difficult to provide a consistent experience (e.g. ensuring users view the same data on mobile devices or web applications at any point in time) across different touchpoints, leading to disjointed interactions and a lack of personalized service. Impact: Customer frustration, lower engagement, and increased likelihood of switching to competitors.
Real-time personalization and engagement	Traditional systems often lack the agility for effective personalization due to lack of real-time support and unified data access. Customers expect personalized offers and support, tailored to their individual needs and preferences. Legacy systems, with their limited data integration and processing capabilities, struggle to analyze customer data in real-time and deliver relevant, timely offers, resulting in missed opportunities to enhance customer relationships and drive revenue. Impact: Missed opportunities for cross-selling and upselling, lower customer lifetime value, and reduced customer loyalty.
Supporting advanced digital features	This often overwhelms older system architectures. Implementing features like multi-currency accounts, instant payments, open banking, etc., requires flexible infrastructure that can handle new data types, integrate with external systems, and scale to support increased transaction volumes. Older system architectures are frequently unable to support these demands. Impact: Inability to innovate and offer competitive services, leading to customer attrition and loss of market share.

Fraud, security, and compliance

Banks focus strongly on protecting themselves and their customers through fraud detection and maintaining tight security. At the same time, they must follow strict industry regulations and be ready for regular audits to show they are meeting legal requirements. These tasks are essential for building trust and operating safely. However, effectively handling these security and compliance responsibilities creates its own set of significant difficulties.

Banking function	Challenge and impact
Data residency and sovereignty	This is complex to manage with traditional databases. Regulations often mandate data storage within specific geographic boundaries, requiring financial institutions to comply with data residency and sovereignty requirements. Traditional databases, with their limited flexibility in data placement, make it challenging and costly to ensure that data is stored and processed in accordance with these regulations. Impact: Potential for non-compliance with data residency laws, leading to fines and legal challenges.
Real-time fraud prevention and security	Legacy systems often lack the speed and data accessibility for effective fraud prevention. Financial institutions face constant threats from increasingly sophisticated cyberattacks and fraudulent activities. Legacy systems, with their limited processing power and data access restrictions, struggle to analyze transaction patterns in real-time and detect anomalies that may indicate fraud, leaving them vulnerable to financial losses and reputational damage. Impact: Financial losses, reputational damage, regulatory fines, and loss of customer trust.
Maintaining regulatory compliance and auditability	Legacy systems struggle with data silos and lack of transparency, making audits difficult. The financial industry is heavily regulated, with requirements for data retention, reporting, and audit trails. Legacy systems, with their data silos and lack of standardized data formats, make it difficult to gather the necessary information for regulatory reporting and audits, increasing the risk of non-compliance and the cost of compliance efforts. Impact: Significant fines, legal penalties, reputational damage, and operational overhead.


Modernization and infrastructure

Many banks are working to update their core technology systems. This often involves moving tasks off older, less efficient database technologies (legacy database offloading) and planning to eventually shut down costly mainframes (mainframe sunseting). The goal is to create a more modern, flexible, and cost-effective technology base for banking operations. However, carrying out these significant infrastructure changes comes with its own specific set of difficulties.

Banking function	Challenge and impact
Modernizing legacy systems and reducing costs	These systems hinder agility and increase costs. Banks need to replace/augment outdated, expensive, inflexible systems (e.g., mainframes, Oracle) to improve agility, reduce operational costs, and enable innovation. These systems often require specialized skills to maintain, are difficult to integrate with modern technologies, and consume significant resources, hindering the ability to adapt to changing business needs and compete effectively. Impact: High operational costs, limited agility, inability to innovate, and difficulty attracting talent.
Operational database consolidation	This leads to data inconsistencies and increased complexity. Many banks have a fragmented landscape of multiple databases, creating silos that hinder a unified view of customer data, increase complexity, and drive up operational costs. This fragmented landscape makes it difficult to gain a holistic understanding of customer relationships, optimize business processes, and implement data-driven initiatives. Impact: Increased management overhead, difficulty gaining a unified view of data, and higher costs.



These challenges highlight the critical need for financial institutions to modernize their infrastructure and embrace cloud technologies to meet customer expectations, ensure regulatory compliance, and maintain a competitive edge.



Spanner: a solution for modern banking challenges

Overcoming these multifaceted challenges requires a robust and modern database solution. Spanner is uniquely positioned to provide the capabilities financial institutions need to thrive in this evolving landscape.

Fully managed ++

Simplified administration, enterprise grade



Horizontally scalable

Distributed database, near unlimited scale



Relational semantics

ACID transactions, SQL, schemas



99.999% uptime SLA



No maintenance downtime



Online, unlimited scaling



Automatic sharding



Automatic failure recovery
RPO = 0, RTO = 0



Strong external consistency



Superior price-performance




Zero-touch global replication



Security and compliance


Spanner is a globally distributed, strongly consistent, and scalable database service developed by Google, designed to handle both relational and non-relational workloads. It combines the strengths of traditional relational databases with the scalability of non-relational systems.



Spanner's multi-model capabilities empower you to build intelligent, AI-enabled applications on top of your operational relational and NoSQL data by leveraging native Vertex AI integration, [Spanner Graph](#) for querying complex relationships, [vector search](#) for semantic search, built-in [full-text search](#) – all with "true ZeroETL" interoperability. This unified approach eliminates data silos, saves costs, reduces operational and security touchpoints, and ensures data consistency across all models.

The following sections detail how Spanner specifically addresses key challenge groups:

Challenge group	Spanner capability
Spanner capability Financial institutions operate in a complex and demanding regulatory landscape, where security breaches and compliance failures can have severe consequences.	• Availability, resilience and data sovereignty • Data protection: access controls and CMEK encryption with optional EKM through trusted partners • Concentration risk: mitigation through open standards • Full end-to-end audit trails with tamper-proof retention locking
Cost efficiency In an environment of increasing competition and evolving customer demands, financial institutions must balance the need to control costs with the imperative to deliver exceptional service and meet stringent regulatory obligations.	• Cost effective elasticity with constant latency • Tiered storage to optimize costs for hot and cold data • Cost-optimized incremental backups for DR • Managed service – low-ops
Evolving customer expectations for modern banking Customers today expect seamless, personalized, and real-time digital banking experiences. Spanner empowers financial institutions to meet and exceed these evolving expectations.	• Always-on availability for seamless customer experiences • Cross-channel consistency • Low-touch integration of data silos • AI-powered personalization



Security, regulatory compliance and resilience


Financial services organizations face significant challenges in managing security, regulatory compliance, and resilience where for instance outages, security breaches and in general compliance failures can have severe consequences.

