

BlueData EPIC® Software Platform

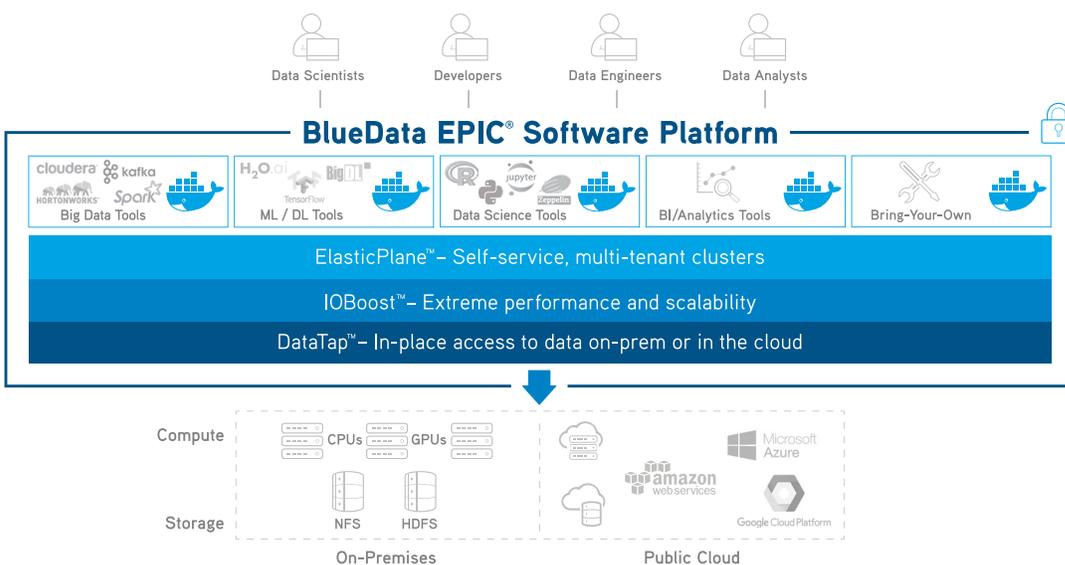
The BlueData EPIC® platform makes it easier, faster, and more cost-effective to deploy large-scale AI and Big Data environments—including TensorFlow, Spark, Kafka, Hadoop, and more—whether on-premises, in multiple public clouds, or in a hybrid model.

Run AI and Big Data Workloads on Containers

With BlueData EPIC (Elastic Private Instant Clusters) software, you can create distributed environments for machine learning, data science, and analytics in minutes rather than months. You can offer a self-service experience with the data and tools that your data science teams need, while providing enterprise-grade security and reducing costs.

The BlueData EPIC platform leverages Docker container technology and patented innovations to deliver faster time-to-value for your digital transformation initiative:

- **ElasticPlane™** enables users to spin up instant clusters on-demand in a secure, multi-tenant environment.
- **IOBoost™** ensures performance on par with bare-metal, with the agility and simplicity of containers.
- **DataTap™** decouples compute and storage resources, eliminating time-consuming data movement.



PRODUCT HIGHLIGHTS

- ▶ **Increase business agility** by empowering data scientists and analysts to create AI and Big Data clusters in minutes, with just a few mouse clicks.
- ▶ **Simplify deployments** with a turnkey container-based solution—including the flexibility to support on-premises, multi-cloud, and hybrid deployments.
- ▶ **Deliver faster time-to-insights** with pre-integrated Docker images for common machine learning, deep learning, BI / ETL / analytics, and data science tools.
- ▶ **Minimize the need to move data** by decoupling compute and storage—with the ability to tap into both on-premises and cloud-based storage.
- ▶ **Maintain security and control** in a multi-tenant environment, integrated with your enterprise's security model (e.g. LDAP, Active Directory, Kerberos).
- ▶ **Achieve cost savings** of 70% or more by improving utilization, controlling usage, eliminating cluster sprawl, and minimizing data duplication.

