HiveIO

White Paper

HiveIO Surge and Disaster Event plan

to AWS EC2

HiveIO Mission Statement

We transform the private datacenter into a simple, cost effective experience that makes it easy to do business today.

Our goal is to build cloud computing that transforms complex technology and layers into secure abstracted operations, ubiquitous client access and easy-to-implement application experience that can be deployed and supported by a single team of generalists.

Table of Contents

Abstract

The purpose of this paper is intended to detail how Amazon Workspace Services EC2 can be used as a surge and failover event architectural tool for HiveIO Customers. Use cases such as non-persist guest workloads (VDI/ VSI) and persist guest workloads (VDI/ VSI) can be either moved on demand or pre-staged for sub minute transitions. HiveIO leaves this at the discretion of the customer.

All steps provided have been tested and verified as of Nov 1, 2018. HiveIO will update this document should there be any changes every six months.

Watch Step by Step video guide - https://vimeo.com/302134944

Introduction

HiveIO provides superior on-premise cloud computing options including all features necessary to host general purpose workloads. All aspects of orchestration, inventory management, KPI trending, Cluster Resource Management ** CRS and on-demand workload management are included. HiveIO offers site-to-site and off premise cloud replication which includes Amazon Web Services (AWS) cloud replication. EC2 architecture capabilities add to HiveIO's cloud continuity with consistent and continuous replication of production workloads from local to AWS, delivering more options for surge events and disaster recovery.

Benefit

Leveraging EC2 for site replication or surge event augmentation provides the following benefits:

- The time needed to switch to and recover during a surge event or disaster event is easy to understand and manage.

- Expensive disaster recovery CAPEX is reduced and replaced with on-demand monthly service cost

Off-site Procedure

This following procedure was utilized to build out a lab on AWS. The goal was to test surge and disaster recovery from HiveIO Fabric Datacenter to AWS EC2. Overall the procedure is broken down into four phases:

Phase 1- Register and initialize billing on AWS EC2

1. Navigate to https://console.aws.amazon.com
2. Register as a new user and login
3. Type search criteria EC2 under the search box at the top of page
4. Follow Amazon steps to setup billing
   https://docs.aws.amazon.com/IAM/latest/UserGuide/tutorial_billing.html

Phase 2- Configure Amazon to enable scripted orchestration and upload image

1. Follow Amazon steps to create an S3 bucket

https://docs.aws.amazon.com/quickstarts/latest/s3backup/step-1-create-bucket.html

2. Validate HiveIO Raw Image is ready to upload Prerequisite checklist

https://docs.aws.amazon.com/vm-import/latest/userguide/vmie_prereqs.html

3. Upload HiveIO Raw image template from your NFS to S3 via the following two options

   a) Drag the file from your desktop to the S3 Bucket (setup in step #1)

   b) Create an S3FS Mount on your linux, OSX or Windows machine and then copy from local NFS to the S# Bucket (setup in step #1)

      i) Centos / Ubuntu Instructions

https://www.interserver.net/tips/kb/mount-s3-bucket-centos-ubuntu-using-s3fs/

      ii) OS X Instructions

https://gist.github.com/jcaddel/730dc58f4653adf33e4f

      iii) Windows Server / Desktop

https://tntdrive.com/mount-amazon-s3-bucket.aspx

   c) Setup an Amazon File Gateway Appliance on HiveIO and copy the HiveIO Raw image to the Amazon NFS Mount.

https://aws.amazon.com/premiumsupport/knowledge-center/file-gateway-ec2/

4)   Download and Install AWS CLI Tools to local laptop/ desktop

https://aws.amazon.com/cli/

5)   Create access keys from AWS for CLI

https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html

6)   Register access keys through CLI

 https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html#Using_CreateAccessKey_CLIAPI

7)   Create security group
        a) aws ec2 create-security-group --group-name workload-na --description "WORKLOAD NA"

8)   Create role = vmimport

        a) Add Permission = "AdministratorAccess"

        b) Edit "Trust relationships"

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "vmie.amazonaws.com"
      },
      "Action": "sts:AssumeRole",
      "Condition": {
        "StringEquals": {
          "sts:ExternalId": "vmimport"
        }
      }
    }
  ]
}
```

Phase 3- Using AWS CLI, import HiveIO RAW Image as AWS EC2 Template

1) Create container.json import manifest

```
[{
        "Description": "HiveIO",
        "Format": "raw",
        "UserBucket": {
                "S3Bucket": "bucketname",
                "S3Key": "hiveimagename.raw"
        }
}]
```

2) Import Image
   a) aws ec2 import-image --description "HiveIO" --disk-containers file://containers.json

3) Verify Status of import
   a) aws ec2 describe-import-image-tasks --import-task-ids import-ami-0705c5f6988c572cb

Phase 4- Launch newly imported image as EC2 Template

1. Find and launch new imported AMIs

   a. https://aws.amazon.com/premiumsupport/knowledge-center/launch-instance-custom-ami/

Conclusion

Utilizing EC2 as an offsite replication tool for surge and disaster recovery is entirely plausible. Based on transition time, HiveIO recommends pre-staging gold image templates on EC2 such that images can be quickly spun up and then stopped on-demand. Most customers utilize EC2 as an offsite replication/ compute DR and surge event tool for storage and virtual server(persistent) or non-persistent virtual desktops.

HiveIO highly recommends that customers run through EC2 transitions from end-to-end monthly in order to verify credentials and data backups so that DR and surge events transition flawlessly.

T-415.340.2089 | F – 415.715.9028

www.hiveio.com | info@hiveio.com | 🐦@HiveIOInc

Learn more at www.HiveIO.com or follow us on Twitter@HiveIO1.