



Colibri Digital

G-CLOUD 15 SERVICE DEFINITION
SNOWFLAKE DATA CLOUD PLATFORM
BUILD ON AZURE

Contents

1. Service Overview	3
2. Features & Benefits	4
3. Technical Requirements & Implementation	5
4. Security & Quality Assurance	5
5. Support & Service Levels	5
6. Constraints	6
7. Onboarding & Offboarding	6
Onboarding	6
Offboarding & Exit	6
Data Portability	6
Knowledge Handover	6

1. Service Overview

Supplier Name	GCI NETWORK SOLUTIONS LIMITED
Service Branding	Colibri Digital
Framework	G-Cloud 15 (RM1557.15)
Service Type	Lot 3: Cloud Support

Colibri Digital's Snowflake Data Cloud Platform Build on Azure, provides a high-performance, secure, and scalable foundation for modern data warehousing, analytics, and AI. By leveraging Snowflake's unique multi-cluster shared data architecture within your Azure environment, we deliver a unified platform that supports consolidation of datasets and scalable performance, subject to workload design and buyer configuration.

Our service focuses on rapid deployment and best-practice configuration, ensuring your platform is governed, compliant, and ready to support data-driven decision-making. We integrate Snowflake seamlessly with native Azure services (such as Azure Data Lake Storage (ADLS) Gen2, Azure Private Link, and Microsoft Entra ID) to create a cohesive ecosystem that reduces operational complexity and the total cost of ownership (TCO).

This Service Definition covers design and build activities only; ongoing operational support and service levels are out of scope for this offering.

Key technical components of the delivery include:

Account & Organisation Setup	Expert configuration of Snowflake Accounts and Organisations within your Azure regions, including VNet integration and Private Link support.
Snowflake Horizon Governance	Implementation of Snowflake Horizon for centralised compliance, data quality monitoring, and fine-grained access control (RBAC).
Automated Data Ingestion	Design and build of automated ingestion patterns using Snowpipe for real-time streaming and Azure Blob/ADLS integration for batch processing.
Data Transformation	Implementation of scalable transformation pipelines using Snowflake Dynamic Tables or dbt (data build tool) to ensure "analytics-ready" data.
Security & Identity	Designed to support public sector security requirements, subject to buyer assurance/accreditation and agreed controls.
Operational Monitoring	Implementation of Snowflake Virtual Warehouses leveraging Account Usage views and Query History for real-time performance tracking and cost observability.

Core Public Sector use cases would include:

Centralised Citizen Data Records	Consolidating disparate datasets into a single "Golden Record" while maintaining strict data privacy and residency controls.
Predictive Policy Modelling	Enabling data scientists to run complex SQL and Python-based models on a unified platform to forecast policy impacts.
Automated Regulatory Reporting	Modernising legacy cycles through automated transformation pipelines, reducing manual effort and improving auditability.
Secure Inter-Departmental Sharing	Facilitating secure, "no-copy" data sharing between government entities using Snowflake Data Exchange and governed shares.
Real-Time Service Monitoring	Leveraging Snowpipe and Dynamic Tables to provide live insights into public service performance and resource allocation.

2. Features & Benefits

The following table outlines the core technical features of the build and the value outcomes they support.

Technical Features	Benefits
Multi-Cluster Shared Data	Strategic Clarity: Supports unlimited concurrent users and workloads without performance degradation for maximum ROI.
Snowflake Horizon	Secure-by-Design: Centralises governance and auditing to ensure compliance with GDPR and NCSC cloud security principles.
Automated Ingestion (Snowpipe)	Operational Excellence: Provides real-time data availability, allowing teams to focus on citizen-facing insights rather than manual loading.
Azure Native Integration	Consistency & Speed: Seamlessly connects with ADLS Gen2, Private Link, and Entra ID for a secure, high-performance cloud ecosystem.
Comprehensive Documentation	Streamlined Accreditation: Built-in security controls and runbooks simplify the process for SIROs to approve live operations.

3. Technical Requirements & Implementation

The implementation is delivered in line with the GDS Service Standard, applying an iterative four-phase model:

The service is delivered by SC-cleared consultants, if required, using a hybrid model of remote and on-site support.

Discovery	Analysis of existing data architecture, technical debt, and sovereignty requirements to define success criteria.
Alpha	Validation of the Snowflake architecture through a Proof of Concept in a sandboxed Azure environment.
Beta	Pilot migration of specific data workloads with limited live traffic to confirm operational readiness and governance.
Live	Full production deployment, integration with live monitoring and transition to steady-state.

4. Security & Quality Assurance

As a specialist in cloud data platforms, Colibri Digital ensures all Snowflake builds meet the highest security standards.

Secure Development:	All project staff are SC-cleared (or clearable) consultants who can be provided, subject to availability and buyer requirements.
Data Handling	Data handling and residency requirements are agreed with the buyer and implemented through the selected Azure region and agreed controls.
Standards	Our delivery workflows are to meet rigorous government and industry security benchmarks.

5. Support & Service Levels

This Service Definition covers design and build activities only; ongoing operational support and service levels are out of scope for this offering.

6. Constraints

The following constraints apply to ensure transparency:

Professional Services Scope	This is a Lot 3 Cloud Support offering; it does not include ongoing Azure or Snowflake consumption costs or third-party licences.
Data Source Accessibility	The feasibility of automated ingestion via Snowpipe is dependent on the availability and connectivity of legacy data systems.

7. Onboarding & Offboarding

Onboarding

Buyers are onboarded via a formal kick-off scheduled at call-off, subject to buyer availability; delivery artefacts and collaboration tooling are agreed with the buyer.

Offboarding & Exit

At completion, we hand over agreed deliverables and confirm secure return/deletion of buyer-provided data in line with agreed handling procedures.

Data Portability

All documentation is provided in open formats (PDF).

Knowledge Handover

All documentation, runbooks, and templates are handed over in open formats (PDF) to ensure operational independence.