INSTALLATION RUNBOOK FOR
Concurrent Aquari Storage

Product Name:	**Aquari Storage**
Driver Version:	**2.2**
MOS Version:	**9.0**
OpenStack Version:	**Mitaka**
Product Type:	**Software-Defined Storage Appliance**

Document History

Version	Revision Date	Description
0.1	12-26-2016	Final version

1. Introduction

This document is a detailed Deployment Guide for Aquari Storage which integrates with Mirantis OpenStack v9.0 Nova, Glance, Cinder and Swift for use as backend storage. It provides a reference architecture for deploying the solution and detailed installation instructions. The document also notes limitations of the solution, describes tests performed during validation, and provides the results of these tests.

1.1 Target Audience

This document is intended to be used by personnel installing and configuring Aquari as a backend storage for Mirantis OpenStack cloud.

2. Product Overview

Concurrent is a global software and solutions company headquartered in Atlanta, GA.
Our business is founded on enabling mission critical applications using our high performance Linux and storage technologies.
Aquari Storage is based on open, flexible and scalable SDS (Software Defined Storage) architecture purposefully designed to fulfill use-cases such as:

- Video Streaming and Online Archiving,
- OpenStack block storage,
- Enterprise File Sync & Share.

Multi-workload flexibility makes Aquari Storage well suited for the enterprises and service providers. Cluster access for individual users or applications is controlled via secure keyrings.

No Single Point of Failure Aquari architecture is designed for access to mission-critical assets in both private and public clouds.

Scalability is an essential factor driving down the amount of capital required for building the new clouds. Aquari seamlessly scales from TB to EB storage clusters.

Resiliency to failures boosts service availability metrics, an important differentiating factor as enterprises and cloud service providers (CSPs) are in the race for guaranteeing higher SLAs.

