



NetSpeak Security & AI Governance Whitepaper

NETSPEEK – FOR EXTERNAL RELEASE

v1.5 - October 2025

NetSpeak Security & AI Governance Whitepaper

EXTERNAL-FACING DOCUMENT NOTICE

This document is provided by NetSpeak Inc. (“NetSpeak”) for **informational purposes** only. It describes general security measures, practices, and approaches within NetSpeak’s platform as of the document’s publication date. NetSpeak reserves the right to update or modify the content at any time to reflect evolving best practices, regulatory requirements, or product enhancements.

The information herein does **not** constitute a legally binding agreement or warranty of any kind. Specific contractual terms and obligations will be governed by a separate written agreement between NetSpeak and its customers or partners.

By accessing this document, you acknowledge that NetSpeak is not liable for any direct or indirect damages arising from its use. If you have questions about particular security or compliance requirements relevant to your organization, please reach out to NetSpeak’s security team or refer to your formal agreement with NetSpeak.

© 2025 NetSpeak Inc. All Rights Reserved.

Date	Author	Version	Changes
March 2025	Osman Bicakci	1.0	Foundational document
April 2025	Osman Bicakci Preston Naclerio	1.0.1	Formatting
October 2025	Spencer Wise Osman Bicakci	1.5	Content updates

Contents

- Future of AI-Driven Pro AV & UC Management **1**
- Infrastructure Overview **2**
- Security Principles and Strategy **3**
- AI Security and Responsible Use **4**
- Operational Security and Governance **7**
- Conclusion **9**

Document Objective & Scope

NetSpeak is an intelligent platform that provides automation, device management & control, and network assistance. We host our software in a secure public cloud environment with robust data protection capabilities. Our approach to security can be summarized as “defense in depth,” reflecting our commitment to safeguarding every layer and component of the system, from user-facing interfaces to back-end infrastructure. We regularly assess and improve this architecture in response to new threats, industry developments, and evolving cybersecurity best practices.

This document provides a high-level overview of how we protect data, manage access, handle compliance obligations, and ensure that our artificial intelligence (AI) features operate safely and responsibly. By design, NetSpeak is flexible enough to run on various modern, enterprise-grade cloud infrastructures.

The Future of AI-Driven Pro AV & UC Management

As modern enterprises and educational institutions scale their collaboration and communication infrastructures, managing professional audiovisual (AV) and unified communications (UC) technologies has become increasingly complex. IT teams and managed service providers (MSPs) must ensure seamless deployment, operation, and monitoring of thousands of conference rooms and classrooms while maintaining security, uptime, and user experience.

At NetSpeak, we've built an enterprise-grade AI orchestration platform designed to scale into the world's largest organizations while maintaining industry-leading cloud and AI security standards. Our platform transforms how businesses and institutions install, operate, and optimize meeting spaces, leveraging the latest advancements in AI Ops.

Meet Lena - Your New Favorite Co-Worker

NetSpeak has developed the first vertically integrated GenAI platform for the AV and UC industries. At its core is Lena, a **Language Enabled Network Administrator**. Lena assists in deploying, managing, and operating collaboration technologies across conference rooms, classrooms, and similar environments.

Built on **NetSpeak's AI Orchestration Platform**, Lena's capabilities are continuously expanded through agents, modular integrations developed in partnership with leading UC and AV manufacturers. These agents are built and maintained by NetSpeak to retain encyclopedic manufacturer-specific or platform-specific knowledge and are programmed to communicate directly with the partner's hardware or software platforms.

NetSpeak's Agent Fleet

NetSpeak is at the forefront of AI-driven orchestration and automation, and our commitment to innovation has led to the development of the **NetSpeak Agent Fleet**, a team of highly specialized AI Agents designed to autonomously work together to manage and operate AV and UC technologies.

