

G-Cloud 15

Cloud support

COMMERCIAL IN CONFIDENCE

By **Hypercube**

15 January 2026

Contents

Section 1 Service definition 3

 1.1 Introduction 3

 1.2 Service description 3

 1.3 Features and capabilities 3

 1.4 Onboarding and offboarding 4

 1.5 Service processes 4

 1.6 Partner credentials 4

 1.7 Roles and responsibilities 5

 1.8 Pricing and commercials 5

 1.9 Data backup, restore and disaster recovery 5

 1.10 Business continuity and resilience 6

 1.11 Onboarding and offboarding support 6

 1.12 Hosting options and data locations 6

 1.13 Data access and exit 6

 1.14 Security 7

Section 1 Service definition

1.1 Introduction

This service definition outlines the cloud support services that Hypercube provides for organisations running workloads on Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP). Our aim is to help customers deploy, migrate, optimise and manage their cloud environments with confidence, leveraging Hypercube's expertise, partner status and toolsets across all three hyperscalers.

1.2 Service description

Hypercube's cloud support service delivers end-to-end assistance across the full cloud lifecycle, from assessment and migration through to optimisation and ongoing managed support. Hypercube's team of cloud engineers, data specialists and AI experts provide best-practice guidance and hands-on support to build secure, scalable and resilient infrastructures. Services are delivered on a cross-platform basis, ensuring consistent governance and performance whether workloads run on AWS, Azure or GCP.

1.3 Features and capabilities

Table 1-1 – Features and capabilities describes the core components of Hypercube's cloud support service. They combine proven delivery patterns with partner-enabled toolsets to provide customers with a high level of flexibility and control over their cloud estates.

Feature	Description
Scalability and flexibility	Cloud resources are scaled up or down quickly and efficiently to meet changing business demand across AWS, Azure and GCP.
Rapid provisioning	Hypercube provisions virtual machines and cloud services on demand, enabling organisations to spin up environments in seconds or minutes.
Architecture, integration and governance	Hypercube's architects design secure, well-governed solutions and integrate services across compute, storage, networking and identity, ensuring compliance with regulatory requirements.
Cross-platform support	The service spans all major hyperscalers, offering consistent processes and tooling across AWS, Azure and GCP.
Regulatory compliance	Support engineers ensure that workloads meet relevant regulatory and industry standards, with controls for data sovereignty and privacy.
Cost optimisation	Hypercube monitors utilisation and implements automation to optimise cloud spend, leveraging tools such as AWS Cost Explorer, Azure Cost Management and GCP Billing reports. Where possible, Hypercube uses funding programmes and partner discounts to reduce costs for end customers.
Security and governance	Hypercube designs and enforces security baselines, including identity and access management (IAM),

Feature	Description
	network security, patching and compliance scanning. This includes alignment with ISO 27001 and other applicable standards.
Automation and infrastructure as code	Infrastructure is deployed and managed through code (e.g., Terraform, CloudFormation, Bicep) to ensure repeatable, auditable deployments across all three clouds.
Observability and support tooling	Hypercube implements monitoring and alerting and provides dashboards and reports to give visibility of performance, availability and costs.
Multi-cloud governance	Hypercube aligns policies across hyperscalers to ensure consistent tagging, cost allocation, identity management and security.

Table 1-1 – Features and capabilities

1.4 Onboarding and offboarding

Hypercube's onboarding process is driven by a seamless, agile methodology. We capture functional and non-functional requirements, define key delivery milestones and agree a clear contractual scope. Training materials and a train-the-trainer approach help customer teams become self-sufficient. At the end of the contract, Hypercube assists with data extraction and ensures a structured handover.

