

IBM Integrated Analytics System

Do data science faster

Highlights

- A high-performance, easy-to-deploy, cloud-ready and purpose-built appliance
 - Advanced analytics model creation with the included IBM® Watson Studio
 - Higher data load speeds & timelier query responses for analytics workloads
 - Federation and advanced querying with the IBM Common SQL Engine
-

Better explore, store and manage your data

In the last few years we have seen a rapid evolution of data. The need to embrace the growing volume, velocity and variety of data from new technologies such as Artificial Intelligence (AI) and Internet of Things (IoT) has been accelerated.

The ability to explore, store, and manage your data and therefore drive new levels of analytics and decision-making can make the difference between being an industry leader and being left behind by the competition. The solution you choose must be able to:

- Harness exponential data growth as well as semi-structured and unstructured data
- Aggregate disparate data across your organization, whether on-premises or in the cloud
- Support the analytics needs of your data scientists, line of business owners and developers
- Minimize difficulties in developing and deploying even the most advanced analytics workloads
- Provide the flexibility and elasticity of a cloud option but be housed in your data center for optimal security and compliance

Benefit from Netezza technology

The IBM Integrated Analytics System meets these needs and includes embedded IBM Netezza Analytics technology with multiple algorithms, including linear regression, decision tree clustering, k-means clustering and Esri-compatible geospatial extensions. The system is designed to work with business analytics and visualization tools, including IBM Cognos, SAP BusinessObjects, Kognitio, Microsoft Excel, QlikView, SAS, Microsoft SQL Server Reporting Services (SSRS) and Tableau. The system also handles model-building and scoring tools such as IBM SPSS, Fuzzy Logix, open source R and SAS.

The IBM Integrated Analytics System drives the insights needed to increase your competitiveness by matching accelerated development and deployment times for your data scientists with a high-performance, optimized and cloud-ready data platform.

As a unified data science solution, the built-in IBM Watson Studio can be used by your data scientists to connect with your organization's data in place. This connection helps data scientists develop machine learning analytics that benefit from a performance-optimized common SQL engine with embedded Apache Spark processing.

From the start, the IBM Integrated Analytics System requires little or no tuning and maintenance to deploy and manage even demanding workloads that require high performance and petabyte-level scalability. The IBM Integrated Analytics System enables machine learning with the Apache Spark processing engine embedded on the system for higher performance analytics. At the same time, this feature can help reduce the complexity of moving analytics and data to separate environments. A common SQL engine shared across the IBM hybrid data management offering family lets you work with your existing on-premises and cloud applications. This flexibility allows you to pick the right environment for the right tasks.

The integrated architecture combines software enhancements, such as asymmetric massively parallel processing (AMPP), with IBM Power technology and flash memory storage hardware. The IBM Integrated Analytics System handles traditional data warehouse workloads and operational mixed workloads. These workloads often require processing queries against large data volumes, quick point queries on small data sets and multiple concurrent operational accesses. As a result, the IBM Integrated Analytics System supports a wide variety of analytics use cases across broad data types and locations on a single solution. This flexibility provides your data scientists with almost endless possibilities.



The IBM Integrated Analytics System

Embedded analytics and machine learning with Spark

The IBM Integrated Analytics System helps simplify data scientists' efforts to train and evaluate predictive models with embedded Apache Spark processing. This feature helps eliminate the need for time-consuming movement and transformation of data to other systems. Once the models are developed using the tools of the data scientists' choice, the testing, deployment and training can be done where the data resides. With each node containing its own Spark executor process, latency is minimized, which helps speed data access and calculations compared to a stand-alone Spark cluster. In those cases where data scientists need to take the workloads off the system, industry-standard tools and the common SQL engine provide the option to seamlessly move models to a Spark cluster.

In addition to streamlining processes, this ability can also provide advanced performance and flexibility for analytics, including machine learning capabilities. Your data scientists can immediately connect to data in the system and begin building models with the five authorized user licenses included with the IBM Watson Studio.

