Security and Resilience of an AI Supercomputer

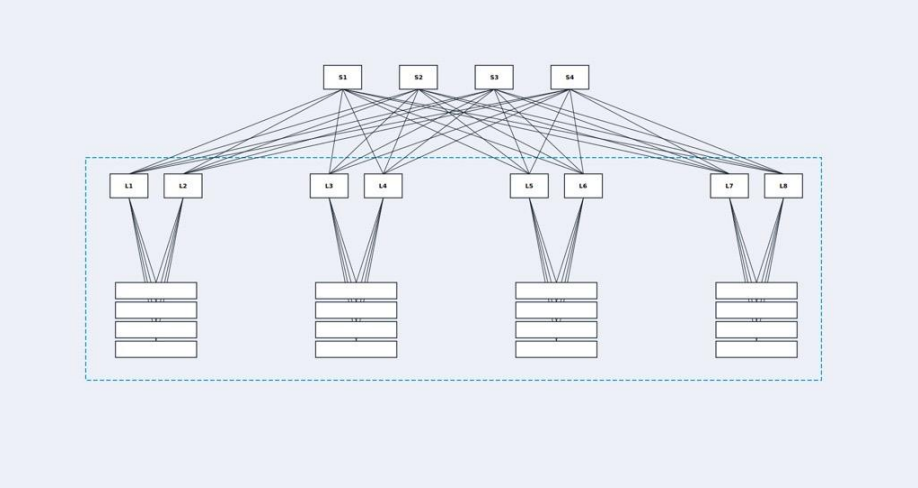

© PHOENIQS 1

The PHOENIQS AI Supercomputer has been designed from the ground up with security, resilience, and continuous availability in mind.

At the core of the architecture is a highly redundant **leaf-spine network design**, where GPU servers connect to multiple **leaf switches (L)**, and each leaf switch connects to multiple **spine switches (S)**. This creates numerous independent communication paths between compute, storage, and network resources, ensuring that no individual server, cable, network interface, or switch becomes a single point of failure. If any component experiences an issue, traffic is automatically rerouted through alternative paths, allowing AI workloads to continue operating without interruption.

Resilience extends across the compute and storage layers. GPU servers are distributed across multiple physical racks, while storage systems are connected through multiple network paths to ensure continuous access to training data, AI models, and checkpoints. This architecture provides the high availability required for enterprise and public-sector AI workloads, where downtime can impact business-critical processes, research activities, and production AI services.

To further strengthen security and digital sovereignty, the PHOENIQS AI Supercomputer can be integrated with dedicated **Hardware Security Modules (HSMs)** to protect cryptographic keys, secrets, and sensitive credentials. HSMs store encryption keys within tamper-resistant hardware, enabling advanced security models such as customer-controlled encryption keys (KYOK/BYOK), secure model protection, and regulatory compliance. Enterprise-grade HSMs are designed to withstand both physical and logical attacks and can automatically and irreversibly destroy protected cryptographic keys if tampering is detected, ensuring that sensitive information cannot be extracted even if hardware is stolen.

Combined with Swiss hosting, sovereign operations, and enterprise-grade infrastructure, PHOENIQS delivers a secure and highly available platform for training, fine-tuning, and running large-scale AI models while maintaining full control over data, models, and workloads.