Challenges	Potential impact	How Spanner can help
Regulatory pressure Financial services organizations are under increased scrutiny with evolving regulations (e.g. DORA, KYC/AML, GDPR, etc.), demanding constant adaptation and incurring substantial compliance costs.	- Escalating operational costs - Heightened legal and reputational risks - Inability to achieve regulatory resilience - Potential for significant financial penalties	99.999% availability SLA - Geo-redundancy with RPO=0 - Data sovereignty - Encryption by default - CMEK and EKM (e.g. Thales) - End-to-end audit trails - IAM and RBAC (FGAC) - Open standards – “Concentration risk”
Complex data management and auditability Maintaining consistent data across distributed systems while ensuring comprehensive audit trails for regulatory scrutiny presents significant operational challenges.		
Cyber threats and data vulnerability Present threat of cyberattacks and data breaches necessitates robust security measures to protect sensitive customer data across global operations, with severe consequences for failures.		
Operational resilience imperative (e.g. DORA) Regulations like DORA mandate heightened IT resilience to withstand disruptions, requiring systems with rapid recovery capabilities and minimized service interruptions.		
Legacy system limitations and amplified risks Legacy infrastructure hinders agility, creates data silos, and introduces single points of failure, exacerbating compliance difficulties and increasing vulnerability to disruptions.		

01 Availability, resilience and data sovereignty


Redundancy

Geo-redundancy, RPO-0, 99.999% availability



Data sovereignty

Geo-redundancy, RPO-0, 99.999% availability



Resilience in financial services is critical not only for maintaining operational stability but also for ensuring a positive customer experience. Customers expect continuous access to banking services, and any disruption can lead to loss of trust, churn, and potential financial harm. Regulations like DORA (Digital Operational Resilience Act) further emphasize the importance of resilience by mandating that financial institutions have the capability to withstand, recover from, and adapt to disruptions. This regulatory focus reflects the systemic risk that operational failures in financial services can pose, highlighting the need for robust systems that minimize downtime and ensure business continuity, thereby safeguarding both the institution and its customers.

Spanner provides robust options for geo-redundancy and data residency to meet the diverse needs of financial institutions.

Geo-redundancy for high availability

[Multi-Region Spanner instances](#) are deployed across multiple designated GCP regions to maximize availability. These instances typically include:

- two read/writes regions (of which one is designated the leader region)
- a witness region
- either standard or additional optional read-only regions.

A multi-region instance has at least 5 replicas of the database distributed across 3 or more regions providing 99.999% of availability.

Dual-region configuration for data residency

To help meet data residency requirements in the countries listed under [available dual region configurations](#), Spanner offers dual-region instances. These instances are deployed across two regions within a single country and replicate data across both zones and regions.

This configuration type:

- Serves reads from two regions in a single country
- Provides higher availability (99.999%) and SLAs than regional configurations
- Can help to meet data residency requirements.

A dual region configuration has six replicas across two regions (three in each region). In contrast to a multi-region configuration, in each region there are two read-write replicas and one witness replica. A minimum of two replicas in each region is required to form a quorum and as such this type of configuration provides up to 99.999% availability.


	Multi-region Spanner	Dual-region Spanner
Availability SLA	99.999%	99.999%
Cross-region RPO	0	0
Automatic and seamless failover	Yes	Yes
Protects against zonal outage	Yes	Yes
Protects against region outage	Yes	Yes
Helps to meet in-country data residency requirements	No	Yes (e.g. Germany, India, Japan)

Geographic distance

Multi-region and dual-region Spanner instances help to protect against both zonal and regional outages with a large enough geographic distance between resources to shield against for instance natural disaster events.

RPO-0 (zero data-loss)


Spanner's multi-region and dual-region configurations employ synchronous data replication across zones and regions, resulting in a Recovery Point Objective (RPO) of 0. This synchronous replication guarantees zero data loss in the event of regional outages and removes the burden of complex data reconciliation during disaster recovery.



02 Data protection: access controls and CMEK encryption with optional EKM through trusted partners

Data encryption is a fundamental requirement to protect sensitive information, maintain customer trust, and comply with regulatory mandates. Encryption safeguards data both at rest (when stored) and in transit (when being transmitted).

Spanner offers both [encryption-at-rest](#) and [encryption-in-transit by default](#) to secure data throughout its lifecycle. This encryption helps to protect data stored on disk, including data in tables, indexes, operational logs, metadata and more.



Cloud KMS

Cloud KMS is a cloud-hosted key management service that allows the management of cryptographic keys. Encryption keys can be used to protect sensitive data residing across GCP with customer managed encryption keys (CMEK). For compliance mandates requiring that keys and crypto operations be performed within a hardware environment, the Cloud KMS integration with [Cloud HSM](#) makes it simple to create a key protected by a FIPS 140-2 Level 3 device. Keys can be generated, used, rotated, and destroyed. KMS integrates with Cloud IAM and Cloud Audit Logging so that one can manage permissions on individual keys and monitor how these are used.

CMEK (Customer Managed Encryption Keys)

Spanner supports CMEK, giving organizations enhanced control over their encryption keys, and management of encryption keys using KMS. With CMEK, organizations can:

- Generate, use, rotate, and destroy encryption keys
- Gain more control and visibility over who has access to data
- Meet compliance and regulatory requirements.

EKM (External Key Management)


For organizations with the most stringent key management requirements, Spanner supports External Key Management (EKM). This option enables customers to use encryption keys managed by a third-party system outside of Google Cloud.

Trusted sovereign partner key management

Spanner's EKM integration allows for the use of partner-managed EKM through a trusted sovereign partner that manages the EKM system. With a partner-managed EKM, the partner creates and manages the keys that are used in Cloud EKM. The partner ensures that your EKM complies with sovereignty requirements. This capability is particularly important for financial institutions that need to adhere to strict data sovereignty and regulatory compliance policies, as it allows them to maintain control over their encryption keys with trusted sovereign partners.

Fine-grained access controls (FGAC)

In the face of complex organizational structures and increasingly stringent regulatory requirements, financial institutions need robust and granular control over access to their data. Spanner addresses this need by providing [fine-grained access control \(FGAC\)](#), which enhances security, simplifies compliance, and provides centralized management.



Enhanced security and control with IAM

Spanner integrates with [Identity and Access Management \(IAM\)](#), enabling administrators to precisely control who can perform actions on specific Spanner resources. IAM provides a centralized platform for managing cloud resource access, offering:

- **Centralized control:** IAM gives administrators a unified view of security policies across the entire organization, simplifying access management even in complex environments with numerous workgroups and projects.
- **Granular access control:** IAM allows for granting access to cloud resources at fine-grained levels, extending beyond simple project-level permissions. This enables organizations to adhere to the principle of least privilege, ensuring that users only have access to the specific resources they need to perform their job functions.
- **Built-in auditing:** IAM includes built-in auditing capabilities that ease compliance processes. A full audit trail history of permission authorization, removal, and delegation is automatically recorded and made available to administrators, providing transparency and accountability.

Combining IAM with Role-Based Access Control (RBAC)

Spanner further enhances security by combining the benefits of IAM with traditional SQL role-based access control (RBAC). This combination enables table and column-level protection, providing an additional layer of granularity and control over data access within the database itself.

By providing fine-grained access control, Spanner empowers financial institutions to enforce strict security policies, meet regulatory requirements related to access control (such as those related to RBAC), and maintain comprehensive audit trails for compliance purposes.

