

RADIANT

WHITEPAPER

Build AI Confidently with a **Security-First** Cloud

Contents

Introduction	1
Weighing Cloud Security: Public vs Sovereign Cloud	2
How the Radiant Cloud Fortifies your AI	3
Isolated and Secure Infrastructure	4
Secure Multi-tenancy for Each Team	4
Single-Pane-of-Glass Monitoring	4
RBAC and Granular Security Controls	5
Built-in Traceability	5
Observability & Resource Tracking for Optimization	5
Checklist for Enterprise Decision Makers	6
Secure your AI journey	7

Introduction

Artificial intelligence (AI) is becoming a core driver of enterprise innovation, but security-first AI infrastructure is essential to protect sensitive data and intellectual property. Technology leaders are increasingly recognizing the need for secure cloud environments for AI, keenly aware that any breach or compliance failure in an AI deployment can be devastating.

82%

of data breaches were tied to data stored in cloud environments.

75%

of CISOs cite gaining control over cloud data as their top priority for managing risk.

Up to **4%**

of Global Turnover as fines for GDPR infractions

The risks of insecure cloud AI deployments include data breaches, IP theft, non-compliance fines, and even compromised AI models. To mitigate these risks, organizations are increasingly seeking isolated, compliant environments for their AI, which is why 79% of enterprises report implementing internal clouds. The Radiant Cloud is the top choice for enterprise AI due to its robust data security, complete control, and flexibility.

The Radiant Cloud is the vertically integrated, security-first alternative to traditional public clouds for AI workloads. Ranging from air-gapped environments to secure multi-tenancy for your internal teams and robust enterprise controls, Radiant makes it easy to build your AI securely and meet the growing security and operational independence needs of your business.

Weighing Cloud Security: Public vs Sovereign Cloud

Not all cloud environments are equal when it comes to security. Public clouds and Sovereign Clouds each have different risk profiles and benefits. Below is a comparison of key security aspects, highlighting how a Sovereign Cloud can eliminate many security and governance risks associated with public cloud deployments:

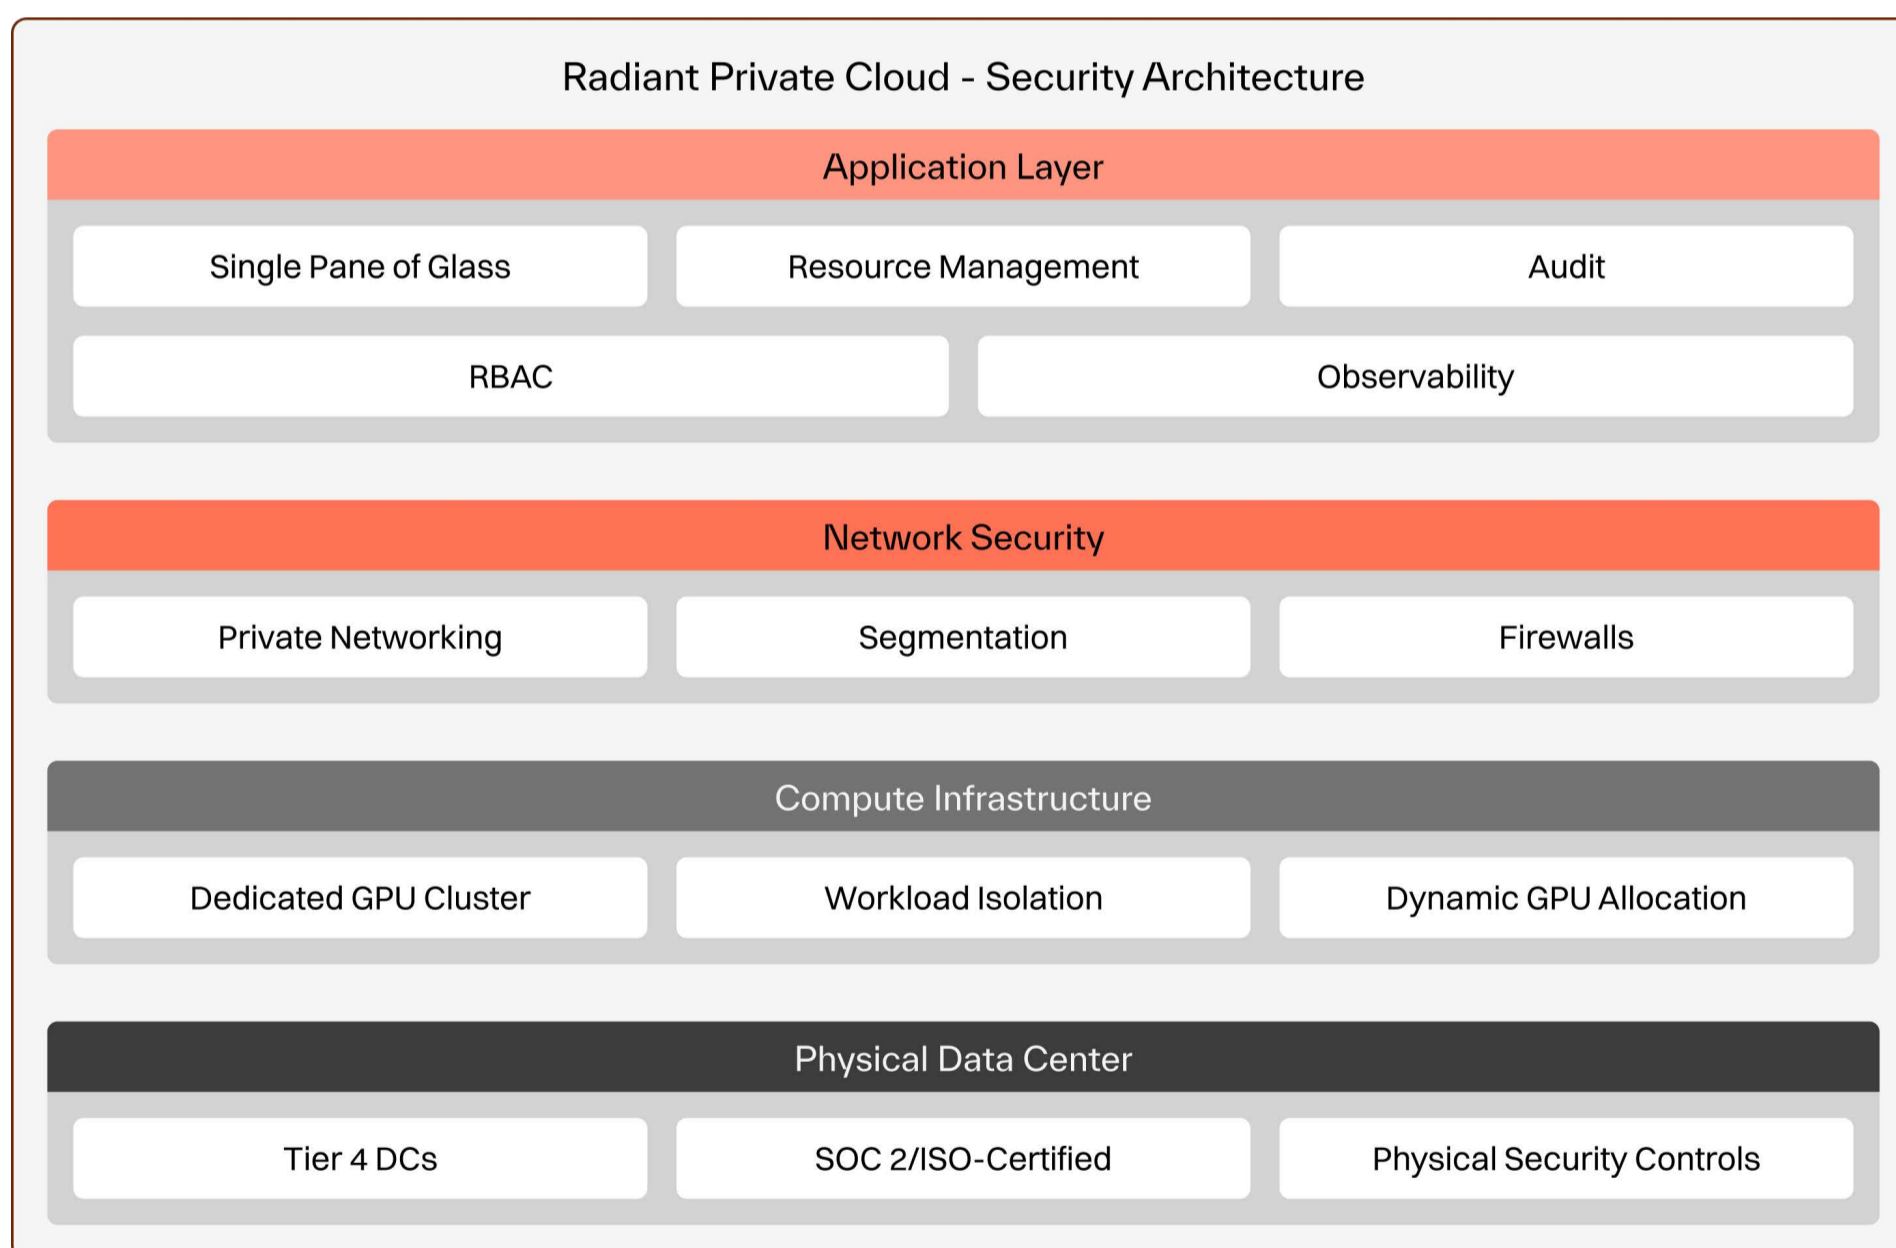
	Public Cloud	Sovereign Cloud
Ownership & Control	Infrastructure is on shared, resources and platform managed by a provider with limited direct control for the customer	Infrastructure is dedicated to one organization and vertically integrated from power to code, giving the enterprise direct oversight and control
Isolation of Resources	Shared environment with other tenants; potential for “noisy neighbors” and cross-tenant vulnerabilities if isolation fails. Data and workloads co-exist with others, raising concerns about data leakage between tenants.	Fully isolated environment with dedicated hardware and network segments for one enterprise, eliminating noisy-neighbor issues and vastly reducing attack surface. Greater ability to implement custom network segmentation and zero-trust architecture.
Compliance & Data Privacy	Must rely on the provider's controls and certifications. It's a shared responsibility to meet regulations – the cloud vendor may be certified, but the customer must ensure their use of the cloud is compliant. Data may reside in global regions by default, raising sovereignty concerns.	Enhanced control over data location and handling , making it easier to comply with strict regulations (GDPR, HIPAA, etc.). Enterprises can choose where data is stored and processed. The environment can be tailored to specific compliance needs (e.g. keeping all data in-country and under specific encryption policies).
Security Responsibility	Follows a shared responsibility model – the provider secures underlying infrastructure (physical, hypervisor) while the customer secures the OS, apps, and data. Misconfigurations by users (e.g. open storage buckets) are a common risk. Security updates for managed services depend on the provider's schedule.	Enterprise (and/or its Sovereign Cloud provider) has end-to-end responsibility for security. With a Sovereign Cloud like Radiant's, the dedicated infrastructure is maintained securely for the client, and the customer retains ultimate control over how data and workloads are secured and patched. This allows a more proactive security posture, everything from network to application can be locked down to the organization's standards.
Visibility & Governance	Relies on provider's tooling and APIs for monitoring. Visibility can be limited – e.g. only high-level logs or abstracted metrics are available, and multi-tenant setups mean less insight into underlying systems. Forensics and troubleshooting can be challenging in an ephemeral, shared environment.	Deep visibility into the stack – full access to logs, network data, and system metrics since the cloud is single-tenant. Enterprises can integrate the Sovereign Cloud with their own SIEM, monitoring, and governance tools for a single-pane-of-glass view. This comprehensive observability simplifies auditing and incident response, as every action and flow in the environment is trackable by the organization.