Through strategic partnerships with the leading manufacturers and platform providers in the industry, NetSpeak's AI Agents integrate directly with hardware and software solutions to:

- **Automate Device Setup & Configuration:** AI-driven workflows ensure that meeting spaces are deployed consistently and efficiently, regardless of scale.
- **Predict and Prevent Failures:** Autonomous detection of early signs of degradation in hardware performance, network stability, and software responsiveness.
- **Optimize Meeting & Learning Environments:** AI-powered utilization insights help organizations maximize space efficiency and improve user engagement.

- **Enable Autonomous Troubleshooting & Resolution:** AI agents proactively diagnose and resolve common issues, reducing the burden on IT and MSP support teams.
- **Ensure Security & Compliance:** Our enterprise-grade AI framework adheres to the highest security standards, ensuring data privacy, cloud security, and AI governance across global deployments.

Infrastructure Overview

NetSpeak's infrastructure is deployed within globally recognized public cloud environments. Our decision to host in these clouds is guided by the availability of secure data centres, strong compliance credentials (for instance, SOC 2 and ISO 27001 certifications), and advanced features for encryption, identity management, and network isolation.

The NetSpeak platform itself is made up of several core components that work together to deliver automation and intelligent interaction:

- **User Interfaces (Web):** NetSpeak platform's web interface is the primary way to interact with NetSpeak. They offer administrative interfaces and day-to-day operational views.
- **API Services:** The back-end logic for NetSpeak runs as a set of modular, scalable web services. These services handle requests from the user interfaces, integrate with external systems, and orchestrate various workflows.
- **Agentic AI Runtime:** This part of the system manages AI-related tasks. It can handle natural language queries, orchestrate advanced logic, and interact with external AI models (e.g., large language models for Natural Language Processing (NLP) or specialized modules for vendor-specific tasks, troubleshooting, room checks, speech recognition).
- **Data Storage:** To support features like conversation history, device configurations, and file storage, NetSpeak uses encrypted repositories ranging from object storage to traditional relational databases.
- **Edge Virtual Appliances:** Customers may require an on-site virtualized appliance to facilitate device management within corporate security frameworks. We offer a lightweight, on-premises virtualized component that securely communicates with cloud-based NetSpeak services, acting as a secure gateway to transmit management commands and information to and from the NetSpeak cloud.

Because our platform is cloud-agnostic, these components remain consistent whether hosted in one cloud provider or another. Our standards for security and reliability, however, never change: every environment must meet the same strict criteria.

Security Principles and Strategy

NetSpeek relies on a few overarching principles that guide all decisions regarding the design and operation of our platform:

- **Zero Trust Networking:** We do not make assumptions about trust based on network location or IP addresses. Every request including the requests from inside our private environment, undergoes proper authentication and authorization checks. Internal communication between services is typically encrypted and is governed by strict firewall rules.
- **Encryption for Data in Transit and at Rest:** Whenever data moves between NetSpeek components or between NetSpeek and users, it travels over secure channels using modern TLS (Transport Layer Security). Data stored within our systems is encrypted at rest, ensuring it is protected from unauthorized access even if physical hardware or storage devices are compromised.
- **Role-Based Identity and Access:** We assign granular roles and permissions so that each user, administrator, or system service can only access the data and operations needed for their duties. This approach prevents accidental overexposure and helps contain potential security incidents.
- **Continuous Monitoring and Response:** NetSpeek collects logs, metrics, and telemetry from all significant components and integrates them into a security monitoring platform. This platform continuously scans for anomalies, suspicious activities, or performance bottlenecks. In the event of an alert, our team follows a standardized incident response procedure.
- **Compliance with Industry Standards:** Our security program aligns with frameworks such as SOC 2, ISO 27001 and ISO 42001. We also strive to comply with applicable privacy regulations, such as GDPR, by enforcing data minimization and providing data subject rights handling.

Data Protection

At NetSpeek, we recognize that data protection is paramount for building and retaining customer trust. We maintain a careful classification scheme to differentiate the sensitivity of data we handle. High-sensitivity data, such as credentials, is subject to the strictest set of controls and is never stored unencrypted.

