



VERINT AI GOVERNANCE PRINCIPLES AND PRACTICES

Responsible and Ethical AI

Prepared for Verint Customers and Partners

This document contains confidential and proprietary information of Verint Systems Inc. and may not be distributed to persons or organizations for which is not intended.

Table of Contents

Introduction	2
Verint AI Strategy	2
Guiding Principles for AI	4
Appropriate: Using AI Appropriately for Our Solutions	4
Secure, Private, and Compliant	5
Fair: Building AI to Avoid Unfair Sensitive Bias	5
Accountable to People: Delivering Trusted Output That Is Accountable to People	5
Processes and Controls	6
Development and Deployment Practices	6
Verint AI Model Management: Data Use Sourcing Rules	7
AI Considerations and Guidelines	8
Transparency	8
Ethical Decision Making	8
Bias and Discrimination	9
Privacy Concerns	9
Security Risks	9
Legal and Regulatory Considerations	10
Misinformation and Manipulation	10
Unintended Consequences	10
AI Monitoring	11
Types of Monitoring	11
Types of Data Monitored	11
Types of Content Guardrails	12
Agent Factory: Governance for Agentic AI	12
Agent and Workflow Tracing	12
Support for AI	13

Introduction

As a provider of Customer Experience (CX) Automation offerings that incorporate artificial intelligence (AI), Verint® is committed to ensuring its solutions not only deliver significant ROI to our customers but also leverage AI in an ethical, responsible manner. The Verint CX Automation Platform™ is architected with data and AI at its core, with best-of-breed applications and AI-powered bots and AI agents that automate a broad range of functions impacting the cost and quality of CX. This document highlights our AI strategy, the principles we follow as we execute on that strategy, and the processes, controls, and guidelines we have put into place to guide our use of AI.

Verint AI Strategy

Verint Da Vinci™ AI is core to the Verint CX Automation Platform and powers the applications our customers use every day. To keep pace with the rapid innovation in the AI industry, Verint's AI model development is an open approach. That strategy incorporates both proprietary and third-party models. Our open approach to AI means customers can be confident in an investment in Verint Da Vinci AI for the long term.


Verint Da Vinci AI maximizes impact by injecting AI directly into business workflows, putting AI at the fingertips of our customers. By injecting AI directly into business workflows, users of our solutions can:

- Convert unstructured data to intelligence and action.
- Use Agentic AI to improve business processes.
- Detect anomalies in data process flows and take appropriate action.


Identify insights, trends, and opportunities with predictive modeling.

We have infused Verint Da Vinci AI into a growing team of specialized bots and AI agents, each of which is built to use AI to do one thing and to do it well. Our bots live on the Verint CX Automation Platform and are deployed for use within the platform. Customers can generally choose to deploy any or all of our bots and AI agents with their Verint solutions. Each bot and AI agent works to augment our customers' human workforces and enhance CX Automation, such as by:

- Workforce Engagement – Customers can ensure the right employees are in the right place at the right time, improving staffing efficiency and service coverage.
- Quality Automation – Customers can automate their evaluations to gain full visibility into quality and compliance.

- 
- Copilots – Customers can double agent capacity by automating manual tasks, increasing throughput without adding headcount.
 - Knowledge Automation – Customers can rely on a single, consistent, and accurate source of their enterprise knowledge artifacts to power both human agents and bots.
 - CX Analytics – Customers can turn raw CX data into actionable insights in minutes, accelerating decision-making.
 - Conversational and Agentic AI – Customers can automate across all CX use cases and channels subscribed for and enabled in the customer configuration, delivering consistent experiences at scale.

Our AI models are trained on public, open source, licensed, synthetic, anonymous, and de-identified data sets generated from the Verint CX Data Hub to make smarter and faster decisions. These trained AI models are not publicly accessible, no customer data is incorporated into enhancements or new offerings, and these models are only available for use by our customers.

- 
- Engagement data may consist of:
 - Interaction data across channels.
 - Experience data from both customers and employees.
 - Workforce performance data.

The Verint CX Automation Platform has unique capabilities to harmonize these diverse data types from multiple applications into a cohesive whole. The AI models within Verint Da Vinci can also be further trained using an organization's specific data to improve their ability to provide better insight and enhance CX Automation specifically for that customer. When such additional training occurs, those uniquely trained AI models are available for use only by that specific organization.