Standards and certifications


The following contains information about Google's certifications and compliance standards it satisfies as well as general information about certain region or sector-specific regulations.

More details can be found here:

[Cloud compliance and regulations resources](#)

Global	Americas		Europe, Middle East and Africa		Asia Pacific	
ISO27001	USA	1.31(c)-(d)	Europe	UK	Australia	Singapore
ISO 27017	HIPAA	FINRA Rule	GDPR	NCSC Cloud	Australian Policy	MTCS Tier 3
ISO 27018	HITrust	4511(c)	EU Model	Security	Principles	
SOC 1	FedRAMP	Canada	contract clauses	Principles	Australian	
SOC 2	FIPS 140-2	Personal	Privacy shield	NHS IG Toolkit	Prudential	
SOC 3	COPPA	Information and	Germany		Regulatory	
PCI DSS	FERPA	Electronic	BSI C5		Authority	
CSA STAR	NIST 800-53	Documents Act	South Africa		Standards	
MPAA	NIST 800-171	Argentina	POPI		IRAP	
Independent Security	Sarbanes Oxley	Personal Data	Spain		Japan	
Evaluators Audit	SEC Rule 17a-4(f)		Esquema Nacional		FISC	
	CFTC Rule		de Seguridad		My Number Act	

03 Concentration risk: mitigation through open standards



Financial institutions are increasingly concerned with concentration risk, which arises from over-reliance on a limited number of technology providers. This risk is driven by factors such as:

- The growing complexity of IT systems,
- The increasing dominance of cloud providers,
- And regulatory scrutiny focused on operational resilience.

Regulations like DORA specifically address concentration risk, pushing financial organizations to demonstrate the ability to avoid vendor lock-in and ensure business continuity even if a key provider fails or becomes unavailable.

Open standards play a crucial role in mitigating concentration risk by promoting interoperability and portability. When systems are built on open standards, it becomes easier to switch providers, migrate applications, and avoid being locked into proprietary technologies.

Spanner's PostgreSQL interface: Spanner addresses concentration risk by providing a PostgreSQL interface. This interface allows applications designed for PostgreSQL to run on Spanner with minimal changes.


Portability: The PostgreSQL interface provides access to the breadth of Spanner features, using schemas, queries, and clients that are compatible with open source PostgreSQL. This simplifies moving an application built on Spanner to another PostgreSQL environment. This portability provides deployment flexibility and supports disaster recovery scenarios, such as a stressed exit.

Reduced vendor lock-in: By supporting PostgreSQL – a widely adopted open standard – Spanner reduces vendor lock-in and provides financial institutions with greater flexibility and control over their future technology choices.

Simplified exit strategy: In the event of a need-to-migrate-away or stressed exit from Spanner, the PostgreSQL interface simplifies the process. The database schema and application code are largely compatible with PostgreSQL, meaning that the primary task becomes data migration. This significantly reduces the time, effort, and cost associated with switching providers, enabling organizations to demonstrate a viable exit strategy to regulators within a reasonable timeframe.

Read more on architecture and data migrations in:
[Failure scenarios and resiliency with Spanner](#)

04 Full end-to-end audit trails with tamper-proof retention locking



Auditing functionality is crucial for financial institutions, serving not only as a critical security measure but also as an essential component for meeting compliance regulations. Comprehensive audit trails provide transparency, accountability, and the ability to reconstruct past events, which is vital for detecting and investigating security incidents, ensuring data integrity, and demonstrating regulatory compliance.

Importance of auditing in finance

- **Security perspective:** Auditing helps in identifying unauthorized access, detecting malicious activity, and investigating security breaches. By logging user actions, data modifications, and system events, audit trails enable security teams to analyze patterns, identify anomalies, and respond effectively to security threats.

- **Compliance regulations:** Financial institutions operate under numerous regulations that mandate robust auditing capabilities. These regulations often require organizations to maintain detailed records of data access, transactions, and system changes for a specified period. Audit logs are essential for demonstrating compliance with these regulations, facilitating regulatory audits, and providing evidence of adherence to internal policies and procedures.

Spanner provides comprehensive auditing capabilities to meet the demands of the financial industry. [Cloud Logging](#) integrates with Spanner and keeps track of different types of audit logs:

- **Admin activity audit logs:** These logs record administrative operations that affect the configuration or metadata of a service. For example, creating, modifying, or deleting a Spanner instance.
- **Data access audit logs:** These logs record data access operations that read the content or metadata of data. For example, querying a Spanner database.

Archiving and analyzing audit logs


Cloud Logging offers flexibility in how audit logs are stored and managed. Audit logs can be stored in various destinations, including:

- [Cloud Storage \(GCS\):](#)
For long-term storage and cost-effective archiving of audit logs.
- [BigQuery:](#)
For powerful analysis and querying of audit logs using SQL. This enables security teams and auditors to efficiently analyze large volumes of log data to identify trends, detect anomalies, and perform forensic investigation.

[Tamper-proof retention locking](#)

To further enhance security and compliance, GCP supports tamper-proof retention locking of audit logs. Tamper-proof retention locking helps financial institutions meet regulatory requirements for data retention and immutability. By preventing unauthorized deletion or modification of audit logs, retention locks provide assurance that the logs are accurate, complete, and reliable for compliance audits and legal proceedings. This capability leverages object retention in Cloud Storage. Retention locks ensure that audit logs cannot be deleted or modified before a specified retention period, even by privileged users.

By providing full end-to-end audit trails with tamper-proof retention locking, Spanner enables financial institutions to strengthen their security posture, meet compliance obligations, and maintain the integrity and reliability of their audit data.



Cost efficiency

Financial services organizations face significant challenges in achieving cost efficiency while simultaneously meeting customer expectations and stringent regulatory requirements. This complex interplay demands solutions that can optimize resource utilization, reduce operational expenses, and enable institutions to deliver exceptional service without compromising compliance.

Challenges	Potential impact	How Spanner can help
Scalability and performance under load Financial systems must handle large and fluctuating transaction volumes. Service outages, slow processing times, declined transactions can lead to customer frustration and lost business opportunities.	High TCO and reduced ROI Customer churn Reputational damage Potential for significant financial penalties Competitive disadvantage	Unlimited scalability Cost-efficient elasticity No overprovisioning compute, storage Economical DR options No-Planned maintenance Low operations
“Always-on” availability Customers and businesses demand instant and 24/7 access to their finances and services. Downtime or delays are unacceptable. Regulatory bodies overwatch a minimum of level of service availabilit		
Data integrity is non-negotiable and quick recovery mandatory Data loss is not acceptable both from a regulatory as well as customer perspective. High data integrity paired with high availability requirements add operational burden and costs.		





05 Cost effective elasticity with constant latency

Financial institutions face the challenge of managing fluctuating workloads while maintaining performance and controlling costs. Traditional database systems often require over-provisioning to handle peak demand, leading to wasted resources and increased expenses during periods of low activity. Spanner addresses this challenge by providing cost-effective elasticity with consistent latency.

