



CATALYST BI Ltd

SERVICE DEFINITION DOCUMENT

Prepared for G-Cloud 15 Framework

TABLE OF CONTENTS

Introduction 4

Catalyst Value Proposition..... 6

Qlik 7

 Data Moved – Qlik Talend Cloud.....8

 Data Analysed – Qlik AI and Analytics Capacity8

 Fusion for Qlik9

Athena NHS Command Centre 10

Chat-A-Lyst – Talk to my Data 11

Snowflake – Data Cloud..... 12

 Quickstart Program For Public Sector12

 Snowflake Platform Architecture13

 Business Continuity15

 Snowflake Cost15

 Data Architecture Patterns With Snowflake16

 Data Lake/Lakehouse17

 Data Mesh17

 Applications.....18

 Data Ingestion and Transformation19

 Openflow Snowflake19

 Native Connectors20

 Snowpipe and Snowpipe Streaming.....20

 Partner ETL tools20

 Unified Platform For End-To-End Ai, MI and Bi.....23

DataOps Live 27

DataRobot..... 28

 Catalyst AI Cloud28

 Use Case Workshop.....28

 AI Readiness Assessment28

Fivetran..... 30

Coalesce..... 30

InsightSoftware – VizLib & Fiplana 31

Consultancy and Education Services 31

Assessment & Strategy..... 32

 BI Modernisation - Data and Analytics Strategy32

Qlik Health Check	32
Data Literacy Adoption Framework	34
Training & Education	35
Data Literacy.....	35
Qlik Sense	35
Qlik NPrinting	36
Qlik GeoAnalytics	37
Qlik APIs and Mashups	37
Qlik App Automations	37
Adaptive Consultancy Services	38
Consultancy and Development	38
Paired Development.....	38
Managed Application Support	39
Catalyst BI Cloud Hosting	40
Qlik Cloud Analytics Migration Services	41
Qlik Sense Deployment & Configuration - Customer Owned	41
Qlik Governance & Security Definition.....	42
Project Management.....	43
Support.....	44
Why Choose Our Services	44
How to Buy Our Services	44

Introduction

Catalyst BI Limited is a UK market leading solution provider of data and AI transformation services based on best of breed technology software partners, enabling us to support our customers on Data Management, Data Analytics and Data Science projects, at any and every stage of an organisation's digital transformation.

We are passionate about successful project deployments, and we aim to ensure you maximise your return on investment with impactful business value.

We are fully committed to delivering world-class outcomes for our customers by helping them unlock the full potential of their data and technology, strengthened by our consulting expertise and value added services.

This service definition document covers Catalyst BI Limited's services as listed on the G-Cloud 15 Framework. It covers our:

Our Data & Analytics Capabilities

Data Management

Our Data Engineering practice delivers modern, scalable, and automated data pipelines built on leading technologies including Qlik, Snowflake, Talend, Fivetran, DataOps.live, and Coalesce. We design and implement end-to-end data ecosystems that enable organisations to ingest, transform, govern, and operationalise data with speed and reliability.

Our offerings span data quality, data governance, pipeline automation, and orchestration, ensuring that data is accurate, trusted, and ready for analytics, AI, and business decision making. We help clients modernise legacy ETL processes, adopt ELT best practices, and build cloud native architectures that scale with business growth whilst driving efficiency.

Data Analytics

Our Data Analytics practice empowers organisations to turn data into actionable insight through modern BI, advanced analytics, and intuitive self-service experiences. Leveraging technologies such as Qlik Business Intelligence, PowerBI, Qlik AutoML, and a full suite of reporting, data analytics, and data integration capabilities, we deliver fast, trusted, and scalable analytics outcomes.

We enhance Qlik Cloud Analytics with Fusion for Qlik, enabling personalised mashup portals for internal and external use cases. These tailored experiences unify dashboards, applications, and curated insights into a single, intuitive interface.

Our solutions are further strengthened by the Vizlib suite- including Finance, Collaboration, Library, Self-service, and Gantt- to extend Qlik's native capabilities and deliver richer visualisation, planning, and workflow experiences. We support organisations with data storytelling, predictive analytics, automated insights, and governed self-service, ensuring every user can access reliable, contextualised information to drive smarter decisions.

AI & Advanced Analytics

Our AI & Advanced Analytics capabilities enable organisations to unlock deeper intelligence, automate decision-making, and accelerate innovation. We combine the strengths of Qlik, DataRobot, and the Snowflake Data Cloud to deliver scalable, governed, and production ready AI solutions. Using Qlik's AI enhanced analytics, including Insight Advisor, Answers and AutoML, we empower business users with automated predictions, natural language insights, and guided analytics. With DataRobot, we deliver enterprise grade machine learning and MLOps, enabling rapid model development, automated feature engineering, model governance, and seamless deployment.

The Snowflake Data Cloud provides the foundation for secure, scalable AI workloads, enabling feature engineering, model execution, and advanced analytics directly within the platform. Our offerings span predictive modelling, automated machine learning, AI-driven analytics, MLOps, and embedded intelligence, helping organisations operationalise AI and create measurable business value.

Associated Services

Our Associated Services provide comprehensive support across the full lifecycle of your data and analytics initiatives. We deliver consultancy, implementation, solution design, and development services to help organisations architect and deploy modern data platforms with confidence.

To ensure long-term success, we offer managed application support, covering optimisation, and continuous enhancement. We also support hosted and SaaS deployments, enabling flexible, secure, and scalable delivery models tailored to your operational needs.

In addition, we provide product support and training, equipping teams with the knowledge and best practices required to maximise the value of their technology investments. Whether adopting new tools or maturing existing capabilities, our experts ensure your organisation is empowered to operate efficiently and innovate continuously.


Our approach goes beyond technology. We support **behavioural change, education, data literacy, and user enablement**, ensuring that solutions are adopted effectively and that people across the organisation feel empowered to use data with confidence. This focus on capability building helps maximise return on investment and ensures long-term success.

Additionally, this document details how our services and company align with the G-Cloud buying process and provides typical information to help you understand how to buy, configure and consume our services, and how to leave our services should the need arise.

Catalyst Value Proposition


Catalyst BI delivers bespoke and end-to-end data-driven solutions, partnering with best-of-breed software and technology to help our customers achieve their data transformation ambitions, across data management, data analytics and data science perspectives.

At a high level, we help organisations to free their data, wherever it resides, improve the decision-making process in your organisation by helping you to make more effective data-driven decisions, and experience the full potential of human and machine intelligence, by seamlessly deploying production-ready, business trusted AI applications.



Qlik

Qlik Cloud sets the benchmark for a new generation of AI, analytics and trusted data products. Empower users at any skill level to freely explore data with powerful AI combined with the industry's most powerful analytics engine. Bring actionable insights into every decision with the industry's most complete for modern BI solution – on our cloud or anywhere you choose. Organizations collect huge amounts of data. But most of the data that can guide decisions and actions isn't analytics-ready or accessible. Qlik operationalizes DataOps for analytics, turning raw data into trusted, actionable data that's easy to find, current, and immediately available to Qlik Sense — on any cloud you choose.



Qlik allows organisations to efficiently move, structure and catalogue data in real time to your cloud of choice with broad support for cloud object stores, data warehouse & lakes, streaming services, relational and non-relational repositories and more.

Hybrid Cloud Analytics means that compute operations can take place entirely in Qlik's cloud or via a hybrid deployment with your cloud of choice. Flexible analytics deployment options protect data location and local governance needs while providing unmatched performance and scale.

Qlik helps to create an end-to-end analytics data pipeline: continuously moving data, creating analytics-ready structures like cloud warehouses and lakes, augmented analytics to generate new insights, API & application integration to immediately act. Automation drives proactive insights to compel action in the moment by directly triggering other processes and downstream activities in relevant systems.

Data is one of your business's most critical assets, which is why Qlik treats it with the utmost care. Through security- and privacy-by-design development and SaaS operational processes, Qlik ensures their solutions align with the latest data protection and privacy laws around the world, such as GDPR and certification with SOC 2, SOC 3 and ISO 27001.

The Qlik platform, which now incorporates Talend easily adapts to your requirements with elastic scalability to meet the peak demands of ever-increasing data, users, and complexity of use cases. Centralized governance and management and data lineage helps to ensure data security, user adoption, accuracy, and overall system reliability.

Qlik's extended ecosystem includes a global user community of over 100,000 participants and a marketplace for sharing applications, data, and knowledge.

Catalyst BI is an Elite Qlik partner and Qlik's no.1 partner in the Public Sector for the last 5 years and is the largest Qlik consulting practice in the UK and is Qlik UK's NHS, Blue Light and Local Govt Strategic partner of choice. We are highly experienced in providing business value through the Qlik platform.

The Qlik capabilities are complimentary and can fit into and adapt to any existing technology eco-system, whether that be Hyperscalers such as Microsoft, Google, Amazon or Best of Breed Solutions including DataBricks, Snowflake and existing BI tools.

Data Moved – Qlik Talend Cloud

Qlik Talend Cloud offers data integration / movement and streaming, quality, security and governance, application integration, open data lakehouse, data product catalog and marketplace, in a single platform that fits with virtually any data architecture/methodology, ensuring the business always works with the most accurate data whenever it's needed.

Through automation and AI augmentation Qlik has eliminated the data barriers between business and IT. No matter your role, you can be confident that our no code/pro code offering will enhance your productivity.

Obtain predictable value from your data while supporting complex security and compliance needs. Data products provide high-quality, curated, readily accessible, and trustworthy data to the right consumers at the right time.

There are 4 editions available with increasing capabilities from Starter, Standard, Premium and Enterprise.

4 Qlik Talend Cloud Editions			
Unifying Heritage Qlik Cloud Data Integration and Talend Cloud Offerings			
Starter	Standard	Premium*	Enterprise*
SaaS Data Movement	Starter +	Standard +	Premium +
	Change Data Capture	Data Transformation	Data Products
		Quality	Enterprise Apps
			Generative AI

* Premium & Enterprise Editions include Heritage Talend Cloud functionality & British Data is a separately licensed offering & Content client managed SDC is separately licensed

Qlik Talend Summary

Trusted data foundation for AI, analytics, and operations


Comprehensive Best-in-class capabilities	Agile For changing requirements	Enterprise-grade Trust for everyone
✓ Universal connectivity from virtually any source to any target ✓ Automated data movement, transformation, and trusted access in one platform ✓ Delivering real-time consumption-ready data for AI, analytics, and operations	✓ Cloud independent and data-platform agnostic ✓ Supports data architect, engineer, scientist and steward workloads ✓ AI-enabled pipelines that control costs and optimize data processing	✓ High-quality, accurate data sharing via data product catalog and marketplace ✓ Supports secure, compliant, and scalable architectures ✓ Foster confidence in your data with the Qlik Talend Trust Score for AI framework
MARKET LEADING CDC & ETL	MARKET LEADING CLOUD INTEGRATION	LEADER IN DATA QUALITY, EASE OF USE AND AUTOMATION

Data Analysed – Qlik AI and Analytics Capacity

Qlik's analytics portfolio includes:

1. Qlik Cloud Analytics: A modern analytics platform for self-service data exploration and visualization.
2. Qlik Predict: Offers no-code machine learning capabilities for advanced analytics.
3. Qlik Answers: Provides smart search functionality using natural language processing.

