



cloud architects

The Server Labs G-Cloud 15 Service Description

G-Cloud 15 Service Definition –
Ultra Compute & HPC Services

Ref: TSL/GCLOUD15/SERVICE/HPC

Issue Number
1.0

Date
Jan 2026



UltraCompute HPC Readiness & Workload Assessment for AWS

Overview

The Server Labs delivers a structured, architecture-led assessment to design high-performance computing services that run effectively and reliably on AWS.

HPC platforms succeed when compute, storage, scheduling, security, and operations are designed together. This service builds a clear understanding of workload behaviour and organisational constraints, then designs the right HPC service model—whether delivered through the UltraCompute product or bespoke HPC services—aligned to AWS-native capabilities and proven reference architectures.

Benefits


- Designs HPC services aligned to real workload behaviour
- Applies proven HPC reference architectures on AWS
- Establishes clear foundations for performance and scale
- Reduces risk during deployment and operational ramp-up
- Creates a direct path from assessment to implementation

What this service provides

This service focuses on designing high-performance computing services that work in production. It examines how workloads execute at scale, how they interact with data, how they are scheduled and secured, and how they must be operated over time. The emphasis is on shaping HPC services that are correctly sized, operationally sound, and aligned to real execution patterns.

The assessment covers:

- Workload execution patterns, including batch, parallel, tightly coupled, and GPU-accelerated jobs
- Compute architecture design across CPU and GPU selection, memory profiles, node topology, and scaling behaviour
- Storage and IO architecture, including shared file systems, scratch usage, throughput, and latency sensitivity



The assessment covers:

- Scheduler and queue design, including prioritisation, fairness, and utilisation efficiency
- Data gravity and movement between compute, object storage, and upstream systems
- Security and tenancy models, including isolation, identity boundaries, and encryption
- Operational service design, including monitoring, performance tuning, and cost visibility
-

The outcome is a designed HPC service model that can be implemented directly, mapped to either UltraCompute or a bespoke HPC architecture depending on workload complexity and organisational constraints

Who is this for

This service is suited to organisations designing or evolving HPC platforms on AWS, including:

- Life sciences and genomics organisations running large-scale pipelines
- Research and engineering teams executing sustained parallel compute workloads
- Space, climate, and physics modelling environments
- Defence and regulated research workloads with strict security requirements
- Enterprises modernising legacy or on-premises HPC estates

It is particularly valuable where multiple workload types must coexist on a shared HPC platform.

Delivery Approach

The assessment is delivered as a structured engagement led by senior HPC and cloud architects.

Work begins with detailed workload discovery to understand execution behaviour, data dependencies, and operational expectations. This is followed by architectural analysis aligned to AWS-native HPC capabilities, including managed schedulers, cluster orchestration, and high-throughput networking where appropriate.

Design decisions are grounded in established HPC reference architectures and real-world delivery experience. Where relevant, AWS ParallelCluster and AWS Parallel Compute Service are considered alongside the UltraCompute platform to ensure the chosen approach reflects both technical requirements and operational reality.

Findings are consolidated into a clear HPC service design that can be implemented directly, avoiding redesign or reinterpretation during deployment.



Key Deliverables

- HPC workload and execution profile
- Designed HPC service model for AWS
- Recommended UltraCompute or bespoke HPC architecture
- Compute, storage, and scheduler design guidance
- Security and tenancy architecture considerations
- Deployment and operational readiness summary

AWS Services & Framework Alignment

This service designs HPC platforms aligned to capabilities provided by Amazon Web Services, including high-performance compute, storage, networking, and orchestration services.

Architectures may reference AWS ParallelCluster and AWS Parallel Compute Service where appropriate, alongside UltraCompute's productised HPC platform and supporting AWS-native services.

Pricing & Commercial Model

- Engagement Model:** Professional Services
- Pricing:** Fixed-price or Time & Materials (depending on complexity)
- Procurement:** Available via G-Cloud 15 and AWS Marketplace



UltraCompute SaaS Deployment & Implementation for AWS

Overview

The Server Labs deploys the UltraCompute HPC platform into a customer's AWS environment, providing a production-grade HPC service in hours rather than months.