Impact of over-provisioning in finance

Financial systems experience significant variations in transaction volumes throughout the day, month, and year. For example, trading platforms might see surges during market open and close, while retail banks might experience increased activity build up organically throughout the day. To ensure consistent performance, many institutions over-provision their database infrastructure, resulting in:

- **Increased costs:** Paying for resources that are not fully utilized
- **Reduced ROI:** Lowering the return on investment in database infrastructure
- **Operational inefficiency:** Managing and maintaining over-provisioned systems.



Spanner's elasticity solution

Spanner is designed to scale elastically and efficiently, allowing financial institutions to:


- **Scale up and down:** Spanner can [automatically adjust resources](#) to match workload demands. This ensures optimal performance during peak times and reduces costs during periods of low activity.
- **Avoid over-provisioning:** Spanner's elastic scaling eliminates the need to over-provision database capacity, significantly reducing infrastructure costs.
- **Maintain constant latency:** Spanner's architecture ensures consistent latency even as the database scales. This is crucial for financial applications that require real-time responsiveness and cannot tolerate performance degradation under load.

Benefits for financial institutions

Spanner's cost-effective elasticity provides several key benefits for financial institutions:

- **Cost optimization:** By eliminating over-provisioning, Spanner helps institutions minimize infrastructure costs and maximize ROI.
- **Improved efficiency:** Spanner's automatic scaling reduces operational overhead and allows IT teams to focus on other strategic initiatives.
- **Enhanced performance (“scale insurance”):** Spanner's ability to maintain constant latency ensures a consistently high-quality user experience, even during peak demand.

AutoScaler



Capacity with auto-scaler

What we see with customers:

~30% - 40% savings on daily organic traffic shops


Serving spikes on inorganic traffic spikes

By providing cost-effective elasticity with constant latency, Spanner enables financial institutions to optimize their database infrastructure, reduce costs, and deliver high-performance applications that meet the demands of modern banking.

For a deeper dive into how Spanner minimizes the need for overprovisioning, refer to this article: [Farewell to overprovisioning: How to unlock cost-effective elasticity with Spanner.](#)

06 Tiered storage to optimize costs for hot and cold data

Traditional databases often require upfront provisioning of storage, leading to overprovisioning and unnecessary costs. Spanner addresses this by decoupling storage from compute, allowing for independent scaling of each layer. This means that organizations only pay for the storage they consume on Spanner, eliminating the need for upfront provisioning and reducing costs.



Furthermore, storing both frequently accessed (hot) and infrequently accessed (cold) data on the same tier can be inefficient. Spanner's tiered storage solution can seamlessly tier data on different storage classes, optimizing costs without sacrificing performance.

By separating hot and cold data, Spanner ensures that:

- Critical data is stored on high-performance SSDs,
- While less frequently accessed data is stored on more cost-effective HDDs.


This approach can significantly reduce storage costs while maintaining the performance and reliability required by financial institutions.

Key benefits of tiered storage

- **Significant cost reduction:** By automatically tiering data to the appropriate storage tier, Spanner can significantly reduce storage costs, especially for large, infrequently accessed datasets.
- **Simplified management:** Spanner's fully managed tiered storage eliminates the complexity of managing multiple storage tiers, reducing operational overhead.
- **Enhanced performance:** By separating hot and cold data, Spanner can optimize query performance and reduce latency for frequently accessed data.
- **Unified data access:** Spanner provides a unified interface for accessing both hot and cold data, simplifying application development and management.

07 Cost-optimized incremental backups for DR

Spanner offers efficient [incremental backups](#) to significantly reduce storage costs for disaster recovery (DR). By capturing only the changes since the previous backup, incremental backups minimize the amount of data stored, leading to lower storage costs.




By leveraging Spanner's incremental backup capabilities, financial institutions can optimize their DR strategy, reduce costs, and improve their overall resilience.

Enhancing disaster recovery

Spanner also offers the ability to copy backups to a different region, providing an additional layer of protection against disasters.

This feature can be particularly valuable for financial institutions that need to meet stringent business continuity and disaster recovery requirements. By storing backups in a separate region, organizations can minimize the impact of unlikely large-scale outages and ensure that they can quickly restore operations in the event of such a major disruption. This capability aligns with regulatory requirements that emphasize the importance of robust disaster recovery plans and the ability to recover critical systems and data in a timely manner.



08 Managed service – low-ops

Spanner is designed to significantly reduce operational overhead compared to traditional databases. As a fully managed database service, Spanner handles many of the complexities of database administration, allowing organizations to focus on their core business and realize cost savings through less operational overhead. To name a few, Spanner eliminates the need for:

- **Server maintenance:** No need to manage physical servers or virtual machines.
- **Hardware lifecycle planning:** Spanner handles the underlying infrastructure, eliminating the need for hardware upgrades and replacements.
- **Software patching and upgrades:** Spanner automatically applies updates and patches, ensuring optimal performance and security.
- **Storage and capacity management:** Spanner automatically scales storage capacity to meet your needs, eliminating the need for manual provisioning.
- **High availability setup and tuning:** Spanner handles high availability and automatic failover, ensuring continuous service availability.
- **Sharding and replication:** Spanner automatically handles sharding and replication, eliminating the complexity of managing distributed databases.

The screenshot shows the 'Configure compute capacity' dialog in the Google Cloud Spanner console. The dialog is titled 'Configure compute capacity' and includes a subtitle: 'Your compute capacity determines the amount of data throughput, queries per second (QPS), and storage limits in your instance. Affects billing.' The dialog is divided into several sections: 'Select unit', 'Choose a scaling mode', and 'Backups'. In the 'Select unit' section, 'Nodes' is selected with a radio button, and a note states 'For display purposes only. One node equals 1,000 processing units.' Below this, 'Processing units (PUs)' is also listed with a radio button and a help icon. In the 'Choose a scaling mode' section, 'Autoscaling' is selected with a radio button, and a note states 'Let Spanner automatically add and remove compute capacity.' Below this, 'Manual allocation' is also listed with a radio button. In the 'Backups' section, 'Enable default backup schedules' is checked with a checkbox, and a note states 'All new databases in this instance are created with a backup schedule that creates full backups every 24 hours. Each of these backups is retained for 7 days. You can edit or delete the default backup schedule once it is created.' At the bottom of the dialog, there are two input fields for 'Minimum' and 'Maximum' values, both set to '100' and labeled 'Nodes'.

By automating many routine tasks, Spanner significantly reduces operational overhead and lowers costs. Organizations no longer need to allocate resources to server maintenance, hardware lifecycle planning, software patching, storage management, high availability setup, or complex sharding and replication strategies. This reduction in operational complexity translates to significant cost savings, both in terms of labor and infrastructure expenses.

Evolving customer expectations for modern banking

The financial industry is undergoing a digital transformation, driven by evolving customer expectations and technological advancements. Customers today demand seamless, personalized, and real-time banking experiences, delivered across multiple channels. To meet these expectations, financial institutions must modernize their infrastructure and adopt innovative technologies. Spanner, with its wide range of capabilities and scalability, is well-positioned to help organizations navigate this evolving landscape.