These solutions work together to unify disparate data sources and support a wide range of analytical use cases, from self-service analytics to governed reporting and operational analytics. Qlik's cloud-native architecture and powerful associative engine allow users to explore data intuitively, uncover hidden relationships, and make data-driven decisions with confidence.



Qlik's approach focuses on delivering connected solutions that span the entire data lifecycle - from ingestion and real-time replication to advanced analytics and AI-driven insights. This comprehensive approach helps organizations avoid fragmented analytics architectures and enables them to close the gaps between data, insights, and action.

There are 4 Capacity editions available with increasing capabilities from Starter, Standard, Premium and Enterprise.

Capacity Types			
	Generative AI - Assistants 200 Questions / 500 Pages Indexed	Generative AI - Assistants 1,000 Questions / 6,000 Pages Indexed	Generative AI - Assistants 5,000 Questions / 20,000 Pages Indexed
	Reporting 1,000 Automation 3rd Party 100	Reporting 5,000 Automation 3rd Party 1,000	Reporting 25,000 Automation 3rd Party 5,000 40GB Large Apps
	QTC Relational & SaaS Sources to Cloud	QTC Relational & SaaS Sources to Cloud	QTC All Sources to Cloud
	Business Intelligence 25GB max 5GB Apps, 1 Tenant Local Region 0 Managed Spaces, 5 Shared Spaces Standard Connectors Standard Automations 50,000	Business Intelligence 5GB Apps, 1 Tenant Local Region 10,000 Managed & Shared Spaces Standard Connectors Standard Automations 100,000	Business Intelligence 15GB Apps, 25 Tenants (Multi-Region) 10,000 Managed & Shared Spaces Standard Connectors + SAP / Uniq Standard Automations 400,000
Starter	Standard	Premium	Enterprise
10 User Pack (1) + Individual User Purchases	25 GB Data Packs	25GB / 250GB Data Packs (min. 50GB purchase)	250 GB Data Packs

Fusion for Qlik

Fusion for Qlik enables organisations to deliver fully personalised mashup portals for Qlik Cloud Analytics, helping to drive adoption and maximise the value of data across the enterprise. By surfacing trusted Qlik insights to employees, suppliers, and customers through a branded, no-code portal, Fusion ensures that analytics are accessible, intuitive, and aligned with your organisation's identity.

Fusion has been successfully deployed within the NHS, simplifying how nurses, clinicians, and frontline staff access and use data to improve patient outcomes. It also supports organisations that need to share data externally in real time, providing a secure, web-based interface that makes analytics easy to consume for third parties.

Whether the goal is to increase engagement with Qlik Cloud, improve data accessibility, or make insights more actionable, Fusion provides a powerful and flexible solution.

Key Features

- Seamless integration with Qlik Sense for a unified analytics experience.
- Personalised dashboards and tailored user journeys to meet the needs of different audiences.
- Secure, scalable access for internal teams and external partners.
- Rapid time-to-value with a proven return on investment.

A free trial is available on request.

Athena NHS Command Centre

The Athena Command Centre is an NHS and Catalyst BI joint collaboration that has successfully reduced triage and discharge times.

It consists of 5 live-view wall-mounted touchscreens being fed with Ambulance, Emergency Department in/out, Live Bed State and Inpatient out live data streams, all interactive with drill down to patient level.

The Athena Command Centre enables rapid analysis for flow coordinators supporting and enhancing on-the-ground decision-making throughout the day.

How does the Athena Command Centre Help?

The ability to see everything in one place; from ambulances enroute to the hospital to individual patients being discharged, enabling clinicians to identify problems before they have occurred.

In addition, the predictive technology helps plan for future capacity issues before they even occur.

Command Centre Consultancy Outcomes

The percentage of patients being cared for on non-medical wards, due to capacity pressures, has been reduced from 8% to 2%

95% of patients now triaged within 15 minutes of arrival at the Emergency Department (up from 67%)

The proportion of patients discharged by midday has grown from 16% to 19%

Chat-A-Lyst – Talk to my Data

Managing vast amounts of product or policy related documentation can be time consuming and inefficient. Employees often waste valuable hours searching through technical manuals, safety data sheet, installation guides and more to answer simple product or policy related queries.

The Chat A Lyst is a smart AI assistant that makes accessing documentation fast, easy and reliable by:

- **Ingesting and processing documents** like manuals, spec sheets, safety data sheets, HR policies and installation guides
- **Breaks documents into manageable sections** for quick and easy retrieval
- **Uses vector embeddings** to allow the bot to search and understand information based on meaning, not just key words
- **Extracts metadata** like product codes for efficient filtering and indexing, enhancing search relevance
- **Interface with users** via chat platforms such as Microsoft Teams, Slack or a Streamlit app in Snowflake
- **Generates real-time accurate** responses to product or policy related queries using Generative AI

Snowflake – Data Cloud

Catalyst is a certified and award winning Snowflake Public Sector Partner of the Year 2025, delivering Snowflake across both public and private sector organisations to scale out the use of data across single and multiple organisations. Catalyst BI's cloud data services provide organizations with the tools they need to unite fragmented data, seamlessly discover and securely share governed data, and execute complex analytic workloads with ease. No matter where your data or users are located, Snowflake's Cloud Data Platform delivers a unified experience that spans multiple public clouds including in the UK.

The Snowflake AI Data Cloud Platform is a globally available platform for all data and AI workloads. Snowflake's cloud-native architecture was built from the ground up to be easy to use, endlessly adaptable, interoperable, and future-proof.

From its inception, Snowflake was designed to be a fully managed, unified software-as-a-service (SaaS) solution, eliminating the administrative overhead and infrastructure management of traditional platforms. This allows teams to focus on deriving value from their data rather than spending time on manual tuning and maintenance. The platform's guiding principles—from separating compute and storage to its multi-cloud, single-governance model—ensure a cost-effective and efficient experience as organizations grow.

Snowflake's adaptable architecture connects disparate data silos, whether they are on-premises or across different cloud providers like AWS, Azure, and GCP. It seamlessly handles structured, semi-structured, and unstructured data, offering interoperable storage that works with data in Snowflake's optimized format or in open table formats within the customers' own cloud storage, eliminating the need for complex and costly data movement processes.

As new technologies like Generative AI emerge, Snowflake integrates them directly into the platform, so that organizations can take advantage of the next disruption without re-architecting their environment. Snowflake designed the platform to be a connected ecosystem, facilitating secure data sharing and collaboration both internally and with external partners. By providing a single secure, governed, and cost-effective platform, Snowflake empowers organizations to build for today and adapt for tomorrow.

Quickstart Program For Public Sector

Snowflake offers Quick Start Packages for the Public Sector. Three offers, Quickstart, Quickstart+, and Accelerator, are available with fixed pricing.

- The Quickstart package serves as the basic entry-level starter offer, and includes:
 - Access to +30 on-demand Snowflake Training Modules
 - 1 seat in Snowflake Fundamentals Course
- The Quickstart+ package is a more enhanced, mid-tier offering for customers, and includes:
 - Access to +30 on-demand Snowflake Training Modules
 - 2 seats in Snowflake Fundamentals Course
 - QuickStart Accelerator
- The Accelerator is the most comprehensive and top-tier package available.
 - Access to +30 on-demand Snowflake Training Modules
 - 3 seats in Snowflake Fundamentals Course
 - QuickStart Accelerator

QuickStart Accelerator is a 2 week series of interactive workshops led by a Snowflake Professional Services Solutions Architect, designed to begin implementing a customer's primary use-case on Snowflake and build customer confidence.

The Snowflake Service is described in more detail in the Snowflake [Documentation](https://docs.snowflake.com/en/) (<https://docs.snowflake.com/en/>). The Snowflake Documentation includes the information you need to access your Snowflake account and perform the administrative and user tasks associated with using Snowflake. The documentation also provides conceptual overviews, tutorials, and a detailed reference for all supported SQL commands, functions, and operators. To protect the Snowflake Service and the customers using our multi-tenant public cloud platform, Snowflake maintains a common [security policy](https://www.snowflake.com/en/legal/addenda/security-addendum/) (<https://www.snowflake.com/en/legal/addenda/security-addendum/>) and UK GDPR based [DPA](https://www.snowflake.com/en/legal/addenda/data-processing-addendum/) (<https://www.snowflake.com/en/legal/addenda/data-processing-addendum/>) which customers will be asked to review when entering into an Order Form. We also anticipate asking customers to review: (i) rights of Service suspension required for Service protection (e.g. a security issue), (ii) a SLA regime covering service failure and loss of use, (iii) aspects concerning deliverable, e.g. that no “Project Specific IPR” will be delivered, (iv) limiting Service usage to “Official” information only (i.e. not Secret/Top Secret), and (v) other relevant clarifications.

Security Secure your environment with continuous risk monitoring and protections, RBAC and granular authorization policies:

- Trust Center (<https://docs.snowflake.com/en/user-guide/trust-center/overview>) allows customers to monitor and manage risks 24/7 across your Snowflake accounts and cross cloud service providers. The Trust Center evaluates your account against recommendations specified in scanners according to a schedule you control. If your account violates any of the recommendations in any of the enabled scanners, then the Trust Center provides a list of security risks, and information about how to mitigate those risks.
- Safeguard data with built-in platform security capabilities across areas like authentication, encryption, connectivity and networking
- Enforce row and column-level security across clouds with RBAC and policies, eliminating the need to manage multiple versions of the same data

Snowflake Platform Architecture

Snowflake is the only cloud data platform with all components tightly integrated. Snowflake is delivered as a service, meaning that Snowflake takes care of managing the underlying cloud infrastructure, data protection, security and service tuning, so that you can focus on what’s important: getting value from the data.

