Deloitte.

SECURITY-COMPLIANT GEMINI ENTERPRISE ON GOOGLE CLOUD

ACCESS CONTROLS LIST (ACL)-BASED GOOGLE GEMINI ENTERPRISE SECURITY

CONTENTS

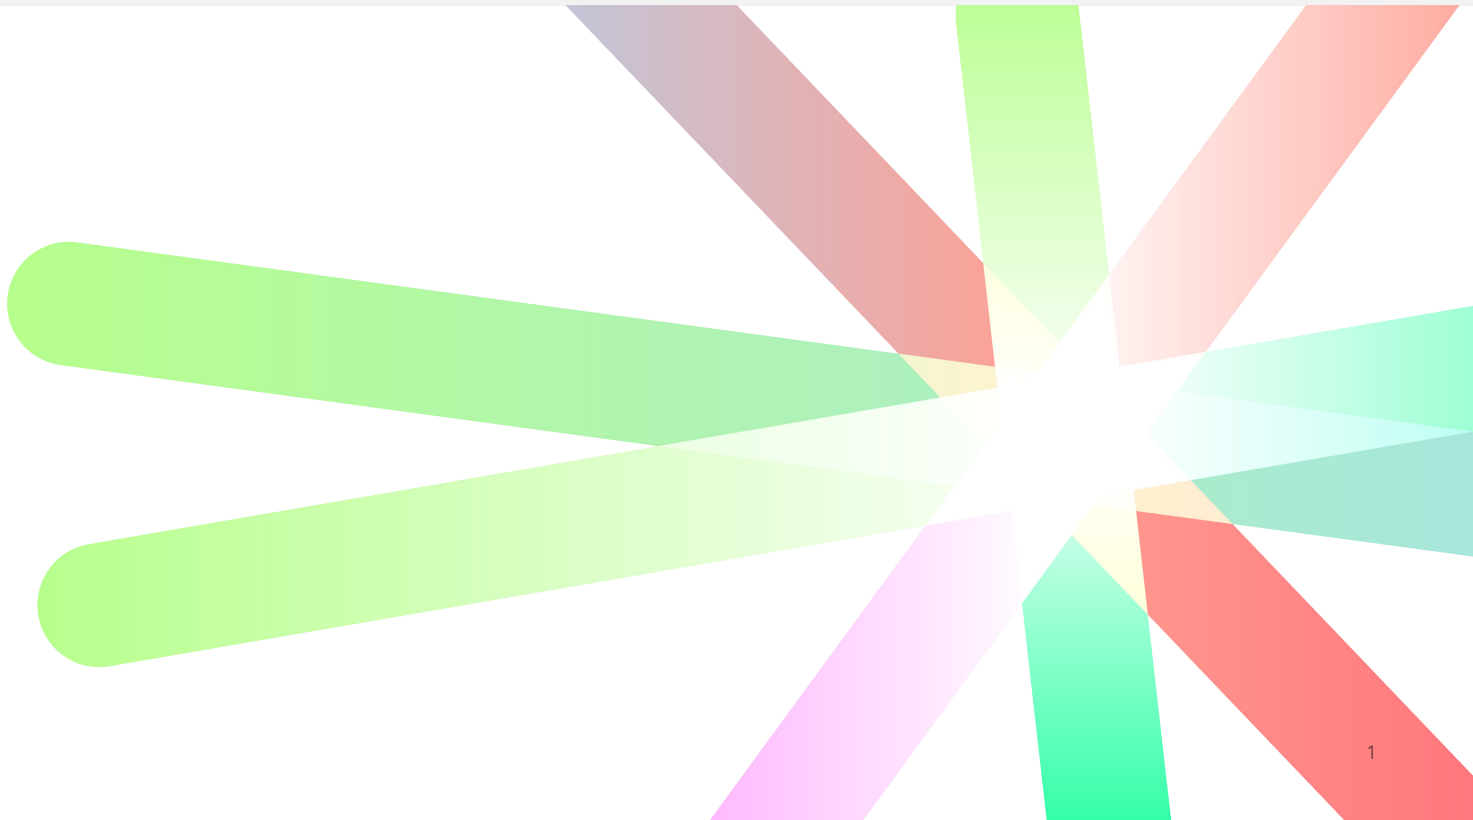

01.
EXECUTIVE SUMMARY

The life sciences industry operates at the intersection of scientific innovation, stringent regulatory oversight, and the fundamental responsibility to safeguard patient safety. Within Good Manufacturing Practice (GMP) and validated environments, digital systems must uphold uncompromising standards of data integrity, reproducibility, and auditability. Regulatory frameworks such as **ICH Q10[1] (Pharmaceutical Quality System)** and **ICH Q9[2] (Quality Risk Management)** establish principles for systematic, risk-based product lifecycle management. Similarly, **21 CFR Part 11[3] (FDA, 1997)** and **Annex 11[4] (EMA, 2011)** outline expectations for electronic records, signatures, and computerized system validation. Together, these standards mandate that any technological adoption in GMP-regulated operations provides traceability, reliability, and verifiable compliance.