Challenges	Potential Impact	How Spanner can help
Seamless omnichannel experience expected Customers expect a consistent and personalized banking experience across all channels. Legacy systems often operate in silos.	Inability to offer competing services Missed opportunities for up and cross-selling Customer attrition Loss of market share	Single source of truth Cross-channel consistency Integration with AI pipelines Feature velocity through simplification Real-time updates and engagement
Real-time personalization and engagement Customers expect relevant and personalized offers. Traditional systems often lack the agility for effective personalization due to lack of real-time support and unified data access.		
Advanced digital features Rich feature set expected to support e.g. instant payments (in EU mandated), digital wallets, P2P payments, virtual credit cards, multi-currency accounts. Velocity to increase features is hampered by older system architectures.		
Intelligent and personalized support Customers expect quick, accurate and personalized support. Integrating AI requires real-time access to consistent and up-to-date data.		

09 Always-on availability for seamless customer experiences

In today's digital age, customers expect seamless and uninterrupted access to financial services. Any downtime or service disruptions can lead to customer dissatisfaction, loss of revenue, reputational damage and in some cases regulatory scrutiny.

Scheduled downtimes might work for physical branches that adhere to opening hours. However, online channels need to work 24/7 to facilitate critical customer journeys like balance inquiries or issuing payment orders.

Similar to customer facing features the back office services and processes supporting these functions can not be taken offline.

- **Real-time notifications and alerts:** push notifications for transactions and fraud alerts.
- **Scheduled payments and recurring billing:** ensuring on-time execution of standing orders.
- **Overdraft protection and limit checks:** real-time overdraft limits and approvals.
- **Opening and onboarding:** managing KYC data and validation processes.

Therefore it is critical to architect systems that minimize both planned (updates, patches, schema changes) and unplanned downtimes (network issues, data center outages).


Spanner's high availability and resilience paired with no need for maintenance downtimes (incl. online schema changes) places it at the core of such architectures. This allows financial institutions uninterrupted service to their customers.



10 Cross-channel consistency

Customers expect a seamless and consistent experience across all channels, whether it's through a mobile app, web portal, or inside a physical branch (e.g. teller machine or customer agent terminal).

Spanner's strong consistency model and global distribution capabilities enable financial institutions to deliver such a real-time and unified customer experience.



Key benefits of Spanner for cross-channel consistency:

- **Single source of truth:** Spanner provides a single source of truth for all customer data, ensuring consistency across channels.
- **Real-time updates:** Changes made to data in Spanner are immediately reflected across all applications and channels.
- **Optional global distribution:** Spanner's global distribution allows low-latency access to data from anywhere in the world.


By providing a single source of truth for operational data with real-time access and strong consistency, Spanner unlocks a frictionless and simple way to build features and use cases that go beyond traditional retail banking core functionality such as:

- offering **digital wallet and prepaid accounts**, or
- extending and integrating banking features into marketing and retail (e-commerce) partners and platforms for **omni-channel loyalty and offers** use cases.


Spanner can set the foundation for of use case in that space as for example:

- **Peer-to-peer payments:** Customers can seamlessly send and receive money to friends and family, regardless of their location or preferred payment method.
- **QR code-based payments:** Enabling quick and convenient payments at physical stores and online merchants.
- **Real-time wallet top-ups and redemptions:** Customers can instantly top up their digital wallets and redeem funds for purchases.
- **Microtransactions and pay-per-use services:** Supporting a wide range of small-value transactions, such as subscription payments and in-app purchases.
- **Split payments and group transactions:** Facilitating shared expenses and group payments, enhancing the social aspect of financial transactions.
- **Real-time rewards and promotions:** Customers can instantly redeem rewards and view personalized offers based on their transaction history.
- **Location-based discounts and notifications:** Financial institutions can trigger targeted promotions and discounts based on customers' location, enhancing the in-branch and online shopping experience.
- **Instant issuance and redemption of gift cards and vouchers:** Customers can easily purchase and redeem gift cards and vouchers across various channels, boosting customer engagement and loyalty.

By providing a consistent and reliable data platform, Spanner empowers financial institutions to deliver exceptional customer experiences across all channels, driving customer satisfaction and loyalty.



11 Low-touch integration of data silos



Financial institutions often operate with numerous historically grown data silos, making it challenging to gain a comprehensive view of customer behavior and risk. Spanner's strong consistency and global distribution capabilities enable seamless integration of these data silos.

Empowering data-driven insights with DataBoost on operational data

[DataBoost](#) and [federated queries](#) are powerful tools that enable financial institutions to unlock the value of their operational data. By seamlessly integrating data from various sources, organizations can gain deeper insights, improve decision-making, and drive innovation.

DataBoost

DataBoost is a fully managed, serverless service that provides independent compute resources for supported Spanner workloads. Spanner's DataBoost, allows to execute analytics queries and data exports without impacting OLTP workloads allowing to tap into operational data managed in Spanner and if required enrich with existing data silos that would otherwise require complex and lagging data integration pipelines.

Key benefits of DataBoost include:

- **Workload isolation:** Ensures that analytics queries do not interfere with critical transactional operations.
- **Improved performance:** Accelerates query execution and reduces latency.
- **Cost optimization:** Avoids over-provisioning of Spanner instances and optimizes resource utilization.

These key capabilities enable a wide range of use cases such as for **fraud detection** or **analytics and reporting** on operational data, such as for instance:

- **Transaction monitoring:** Detecting anomalies based on real-time behavioral patterns.
- **Transaction risk scoring:** Flag suspicious transactions in real time.
- **Account takeover protection:** Identifying unusual access attempts and locking accounts.
- **Card chargeback processing:** Managing dispute resolution workflows.
- **KYC and AML compliance:** Storing flagged transactions and ensuring reporting compliance.


Spanner columnar engine: Spanner columnar engine lets you run analytics with significantly improved performance on the latest operational data by creating a columnar representation of your transactional row-based data and using vectorized execution engine to run analytical queries. Spanner columnar engine increases scan performance by up to 200 times, eliminating the need for ETL while maintaining strong consistency and having zero impact on your transactional system.

Workloads that would benefit from using columnar engine include the following:

- Operational reporting extracts up-to-the-second business intelligence from the latest operational data.
- Served analytics power dashboards and custom drill-downs with interactive latency.
- **Federated analytics** seamlessly combine data from Spanner and other sources in BigQuery.



12 AI-powered personalization



Having a database such as Spanner which truly enables a single source of truth for operational data can serve as a foundation to power machine learning life-cycles from training to operationalization. Spanner, in combination with [Vertex AI integration](#) or [LangChain](#), creates a powerful combination for building intelligent and personalized solutions. By leveraging the power of AI and machine learning, organizations can enhance customer satisfaction, reduce operational costs, and gain a competitive edge.

Spanner's real-time data capabilities, coupled with machine learning capabilities, enable the creation of, for instance, intelligent chatbots that can provide personalized and contextually relevant responses to inquiries both to the customer as well as back-office agents.

