Informatica

White Paper

Informatica Intelligent Data Management Cloud

CLAIRE Security, Privacy and Compliance Overview

Where data & AI come to LIFE

Contents

Summary

This paper outlines Informatica's approach to responsible AI development, comprising AI safety, security, privacy and ethics for the **CLAIRE® AI engine** within the **Intelligent Data Management Cloud™ (IDMC)** platform. Key points covered in the paper include:

1. Development and design principles

- CLAIRE focuses on enhancing human productivity in data management.
- Informatica is committed to data security, transparency and responsible AI democratization.
- Informatica avoids AI applications that could cause harm or violate rights.

2. Metadata and data usage for CLAIRE training

- CLAIRE may use customer data and metadata for training machine learning models within the customer environment. Customer data is never used for training across different customer environments. However, subject to the customer's opt-out rights, customer metadata in anonymized form may be used to train models across customer environments, both now and in the future.
- CLAIRE implements both customer-specific and customer agnostic models.
- Training uses anonymized metadata, public datasets and Informatica documentation. Customers can opt out of metadata sharing for CLAIRE training at any time from the Informatica Administrator interface.

3. Geographical deployment model and inference

- Informatica's policy of using multiple locales ensures that customer metadata and data does not cross geographical boundaries for CLAIRE training.
- The Informatica IDMC platform complies with data residency and AI regulations across the globe, and only necessary metadata and data may be transmitted internationally for non-US customers at inference time, secured with encryption during transfer.

4. Opt-in, opt-out and disposition of data

- Customers can enable or disable the use of CLAIRE GPT and CLAIRE Copilots within an Informatica IDMC organization at any time.
- Customers can opt out of sharing metadata for training CLAIRE at any time.
- Informatica retains customer data for at least 30 days post-subscription termination, with deletion within 60 days.

5. Third-party software and models

- CLAIRE uses open-source and third-party proprietary large language models after thorough testing, fine-tuning and with Informatica-developed guardrails. Informatica routinely monitors known vulnerabilities, updates software, hardens code and follows its own security and governance policies.

6. Secure development lifecycle

- Informatica conducts multiple layers of security testing including automated scans and internal and third-party assessments.
- Internal threat modeling and security architecture reviews are integral to Informatica's software development lifecycle (SDLC) processes.

7. AI governance

- Informatica's AI Governance Council establishes ethical frameworks, policies and processes to guide AI initiatives. It helps ensure that CLAIRE systems operate reliably and safely while aligning with Informatica's responsible AI principles. Additionally, the council assists in applying responsible AI principles and best practices.

8. Compliance and certifications

- Informatica takes data security and compliance seriously and has achieved a number of compliance certifications. Details on relevant certifications can be found at **https://trust.informatica.com/compliance.html**.

This comprehensive approach demonstrates Informatica's commitment to secure, responsible and innovative AI-powered data management.

Introduction to CLAIRE

CLAIRE is a metadata-powered AI engine that is a key component of the Informatica Intelligent Data Management Cloud (IDMC) platform. It leverages artificial intelligence and machine learning to enhance and automate data management capabilities.

CLAIRE GPT is a generative AI (GenAI)-powered IDMC service that provides a natural language interface designed to help data analysts, data stewards, data scientists and data engineers discover data for analytics, explore metadata and assess data quality to determine whether the asset is fit for use, find data insights, explore master data management (MDM) business entities and automate ELT pipeline generation to build data products.

CLAIRE Copilots provide natural language based integrated development experience for IDMC applications designed to boost productivity. Copilot generates ingestion, replication, integration pipelines, business processes, business and technical summaries.

CLAIRE Security, Privacy, and Compliance Overview

Informatica takes the security, privacy, and compliance of CLAIRE seriously by carefully considering its approach to customer metadata and data, maintaining secure development practices, using third party AI models, training AI models and adhering to ethical and responsible AI principles.

This paper describes the types of metadata and data used by CLAIRE for training and during inference, the separation between customer-specific and agnostic models, data residency practices and options for customers to control their data usage. Additionally, this paper outlines Informatica's secure AI model development lifecycle, risk mitigation strategies and compliance certifications all aimed at ensuring the responsible and secure use of AI within IDMC.

By providing transparency into these practices, Informatica aims to build trust with our customers and demonstrate our commitment to data privacy, security, regulatory compliance and ethical AI development as we continue to push the frontiers of intelligent data management.