Traditional automation and rule-based infrastructures, while effective for predictable workflows, lack the adaptability required to manage today's heterogeneous datasets and dynamic business processes. In contrast, the emerging paradigm of agentic AI introduces autonomous, context-aware systems capable of reasoning, learning, and executing tasks—all while embedding compliance requirements into their operational design. Unlike static digital systems, AI agents can integrate risk-based control strategies[2] (ICH Q9(R1), 2023), generate contemporaneous audit trails, and support adaptive validation approaches aligned with evolving regulatory expectations for artificial intelligence/machine learning (AI/ML) technologies (FDA's AI/ML Plan[5] (2021); EMA Reflection Paper on the Use of AI[6]).

Cloud-native infrastructures provide a foundation for realizing this paradigm shift. Leveraging Google Cloud Gemini Enterprise—formerly known as Google Agentspace—we designed, confirmed, and deployed an agentic AI solution within a GxP-regulated production environment.

Deloitte leveraged Google Cloud—with its integrated security, compliance, and audit capabilities—to design, confirm, and deploy a Gemini Enterprise solution within a GxP-regulated production environment. The platform's adherence to regulatory standards such as ISO/IEC 27001, HIPAA, and 21 CFR Part 11 facilitated end-to-end compliance, while its native auditability enabled continuous monitoring and traceability of system operations. Effective deployment in a live production setting demonstrated not only the scalability and broad nature of the solution but also its ability to help an organization maintain regulatory alignment under real-world operating conditions.

This adoption underscores the feasibility of embedding Gemini Enterprise within confirmed environments and highlights the potential for broader application across the life sciences value chain.

This solution demonstrates:

End-to-end compliance through alignment with ISO/IEC 27001, HIPAA (Health Insurance Portability and Accountability Act of 1996), and 21 CFR Part 11.

Continuous auditability and monitoring via integrated Google Cloud security and logging.

Comprehensive scalability for various use cases in live production environments.

A specific innovation in our implementation was the design of an ACL-driven data security model. This means that users only receive responses from documents they are explicitly authorized to access in the source system, maintaining compliance with regulatory requirements for controlled access and data integrity.

By integrating regulatory rigor with advanced cloud-native AI, our work illustrates a pathway for embedding compliance-by-design into enterprise AI adoption. As the scale and complexity of life sciences data grows, agentic AI on secure, compliant platforms like Google Cloud provides a framework for harmonizing innovation with oversight—ultimately helping to safeguard product quality and patient outcomes.

02.

DESIGN ARCHITECTURE, DATA SOURCES AND ACL

The foundation of our agentic AI deployment was a **hybrid integration architecture** designed to bring together both structured and unstructured data from cloud-native and on-premises systems. This architecture enabled not only efficient interoperability across platforms but also **strict preservation of security and compliance requirements** central to GxP-regulated operations.