UltraCompute is a licensed, productised HPC platform that removes the need to design, build, and harden an HPC environment from first principles. It allows organisations to launch a new HPC service, replace ageing legacy clusters, or add burst capacity on demand—while running entirely inside the customer's AWS account and security perimeter.

Benefits

- Delivers a full HPC service in hours, not months
- Eliminates months of bespoke HPC platform engineering
- Runs entirely within the customer's AWS account
- Keeps all data inside existing security and compliance boundaries
- Supports new, replacement, and burst HPC use cases
- Continuously improved through a managed product control plane

What this service provides

This service delivers a deployable HPC product, not a custom-built platform.


UltraCompute encapsulates years of proven HPC experience in architectures, scheduler configurations, storage patterns, and operational controls into a licensed SaaS product. Deployment installs the UltraCompute platform into the customer's AWS environment, where it operates as a native HPC service under the customer's governance, security, and network controls.

UltraCompute can be used as:

- A **new shared HPC service** for research, engineering, or data teams
- A **drop-in replacement** for ageing on-prem or legacy HPC clusters
- A **burst and overflow platform** to extend existing HPC capacity

Because the platform executes entirely within the customer's AWS account, there is no data movement to external environments. Security, identity, encryption, and data residency are enforced using the customer's existing AWS controls.

The result is a production-ready HPC service that teams can adopt immediately, while retaining full control over security, cost, and operations.



Who is this for

This service is suited to organisations that want HPC capability delivered quickly, predictably, and safely, and without the overhead of building and operating a bespoke platform.

It is particularly valuable where:

- HPC services must be available rapidly to support programmes or research
- Legacy HPC platforms are expensive, fragile, or operationally constrained
- Capacity demand is variable and requires elastic scaling
- Security or regulatory requirements prohibit external compute platforms
- Multiple teams need access to a governed, shared HPC service

Delivery Approach

UltraCompute is deployed using a scripted, repeatable installation process into the customer's AWS account.

The deployment establishes UltraCompute as a licensed SaaS product operating within the customer environment. Core platform components, schedulers, compute pools, and storage integrations are installed and configured automatically, aligned to the agreed service design.

UltraCompute is managed through a product control plane that enables continuous platform improvement, upgrades, and feature evolution without requiring customer intervention or exposing customer data. All workload execution, data access, and runtime operations remain fully contained within the customer's AWS account at all times. Initial workloads are onboarded to validate platform behaviour, performance, and access, after which the service is ready for immediate production use.



Key Deliverables

- Licensed UltraCompute HPC platform deployed in the customer's AWS account
- Automated installation and configuration of core HPC services
- Configured schedulers, compute pools, and scaling policies
- Integrated storage and data access paths optimised for HPC
- Secure access aligned to existing IAM and network controls
- Initial workload onboarding and operational validation

AWS Services & Framework Alignment

UltraCompute operates on Amazon Web Services, integrating with native AWS networking, identity, compute, and storage services.

The platform aligns with established AWS HPC patterns, including AWS ParallelCluster and AWS Parallel Compute Service where appropriate, while providing a higher-level, productised HPC service that removes the need for customers to engineer and maintain these components themselves

Pricing & Commercial Model

UltraCompute is delivered under a software licence, deployed into the customer's AWS account.

- Licence pricing and usage costs are provided on request
- Capacity scales on demand within the customer environment
- Optional Professional Services may be used for onboarding or integration
- Optional managed services are available for ongoing platform operation

This model provides predictable access to a continuously improving HPC platform without long-term infrastructure commitment.



UltraCompute Extension & Custom Configuration for AWS

Overview

The Server Labs extends and configures the UltraCompute HPC platform to support advanced workloads, sector-specific controls, and complex operational requirements.

UltraCompute is designed to be usable out of the box, but real production environments often demand more. This service adapts the UltraCompute product to fit specialist storage patterns, regulatory constraints, legacy integrations, and non-standard scheduling or scaling behaviour—while retaining the benefits of a productised, upgradeable HPC platform.