Responsibility Matrix: Snowflake vs Other Data Lakes & Data Warehouses				
	Task	Ordinary on-prem DW/DL	Ordinary cloud DW/DL	Snowflake
Infrastructure	Hardware	Customer	Cloud vendor	Cloud Data Platform as a Service
	Software	Distribution	Cloud vendor	
	Hardware cluster setup	Customer	Customer	
	Software provisioning	Tools	Tools	
Data & service protection	Data protection & retention	Platform or Add-on	Platform or Add-on	
	Node failure protection & recovery	Customer	Customer	
	Disaster recovery	Customer	Customer	
	Service monitoring & alerting	Customer	Customer	
Security	Physical security	Customer	Cloud vendor	
	Deployment security	Customer	Customer	
	Security monitoring	Customer	Customer	
Database management & tuning	Compute scaling	Customer	Customer	
	Index management	Customer	Customer	
	Data partitioning	Customer	Customer	
	Metadata & statistics management	Customer	Customer	
	Query optimization	Customer	Customer	

Snowflake implements a multi-cluster, shared data platform on Amazon Web Services, Microsoft Azure, and Google Cloud Platform. At its core, Snowflake is a columnar-store ANSI-SQL relational database, but its unique architecture enables a variety of workloads beyond traditional data warehousing, including lakehouse analytics, data engineering, AI, applications and collaboration.

Core components of the Snowflake architecture:

- Inside the Snowflake Platform architecture, the Interoperable Storage layer gives customers the choice whether to utilize Snowflake-managed or customer-managed cloud storage. When data is stored in Snowflake-managed storage, Snowflake automatically manages encryption, compression, partitioning, and clustering of the data in the storage layer. For customers who choose to store their data externally in open table formats, Apache Iceberg™ tables for Snowflake combine the performance and query semantics of typical Snowflake tables with external cloud storage that you manage.
- The Multi-Cluster Compute layer represents the independently scalable compute clusters that process queries and updates against the underlying Snowflake database. Snowflake compute clusters elastically scale vertically, horizontally and across workloads and are configured by default to shut down when demand falls to zero for a specified period of time. The Adaptive Compute service automatically adjusts compute resources based on the workload demands. Snowflake's compute can execute workloads in SQL, Python, Java, Scala and Javascript, so there is no need to move the work and data between different compute engines. Both CPU and GPU compute resources are available based on the workload demands.
- Cortex AI delivers a fully integrated layer of Generative AI services, including state-of-the art LLM models hosted on scalable GPU infrastructure, and easy-to-use tools to leverage them inside your pipelines and workflows. Use AI SQL for powerful multimodal LLM-based analytics directly inside your SQL queries; process documents at scale with Document AI; leverage Snowflake Intelligence for AI agent-powered conversational analytics over your structured and unstructured data with unparalleled accuracy.
- The Cloud Services layer sits over the storage, compute, and AI layers to manage and optimize the platform's operations. As workloads are added or removed from the Snowflake platform, the cloud services layer automatically scales to support optimization, security, transactional consistency, and

metadata management with near-zero administration required. All you have to worry about is designing, loading, and querying your data, and Snowflake takes care of the rest.

- Finally, the Snowgrid is a cross-cloud technology layer that allows you to distribute your data and operations across the regions and cloud providers of your choice while maintaining the same Snowflake experience regardless of complexity. The Snowgrid enables a multi-cloud strategy, including a cross-cloud approach to failover and disaster recovery for unmatched business continuity, and to mix and match clouds as you see fit. Powered by Snowgrid, Snowflake Marketplace provides instant access to thousands of datasets and applications that organizations rely on – the most relevant and valuable content is ready to be mobilized to drive better decisions, products and services, and business outcomes.

Business Continuity

The Snowflake Data Cloud was designed from the ground up to be a continuously available cloud service, resilient to failures, disruption, or data loss. Its continuous data protection (CDP) capabilities protect against (and provide easy self-service recovery from) accidental errors, system failures, and malicious acts.

- Customer error. To mitigate accidental data deletion or corruption, Snowflake Time Travel enables accessing historical data (i.e. data that has been changed or deleted) at any point within up to 90 days. Historical data can be retrieved and restored with a simple SQL query. Fail-safe provides an additional 7 days period during which historical data may be recoverable by Snowflake.
- Infrastructure failures – from a single instance of virtual machine failing to an Availability Zone outage. The Snowflake architecture includes triple redundancy of all critical services, including built-in redundancy across three availability zones within a cloud region.
- Regional or multi-region failures. Customers with critical workloads can set up cross-region/cross-cloud replication and failover/failback including client redirect with just a few SQL commands - no need to set up complex processes and HA/DR infrastructure.

Snowflake Cost

Snowflake is a fully managed and fully integrated service. With other solutions, customers receive separate bills for the data warehouse/data lake service and for the cost of the underlying infrastructure such as storage, compute, file access and support. Snowflake customers receive one bill, with all costs billed directly by Snowflake.

The total cost of using Snowflake is the aggregate of the cost of using compute resources, storage and optional services such as snowpipe and materialized views. There are no data egress charges with Snowflake as long as customers use one of the many drivers and connectors to access the data and do not export files directly from Snowflake to another cloud or region object storage. Snowflake Support is also included in the cost of a Snowflake credit. To learn more, visit <https://docs.snowflake.com/en/user-guide/cost-understanding-overall>.

Snowflake provides many built-in cost management capabilities to help customers understand, control and optimize costs:

1. Visibility – understand, attribute, monitor spend:

- The Cost Management UI provides an account-level consumption and spend, consumption trend for virtual warehouses, most expensive queries and more.
- Customers can set up account- or project-specific budgets that define a monthly spending limit for a group of Snowflake objects. When the system projects the spending limit will be exceeded, Snowflake sends a daily notification email to the Snowflake admin.
- Snowflake automatically monitors consumption trends for anomalies and can be set up to send alerts when unusual activity is detected.
- Customers can query the usage views to drill down into cost data and generate custom reports and dashboards. The Query Insights view provides actionable recommendations for query optimization - such

as pointing out an incomplete join condition generating a cartesian join, or spillage due to an undersized compute resource.

2. Control – Limit and control spend.

- Snowflake virtual warehouse (compute) sizes provide implicit upper bounds for the credit burn rate and cost. In addition, virtual warehouses auto-suspend and auto-resume by default so customers only pay for what they use.
- Resource Monitors provide alerting and hard limits for individual virtual warehouses during a specific time interval. The virtual warehouse can be configured to suspend once the credit limit is hit.

3. Optimization – Identify inefficiencies and optimize spend.

- Since Snowflake's founding in 2012, we strive to be the market leader in fast performance for all workloads, passing the optimizations to the customers automatically and increasing the value of a Snowflake credit. The Snowflake Performance Index measures the improvement in query runtime and other metrics based on these improvements. To learn more, visit <https://www.snowflake.com/en/pricing-options/performance-index/>.
- Examples of optimization features include Snowpark-optimized warehouses, which reduce the time and cost of memory-intensive workloads; serverless capabilities such as serverless tasks, query acceleration service, automatic clustering, materialized views and search optimization, which deliver performance improvements for specific query patterns and help save on costs over time.
- Recently introduced Generation 2 (Gen2) virtual warehouses run on upgraded hardware and deliver 2.1x faster performance YoY, with 2.3x higher throughput for Concurrent BI workloads and up to 4.4x faster Delete, Update and Merge operations.

Data Architecture Patterns With Snowflake

Organizations in every industry acknowledge that data is one of their most important assets. However, many are falling short of realizing the potential of their data for three common reasons:

- Even with the increasing adoption of the cloud to store data at virtually infinite volumes, organizations still struggle to centralize their data. Silos are everywhere—architectures, clouds, business entities, languages, data types, and table formats.
- To account for these silos, data engineers are tasked with creating complex, error-prone pipelines. An organization may be able to store massive amounts of data in their data lake, but processing and querying that data is still either slow, expensive, or both. And that's just first-party data. Incorporating shared, third-party data can be just as difficult and expensive: requiring making a copy to ETL to your data lake.
- And with lots of systems and lots of copies, governance is a nightmare and collaboration is difficult. This presents risk for compliance, and leads to unknown or poor data quality. Users can't easily find data, trust it's what they need, and then share to collaborate.

Snowflake's adaptable architecture allows you to truly manage all data together and unify the silos:

- Cross-region, cross-cloud, cross-entity
- Across programming languages
- Across structured & unstructured data
- Across compute engines and analytical/transactional workloads


Having flexibility in architectural patterns is critical to not only address organizations' business needs today but also will let them adapt to changing needs in the future. With Snowflake, you get one platform that can support a wide range of architectural patterns, including:

- Data warehouse: Centralized analytics for structured business data
- Data lake: Unlimited storage for versatile data types and workloads

- Data lakehouse: Transactional data lake for unified analytics, AI/ML, collaborative workloads
- Data mesh or fabric: Distributed and governed, domain-oriented collaboration
- Transactional applications: Build and run applications directly on Snowflake with Unistore's Hybrid Tables or leverage with Snowflake PostgreSQL for low-latency, transactional needs.

Data Lake/Lakehouse

The purpose of data lakes is to liberate business data from applications that generate it, and make it widely available for analytics. The standard pattern replicates data from transactional systems into a raw or staging area in Snowflake and uses Snowflake's compute power to transform and enrich the data to create curated tables and data marts that end users use to run reports and create predictive models. This pattern delivers the most benefits for best performance, security, and automatic management. Snowflake-managed storage offers a reduced total cost of ownership, better time to value, and limits the risk of undesired data exposure and compliance risks. With Snowflake managing the secure data layer, customers can focus on business objectives over a myriad of technical decisions (e.g. partition size, encryption, access control).



Snowflake Horizon Catalog is embedded in the Snowflake service and provides out-of-the-box transactional consistency, unified security, governance and discovery across regions and clouds. Whenever you load data or run queries in Snowflake, Horizon operates behind the scenes without any configuration or intervention from you.

Some organizations prioritize interoperability with other compute engines such as Apache Spark, Flink and Trino and therefore don't mind the additional administrative burden of managing their data in customer-managed storage (e.g. partition size, encryption, access control, etc). For such customers, Snowflake's extensive support of Apache Iceberg™ table format delivers this interoperability. Iceberg tables for Snowflake combine the performance and query semantics of typical Snowflake tables with external cloud storage that you manage.

Horizon Catalog federates across Iceberg REST catalogs (including the open source Apache Polaris catalog) through Catalog-linked Databases, which automatically syncs Iceberg objects managed by remote catalogs in Horizon Catalog. Data engineers can act on their entire Iceberg ecosystem from a single pane of glass, ensuring all read and write operations stay in sync and are protected by Horizon Catalog's leading governance and discovery capabilities.

