

White paper

Sharon AI and VAST Data

Solving the data challenge
behind Sovereign AI

Table of Contents

Executive Summary	3
Introduction	4
Sharon AI + VAST Data	5
Delivering Sovereign AI with VAST Data	5
VAST AI OS Foundational Capabilities	6
Conclusion	7
Learn how Sharon AI and VAST Enable Sovereign AI	7

Executive Summary

AI priorities have shifted. Regional governments, regulated industries, and enterprises now prioritise control over sensitive data, intellectual property, and critical infrastructure systems. In this climate, the goal is to safeguard these assets against foreign legal claims and cross-border data access.

Sovereign AI is the ability for a nation or regulated enterprise to develop, operate, and govern AI systems within its own jurisdiction, legal framework, and security controls. It is not merely an “AI stack deployed domestically.” It needs to demonstrate control over data, models, infrastructure, identity management, operational processes, and AI-generated outputs that all form the AI lifecycle.

The distinction matters, as AI generates many forms of sensitive data beyond the original corpus, including prompts, metadata catalogues, retrieved context, embeddings, vector indexes, checkpoints, inference logs, and outputs. While a primary dataset may remain within a sanctioned geographic area, data sovereignty fails the moment these derived artefacts or operational logs cross jurisdictional boundaries.

The Sovereign AI Challenge: Modern AI Pipeline Complexity

As organisations move from simple content generation to complex, agentic AI workflows capable of planning, retrieving data, executing tools, and persisting memory, traditional infrastructure models fall short. This creates three critical operational bottlenecks:

- **Hidden Artefact Leakage:** Data sovereignty is rarely broken by primary raw datasets; it fails across derived assets. Modern AI continuously generates secondary artefacts, prompts, vector embeddings, retrieved context, checkpoints, and inference logs. When standard clouds move these assets across borders or third-party services, regulatory boundaries are breached.
- **Loss of Audit Lineage & Control:** Activating proprietary enterprise data requires copying and transforming information across disparate vector databases and GPU clusters. This data movement obscures where sensitive records reside and how models interact with them, making compliance auditing nearly impossible.
- **Regulatory Fragmentation:** Across the Asia-Pacific region, regulations range from strict legislative acts to complex content and algorithm registration rules. Legacy architectures force organisations to manage security and policy manually across isolated silos, creating massive compliance overhead.

Together, Sharon AI and VAST Data address this sovereignty challenge. Through our strategic partnership, we offer purpose-built, high-performance GPU infrastructure tailored for organisations with strict compliance requirements, keeping sensitive data and AI workloads within our sovereign jurisdiction. Coupled with the VAST AI Operating System (VAST AI OS), our sovereign supercluster runs real-time reasoning and AI workflows on a unified data platform. Organisations know where their data is stored, where models execute, who operates the infrastructure, and which jurisdiction governs their information.

Introduction

Enterprises in regulated industries, regional governments and institutions increasingly want the ability to create artificial intelligence using local computing infrastructure, datasets, talent, business ecosystems, languages, and intellectual property rather than being completely dependent on foreign platforms or models. It permits nations to codify and reflect their own language, culture, history, and societal values in AI outputs, bypassing reliance on external narratives. NVIDIA describes sovereign AI as a nation's capability to produce and own its artificial intelligence using its own domestic infrastructure, data, workforce, and business networks.

Sovereign AI is not just an Australian concern. Across the Asia-Pacific region, jurisdictions are advancing their own AI strategies in different directions. South Korea's AI Basic Act took effect in January 2026. China already requires algorithm registration and the labelling of AI-generated content. Others, including Japan, Singapore, and India, rely largely on voluntary or sector-specific frameworks. This compounds the sovereignty challenge; the same workload can satisfy one market's requirements and fall short of another's.

Sovereign AI requires a shift in thinking. Instead of treating each domestic cloud, GPU cluster, or AI service as an independent governance domain, organisations need a sovereign AI provider that has built policy enforcement directly into the data layer. The goal is for governance to travel with data. This requires consistent management of identity, access, encryption, auditing, and lineage semantics across the environments in which data is stored and processed.

Sharon AI + VAST Data

Sharon AI and VAST Data are designing that together. The VAST AI OS brings unified storage, database services, global data access, event-driven processing, vector search, AI data pipelines, and governance within a common framework inside Sharon AI. At its foundation, organisations can govern access, reduce unnecessary data movement, and support AI-ready operations.