Benefits

- Extends UltraCompute without breaking the product model
- Supports specialist and regulated HPC workloads
- Integrates with existing tools, pipelines, and estates
- Enables complex scheduling and scaling requirements
- Preserves upgradeability through the UltraCompute control plane
- Avoids long-term bespoke platform lock-in

What this service provides


This service focuses on adapting UltraCompute to real-world complexity.

UltraCompute provides a strong, standardised HPC foundation. Extension & Custom Configuration builds on that foundation to meet advanced technical and organisational requirements while keeping the platform aligned to the core product.

Typical areas of extension include:

- Specialist storage architectures, such as high-throughput shared file systems, tiered scratch patterns, or data-intensive workflows
- Advanced scheduling and queue models, including workload isolation, priority handling, pre-emption, or fairness controls
- Sector-specific controls, including data residency, access constraints, or compliance-driven
- Integration with existing toolchains, pipelines, identity providers, or data platforms
- Hybrid or transitional models, where UltraCompute operates alongside legacy HPC estates

Extensions are designed to work with the UltraCompute platform, not around it. Wherever possible, configuration is implemented through supported patterns, extension points, and automation rather than one-off custom code.



Who is this for

This service is suited to organisations using UltraCompute where standard configurations need to be adapted to meet production realities.

It is particularly valuable where:

- HPC workloads have specialist IO, scheduling, or performance characteristics
- Regulatory or sector constraints require additional controls
- Existing HPC tools or workflows must be preserved
- Multiple workload classes need to coexist on a shared platform
- The platform must evolve without losing upgradeability

Delivery Approach

Extension work is delivered as a **controlled enhancement of the UltraCompute product**.

The Server Labs works with the customer to identify where standard UltraCompute capabilities are sufficient and where extension is required. Enhancements are then implemented using supported configuration models, automation, and integration patterns, ensuring compatibility with the UltraCompute external control plane. Where necessary, extensions are coordinated with the product roadmap to ensure future upgrades remain smooth and non-disruptive. This approach allows customers to meet advanced requirements today while continuing to benefit from ongoing platform improvements.



Key Deliverables

- Configured UltraCompute extensions aligned to workload needs
- Advanced scheduler, storage, or scaling configurations
- Integrated security and sector-specific controls
- Toolchain and platform integration components
- Updated operational and configuration documentation
- Validated upgrade and support compatibility

AWS Services & Framework Alignment

UltraCompute extensions continue to operate on Amazon Web Services, integrating with native AWS compute, storage, networking, and security services.

Configurations may align with AWS ParallelCluster and AWS Parallel Compute Service patterns where appropriate, while remaining fully compatible with the UltraCompute product architecture.

Pricing & Commercial Model

This service is delivered as Professional Services alongside the UltraCompute licence.

- Pricing is provided on request
- Scope reflects the complexity and depth of required extensions
- Delivered in a time-boxed, well-defined engagement

This model allows customers to extend capability without compromising the core product.



UltraCompute Managed HPC Services for AWS

Overview

The Server Labs provides managed services for the UltraCompute HPC platform, ensuring it remains reliable, performant, secure, and cost-effective as workloads scale and evolve.

This service allows organisations to consume UltraCompute as a fully managed HPC capability, without taking on the operational burden of running and tuning a complex compute platform. The focus is on platform health, performance, and continuous improvement—while customers retain full ownership of workloads, data, and execution.

Benefits

- Operates UltraCompute as a production HPC service
- Maintains performance, reliability, and availability at scale
- Reduces operational load on internal engineering teams
- Optimises utilisation and cost over time
- Preserves security, governance, and data residency
- Aligns HPC operations with Trusted Foundations


What this service provides

This service focuses on operating UltraCompute as a dependable, long-running HPC service.

UltraCompute Managed HPC Services cover the ongoing stewardship of the platform rather than the workloads running on it. The Server Labs takes responsibility for platform operation, ensuring that schedulers, compute pools, storage integrations, and supporting services continue to perform as intended under changing demand.