Aquari Storage Cluster Architecture

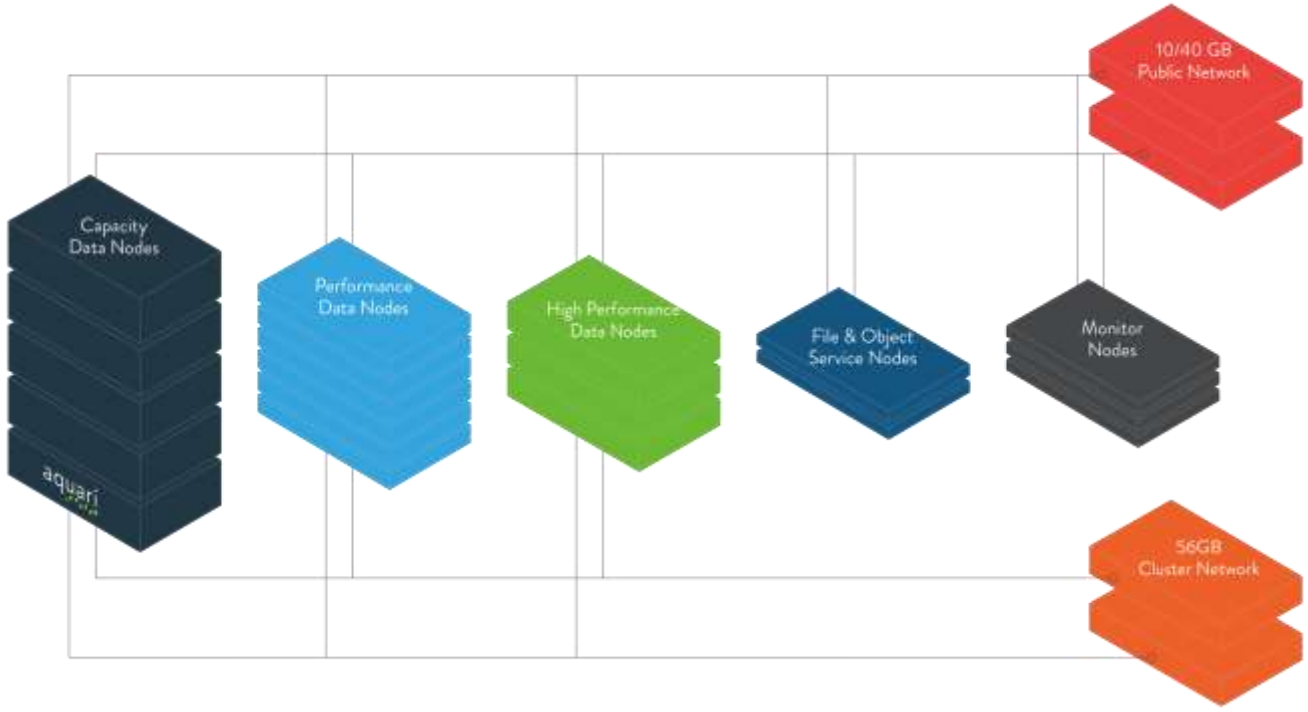

The Aquari Storage cluster consists of Data Nodes interconnected with fast 56GB network and connected to the Monitor Nodes via dedicated 10/40GB network. File (SMB, NFS) & Object (S3, Swift) services are enabled through the FOS Node (File & Object Services).
Depending on the application, we offer following Data Nodes:
- *High Performance* for nDVR, video streaming and other applications demanding high bandwidth.
- *Performance* are designed for OpenStack related use-cases.
- *Capacity* are the cost efficient storage for archives and backups.

Multiple Monitor Nodes continuously keep track of data integrity and health of various parts of the Aquari cluster.

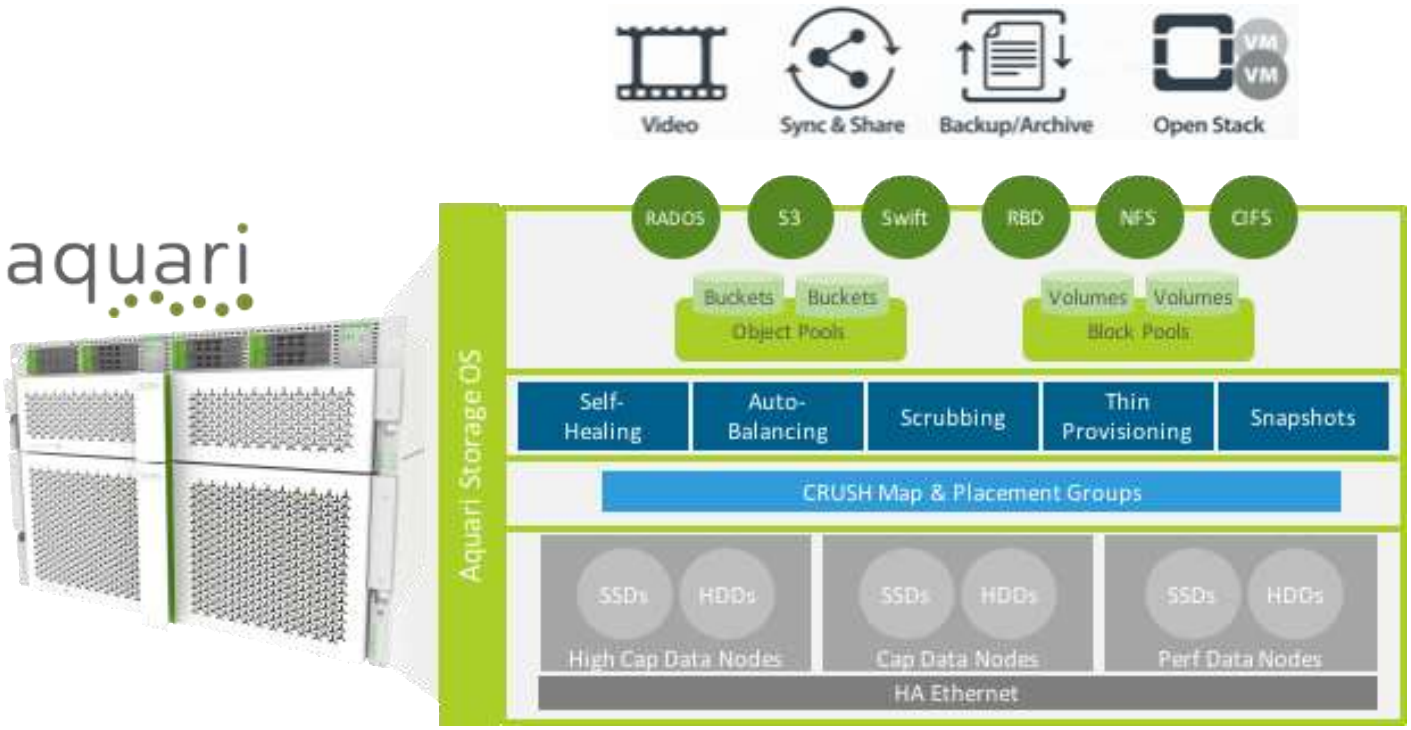

Aquari Storage integration with Mirantis OpenStack and services versatility opens support for numerous use-cases including but not limited to:

- OpenStack storage for virtual machines,
- Shared file storage for file-based applications,
- Object storage for images and snapshots.