By leveraging customer data stored in Spanner, these solutions can:

- **Deliver tailored responses:** Understand customer needs and preferences to provide accurate and helpful information.
- **Maintain contextual awareness:** Remember past interactions and tailor future responses accordingly.
- **Proactively offer assistance:** Anticipate customer needs and offer proactive solutions.

LangChain, in conjunction with Spanner's [vector search capabilities](#), empowers financial institutions to implement advanced search functionalities. By embedding documents and queries into vector space, organizations can:

- **Semantic search:** Find relevant information based on the semantic meaning of queries, rather than exact keyword matches.
- **Personalized recommendations:** Recommend products and services based on customer preferences and past behavior.
- **Knowledge base search:** Quickly access and retrieve information from a vast knowledge base.

This enables a wide range of key use cases:

Grounded interactions

- **Personalized responses:** AI-powered chatbots can access real-time customer data from Spanner to provide tailored responses to inquiries.
- **Context-aware conversations:** Chatbots can maintain context across multiple interactions, leading to more natural and efficient conversations.

Proactive support and recommendations


- **Personalized recommendations:** AI can analyze customer behavior and preferences to suggest relevant financial products or services.
- **Proactive issue resolution:** AI-powered systems can identify potential issues and proactively reach out to customers with solutions.

Multi-channel support consistency

- **Seamless transitions:** Ensure a consistent customer experience across different channels (e.g., web, mobile, voice) by synchronizing information in real-time.



Use case example: digital online banking platform



Our financial customers decide to build modern online banking solutions with Spanner at its core for five reasons:



Unlimited scalability and cost efficient elasticity

- Provides them “scale insurance” for future growth scenarios
- Enables to elastically adjust resources to the given demand to cost-effectively meet peak usage patterns
- Customer experience considerations call for consistent and low latency response times
- Strongly consistent transactions are a hard requirement space across regions



Compliance and security

- Geo-distribution across data centers with RPO=0
- Encryption at rest and in-transit both with Google as well as customer managed encryption keys
- Comprehensive audit logging and rich access controls



Industry leading availability

- Provides solutions with high availability to power their applications (SLA 99.999%)
- Always on/zero downtime



Simplified operations


- Allows to reduce operations significantly in order to focus on resources on innovating (no need to deal with upgrades, patches, hardware, configurations etc.)



Predictability of costs and investments

- Enables to forecast and control expenses as Spanner scales linearly.

Architecting a spanner-powered online banking system of engagement



Building upon the previously established requirements for a modern online banking platform, this section delves into the key architectural components and a representative system design using Spanner. We will examine how a Spanner-powered "System of Engagement" effectively addresses critical needs such as scalability, strong consistency, robust security, and high availability.

The following discussion outlines the core building blocks of such a system, demonstrating how they interact to deliver a seamless and reliable user experience while adhering to stringent regulatory and operational demands. The example architecture illustrates how Spanner's capabilities are leveraged to create a resilient and future-proof foundation for online retail banking.


Multi-channel support

A wide range of channels need to be supported ranging from mobile, web, voice (inc. chatbots) and up to 3rd party integrations such as OpenBanking.

Key considerations are:

- Security (e.g. DDoS protection, encryption)
- Performance for customer experience (CDN, caching)
- Device support (e.g. Android, iOS)
- Analytics (web, customer pathways, performance metrics, messaging, notifications).

[Apigee API management](#) platform presents a powerful, globally distributed and managed turn-key solution for all internal and external APIs, with the option for hybrid runtimes in closed-network environments. The Apigee platform is a full-lifecycle API management platform, and can be automatically scaled from small to enterprise-wide large scale API deployments. It can model both centrally and locally managed API projects. Apigee specializes in every aspect of APIs, from technical security and CI/CD deployment standards, to productization and monetization of API assets in digital marketplaces.



Operational apps

Containerization and Kubernetes form the basis for service deployment, covering larger legacy services, microservices, and event-driven operational services.

Managed cloud services reduce operational effort, and accelerate the time-to-market and deployment pipelines for APIs, services and data, in a uniquely holistic approach.


Loosely coupled and event-driven architectures provide an elegant solution and many benefits, but can introduce several challenges. Specifically, they can cause data consistency challenges in complex update scenarios to keep different models in sync.

The requirements for an operational data layer for retail banking solutions must therefore:

- Be able to handle the transactions with high availability
- Have deterministic failure modes
- Include a mechanism to notify dependent services that a system write, that they were interested in, was made.

Databases that support ACID transactions are required, but making databases highly available, especially for writes, is difficult. Replication can cause inconsistencies in data and can add cost and complexity to your architecture.

Additionally, typical high-availability database configurations cannot handle failures across availability zones without incurring replication delays. As those delays increase, so does the probability that additional failures might lead to the loss of all of the writes that are in transit to the replica node in another availability zone. These additional failures mean that the writes aren't fully written to the replica prior to the failure in that zone.



Spanner as the central data management component is the choice to manage transactional workloads.

Spanner is at the core of powering the “operational apps” meaning handling transactions, payment order, powering views such as checking balances and managing the data for all related customer journeys on the frontend incl. its master data.

Spanner is a scalable, enterprise-grade, globally-distributed, and strongly consistent database service built for the cloud specifically to combine the benefits of relational database structure with non-relational horizontal scale. This combination delivers high-performance transactions and strong consistency across rows, regions, and continents with an industry-leading 99.999% availability SLA, no planned downtime, and enterprise-grade security.



Scale: Horizontally scalable across rows, regions, and continents, from 1 to hundreds or thousands of nodes.



Fully Managed: Ease of deployment at every scale and every stage. Synchronous replication and maintenance are automatic and built-in.



Relational Semantics: Schemas, ACID transactions, and SQL queries (ANSI 2011).



Multi-Language Support: Client libraries in C#, Go, Java, Node.js, PHP, Python, and Ruby. JDBC driver for connectivity with popular third-party tools.



Transactional Consistency: Purpose-built for external, strong, global transactional consistency.



Enterprise Grade Security: Data-layer encryption, IAM integration for access and controls, and audit logging.



Highly Available: Up to 99.999% availability.

Some customers choose to segment the data into different Spanner instances along the function they are serving, others segment by tenancy – such as by sub-brands, entities or by splitting systems of engagement and record.


Key design decisions need to align with the consideration that Spanner provides a highly available and multi-regional database powering the applications.

An application that is deployed into a single-region or availability zone would limit the overall availability of the solution. The database layer with Spanner provides up to 99.999% of availability, an application that is deployed into, for instance a zonal Kubernetes cluster would lower the overall availability of the end-to-end solution to 99.5%.

This is not just a platform or infrastructure consideration, but also an application design decision. Stateless applications that can scale out horizontally face the challenge of state management. This problem can also be solved with Spanner as it can be deployed in multi-region configuration eliminating complex statement management orchestrations across scalable stateless application services.

Google Kubernetes Engine (GKE) is a managed, production-ready environment for deploying containerized applications. Launched in 2015, Kubernetes Engine builds on Google's experience of running services like Gmail and YouTube in containers for over 12 years. GKE completely eliminates the need to install, manage, and operate Kubernetes clusters.

