Azure Synapse Analytics vs. Other Leading Solutions

Table of Contents

COMPARISON GUIDE:
Azure Synapse Analytics vs.Other Leading Solutions
Make data-driven decisions more effectively and efficiently

01 Introduction to Modern Analytics Solutions.....4
02 What to Consider When Choosing an Analytics Solution.....5
03 Overview of Azure Synapse Analytics.....6
04 Comparison GuideAzure vs. Other Leading Solutions.....7
05 Why Companies Make the Move.....8
06 Preparing to Move to Synapse?.....9

Chapter 1:
Introduction to Modern Analytics Solutions

Analytics solutions are now foundational tools across industries, empowering organizations to transform enterprise data into meaningful, actionable insights. But while teams have been onboarding these solutions at a fast clip, not everyone is tapping into the technology's full potential.

Data Driven Organizations Do Better

organizations that intensively use customer analytics are:

23x more likely to acquire new customers

6x more likely to retain their customers

19x more profitable

87.8% of executives report increases in corporate data and analytics investments.
Source: Forbes.

In fact, most aren't. Fewer than 20% of companies have achieved advanced analytics at scale.1 Those who aren't effectively leveraging their analytics solutions to steer decision-making are missing out on opportunities to minimize costs, project future revenue and resource needs and reduce complexity, to name a few.

There's a lot to gain from putting data at the center of everything you do, but finding the right path to get there isn't always obvious. The challenge begins with picking the right analytics solution—one that aligns with your organization's needs and is both time- and cost- efficient to deploy.

Azure Synapse Analytics (Synapse) is an analytics service that brings together SQL technologies used in enterprise data warehousing, Apache Spark technologies for big data and Azure Data Explorer for log and time series analytics.

Ideal for mid-and large-sized enterprise organizations, as it accelerates time to insight across data warehouses and big data systems, Synapse provides business and technical benefits for both high-level decision-makers as well as data engineers and data scientists.
In this guide, we'll look at Azure Synapse Analytics next to several of the leading modern analytics solutions—Snowflake, BigQuery and Redshift—to help you determine the best fit for your environment.

Chapter 2: What to Consider When Choosing an Analytics Solution

When you start comparing analytics solutions, they may not appear overwhelmingly different. After all, every solution is more or less focused on helping you complete the same tasks. This can make it even harder to isolate the best solution for your needs.

As you shop around for an analytics solution, keep these things in mind:

Business objectives:
Start by naming your organization's overarching goals, and then break these down into measurable analytics tasks. The right solution should not only meet your analytics requirements but also offer data visualization and reporting tools so you can effectively deliver insights to both technical and non-technical users.

Use cases:
Understand how specifically will your organization leverage and analyze data. This should be the main driver in your decision-making. For example, organizations aiming to analyze user behavior may need immediate insights, requiring a solution that can deliver real-time analytics. But if large-scale analyses aren't time-sensitive, a batch processing-oriented solution may be a better, more cost-efficient choice.

Data integrations:
The best analytics solution for your organization is one that can easily integrate with your current infrastructure. It's also important to consider how a solution can connect with third-party data sources and applications, both those you use today and those you may onboard in the future.
Security: Data privacy and security are paramount to today's business success. Choosing a solution that adheres to the most stringent industry and regulatory standards with advanced data encryption (both at rest and in transit), managed access controls and automated threat detection are essential for compliance and peace of mind.

Flexibility:
Analytics platforms should grow as your organization does, both in terms of data volume and complexity. Importantly, a flexible analytics solution should keep pace with not only your organization's growth but the industry at large with the ability to adapt to future business and technology innovations.

Chapter 3:
Overview of Azure Synapse Analytics

Synapse is an enterprise data analytics platform that accelerates time to insight across data warehouses and big data systems. It's the prime analytics tool in Microsoft, where it lives among the 14 other Azure analytics products, including Data Factory and Data Lake Analytics.

More specifically, Synapse is a powerful data analytics service that combines data integration, enterprise warehousing and big data analytics to power data transformation, BI and machine learning processes. Synapse provides data engineers with authoring experiences that support data analysis and collaboration. It gives data scientists a platform to run AzureML jobs through their Spark notebooks and pipelines.

Synapse brings together the best of SQL technologies, Apache Spark technologies and Azure Data Explorer, as well as Power BI and Azure Machine Learning.

Key Functionalities

Azure Synapse contains multiple functionalities, including pipelines, data lakes and Apache Spark.

