

Grid Computing on AWS for Financial Services

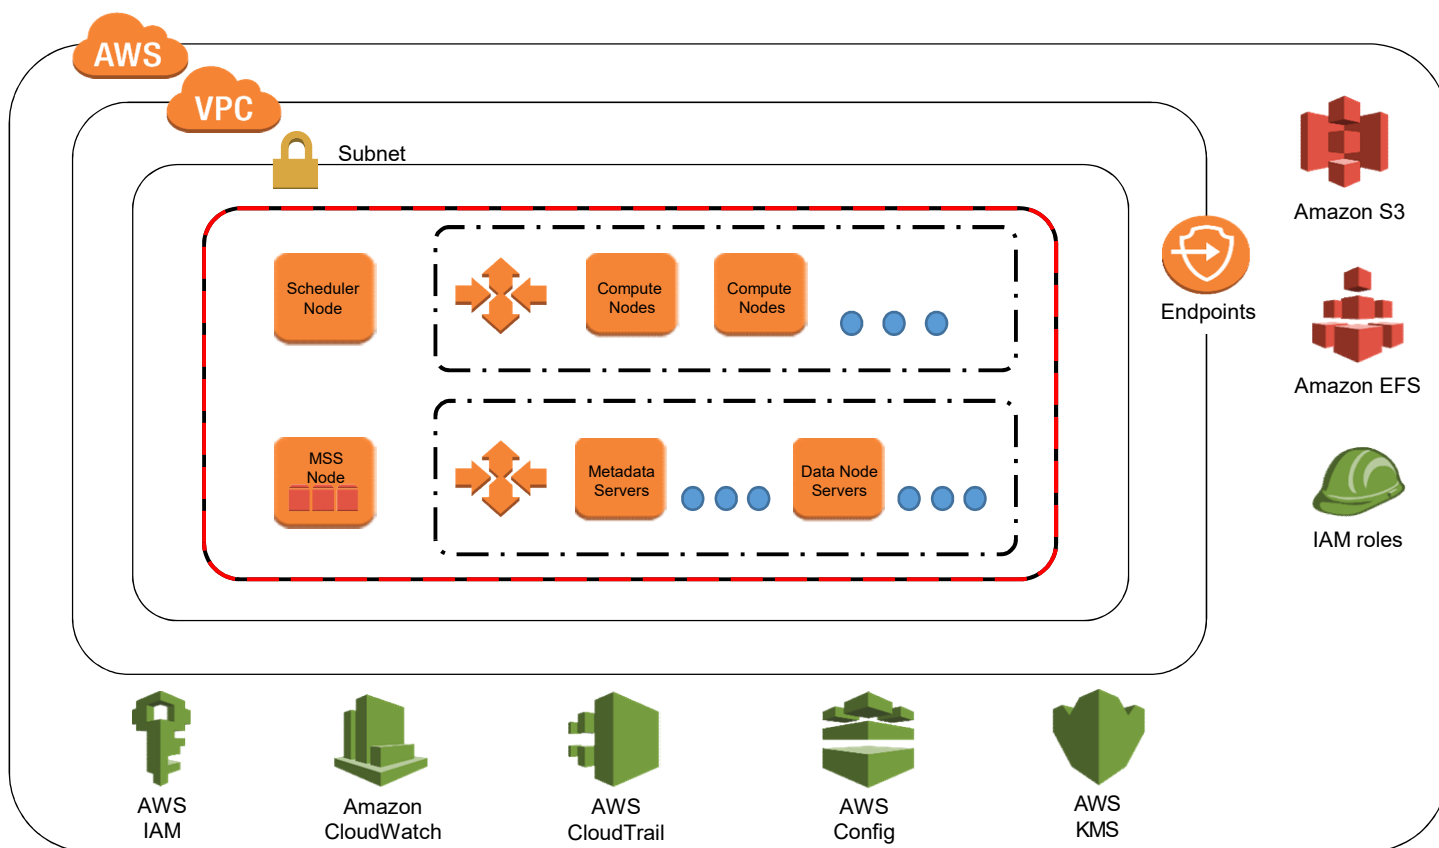


Financial Services

Build better grid computing capabilities with auto-scaling, unlimited storage, different instance types, enhanced security, and expanded big data capabilities.

