MESOSPHERE

SOLUTION BRIEF

Mesosphere Data and Analytics Service Automation

Overview

Organizations today have access to massive amounts of data, but these useful, but unstructured bits of information arrive in a variety of formats, making it impossible to make sense of. Competitive advantage is found by quickly accessing and monetizing huge pools of information to improve customer experiences and ensure rapid adaptation to evolving market demands. In many industries gaining even a few milliseconds of decision-making power can mean the difference between winning and losing.

The need for fast access to useful, accurate data is forcing organizations to search for ways to more effectively make data-driven decisions. As a result, many businesses are finding tremendous promise in a new world of data and analytics technologies that are able to deliver the quality experiences that customers are looking for today.

Cloud-native data services, such as Kafka, Cassandra, and Spark are routinely used to access, process, and scale high volumes of data, however, they all require significant effort and skill to deploy and operate. Making matters worse, organizations are then faced with the decision to deploy on a particular cloud service or use on-premises resources. These challenges and endless debates prevent speed and agility—the very things your data and analytics projects are supposed to deliver. To solve these problems, businesses are looking for easier ways to automate these processes to move faster, while enabling greater flexibility.

Key Benefits

Simplified experience with push-button deployment and management of data and analytics services

Flexible deployment options to empower deployment in any public or private cloud infrastructure

Enterprise-grade supportability for the most common prevailing open source data services

With the data-driven architecture enabled by Mesosphere DC/OS, data queries will go from minutes to milliseconds. This opens up a whole new world of possibilities when it comes to patient and provider services."

—Shailan Lala, Director,
Cloud Engineering at athenahealth

Mesosphere Data and Analytics Service Automation:

- Automates the full lifecycle of data and analytics technologies--including installation, scaling, zero-downtime upgrades, and automatic recovery from failure in a consistent way across any private or public cloud, and edge infrastructures
- Provides full visibility, uniform control, enterprise grade security, and extremely low operational overhead
- Accelerates deployment and customization for your data project with a JumpStart Accelerator engagement

Mesosphere Data and Analytics Service Automation

Mesosphere Data and Analytics Service Automation is a standardized and modular approach to delivering simplified installation, scalability, upgrades, and application resilience for a wide variety of open source data analytics and machine learning technologies. Mesosphere automates the Day 2 operations of data technologies such as Kafka, Cassandra, and Elastic, and analytics engines such as Spark, Hadoop, and Tensorflow—providing a consistent experience across any private or public cloud, and edge infrastructures.

With our wide-ranging support of cloud-native data services, powered by Mesosphere DC/OS, Mesosphere Data and Analytics Service Automation ensures that developers and data scientists can leverage the very best open-source technologies while maintaining enterprise grade operability for their initiatives.

Mesosphere tests the latest open-source packages to validate interoperability, so your development and data science teams can focus on critical capabilities such as building secure and insightful applications rather than infrastructure.

Line-of-Business Owners

In the chase to capture new markets and provide increased value to their customers, line-of-business teams are rapidly building applications which require a wide range of advanced data analytics services. These services are complex and involve a variety of distributed technologies for data transport, persistence, and analytics—often requiring portability across infrastructures that public cloud providers alone can't offer.

Mesosphere DC/OS provides advanced Service Automation enabling organizations to adopt new and emerging data services across any provider whether on-premise, public cloud, or edge infrastructure. Leveraging Mesosphere's automation accelerates the pace of application innovation, enabling the business to dynamically prototype, build, and implement new services as rapidly as demands evolve. Most critically, Mesosphere allows line-of-business owners to apply common standards which do not compromise rapid innovation or enterprise requirements for security and compliance.

Data Architects

Mesosphere DC/OS helps data architects rapidly adopt the prevailing open source technologies they require, without the need to employ deep domain operational expertise to deploy and manage these new technologies. As a result, application design teams can concentrate on building scalable data pipeline in support of business objectives, instead of focusing on complicated, but rudimentary tasks that these services require to function.

Building data pipelines with predictable costs as you scale, independent of API call volume, is crucial to your competitive advantage. Using Mesosphere DC/OS, data architects are able to deploy across any infrastructure, while leveraging unique data density capabilities to gain greater scale across a smaller footprint of physical or virtual resources.