Sovereign Clouds deliver the non-negotiable structural control and operational independence at the cost of additional management overhead. In contrast, public clouds provide convenience and rapid scalability but require trusting a third-party platform and diligently managing shared security responsibilities. Notably, using a public cloud means closely watching security configurations, since the slightest oversight (e.g. misconfigured access rules) could expose data to the Internet.

On the compliance side, Sovereign Clouds are the preferred choice for keeping sensitive data local and achieving strict regulatory and national requirements, whereas public clouds might offer more built-in services that ease daily management but could complicate compliance audits. In summary, for enterprises and nations with low risk tolerance and high governance needs, the Radiant Cloud eliminates the security risks inherent in public cloud services by delivering true operational independence.

How the Radiant Cloud Fortifies your AI

The Radiant's Cloud is purpose-built to address enterprise AI security priorities, providing a secure and controlled environment without sacrificing performance.



Here's how the Radiant Cloud drives a security-first approach to running AI workloads:

Isolated and Secure Infrastructure

The Radiant Cloud runs on dedicated infrastructure for each customer. Your AI workloads operate on hardware exclusively allocated to your organization, ensuring complete isolation from any other tenant or workload. This dedicated approach means no shared resources, no noisy neighbors, and strong isolation at the compute, storage, and network layers - dramatically reducing the risk of cross-tenant data leakage or interference. Enterprises gain full control over their AI environment: you decide where data resides and how it's secured, meeting data sovereignty requirements.

