



Enabling the Creation of Dynamic Globus Endpoints on AWS via CloudyCluster

Brandon Posey
Omnibond



Outline

- Motivation
- AWS Introduction
- CloudyCluster Overview
- Globus Dynamic Endpoint Generation
- Demo
- Conclusion
- Questions



Motivation

- Many researchers do not have adequate access to HPC resources exactly when they need it
- Steep learning curve associated with Cloud Providers researchers do not have time to learn how to use these services
- Data transfer to Cloud Provider Resources can be complex for researchers with little system administration experience



Amazon Web Services (AWS)

Overview

- Amazon Web Services (AWS) are a collection of services that allow users to allocate resources within the Amazon cloud
- Provide a number of different computing services such as database systems, virtual servers, identity and access management, data storage, and many others
- Pay-as-you-go services so that the user only pays for what they use
- Most services are charged on an hourly basis



CloudyCluster Overview

- Simple Web Based UI for creating dynamic fully operational HPC Clusters on demand within AWS without in depth knowledge of AWS
- Medium size clusters can be created within 20 minutes
- Standard suite of HPC Software pre-configured and installed
- Automated Globus Endpoint Creation for easy data transfer to and from the Cluster
- High availability OrangeFS Parallel Filesystem and/or a Scratch Filesystem preconfigured and automatically mounted

New Cluster Setup

Provide the following information and then select a cluster creation method.

Select the cluster region

us-west-2 *

Select the instance ssh key to use for instances in this cluster

ccDemo *

Select the Cluster availability zone

us-west-2a *

Cluster Information

Select your preferred CloudyCluster version.

Current Version *

Enter a Name for the Cluster.

globusWorldCluster *

Next you will create three sets of instances

- Utility (Scheduler, Access Instances, etc.)
- Working Storage (Scratch)
- Compute Groups

Select a method to create a cluster:

- Quick Start Clusters >
- Advanced Configuration >

Quick Start

Select one of the following cluster configurations.

- 1 Job-Scaling Compute Group (ccq and Torque Scheduler)
4 OrangeFS File System Instances, 1 Standby
1 Login Instance, 1 Scheduler, and 1 NAT Instance
- 4 Dedicated Compute Instances in 1 Compute Group
Amazon Elastic File System - EFS - (configured below)
1 Login Instance, 1 Scheduler, and 1 NAT Instance
- 4 Dedicated Compute Instances in 1 Compute Group
4 OrangeFS File System Instances, 1 Standby
1 Login Instance, 1 Scheduler, and 1 NAT Instance
- 8 Dedicated Compute Instances in 1 Compute Group
4 OrangeFS File System Instances, 1 Standby
1 Login Instance, 1 Scheduler, and 1 NAT Instance

Amazon Elastic File System:

Amazon EFS provides a low-latency, shared access fully-managed NFS file system.

Enable EFS

EFS Name:

Shared Software Installation Mount:

Shared Software Mount is a place to install software that can then be mounted on a cluster for



Dynamic Globus Endpoint Creation

- UI driven Dynamic Globus Endpoint Creation
- All the user has to do is enter their Globus credentials in the CloudyCluster UI and activate the endpoint in Globus
- Endpoints are created within 3-5 minutes
- Allows transfer of data directly to the Cluster's shared filesystem(s)
- Allows researchers with little Globus administration experience to utilize all of the Globus features on their own Cluster



Dynamic Globus Endpoint Creation

- Globus Endpoints are created on a per Cluster basis
- Extra Globus transfer nodes can be added for faster parallel transfers
- All AWS Security Group and firewall entries are created dynamically
- Utilizes OAuth for Globus Endpoint Activation and all authentication is done locally on the Cluster

Demo

The screenshot displays the CloudyCluster management console. The interface is organized into a sidebar on the left and a main content area on the right. The sidebar contains navigation options: OVERVIEW (selected), CLUSTER ACCESS, COLLABORATIONS, SETTINGS, and ADMINISTRATION. The main content area shows the 'Overview of globusWorld' page, which is divided into several sections:

- Utility - Utility:** Contains four components: Control (IP: 52.37.97.209, status: running), myScheduler (IP: 10.0.2.6, status: running), myLogin (IP: 52.38.67.16, status: running), and EFS (IP: 10.0.2.242, status: connected).
- Compute - cg339:** Contains four HPC instances (IPs: 10.0.4.102, 10.0.4.104, 10.0.4.105, 10.0.4.103, all status: running).
- Working - oranges:** Contains four FS instances (IPs: 10.0.3.6, 10.0.3.166, 10.0.3.205, 10.0.3.15, all status: running).
- Network - VPC Info:** Contains four network components: NAT (IP: 10.0.1.236, status: running) and three Subnets (IPs: 10.0.4.0/24, 10.0.1.0/24, 10.0.2.0/24, all status: connected).

Each section includes a 'More' button for additional details. The interface also features a top navigation bar with 'HELP' and 'LOGOUT' links, and a search bar in the main content area.



Conclusion

- Allow researchers access to on demand HPC clusters exactly when they need it
- Clusters and Globus Endpoints are available in minutes
- Minimal AWS knowledge required for creating HPC Clusters within AWS
- Dynamic Globus Endpoint Generation allows for easy data transfer to and from the Cluster
- Researchers can now focus more on their research instead of how to obtain the HPC environment that they need



Questions?

Thank you!

Brandon Posey

brandon@omnibond.com