Crusoe

# The multicloud playbook for AI builders

CONTENTS

# The cloud infrastructure playbook that served the Web 2.0 era is no longer sufficient for the AI revolution.

While hyperscalers remain essential for general-purpose application hosting and storage, the unique demands of AI and high-performance computing (HPC) require a more specialized approach.

For AI-native companies and scaling enterprises, the future is not about choosing a single provider, but about adopting a strategic multicloud AI architecture. This paper outlines the case for a multicloud solution: maintaining your application and data layers on hyperscalers while offloading intensive training and inference to a specialized AI infrastructure provider like Crusoe Cloud.

# 1

# The availability imperative: Overcoming capacity constraints

For the last decade, consolidation was the goal for CIOs seeking simplified management. However, AI workloads, specifically large language model (LLM) training and high-throughput inference, require massive, immediate scale.

In the AI race, speed is defined by “time-to-compute.” The difference between starting a project today versus waiting months for a hardware allocation is the difference between market leadership and obsolescence.

## Hyperscaler

### The hyperscaler reality: Generalist infrastructure

Hyperscalers are architected to serve millions of workloads across every industry. This breadth is their strength, but it creates tradeoffs for specialized use cases. AI workloads (particularly LLM training and high-throughput inference) have unique requirements around GPU density, network topology, and power delivery that general-purpose infrastructure wasn't originally designed to prioritize.

## Crusoe Cloud

### The Crusoe approach: Purpose-built infrastructure

Crusoe delivers vertically-integrated and purpose-built infrastructure for AI. This means we build data centers engineered specifically for the demanding requirements of frontier computing and have a commercially-available Cloud platform that provides the infrastructure backbone for AI workloads.

**Best-of-breed architecture:**
Unlike hyperscalers that offer general solutions across their massive portfolios, Crusoe partners with category leaders (from NVIDIA and AMD for cutting-edge silicon to VAST Data for storage and networking needs), to deliver optimized performance at every layer. We're also building next-generation data centers from the ground up, where our team engineers facilities like Abilene with hardware, software, power, and cooling systems co-designed to maximize AI performance.

**Strategic takeaway:**

High-growth teams are gaining a competitive advantage by partnering with specialized providers to secure GPU infrastructure that can be deployed faster than industry-standard timelines.


# 2

## De-risking AI scaling: The economic advantage

For AI-native organizations, compute is often the single largest driver of burn rate. The economics of your cloud choice directly impact your runway.

## Hyperscaler

### The “everything store” tax

Hyperscalers operate on a model that supports all workloads, which requires maintaining hundreds of legacy services and complex support and pricing tiers. While this is critical to business infrastructure, this legacy overhead is baked into every GPU hour. Many providers also utilize egress and ingress fees to create financial barriers that lock your data inside their ecosystem.

## Crusoe Cloud

### The Crusoe approach: Focused on AI performance

Crusoe Cloud infrastructure is designed for high-performance computing, stripped of bloat to focus entirely on raw performance and speed.

**Optimized pricing:**
By removing the overhead of extraneous legacy services, we work to offer high-performance GPU instances at a competitive price.

**Zero-friction portability:**
We believe your data belongs to you. Crusoe charges zero ingress and zero egress fees, providing the architectural freedom to move training data as needed.

**Strategic takeaway:**

Specialized AI infrastructure providers deliver both economic and operational advantages — reducing compute costs while giving you the architectural freedom to move workloads as your needs evolve.