Data Engineers and Scientists

Data engineers are often challenged with model building, training, prototyping, and eventually iterating in production new environments for developers and testers. Because all of the technologies data engineers use require separate installations, scaling, upgrading, and monitoring, subject matter experts are required to build, maintain, and automate each platform in a repeatable way.

Mesosphere DC/OS provides built-in Service Automation to make the deployment of these technologies easy, allowing engineers and scientists to focus on application construction, testing, and deployment instead of basic operational needs.

Accelerate Machine Learning with Dramatically Lower Costs

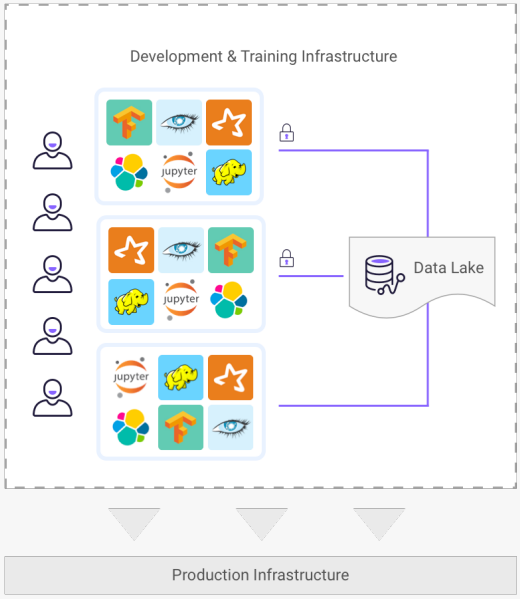

Time to Production: Weeks/Months

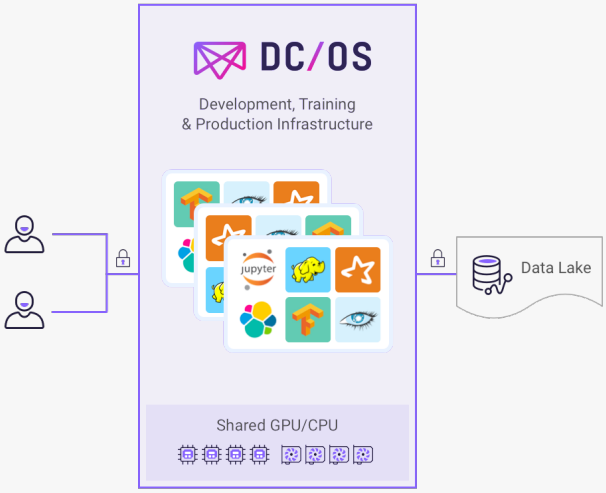

Time to Production: Minutes/Hours

Features and Benefits

Automated Operations	Automation platform with built-in lifecycle management and operational expertise
High Density Services	Increased utilization by deploying multiple data and analytics services on the same infrastructure without requiring virtualization
Performance and Health Monitoring	Responsive metrics that integrate with your existing enterprise or cloud SIEM and monitoring systems, provide insight into performance and data health
Centralized Control Plane	Reduced operational overhead and FTE requirements by centrally managing all data and analytics services deployments
Self-healing	Self-healing ensures data and analytic services health meets your requirements in the event of a mal-function or provider outage self-healing
Automated Failover	Multi-master design for control plane components and selected workloads
Supported Data Frameworks	High-level SLAs, emergency patching, unlimited incidents
Built-in Security and Encryption	Increased security with automatic configuration of TLS encryption for communication between services
Role Based Access Control	Ensure separation of concerns with built-in Service Accounts that isolate data and analytics services, administration, and integration with your enterprise or cloud services user directory
Advanced Networking	Safeguard data and analytics service network isolation and control through policy driven software-defined networking. Integrate with your enterprise or cloud services network load balancer for traffic management
Templates for Repeatability	Guarantee consistency across deployments and save time by creating a configuration template based upon your architecture, policies, and infrastructure choices
Supported Frameworks	Enterprise grade support and services for prevailing versions of data and analytics services running on Mesosphere DC/OS

About Mesosphere

Mesosphere is leading enterprise transformation toward distributed computing and hybrid cloud. We combine the rich capability you get from public cloud providers with the freedom and control of choosing your own infrastructure. Mesosphere is the premier platform for building, deploying, and elastically scaling modern applications and big data services.

mesosphere.com