Pipelines
ETL (extract, transform, load) or ELT (extract, load, transform) pipelines in Synapse give data engineers a low/no-code visual environment for managing data. These pipelines ingest and clean log data and can then either kick off a mapping data flow to analyze that data or start Spark notebook jobs to analyze and transform it.

Azure Data Lake
Azure Data Lake — a service of built on Azure data Storage — stores ingested data. This data can then be processed and used as a basis for a variety of analytic needs. Azure Data Lake also provides core data consistency across applications, powering big data analytics, machine learning, predictive analytics and other forms of intelligent action.

Apache Spark compute capabilities
Apache Spark (Spark) is a parallel processing framework that supports in-memory processing.Synapse makes it easy to create and configure a serverless Spark pool in Azure, where you can run Spark jobs to manipulate and analyze data. For example, you can take data from four or five sources and then use Spark pools to write delta tables. These delta tables are a part of the medallion architecture, wherein the gold layer houses the aggregated data. You can then use an on-demand SQL pool to read the data. Note that a Data Lakehouse comprises delta-tables on Data Lake with the medallion architecture. This way, a Data Lakehouse is the modern Data Warehouse.

Synapse Dedicated SQL Pools
Dedicated SQL Pools, which are for traditional Data Warehousing use cases, allow organizations to choose a level of SQL compute and storage. While they are more expensive than On-Demand SQL pools, Dedicated SQL Pools can, unlike Azure SQL, be turned off when not in use to save on costs.

How Azure Synapse Analytics Differs from Other Analytics Solutions

It's part of the Microsoft ecosystem.

Synapse is native to Microsoft, meaning that teams already working in Microsoft Azure can quickly and easily connect dozens of tools and solutions for a comprehensive — and efficient — data insights environment.

It has serverless capabilities.

Synapse offers both dedicated SQL pools and on-demand, serverless SQL pools.
A serverless SQL pool is a distributed data processing system that looks and acts like an SQL server — but it's not. It's built for large-scale data and computational functions, enabling you to run queries in seconds to minutes.

With Serverless SQL pools, you can still use SQL Server Management Studio to view your published tables. These published tables can then be consumed by Power BI. Ultimately, you can work as you would in a traditional data warehouse, but you don't have to set up and maintain infrastructure or clusters — it's all handled for you by the Synapse service. Now, if teams make a mistake, they can easily roll back delta tables in serverless SQL pools. Synapse is also a cost-saver. The meter in Synapse only logs time on an on-demand table when you're actually reading or writing — when you're finished, it shuts off, which means you're not paying for a data warehouse 24/7.

Additionally, your data pipelines are low-code or no-code, which requires very little compute resources, saving you additional budget.
This is in stark contrast to the dedicated SQL pools of many other analytics solutions, where you're responsible for provisioning your own server. They're essentially combinations of memory and compute — and they're expensive because you either have to let them run non-stop or manually shut them down during quiet periods.

Supporting Tech Teams While Improving Your Bottom Line

Azure Synapse's flexibility and stability help enterprises achieve their strategic business objectives faster, while bolstering their security posture, enhancing service availability and supporting effortless scale.

With Synapse, organizations can generate critical insights about the health of their business and the state of the market with data from virtually any warehouse, data lake, operational database or big data analytics system in their environment. This wide array of data stores — all easily accessible via pre-built connectors — empowers teams to enrich existing apps and operations, promoting greater efficiency and profitability.

Plus, because it's all engineered to suit the changing demands of modern businesses, Synapse helps teams safeguard their data from unauthorized access and use, creating a downtime-resistant environment, while providing them with a unified user experience to increase organization-wide efficiency and lead the way for limitless scale.

Data engineers will be able to ramp up quickly in Synapse because it uses languages they're already familiar with — Python, Scala, R and T-SQL. In other words, if they're transitioning to Azure from a different analytics tool, they won't be stunted during migration by any complicated configuration challenges, since the familiar language makes onboarding intuitive and straightforward.

Meanwhile, data scientists will welcome Synapse's various machine learning capabilities, like Azure Machine Learning, which they can use to enhance intelligent apps without any data movement.ategic Differentiator:

Chapter 4: Comparison Guide- Azure vs. Other Leading Solutions

Synapse delivers powerful insights at scale with advanced security features, cost optimization capabilities and a unified experience. How does it compare to competing analytics solutions?

Let's start by looking at the competitors to Synapse — namely, Snowflake, Redshift, and BigQuery.