Verint Da Vinci AI services and models are built in Verint Labs, in combination with data from the Verint CX Data Hub within the bounds of the Verint CX Automation Platform. Verint Labs are secure, 'least privilege' access controlled regional environments aligned to the same environments where production data is hosted. These are managed by a team of data and AI scientists actively collaborating with customers on specific projects, with the infrastructure to securely incorporate tools and algorithms from technology partners, external research organizations, and universities. This team drives the research-to-product pipeline of Verint Da Vinci and the Verint CX Automation Platform by focusing on improving our AI and machine learning capabilities.

Guiding Principles for AI

AI is integral to the Verint CX Automation Platform and powers the applications our customers use every day. Verint realizes that advanced technologies like AI can raise important challenges that must be addressed clearly, thoughtfully, and directly. The principles that guide our use of AI in our offerings include:

- Deploying AI only for appropriate uses.
- Using AI securely and in compliance with privacy and other regulations.
- Working to ensure AI operates in a fair and safe manner.
- Designing AI to be controlled and accountable to people.

These principles inform our commitment to developing AI technology responsibly and ethically. They are incorporated into our company culture and are ingrained in the processes used to deliver the technologies and solutions we provide our customers. Workforce Engagement – Customers can ensure the right employees are in the right place at the right time, improving staffing efficiency and service coverage.



Appropriate: Using AI Appropriately for Our Solutions

Verint works to limit potentially harmful applications as we design and deploy AI technologies. We review the AI technologies we deploy, their use cases, and the related risks. Verint also evaluates how the AI technology will scale across billions of interactions and global deployments to help ensure the AI technology we offer can provide the value we intend.

Secure, Private, and Compliant

AI technologies are often exposed to large amounts of personal data, raising issues related to data privacy and security. Verint designs and develops AI technologies using research and development best practices. Verint works to develop its AI systems in compliance with applicable laws and regulations and designs such systems in a manner that allows compliant use by customers. This includes respecting third-party intellectual property rights in the design and development process.

Verint also supports the adoption of data protection, compliance, and privacy regulations to address risks associated with AI, as well as the implementation of safe data handling practices, and undertakes annual policy reviews and updates to keep pace with changes.

Fair: Building AI to Avoid Unfair Sensitive Bias

AI algorithms and datasets can reflect and/or reinforce biases. Verint understands that distinguishing fair from unfair biases is not always simple and may differ across cultures and societies. We seek to avoid the use of our AI solutions resulting in bias based on race, color, religion, sex, age, national origin, genetic information, disability, veteran status, and any other classification or characteristic protected by applicable government, state or local laws.

Accountable to People: Delivering Trusted Output That Is Accountable to People

The AI technologies and use cases Verint offers provide appropriate opportunities for assessments, relevant transparent explanations, and value. Our AI technologies are subject to appropriate human direction and control. This helps to manage risks related to ethical use cases, unintended consequences, and transparency, and provides mechanisms to help alleviate misinformation and manipulation.

Processes and Controls

Verint has established and maintains standard practices across its research, engineering, development, and operational activities when AI is being offered in a solution, including:

- Development and deployment practices in accordance with compliance standards.
- Model management and data use and sourcing rules.

Risk assessments built into processes and reviewed by compliance teams.

- Service and AI monitoring.
- Support through our Verint Connect forums.


Development and Deployment Practices

Verint's AI capabilities are developed and reside within the Verint CX Automation Platform. The Verint CX Automation Platform includes secure regional cloud environments wherein Verint maintains industry-standard security practices that follow applicable guidelines set forth by PCI, GDPR, ISO, SOC2, and HIPAA, as examples.

Additional Security and Privacy Resources:

- [Verint Trust Center](https://trust.verint.com/) (https://trust.verint.com/)
- [Verint Cloud Security and Compliance Overview](https://www.verint.com/verint-cloud-security-and-compliance/) (https://www.verint.com/verint-cloud-security-and-compliance/)
- [Verint Certification Product Matrix](https://www.verint.com/certification-product-matrix/) (https://www.verint.com/certification-product-matrix/)

Verint has implemented and maintains appropriate technical and organizational measures intended to protect personal data against accidental, unauthorized, or unlawful access, disclosure, alteration, loss, or destruction. Our security practices govern assets and data within our control, as well as our operating environments.




Verint AI Model Management: Data Use Sourcing Rules

Verint has also established practices for selecting and/or building, deploying, and managing AI models and data use restrictions:

Customer data is always stored in the selected Verint secure regional cloud environment when using our AI services. Some Verint Da Vinci AI services are provided by third parties, as published in the Verint Supported Commercial and FOSS AI Models guide, and may be dynamically allocated by the third party based on service loads. This may include regional inference endpoints (e.g., AWS CRIS) where processing may occur elsewhere within the same geographic area. However, customer data is never stored outside the regional environment in which it originated, and Verint manages the secure API credentials and configuration for these services throughout, so customer data remains under Verint's control.