Core service areas include:

- **Platform operations**, including monitoring, alerting, and incident response
- **Performance optimisation**, covering scheduler behaviour, queue efficiency, and resource utilisation
- **Capacity and scaling management**, ensuring compute resources match workload demand
- **Patch and upgrade coordination**, delivered through the UltraCompute control plane
- **Operational assurance**, including validation of availability, resilience, and security posture



- **Cost and utilisation optimisation**, aligned with agreed service objectives

The service ensures that UltraCompute remains a stable and trusted platform as usage grows, workloads diversify, and requirements change.

Who is this for

This service is suited to organisations that want to consume HPC as a managed capability rather than operate it themselves.

It is particularly valuable where:

- HPC platforms support mission-critical or regulated workloads
- Internal teams want to focus on research, engineering, or delivery rather than platform operations
- Workload demand is variable and requires active tuning
- Operational risk must be minimised as scale increases
- Platform reliability and predictability are essential

Delivery Approach

UltraCompute Managed HPC Services operate alongside the UltraCompute SaaS platform and external control plane.

The Server Labs monitors and manages the platform continuously, using telemetry from the customer's AWS environment combined with platform-level signals from UltraCompute. Operational activities are performed without exposing customer data or workloads, which remain fully contained within the customer's AWS account.

Service operation follows a defined cadence, covering routine health checks, performance reviews, optimisation activities, and upgrade coordination. Findings and actions are shared transparently, ensuring customers retain visibility and control while avoiding day-to-day operational overhead.



Key Deliverables

- Ongoing operation of the UltraCompute HPC platform
- Platform monitoring, alerting, and incident management
- Performance and utilisation optimisation activities
- Coordinated upgrades and platform improvements
- Operational reporting and service reviews
- Clear escalation and support processes

AWS Services & Framework Alignment

UltraCompute Managed HPC Services operate on **Amazon Web Services**, integrating with native AWS monitoring, security, and operational tooling.

The service aligns with Trusted Foundations, including Resilience & SRE best practices, Zero Trust security principles, and the FinOps Operating Model, ensuring HPC operations remain robust and sustainable.

Pricing & Commercial Model

This service is delivered as a managed service alongside the UltraCompute licence.

- Pricing is provided on request
- Typically structured as a recurring service aligned to platform scale
- Optional enhancements available for extended support or optimisation

This model provides predictable operational coverage while allowing capacity and usage to scale on demand.



HPC Architecture & Migration Services for AWS

Overview

The Server Labs designs and delivers bespoke high-performance computing architectures on AWS for organisations with complex, regulated, or mission-critical workloads.

This service exists for environments where HPC is central to delivery, not an experiment. It supports large-scale migration and modernisation programmes, including national research platforms, defence and security workloads, and highly regulated scientific and engineering estates. The focus is on building HPC platforms that are performant, secure, operable, and credible at scale.

Benefits

- Designs HPC platforms tailored to complex, real-world constraints
- Supports large-scale migration from legacy and on-prem estates
- Handles regulated, sovereign, and hybrid operating models
- Validates performance and scalability through benchmarking
- Aligns architecture with long-term operational sustainability


What this service provides

This service focuses on **architecting HPC platforms where standard solutions are insufficient**.

The Server Labs works with clients to design HPC environments that reflect the realities of their workloads, governance obligations, and operational maturity. Architecture is developed holistically across compute topology, scheduler strategy, storage and IO design, data movement, identity, security boundaries, and resilience.

Engagements typically include:

- End-to-end HPC architecture design for AWS
- Migration strategy from on-prem or legacy clusters
- Hybrid and phased transition models where estates cannot move in one step
- Scheduler and workload orchestration design aligned to workload behaviour
- Storage and IO architecture for data-intensive and performance-sensitive workloads
- Security and tenancy models suitable for regulated and sovereign environments
- Performance benchmarking to validate architectural decisions



The result is an HPC platform design that is technically sound, defensible to stakeholders, and ready for implementation at scale.

Who is this for

This service is designed for organisations where HPC is a core capability rather than a supporting function.