Faced with the limitations of on-premise compute and storage capabilities, Financial Services institutions of all sizes—from fintech startups to the global investment banks—are working with AWS to extend their on-premise grids to the cloud or build cloud-native grids. **The following is a reference architecture of a grid computing solution that many of our Financial Services customers are deploying in their respective environments.**

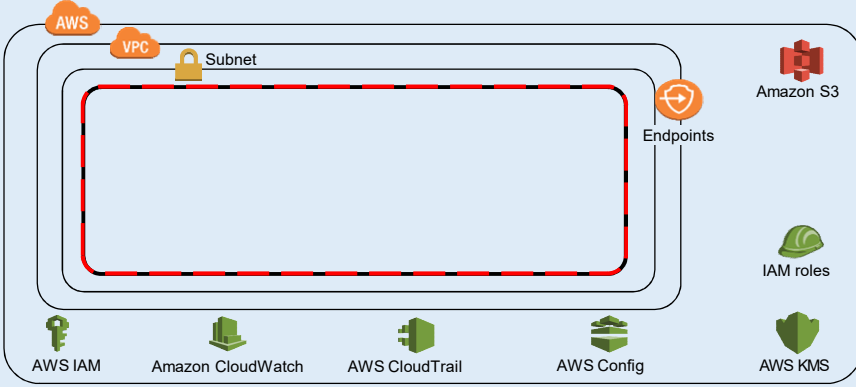

The AWS Cloud enables on-demand capacity so users can shut down their workloads when jobs are completed. Customers can also take advantage of advanced features—such as transient clusters or auto-scaling clusters that enable environments that can expand and contract depending on your job load.



- 1 Launch a virtual private cloud (VPC) with a subnet large enough to support your largest cluster.
- 2 Set up foundational services for storage, monitoring, and security, including: Amazon S3, AWS IAM and IAM roles, Amazon CloudWatch, AWS CloudTrail, AWS Config, and AWS KMS.
- 3 Set the storage layer by provisioning EFS and initiating and configuring servers to provide high-performance parallel file systems, such as IBM GPFS or Intel Lustre.
- 4 Set up the scheduling layer by initiating and configuring an EC2 instance with the appropriate scheduling software.
- 5 Set up an auto-scaling group to automate the creation and setup of EC2 instances for compute, based on the jobs pending in the queue.

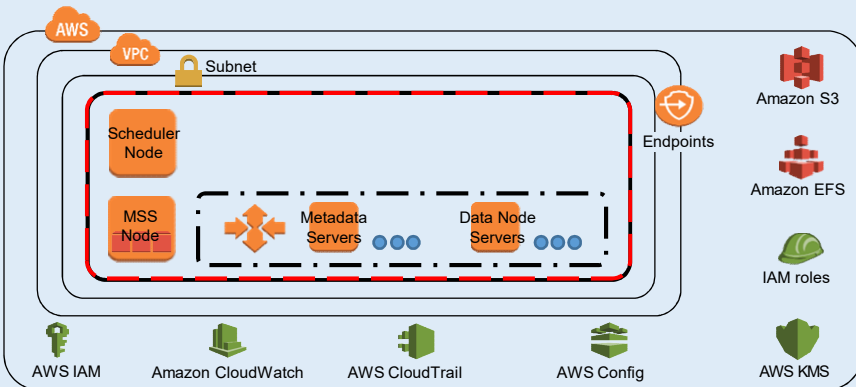
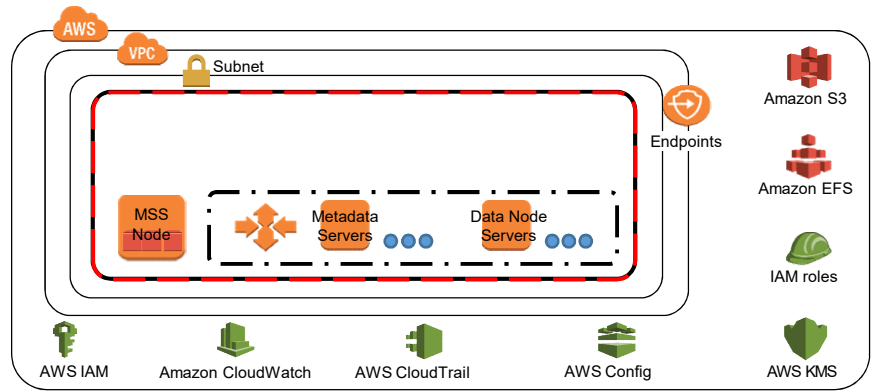
For more information on grid computing and other ways AWS can help your organization, visit us at: <http://aws.amazon.com/financial-services>

1 Launch a virtual private cloud (VPC) with a subnet large enough to support your largest cluster.



2 Set up foundational services for storage, monitoring, and security, including: Amazon S3, AWS IAM and IAM roles, Amazon CloudWatch, AWS CloudTrail, AWS Config, and AWS KMS.

3 Set the storage layer by provisioning EFS and initiating and configuring servers to provide high-performance parallel file systems, such as IBM GPFS or Intel Lustre.



4 Set up the scheduling layer by initiating and configuring an EC2 instance with the appropriate scheduling software.

5 Set up an auto-scaling group to automate the creation and setup of EC2 instances for compute, based on the jobs pending in the queue.