Verint may access, process and use customer data to provide, maintain, and assess improvements to the SaaS Service provided; no customer data may be incorporated into product enhancements or new offerings.

- Verint's AI provision in its Master Service Agreement (MSA) makes clear:
 - » Inputs and outputs are confidential information of our customer
 - » Without a customer's written consent, Verint will not directly train AI models with customer data



Note that third-party AI services operate within the Verint Secure Regional Cloud environments. These services run in a private, not public, configuration and are fully controlled by Verint. Verint does not permit appointed third-party model providers to use customer data to train those third-party AI models and does not provide access to inputs and outputs to those third parties.

AI Considerations and Guidelines

Verint has established a framework of considerations and guidelines for key factors that are involved in the development, deployment, and use of AI technologies. The framework addresses some common questions about how AI impacts these solutions. For example, is the service meant only to provide information for humans to act on, or is it meant for making automated decisions? How explainable are the model's predictions? Does the implementation provide any means to understand why the engine made the prediction it did? Is AI used in a way that can impact human employment or pay? Verint seeks to address these considerations across our solution offerings, as discussed above, and evaluates them from the following perspectives:

Transparency

Lack of transparency in AI systems, particularly in deep learning models, is a pressing issue. This opaqueness obscures the decision-making processes and underlying logic of these technologies. To help provide insight into those opaque areas, we seek to provide the context and scope of the information that the underlying algorithms and processes are using. This includes:


- Display transparency: Provides a real-time understanding of the task being performed by the AI, and of the actions the AI system is taking.
- Explainability: Provides information in a backward-looking manner on the logic, process, factors, or reasoning upon which the system's actions or recommendations are based.

Ethical Decision Making

Instilling ethical values in AI systems, especially in decision-making contexts with significant consequences, presents a considerable challenge. Researchers and developers must prioritize the ethical implications of AI technologies to avoid negative societal impacts.

We seek to address this by:

- Performing an AI ethics review for new projects as part of the research scope and definition process. The AI ethics review process explores the possible harms that can be caused by normal use of the AI: with high-quality results; with incorrect or low-quality results; how the technology could be misapplied; and how the AI can correctly or incorrectly learn as it is used. The Verint ethics review includes the following questions, and a discussion of how each issue may be mitigated:
 - » What possible harms exist when the technology is being used as intended and functioning correctly?
 - » What possible harm exists when the technology is being used as intended but giving incorrect results?
 - » What possible harms exist following from potential misuse of technology?
 - » If the system learns from user input once deployed:
 - What checks and limitations are intended with respect to system learning?
 - Does that learning incorporate any customer-specific data?
 - Is the learning applicable to the offering generally, or only at the tenant level?



Bias and Discrimination

AI systems can inadvertently perpetuate or amplify societal sensitive biases due to biased training data or algorithmic design. To minimize discrimination and ensure fairness, it is crucial to invest in the development of unbiased algorithms and diverse training data sets.

We seek to address this by:

- Being aware of commonly known biases that may be present in AI systems, such as data bias, algorithmic bias, and confirmation bias.
- Periodically reviewing, evaluating, and addressing AI-generated outputs during the development process for potential biases and inaccuracies.
- Addressing bias or accuracy issues reported through our Incident Management System as a bug/defect.
- Using AI systems provided by vendors with transparent methodologies and documentation to better understand their decision-making processes.
- Documenting and communicating any identified biases and mitigation efforts.

Privacy Concerns

AI technologies often collect and analyze large amounts of personal data, raising issues related to data privacy and security. To mitigate privacy risks, we endeavor to adhere to safe data handling practices.

We seek to address this by:

- Complying with privacy and other laws applicable to Verint as a provider of AI.
- Adhering to Verint's Code of Conduct, Global Information Security Policy, and privacy policies when accessing and processing personal data.
- Complying with our contractual obligations concerning the handling of customer data, including with respect to personal data.

- Using secure and data privacy linked environments, such as Verint Labs, to research, develop, and test AI systems or services.

Security Risks

As AI technologies become increasingly sophisticated, the security risks associated with their use and the potential for misuse also increase. Hackers and malicious actors can harness the power of AI, including frontier AI systems that introduce new attack vectors, to develop more advanced cyberattacks, bypass security measures, and exploit vulnerabilities in systems.

We seek to address this by:

- Adhering to Verint's internal AI Use Policy, these guidelines, and our Global Information Security Policy when using and developing AI systems.