Data Mesh

Snowflake's platform provides a robust foundation for implementing a data mesh architecture by directly supporting its core principles. It enables domain-oriented ownership by allowing individual business domains to manage their data and infrastructure independently within a single, unified platform. The cornerstone of its support is Secure Data Sharing, which facilitates the creation of data as a product. This feature allows domains to share live, ready-to-query data directly from their accounts with other domains without creating complex ETL pipelines or data copies, ensuring consumers always have access to fresh, governed data.

Furthermore, Snowflake enables federated computational governance through features like dynamic data masking, row-access policies, and object tagging. These tools empower domain teams to apply fine-grained security and governance controls locally on their data products while adhering to global organizational standards. With the integration of Snowpark and Streamlit in Snowflake, data products can extend beyond simple datasets to include interactive applications and complex data processing logic, all discoverable and accessible within the same ecosystem. This combination of secure, live sharing and decentralized governance makes Snowflake an effective platform for building and scaling a data mesh.

Applications

For decades, transactional and analytical data have remained separate, significantly limiting how fast organizations could evolve their businesses. Snowflake is changing this by unifying these disparate worlds, allowing you to build and run applications directly on the same platform as your analytics. This transformation is powered by a suite of foundational innovations: Unistore for hybrid transactional-analytical tables, a fully managed Postgres service (based on the Crunchy Data acquisition) for OLTP applications, Snowpark Container Services for running any containerized logic, and the Snowflake Native App Framework to package and deliver these new solutions securely to customers.

With Unistore, organizations can use a single, unified data set to develop and deploy applications, and analyze both transactional and analytical data together in near-real time. Underpinning the Unistore workload is a hybrid table - a Snowflake table type that is optimized for hybrid transactional and operational workloads that require low latency and high throughput on small random point reads and writes. Hybrid tables accelerate application development by accessing analytical and transactional data with a single query and provide the following key benefits:

- **Fewer Databases:** Hybrid tables run double-digit millisecond point operations alongside your analytical queries, all within Snowflake. Reduce the need to operate and synchronize across multiple databases. With a single governed dataset for both your transactional and historical data, your application data is ready to query, no ETL required.
- **Get started quickly:** Query hybrid tables with the syntax you are used to in Snowflake. Hybrid tables enforce primary keys and referential integrity constraints, as well as provide secondary indexes for accelerated lookups. Snowflake governance capabilities like masking policies and row based access policies work out of the box.
- **Easily bring transactional data to your analytical workloads:** With your transactional and analytical data in a single platform, you can enrich your analytics and create new experiences.

In addition to Unistore, Snowflake is in the process of building a fully managed Postgres service based on the recent acquisition of Crunchy Data, a leading open source Postgres provider. This service is engineered to run OLTP workloads natively within the Snowflake AI Data Cloud and will feature full wire-protocol compatibility with open-source PostgreSQL. This compatibility is critical, as it allows developers to migrate existing applications with minimal to no code refactoring. Architecturally, the service facilitates seamless, high-performance data interoperability between the managed Postgres instances and Snowflake's core analytical engine, eliminating the need for separate ETL pipelines. Crucially, this integration extends Snowflake's unified governance and security model to Postgres, allowing a single control plane for applying granular access policies, data masking, and object tagging across both transactional and analytical data estates.

The Snowflake Native App Framework is a developer framework that enables providers to build, distribute, and monetize Snowflake Native Apps on Snowflake Marketplace. Snowflake Native Apps run inside the end-consumer's Snowflake account.

Using the framework, providers can build applications with Snowflake functionalities like UDFs, stored procedures, and Snowpark. Functionalities have been extended with the capability to provide low-code UX framework through

the Streamlit Integration and to empower providers to craft more complex apps via the integration with Snowpark Container Services.

Providers can deploy and operate Snowflake Native Apps with out-of-the-box versioning and monitoring support, distribute to Snowflake Marketplace through public or private listing, and leverage billing infrastructure to monetize apps using customizable billing models.

Consumers can purchase apps with a credit card or their Snowflake credit, install and run these apps within their Snowflake accounts, running directly against their data securely in the Data Cloud.

Snowpark Container Services is a fully managed container offering designed to facilitate the deployment, management, and scaling of containerized applications within the Snowflake ecosystem. This service enables users to run containerized workloads directly within Snowflake, ensuring that data doesn't need to be moved out of the Snowflake environment for processing. Unlike traditional container orchestration platforms like Docker or Kubernetes, Snowpark Container Services offers an OCI runtime execution environment specifically optimized for Snowflake. This integration allows for the seamless execution of OCI images, leveraging Snowflake's robust data platform.

As a fully managed service, Snowpark Container Services streamlines operational tasks. It handles the intricacies of container management, including security and configuration, in line with best practices. This ensures that users can focus on developing and deploying their applications without the overhead of managing the underlying infrastructure. Snowpark Container Services is fully integrated with Snowflake and with third-party tools like Docker. Seamless integration makes it easier for teams to focus on building data applications.

Snowpark Container Services unlocks a wide array of new functionality, including these features:

- Create long-running services.
- Use GPUs to boost the speed and processing capabilities of a system.
- Write your application code in any language (for example, C++).
- Use any libraries with your application.

The integration of Snowpark Container Services with the Snowflake Native App Framework fundamentally expands the types of applications providers can build and monetize. It allows developers to package complex logic, written in any language (e.g., C++, Python, Java) and using any library, into a container. This container can then be deployed as part of a Native App, enabling features like long-running services or GPU-accelerated processing to run directly and securely against a consumer's data within their own Snowflake account. All of this comes with Snowflake platform benefits, most notably ease-of-use, security, and governance features. And you now have a scalable, flexible compute layer next to the powerful Snowflake data layer without needing to move data off the platform.

Data Ingestion and Transformation

Data ingestion & transformation is a key component of a robust data platform. A data platform has limited value if you can't bring in the data. Snowflake provides an extensive array of data ingestion and transformation services, giving you the choice and convenience to build reliable, scalable data pipelines tailored to any source, format, or latency requirement.

Openflow Snowflake

Openflow is a new managed data integration service within the Snowflake platform, built on Apache NiFi, designed to connect and process data from virtually any source into Snowflake and other destinations. Openflow unifies the entire data pipeline — from ingestion of batch, streaming, and unstructured data to transformation — within a single, governed platform, eliminating the need for separate ETL tools and accelerating the development of AI applications.

Key benefits of Openflow:

- Open and extensible: An extensible managed service that's powered by Apache NiFi, enabling you to build and extend processors from any data source to any destination.
- Unified data integration platform: Openflow enables data engineers to handle complex, bi-directional data extraction and loading through a fully managed service that can be deployed inside your own VPC in the cloud or on-premises.
- Enterprise ready: Openflow offers out-of-the box security, compliance, and observability and maintainability hooks for data integration.
- High speed ingestion of all types of data: One unified platform that lets you handle structured and unstructured data, in both batch and streaming modes from your data source to Snowflake at virtually any scale.
- Continuous ingestion of multimodal data for AI processing: Near real-time unstructured data ingestion, so you can immediately chat with your data coming from sources such as Sharepoint, Google Drive, and so on.

Architecturally, Openflow is comprised of:

- Control plane: runs in Snowflake as a managed service, allowing customers to provision, command, control data plane.
- Data Plane: Openflow connectors and pipelines run in the data plane environment. You can choose to run the Dataplane in your own VPC via Bring Your Own Cloud (BYOC) or in Snowflake as a fully managed service. Data planes can scale horizontally, support multi-node and multi-cluster deployments, and DR resilience. You can have multiple data planes tied to the control plane, serving different use cases.
- Openflow Connectors: Highly curated, fully managed, first-class, secure, governed experiences against key systems (unstructured data, SaaS systems, databases, streaming etc). Connectors are extensible - you can create custom connectors and run them in the data plane.
- OSS Nifi: Openflow is powered by the open source Nifi data integration framework that is trusted by thousands of enterprise customers. Openflow super charges Nifi.

Native Connectors

Snowflake and partner-delivered connectors to a variety of data sources can be provisioned from the Snowflake Marketplace so that you can start replicating data in minutes.

Snowpipe and Snowpipe Streaming

Snowflake supports low-latency, continuous data ingestion via micro-batch ingestion (using Snowpipe) and streaming (using Snowpipe Streaming). Data is ingested automatically from sources such as cloud object storage (AWS S3, Azure blob, Google Cloud Storage) and streaming platforms like Kafka. Loading happens continuously as new data is generated, eliminating the need for scripts and scheduling tools. The COPY INTO command can be used to perform the initial bulk load.


Partner ETL tools

Snowflake supports a vast community of extract, transform, and load tools (ETLs) to streamline the continuous movement of data into and out of the Snowflake cloud-native platform. Companies such as Fivetran, Matillion, Qlik and Informatica have created tight integrations with Snowflake through dedicated connectors. See the Snowflake ecosystem for more details: <https://docs.snowflake.com/en/user-guide/ecosystem>.

Data Sharing By far the easiest way to get data into Snowflake is through secure data sharing between accounts of Snowflake. With secure data sharing, no actual data is copied or transferred between accounts. This unique architecture provides the foundation for privacy-preserving collaboration, as organizations can grant secure access to live data while maintaining full control and applying governance. Because no data is copied or exchanged,

secure data sharing setup is quick and easy for providers and access to the shared data is near-instantaneous for consumers. You can share data, semantic views, AI models, and applications.

Data Clean Rooms leverage Snowflake's data sharing capabilities and allow you to combine and analyze data from different parties with privacy-preserving configurations that help protect the underlying data. Delivered as a native solution, clean rooms makes it easy to build, connect, and use data from multiple parties, while applying enhanced privacy and increased security.



Snowflake Marketplace - another feature that leverages Snowflake's data sharing capabilities - provides access to data products from hundreds of data providers for categories ranging from weather to demographics to economics to healthcare and energy.

The Internal Marketplace in Snowflake is a curated, secure space for collaborative data sharing within an organization. It centralizes internally-available data products, allowing teams to easily discover and use vetted data shared within your internal teams. Access can be managed by account targeting and Role-Based Access Control (RBAC), ensuring that data remains secure and accessible only to authorized users.

dbt Projects

dbt Core is an open-source data transformation tool and framework that many customers use to define, test, and deploy transformations. A favorite among data teams, dbt offers the ability to bring software engineering best practices and improved efficiency in SQL and Snowpark data transformation workflows right in Snowflake. dbt projects in Snowflake enable data teams to build, run and monitor dbt projects directly in Snowsight UI, reducing context switching, simplifying setup and accelerating the development lifecycle of data pipelines.

With dbt Projects support, you can:

- Enable new teams to build and deploy pipelines with uniform governance: Accelerate onboarding and empower new teams to create the pipelines they need through an intuitive interface with a uniform governance and security model for data and pipelines.
- Consolidate systems to reduce administration and improve debugging: Run dbt natively on Snowflake and streamline the pipeline development lifecycle to improve developer productivity and quickly spot and address problems.