1.5 Service processes

Hypercube's engagement model follows four key stages:

1. Assessment and planning – Evaluate existing infrastructure, identify gaps, perform workshops and design the target state.
2. Migration and deployment – Plan and execute migrations with minimal disruption, establishing cloud foundations, landing zones and security controls.
3. Optimisation and automation – Introduce automation and infrastructure-as-code, monitor usage and optimise resources to ensure cost-efficient operations.
4. Ongoing managed support – Provide proactive monitoring, maintenance, patching, incident management and continual service improvement to keep customer environments secure and reliable.

1.6 Partner credentials

Hypercube's cloud expertise is recognised by leading hyperscalers through official partner accreditations. These credentials unlock additional funding, resources and priority support for our customers.

- AWS Advanced Tier Services Partner: This accreditation is awarded to partners who demonstrate a high level of AWS expertise, employ AWS-certified professionals and have a proven track record of delivering successful projects. It enables Hypercube to access partner funding programmes and cost-optimisation tools for proof-of-concepts and migrations.
- Azure Infrastructure Solutions Partner: Hypercube has achieved Microsoft's Azure Infrastructure Solutions Partner designation. This status recognises Hypercube's broad capability to help customers accelerate migration and modernisation of key infrastructure workloads to Azure, providing secure and scalable cloud foundations.

- Google Cloud Partner: As a member of the Google Cloud Partner Advantage programme, Hypercube's engineers hold Google Cloud certifications and have experience designing and delivering GCP solutions. Our partner status gives us access to training resources, co-funding opportunities and priority support channels, enabling us to accelerate migrations and optimise GCP environments.

1.7 Roles and responsibilities

Clear responsibilities ensure successful engagements:

- Hypercube responsibilities: Provide skilled resources, design and implement cloud architectures, manage migrations, monitor and maintain environments, and ensure compliance with agreed service-level agreements (SLAs).
- Customer responsibilities: Supply timely information, grant necessary access, review and approve deliverables, and maintain any on-premises elements within their control.

1.8 Pricing and commercials

Pricing is tailored to each engagement based on scope, complexity, required SLAs and hyperscaler consumption. Hypercube offers flexible pricing models including time-and-materials, fixed-price packages and consumption-based services. Funding or credits may be available through our partner programmes to offset migration and proof-of-concept costs.

1.9 Data backup, restore and disaster recovery

Protecting customer data and ensuring rapid recovery are central to Hypercube's managed service. Hypercube designs backup and disaster-recovery (DR) strategies that leverage the native services of AWS, Azure and GCP. Key elements include the following:

- Integrated backup and DR platform: Rather than adopting separate tools for backup and low-RPO/RTO disaster recovery, Hypercube advocates a single solution that provides both backups and fast fail-over from one platform. Using a one-stop solution reduces complexity and cost.
- Multi-cloud and hybrid protection: Hypercube's backup solution is flexible enough to protect workloads running across on-premises data centres as well as hybrid AWS, Azure and GCP environments. It allows administrators to manage on-premise and cloud workloads from a single dashboard.
- Database-aware backups: To support mission-critical workloads, Hypercube ensure that backup processes are application-consistent for databases like Oracle, SQL Server, MySQL and PostgreSQL. Application-consistent, incremental backups help achieve low recovery point objectives.
- Cross-region replication: For resilience against regional outages, Hypercube replicates backups to other regions or clouds and implement cross-region DR patterns. This ensures that backups remain accessible even if the primary region is unavailable.
- Automated recovery orchestration: Where required, Hypercube utilises tools such as AWS Elastic Disaster Recovery or similar services to orchestrate fail-over of entire workloads. These tools minimise RPO and RTO for critical applications.
- Cost-efficient storage: Backups are stored in object storage (e.g. Amazon S3) or equivalent services. Using native cloud object storage keeps costs low and allows direct recovery from the backup store.
- Regular testing: Hypercube schedules regular DR drills and restoration tests to verify the integrity of backups and the effectiveness of recovery procedures.