Development and Design Principles

Ethical and Responsible AI

Many of Informatica's products and services feature technology that enables our users to process information with increasing autonomy. At Informatica, we understand the profound impact of artificial intelligence (AI) that makes this automation possible, and we guide our AI development with an ethical, responsible and comprehensive set of principles.

These principles are designed to ensure that the AI technologies we create and deploy are developed and used in a way that respects human rights, contributes to societal benefits, upholds privacy and security and prioritizes transparency and explainability.

We aim to democratize AI, providing tools accessible to all users regardless of technical expertise.

Our commitment also extends to designing AI for deployment in ways that will not harm or undermine our values.

Under our standard contract terms, Informatica's intellectual property indemnity applies to our AI features.

CLAIRE Principles

Informatica's key principles when designing and developing CLAIRE applications and features are as follows:

- **Focus on enhancing human productivity in data management:** We aim to develop AI technologies for data management, making it easy for data teams and business users to manage their data effectively. By narrowing our focus, we aspire to deliver impactful solutions while tailoring our technologies to the unique needs and challenges within this area.
- **Ensure data security and accountability:** We pledge to create AI technology that prioritizes data privacy and security and balances them with the functionality of our product features. AI development oversight includes third-party audits, robust feedback mechanisms and a dedicated oversight team. We will maintain documentary evidence of how our AI is trained. This will ensure transparency in our processes and trust in our operations.
- **Provide transparency and explainability:** We aim to create AI models that are effective and understandable. We leverage advanced explainability frameworks and tools to provide insights into how our models make decisions, allowing users to understand how our AI application reaches conclusions.
- **Design delightful user experiences:** We aim to harness AI to augment human productivity by crafting thoughtfully designed user experiences that delight end users.
- **Democratize AI responsibly:** We are committed to making AI accessible to a broad range of users while maintaining a strong focus on ethical and privacy considerations. We balance openness with robust control mechanisms designed to prevent technology misuse and protect data privacy.

Risk Mitigation

Informatica acknowledges the potential risks and ethical issues associated with AI applications. Accordingly, Informatica will not develop or deploy AI:

- In ways designed to cause or create an undue risk of harm to individuals or society.
- For purposes that are prohibited in the major jurisdictions in which we do business or otherwise constitute a clear threat to the safety, livelihoods and rights of people.

Informatica is committed to using AI to make the world a better place, without harming or causing disadvantage to any individual or group. We believe in using AI responsibly, and we are dedicated to following these principles as we develop and deploy AI technologies. We understand that the field of AI is rapidly evolving. Thus, we will reassess and update these principles to keep pace with technological advancements and emerging ethical considerations. We firmly believe that by adhering to these principles, we can drive progress while ensuring the responsible use of AI.

Transparency and Ethical Considerations in CLAIRE

To better incorporate trustworthiness considerations into the design, development, use and evaluation of AI products, services and systems, Informatica is adopting the NIST AI Risk Management Framework in its AI software development lifecycle and risk management practices.

Being an enterprise-grade AI system, CLAIRE puts transparency and ethical considerations on its top priority list. Informatica ensures non-maleficence and beneficence practices are followed throughout SDLC, model training and operations of CLAIRE.

CLAIRE security assessments cover key risk categories that include, but are not limited to, data breaches, model manipulation and adversarial attacks. Risk assessments are performed by reputable internal and third-party assessment providers.

By design, CLAIRE has complete traceability and transparency for data sources and collection methods. The training dataset is reviewed periodically for AI fairness and any such findings are remediated in accordance with internal service-level agreements based on severity. The dataset doesn't identify subgroups, such as age or gender. Informatica performs ongoing monitoring of the responses generated from its system for integrity.

With user consent-driven minimal data collection, encryption, bias detection, fairness audits, access controls, continuous monitoring and responsible AI practices, Informatica takes steps to ensure security, fairness and transparency.

Potential abuse detection and remediation is a continuous process in CLAIRE. As a part of prerelease assessments, data privacy review and threat modeling are performed. All end-user prompts pass through a guardrail in the CLAIRE architecture. This layer tracks and differentiates potential abuse attempts from genuine requests in the data management context.

CLAIRE conducts ongoing monitoring and learns from explicit user feedback to protect against misuse in place.

Informatica uses the "Principle of Least Privilege". Informatica employees and applications only have access to the data and operations they need to perform their jobs. Only designated personnel in the production operations team have access to the models in production.

Calls to the model can originate only from the designated production product clusters, which are, in turn, controlled by the production operations team of the specific product.