Applying layered controls designed to protect AI services against prompt injection, unauthorized data disclosure, model misuse, tool abuse, malicious agent behavior, and AI supply-chain threats. With this consideration in mind, AI-enabled services undergo structured security review and testing, including adversarial testing where appropriate, prior to production deployment.

- Using secure and data privacy linked environments, such as Verint Labs, to research, develop, and test AI systems or services.

Legal and Regulatory Considerations

New legal frameworks and regulations are being adopted around the world to address the unique issues arising from AI technologies, including allocation of responsibilities, liabilities, and intellectual property rights. Legal systems are evolving to keep pace with technological advancements and protect the rights of individuals and businesses.

The rights and laws that may apply to Verint offerings include:

- Data protection and privacy laws.
- Intellectual property laws.
- Laws applicable to employment, the protection of protected classes, or other human interest related laws and regulations.
- Laws regulating the use of AI itself.

We seek to address this by:

- Developing AI systems designed to be in compliance with applicable laws and regulations and designing such systems in a manner that allows compliant use by customers.

Remaining vigilantly informed by industry frameworks including the NIST AI Risk Management Framework, OWASP guidance for LLM applications, ISO AI governance standards, and applicable regulatory requirements.

- Respecting third-party intellectual property rights in the design and development process.
- Adhering to Verint's internal AI Use Policy and these guidelines when developing AI services or systems that are used by our customers.
- Monitoring and reviewing applicable updates to legal and regulatory requirements and best practices.

Misinformation and Manipulation

AI-generated content, such as deepfakes, contributes to the spread of false information and the manipulation of public opinion. Efforts to detect and combat AI-generated misinformation are critical in preserving the integrity of information in the digital age. Data from

individuals (such as customers of our customers) may be collected and fed into AI services for further enrichment or automation and often includes the type of information (e.g., personal sensitive information, voice audio recordings, pictures, account information, etc.) that can be used by malicious actors to falsely represent an individual. Scrutinizing the security of each AI service we use and ensuring the AI service does not act on behalf of or represent itself as the consumer are core focuses for us.

We seek to address these concerns by:

- Ensuring the data inputs going into an AI service are protected, and there are safeguards to prevent interception or use by third parties for alternative purposes.
- Ensuring the AI service does not act on behalf of or represent itself as the consumer.

Unintended Consequences

AI systems, due to their complexity and lack of constant human oversight, might exhibit unexpected behaviors or make decisions with unforeseen consequences. This unpredictability can result in outcomes that negatively impact individuals, businesses, or society. Robust testing, validation, and monitoring processes can help developers and researchers identify and fix these types of issues before they escalate.

We seek to address this by:

- Designing AI models used in the service to provide explanations displayed to the human decision maker.
- Including in product documentation descriptions of the AI models used with the Verint CX Automation Platform.
- Establishing a development approval process that requires review of how AI capabilities are combined with other Verint functionality, and how the AI will be adopted within the Verint offerings to restrict further development use cases without required approvals.

AI Monitoring

The Verint CX Automation Platform provides capabilities to both actively (e.g., continually observe and analyze) and passively (e.g., issues reported through our standard support tickets) monitor AI services deployed for customers in regional clouds around the world. Model monitoring serves as a measure to maintain the health, efficiency, and effectiveness of AI applications, AI Agents, or bots.

Types of Monitoring

- **Passive Monitoring:** Verint provides a mechanism for clients, partners, and Verint employees (e.g., professional services team members working on behalf of clients) to submit support tickets about application or bot behaviors that are unexpected or should be investigated. Verint's support organization will coordinate with the appropriate product, engineering, and research teams behind the scenes to address the issue or to answer a question.
- **Active Monitoring:** Active model monitors may be put in place as part of the core project definition across engineering/architecture, research, and development depending on the use case. Each application or bot that deploys AI can monitor performance and incorporate enhancements back into the core model, provide additional training data, perform additional model training events, or implement other remediation tasks to address problems identified by the monitor.

Types of Data Monitored

Active monitoring on the Verint Cloud Platform draws on three data capture categories: Logging, Metering, and AI Monitoring.

- **Logging:** Logging data captures the operational signals generated as requests flow through the platform, providing the visibility needed to monitor service health and troubleshoot issues (e.g., status codes, heartbeats, error traces, timing, and call metadata).
- **Metering:** Metering data tracks the volume and type of platform usage generated by each customer, forming the basis for license compliance verification and billing (e.g., API call counts, transaction/session volumes, entitlement usage, and billing units).
- **AI Monitoring:** AI monitoring data captures the prompts, inputs, outputs, and quality signals generated when AI models are invoked, enabling Verint to evaluate model performance and adherence to Responsible and Ethical AI standards (e.g., prompt text, model inputs/outputs, confidence/quality scores, and compliance flags for toxicity, bias, prompt injection, and hallucination and drift).