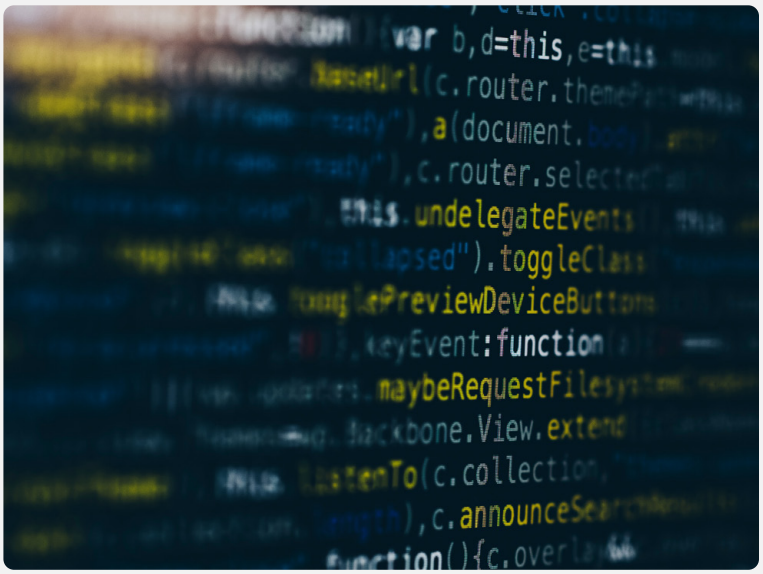

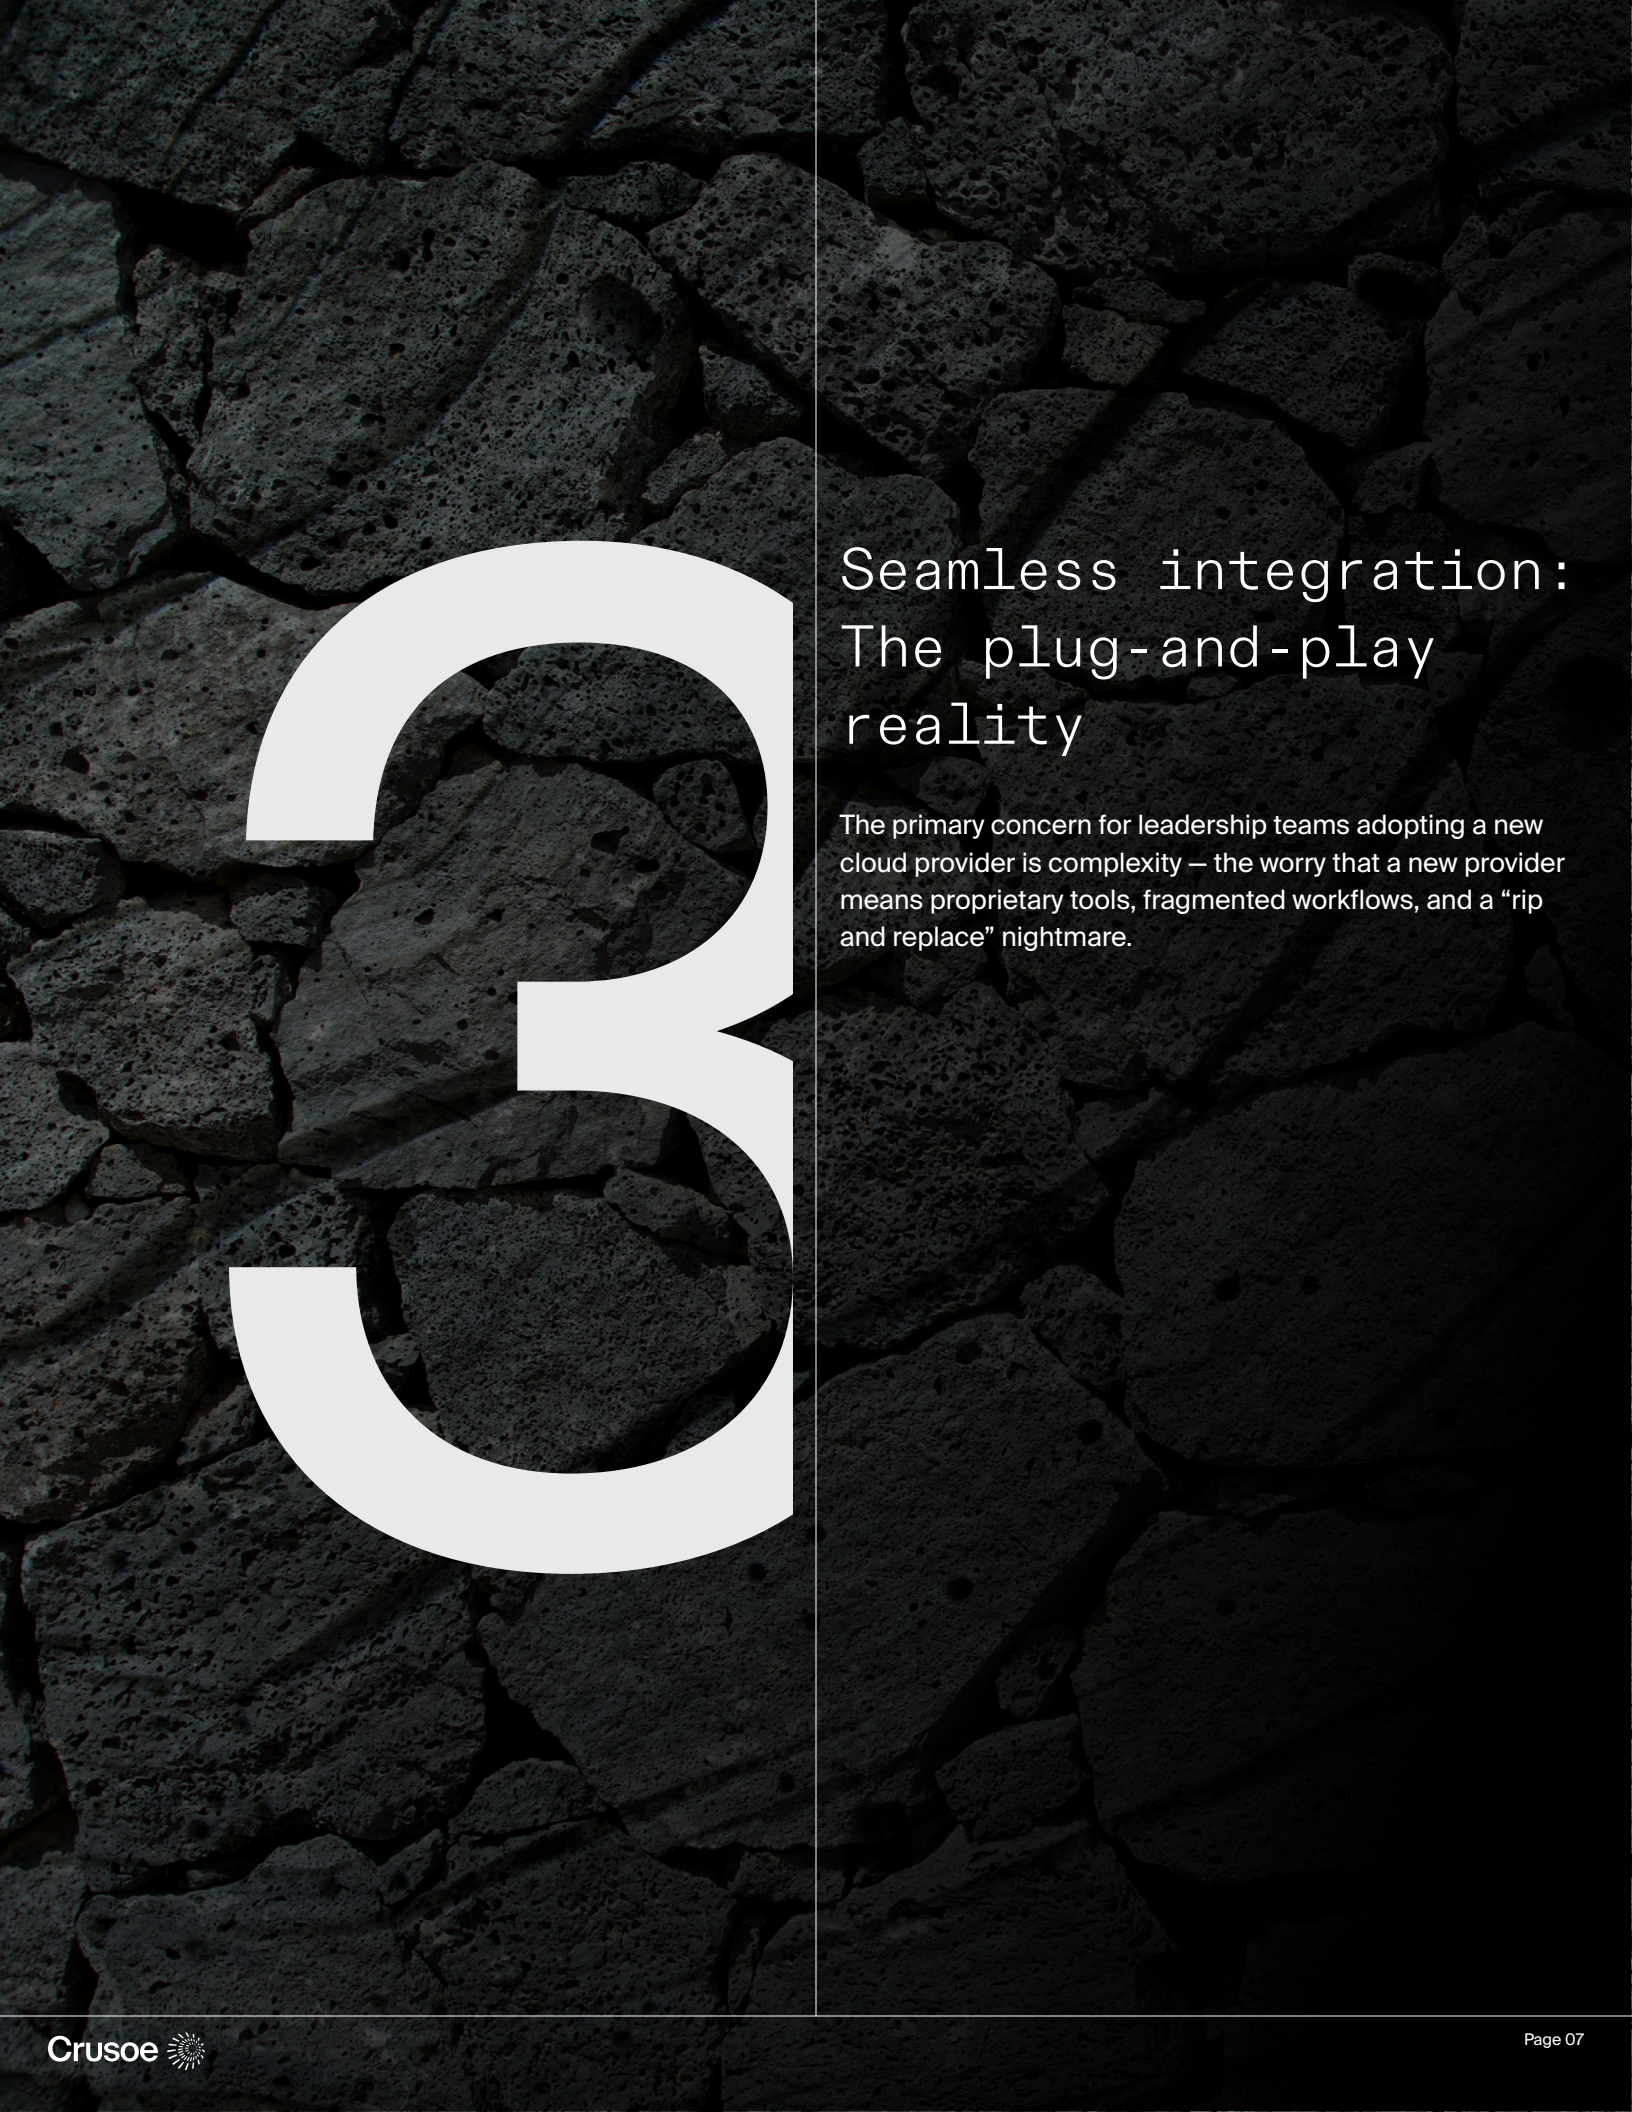

# 03

## Seamless integration: The plug-and-play reality

The primary concern for leadership teams adopting a new cloud provider is complexity — the worry that a new provider means proprietary tools, fragmented workflows, and a “rip and replace” nightmare.

## Hyperscaler

### The hyperscaler truth: The walled garden

Hyperscalers are often incentivized to build “walled gardens,” creating proprietary versions of open-source tools and custom APIs that can make it difficult to connect with outside resources. Building on proprietary Platform-as-a-Service layers often means migrating or connecting to another cloud requires significant, costly re-engineering.

## Crusoe Cloud

### The Crusoe approach: Open standards

Crusoe removes this friction by adhering to open standards and ensuring your team can get running quickly.

**No refactoring:**
We utilize standard Linux environments and enhanced Kubernetes clusters. If your team can run a container on a hyperscaler, they can point that same workload to Crusoe Cloud with Crusoe Managed Kubernetes (CMK).

**Direct connectivity:**
We support standard networking protocols, allowing you to create secure tunnels between your existing data lakes and Crusoe Cloud’s compute clusters.

**Managed Inference at production scale:**
Scaling inference workloads presents distinct challenges from training. As request volumes grow, teams face capacity planning complexity, cold-start latency issues, and the operational burden of managing distributed serving infrastructure. Crusoe Managed Inference eliminates this overhead entirely — deploy your models via API and scale seamlessly without expanding your DevOps team or managing underlying infrastructure. Focus your engineering resources on improving your models, not maintaining serving infrastructure.

**Strategic takeaway:**

Adding AI-specialized infrastructure doesn’t require rearchitecting your stack. Crusoe fits into what you’ve already built.