Dynamic Tables

As the raw data lands in Snowflake, dynamic tables are the building blocks of declarative data transformation pipelines. They significantly simplify data engineering in Snowflake and provide a reliable, cost-effective, and automated way to transform your data incrementally using embedded change data capture (CDC). Instead of defining data transformation steps as a series of tasks and then monitoring dependencies and scheduling, you can

simply define the end state of the transformation using dynamic tables and leave the complex pipeline management to Snowflake.

Streams and Tasks

Snowflake streams and tasks, in combination with stored procedures and UDFs, enable you to design sophisticated data pipelines, particularly when the logic requires the use of Python, Java or Scala libraries.

The Snowflake UI makes monitoring data transformation pipelines easy by providing graphical observability interfaces displaying the pipelines, their status, dependencies and refresh history.



Migrating to Snowflake

Traditional data platforms and big data solutions often fail to achieve their main goal: allowing users to work with data without restrictions on size, speed, or adaptability. Snowflake provides a robust migration framework designed to help organizations seamlessly transition from legacy data platforms and modernizing their data analytics workloads.

SnowConvert is a user-friendly tool that helps you modernize your traditional data platform by migrating it to Snowflake. SnowConvert analyzes your existing SQL code (from platforms like Oracle, SQL Server, and Teradata) and converts it to Snowflake SQL. After conversion, you can immediately take advantage of all Snowflake's features and capabilities. Currently, SnowConvert supports converting code from these source platforms:

- Teradata
- Oracle
- SQL Server
- Sybase IQ
- Redshift
- Azure Synapse

While SnowConvert automates a high percentage of the initial code conversion, some complex or non-standard code may still require manual review. SnowConvert AI extends the power of the core tool by tackling this final stage of the migration.

Integrated into Visual Studio Code as the "SnowConvert Migration Assistant," it acts as an AI-powered assistant that helps developers resolve the issues that SnowConvert could not automatically fix. Powered by Snowflake Cortex AI, it intelligently analyzes the remaining code, providing natural language explanations for conversion problems and suggesting specific code fixes. By streamlining this often complex and manual cleanup process, SnowConvert AI dramatically reduces the overall complexity, costs, and timelines associated with data warehouse, BI, and ETL migrations, making the journey to the Snowflake AI Data Cloud more accessible than ever.

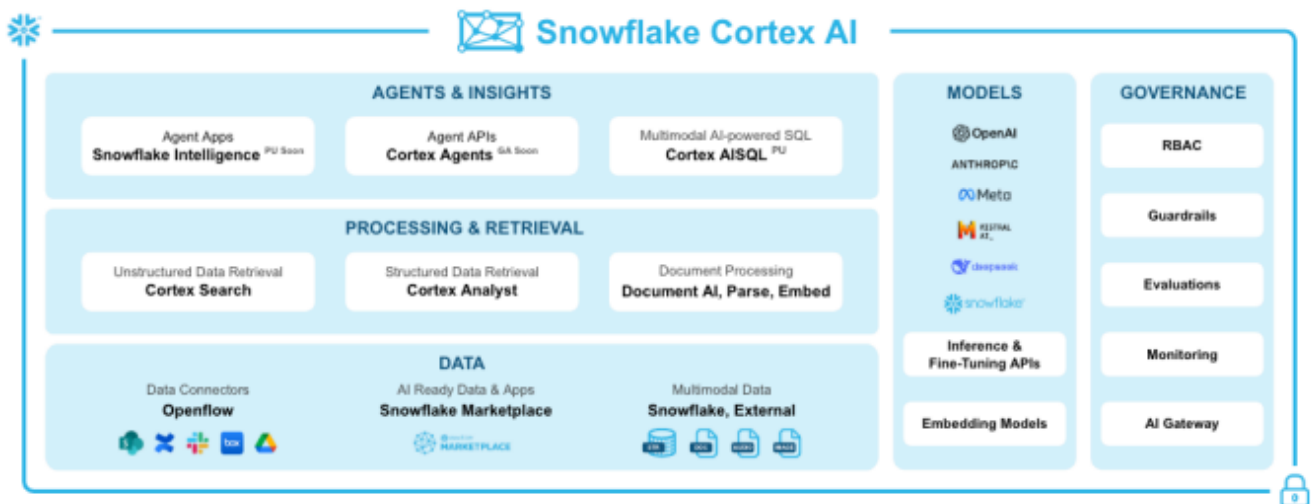
Unified Platform For End-To-End Ai, ML and Bi

Snowflake delivers a comprehensive, end-to-end platform that supercharges AI and ML by unifying its powerful capabilities into a single, fully managed service. The platform combines Snowflake ML, which supports the complete predictive modeling lifecycle (from feature engineering and model training to deployment and monitoring) with Snowflake Cortex AI, which provides ready-to-use, fully managed LLM models and features that make it easy to deploy generative AI applications on your data. Because these capabilities are fully managed and deeply integrated, customers can build, deploy, and govern everything from custom models to advanced AI agents directly on their data without the complexity of moving data or stitching together separate tools. This unified approach is built on three core value propositions: it's easy, reducing development costs by eliminating infrastructure management; it's connected, enabling better quality and operational efficiency by having all components work together seamlessly; and it's trusted, extending Snowflake's industry-leading security and unified governance to all data and models within one secure perimeter.

In addition, Snowflake delivers native data exploration and visualization experiences via Workspaces and Streamlit. For enterprise BI, Snowflake seamlessly integrates with third-party BI solutions such as Qlik, Power BI, Tableau, Sigma, Quicksight and more.

Snowflake Cortex

The world of business analytics is experiencing its most significant transformation in generations. Generative AI is at the heart of this evolution, fundamentally democratizing enterprise AI and redefining user experiences. The focus is shifting from complex dashboards to intelligent, conversational AI agents that allow anyone to ask questions in natural language and get trusted answers. Snowflake is leading this charge with an easy, efficient, and trusted platform where these capabilities are fully managed and integrated. Snowflake Cortex - a fully managed AI layer within Snowflake's secure perimeter - enables business users and data scientists to turn structured and unstructured data into actionable insights, without complex tooling or infrastructure. Whether customers are analyzing tables, processing documents, deploying AI agents, or training models, Snowflake integrates these into its secure platform with unified governance, eliminating the need for infrastructure management or complex tooling.



Snowflake Intelligence provides a unified conversational AI interface integrated directly within the Snowflake platform, enabling business users to analyze, summarize, and act on both structured and unstructured data using natural language. This interface eliminates the need for traditional dashboards or SQL, allowing business users to access and derive insights easily and reliably. For developers, it reduces technical friction by providing the flexibility to quickly customize experiences, integrate new data sources, and expand capabilities, ultimately helping them serve business users more efficiently and effectively.

Cortex Agents power the conversational experiences within Snowflake Intelligence by coordinating tasks across both structured and unstructured data sources to provide valuable insights. Cortex Agents plan their tasks, utilize various tools to execute them, and generate responses. They can leverage Cortex Analyst, Cortex Search, and large language models (LLMs) as tools for data analysis.

Cortex Agents use Cortex Analyst and Cortex Search as tools, along with LLMs, to analyze data. Cortex Search extracts insights from unstructured sources, while Cortex Analyst generates SQL to process structured data. A comprehensive support for tool identification and tool execution enables delivery of sophisticated applications grounded in enterprise data.

REST API make it easy to call Cortex Agents from external applications, and integration with Microsoft Teams and M365 Copilot allows users to fit the agents directly into their daily workflow, while staying within Microsoft and Snowflake's trusted security boundaries and working with your existing semantic models and data governance policies.


Cortex Analyst is a fully-managed, LLM-powered feature that helps you create applications capable of reliably answering business questions based on your structured data in Snowflake. With Cortex Analyst, business users can ask questions in natural language and receive direct answers without writing SQL.

Building a production-grade conversational self-service analytics solution requires a service that generates accurate text-to-SQL responses. For most teams, developing such a service that successfully balances accuracy, latency, and costs is a daunting task. Cortex Analyst simplifies this process by providing a fully managed, sophisticated agentic AI system that handles all of these complexities, generating highly accurate text-to-SQL responses. It helps you accelerate the delivery of high-precision, self-serve conversational analytics to business teams, while avoiding time sinks such as complex RAG solution patterns, model experimentation, and GPU capacity planning. The generated SQL queries are executed against the scalable Snowflake engine, ensuring industry-leading price performance and lower total cost of ownership (TCO).

Generic AI solutions often struggle with text-to-SQL conversions when given only a database schema, as schemas lack critical knowledge like business process definitions and metrics handling. Cortex Analyst overcomes this limitation by leveraging semantic models, which define business metrics and model business entities and their relationships. Cortex Analyst reads the information captured in the semantic model definition and generates the SQL against the physical tables directly.

Automatic semantic model generation makes creating semantic models for Cortex Analyst easier and faster. By analyzing schema metadata, query history and dashboards, it constructs performant, reusable models, eliminating weeks of manual work. Cortex Search is a feature that uses state-of-the-art models for fast retrieval and ranking of information from unstructured documents.

Cortex Search enables low-latency, high-quality “fuzzy” search over your Snowflake data. Cortex Search powers a broad array of search experiences for Snowflake users including Retrieval Augmented Generation (RAG) applications leveraging LLMs. Cortex Search gets you up and running with a hybrid (vector and keyword) search engine on your text data in minutes, without having to worry about embedding, infrastructure maintenance, search quality parameter tuning, or ongoing index refreshes. This means you can spend less time on infrastructure and search quality tuning, and more time developing high-quality chat and search experiences using your data. Check out the Cortex Search tutorials for step-by-step instructions on using Cortex Search to power AI chat and search applications.



AISQL

Cortex AISQL embeds AI capabilities directly into SQL environments, introducing native AI operators like `AI_FILTER`, `AI_JOIN`, and `AI_AGG` within the Snowflake ecosystem. This allows users to build sophisticated analysis pipelines across multiple data types (text, image, and soon audio) without needing a separate AI stack. It empowers data analysts to perform advanced tasks, such as semantic joins and multi-label classification, all within their familiar SQL interface.


Document AI is a fully managed service that uses the powerful Arctic-TILT, a proprietary large language model to automatically extract information like text, table values, and handwritten content from documents such as PDFs. It is designed to improve efficiency and reduce the costs and errors associated with manual data entry. Business users can take advantage of a no-code interface to test and refine the AI's extraction capabilities, while data teams can use SQL to extract data from multiple documents at once and load it directly into Snowflake tables.