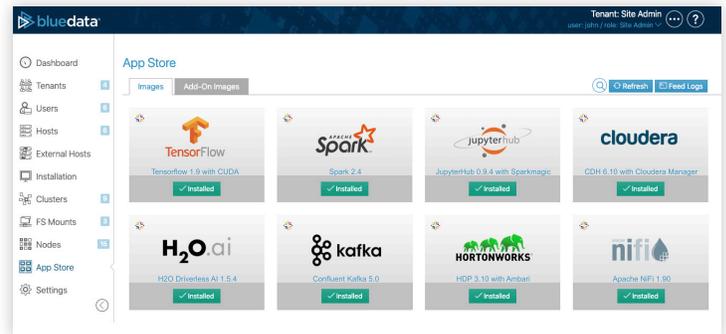
Accelerate AI and Big Data Deployments

Today, many AI and Big Data initiatives meet the needs of multiple user groups (or tenants) with multiple, dedicated physical servers on-premises. This approach can be complex, expensive, and time-consuming. For example, it can take several weeks to provision a new Hadoop cluster or distributed TensorFlow environment.

In some cases, these organizations turn to public cloud services to jumpstart their deployments. But deploying a public cloud environment for AI or Big Data workloads still requires the provisioning and configuration of multiple components, so a dedicated DevOps team is often necessary. And most public cloud services don't meet the needs of different use cases for multiple teams and tenants.

The BlueData EPIC software platform is designed to simplify and accelerate multi-tenant AI and Big Data deployments, whether on-premises and/or in multiple clouds. Enterprise deployments that may have taken months can now be completed within a couple days. Data scientists and analysts can create their own environments on-demand within minutes—while accessing data in on-premises or cloud-based storage.

BlueData EPIC installs as a container-based software layer between the underlying infrastructure (either on-premises and/or public cloud, CPU and/or GPU) and the applications (for machine learning, analytics, and other data-driven workloads). BlueData EPIC delivers faster time-to-value, lower costs, and ultimate flexibility for AI and Big Data deployments—for multiple user groups with different use cases and tools—without the need for DevOps skills or a specialized infrastructure team.



Common AI and Big Data tools are available via one-click deployment.

KEY FEATURES

Self-Service Environments

Users can get up and running quickly with BlueData EPIC's ElasticPlane functionality. New containerized environments are provisioned on-demand with just a few mouse clicks—whether they're transient for development and testing, or long-running for a production workload. Data scientists and analysts can now quickly respond to dynamic business requirements for a variety of use cases ranging from deep learning with TensorFlow to analytical SQL on Hadoop.

Secure Multi-Tenancy

Different project teams, groups, or departments across the enterprise can share the same infrastructure—and access the same data sources—for their AI and Big Data analytics workloads. BlueData provides multi-tenancy and data isolation to ensure logical separation between each project, group, or department within the organization. The platform integrates with enterprise security and authentication mechanisms such as LDAP, Active Directory and Kerberos.

Containers and Bare-Metal Performance

BlueData has developed patented I/O optimization innovations to deliver the agility benefits of Docker containers for AI and Big Data, while ensuring performance comparable to that of bare-metal servers. EPIC's IOBoost functionality provides application-aware caching and elastic resource management that adapts dynamically to changing application requirements, delivering the best possible performance.

Flexibility for Tools of Choice

The BlueData EPIC platform is pre-integrated with many of the most common AI and Big Data tools. You can easily spin up Cloudera or Hortonworks clusters; create standalone Spark and Kafka environments; deploy machine learning with TensorFlow or H2O; and install your choice of data science tools like Jupyter or Zeppelin notebooks. You can even add your own preferred tools and quickly upgrade to new versions—providing the ultimate in agility and configurability.

Data Access from Any Storage

With EPIC's DataTap capability, you can access data from any shared storage system (including HDFS as well as NFS) or cloud storage (e.g. Amazon S3). That means you don't need to make multiple copies of data or move data, before running your analysis. Sensitive data can stay in your secure storage system with enterprise-grade data governance, without the cost and risks of creating and maintaining multiple copies or moving large-scale data.

Compute and Storage Separation

BlueData disconnects analytical processing from data storage, giving you the ability to independently scale compute and storage on an as-needed basis. With BlueData EPIC, you also have unparalleled flexibility to mix and match infrastructure. You can run the BlueData platform on-premises, on multiple public clouds (Amazon Web Services, Google Cloud Platform, or Microsoft Azure), or in a hybrid model. This enables more effective utilization of resources and lower operating costs.