LeanDNA Data and Security Policy for AI

Executive Summary

LeanDNA is committed to maintaining the highest security and governance standards to protect customer data used by its AI and ML systems. This policy explains that customer data remains the property of the customer, is never used for training publicly accessible Large Language Models (LLMs), and is strictly processed within secure cloud environments. All AI components are governed by the same rigorous, SOC 2 Type II certified platform-wide security principles, including data encryption, access controls, and technical safeguards against AI-specific threats.

1. Introduction and Purpose

LeanDNA understands the importance of security and is committed to maintaining the highest standards to ensure that customer data is fully protected and managed with strict governance controls. This policy establishes the minimum requirements for the secure development and operation of Artificial Intelligence (AI) and Machine Learning (ML) systems and combines with the overall security policy to cover all aspects of the platform. LeanDNA continues to invest in security infrastructure and processes to ensure that all AI capabilities are safe and secure.

2. Key Definitions

To ensure clarity, the following terms and acronyms used in this policy are defined:

- **LLM (Large Language Model):** A type of general-purpose Artificial Intelligence (AI) model that is capable of understanding and generating human-like text.
- **ML (Machine Learning) Model:** A mathematical model that is trained on a specific dataset to recognize certain patterns. Unlike general-purpose models, this pre-trained model is used to make highly targeted decisions or predictions (inference) for a dedicated task without explicit programming.
- **VPC (Virtual Private Cloud):** A virtual, isolated section of the Amazon Web Services (AWS) cloud where LeanDNA hosts its infrastructure. It provides a highly secure and private environment for processing and storing customer data.
- **Inference:** The specific stage during which an already-trained Machine Learning (ML) model processes new, unseen input to generate a prediction, insight, or response. Data used during inference is strictly for prediction, not for model training.
- **Principle of Least Privilege (PoLP):** An established security concept that requires every user, account or system be granted only the minimum access rights necessary to perform its authorized function, and nothing more.
- **MFA (Multi-Factor Authentication):** A security process that requires a user to provide two or more pieces of evidence (factors) to gain access to a resource.
- **SOC 2 Type II:** An independent audit and certification that verifies a service organization's controls are designed and operating effectively to meet the Trust Services Criteria (e.g., Security and Confidentiality) over a specified period.

3. Data Governance and Customer Rights

This section establishes the definitive rules for the use, access, and ownership of all customer data utilized by LeanDNA AI.

3.1. Data Ownership and Usage Principles

LeanDNA infrastructure and application stack are hosted on **Amazon Web Services (AWS)**. All customer data is processed, stored, and protected with encryption within **isolated virtual private clouds (VPCs)** under LeanDNA's exclusive control, strictly **prohibiting its transfer outside** of this secured cloud environment.

Customer **owns the source data, user inputs, and the resulting outputs** generated by LeanDNA AI. Conversely, LeanDNA retains all rights and ownership in the underlying LeanDNA AI platform and its proprietary Large Language Model (LLM) technology. LeanDNA guarantees that **no customer data will be utilized for the training, fine-tuning, or improvement** of any publicly accessible or third-party LLMs.

3.2. Data Access Control

Access to and review of Customer data by LeanDNA AI is strictly controlled by the **site access permissions** that the individual user is authorized to access within the LeanDNA platform. All data for authorized sites is accessible to LeanDNA AI.

Customer data access for LeanDNA personnel is governed by the **Principle of Least Privilege (PoLP)**. Access is **governed by SOC 2 controls**, is strictly role-based, and is necessary only for the completion of a specific, authorized task. **Multi-Factor Authentication (MFA)** is required and only secure, auditable connections are permitted for access.

3.3. Data Retention

Customer data is retained throughout the term of the active contractual and business agreement. Upon termination, LeanDNA is committed to the **secure and permanent deletion** of all customer data to be completed within a **30-day period**.

3.4. Operational Logging and Monitoring

LeanDNA maintains **secure, comprehensive, and auditable logs** to ensure the integrity, reliability, and security compliance of the LeanDNA AI platform. This log data is used **exclusively** for troubleshooting, internal product operations, usage analysis and improving platform quality and performance following the data access control requirements.

4. Platform-Wide Security Principles

LeanDNA ensures that the security posture for all AI components is consistent with the rigorous standards applied to the entire LeanDNA platform. The **same governance principles** for access control, incident response, vulnerability management, and auditability applicable to the entire LeanDNA platform are extended to all AI systems and their underlying infrastructure.

All data is **encrypted in transit and at rest** using the industry-leading standards. LeanDNA maintains **SOC 2 Type II certification** for the **Security and Confidentiality** principles, demonstrating the continuous effectiveness of internal controls.

LeanDNA implements and maintains safeguards against prompt injection and other AI-specific threats through technical controls, including **prompt sanitization, context filtering, and access validation layers**. Furthermore, LeanDNA is committed to developing and integrating new strategies to combat emerging adversarial attacks.

5. Security Controls by Type

LeanDNA defines distinct security requirements based on whether **traditional Machine Learning (ML) models** or **Large Language Models (LLMs)** are used in the platform.

5.1. Traditional Machine Learning (ML) Models

LeanDNA leverages **traditional ML models** built for advanced statistical analysis and hosted in **isolated Virtual Private Clouds (VPCs)** under its exclusive control. These **small-scale models** are inherently distinct from Large Language Models (LLMs) and are utilized to deliver highly **targeted product capabilities** for individual tasks and decisions.

LeanDNA trains these mathematically-founded models using **anonymized and aggregated customer data**. This strategy enables the models to provide **valuable insights**, generate **accurate predictions** (inference), and identify **actionable optimization opportunities** specific to the customer's operations.

5.2. Large Language Models (LLMs)

LeanDNA utilizes third-party LLMs through **AWS Bedrock** and applies enhanced controls to guarantee data security. All model invocation occurs within an **isolated, secure AWS environment** controlled by LeanDNA, ensuring that customer data **is never sent to the model provider for any purpose**. Furthermore, LeanDNA maintains **assurance of security controls** by relying on and reviewing the results of independent security and compliance assessments conducted on AWS Bedrock.