It is particularly suited to:

- National or large-scale research programmes
- Life sciences, genomics, and bioinformatics platforms
- Space, climate, physics, and engineering simulation environments
- Defence and security-sensitive workloads
- Enterprises retiring or modernising large on-prem HPC estates

These engagements typically involve multiple stakeholders, long delivery timelines, and material operational risk if executed poorly.

Delivery Approach

Engagements are led by senior HPC architects with experience delivering large, complex platforms.

Work begins with deep discovery across workloads, data flows, security constraints, and existing estate limitations. Architecture is then developed iteratively, balancing performance requirements with operational, security, and cost considerations.

Migration strategies are shaped to minimise disruption and risk, often involving staged transitions, parallel running, and controlled cutovers. Performance is validated through benchmarking and testing to ensure cloud-based HPC meets or exceeds existing capability.

Designs are delivered at a level of detail suitable for direct implementation and long-term operation



Key Deliverables

- Bespoke HPC architecture design
- Migration and transition strategy
- Scheduler, compute, and storage architecture
- Security and tenancy design
- Performance benchmarking results
- Implementation-ready design artefacts

AWS Services & Framework Alignment

Architectures are designed on Amazon Web Services, using AWS compute, storage, networking, and security services appropriate to high-performance workloads.

Designs align with Trusted Foundations, including Zero Trust security, Resilience & SRE practices, and FinOps principles.

Pricing & Commercial Model

This service is delivered as a major Professional Services engagement.

Pricing is provided on request and typically reflects the scale, duration, and complexity of the programme. Engagements are commonly structured around milestones or phases and often exceed £1m in value.

Eligible for AWS Migration Acceleration Programme (MAP) co-funding where applicable.



TSL Company Profile

Founded in 2004, TSL are cloud architects and integrators.

We offer a full range of lifecycle cloud solutions and services, from consultancy, software/application development, to cloud architecture design, and integration, to fully managed services.

Our specialisations are High Performance Computing and AI/ML deployed in highly secure cloud environments, with best practice built in from day one.

Our clients include private and public organisations, with an emphasis on those processing large amounts of data in mission critical environments, particularly in the Lifesciences, Utilities, Banking, Space, Metrology and Telecoms sectors.




The Server Labs 3 Layer Service Map





AI Readiness




Cloud Maturity Accelerators



Trusted Foundations





The Company's mission is:

- To provide expert services in the field of cloud architectures, High Performance Computing, and advanced cloud software engineering
- To radically improve cloud software development processes
- To help organisations achieve better business results through the correct use of cloud technologies
- To create high quality innovative cloud software solutions, providing added value to our customers

The specific value, experience and expertise that The Server Labs can provide are:


1. **Technical excellence** and capability to act as lead on architectural decisions and as technology expert in software and system subjects.
2. **Proven experience** in implementing HPC, FinOps, Cloud Security and Big Data Projects.
 - FinOps sustainable change management to enable innovation
 - Industry leading experts in high performance computing
 - Cloud security and sensitive, big data management and processing
3. **Cloud migration:** Proven track record in migrating large, complex and sensitive dataset from on-premise to cloud environments and from legacy to target cloud.
4. **Architecture experience** at software and system level.

AWS MAP and WAF certified

TSL have achieved AWS MAP competence, which is a rigorous and demanding standard to achieve and few boutique consultancies like TSL have achieved this.

Our architects and engineers are experts with an average of 10 years' experience in the planning, design and development of complex software systems. Our multinational team has been a pioneer in cloud native service, serverless computing, automation, infrastructure as code, and distributed architectures, and has the required hands-on experience in many state of the art technologies.

The Server Labs is a leader in cloud computing solutions and services, helping organisations to move to the cloud at all levels.



Ready to lead the future?

Explore what's possible with TSL—
where innovation meets execution.
Contact us today to start your journey.

[Visit our website](#)

[Email us](#)



United Kingdom

The Server Labs Ltd.
10 Bloomsbury Way London
WC1A 2SL United Kingdom

+44 (0)203 948 1082

Germany

The Server Labs Berliner
Allee 47, 64295 Darmstadt,
Germany

+49 6151 277 6037

Spain

The Server Labs S.L.
C/María de Molina, 39
28006 Madrid, España

+34 91 745 68 77