Customer questions and answers are stored in the database in an encrypted format and aren't present anywhere in plain text.

[1] https://www.nist.gov/itl/ai-risk-management-framework

Metadata and Data Usage for CLAIRE Training

Informatica might use customer metadata for training CLAIRE customer agnostic models. Informatica might also use customer data for training customer specific models as described in detail above. As part of IDMC operations, Informatica might store customer metadata and data for providing data management services.

Metadata and Data Stored in Intelligent Data Management Cloud™

Informatica Intelligent Data Management Cloud™ (IDMC) stores different types of information:

Metadata

IDMC stores the following types of metadata:

- **Operational and usage metadata:** This includes information extracted from service and activity logs, such as service events (e.g., starting and shutting down of services), service telemetry (e.g., resource usage) and user activity events (e.g., logins and logouts to IDMC services). It also includes information about cloud service usage events, such as user activities like the creation and deletion of data collection in the customer's data marketplace.
- **Technical metadata:** This includes data schemas, rules, data profile statistics and design metadata that define integration tasks and processes, such as data sync, data replication, mappings and templates, task flows, process definitions and data lineage.
- **Business metadata:** This includes information related to customer data, including data classifications and glossaries designated by the customer, as well as crowdsourced information like data product ratings and comments.

Metadata collection by IDMC is necessary to provide data management services on the cloud and cannot be disabled. However, customers can select to opt out of having their metadata used for cross-tenant CLAIRE training from the Informatica Administrator interface at any time.

Data

Depending on how IDMC features are configured and enabled, customers may need to store partial or full business data records from applications and endpoints accessed through IDMC.

However, CLAIRE does not use business data stored in applications for customer-agnostic AI or large language model (LLM) training. For more information, see the section "Assets Used to Train CLAIRE Models," below.

Encryption for Data and Metadata at Rest

Any data or metadata persisted in the IDMC multi-tenant data repository is encrypted using the AES encryption algorithm with a 256-bit key. By default, the key is rotated once a year, but Informatica administrators can choose to rotate it every 90, 120 or 180 days.

Customer Managed Encryption Keys

IDMC supports customer-managed keys. While Informatica uses strong encryption practices, customers can utilize their own encryption keys to safeguard data. A customer can control their encryption keys and maintain the authority to encrypt and decrypt their data within the Informatica cloud environment. This gives customers confidence that their data remains confidential and protected, even from Informatica as the service provider. For more information, see "**Customer managed encryption keys**" on the Informatica Documentation Portal.

Customer-Specific vs. Customer-Agnostic CLAIRE Models

IDMC implements various AI models with different scopes to deliver CLAIRE functionality:

- **Customer-specific models:** CLAIRE models are tailored to meet the unique needs of individual customers. They are trained using data and metadata specific to each customer, allowing Informatica to provide services customized for that particular customer. These models are not shared with any other customers.
- **Customer-agnostic models:** CLAIRE models serve all customers. They are trained on Informatica product documentation, public data sets and anonymized metadata with customer consent.

Assets Used to Train CLAIRE Models

CLAIRE customer-agnostic models are trained on metadata, including technical metadata, operational and usage metadata and customer business metadata.

This metadata is anonymized before being used for training, ensuring that no sensitive information is exposed during the process.

Informatica never uses customer data stored in IDMC for training CLAIRE customer-agnostic models. This approach ensures that the customer's data remains private and secure.

Informatica does not access customer data endpoints (such as Salesforce, SAP or Snowflake) or other cloud applications, databases or data warehouses for CLAIRE or LLM training.

For customer-specific models, data from a customer's tenant can be used to fine-tune a model within the same tenant. For example, in Informatica Master Data Management (MDM), fine-tuning of match-and-merge customer-specific models is explicitly performed by the customer to de-duplicate MDM business entities. This data is not used to cross-train CLAIRE models and is only utilized by customer-specific CLAIRE models.

Informatica does not use prompt-response metadata or data for training CLAIRE models. However, Informatica might use feedback provided by users for responses to the prompt conversation, such as thumbs up/down or comments, for further model training, provided the customer has not opted out of allowing CLAIRE to use customer metadata for training.