- **Data at Rest:** For all storage solutions, from relational databases to object storage, our default setting is to encrypt data at rest. This encryption uses proven cryptographic algorithms, and the keys are managed in hardware-backed vaults. If customers have their own encryption key requirements, we can often integrate with bring-your-own-key (BYOK) models as well.
- **Data in Transit:** We use TLS (version 1.2 or above) to protect all HTTP-based traffic. This includes connections between user interfaces and our API services, as well as inter-service calls and communication with third-party integrations. We maintain policies

to rotate and renew certificates proactively, reducing the risk that old or compromised certificates remain active.

- **Minimization and Retention:** We only collect the minimum information needed to fulfill the purpose of the service. Whenever possible, personally identifiable data is anonymized. We also enforce clear retention policies that purge or archive data after specified timelines. Customers can request data deletion in compliance with privacy regulations.

Identity and Access Management

NetSpeak relies on a centralized identity system that authenticates both end users and administrators. Depending on the deployment, this system may integrate with industry-standard protocols such as OAuth 2.0 or OpenID Connect. Through this central system, we issue and validate secure tokens for user sessions and inter-service communication.

All high-level or administrative actions require multi-factor authentication (MFA). We also apply the principle of least privilege to ensure no user or system account is granted more privileges than strictly necessary. This “locked down by default” stance drastically reduces the attack surface within NetSpeak.

AI Security and Responsible Use

NetSpeak incorporates AI in a variety of ways, such as understanding spoken commands (speech recognition) and text-based interactions using large language models (LLMs). Because these technologies carry unique risks from data privacy concerns to the possibility of malformed or misleading outputs, we place significant emphasis on proactive security measures and responsible usage guidelines. The following sections describe how we protect user information, preserve the integrity of AI-driven processes, and maintain the high standards expected of NetSpeak’s platform.

Comprehensive AI Gateway

We employ an AI gateway designed to enforce a zero-trust framework for LLM-related traffic. By applying refined security and quota rules across every AI endpoint, we can block or throttle abusive requests in real time, safeguarding system resources against unauthorized consumption or exploitation. This gateway also centralizes enterprise-wide policies to ensure that any generative AI invocation aligns with internal standards and relevant regulations. With this approach, users receive consistent, policy-compliant results while the gateway intercepts potential misuse before it ever reaches the deeper layers of our infrastructure.

Automated Red Teaming for AI

In parallel with our gateway, we continuously conduct adversarial testing on our AI components to surface issues like prompt injection, hallucinations, or biased outputs. By running automated “red team” simulations powered by advanced algorithms and an up-to-date threat intelligence database, NetSpeak remains proactive in detecting novel attack vectors. This automated red teaming framework allows us to identify and neutralize potential vulnerabilities, such as unauthorized instruction overrides or content drift, long before they can harm real users or degrade system reliability.

Real-Time Monitoring and Alerting

Our platform monitors AI activity around the clock, tracking metrics such as response latency, error rates, and resource utilization in real time. If the system detects anomalies, policy violations, or suspicious patterns (for instance, repeated forced context changes), it triggers immediate alerts for rapid intervention. Meanwhile, every transaction is logged with sufficient detail to facilitate prompt debugging and compliance verification, supporting a clear audit trail for internal oversight or external investigations. This continuous assessment also covers multi-language testing scenarios, allowing NetSpeak to maintain consistent security and reliability in diverse linguistic contexts.

Scalable Observability

To address large volumes of AI requests without bottlenecks, NetSpeak’s infrastructure supports robust analytics, including specialized evaluators that measure elements such as completeness or tone in generated responses. This capability helps us track risk indicators or drift trends over time and ensures that even in high-traffic environments, performance remains stable and user satisfaction remains high.

By blending automated red teaming with a real-time AI gateway, continuous observability, and other proactive measures, NetSpeak significantly reduces the likelihood of malicious inputs, hallucinations, privacy breaches, and similar threats. These capabilities enhance our existing commitment to secure model hosting, privacy-focused LLM chat, and prompt injection defenses, delivering a comprehensive, defense-in-depth approach to generative AI.