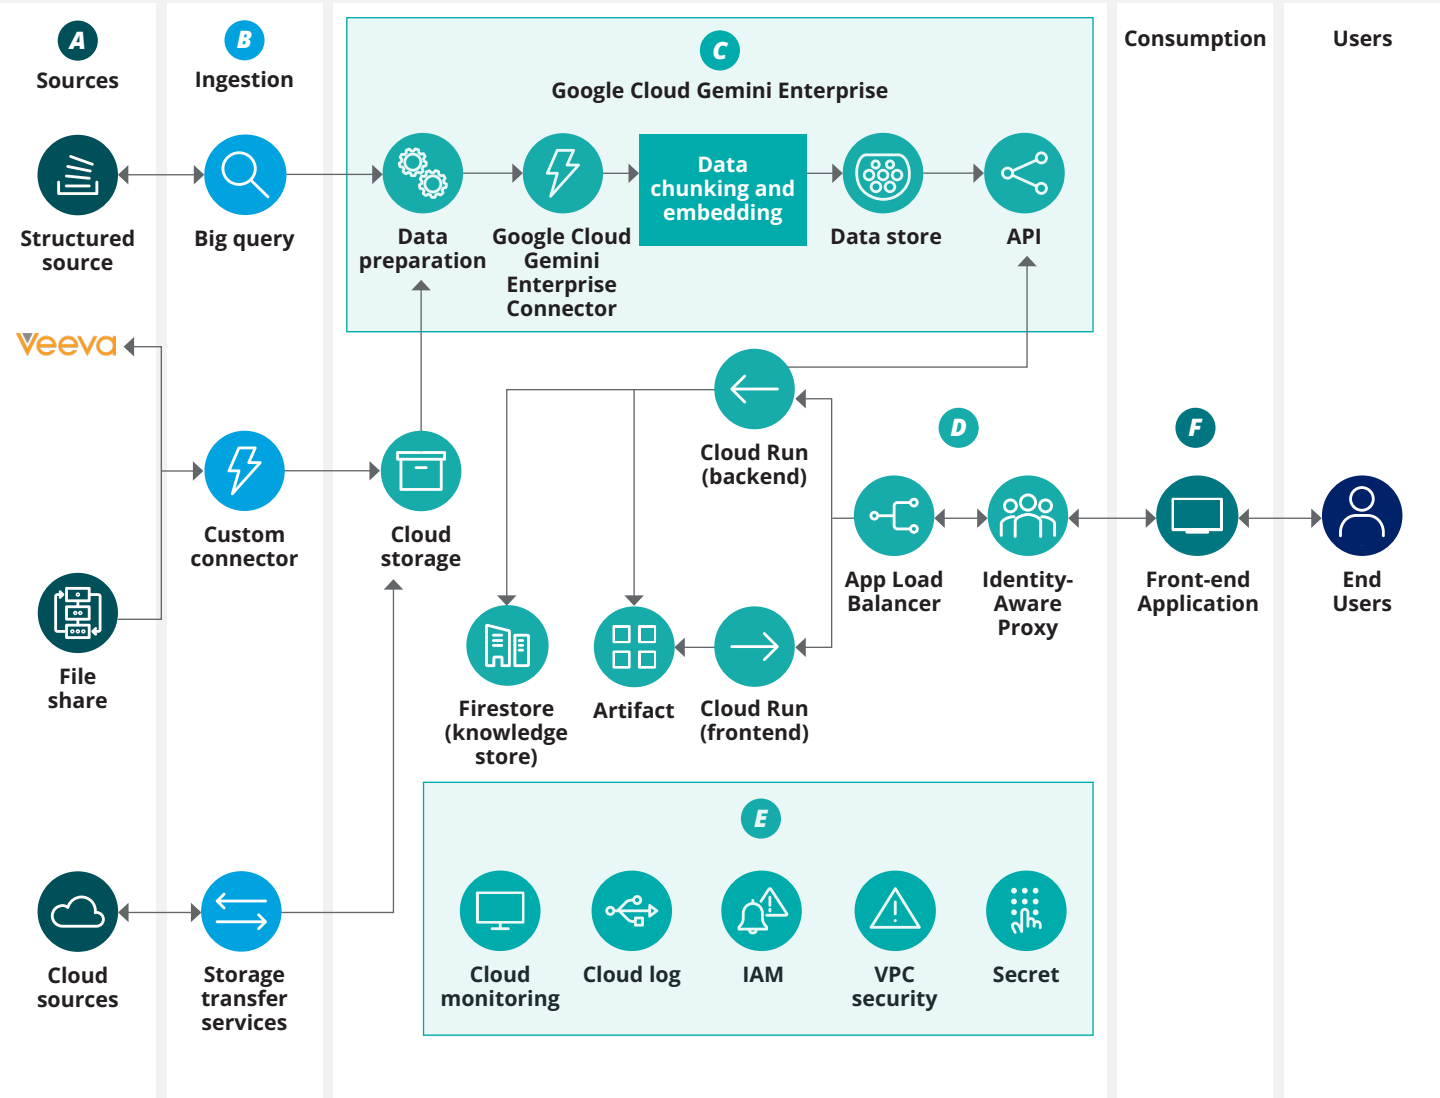

Figure 1: Hybrid data integration design for Gemini Enterprise deployment within Google Cloud Platform

In Figure 1, data is ingested, normalized, and routed securely across platforms without long-term storage, supporting scalability and compliance. ACLs from Veeva Vault QualityDocs (QDocs) are retrieved via application programming interface (API) calls and mirrored into Google Cloud's Identity and Access Management (IAM) system, enabling enterprise-grade security, role-based access, and regulatory alignment. The resulting harmonized data layer enables compliant, auditable, and scalable agentic workflows. Here are more details about different layers shown in Figure 1.