3. Joint Reference Architecture

Aquari Storage integrates with OpenStack projects such as Nova, Cinder, Manila, Glance and Swift.

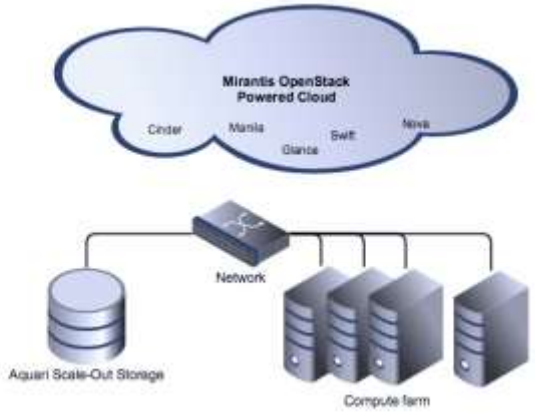

4. Physical and Logical Network Topology

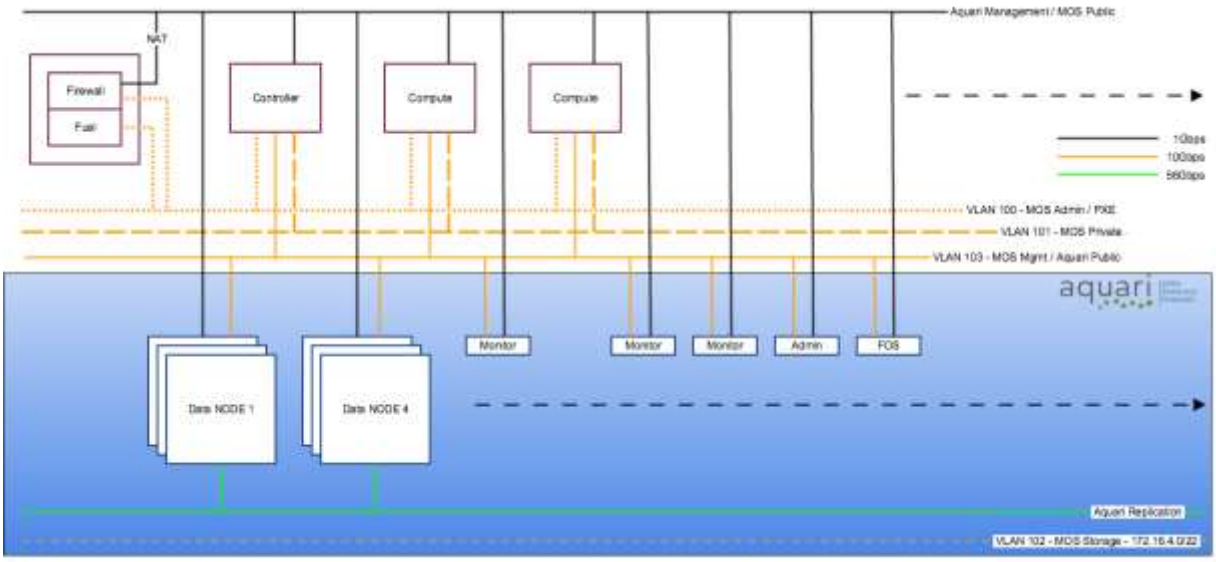

Picture above depicts logical network topology of the Mirantis OpenStack storage backend based on the Aquari Storage.

5. Installation and Configuration

5.1 Environment Preparation

Aquari UI driven installation process goes through the following phases:

a. Preparation of the configuration file
b. Bootstrapping the Administration and File Object Service (FOS) nodes
c. Provisioning monitors and storage nodes for the Aquari cluster topology. Aquari installer access bare metal hosts via PXE and deploys Aquari software on the bare metal hosts.