Types of Content Guardrails

Active Monitoring also includes a Content Guardrail service that provides text-based content safety analysis for AI-powered bots, protecting against toxic content and prompt injection attacks. In its current state (V1), Guardrails scans outbound prompts for injection risk and scans model responses for toxicity, returning probability scores across categories such as overall toxicity, obscenity, insults, sexually explicit content, identity-based attacks, and threats. While Content Guardrails actively identifies and scores harmful content, it does not block it or take protective action at that time. The guardrails are actively logging scan information and returning that information to the Verint bot or AI agent, and the review and analysis of that information takes place in the regional Verint lab by Verint research and data scientists to assess the potential for corrective actions to be taken later. Content Guardrails is deployed across a broad set of Verint Da Vinci AI engine services and bots in all production regions, with scans running synchronously on every request and a latency cut-out threshold to limit performance impact.

Agent Factory: Governance for Agentic AI

As Verint expands the Da Vinci AI platform to include Verint Agent Factory™ — a capability that allows customers to build AI agents using Verint Services and connect those agents to external tools and systems via the Model Context Protocol (MCP), as well as manually define node-based process workflows that orchestrate AI Agents alongside deterministic steps — Verint extends its Responsible and Ethical AI principles to address the unique risks that agentic, tool-connected AI introduces.

Agent and Workflow Tracing

Agent Factory provides AI agent and workflow tracing across every agent and node-based workflow built on the platform. This tracing captures each step, tool invocation, and decision an AI Agent or workflow makes during execution, enabling a user to see not only the outcome but the reasoning path that produced it.

We seek to address transparency and accountability for agentic AI by:

- Recording each decision point, tool or MCP server call, and data handoff within an agent's or workflow's execution as a discrete, reviewable step.
- Making the full execution trace available to the user who owns the agent or workflow, so the “why” and “how” behind a given outcome can be reconstructed after the fact rather than treated as a black box.
- Retaining trace records in a manner consistent with Verint's data handling and retention practices, supporting both internal quality review and customer-facing audit needs.

Policy Gate: Bounding Agent and MCP-Based Actions

Because each MCP, A2A, or third-party tool connection introduces new actions, data, and potential emergent behaviors, Agent Factory applies a policy gate as a decision boundary in front of every tool invocation or task delegation. The policy gate evaluates an agent’s intent, context, and permissions before allowing it to act, which helps prevent goal drift — where an agent optimizes for whatever capabilities happen to be available rather than the user’s original objective — and blocks unintended tool chains or policy-violating actions that can emerge from autonomous, multi-step reasoning.

We seek to address the compliance and scope risks of agentic AI by employing:

Purpose-aware egress gates: Before every external tool or MCP call, these gates validate that the requested data use is compatible with the declared purpose, that only necessary fields are used, and that the destination region and transfer mechanism are permitted — failing closed (deny by default) unless explicitly allowed.

Proportional safeguard tiers: Verint applies baseline, heightened, or strict controls based on the risk of the action, ranging from read-only tools with ephemeral memory at baseline, to human review before consequential actions at the strict tier.

Memory governance: Verint defaults agent memory to off, limiting session memory to ephemeral, task-scoped execution, and requiring explicit opt-in with time-boxed retention and cascade deletion for any account-level memory.

End-to-end traceability: Verint records, for every workflow step, the tool used, fields accessed, purpose, lawful basis, region, decision (allow, redact, or block), and outcome, tying each action back to a single audit trail.

Together, agent and workflow tracing and the policy gate ensure that AI Agents and node-based workflows built through Agent Factory remain aligned to their defined task, operate within customer- and region-appropriate boundaries, and produce a demonstrable, auditable record suitable for global regulatory frameworks such as the GDPR and EU AI Act.

Support for AI

Verint Connect (<https://connect.verint.com/>) is the Verint extranet for customers and partners and is the primary support mechanism for these constituents to contact Verint, ask questions, find information, report problems, and receive training materials. If any question, issue, or problem with an AI component of a Verint solution is encountered, it can be reported through Verint Connect with the appropriate severity rating for a response by Verint.

Americas

info@verint.com
+1 770 754 1900
1-800-4VERINT

Europe, Middle East & Africa

info.emea@verint.com
+44 (0) 1932 839500

Asia Pacific

info.apac@verint.com
+ (852) 2797 5678



verint.com



x.com/verint



linkedin.com/company/verint



verint.com/blog