A 2.1. DATA SOURCES

Enterprise data originated from diverse systems such as Veeva Vault QDocs, File Share, structured repositories, and cloud object storage. Each source served as a system of record, maintaining not only the content but also the associated ACLs that dictate entitlements and permissions. Document classification and tagging policies were leveraged to confirm that data handling was aligned with enterprise security standards.

B 2.2. DATA INGESTION

Ingestion pipelines extracted content, metadata, and ACLs, then normalized them into JSONL batches for secure transfer into Google Cloud. This staging layer was intentionally **non-persistent**, preventing end-user access and reducing compliance risk. By embedding ACL data during ingestion, we created a strong foundation for preserving **source-driven security entitlements** across the lifecycle of the data.

C 2.3. GEMINI ENTERPRISE DATA PLANE

Within Gemini Enterprise, the ingested data was indexed with ACL metadata embedded at the document level. This enabled **real-time enforcement of entitlements**—when a user queries the system, results are dynamically filtered to make sure that only authorized documents are retrieved.

Two approaches to agent development were supported:

PREBUILT AGENTS

for rapid deployment of common use cases (e.g., Deep Research, NotebookLM Enterprise).

CUSTOM AGENTS

built using Agent Designer and powered by Gemini, capable of reasoning over ACL-enforced indexes.

D 2.4. INTERACTION LAYER

End-users interacted with the system via a thin, UI-centric frontend hosted on Cloud Run and secured with Identity-Aware Proxy (IAP). The frontend delegated operations to backend APIs, enabling identity context and entitlements to flow consistently from authentication through to query resolution.

E 2.5. APPLICATION AND ACCESS

Users could access agent functionality either through the out-of-the-box Gemini Enterprise web app or through custom enterprise applications. Both were secured via load balancing and identity controls, allowing only authorized cohorts to engage with the system.

F 2.6. SECURITY AND GOVERNANCE

A defense-in-depth model was applied throughout the pipeline:

IAM INTEGRATION

Enterprise identity provider (IdP) roles and group claims mapped directly into Google Cloud IAM.

SECRET MANAGEMENT

Tokens and keys secured in Secret Manager with rotation policies.

ENCRYPTION AND PERIMETER SECURITY

Customer-managed encryption keys (CMEK) and Virtual Private Cloud (VPC) Service Controls safeguarded data in motion and at rest.

AUDIT AND MONITORING

Immutable logging and real-time monitoring provided transparency, forensic readiness, and continuous compliance validation.

2.6.1. GEMINI ENTERPRISE MODEL ARMOR

GOVERNANCE TO GUARDRAILS: MODEL ARMOR-ENABLED ADVERSARIAL DEFENSE FOR ENTERPRISE GENERATIVE AI

The deployment of large language models (LLMs) in mission-critical, regulated environments—specifically, a commercial drug manufacturing solution leveraging multi-source data from **Veeva QDocs**, **SAP batch systems**, and structured **Google Cloud BigQuery** repositories—necessitates an uncompromised security posture against sophisticated adversarial attacks. The primary threat vector is the **Prompt Injection (PI)** attack, particularly the **Indirect PI** variant, where malicious instructions are covertly embedded within retrieved context data, thus compromising the integrity and confidentiality of manufacturing records and intellectual property.

> The integration of **Model Armor** on Google Cloud is posited as a fundamental, model-agnostic **Generative AI (GenAI) Security Gateway**.

This architectural placement facilitates a critical, real-time security inspection layer positioned between the client application and the Vertex AI LLM endpoint. The dual-phase operation of Model Armor—inspecting both the **input corpus** (user prompt, retrieved data context) and the **LLM output**—is paramount for securing the retrieval-augmented generation (RAG) workflow. Crucially, the input inspection provides a pre-emptive defense against Indirect PI by scanning the entire context, including unstructured text extracted from Veeva QDocs and structured data from SAP via BigQuery, before it is processed by the model.