Secure Model Hosting and Data Isolation

We deploy our AI workloads in secure, isolated environments that meet strict standards for data protection, compliance, and access control. Our architecture combines self-hosted AI models with trusted large language model (LLM) service providers, all operating under our unified governance and security framework. We carefully vet each LLM service provider for enterprise-grade security, compliance alignment, and data residency requirements. All outbound

connections to these providers are tightly managed, and no customer data is ever used for model training or persistent storage beyond the processing session.

All AI workloads, including LLMs and related subsystems, reside in private subnets within our cloud infrastructure and are shielded from the public internet. We maintain network-level segmentation, encryption in transit (TLS 1.3+), and authenticated communication through private API gateways to ensure controlled, auditable data flow at all times.

To provide flexibility across customer tiers and regions, we support multiple model-hosting configurations. Enterprise and regulated customers can choose fully self-hosted or dedicated model instances within our infrastructure, while other tiers leverage pre-approved, private-endpoint integrations to ensure optimal performance, scalability, and cost efficiency without compromising security.

We apply the principle of least privilege throughout our AI ecosystem, granting each service only the minimal permissions required to perform its role. Separate workloads and customer environments remain fully isolated, preventing any form of data commingling unless explicitly authorized under an agreed data-sharing policy. All logs, prompts, and model outputs are securely stored within our private cloud environment, following our encryption, retention, and access-review policies.

Prompt Injection Defenses

Emerging threats in LLM systems often revolve around “prompt injection”, where maliciously crafted inputs attempt to rewrite or override an AI model’s intended behavior. NetSpeak confronts this risk by inspecting all inbound queries via an automated AI compliance platform for unauthorized commands or embedded tokens that could compromise system logic. Our multi-layer approach factors in context analysis and abnormal usage patterns, rejecting or sanitizing any suspicious inputs before they reach the AI engine. This system evolves in response to newly discovered vulnerabilities, mirroring the broader security community’s research efforts to keep pace with changing attack tactics.

Output Filtering and Content Moderation

Even advanced models can produce inappropriate, sensitive, or outright dangerous text when given misleading inputs or taken out of context. To address that, NetSpeak has implemented a moderation layer that analyzes the AI’s output for personal information, offensive language, or violations of our Acceptable Use Policy. Any flagged sections are either redacted or automatically modified before being shown to the end user. Potentially problematic outputs are also logged for further investigation, which helps our team refine future moderation rules and maintain a safe, respectful conversational environment for every stakeholder.

Detailed Governance and Lifecycle Management

All NetSpeek AI systems fall under a rigorous governance program designed to ensure accountability in development, deployment, and maintenance. Every model is registered with a unique identifier, making it easy to track release history, training sources, and any relevant updates or fine-tuning processes. We emphasize data provenance as part of the training pipeline, prioritizing anonymized or masked information wherever possible. Performance metrics are continually measured in areas like accuracy, latency, and error rates, and we also keep track of how frequently disallowed outputs occur to gauge a model's safety profile. Regular internal reviews, led by product and security specialists, further validate that anonymization policies, retention guidelines, and access controls remain aligned with corporate standards and regulatory mandates.

Data Minimization and Customer Control

Our approach to AI is shaped by a philosophy of data minimization: we store and process only what is necessary for effective service delivery. User-generated conversations are never repurposed for broader analytics or training unless the customer specifically opts in. Moreover, organizations can specify how long transcripts or logs should be kept, granting them flexibility to meet internal data governance or compliance requirements. If a user or organization requests an immediate redaction or deletion of content, NetSpeek follows documented procedures to ensure references to that content are removed from active systems and backups where technologically feasible. This way, customers maintain control over their data at all times.

Regulatory Considerations

NetSpeek's handling of AI features is carefully attuned to various regulatory obligations, including GDPR, which grants European data subjects rights like erasure or portability. We strive to build tools and processes that respect these rights, even where AI-driven functionalities are involved. Meanwhile, our alignment with responsible AI frameworks, such as those advocated by standards bodies like ISO or research initiatives within the security community, ensures we stay informed about emerging best practices for bias detection, interpretability, and safe model deployment.