Governments and regulated industries keep authoritative data under organisational control and selectively make governed data available to compute when and where appropriate.




Figure 1. How VAST Enables Sovereign AI

Delivering Sovereign AI with VAST Data

Together, we treat sovereignty as a control-plane problem solved directly at the data layer. Powered by the VAST Disaggregated Shared Everything (DASE) architecture, the VAST AI OS replaces fragmented environments with a single, high-performance data foundation.

By abstracting metadata into the VAST Catalog, regulated industries can streamline data access and preparation. It acts as a centralised, searchable index that organises the data on a cluster using metadata, tags, and user-generated annotations. In clusters with multiple tenants, data access is restricted to queries made inside the tenant. As a result, data engineers can quickly find relevant data they are authorised to see, without sifting through siloed datasets.

As Table 1 shows, VAST AI OS inside Sharon AI is not your typical storage solution, it combines file, object, block, structured data, vector information, and AI processing into a common architecture, collapsing what are usually distinct products with individual security boundaries into a secure, trusted unified solution. Data governance improves because vector search aligns with file-system permissions. AI engineers and developers can only retrieve AI results from data they are authorised to access. Regional governments and regulated industries can then ingest data, query vectors, fine-tune models, and run inference securely on the same platform.

VAST AI OS Foundational Capabilities

Products	Sovereign AI relevance
VAST DataStore	The persistent data foundation supporting file, object, block, and table elements.
VAST DataBase	A database architecture for structured data, metadata, analytical information, vectors, streams, catalogues, and AI context.
VAST DataSpace	A globally distributed data layer that extends data access across VAST environments while maintaining a unified namespace.
VAST DataEngine	An event-driven compute framework designed to execute processing near data and orchestrate data-driven workflows.
VAST SyncEngine	Discovers, indexes, and synchronises external data sources into the data environment.
VAST InsightEngine	An AI data-processing framework for transforming unstructured information into searchable and AI-ready context, including vector embeddings and retrieval workflows.
VAST AgentEngine	An execution and orchestration environment designed for production AI agents and persistent AI workflows.

Using VAST DataSpace, companies can strike a balance between data sovereignty and the availability of Sharon AI GPUs. When data is first accessed, DataSpace caches it close to the compute and prefetches related data to accelerate subsequent reads. By controlling caching behaviour and policies related to global folders, organisations can minimise data movement and generate the precise audit evidence required to prove operational controls.




Figure 2. VAST DataSpace to Sharon AI

This approach keeps sensitive information secured within approved locations, while providing Sharon AI GPU clusters with fast, governed access through a unified global namespace.

Conclusion

As AI systems grow more dependent on proprietary data, data sovereignty becomes a core requirement. Together, Sharon AI and VAST offer a unified data foundation that allows organisations to keep strategic data under their control while securely connecting it to GPU clusters. Governance, identity, data access, and auditability stay linked to the data, helping organisations innovate across distributed infrastructure without compromising sovereignty.

Learn how Sharon AI and VAST enable Sovereign AI Sharon



One Data Fabric

Sovereign AI workloads run where compute capacity is available while maintaining control of data and security policies.



Maximum Utilisation

GPUs can access and load models from VAST storage with performance consistent with local NVMe devices.



Consistent Policies

VAST DataSpace enforces a single trust zone across environments. Data governance remains consistent everywhere.

Sharon AI and VAST Data deliver the foundation for sovereign AI, enabling governments, enterprises, and regulated industries across Asia-Pacific to build AI infrastructure without sacrificing scale or speed.

To explore further, visit

[Sharon AI](#) | [VAST Data and Sovereign AI](#) | [VAST DataStore](#) | [VAST DataSpace](#)

About Sharon AI

Sharon AI (NASDAQ: SHAZ) is a leading Australian neocloud delivering trusted sovereign AI infrastructure. Through its AI Factory platform and world-class ecosystem of technology and co-location partners, Sharon AI expands access to the scalable capabilities that organisations need to build, train, and run AI, from model training through to inference and agentic AI. Serving customers globally, Sharon AI helps organisations move faster from AI potential to measurable value. For more information, visit www.sharonai.com