Snowflake ML

Snowflake ML is an integrated set of capabilities for end-to-end machine learning in a single platform on top of your governed data. This is a unified environment for ML development and productionization that is optimized for large-scale distributed feature engineering, model training and inference on CPU and GPU compute without manual tuning or configuration.

Scaling end-to-end ML workflows in Snowflake is seamless. You can do the following:

- Prepare data
- Create and use features with the Snowflake Feature Store
- Train models with CPUs or GPUs using any open-source package from Snowflake Notebooks on Container Runtime.
- Operationalize your pipelines using Snowflake ML Jobs.
- Deploy your model for inference at scale with the Snowflake Model Registry
- Monitor your production models with ML Observability and Explainability.
- Use ML Lineage provides to track the source data to features, datasets, and models throughout the ML pipeline



Snowflake ML is also flexible and modular. You can deploy the models that you’ve developed in Snowflake outside of Snowflake and externally-trained models can easily be brought into Snowflake for inference. **Snowflake Notebooks on Container Runtime** provide a Jupyter-like environment for training and fine-tuning large-scale models in Snowflake, without infrastructure management. Start training with preinstalled packages such as PyTorch, XGBoost, or Scikit-learn, or install any package from open-source repositories like HuggingFace or PyPI. Container Runtime is optimized to run on Snowflake’s infrastructure to provide you with highly efficient data loading, distributed model training, and hyperparameter tuning.

Snowflake’s Data Science Agent represents a significant advancement in machine learning workflow automation, designed to enhance the productivity of data scientists through natural language interactions. This innovative agentic companion streamlines the entire machine learning pipeline by automating crucial steps that traditionally require extensive manual coding and configuration. The agent assists data scientists by facilitating data preparation processes, automating feature engineering tasks, and streamlining model training workflows through intuitive natural language commands.

The Snowflake Feature Store is an integrated solution for defining, managing, storing and discovering ML features derived from your data. The Snowflake Feature Store supports automated, incremental refresh from batch and streaming data sources, so that feature pipelines need be defined only once to be continuously updated with new data. Use

Snowflake ML Jobs to develop and automate ML pipelines. ML Jobs also enable teams that prefer working from an external IDE (VS Code, PyCharm, SageMaker Notebooks) to dispatch functions, files or modules down to Snowflake’s Container Runtime.

The Snowflake Model Registry allows for the logging and management of all your ML models, regardless of whether they’re trained on Snowflake or other platforms. You can use the models from the model registry to run inference at scale. You can use Model Serving to deploy the models to Snowpark Container Service for inference.

ML Observability provides tools to monitor model performance metrics in Snowflake. You can track models in production, monitor performance and drift metrics, and set alerts for performance thresholds. Additionally, use the ML Explainability function to compute Shapley values for models in the Snowflake Model Registry, regardless of where they were trained.

ML Lineage is a capability to trace end-to-end lineage of ML artifacts from source data to features, datasets, and models. This enables reproducibility, compliance, and debugging across the full lifecycle of ML assets.

For business analysts, use ML Functions to shorten development time for common scenarios such as forecasting and anomaly detection across your organization with SQL. Snowflake AI & ML Studio provides intuitive graphical interfaces to many of the Snowflake Cortex and Snowflake ML components.

DataOps Live

DataOps.live is a modern DataOps platform designed to bring DevOps-style automation, governance, and agility to the entire data lifecycle. It enables organisations to build, test, deploy, and operate data products with the same discipline and speed that software engineering teams have enjoyed for years. By orchestrating everything from ingestion and transformation to testing and deployment, DataOps.live ensures that data pipelines are consistent, reliable, and production-ready.

At its core, DataOps.live provides a single, automated framework for managing complex data environments- particularly those built on the Snowflake Data Cloud. It eliminates manual processes, reduces operational risk, and accelerates delivery by introducing CI/CD, version control, automated testing, and environment management into the data workflow.

Working in partnership with Catalyst DataOps.live helps organisations modernise their data operations, improve governance, and deliver trusted data products faster and more efficiently.

Key Capabilities

- End-to-end automation of data pipelines, transformations, and deployments.
- CI/CD for data, enabling rapid, repeatable, and controlled releases.
- Environment management that allows teams to spin up, test, and promote changes safely.
- Automated testing and validation to ensure data quality and prevent production issues.
- Deep integration with Snowflake, optimising performance, governance, and scalability.
- Version-controlled development, ensuring full traceability and auditability.
- Data product lifecycle management, supporting modular, reusable, and governed data assets.

Together, Catalyst BI and DataOps.live empower organisations to transform how they build and operate data platforms- delivering trusted, governed, and production-ready data at speed and scale.

DataRobot

Catalyst BI has extensive experience of delivering successful data-driven projects in the public sector, and with specific focus on the National Health Service. Catalyst BI is the only UK based Managed Services (MSP) of DataRobot, the industry-leading automated machine learning platform, and Catalyst is proud to hold the highest Platinum level partner status.

Catalyst BI can deliver DataRobot Enterprise solutions (both Predictive and Generative AI) in conjunction with DataRobot to organisations that already exhibit data science expertise and may even employ data scientists or data analysts. The DataRobot platform allows you to scale your data science capabilities, build and deploy machine learning models, and automate the delivery of models, ensuring the best models are deployed for any data science problem.

Once you've built your model, it's ready for production and can be deployed with a single click or line of code to the environment of your choice. With Machine Learning Operations (MLOps), you can have a single place to manage and monitor any model, regardless of where it was deployed. With Continuous AI, you can create multiple MLOps retraining strategies to refresh your production models based on the schedule of your choosing—like when accuracy drops below a predetermined threshold or data drift occurs. After you deploy your model, you can interpret the results in an AI-powered app, such as Qlik. This helps to simplify AI consumption and empower frontline decision-makers, resulting in trusted, fully operationalised data science solutions. “Explainability” is critical in healthcare, and DataRobot's platform makes it easy for clinicians to understand DataRobot's predictions and act on those insights.

Catalyst AI Cloud

Additionally, Catalyst offers our Managed Service to support and drive early-stage adoption of data science outcomes within your organisation. This is called Catalyst AI Cloud, where Catalyst manages the development, deployment, and maintenance of the DataRobot data science service. We partner with our customers to help you gain traction with AI projects, with a plan to help you develop your in-house AI maturity. Catalyst BI works with you to develop, identify and rank use cases for data science, then ensures your models are deployed and value is achieved, beginning with pilots, but then supporting you to scale across your organisation as value, trust and expertise is gained.

Use Case Workshop

Catalyst offers a ½ day on-site workshop with various stakeholders (this can be hosted virtually if this is a preferred option for your organisation). The workshop will cover AI and ML and how to identify potential use cases within your organisation, to help stakeholders understand machine learning, and how to use it strategically. Typically, 8-12 use cases are identified, and these are ranked by Value, feasibility, and status.

These values will be rough estimates, that need an AI Maturity check and Use Case POV to confirm values

Typically after the workshop, Catalyst recommends an AI Maturity Assessment to help identify the most effective route forward for the organisation.

AI Readiness Assessment

To achieve value through data science projects, organisations must achieve a state of readiness from both an organisational and technical perspective. It is considered best practice to lay strong foundations, ensure that your data is in a healthy state, address organisational and behavioural change, and adopt a culture that embraces the value of data.

The Catalyst AI Readiness Assessment provides an expansive and thorough assessment of the current use of predictive analytics. Our experienced consultants will review and document your current adoption and models, identify any issues/areas to address and provide clear prioritised recommendations on how these can be improved.

Our assessment can be broken down into 4 key areas:

1. Technology and Data
2. Strategy and People
3. Organisation and Products
4. Ethics and Governance

Each area is scored and scores are rolled up per area:

- None
- Exploring
- Experimenting
- Formalising
- Integrating

The outcome of the assessment is a comprehensive report which documents your current environment, our findings in each of the 4 reviewed areas and our prioritised recommendations.

Within each section of the report, we provide clear instructions on developing your readiness, extracting more value from your deployment, and ensuring the best user experience. Catalyst BI then help you to assemble a custom Go-Live Plan to enable you to meet your goals, and progress into our first projects, either with Catalyst AI Cloud (DataRobot Catalyst Managed Service), or the enterprise DataRobot Platform.

Fivetran

Catalyst has deep, hands-on experience implementing Fivetran as part of modern data platform transformations. Our expertise typically includes:

End-to-end Fivetran implementation

Catalyst configures connectors, optimises sync schedules, and ensures data lands in the warehouse in a clean, structured format.

Designing scalable data architectures

Our team are experienced in integrating Fivetran into broader data ecosystems - Snowflake, Databricks, Azure, dbt, and BI tools- ensuring the whole stack works seamlessly.

Performance and cost optimisation

Catalyst in partnership with Fivetran helps organisations:

- Reduce unnecessary sync volume
- Optimise connector usage
- Tune warehouse consumption
- Implement governance and monitoring

Migration from legacy ETL

Catalyst often replaces brittle, hand-coded pipelines with Fivetran's automated connectors, reducing engineering overhead and improving reliability. Fivetran provides the **automation, reliability, and simplicity** that modern data teams need. Catalyst brings the **expertise, architecture, and delivery capability** to make Fivetran work at scale inside real organisations.

Coalesce

Catalyst has practical experience implementing Coalesce as part of Snowflake-centric data transformations. Our expertise typically includes:

Designing transformation frameworks

Catalyst builds reusable Coalesce templates and modelling patterns tailored to each organisation's needs, ensuring:

- Consistency
- Maintainability
- Faster onboarding for new engineers

Accelerating Snowflake modernisation

Coalesce can be used for replacing:

- Legacy ETL tools
- Hand-written SQL pipelines
- Fragile transformation scripts

Resulting in cleaner, more governed data models.

InsightSoftware – VizLib & Fiplana

Catalyst is a certified VizLib partner and delivers the full library of VizLib extensions/Add-ins to the Qlik portfolio as follows:

- VizLib Library - Build better dashboards in Qlik
- VizLib Self Services - Empower users with next-level independent data exploration
- VizLib Finance - A single source of truth for modern financial reporting
- VizLib Gantt - Add more flow to your planning in Qlik
- VizLib Collaboration - Writeback directly in Qlik

Fiplana is a financial planning and analysis (FP&A) tool designed to integrate directly with Qlik, providing real-time budgeting, forecasting, and data analysis. It eliminates manual, spreadsheet-based processes by allowing users to create, edit, and write back data directly within Qlik dashboards. Key features include improved collaboration, automated reporting, and enhanced financial modelling.

Consultancy and Education Services

Having a Catalyst BI consultant involved early in your planning can ensure that you avoid costly re-work as your project progresses, ensuring you get the foundations right from day one.