Finally, other sources of public data are used to train Informatica customer-agnostic large language models, such as:

- General world knowledge: wiki data, books, public articles
- Verticalized domain knowledge: industry terms, industry metrics, industry systems
- Informatica documentation and knowledge base articles

The following table shows which assets are used to train CLAIRE:

Asset Type	Can Customers Opt-Out of Training Customer-Agnostic AI Models	Used in Training Customer-Agnostic AI Models	Used in Training Customer-Specific Models
Technical metadata (anonymized)	Yes	Yes	Yes
Business metadata (anonymized)	Yes	Yes	Yes
Customer data stored on IDMC	N/A*	No	Yes (User explicit)
Customer data on applications	N/A*	No	No
Feedback on the prompt responses on CLAIRE AI	Yes	Yes	Yes
Operational and usage metadata (anonymized)	No	Yes	Yes

* Customer data is never used for training customer-agnostic AI models.

CLAIRE Model Fine-Tuning

Customers can refine, train and fine-tune selected CLAIRE customer-specific models on their own data. The data and the refined models are unique and are not shared with other customers. They are stored within the context of the encrypted customer tenant and are available only to that specific customer tenant where the model was trained.

CLAIRE Model Evaluation

Informatica uses human-curated validation datasets to evaluate models for coherence, consistency, accuracy and completeness. Informatica scores test prompts based on these criteria and benchmarks them against established thresholds. Informatica also evaluates several natural language variations of prompts to ensure the model is generalized for language understanding.

Results are evaluated by engineering, including development and quality engineering (QE), as well as through an automated (LLM-assisted and LLM-as-a-judge) evaluation process in different phases of the model lifecycle.

Informatica also collects user feedback in production where users can signal if the generated response meets their expectations.

Informatica's model evaluation frameworks include, but are not limited to, the following KPIs:

- **Model development:** Accuracy, completeness, coherence, faithfulness, precision, recall, relevancy, F1-score, BLEU and ROUGE
- **Model serving and operation:** Inference time, scalability and time to first token
- **Non-functional based on business operation needs:** Acceptance rate, click-through rate and satisfaction rate

These KPIs are selected based on the technical problems being solved. For example, for classification tasks engineering uses accuracy, precision, recall, F1-score. For natural language tasks such as summarization, they use relevancy, accuracy, completeness, coherence and BLEU and ROUGE.

CLAIRE Model Usage and Scaling

Models in production are operated under an observable system to check for the following kinds of outcomes:

- Infrastructure aspects such as auto-scaling, metrics and availability
- Functional aspects such as user feedback, intent distribution, usage stats and model errors, which are captured via telemetry

When the application starts to scale, risks might arise related to the scalability of models that it uses. LLMs can have scalability challenges associated with functional capabilities, such as accuracy and hallucinations, as well as non-functional challenges like memory usage, inference latency and computational cost.

CLAIRE is not designed for high-risk use cases. Examples of high-risk use cases include:

- Real-time and post-biometric identification
- Addressing safety components of products, including the management and operation of road traffic and the supply of water, gas, heat and electricity
- Determining access to education and vocational training or assessing students in the education or training
- Recruiting, selecting, screening, filtering or evaluating employment candidates, or making decisions on promotion, termination, allocating tasks or evaluating performance
- Evaluating eligibility for public benefits and services or dispatch first-response services by or on behalf of public authorities
- Evaluating creditworthiness
- Assessing risk and pricing for insurance
- Certain law enforcement use cases
- Certain migration, asylum and border control management use cases
- Assisting a judicial authority in researching and applying law and facts

Accuracy

Hallucinations cannot be completely avoided in LLMs today. CLAIRE minimizes hallucinations by using the following methods:

- Leveraging high-quality training data
- Fine-tuning with domain-specific data
- Leveraging retrieval-augmented generation (RAG) using the tenant's metadata and other data sources to ground responses in enterprise context
- Carefully prompting and using a response verification pipeline

Informatica evaluates CLAIRE GenAI features and models using curated test cases to help ensure that the output meets the requirements of business use cases. Informatica also monitors failure cases in production and periodically reviews them to determine common failures, which are then used for model improvement. In case of failure in generating responses or an unexpected result, Informatica informs users when CLAIRE is unable to handle the requests. Additionally, users can communicate this to Informatica through the user interface. Such feedback will be used to further improve the GenAI features and models.

Confidentiality

CLAIRE stores all customer prompts and respective responses in isolated storage partitions specific to the customer, excluding customer data in the responses, which is not stored by CLAIRE.

Geographical Deployment Model and Inference

CLAIRE uses global and regional deployment stacks to strategically manage resources and adhere to security and compliance as per regional requirements.

