

BIG DATA ON AMAZON WEB SERVICES

Offices in San Francisco, Los Angeles, San Diego, Seattle and Portland, United States | www.dualspark.com



We bring modern services, automation and cutting edge platforms to our clients. DualSpark specializes in making a move to AWS successful at any scale.

FINDING BIG DATA ENGINEERS with hands-on AWS experience is a challenge. DualSpark brings you the expertise to set up your Big Data ecosystem to analyze data at massive scale. We collaborate with those who know your internal processes best, fine-tuning controls to allow your system to run as close to operational efficiency as possible. We also identify AWS services to support delivery of actionable insight from your business drivers. Delivering Big Data and HPC solutions for your business.

AWS BIG DATA OPTIMIZATION

- » Resize dynamically and efficiently
- » Streamline analysis for insights
- » Build for durability & elasticity
- » Identify ideal usage patterns



amazon
web services

Partner
Network

ADVANCED CONSULTING PARTNER

SOLUTIONS FOR THE ENTIRE BIG DATA LIFECYCLE ON AWS
Establishing the right environment to extract actionable intelligence and analyze dynamic data at massive scale through proper architecture, engineering, and automation.

Whilst AWS offers all you need to handle Big Data workloads - from real-time streaming analytics that create customized content for millions of users, to analyzing global sales for multiple products - selecting the optimal usage pattern to balance performance, cost and automation is a common difficulty. DualSpark has the expertise you need to optimize the designs for performance, availability and cost.

As an Advanced AWS Consulting Partner, we understand the AWS ecosystem of analytical solutions that can be leveraged to address the volume, velocity and variety of data for multiple use cases. We work closely with your team to extract your requirements to support data collection, streaming, storing, analysis and archiving so you can:

Deliver customized content to millions of users

Create personalized recommendations for medicine based upon patient genetic information

Analyze billions of financial transactions

Process operational logs to automate your deployment

Our DualSpark engineers work with your team to examine your specific requirements, identify ideal usage patterns, and validate cost, performance and scalability to help drive strategic decisions, in real time.

We share our expertise with the primary AWS Big Data and High Performance Computing options, including Redshift, Kinesis, Elastic MapReduce (EMR), and DynamoDB. Where custom applications or compliance require a self-managed environment, we can help you design and deploy a customer platform on AWS EC2.

Let us talk about your unique challenges.

Contact us at (800) 501-3275 or info@dualspark.com to learn more about our work and how we can help.