SPECIFIC SECURITY POLICIES ENFORCED BY MODEL ARMOR

The defense strategy relies on granular, customized security policies to address the unique risk profile of the pharmaceutical manufacturing domain:

POLICY DOMAIN	MECHANISM AND RATIONALE	MATHEMATICAL AND LINGUISTIC FILTERING
 **PROMPT INJECTION MITIGATION**	**DIRECT/INDIRECT PI DETECTION** Utilizes advanced linguistic analysis to identify adversarial instructions and attempts to override the system prompt.	Filtering based on high cosine similarity to known **jailbreak prompts** and detection of meta-commands (e.g., [SYSTEM_OVERRIDE], IGNORE PREVIOUS INSTRUCTIONS).
 **SENSITIVE DATA PROTECTION (SDP)**	**CONFIDENTIALITY ASSURANCE** Integrates with Google Cloud DLP (Sensitive Data Protection) to redact proprietary information prior to LLM ingestion.	Custom InfoTypes defined for **proprietary chemical structures**, **Regulatory Submission IDs (e.g., NDA/BLA codes)**, and highly specific **Batch Release Criteria**. This prevents Indirect PI via data exfiltration or context manipulation involving classified text.
 **OUTPUT GUARDRAILS**	**DISCLOSURE PREVENTION** Scrutinizes the LLM's final response for inadvertent or malicious leakage of internal data or system prompts.	Blocking outputs that reveal **system prompt configuration**, internal API formats, or **specific sensor calibration constants** (time-series data) not intended for general user consumption.
 **HARMFUL CONTENT/ABUSE**	**RESPONSIBLE AI COMPLIANCE** Standard filtering for generating harmful, biased, or unregulated advice, especially crucial in a health/drug manufacturing context.	Strict adherence to Vertex AI Safety Settings policies (Harm Categories: Hate, Harassment, HCRT, Dangerous Content) with zero-tolerance thresholds.

The implementation of Model Armor within the Google Cloud framework, coupled with these specific, regulatory-aware policies, transforms the GenAI application from a security vulnerability into a securely governed component of the commercial drug manufacturing infrastructure. This architecture establishes a necessary boundary condition for safely leveraging LLMs with highly sensitive enterprise data.

2.7. WHY ACL MATTERS

A central design innovation was the mirroring of Veeva Vault ACLs into Google Cloud IAM. This meant that users accessing knowledge through Gemini Enterprise only received content they were entitled to in the original source system. By enforcing security at both the ingestion stage and the retrieval stage, the system guaranteed regulatory alignment and mitigated the risk of unauthorized disclosure.

This architecture illustrates how compliance-by-design principles can be embedded into AI adoption. By harmonizing data from multiple sources under a secure, transient, ACL-driven integration model, enterprises can accelerate innovation while preserving control, auditability, and patient safety.

03.

PLATFORM OVERVIEW: GEMINI ENTERPRISE

Google Cloud's Gemini Enterprise provides regulated enterprises with a secure, governed, and scalable environment to design, deploy, and operate agentic AI solutions. Unlike traditional AI platforms that prioritize flexibility over control, Gemini Enterprise is purpose-built for businesses where **data governance, compliance, and auditability are non-negotiable.**

3.1. WHAT GEMINI ENTERPRISE PROVIDES

AGENT DESIGNER

- Intuitive no-code/low-code interface for building and managing agents.
- Supports granular workflow orchestration, policy enforcement, and lifecycle management.
- Embeds compliance and governance requirements directly into agent design.

PREBUILT AGENTS

- Ready-to-deploy agents accelerate time-to-value by addressing common enterprise needs:
  - **Deep Research:** Context-aware synthesis of information across multiple sources.
  - **Idea Generation:** Facilitates collaborative innovation and brainstorming.
  - **NotebookLM Enterprise:** Secure, enterprise-ready adaptation of Google Cloud's NotebookLM for regulated contexts.

ENTERPRISE CONNECTORS

- Out-of-the-box integrations with Google Drive, CMS platforms, and more.
- APIs and SDKs for building custom connectors, enabling integration with proprietary or industry-specific systems.

GOVERNANCE AND ORCHESTRATION

- Centralized Agent Gallery for discovery, cataloging, and lifecycle oversight.
- Policy and asset discovery tools to map, classify, and control enterprise data sources.
- Threat monitoring and anomaly detection to facilitate secure agent activity.
- Built-in audit trail and reporting to support regulatory requirements.

SECURITY AND COMPLIANCE POSTURE

- Backed by Google Cloud's enterprise-grade security, including IAM, encryption at rest/in transit, and continuous compliance monitoring.
- Aligned with regulatory frameworks such as HIPAA, GDPR, SOC 2, and 21 CFR Part 11.
- Provides immutable logging for end-to-end traceability of agent actions and data flows.

3.2 POSITIONING GOOGLE CLOUD GEMINI ENTERPRISE TO AID IN COMMERCIAL DRUG MANUFACTURING

For our deployment, we leveraged the integration of Google Cloud's Gemini Enterprise ecosystem—as shown in Figure 2—with Vertex AI Agent Builder to provide a broad foundation for building and deploying AI agents within GxP-qualified pharmaceutical manufacturing environments.