CLAIRE uses the following deployment stacks:

- **CLAIRE global intelligence stack:** A combination of expert LLMs, hosted by Informatica, used for inferences for CLAIRE Copilot capabilities. The global stack resides in the US.
- **CLAIRE regional full stack:** Includes provisioned third-party LLMs for CLAIRE GPT capabilities based on availability from the model provider. This stack serves end users and includes metadata, a conversation store and a moderator framework. A regional full stack is available in the US and EU.
- **CLAIRE regional partial stack:** Serves end users and only includes metadata, a conversation store and a moderator framework. A regional partial stack is available in APJ, Canada, Dubai, and the UK.

For copilots, CLAIRE is deployed in a centralized location with the global intelligence stack.

For CLAIRE GPT, CLAIRE is deployed in a regional full stack and in regional application stacks. This approach ensures that models are accessible by the regional application stacks while adhering to regional security and compliance requirements.

When a user initiates a request, CLAIRE determines which stack to use based on the user geography and whether the user enters the request in CLAIRE GPT or CLAIRE Copilot. The CLAIRE GPT and Copilot functionality available depends on the customer geography.

For example:

- A request from a user through CLAIRE GPT in the European Union will only use the CLAIRE regional full stack available in the EU because models required for metadata and data processing are present.
- A request from a user through CLAIRE GPT in the UK will use the CLAIRE regional partial stack in the UK for data inferencing and the CLAIRE regional full stack in the EU for metadata processing.
- A request from a user in Data Integration using CLAIRE Copilot will use the CLAIRE Global intelligence stack regardless of the location. In this case, only metadata is processed.

Figure 1 shows the CLAIRE geographical deployment model and inference as of November 2025.

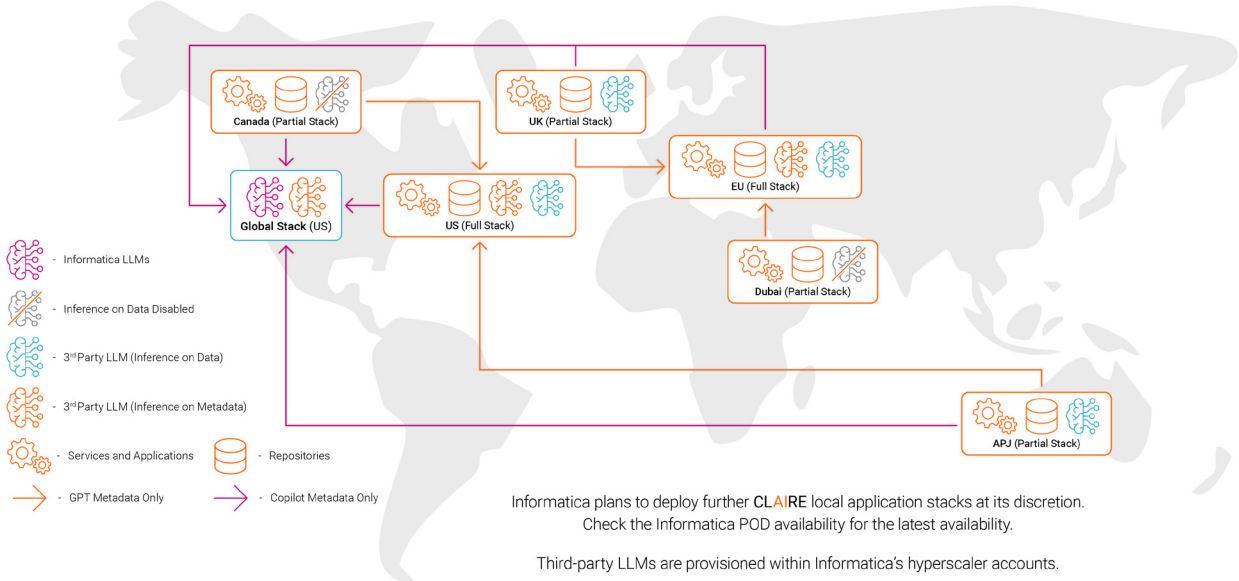

Figure 1. CLAIRE Geographical Deployment Model and Inference

The following table summarizes geographical distribution, data persistence, data inference and model training aspects of this deployment model:

	CLAIRE Global Intelligence Stack	CLAIRE Regional Full Stack	CLAIRE Regional Partial Stack
Geographical deployment	Global (deployed in the US only)	Local (e.g. US, EU)	Local (e.g. UK, APJ, Canada, Dubai)
Persistence	None (stateless)	Metadata and conversation	Metadata and conversation
Inference (processing)	Metadata for copilot capabilities	Metadata for CLAIRE GPT capabilities Data for CLAIRE GPT capabilities	Metadata for GPT capabilities (sent to the nearest regional full stack) Data for GPT capabilities (disabled if the third-party LLM isn't available locally)

Opt-In, Opt-Out and Disposition of Data

Enabling or Disabling GenAI Services

Customers can enable or disable CLAIRE GPT and CLAIRE Copilots or any other GenAI service from Informatica at any time through the Informatica Administrator interface. By default, these options are disabled, giving customers complete control over the activation of CLAIRE GPT and CLAIRE Copilots or any other GenAI service in a specific organization.