To complement GKE, specifically as an example for short running or for stateless services, Cloud Functions or Cloud Run are suitable additions to GKE as a viable option to provide such flexibility and range of required functionality.



Core banking


In some organizations the actual clearing of transactions is conducted in the bank's core banking system.

These core banking systems are oftentimes complex solutions with a large legacy such as mainframe powered applications. These platforms can't be quickly replatformed and hence often the design decision to split the platforms into:

- System of Engagement (SoE), and
- System of Record (SoR).



The SoE is powering the customer journeys and is powered by Spanner. The actual clearing of transactions is then handled by the SoR layer powered by a core banking system.



Integration and distribution

Spanner is a database primarily optimized to power OLTP workloads. This means that large scale analytical workloads, such as for instance performing reporting, fraud analytics, AML, KYC use cases are best deployed on [BigQuery](#) in Google Cloud's ecosystem.

Learned patterns or trained ML models can be operationalized directly in Spanner utilizing its [Vertex AI integration](#) capabilities.

For this purpose, data managed on the operational stack needs to be replicated to the analytics stack.

An analytics stack architecture must be able to ingest varying volumes of data from different sources whether it is a stream, batch, or a unified pipeline. Furthermore, a modern architecture should have a unified model of ingesting, processing and transforming data whether it is a batch or streaming workload.

[Cloud Pub/Sub](#) provides a simple and reliable staging location for event data on its journey towards processing, storage, and analysis. Pub/Sub can scale without provisioning, partitioning, or load isolation worries, and expand applications and pipelines to new regions simply with global topics. Pub/Sub includes end-to-end encryption, IAM, and audit logging, as well as NoOps, fully automated scaling and provisioning with virtually unlimited throughput. It also provides extreme data durability and availability with synchronous cross-zone replication, plus native client libraries in major languages and an open-service API.

Cloud Dataflow is a fully-managed service for transforming and enriching data in:

- stream (real time) and
- batch modes.

[Cloud Dataflow](#) has a serverless approach which removes operational overhead with performance, scaling, availability, security and compliance handled automatically so users can focus on developing pipelines instead of managing server clusters.

[Cloud Composer](#) is a fully managed workflow orchestration service that empowers to author, schedule, and monitor pipelines that span across clouds and on-premises data centers. Built on the popular Apache Airflow open source project and operated using the Python programming language, Cloud Composer is free from lock-in and easy to use. Cloud Composer pipelines are configured as directed acyclic graphs (DAGs) using Python, making it easy for users of any experience level to author and schedule a workflow. One-click deployment yields instant access to a rich library of connectors and multiple graphical representations of workflows in action, increasing pipeline reliability by making troubleshooting easy. Cloud Composer is deeply integrated within the Google Cloud Platform, giving users the ability to orchestrate their full pipeline.

Analytics and warehouse

The analytics stack offers online retail banking solutions the flexibility to capture all aspects of the business operations for analytical purposes to power use cases such as KYC, fraud, AML, credit scoring, auditing, data lineage, reconciliation and many more.

Over time, this data can accumulate into the petabytes, but with the separation of storage and compute, it's now more economical than ever to store all of this data.

After capturing and storing the data, a variety of processing techniques to extract insights from it can be applied. Data warehousing has been the standard approach to doing business analytics.

This approach requires fairly rigid schemas for well-understood types of data, such as orders, order details, and inventory. Analytics that are built solely on traditional data warehousing make it challenging to deal with data that doesn't conform to a well-defined schema, because that data is often discarded and lost.

Moving from data warehousing to the "store everything" approach of a data lake is useful only if it's still possible to extract insights from all of the data. Furthermore, this "store everything" approach facilitates both descriptive and predictive analytics – in contrast to the traditional data warehouse-only approach. Data scientists, engineers, and analysts often want to use the analytics tools of their choice to process and analyze data in the lake. In addition, the lake must support the ingestion of vast amounts of data from multiple data sources.

The analytics stack collects and stores online retail banking relevant data in a central repository to provide:

- Central access to data
- Advanced analytics
- Operational services
- Data science workbench.

Google follows six pillars – as critical analytics stack factors – powered by multiple products:



Easy
data migration



Scalable, strongly consistent
cloud storage



Data processing
on demand



Serverless
data storage



Innovative
AI and ML



Security and
governance at scale

[Cloud Storage](#) is well suited to serve as the central storage repository to ingest and store varying volumes of data, from different sources and various formats for many reasons. It will retain all data in its native format, support all data types and all users and adapt to changes easily.

[BigQuery](#), Google's serverless, highly scalable enterprise data warehouse, is designed to make analytics more productive with unmatched price-performance. Google does all resource provisioning behind the scenes, No need to manage upgrading, securing, or managing the infrastructure. BigQuery's high-speed streaming insertion API provides a powerful foundation for real-time analytics, making the latest business data immediately available for analysis. It supports a standard SQL dialect which is ANSI:2011 compliant, which thereby reduces the need for code rewrites.

[Data Catalog](#) is a fully managed and scalable metadata management service that empowers organizations to quickly discover, manage, and understand all their data in Google Cloud. It offers a simple and easy-to-use search interface for data discovery, a flexible and powerful cataloging system for capturing both technical and business metadata, and a strong security and compliance foundation with [Cloud Data Loss Prevention \(DLP\)](#) and [Cloud Identity and Access Management \(IAM\)](#) integrations.

[VertexAI](#) is a fully managed, unified machine learning (ML) platform. It allows developers and data scientists to build, deploy, and scale ML models, including those powered by generative AI. It offers tools ranging from data preparation and model training to deployment and monitoring. It also gives access to Google's advanced models like Gemini, enabling the creation of innovative AI applications. Spanner, in combination with [Vertex AI integration](#) or [LangChain](#), creates a powerful combination for building intelligent and personalized solutions. Learned patterns and or trained ML models can be operationalized directly in Spanner utilizing its [Vertex AI integration](#) capabilities.



Governance and tools

[Cloud Monitoring](#) aggregates metrics, logs, and events from infrastructure, giving developers and operators a rich set of observable signals that speed root-cause analysis and reduce mean time to resolution (MTTR).

[Data Loss Prevention API \(DLP\)](#) helps to better understand and manage sensitive data. It provides fast, scalable classification and redaction for sensitive data elements like credit card numbers, names, social security numbers. Cloud DLP classifies this data using more than 90 predefined detectors to identify patterns, formats, and checksums, and even understands contextual clues. DLP can optionally redact data,, using techniques like masking, secure hashing, tokenization, bucketing, and format-preserving encryption.

[Identity and Access Management \(IAM\)](#) lets administrators authorize who can take action on specific resources, giving full control and visibility to manage cloud resources centrally. For established enterprises with complex organizational structures, hundreds of workgroups, and potentially many more projects, Cloud IAM provides a unified view into security policy across the entire organization, with built-in auditing to ease compliance processes.

Cloud IAM enables to grant access to cloud resources at fine-grained levels, well beyond project-level access. A full audit trail history of permissions authorization, removal, and delegation gets surfaced automatically for admins.