This ecosystem specifically combines low-code tools for rapid prototyping with high-code development capabilities that empower data scientists and engineers to implement custom logic, integrate advanced ML models, and orchestrate complex workflows.

Central to this architecture is the open-source Agent Development Kit (ADK), which enables the construction of multi-agent systems capable of specialized functions such as predictive maintenance, yield optimization, real-time quality monitoring, and supply chain coordination, while maintaining alignment with regulatory frameworks such as FDA 21 CFR Part 11 and EMA Annex 11.

By leveraging audit-ready pipelines, traceable deployment workflows, and compliance-aware agent design, the platform supports the development of scalable, transparent, and adaptive AI agents that not only enhance manufacturing efficiency but also maintain strict adherence to validation and regulatory requirements. This approach advances the pharmaceutical industry toward Pharma 4.0, where autonomous, explainable, and GxP-compliant AI agents play a critical role in enabling adaptive control and decision-making across biopharmaceutical production pipelines.

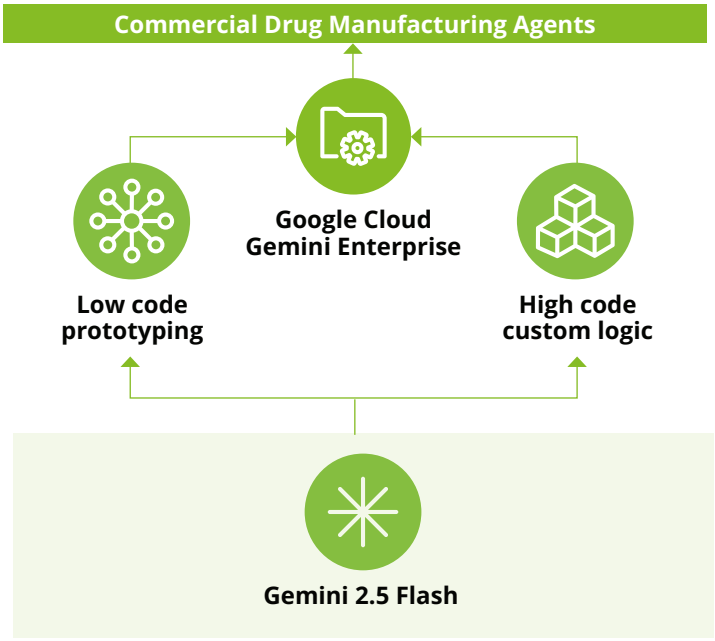

Figure 2: Conceptual architecture for Google Cloud's Gemini Enterprise ecosystem

Figure 2 showcases a conceptual architecture that illustrates the integration of Google Cloud's Gemini Enterprise ecosystem for deploying AI-driven agentic systems in commercial drug manufacturing. This diagram demonstrates how Gemini Enterprise can provide the overarching environment that enables both low-code and high-code development environments to drive efficiency in commercial drug manufacturing.

3.3 WHY GEMINI ENTERPRISE CAN BE BENEFICIAL FOR REGULATED ENTERPRISES

For industries like life sciences, financial services, and health care, adopting AI requires balancing **innovation with compliance**. Tools like Gemini Enterprise enable organizations to:

ACCELERATE AI ADOPTION
without building custom compliance frameworks from scratch.

EMBED ACL-DRIVEN DATA SECURITY
into agent workflows, so users only see content they are entitled to.

SCALE CONFIDENTLY
across business functions, knowing that governance, security, and auditability are built in.

By combining intuitive design, enterprise connectors, and regulatory alignment, Google Cloud Gemini Enterprise enables enterprises to move beyond experimentation and deploy **production-grade AI agents** that can withstand regulatory scrutiny.

04.

DATA-LEVEL SECURITY AND ACLs

Ensuring document-level security is foundational when deploying AI in regulated enterprises. For life sciences, health care, and other GxP-regulated industries, unauthorized disclosure of information is not just a technical risk—it is a regulatory violation with potential consequences for compliance and patient safety.

In our implementation, Google Cloud Gemini Enterprise was extended with an ACL-based data security model that mirrors the entitlements of the original source system. This embeds safeguards that a user accessing knowledge through an agent sees only the content they are authorized to see at the source.

4.1 IDENTITY AND ACCESS MANAGEMENT (IAM)

Google Cloud Gemini Enterprise leverages Google Cloud IAM to manage who can create, manage, and query datastores and applications.

> Principals (individuals, groups, or service accounts) are assigned roles that strictly control actions such as datastore creation, query execution, and application use.