Opting out of Metadata Sharing with Customer-Agnostic Models

Customers can opt-out of providing their metadata for training customer-agnostic CLAIRE models at any time through the Informatica Administrator interface, which allows customers complete control over the use of their metadata at any point in time.

Note: Opting out might result in the loss of non-critical CLAIRE functionality.

Disposition of Data

Informatica's policy is to retain processed customer data and customer-specific metadata for at least thirty (30) days after termination or expiration of a customer's subscription to the cloud service.

Informatica de-identifies and deletes processed customer data and customer-specific metadata within sixty (60) days of termination or expiration of the customer's subscription to the cloud service.

Informatica will promptly comply to the extent practicable with written requests to destroy processed customer data within shorter time periods than those indicated above and provide written certification of destruction of processed customer data upon the customer's written request.

For more details on Informatica data retention and data destruction policies, please visit **Informatica Security Addendum**.

Data Access in CLAIRE GPT

Informatica recognizes and understands the importance of the security and privacy of customer data. The data exploration capabilities of CLAIRE GPT entail fetching data in real time from a data source. However, the data isn't persisted in any Informatica repository and is only viewable in and downloadable from a session while CLAIRE GPT is open. If a user exits CLAIRE GPT, the data will no longer be visible or available.

CLAIRE GPT honors metadata access control permissions for data assets as established by the Informatica administrator. For example, users can only discover the assets on which they have permissions and will only be able to perform data exploration over connections for which they have permissions.

Localization and Residency

Adhering to data residency and artificial intelligence regulations is of paramount importance to Informatica.

CLAIRE retains end-user prompts and generated responses. This retention is necessary to resume previous interactions and retrieve contextual information, which is vital for large language models (LLMs) interacting with end users.

Conversation data older than 180 days is stored but designated as inactive. Informatica deletes conversation data after 13 months. If a customer disables the CLAIRE GPT and Copilot services, Informatica deletes the conversation data after two months. For customers who have consented to metadata sharing, Informatica doesn't delete the metadata after 13 months or two months following service deactivation.

Informatica CLAIRE GPT and CLAIRE Copilots might transfer metadata from the CLAIRE regional stack to either the global intelligence stack or the nearest CLAIRE regional stack based on third-party LLM availability. Informatica doesn't transfer any data outside of the tenant's deployment region. Transfers are always conducted in accordance with governing law and, where applicable, with Informatica's binding corporate rules.

The specific metadata that might be transferred includes:

- End-user prompts and embeddings.
- Embeddings derived from the tenant's Data Governance and Catalog, Data Marketplace, or Master Data Management instance to provide context relevant to the conversation metadata. This context is utilized as input for an NL2SQL LLM model, generating an SQL query.
- Natural language prompts containing metadata, which are queried per the end-user prompt and are used to generate the final output to the end user.

During transfer between the regional and global stacks, all information exchanges occur through protected channels wherein data is encrypted in transit utilizing HTTPS and mTLS encryption mechanisms.

Third-Party Software and Models

CLAIRE uses open-source and proprietary large language models after thorough testing, fine-tuning and with Informatica developed guardrails. Informatica also routinely monitors known vulnerabilities, updates software, hardens code and follows its own security and governance policies. CLAIRE uses the following third-party LLMs:

- Microsoft Azure OpenAI provisioned under Informatica's Azure account.
- Anthropic Claude provisioned using Bedrock under Informatica's Amazon Web Services account.

Secure Development Lifecycle

Informatica follows a thoughtfully crafted secure development lifecycle to preemptively pinpoint and resolve security concerns in our products. This, in turn, cultivates a vigilant, security-conscious mindset among our product engineering teams and significantly reduces the likelihood of security breaches.

Informatica's approach to CLAIRE AI security spans several key stages, each designed to integrate security measures seamlessly within the development process. This approach ensures that security considerations are embedded from the outset and throughout the lifecycle. Vulnerability discovery is conducted through multiple layers of identification sources, or what we internally refer to as "sources of truth, as illustrated in Figure 2."