This interactive, collaborative, cloud-based environment allows data scientists to use multiple tools to activate their insights. Data scientists also have the option of using Python, R or Scala using Jupyter Notebook with a Jupyter Notebook container included on the system. Jupyter can be used to execute interactive code with one-click deployment that transforms the code into a compiled and deployed Spark application.

In addition to prebuilt functions for data mining, prediction, transformations, statistics, geospatial data and data preparation, the Spark capability embedded in the IBM Integrated Analytics System supports open source R and other programming languages like Python, Java, C, C++ and Lua.

Integration and performance

IBM helps simplify the deployment and management of the analytics system using a design based on more than 20 years of experience with thousands of clients across multiple industries and regions. The software and hardware arrive at your data center configured to work together as a single performance-optimized solution. Within hours, you can load data without creating database indexes or struggling to tune and retune the data warehouse once it's operational.

Clients using the IBM Integrated Analytics System and the included IBM Db2 Warehouse technology should immediately recognize the common SQL engine used across the entire IBM hybrid data management solution portfolio. The IBM Db2 Warehouse is designed for data warehouse and analytics workloads. The common SQL engine uses dynamic in-memory columnar technologies for multi-workloads based on IBM Db2 and IBM BLU Acceleration technology. BLU Acceleration massively parallel processing (MPP) architecture is designed for rapid and deep analysis of data that can scale into the petabytes. With query response times up to 100 times faster than earlier systems,¹ BLU Acceleration columnar tables can coexist with traditional row tables in the same schema, storage and memory so you can query both row and BLU Acceleration columnar tables at the same time. Adding BLU acceleration technology to traditional in-memory capabilities can accelerate performance even when data sets exceed the size of the memory. The dynamic in-memory columnar technologies of BLU Acceleration with data skipping offer an efficient method to scan and find relevant data even when the data is compressed.

The IBM Integrated Analytics System leverages IBM Power Systems and IBM FlashSystem technology to improve reliability and performance at the hardware level. Today's IBM Power architecture enables denser systems that can achieve similar performance with less nodes than previous offerings. As the default storage for the system, IBM FlashSystem offers ultra-low latency and high near-in-memory I/O speeds with outstanding reliability.

While the analytics applications run at peak performance, the IBM Integrated Analytics System also brings new levels of reliability to help you meet or exceed your service level agreements. Power Systems and IBM FlashSystem storage is rated with increased uptime thanks to the fault tolerant design that helps eliminate a single point of failure. According to a 2017 Information Technology Intelligence Consulting (ITIC) survey, IBM Power Systems has the least amount of unplanned downtime — with 2.5 minutes per server per year — of any mainstream Linux server platform.²

Redundancies are built into components throughout the system, helping ensure continued operation in case of a hardware failure. The system also includes additional built-in, high-availability features to provide automated failovers for performance continuity. Monitoring and management for all components — hardware and software — is provided by a built-in console powered by IBM Data Server Manager that's used across the Db2 family.

A hybrid approach to the cloud and your data

When it comes to your data, a one-size-fits-all approach rarely works. The IBM Integrated Analytics System is built on the common SQL engine, a set of shared components and capabilities across the IBM hybrid data management offering family that helps deliver seamless interoperability throughout your infrastructure.

For example, a data warehouse that your team has been using might need to be moved to the cloud to meet seasonal capacity demands. Migrating this workload to IBM Db2 Warehouse on Cloud can be done seamlessly with tools like IBM Bluemix Lift. The common SQL engine helps ensure no application rewrites are required on your part.

The common SQL engine provides a view of your data regardless of where it physically sits or whether it's unstructured or semi-structured data. The system's built-in data virtualization service in the common SQL engine helps unify data access across the logical data warehouse allowing you to federate across Db2, Hadoop and even third-party data sources.

Replication, scalability and expansion options

IBM Data Replication for Db2 Continuous Availability, a new, optional service for IBM Integrated Analytics System customers is also available. It supports highly available Db2 Warehouse environments by synchronizing data over both row and columnar organized tables and schemas, whether on the same platform, across the data center, or around the world. This software replication offering also supports active and stand-by replicas for workload balancing, shifting workloads during planned outages while also dramatically reducing the time to recovery for unplanned outages. This offering is pre-integrated into IAS and lets you get started quickly with a 90-Day 'Try it Now' license.