Assessment & Strategy

BI Modernisation - Data and Analytics Strategy

Our BI Modernisation Consultancy services will help you to build a roadmap and strategy to fully leverage the capabilities of your data.

Empowering fluency in data: our education program is designed to fill a critical business skills gap by extending data literacy capabilities to a broad audience, resulting in a data literate workforce that is empowered and more effective.

Driving success in digital transformation: we support your digital transformation journey by providing deep expertise and best practices around data and analytics. We offer consulting solutions to ensure adoption of data literacy across the enterprise, supported by the right culture.

Enabling better data-informed decisions: through the power of Qlik Sense, you can easily visualize, explore, and collaborate around your data, in a smart intuitive fashion. With our associative engine and new augmented analytics capabilities, users of all skill levels are guided to ask the right questions and spot hidden insights.

The Modernisation Review provides an assessment of your existing Qlik applications including key functions, data sources, data modelling practices and usage. The review will look to identify any issues to ensure these are not carried forwards.

As part of the engagement Catalyst BI will advise on the best approach to migration, including strategies for efficient re-use of applications, data sources and QVDs across the Qlik platforms

The outcome of the review is a tailored set of recommendations to ensure an efficient and robust migration of QlikView applications and a roadmap for the rollout of self-service analytics in Qlik Sense.

Our Application Review is recommended for those who plan to migrate existing applications to Qlik Sense.

Qlik Health Check

Our Qlik Health Check provides an expansive and thorough assessment of your Qlik deployment. Our experienced consultants will review and document your current environment and processes, identify any issues or areas to address and provide clear prioritised recommendations on how to improve.

The Health Check covers twenty key areas which span five core themes, including everything from hardware and security to how successful your implementation has been at informing and engaging users. A simple 5-point scale is applied to each of the 25 points we assess.

- **Vision, Roadmap and Integration**

We review how well the Qlik deployment aligns to the organisation's overarching data and analytics strategy, as well as the level of innovation. We also look at how well the analytics platform is embedded into the wider business from both a technical and operational perspective, and how successful the internal communication is (branding and messaging) at engaging users.

- **Systems & Performance**

We review system performance including CPU, RAM, and disk and check for optimal configuration of the Qlik environment. We also look at your testing environment and approach, disaster recovery, and upgrade management.

- **Application Deployment**

We examine your deployment structure, how applications are organised and presented, scheduling and performance of data reloads, security and governance, testing and release. We also review licence allocation to ensure efficient usage and identify any shortcomings.

- **Application Development**

We review your development processes and assess applications against best practice development. We consider the data structure, application design, performance optimisation and documentation of deployed applications.

- **Adoption & Integration**

We evaluate additional factors which influence the overall performance of your Qlik deployment and user experience. These include data and analytics maturity, capability and skills, data literacy, and user engagement.

As part of the engagement Catalyst BI will also advise on the best use of QlikView and/or Qlik Sense, including strategies for efficient re-use of data sources and QVDs across the Qlik platforms, recommendations for an efficient and robust migration of any QlikView applications and a roadmap for the rollout of self-service analytics in Qlik Sense.

The outcome of the Qlik Health Check is a comprehensive report which documents your current environment, our findings in each of the 25 areas reviewed and recommendations. With each point we leave clear instructions on how to improve performance, extract more value from your deployment and ensure a good user experience.

Data Literacy Adoption Framework

Data literacy is the ability to read, work with, analyse and communicate with data.


In our information-driven world, data literacy is just as important as reading or writing. Every individual needs to know how to confidently work with data and analytics to generate actionable insights fast, regardless of your role, skill level, or the BI tools you use. Not everyone needs to be a data scientist... but everyone needs to be data literate. Improving data literacy hones your decision-making skills. You learn to ask the right questions of your data, interpret your findings, and take informed action.

What does a data literate workforce mean for organisations?

- **Get quicker insights** because you don't have to rely on just a few data scientists who can only fulfill so many requests and may not have a good understanding of the business questions you need them to answer. Everyone is empowered to get insights on their own.
- **Make better decisions** because your workforce has been trained to ask the right questions of their data, feels comfortable using data to support important decisions, and has the ability to question data with an intelligent level of data skepticism.
- **Stay competitive** because you can quickly make sense of your data in order to adapt to what's coming and stay ahead of your competitors.
- **Foster employee engagement** where your staff feels empowered in the decision-making process and benefits from a work environment that encourages collaboration.

Our consultants will work with you to develop and implement a six-step **Data Literacy Adoption Framework** adapted specifically to suit the needs of your organisation.

- **Driving success in digital transformation:** we will support your digital transformation journey by providing deep expertise and best practices around data and analytics. Our consulting services and adoption framework will help to ensure adoption of data literacy across the enterprise, supported by the right culture.
- **Empowering fluency in data:** our Data Literacy education program is designed to fill a critical business skills gap by extending data literacy capabilities to a broad audience, resulting in a data literate workforce that is empowered and more effective.
- **Enabling better data-informed decisions:** through the power of Qlik Sense, you can easily visualize, explore, and collaborate around your data, in a smart intuitive fashion. With the associative engine and new augmented analytics capabilities, users of all skill levels are guided to ask the right questions and spot hidden insights.



Training & Education

Our Training & Education Services provide learning opportunities and resources to maximise the success of your Qlik deployment. Whether it's improving your Data Literacy skills, learning how to get started with Qlik Sense, using R and predictive technologies or producing content rich reports with Qlik Nprinting we have a course for everyone.

All of our courses can be delivered in person or remotely with supporting course material classroom style or workshops based on your own data.

Data Literacy

Data literacy is the ability to read, work with, analyse and communicate with data.

In our information-driven world, data literacy is just as important as reading or writing. Every individual needs to know how to confidently work with data and analytics to generate actionable insights fast, regardless of your role, skill level, or the BI tools you use. Not everyone needs to be a data scientist... but everyone needs to be data literate. Improving data literacy hones your decision-making skills. You learn to ask the right questions of your data, interpret your findings, and take informed action.

Data Informed Decision Making

Learn about the importance of Data Literacy in the digital age. What it means to be Data Literate, why analytics are important and techniques to help you make data informed decisions.

Foundational Data and Analytics

Understand core principles of data, including data types, structures, dimensions and measures. Explore some key analytical concepts including distributions, correlation and causation, simple linear regression and more.

Qlik Sense

Our Qlik Sense courses are designed to enable your teams to unlock the power of Qlik. Throughout these courses we introduce best practice techniques and time saving methodologies developed over many years with the experience and lessons taken from hundreds of Qlik deployments.

Qlik Sense for Professional Users

Qlik Sense for Professional Users is a short Training Workshop designed to give end users the skills and knowledge needed to confidently work with Qlik Sense applications and take full advantage of the self-service analytics capabilities.

The workshop will provide an introduction to Qlik Sense, including how to navigate the platform and applications, finding and generating insights, and building Story Boards. In addition, users will learn how to develop a simple application from scratch using the Data Manager and Cognitive engine.

Qlik Sense Foundation

Qlik Sense Foundation is an introductory workshop designed to impart the skills and knowledge needed to quickly get up and running with Qlik Sense.

In this workshop developers will learn how to navigate in Qlik Sense, connecting to data sources using both the Data Manager and Script Editor, creating visualisations, master dimensions and measures and using the Insights Advisor.

Qlik Sense Advanced Application Development

Qlik Sense Advanced Application Development is a deep dive session which looks at data modelling concepts, important scripting techniques, application design principles and analytical approaches. The course introduces the Catalyst BI Methodology, a best practice approach to development with a defined set of standards and processes which when followed ensure high quality, sustainable and extensible applications. You will learn key principles that communicate information to your users in the most effective way and find out how to perform comparative analysis e.g. This Year versus Last Year or Department A versus Department B, without having to redesign your data structure.

Areas covered include

Data modelling and scripting techniques

Creating data connections

Simplifying development through structured scripting

Effective use of variables

Working with QVD files

Transforming and joining data

Managing dates

Script debugging

Design best practice

Advanced expressions including Set Analysis and advanced aggregation

Qlik Sense for QlikView Developers

Qlik Sense for QlikView Developers is an Introductory Training and Best Practice Workshop designed to give experienced QlikView developers the skills and knowledge needed to quickly adapt and get up and running with Qlik Sense.

In this workshop developers will learn how to navigate in Qlik Sense, connecting to data sources using both the Data Manager and Script Editor, creating visualisations, master dimensions and measures and using the Insights Advisor.

The workshop will provide an overview of the key differences between QlikView and Qlik Sense and focus on best practice for developing Qlik Sense applications that are designed to facilitate self-service analytics.

Administer and Maintain Qlik Sense

Administer and Maintain Qlik Sense is a technical course that gives you the information, tools, techniques, and skills needed to set up, configure and maintain the Qlik Sense Enterprise platform.

Learn how to manage users and licences, publish applications, create custom security rules using Attribute Based Access Control, manage Data Connections and monitor usage and performance.

Qlik NPrinting

Learn how to create Qlik NPrinting scheduled reports with rich content including Excel, Word, PowerPoint and HTML. This interactive workshop provides an overview of NPrinting and covers topics including how to connect to

Qlik applications, designing reports, creating and applying filters, creating and applying conditions, managing users, distributing reports to the newsstand and distributing reports via email.

Qlik GeoAnalytics

Qlik GeoAnalytics is an advanced suite of location intelligence tools. In this course you will learn about how Qlik GeoAnalytics differs from the basic Qlik Sense mapping objects. You will learn how to use map layers and drill down capabilities as well as exploring several GeoAnalytics operations, including drive times and routes.

Qlik APIs and Mashups

In the context of web development, a mashup is a web page or web application that uses content from more than one source to create a single new service displayed in a single graphical interface.

Qlik's Capability APIs enable you to easily and quickly integrate with your Qlik Sense objects to produce enriched results that were not necessarily the original reason for producing the source application in Qlik Sense. You can reuse Qlik Sense visualisations, including your custom extensions, and you can also make use of Qlik Sense data and calculations.

Our Qlik APIs and Mashups course will kick start your developers on their journey with Qlik Sense mashups. The course will introduce the Qlik Sense Dev Hub, work through how to create a simple mashup and hypercube as well as chart JavaScript libraries, object interaction and the Selections API, Mashup deployment and authentication. Whilst the Qlik Sense Dev Hub simplifies development delegates on this course will need to understand the rudimentary concepts of JavaScript, HTML (Hyper Text Mark-up Language), and CSS (Cascading Style Sheets).