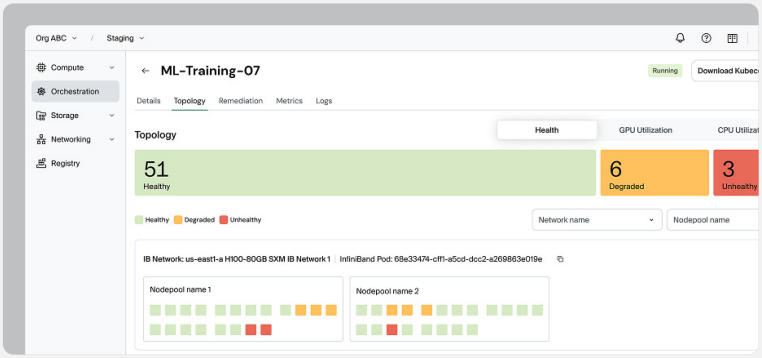


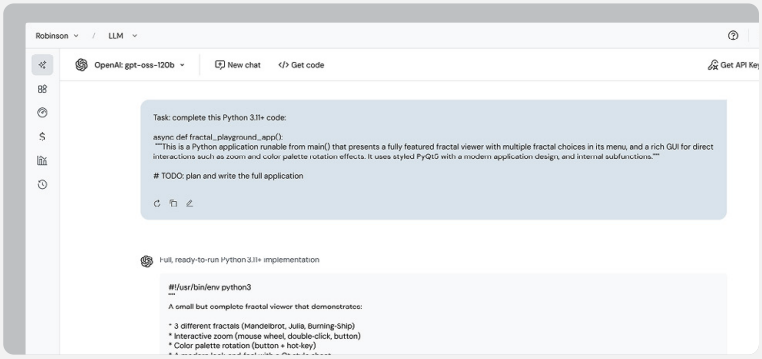

# 4

## The partnership advantage: Engineering expertise, not automated support

In a race defined by iteration speed, you cannot afford to be a ticket number in an automated queue.

## Hyperscaler

### Transactional utility vs. strategic partner

Hyperscalers often route customers through automated scripts and chatbots. At their scale, supporting hundreds of thousands of customers across dozens of product lines, it's an operational necessity, not a choice. Standardized support tiers and self-service tooling are how they manage volume.

## Crusoe Cloud

### Engineering leadership engagement

Beyond standard support, Crusoe offers deep technical partnerships for customers with complex AI infrastructure needs. Our solutions engineers can work as an extension of your team, collaborating on distributed training frameworks, MLOps, and cluster optimization. For customers running demanding workloads at scale, we offer an engineering engagement program that provides direct access to our technical team and the ability to inform our product roadmap.

At Crusoe, white-glove support is the standard. You gain direct access to solutions engineers who specialize in architecting complex AI workloads. With consistent 100% customer satisfaction, we provide the human layer of expertise necessary for production-grade AI.

**Strategic takeaway:**

Direct access to AI infrastructure experts helps you move faster and avoid costly mistakes during critical training and deployment phases.

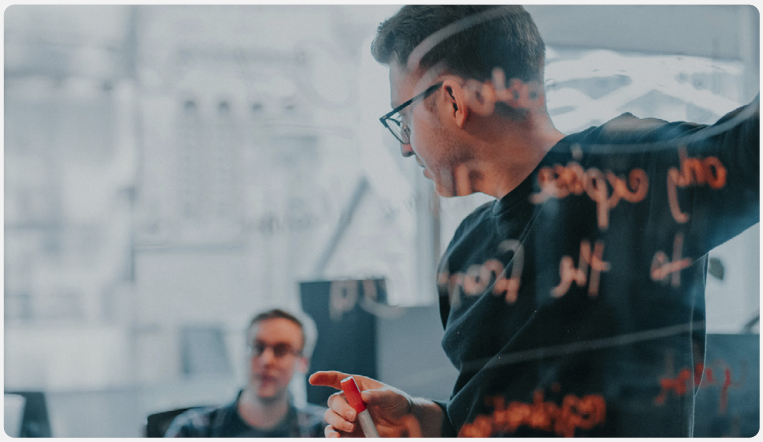

**Embrace the multicloud solution**
Just as you choose the right tool for a specialized job, it's time to optimize your stack with an AI-native cloud.

| | Crusoe Cloud |
|---|---|
| **Stay on Hyperscaler** | **Move to Crusoe Cloud** |
| **Legacy apps:** Customer-facing interfaces | **LLM training:** Large-scale development |
| **Databases:** SQL/NoSQL user data | **Inference:** Massive production clusters |
| **Business logic:** ERP, CRM, and Auth | **HPC:** Complex physics and bio-simulations |

Ready to build the future faster?

Schedule a consultation with our AI experts today.