


# Microsoft Analytics Platform System



## Enterprise-ready Big Data

- Integrated relational and non-relational data for a turnkey Big Data analytics appliance
  - Security, manageability, and high availability for Hadoop
  - Seamless querying across relational and non-relational data using simple T-SQL
- Insight into virtually all data types with familiar tools through native business intelligence

## Next-generation performance at scale

- Linear scale-out architecture of up to 6 petabytes of data
- Up to 100 times the performance and 15 times more data compression than traditional data warehouses with In-Memory Columnstore
- Removal of traditional data warehouse bottlenecks with the MPP capabilities of SQL Server PDW
- Scalable user accessibility via running of mixed workloads at high concurrency

## Engineered for optimal value

- Lowest price per terabyte for a data warehouse appliance in the industry
- Choice of hardware through HP, Dell, and Quanta

For decades, the data warehouse has been at the center of the enterprise's decision support infrastructure, acting as the system of record for data analysis. Now the traditional data warehouse has reached a critical point, requiring major business-driven changes to the systems in place today. Key contributing factors include:

- **Data growth:** Databases designed with traditional symmetric multiprocessing (SMP) architecture cannot scale to keep up with the amount of data that is expected to grow tenfold over the next five years without major investments in hardware, tuning, support, and maintenance.
- **Non-relational data:** Organizations are using Apache Hadoop to store existing data and process new data types from sources like blogs, sensors, social media, and devices. This data can end up isolated from users because it is not integrated with data in the traditional data warehouse.
- **End-user expectations:** End users need results in near real time, and they expect their internal systems to match the speed of an Internet search engine.

The modern data warehouse needs to enable users to collect and analyze virtually all data, regardless of its size or type. It also needs to deliver performance, scale, and user accessibility to keep up with enterprise demand in this world of Big Data.

## Microsoft Analytics Platform System

Microsoft is ready to help organizations transition to a modern data warehouse with the Microsoft Analytics Platform System (APS), a no-compromise solution that unifies non-relational data from HDInsight, Microsoft's 100-percent distribution of Hadoop based on the Hortonworks Data Platform, with relational data from Microsoft SQL Server Parallel Data Warehouse (PDW), a massively parallel processing (MPP) relational data

warehouse, into a single integrated appliance. APS provides tier-one performance, low TCO, and accessibility to all users through some of the most widely used business intelligence tools in the industry such as Microsoft Excel.

## Solution advantages

Through seamless integration of HDInsight and PDW, APS offers the following benefits:

**Enterprise-ready Big Data:** HDInsight gives organizations the most enterprise-ready Hadoop distribution on the market with end-user authentication through Active Directory, management and monitoring through Microsoft System Center, high availability with Windows Server Failover Clustering, and performance tuning for APS hardware. PolyBase, the query tool available only in APS, enables users to easily query PDW and HDInsight data using T-SQL, without investing in Hadoop-based skills or training. With PolyBase, organizations can take advantage of flexible hybrid Hadoop solutions and query across Hortonworks, Cloudera, and even into the cloud with Microsoft Azure HDInsight.

**Next-generation performance at scale:** With true scale-out architecture from HDInsight and PDW, APS is one of the most powerful distributed computing and scale-out systems on the market. As an organization's need for compute or storage grows, they can simply add more resources to scale up to 6 petabytes. Additionally, they can eliminate common data warehouse bottlenecks with the MPP capabilities of PDW. The In-Memory Columnstore feature of PDW changes how data is stored and can help organizations gain up to 15 times more compression than is typically possible and up to 100 times the performance in I/O and query processing over traditional rowstore processing.

**Engineered for optimal value:** Microsoft streamlines the hardware footprint and optimizes value for the modern data warehouse investment. Windows Server 2012 reduces cost through Storage Spaces by allowing SAN levels of performance with commodity storage drives. In-Memory Columnstore helps organizations reduce storage usage by up to 70 percent. As a result of these innovations, PDW offers the lowest price per terabyte in the industry for a relational data warehouse appliance. By integrating Hadoop into the same rack as the relational data warehouse, organizations can save on consulting/configuration costs for Hadoop as well as


hardware, energy, and general data center costs with an integrated appliance.

## Customer success story

The Royal Bank of Scotland—the leading UK provider of corporate banking services—needed a powerful analytics platform to improve performance and customer services. To handle multiple terabytes of data and an unprecedented level of query complexity more efficiently, the bank implemented with Analytics Platform System with SQL Server PDW. As a result, it gained near-real-time insight into customers' business needs as well as emerging economic trends, cut a typical four-hour query to less than 15 seconds, and simplified deployment.

## Choice of hardware vendors

APS offers the flexibility to choose hardware from Dell, HP, or Quanta—each built for optimal performance.



## Conclusion

APS is a no-compromise modern data warehouse solution that seamlessly combines a best-in-class relational database management system, in-memory technologies, Hadoop and cloud integration in a turnkey package built for Big Data analytics.

## Call to action

For more information on APS, go to <https://www.microsoft.com/sqlserver/en/us/solutions-technologies/data-warehousing/pdw.aspx>, or contact your Microsoft representative.



The PolyBase technology and HDInsight are powered by Java.



Java is a registered trademark of Oracle and/or its affiliates.