Operational Security and Governance

Network Architecture and Segmentation

NetSpeek is deployed in a segmented environment where services are grouped into different subnets or virtual networks based on their role. This segmentation means that, for example, only specific application subnets can communicate directly with the database tier, and only specific management workstations can communicate with administrative interfaces.

We pair these logical controls with firewall rules and security groups that strictly govern inbound and outbound traffic. This zero-trust approach ensures that even if an attacker gains a foothold in one segment, lateral movement is sharply limited.

Monitoring and Incident Response

We have built a robust security monitoring framework to detect and alert unusual activity. For instance, we aggregate logs from our application servers and system events into a centralized platform. Automated rules can trigger alerts if certain thresholds or patterns (like multiple failed login attempts) are exceeded.

When an alert is raised, our incident response procedure guides the following steps, including containment of the threat, collection of forensic evidence, notification to any affected customers, and eventual eradication of the root cause.

Regulatory Compliance and Certifications

NetSpeak's security posture is aligned with recognized standards and best practices. Although the underlying cloud environments hold numerous certifications (such as ISO 27001, SOC 2, PCI DSS, and others), we also maintain our own internal controls and undergo external audits where appropriate. We strive to be transparent about our security, privacy, and operational measures so that customers are confident in meeting their compliance obligations while using NetSpeak. If your organization has specific governance or data residency requirements, our solutions consultants can collaborate on a deployment model (public, private, or hybrid cloud) that meets those needs without sacrificing security.

Third-Party Integrations

To enhance and extend NetSpeak's capabilities, we integrate with certain third-party services or allow customers to connect their own SaaS platforms. Each of these integrations is protected by secure authentication methods (such as OAuth, token-based access, or certificate-based trust). We carefully review third-party security documentation, API usage patterns, and data flow diagrams to ensure they meet our standards for safety, encryption, and data handling.

Wherever possible, we encourage separation of duties so that each integration only has the privileges necessary for the tasks at hand. In addition, every integration is subject to the same logging and auditing process that applies to NetSpeak's internal components.

Secure Development Lifecycle

Our engineering teams follow a secure development lifecycle (SDL) that includes threat modelling, peer reviews, static and dynamic testing, and frequent penetration tests by external specialists. We believe that security cannot be just a final checkbox before release but rather a continuous practice that starts with end-to-end solution design. Key SDL highlights include:

- **Threat Modeling and Requirements:** Early in a project, we identify potential abuse cases and define security requirements alongside functionality.
- **Automated Testing:** We embed static analysis and dependency scanning into our CI/CD pipelines, ensuring any vulnerabilities in libraries or code are flagged immediately.
- **Manual Review:** For sensitive code changes—especially those dealing with authentication, encryption, or data handling—team members with security expertise perform a deeper manual review.
- **Regular Assessments:** Periodic third-party penetration tests help us discover blind spots or newly emerging threats so we can remediate them quickly.

Conclusion

Security is at the heart of NetSpeak’s mission to deliver a reliable, trustworthy platform for automation and AI-driven assistance. We have designed our technology stack around leading security practices to ensure the confidentiality, integrity, and availability of customer data. Although the specifics of our cloud deployment may evolve over time, our commitment to rigorous security standards does not.

We believe in transparency, constant innovation, and active customer collaboration to meet their unique regulatory and operational needs. If you have questions about any of the topics in this document or wish to learn more about how NetSpeak can address particular compliance scenarios, our Solutions Consulting team is happy to provide additional information.

This overview is provided for informational purposes only and describes our general security practices. As NetSpeak evolves with new features or infrastructure improvements, the fundamental principles outlined here remain consistent: robust security, privacy, and responsible innovation.



© 2025 NetSpeak Inc. All Rights Reserved.

Document Version: 1.5

Last Updated: October 2025