For example, users must hold specific roles (e.g., *Vertex AI User, Discovery Engine Admin*) to build or query datastores. These permissions are provisioned via enterprise identity providers (SAML/OIDC) and enforced consistently across ingestion and retrieval.

4.2 ACCESS MODELS FOR KNOWLEDGE BASES

DOWNSTREAM-DRIVEN ACCESS (TRANSFORMED DATA)

SCENARIO

Structured sources undergo transformation in data lakes or ETL pipelines before ingestion.

MODEL

Access is governed by downstream roles and user personas, simplifying security into role-based groups (e.g., Gemini Enterprise Consumer).

BENEFIT

Transformation standardizes entitlements, making enforcement straightforward.

SOURCE-DRIVEN ACCESS (RAW/UNSTRUCTURED DATA)

SCENARIO

Unstructured sources such as regulatory PDFs, SOPs, or clinical reports retain their native ACLs from systems like Veeva Vault QDocs.

COMPLEXITY

Access depends on a combination of factors such as user role (author, reviewer, consumer), document classification (public, restricted, internal), and active directory group membership.

CHALLENGE

These layered entitlements need to be preserved at ingestion to prevent exposure of sensitive information.

EXAMPLE: SECURITY IN VEEVA VAULT QDOCS

Document-level security in Veeva Vault QDocs typically follows a three-tier model.

SECURITY LEVEL	SECURITY CONTROL	DESCRIPTION
Level 1	Group name	Defines user groups with document access
Level 2	Classification	Document sensitivity (internal, public, etc.)
Level 3	User role	Role-based visibility (author, reviewer, consumer)

In practice, this creates **composite entitlements** such as R&D-Internal-Consumer, blending group, classification, and role.

4.3 MAPPING ENTITLEMENTS TO GEMINI ENTERPRISE

To handle this complexity, we introduced a **lookup table mapping strategy**. This aligns source groups, roles, and classifications with corresponding Gemini Enterprise AD groups.

VEEVA AD GROUP	MAPPED AD GROUP	DESCRIPTION
R&D-INTERNAL-CONSUMER	Mapped-R&D-Internal-Consumer	 Holds all the Veeva AD group for R&D group, with doc classification as Internal and that has consumer access - as child/member group under Mapped-R&D-Internal-Consumer
R&D-INTERNAL-AUTHOR	Mapped-R&D-Internal-Consumer	
R&D-INTERNAL-PUBLISHER	Mapped-R&D-Internal-Consumer	
R&D-PUBLIC-CONSUMER	Mapped-R&D-Public-Consumer	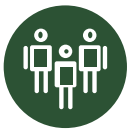 Holds all the Veeva AD group for R&D group, with doc classification as Public and that has consumer access - as child/member group under Mapped-R&D-Internal-Consumer
R&D-PUBLIC-AUTHOR	Mapped-R&D-Public-Consumer	
R&D-PUBLIC-PUBLISHER	Mapped-R&D-Public-Consumer	

This model enables **flattened ACLs** to be enforced at query time. Users querying Gemini Enterprise indexes receive only the documents their mapped group entitles them to—nothing more.

METADATA FOR UNSTRUCTURED DATA

To enforce these rules, each unstructured document ingested into Gemini Enterprise needs to include metadata such as:

- Unique ID
- Content URI (e.g., gs://bucket/path/file.pdf)
- MIME type
- ACL details (groups/users with access)
- Business/technical metadata in JSON format

Here is the sample JSON –

```
{
    "id": "<your-id>",
    "jsonData": "<JSON string>",
    "content": {
        "mimeType": "<application/pdf or text/html>",
        "uri": "gs://<your-gcs-bucket>/directory/filename.pdf"
    },
    "acl_info": {
        "readers": [
            {
                "principals": [
                    { "group_id": "group_1" },
                    { "user_id": "user_1" }
                ]
            }
        ]
    }
}
```

Note: Gemini Enterprise currently supports ingestion of common enterprise document formats (PDF, DOCX, PPTX, XLSX, HTML, XML, etc.), enabling broad coverage for regulated repositories.

4.4 KEY PRINCIPLE: SECURITY BY DESIGN

The overarching principle is simple yet critical:

> "Users see only what they are entitled to in the source system. Nothing more, nothing less."

By embedding ACL enforcement at both ingestion and retrieval, our approach transforms document-level security from an afterthought into a design cornerstone, aligning with regulatory expectations for controlled access, data integrity, and auditability.