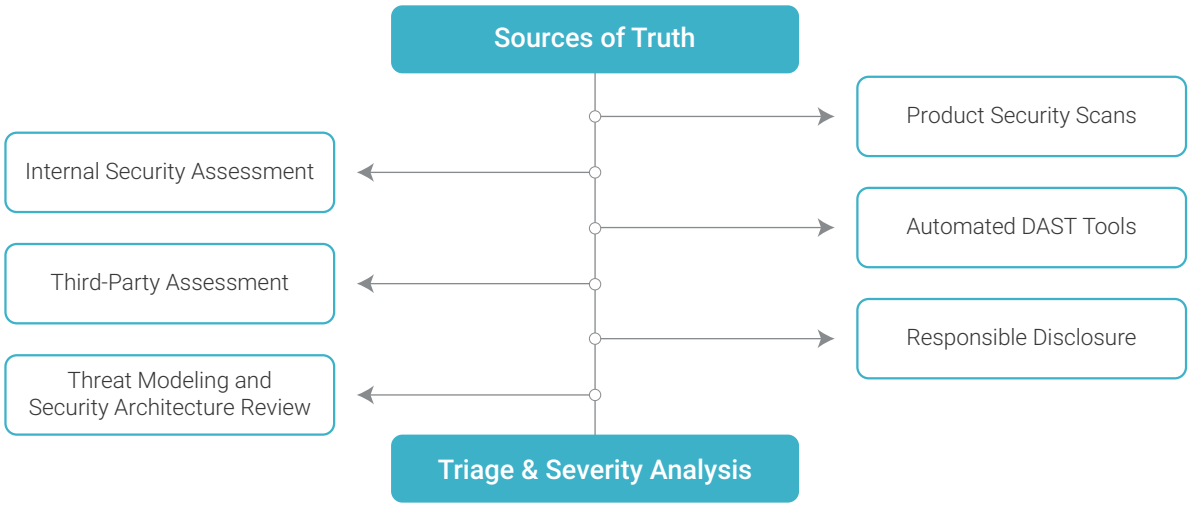

Figure 2. Vulnerability Identification — "Sources of Truth"

Product Security Scans

Informatica products undergo automated scans driven by continuous integration and continuous delivery (CI/CD) to identify vulnerabilities at various levels, from third-party components to application source code. This foundational aspect of our shift-left security strategy enables early detection and remediation during the SDLC.

Internal Security Assessment

Informatica's team of cybersecurity engineers collaborate closely with development teams on each epic and story created during development sprints, ensuring continuous application security assurance. Informatica follows OWASP guidelines.

Automated DAST Tools

Integrated with CI/CD pipelines, dynamic application security testing (DAST) tools automate repetitive tasks in vulnerability discovery and analysis, providing rapid feedback to developers.

Third-Party Security Assessment

Informatica engages reputable third-party cybersecurity assessment providers to ensure continuous, industry-standard assessments of our products.

Responsible Disclosure Program

Through a responsible disclosure program, Informatica invites professional hackers to test our external attach surfaces for vulnerabilities and recognizes those hackers for their efforts.

Threat Modeling and Security Architecture Review

Informatica employs threat modeling (TM) and security architecture review (SAR), as shown in Figure 3, to systematically identify and assess potential threats, vulnerabilities and risks during the design phase of systems or applications.