Moreover, Radiant's data centers meet rigorous security standards (SOC 2 and/or ISO 27001 certified), so the physical infrastructure and facilities are audited and compliant. In short, the Radiant Cloud gives you a cloud that's yours alone - tightly secured from the ground up with full control over your most sensitive AI data and models.

Secure Multi-tenancy for Each Team

Designed for distributed teams, the Radiant Cloud enables every team to innovate in an isolated workspace while maintaining central oversight. By providing distinct workspaces for each team (a secure, multi-tenant architecture for internal teams), organizations can segregate projects and data per team.

Each team's data, models, and pipelines are isolated from the ground up, at the software, network and compute levels. This data isolation guarantees that even though teams within the organization share the underlying cloud platform, their data does not accidentally mix. One team cannot access another team's datasets or models appropriate privileges. Isolation not only protects sensitive information but also reduces the risk of cross-team errors and aids compliance with data sovereignty rules.

Single-Pane-of-Glass Monitoring

Radiant provides a unified monitoring and management interface that offers real-time **visibility into your AI workloads spanning teams across your organization.** This "single pane of glass" approach means your team can monitor GPU/CPU utilization, memory, storage I/O, and network performance for every AI job or service.

Centralized logs and metrics allow security and ops teams to watch for anomalies or unauthorized activities across the entire AI stack. Because the Radiant Cloud environment is fully under your control, you can integrate with your existing security monitoring tools - by feeding Radiant's telemetry into your SIEM or analytics platforms to extend enterprise-wide observability.

The result is end-to-end situational awareness: from model training runs to inference endpoints, you have real-time insights at your fingertips. This level of visibility not only helps detect potential security issues early but also aids in optimizing AI performance by catching bottlenecks or inefficiencies quickly.

RBAC and Granular Security Controls

The Radiant Cloud includes robust Role-Based Access Control (RBAC) capabilities and fine-grained permission settings. You can define roles and assign privileges so that users only access the data and functions they require. This principle of least privilege reduces risk by minimizing unnecessary access. For example, a data scientist might deploy and monitor models but not have rights to alter core infrastructure settings. Radiant's platform enforces these controls consistently across all services.

In addition, Radiant supports two-factor authentication, ensuring that only authenticated, authorized personnel can launch AI workloads or view sensitive data. These granular controls form the backbone of AI security governance - together with monitoring, they allow enterprises to confidently share an AI platform among teams without fear of unchecked access.

Built-in Traceability

To help customers meet regulatory and internal compliance requirements, the Radiant Cloud provides comprehensive audit and compliance tooling. Every action in the environment - from data uploads to model deployments, from user logins to configuration changes - can be logged and tracked.

Audit logs help keep a record of all activities, showing who accessed what and when. This level of traceability creates a verifiable chain of custody for data and models, which is invaluable for compliance audits and incident investigations.

Radiant's compliance-oriented features don't stop at logging; the platform is designed in alignment with major security standards. Radiant undergoes annual ISO 27001 audits and SOC 2 compliance, reflecting a commitment to security best practices. For customers, this means the Sovereign Cloud environment inherently meets high security benchmarks, and built-in tools like compliance dashboards or reporting templates can simplify aligning your usage with frameworks like GDPR or industry-specific rules.

Observability & Resource Tracking for Optimization

Radiant's platform offers deep observability into AI workloads and resource utilization, which not only bolsters security but also helps optimize performance and costs. Through detailed metrics and resource tracking, you can see how every model training or inference job consumes GPUs, CPUs, memory, and storage. This transparency enables optimization of AI workloads: teams can identify under-utilized resources, spot performance anomalies, and right-size their clusters for efficiency. In practice, such observability leads to optimal training times and higher resource utilization, translating to cost savings via greater efficiency.

Real-time observability thus serves a dual purpose - strengthening security (through continuous oversight of system behavior) and driving cost-effective scaling of AI by ensuring you're getting the most out of your dedicated infrastructure. With the Radiant Cloud's rich telemetry, enterprise AI teams can confidently innovate, knowing they have the data to keep the environment both secure and optimized.