05.
LESSONS LEARNED AND LEADING PRACTICES

Our deployment of Google Cloud Gemini Enterprise with ACL-driven data security surfaced several important lessons that can guide regulated enterprises on their own AI adoption journey. These insights are applicable across industries where compliance, governance, and data integrity are critical.

01 SECURITY MUST BE DESIGNED IN, NOT BOLTED ON

- Enforce **document-level ACLs from the beginning**. Retrofitting security into ingestion pipelines creates complexity and risk.
- Treat ACL flattening, validation, and mapping as **first-class architectural components**.

02 ALIGN EARLY WITH SOURCE SYSTEM OWNERS

- Security is only as strong as the **metadata and ACL fidelity** provided by the source system (e.g., Veeva Vault).
- Engage early with source system owners to confirm which business and access metadata can be shared. This avoids downstream rework and approval delays.

03 CROSS-FUNCTIONAL COLLABORATION IS ESSENTIAL

- Involve **security, identity management, and networking teams** early in the design phase.
- ACL mirroring requires coordination across enterprise directories, active directory groups, and cloud IAM mappings. Early alignment accelerates approvals and helps prevent architectural roadblocks.

04 CONFIRM WITH PROOF-OF-CONCEPT BEFORE SCALING

- When using **lookup tables or group-mapping strategies**, confirm them with a proof-of-concept (POC).
- POCs help uncover gaps in role definitions, metadata consistency, or system interoperability before committing to full-scale deployment.

05 BUILD FOR AUDIT READINESS

- Determine whether logging and monitoring are integrated with ACL enforcement. Every query, ingestion event, and document retrieval should leave an **immutable trail.**
- This not only supports compliance but also strengthens enterprise confidence in deploying AI at scale.

06.
TAKEAWAYS

The most important lesson is clear: **effective agentic AI adoption in regulated industries should involve security-by-design, not security-by-configuration**.

By embedding ACL enforcement into the ingestion, indexing, and retrieval layers, enterprises can move forward with innovation that does not come at the expense of compliance.

This approach lays a foundation for **scalable, auditable, and regulatorily aligned AI systems**—where users always see only what they are entitled to see, and enterprises can move forward with confidence.

07. REFERENCES

1. ICH. (2008). *ICH Q10: Pharmaceutical Quality System*. International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use.
2. ICH. (2023). *ICH Q9(R1): Quality Risk Management.* International Council for Harmonisation of Technical Requirements for Pharmaceuticals for Human Use.
3. U.S. Food and Drug Administration (FDA). (1997). *21 CFR Part 11: Electronic Records; Electronic Signatures.* U.S. Department of Health & Human Services.
4. European Medicines Agency (EMA). *(2011). Annex 11: Computerised Systems.* Retrieved from https://health.ec.europa.eu/system/files/2016-11/annex11_01-2011_en_0.pdf
5. U.S. Food and Drug Administration (FDA). (2021). *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device Action Plan.* Center for Devices and Radiological Health.
6. European Medicines Agency (EMA). (2021). *Reflection Paper on the Use of Artificial Intelligence (AI) in the Medicinal Product Lifecycle.*
7. Google Cloud. (2024). *Compliance Resource Center.* Retrieved from https://cloud.google.com/security/compliance
8. Google DeepMind. (2023). *Gemini: Google DeepMind's Multimodal Foundation Model.* Retrieved from https://deepmind.google
9. Google Cloud. (2024). *Prepare Data for AgentSpace Enterprise.* Retrieved from https://cloud.google.com/agentspace/agentspace-enterprise/docs/prepare-data
10. Google Cloud. (2024). *Identity and Access Control in AgentSpace Enterprise.* Retrieved from https://cloud.google.com/agentspace/agentspace-enterprise/docs/identity#acl-storage-unstructured
11. Google Cloud. (2024). *Google AgentSpace Overview.* Retrieved from https://cloud.google.com/products/agentspace
12. Promevo. (2024). *Agentspace vs Vertex AI vs Gemini.* Retrieved from https://promevo.com/blog/agentspace-vs-vertex-ai-vs-gemini

AUTHORS

Varun Kumar*
AI & Data Engineering
Offering and Specialist Master
Deloitte Consulting LLP
varunkumar6@deloitte.com

Kashinath Yadav*
AI & Data Engineering
Offering and Manager
Deloitte Consulting India Pvt Ltd
kasyadav@deloitte.com

Dalveer Rajput
AI & Data Engineering
Offering and Managing Director
Deloitte Consulting LLP
drajput@deloitte.com

Vikranth Gudala
AI & Data Engineering
Offering and Managing Director
Deloitte Consulting LLP
vigudala@deloitte.com

*Equal Contribution Authors