1.10 Business continuity and resilience

In addition to technical backup and DR services, Hypercube works with clients to develop business continuity plans that address wider operational resilience:

- Resilience by design: Hypercube architects solutions across multiple availability zones and regions to mitigate single points of failure. Reliability is bought from hyperscalers but resilience must be built into the design; multi-site deployments provide rapid fail-over when individual regions experience issues.
- Multi-cloud strategy: Where appropriate, Hypercube distributes workloads across AWS, Azure and GCP to avoid dependency on a single provider. This reduces concentration risk and allows services to continue if one hyperscaler experiences an outage.
- RPO/RTO definition: Recovery time objectives (RTO) and recovery point objectives (RPO) are agreed with the customer and inform the choice of backup frequency, replication strategy and fail-over architecture.
- Incident playbooks and “break-glass” procedures: Hypercube develops runbooks for common failure scenarios, ensuring teams know how to trigger fail-over and restore services quickly.

1.11 Onboarding and offboarding support

Hypercube's onboarding and offboarding processes are designed to ensure a smooth start and finish to each engagement:

- Onboarding: Hypercube runs discovery workshops to gather functional and non-functional requirements, assess existing architectures and design target landing zones. Detailed runbooks, training material and a train-the-trainer approach help customer teams become self-sufficient. For large migrations Hypercube provides project management, technical design and additional training; any extra effort beyond standard onboarding is agreed and costed separately.
- Offboarding: At the end of the contract, customers retain full ownership of their data. Hypercube provides tools and guidance to export workloads (for example, exporting virtual machines to an open virtualisation format and transferring data to another provider). A project is initiated to manage decommissioning, certificate revocation and removal of trust relationships; additional offboarding work may be chargeable. Customers are responsible for planning and executing migration to a new platform; Hypercube assists with data extraction and documentation.

1.12 Hosting options and data locations

Hypercube's service supports workloads hosted on AWS, Azure and GCP. Customers choose the region(s) in which their data resides based on latency, regulatory and sovereignty requirements. For example, Hypercube deploys services exclusively within UK sovereign data centres or across multiple regions globally. Customers have control over data storage and processing locations and enforce regional restrictions as needed.

1.13 Data access and exit

Customers always retain ownership of their data. At contract end, Hypercube provides mechanisms to extract or migrate data and workloads:

- Data export tools: We enable customers to export virtual machines in open virtualisation formats (e.g., OVF) and download or transfer them to new environments.
- Project-based offboarding: Where complex migrations are required, Hypercube initiates a project to manage transfer of data, communication lines, trust

relationships and certificates. Such projects are agreed with the customer and may incur additional fees.

- Retention periods: Any residual data in Hypercube systems after contract termination is securely erased according to agreed retention schedules.

1.14 Security

Security is foundational to our cloud support service. Hypercube adopts a multi-layered security model aligned with the shared responsibility principles of AWS, Azure and GCP:

- IAM: Hypercube controls access to resources and what actions are performed. IAM structures define identities, authentication and authorisation, ensuring that access is granted safely, on purpose and tracked.
- Policy enforcement: Hypercube defines and enforces policies that specify the operations users perform and the conditions under which resources are used. Policies implement intent and ensure consistent enforcement across clouds.
- Encryption: Data is encrypted at rest and in transit. Hypercube supports multiple encryption models and ensures that even if unauthorised access occurs, data remains unreadable.
- Shared responsibility model: Hypercube helps customers understand and fulfil their responsibilities. Cloud providers handle physical security, network infrastructure and core services, while customers are responsible for data security, application configuration, IAM, policy definitions and encryption key management.
- Compliance and certifications: Hypercube's processes align with recognised frameworks such as ISO 27001 and the Cyber Essentials scheme. Hypercube implements controls to meet sector-specific regulations. Customers integrate their SIEM and SOC services to ensure visibility and auditability.
- Continuous monitoring: Hypercube deploys monitoring and alerting tools across all cloud environments to detect security incidents, vulnerabilities and misconfigurations promptly. Automated remediation and incident response runbooks ensure rapid containment and resolution.