Perhaps the most well-known analytics solution, Snowflake is a traditional data warehouse. It delivers the Data Cloud, a global network where organizations can mobilize data, unite their siloed data, discover and share governed data and execute diverse workloads. It offers features like real-time data collaboration, in-database machine learning and BI Tool integration.

Redshift, meanwhile, is the data warehouse native to Amazon Web Services (AWS). It's a fast, fully-managed, AI-powered, petabyte-scale data warehouse in the cloud with massively parallel processing (MPP) architecture. This MPP architecture enables organizations to quickly run complex queries operating on large amounts of data.

Finally, BigQuery is the data warehouse in Google. It's a fully-managed, serverless, cost-effective enterprise cloud data warehouse with built-in machine learning that works across clouds. Interestingly, BigQuery features Duet AI, an AI collaborator directly integrated into the warehouse for ML/AI insight at scale.

Performance

Azure Synapse combines SQL data warehousing with big data analytics services, making it a limitless analytics service that can support organizations who need an integrated analytics solution. Its advanced analytics capabilities can efficiently handle both large volumes of data and complex data structures from any data warehouses and big data analytics systems in your environment, empowering your organization to optimize the performance of all queries to rapidly gain insights and, ultimately, reduce overall project development time.

Synapse vs. Redshift
Redshift is the ideal analytics solution for AWS users. It's simple to set up, maintain and integrate with other AWS services, enabling organizations to move, transform and load data quickly and reliably. This integration, however, is necessary since Redshift doesn't have any advanced analytics capabilities built in. For AWS users who want to stay in the AWS universe, Redshift is a simple, natural solution, but it could pose unnecessary clumsiness for those working in other environments.

Ease of Use

Synapse is a PaaS (Platform-as-a-Service) and, as a part of the broader Microsoft ecosystem, is feature-dense yet very user-friendly.

Particularly, since Synapse lives in Azure, you don't have to worry about managing any additional third-party services. This not only accelerates time to start but makes computing easier than it is when navigating two different systems.

Synapse vs. Snowflake
Snowflake, on the other hand, is a SaaS (Software-as-a-Service) platform. This means it must run on top of Azure, Google Cloud, or AWS.

Because it's SaaS, Snowflake is one of the most intuitive and easy-to-use analytics solutions. It was also designed to offer users considerable control and flexibility, (e.g., supporting discretionary access control and integrating with many different data lineage tools). However, it lacks Synapse's native integration with the rest of the Azure universe. It's thus an extra step for teams to connect with the Microsoft ecosystem and take advantage of its advanced analytics features.

Security

Synapse is equipped with advanced, enterprise-grade security and privacy features, including:

- Column-level encryption
- Column- and row-level security
- Dynamic data masking
- Automated threat detection
- Always-on encryption
- Managed access control

Together, these features cover the five layers of security: data protection, access control, authentication, network security, and threat protection.

Synapse vs. BigQuery
Like Synapse, BigQuery enables administrators to manage users' roles and permissions, an important step in managing privacy and safeguarding data. But while both analytics solutions use encryption, only BigQuery turns this feature on by default; Synapse doesn't. BigQuery also allows you to configure a network security perimeter with Google Cloud Platform's Virtual Private Cloud (VPC) Service Controls, though Microsoft offers similar functionalities with its virtual networks. happen easily from the cloud-based control pane, forgoing the need for companies to deploy, patch, or manage an on-premise solution.

A downside of Commvault Cloud is a lack of feature parity with competitive solutions — and even the previous on-premises Commvault — that often matters a great deal to enterprise organizations. It also doesn't have all the porting functionality that the previous Commvault enterprise-level solution had.

Integration with Third-Party Databases and Tools

Stacked against the competition, Synapse really shines with its wide breadth of integrations.

For one, its native integration with Microsoft gives you easy access to the brand's larger ecosystem. But Synapse also extends well beyond the Azure universe, empowering teams to perform even more efficient analytics and reporting with 95+ pre-built connectors for popular databases and tools, including Salesforce, PayPal and Oracle.

Synapse vs. Redshift
In comparison, Redshift is notably lacking in this area.

While the AWS analytics solution does offer direct integrations with several AWS services (e.g., Amazon S3, Amazon DynamoDB, and Amazon EMR), it doesn't give you any pre-built connectors for other third-party databases or tools.

Cost

Across the board, analytics solutions are notorious for their opaque billing. But Synapse is starting to make things a bit easier. Its serverless SQL pool means you're not running a data warehouse 24/7, so you're likely to see a smaller bill at the end of the month. Plus, if you're already working in Microsoft, there's even more room for cost savings.