Checklist for Enterprise Decision Makers

When evaluating a security-first cloud provider for AI workloads, enterprise decision-makers should consider the following checklist. Use these criteria to ensure any solution will meet your organization's security and compliance requirements:

- Compliance & Certifications:** Does the provider demonstrate adherence to industry security standards? Verify they have relevant certifications (e.g. ISO 27001, SOC 2, GDPR readiness) and can support your compliance needs (such as data residency or specific regulatory frameworks in your industry).
- Dedicated Infrastructure:** Confirm that the cloud environment offers private, isolated resources for your workloads. True isolation means your data and computations won't share hardware with other customers, greatly enhancing security.
- Secure and distinct workspaces for each team:** Does the cloud provide a multi-tenant access for your teams to segregate projects and data per team?
- Flexibility to Bring your own Compute:** If you already own Compute or prefer your Compute for certain projects, consider a cloud that lets you to Bring your Own Compute (BYOC) or integrate it with the provider's compute infrastructure.
- Access Controls & Data Protection:** Evaluate the strength of identity and access management. Look for robust RBAC (role-based access control) to restrict user permissions and integration with your Identity Provider (SSO/ MFA). The provider should enforce least privilege and protect sensitive data through every layer. The provider should enforce least privilege and protect sensitive data through every layer.
- Monitoring & Visibility:** Ensure the solution provides real-time monitoring, logging, and observability into your AI processes. You should have a **single pane of glass** to oversee all AI jobs, with alerts for any anomalies. This visibility is key to quickly detecting issues like unauthorized access or data misuse in AI workflows.
- Audit Logging & Reporting:** Check that comprehensive audit logs are available and easily accessible. Built-in logs and reporting will simplify proving adherence to policies and investigating incidents if they occur.

This checklist will help you identify a cloud partner that treats security as a first-class priority, not an afterthought. A provider like Radiant, which meets all these criteria, can empower your enterprise to pursue AI initiatives without compromising on protection.

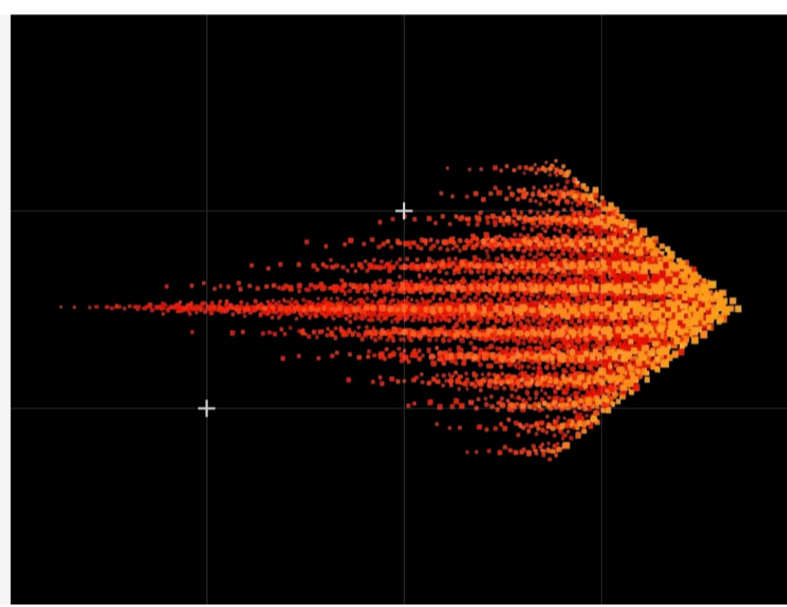
Secure your AI journey

Enterprises must prioritize a security-first cloud infrastructure for AI to protect their data, ensure compliance, and build AI solutions with confidence. AI initiatives can deliver immense value, but without the right security foundation, they also introduce new risks. A security-first AI cloud means you can accelerate innovation without compromising on governance or peace of mind.

Build AI confidently knowing that your models and sensitive data are running in a hardened environment designed for enterprise needs. The Radiant Cloud delivers the performance and scalability for AI, plus the security and control that public clouds often lack. It empowers your organization to harness AI's full potential while keeping your risk low and your compliance on track.

Don't leave the security of your AI to chance. Take the next step towards building AI confidently with the **right cloud foundation**.

TALK TO A SOVEREIGN CLOUD EXPERT



About Radiant

Radiant is building the utility model for compute - ubiquitous, always on and offering superior economics. Through vertical integration of capital, powered land, compute, and software, we deliver sovereign-grade control for nations, turnkey certainty for enterprises, and limitless capacity for builders. This model will democratize AI, equipping every individual with the ability to unlock their creative potential and tackle world-changing problems.

LEARN MORE AT [RADIANT.CO](https://radiant.co)

RADIANT

We Are Powered Intelligence