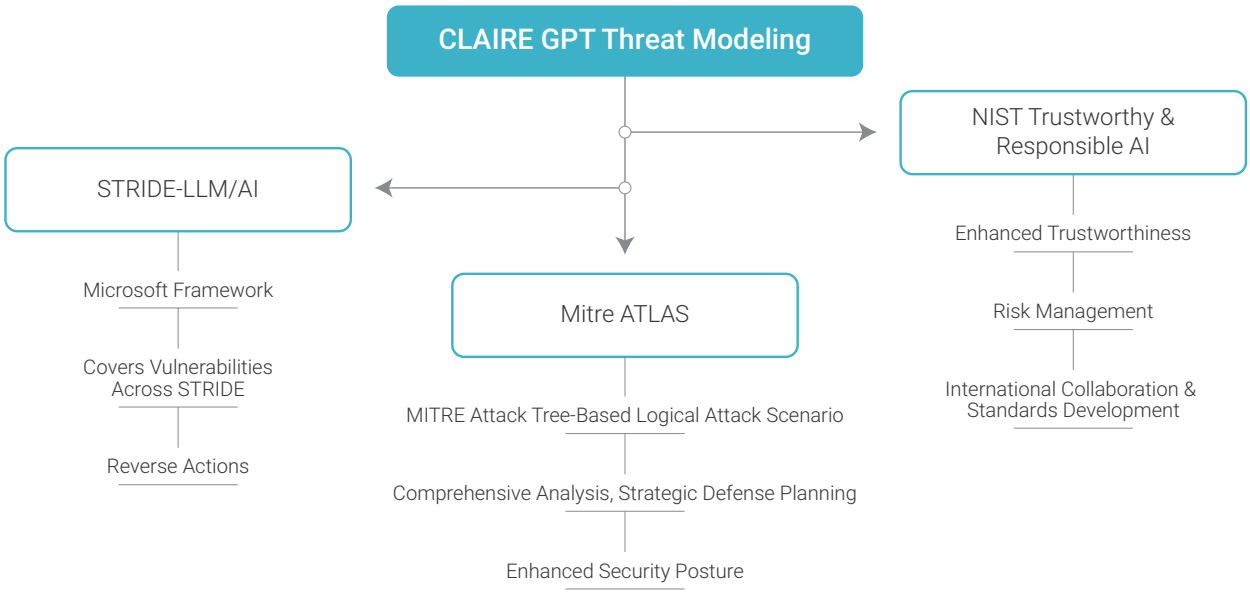

Figure 3. CLAIRE GPT Threat Modeling Framework

Triage and Severity Analysis

Informatica utilizes industry-standard practices (CVSS, Mitre Attack Framework, OWASP, EPSS, NIST) for classifying, risk-rating and describing cyberattacks and intrusions. Our product security team rigorously assesses each vulnerability identified to understand its potential exploitability, consequences, severity and impact.

AI Governance Council

Informatica's AI governance framework is structured around the AI Governance Council, which oversees our adherence to ethical AI development and deployment. The council includes representatives from research and development, product, security, compliance, legal, architecture and program teams dedicated to responsible AI.

The council's responsibilities include:

- Establishing ethical frameworks, policies and processes to guide AI initiatives driven by changes in industry practice, new laws and new desired use cases in Informatica products
- Ensuring that CLAIRE systems operate reliably and safely and align with Informatica's Responsible AI principles
- Assisting Informatica internal AI practitioners and users in applying responsible AI principles and best practices

Informatica IDMC Certifications and Compliance

Certifications and Compliance

The security of customer data is a critical objective of the IDMC platform. Informatica has established a risk-based information security program to protect Informatica and our customers' data security and privacy.

Informatica has voluntarily undertaken and/or is required by contractual obligation to perform in accordance with the standards listed below. These standards are measured by internal security teams and champions, third parties and external assessment partners such as AICPA-accredited external audit firms.

Among others, IDMC has achieved the following certifications: SOC 1, SOC 2, SOC 3, TX-RAMP Level 1 and HIPAA/HITECH. Some IDMC services and corporate environments are certified to other industry-specific compliance standards, such as GxP, U.K. Cyber Essentials Plus and U.S. Government FedRAMP Moderate.

For a complete list of certifications, assessments and standards for IDMC, please visit the **Informatica Trust Center**.

About Us

Informatica (NYSE: INFA), a leader in AI-powered enterprise cloud data management, helps businesses unlock the full value of their data and AI. As data grows in complexity and volume, Informatica's Intelligent Data Management Cloud™ delivers a complete, end-to-end platform with a suite of industry-leading, integrated solutions to connect, manage and unify data across any cloud, hybrid or multi-cloud environment. Powered by CLAIRE® AI, Informatica's platform integrates natively with all major cloud providers, data warehouses and analytics tools — giving organizations the freedom of choice, avoiding vendor lock-in and delivering better ROI by enabling access to governed data, simplifying operations and scaling with confidence.

Trusted by about 5,000 customers in nearly 100 countries — including over 80 of the Fortune 100 — Informatica is the backbone of platform-agnostic, cloud data-driven transformation.
Informatica. Where data and AI come to life.™

Worldwide Headquarters
2100 Seaport Blvd.
Redwood City, CA 94063, USA
Phone: 650.385.5000
Fax: 650.385.5500
Toll-free in the US: 1.800.653.3871

informatica.com
linkedin.com/company/informatica
x.com/Informatica

CONTACT US

Where data & AI come to

IN09-4999-1025