Detailed installation instructions can be found in Aquari Installation & Configuration Guide.

5.2 MOS Installation

Aquari Storage supports SMB and NFS shares, Swift and S3 object storage as well as block data endpoints.

Configuration steps required to connect MOS to Aquari Storage inlude:

- Install Fuel Ceph plugin;
- Provision pools for storing images (Glance), volumes (Cinder) and Ephemeral Drives (Nova) in Aquari UI;
- Configure Fuel Ceph plugin to use newly provisioned pools as per image bellow.

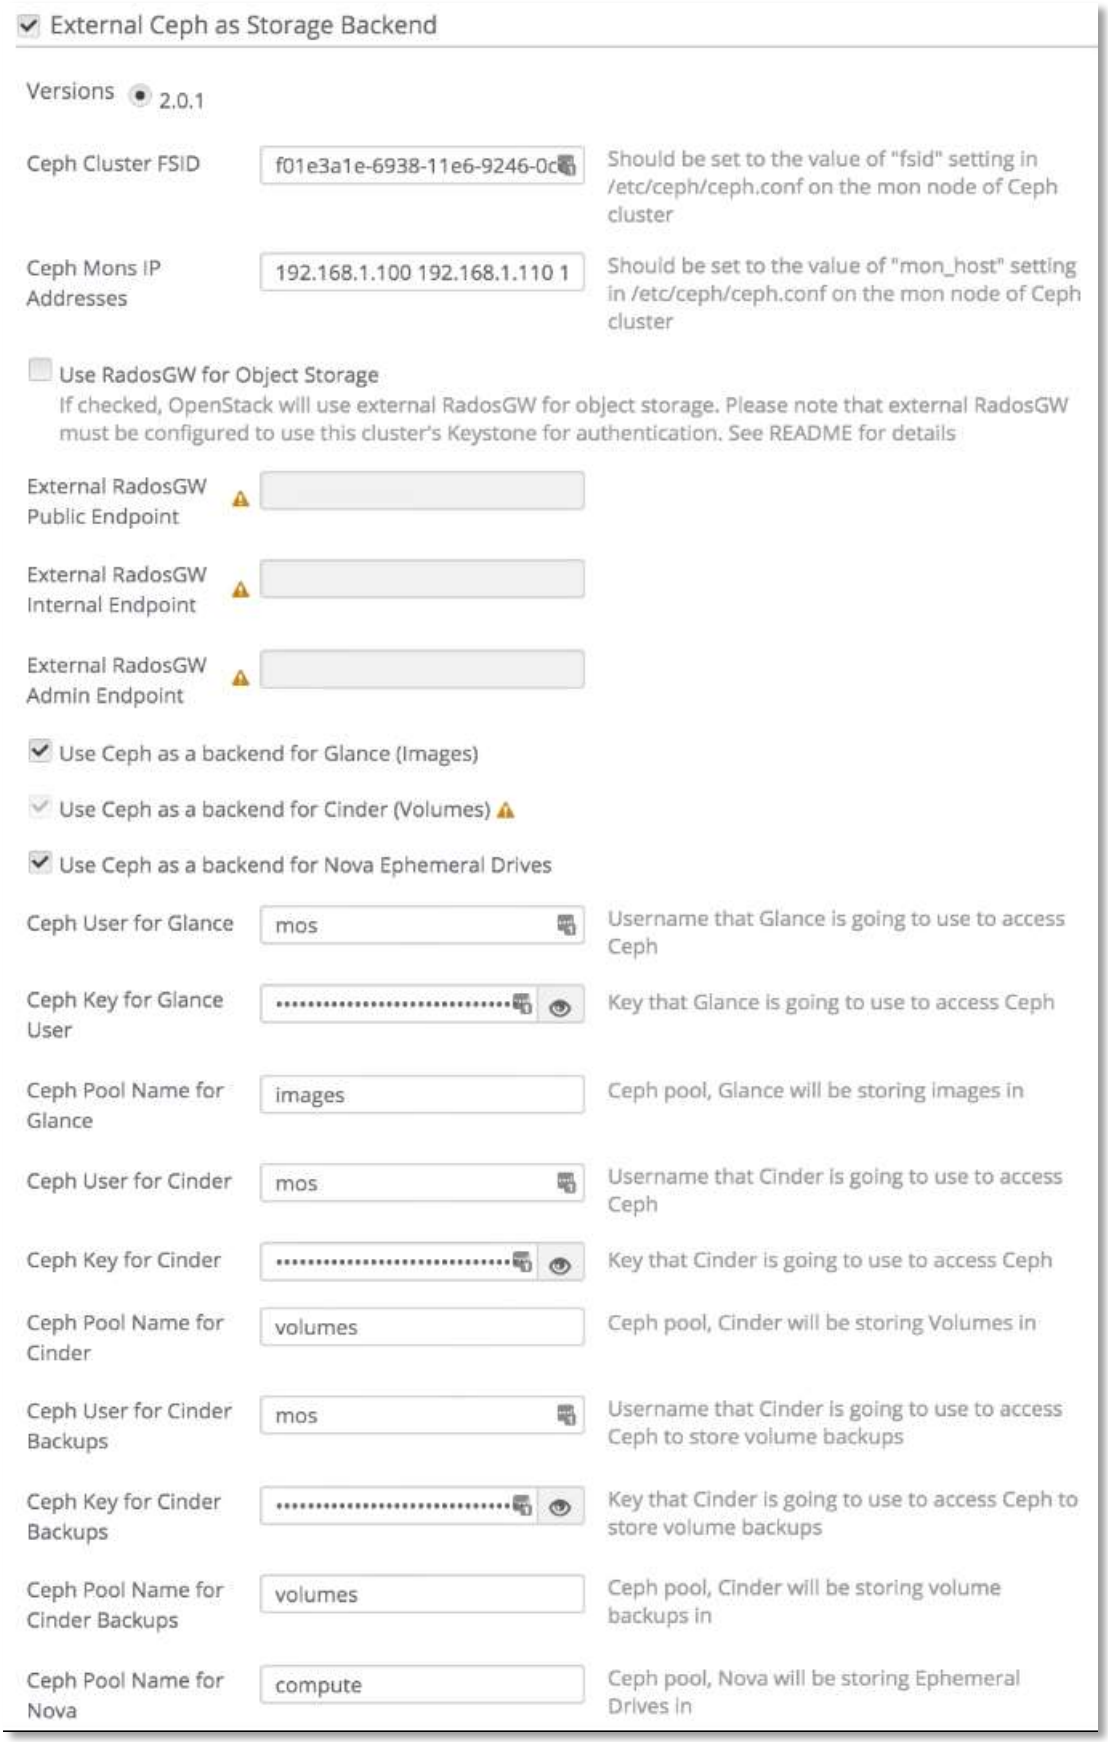
External Ceph as Storage Backend
Versions
2.0.1
Ceph Cluster FSID
f01e3a1e-6938-11e6-9246-0c
Should be set to the value of "fsid" setting in /etc/ceph/ceph.conf on the mon node of Ceph cluster
Ceph Mons IP Addresses
192.168.1.100 192.168.1.110 1
Should be set to the value of "mon_host" setting in /etc/ceph/ceph.conf on the mon node of Ceph cluster
Use RadosGW for Object Storage
If checked, OpenStack will use external RadosGW for object storage. Please note that external RadosGW must be configured to use this cluster's Keystone for authentication. See README for details
External RadosGW Public Endpoint
External RadosGW Internal Endpoint
External RadosGW Admin Endpoint
Use Ceph as a backend for Glance (Images)
Use Ceph as a backend for Cinder (Volumes)
Use Ceph as a backend for Nova Ephemeral Drives
Ceph User for Glance
mos
Username that Glance is going to use to access Ceph
Ceph Key for Glance User
Key that Glance is going to use to access Ceph
Ceph Pool Name for Glance
images
Ceph pool, Glance will be storing images in
Ceph User for Cinder
mos
Username that Cinder is going to use to access Ceph
Ceph Key for Cinder
Key that Cinder is going to use to access Ceph
Ceph Pool Name for Cinder
volumes
Ceph pool, Cinder will be storing Volumes in
Ceph User for Cinder Backups
mos
Username that Cinder is going to use to access Ceph to store volume backups
Ceph Key for Cinder Backups
Key that Cinder is going to use to access Ceph to store volume backups
Ceph Pool Name for Cinder Backups
volumes
Ceph pool, Cinder will be storing volume backups in
Ceph Pool Name for Nova
compute
Ceph pool, Nova will be storing Ephemeral Drives in