Qlik App Automations

Eliminate repetitive work with a visual interface to build automations. No need to write code or start from scratch, use drag-and-drop to automate reports, data reloads, and process flows with ease. In this course Catalyst will provide an overview of use cases in Qlik App Automations including integration with Teams. Delegates will learn to create an App Automation, manage an App Automation, use task chaining and reloads, send email messages, automate administrative tasks and highlight best practice.

Adaptive Consultancy Services

Our Adaptive Consultancy Services offer a flexible approach to consultancy and development which enables us to support your Qlik projects in a way which is tailored to your organisation or specific project challenges.

We use our extensive experience to ask the right questions, in order to understand your organisation, your challenges and your requirements. We then work in partnership with you to design and develop robust and scalable solutions.

Whether it is in the design and implementation of a new Qlik deployment, development of new applications, or re-design and conversion of QlikView apps to Qlik Sense or our Adaptive Consultancy Services will accelerate your data and analytics strategy.

Our Consultancy Services

- Business Analysis
- Solution Architecture
- Data Modelling and Scripting
- Application Design and Development
- AI Model Development

When you work with Catalyst BI's expert team of Qlik consultants you can choose the delivery mode (Outsourced, Paired, Independent or any combination) that best suits your specific needs.



Consultancy and Development

Catalyst BI's expert team provide Consultancy and Development services to accelerate your Qlik projects. When you need to develop or migrate applications quickly or have limited capacity within your team our expert consultants are on hand to take the lead and deliver.

Paired Development

A truly collaborative approach, our Paired Development workshops allow your developers to work alongside expert Catalyst BI consultants to enable rapid knowledge transfer and adherence to Qlik Sense development best practices.

Paired Development workshops allow you to develop or migrate applications quickly whilst developing skills and capability.

Managed Application Support

With our Managed Application Support your team can be confident in developing applications, safe in the knowledge that Catalyst BI's expert team are on hand to provide help and support when needed.

Access experienced Qlik consultants via the Catalyst BI support desk for quick and efficient resolution to challenging problem areas. Hours are consumed in increments of 15 mins and can be used anytime within 12 months.

Infrastructure & Architecture

Catalyst BI Cloud Hosting

The Catalyst BI Qlik Cloud hosting provides a fully hosted, secured, and fully serviced Qlik Sense or QlikView Business Intelligence environment hosted within the Microsoft Azure Cloud. This can be a simple hosted environment, clustered or multi-product encompassing the entire Qlik product set including NPrinting, Alerting, Mobile, GeoAnalytics and Insight Bot.

Hosted within a Microsoft Azure environment, the service is a fully serviced, patch and upgraded environment for Qlik deployments. Due to this being a flexible platform the offerings that this fully managed service can bring are to prove a custom solution with the likes of VPN tunnels, port securing, jump boxes, network load balances and high availability provided over multiple datacentres.

- Compute and application hosting
- Container service
- Data warehousing
- Relational database
- Other database services
- Infrastructure and platform security
- Platform as a service (PaaS)
- Other storage services

Service features

- Data visualisation, analytics and data insights
- Collect data over multiple sources and locations
- Custom and feature-rich dashboards
- Mobile-enabled
- Analyse hundreds to billions of rows (Big data)
- Data Searching over tens of millions of data points
- Standard and customised training
- Build analytical applications
- Drag and Drop visualisation creation
- Automated data loading

Service benefits

- Rapid deployment of advanced analytics, deployed in days not weeks
- Fully hosted with no infrastructure or technical overhead
- Access anywhere data insights, available on any device
- Secure environment to safeguarded data with a proven track record
- High availability with 99.6% up-time guarantee
- Easy to use with little or no end user training
- Powerful administration and auditing built in

Architecture Design and Deployment

Catalyst BI will provide full design, deployment and if needed migration to Qlik Sense, we can provide to fully resilient high availability Qlik clusters to single node installations and as you expected all the relevant configuration and testing.

We approach this in the following manner:

- Design of the Qlik Sense architecture considering

- Sizing
- Architecture
- Future Proofing
- Integration with authentication needed including SSO e.g. Azure AD, Google, local AD with ADFS
- Testing including post-performance and future scalability testing

Sizing Design

Every deployment of Qlik Sense is different, and the following is an example for the expected deployment you're your use and environment. This requires detailed and careful discussions with our skilled architects skilled in infatuate and Qlik technologies

Qlik Cloud Analytics Migration Services

Moving to Qlik Cloud from QlikView or Qlik Sense opens up capabilities for AI, Automations and sharing data with unlimited users. Optimising your current on premise or Qlik Cloud User set-up to move to the Qlik Cloud Analytics capacity model is one of Catalyst's strengths, supported by the Qlik Cloud migration toolkit. We deliver the following services:

- Migration Planning Workshop
- Set up SaaS Tenant
- Installation of Data Gateway
- Conversion of Apps Workshop
- Train the Trainer
- Knowledge Transfer
- Qlik Cloud Training Courses
 - Qlik Cloud 101 (for end users)
 - Qlik Cloud Administration
 - Insight Advisor / Collaboration Platforms
 - Qlik Cloud Analytics Reporting
 - Qlik App Automations
 - Replicating Data to Cloud Storage
 - Qlik AutoML
 - Data Flows
- Technical Project Management
- Post-Implementation Review
- Managed Application Support (100hrs)
 - Consultancy and Development services to accelerate your analytics modernisation program.
 - Choose the delivery mode (Outsourced, Paired, Independent or any combination) that best suits your specific needs
- Outsourced App Migration (optional)
 - Analysis of on-premise apps
 - Project plan for app migration
 - Time estimate to migrate all or some of your apps

Qlik Sense Deployment & Configuration - Customer Owned

Where you would prefer to host your own Qlik environment, Catalyst BI will work directly with your own internal or external IT provider providing full design, deployment with full architecture design and sizing.

This can be either a Single or Multi node environments providing high reliance, testing and development deployments. Like the hosted approach we approach this in the following manner:

- Design of the Qlik Sense architecture considering
 - Sizing
 - Architecture
 - Future Proofing
- Integration with authentication
- Testing including post-performance and future scalability testing

Qlik Sense Multi-Cloud Setup and Configuration

Where multiple sites internal, cloud or externally hosted, we can also deploy and support clustered Qlik Sense deployments with a mix of cloud and internal deployments.

This service will take elements of the cloud and customer owned environments to provide a hybrid deployment providing the efficiency of an on-premises solution to access localised data and the flexibility and resilience of the cloud.

Qlik Governance & Security Definition

The Qlik Sense platform allows organisations to flexibly define access rights and permissions to control not just to what users see but what they can do. On the journey to enabling self-service analytics across an organisation, processes, governance structures and responsibilities need to be clearly defined.

We will work with you to collaboratively develop and document a governance framework for your Qlik self-service reporting and analytics solution. The Governance & 'Ways of Working Framework' will define the governance structure and the accompanying processes required to mitigate risks and ensure a well-controlled, robust platform for devolved development and administration.

The Framework will cover:

- Administration of licences, users, and access permissions
- Access to and administration of data connections
- Access to levels of development capabilities and functionality
- Devolved administration rights for application development and deployment
- Enhanced role-based data access and security (Section Access)




Figure 1: Example of Custom Access Roles & Permissions

Project Management

Our process-light Project Management services are designed to ensure your Qlik project is delivered smoothly on time and to budget.

You will be assigned a dedicated Project Manager who is experienced at not only managing project deliveries but will also have at least ten years of practical project delivery experience.

To be successful, Business Intelligence projects often require a blended project approach, incorporating resource from both business and technical functions. Bringing these sometimes disparate functions together to work in a cohesive manner and deliver challenging and time critical projects can be difficult. Our technical Project Managers have the experience and knowledge to successfully bring together different stakeholders, avoid common pitfalls and keep the project on track, applying lessons learnt from hundreds of previous projects.

We will work closely with you to set clearly defined objectives for each project and ensure consistent progress with regular communication throughout to keep the project on track. We will make sure roles and responsibilities are clearly defined and any risks and issues are closely monitored and proactively managed.

Our Project Management service provides:

- Dedicated, experienced Project Manager with a minimum of 10 years practical project delivery experience who can act as a conduit between the different business functions
- Clearly defined Project Initiation Document (PID) detailing the background, objectives and expected outcomes of the project, roles and responsibilities, and risks
- Project plan and resourcing schedule
- Regular project update meetings
- Regular project highlight report to monitor progress, risks and issues

Support

The Catalyst BI team is on hand to support you at every stage. Our accredited expert team of Qlik/Snowflake/DataRobot/Talend consultants will help get up and running quickly, and with our **Education** and **Adaptive Consultancy** services will ensure you maximise performance and realise value.

Our approach at Catalyst BI is one of collaboration and partnership. From the very first engagement we will work closely with you to understand your goals, challenges, organisation and teams. We work hard to develop long term relationships and aim to be a strategic part of your business so we can help you get to the next level of your data and analytics strategy.

We provide continued support and maintain open lines of communication at all times should you need it. You can contact the Catalyst BI Support Centre by emailing support@catalyst-it.co.uk.

1st line Product Support

As a Catalyst BI Qlik customer you will automatically receive 1st line support for your products. We will work with you to identify issues and escalate where needed to support.

Adaptive Consultancy - Managed Application Support

With our Managed Application Support service your team can be confident managing the Qlik platform and developing applications, safe in the knowledge that Catalyst BI's expert team are on hand to provide help and support when needed.

Whether you have a question on how to load data, want to configure new access rules on the server, upgrade your software, need a helping hand with some complex analysis or just want us to take on some developments, our Managed Support service has you covered.

You can access our experienced Qlik consultants via the Catalyst BI Support Centre for a quick and efficient resolution to challenging problem areas when needed. Hours are consumed in increments of 15 mins, deducted from your 'bank' of hours, and can be used anytime within 12 months. The minimum contract period is 40 hours.

As a Managed Application Support customer, you will also get regular update calls with your Project Manager so that you understand the status of your implementation and projects, and we can help organise your requests. Your Project Manager will help drive and deliver business change, ensure that you are always addressing your business issues on a priority basis and keep you abreast of important product updates.

Why Choose Our Services

Catalyst BI is Qlik's largest Elite partner working in the public sector and we work closely with their technical, sales and account management teams.

We provide organisations with the opportunity to access an extremely powerful and flexible solution that is easy to customise and will help drive your business operations.

Our culture is all about helping our clients to be more successful and to partner with our clients over the long term. We also have a long and proven track record in delivering public sector projects on time and within budget.

How to Buy Our Services

If you are interested in buying our services, please contact Catalyst BI Limited to discuss your requirements in detail.

Contact Details

Catalyst BI Limited

Wellington House,

Aviator Court,

Clifton Moor,

York

YO30 4UZ

0330 380 1609

bi@catalyst-it.co.uk