To add some clarity to opaque billing, Azure also offers a pricing calculator to help organizations run general cost estimates before launching a project.

Synapse vs. Snowflake
While Synapse offers serverless SQL pools, dedicated SQL pools, and Spark pools, Snowflake uses virtual warehouses, (i.e., it functions like a traditional data warehouse). Snowflake mitigates 24/7 charging with its auto-scaling and auto-suspend feature, which allows you to isolate individual workloads during busy or idle periods. It's worth noting that Synapse's serverless SQL pools and Spark pools have automatic scaling by default, but its dedicated SQL pools require manual adjustments.

Chapter 5: Use Cases for Azure Synapse Analytics-Why Companies Make the Move

Compared to competing analytics solutions with bigger name recognition and a wider talent pool, why would companies make the move to Synapse? Let's look at some of the major reasons our customers have made the choice to switch:

They're on a Data Modernization Journey

Many businesses make the move to Azure Synapse Analytics because they want to upgrade an older technology.

For example, one customer had old servers — IBM mainframe DB2 for iSeries — and wanted to move to a PaaS. Since they already used Microsoft, they went with Azure to tap into the solution's native integrations.

Their data modernization journey began with Azure SQL Managed Instance, where they migrated their data using Data Factory. This move allowed them to get rid of their old, expensive machine and decommission the data center so they no longer had to pay for it. In addition to eliminating their technical debt and modernizing their infrastructure, they added Azure Machine Learning and PowerBI. This gave them valuable, timely insights that enabled them to uncover opportunities to build and expand their business.

They're Dealing with Siloed Data

Thirty-nine percent of organizations that identify as being highly data-driven have 50+ data silos. Data silos stifle your organization's ability to turn data into actionable insights — but Synapse makes it possible to pull that disparate data together and turn it into business value.

In fact, using Synapse to break down silos is our number one customer use case. You can connect different databases from different sources, feed it all into Synapse, and then layer machine learning, Power BI or other integrations on top.

They Need to Process Analytics Data Efficiently

Many organizations initially attempt to accomplish both the processing and analysis of their data through Power BI. However, they quickly encounter inefficiencies, particularly with Power BI's computational capabilities. Power BI struggles with heavy compute tasks, leading

to prolonged wait times for data processing, especially on standard business laptops. Additionally, the complexity of Power BI's primary programming languages can be a hurdle for developers, who find it significantly easier to work with Spark.

Azure Synapse Analytics integrates seamlessly with Power BI, effectively separating computing tasks from data visualization and significantly reducing the time to obtain actionable insights. This distinction is crucial for businesses experiencing slow Power BI report generation due to simultaneous data computations and loading. By shifting compute-intensive tasks to Synapse, organizations can access and act upon critical data more swiftly, enhancing their decision-making processes.

They Want Everything in One Service

Choosing Azure Synapse Analytics is an intuitive decision for those already embedded within the Microsoft ecosystem due to its seamless integration and comprehensive data analytics capabilities. However, its appeal extends far beyond those who are exclusively committed to Microsoft products. Even for users operating in diverse technological environments, Synapse offers a compelling suite of features that can significantly enhance data analytics practices.

Synapse consolidates critical data analytics functions into a single, unified platform, including:

Pipelines for Data Ingestion:
Synapse facilitates the seamless integration of data from Microsoft and non-Microsoft sources, ensuring that data is readily available for analysis.

Serverless SQL Pools:
These allow for scalable, pay-as-you-go data query execution, enabling users to manage and analyze large datasets without the need for upfront provisioning.

Spark Compute Capabilities:
Synapse incorporates Apache Spark, offering powerful data processing and analytics capabilities that can handle complex computations and large volumes of data.

For organizations with limited operational resources, Synapse simplifies the management of data analytics infrastructure, reduces the need for specialized training across multiple tools and streamlines the entire data analytics workflow. It's important to note that any projects you do in Synapse today can easily be migrated in the future to Microsoft Fabric, the recently released all-in-one data and analytics solution.

Chapter 6: Preparing to Move to Synapse? Get the Support You Need

Almost any infrastructure is well positioned to adopt Synapse—even organizations at lower maturity stages. Plus, it's a low-pressure move. Depending on how you architect it, you don't have to worry about being locked into the solution forever.

The best route to Synapse is the one you take with a partner.

As a Microsoft Solutions Partner, US Signal® IT Solutions can help you easily transition to Azure Synapse Analytics and fully benefit from the vast capabilities of the Microsoft ecosystem.

Contact our team to talk about making the move.