Both the software and hardware architecture have been designed to grow and scale as you bring more workloads to support your business onto the system. Compute and storage capacity can be expanded independently, providing almost cloud-like levels of flexibility and elasticity. Hardware expansion is non-disruptive to your business and can be done in place on the system.

The IBM Integrated Analytics System also supports multi-temperature tiered storage to help ensure the highest levels of performance, even with large volumes of data. The system manages the most recently used and active hot data directly on the system storage nodes, while older, less active cooling data resides on more cost-efficient, high-density IBM Storwize storage devices.

Use cases

The common SQL engine used in the IBM Integrated Analytics System lets you match the right workload with the right deployment platform, while helping ensure that data is accessible regardless of type, location or size. The following are a few use cases to inspire you in getting started:

- Make operational excellence the new normal by creating a logical data warehouse using Hadoop, data marts or other associated deployments that all interact and offer a unified view of your data whether they sit on premises or in the cloud.
- Create personalized customer experiences in real time by using your internal data with analytics that requires high performance and simplified scalability.
- Improve time to market for new product innovations by using embedded machine learning to accelerate performance on complex analytics and deliver insight to your users.
- Deliver the requested workloads to your business users more quickly with a system that loads data faster and requires little or no tuning, nor extensive levels of configuration.
- Expand the workload options available to business users with the addition of operational mixed workloads.
- Help ensure operational compliance with more accurate views of your data.
- Offer new levels of insights to grow and expand your business by building analytics across different data types and data sets using data virtualization to unify data access across the logical data warehouse and other technologies.

Specifications

The IBM Integrated Analytics System integrates and optimizes all compute, storage and networking resources with analytics and data warehouse software. It's available in rack configurations as shown in the table.

Specifications

IBM Integrated Analytics System model number	M4002-003	M4002-006	M4002-010	M4002-020	M4004-040
Active nodes	3	5	7	13 + 1 Standby	25 + 3 Standby
CPU cores	72	120	168	312	600
Memory (TB)	1.5	2.5	3.5	6.5	12.5
Approximate uncompressed user data (TB) ¹	27	54	81	162	324
Maximum power per rack (Watts)	4,300	7,000	9,600	19,200	38,400
Cooling (BTU/hour)	14,700	24,000	33,000	66,000	132,000
Rack weight (Kg)	530	620	710	1,420	2,840
Height (cm) ²	202				
Depth (cm) ²	109.2				
Width (cm) ²	64.8				
Power cords	For North America and Japan—200 - 240 V, Single Phase, 30 A For EMEA—220 - 240 V, single phase, 32 A				
Drops/rack	4				
Safety	UL/CSA 60950-1 and, also, have EN 60950-1 and IEC 60950-1				
Emissions	FCC part 15, ICES-003, AUS/NZ CISPR 22, VCCI, EN55022/55032 Class A; European immunity: EN 55024, EN 61000-3-2 and EN 61000-3-3				

1. Assume up to 4x compression to calculate user data based on approximate uncompressed user data. For example a full rack user data capacity would be 4 x 81 TB resulting in 324 TB
2. Dimensions are given per rack

Why IBM?

IBM provides one of the most comprehensive portfolios of hybrid data management solutions to help you maximize the value of your information assets and discover new insights to make better and faster decisions, and optimize business outcomes.

For more information

To learn more about IBM Integrated Analytics System, contact your IBM representative or IBM Business Partner, or visit the following website: ibm.biz/BdjfvP.

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing

© Copyright IBM Corporation 2019.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:
IBM Netezza® , IBM Cognos® , IBM SPSS® , IBM Power® , IBM Db2® Warehouse , IBM Db2 and IBM BLU Acceleration® , IBM Power Systems™ , IBM FlashSystem® , IBM Bluemix® Lift , IBM Storwize®



Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.