[Cloud KMS](#) is a cloud-hosted key management service that allows the management of cryptographic keys. Encryption keys can be used to protect sensitive data residing across GCP with customer managed encryption keys (CMEK). For compliance mandates requiring that keys and crypto operations be performed within a hardware environment, the Cloud KMS integration with [Cloud HSM](#) makes it simple to create a key protected by a FIPS 140-2 Level 3 device. Keys can be generated, used, rotated, and destroyed. KMS integrates with Cloud IAM and Cloud Audit Logging so that one can manage permissions on individual keys and monitor how these are used.

Real-time fraud and risk analytics

For banks, the ability to analyze vast streams of transaction data in real time is no longer a luxury, but a core requirement for security and risk management. Traditional architectures, which rely on moving data via slow ETL (Extract, Transform, Load) pipelines from transactional databases to analytical warehouses, introduce a critical information lag. This latency means that fraud is often detected after the fact, and risk exposure is calculated on stale data, leaving the bank vulnerable.

Spanner's integrated [columnar engine](#) creates a new paradigm by enabling true Hybrid Transactional/Analytical Processing (HTAP). It allows banks to run complex, ad-hoc analytical queries directly on live, transactionally consistent data. This collapses the traditional OLTP and OLAP silos into a single, unified platform, turning the bank's operational data into a source of real-time intelligence.


This provides transformative benefits for modern banking:

- **Instant fraud detection:** By running analytical models directly against the stream of incoming transactions, fraudulent activity can be identified and stopped as it happens, not minutes or hours later, dramatically reducing financial losses and protecting customers.
- **Up-to-the-second risk analysis:** Risk management teams can get an immediate and accurate view of the bank's global exposure to a specific market event or counterparty. This enables more proactive and effective risk mitigation strategies based on the most current data available.
- **Real-time AI-powered personalization:** With direct analytical access to a customer's live transaction history and account state, banks can power sophisticated recommendation engines and AI models to deliver context-aware, personalized offers and advice at the precise moment of customer interaction.
- **Accelerated compliance and reporting:** When regulators require ad-hoc analysis, teams can query the single source of truth directly, eliminating data validation cycles and dramatically reducing the time required to generate accurate reports.

Use case example: CIAM (Customer Identity and Access Management)

Customer Identity and Access Management (CIAM) are core service components in modern architectures managing user identities and profile data enabling personalized experiences and efficiently managing digital identities. Many organizations design and develop CIAM solutions with Spanner at its core to custom tailor user experiences along customer journeys both in the consumer space and enterprise segments. These solutions call for modern, seamless, friction-free and personalized customer experiences across multi-channel (web, mobile, phone).

As an example in online banking and payments, CIAM data structures require managing customers, partners, accounts and product master data which might be centrally consolidated from federated identity systems in an event-driven architecture. This master data is then enriched in CIAM with organization specific domain entities and workflows such as catering for regulatory driven entitlements as part of the authorization scopes.




On one hand the access scopes and hierarchies need to cater for channel-specific customer experience related authorizations. These include simple things such as enabling/disabling certain online banking features, views or workflows. As an example, default permissions can be defined and applied to customer segments, e.g.:

- online banking users have by default access to mobile banking and balance views
- and for instance these users cohorts who have granted online banking authorization can change profile data such as their postal address.

But it can be also used centrally to manage access to specific products through certain channels in defined postal code areas.

These hierarchies and scopes extend to regulatory privilege consent workflows such as a natural person needs to consent and sign waivers to be able to trade derivatives if they activate a brokerage account. The workflow status needs to be tracked in the CIAM system and is often powered by 3rd party applications such as a document management or workflow system.



The challenges

CIAM requires secure, highly reliable and flexible solutions to support such workflows. Specifically in the finance sector compliance rules drive the requirement for such systems (e.g. entitlement management) as they need auditable authorization models to resources based on verification and consent processes. CIAM systems often need to be designed for multi-tenancy to be able to centrally manage and integrate acquisitions and subsidiaries.

The reasons why Spanner is attractive for this type of use case is due to these requirements:

High availability: CIAM is a core gatekeeper component to all other services. If this component is down no other dependent service can function.

Single (global) consistent view: A single consistent source of truth is desirable to store and manage all aspects of a CIAM component. This brings scalability and geo-distribution challenges with traditional database systems and would require complex sharding and replication topologies to be built and maintained. Reconciliation of dispersed silos or synchronization between different systems in case of an authorization change are not acceptable both due to potential lag and inconsistent states in between model changes.

Catering for global user base and compliance: If CIAM serves a global user base, a globally available datastore is required. Applications close to the customer need to serve data with relatively low latency. This requires the option to place consistent replicas close to the data consumer. At the same time regulatory bodies call for isolating customer data within jurisdictional regions.

Very high durability: User data, authorization rules, intermediate workflow states etc. need to have a very high durability not just from the sake of a good user experience (you don't lose data), but also in finance this is driven by regulatory frameworks. Many jurisdictions require data to be redundantly stored in geo-separated regions with RPO=0 in case of a disaster.

Scalable and cost efficient: CIAM datastores require scalability to be able to serve data with consistently low latency times in an economical way. This requires to efficiently deal (i.e. elastically scale up and down) with weekend, day/night patterns as well as handling peak events (launches, promotional events). CIAMs power a broad range of services. When new functions are connected, the database needs to scale with the additional demand.


In finance a CIAM platform might power initially web online banking. Later the platform might get extended to mobile, chatbots, integrate into call centers or extend to 3rd party open banking applications. Another example is also if these banking platforms plan for multi-tenants such as onboarding acquisitions or subsidiaries.

Audit: Changes to authorization rules, hierarchies and respectively authentication and authorization requests and responses need to be tracked. These include granted scopes, attempts, violations, redirect URIs. These logs can be used for debugging or fraud and abuse analytics purposes. In some industries this is also a regulatory requirement with immutable and tamper proof retention locks.

Security (CMEK, EKM): Some industries or specific features (e.g. handling financial transactions and payments) require a very high standard in terms of security. The baseline is that the database encrypts data at-rest and in-transit. Some organizations require in addition customer managed encryption keys (CMEK) or integration of external key management systems (EKM).

Architecture example

The architecture diagram illustrates CIAM services in the context of a multi-channel online banking application along with various related (micro) services.



The core application is deployed across two regions: europe-west1 (Belgium) and europe-west2 (London). In each region, there is a Kubernetes cluster (GKE) hosting various microservices or components, such as viewing a balance, or creating transactions.

Each region has a CIAM (Customer Identity and Access Management) component, which is one of the most critical components as all services depend on its availability.

Therefore the architecture employs a geo-partitioning strategy to manage data and services across different regions and cater for differently sized user populations.

- UK Partition in europe-west1 London, and
- EU Partition in eur3 (Belgium, Netherlands and witness in Finland).

The UK partition is a regional instance where the resources are spread across one data center in London adhering to DR regulatory requirements which are satisfied by the zone separation of that particular data center campus.

The EU partition is spread out in a multi-region deployment to provide geo-redundancy with RPO-0 across a larger geographical region.

Google Cloud

