



# **Mellanox OpenStack Solution Reference Architecture**

Rev 1.0

April 2013

[www.mellanox.com](http://www.mellanox.com)

NOTE:

THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT (“PRODUCT(S)”) AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES “AS-IS” WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER’S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON INFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.



Mellanox Technologies  
350 Oakmead Parkway Suite 100  
Sunnyvale, CA 94085  
U.S.A.  
[www.mellanox.com](http://www.mellanox.com)  
Tel: (408) 970-3400  
Fax: (408) 970-3403

Mellanox Technologies, Ltd.  
Beit Mellanox  
PO Box 586 Yokneam 20692  
Israel  
[www.mellanox.com](http://www.mellanox.com)  
Tel: +972 (0)74 723 7200  
Fax: +972 (0)4 959 3245

© Copyright 2013. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, ConnectX®, CORE-Direct®, InfiniBridge®, InfiniHost®, InfiniScale®, MLNX-OS®, PhyX®, SwitchX®, UFM®, Virtual Protocol Interconnect® and Voltaire® are registered trademarks of Mellanox Technologies, Ltd.

Connect-IB™, FabricIT™, Mellanox Open Ethernet™, MetroX™, MetroDX™, ScalableHPC™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners

# Contents

<b>1</b>	<b>Overview .....</b>	<b>7</b>
<b>2</b>	<b>Storage Acceleration .....</b>	<b>9</b>
<b>3</b>	<b>Network Virtualization .....</b>	<b>11</b>
3.1	Performance .....	12
3.2	Quality of Service (QoS) .....	12
3.3	Seamless OpenStack Integration .....	14
<b>4</b>	<b>Setup and Installation .....</b>	<b>15</b>
4.1	Basic Setup .....	15
4.2	Hardware Requirements .....	15
4.3	Software Requirements .....	16
4.4	Prerequisites .....	16
4.5	Software installation .....	16
<b>5</b>	<b>Setting Up the Network .....</b>	<b>17</b>
5.1	Configuration Examples .....	17
5.1.1	Creating a Network .....	17
5.1.2	Creating an Instance (Para-Virtualized vNIC) .....	18
5.1.3	Creating an Instance (SR-IOV) .....	19
5.1.4	Creating a Volume .....	20
5.1.5	Binding a Volume .....	21
5.2	Verification Examples .....	22
5.2.1	Instances Overview .....	22
5.2.2	Connectivity Check .....	22
5.2.3	Volume Check .....	22
	<b>Appendix A: Scale-out Architecture .....</b>	<b>23</b>

## List of Figures

Figure 1: Mellanox OpenStack Architecture .....	8
Figure 2: OpenStack Based IaaS Cloud POD Deployment Example .....	9
Figure 3: RDMA Acceleration .....	10
Figure 4: eSwitch Architecture.....	11
Figure 5: Latency Comparison .....	12
Figure 6: QoS, Setup Example.....	13
Figure 7: QoS, Test Results .....	13
Figure 8: Network Virtualization.....	14
Figure 9: Mellanox MCX314A-BCBT, ConnectX-3 40GbE Adapter.....	15
Figure 10: Mellanox SX1036, 36x40GbE .....	15
Figure 11: Mellanox 40GbE, QSFP Copper Cable .....	16
Figure 12: Quantum net-create/subnet-create Commands.....	18
Figure 13: OpenStack Dashboard, Instances .....	18
Figure 14: OpenStack Dashboard, Launch Instance .....	19
Figure 15: OpenStack Dashboard, Launch Interface - Select Network .....	19
Figure 16: Quantum port-create Command .....	20
Figure 17: Using the nova boot Command.....	20
Figure 18: OpenStack Dashboard, Volumes.....	20
Figure 19: OpenStack Dashboard, Create Volumes .....	21
Figure 20: OpenStack Dashboard, Volumes.....	21
Figure 21: OpenStack Dashboard, Manage Volume Attachments .....	21
Figure 22: VM Overview .....	22
Figure 23: Remote Console Connectivity .....	22
Figure 24: OpenStack Dashboard, Volumes.....	23
Figure 25: OpenStack Dashboard, Console.....	23
Figure 26: Scale-out Architecture .....	24

# Preface

## About this Document

This reference design presents the value of using Mellanox interconnect products and describes how to integrate the OpenStack solution with the end-to-end Mellanox interconnect solution.

## Audience

This reference design is intended for server and network administrators.

The reader must have experience with the basic OpenStack framework and installation.

## Document Conventions

The following lists conventions used in this document.



**NOTE:** Identifies important information that contains helpful suggestions.



**CAUTION:** Alerts you to the risk of personal injury, system damage, or loss of data.



**WARNING:** Warns you that failure to take or avoid a specific action might result in personal injury or a malfunction of the hardware or software. Be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents before you work on any equipment.

## References

For additional information, see the following documents:

**Table 1: Related Documentation**

Reference	Location
Mellanox OFED User Manual	<a href="http://www.mellanox.com">www.mellanox.com</a> > Products > Adapter IB/VPI SW > Linux SW/Drivers <a href="http://www.mellanox.com/content/pages.php?pg=products_dyn&amp;product_family=26&amp;menu_section=34">http://www.mellanox.com/content/pages.php?pg=products_dyn&amp;product_family=26&amp;menu_section=34</a>
Mellanox software source packages	<a href="https://github.com/mellanox-openstack">https://github.com/mellanox-openstack</a>
OpenStack Website	<a href="http://www.openstack.org">www.openstack.org</a>
Mellanox OpenStack wiki page	<a href="https://wiki.openstack.org/wiki/Mellanox-OpenStack">https://wiki.openstack.org/wiki/Mellanox-OpenStack</a>
Mellanox approved cables	<a href="http://www.mellanox.com/related-docs/user_manuals/Mellanox_approved_cables.pdf">http://www.mellanox.com/related-docs/user_manuals/Mellanox_approved_cables.pdf</a>
Mellanox Ethernet Switch Systems User Manual	<a href="http://www.mellanox.com/related-docs/user_manuals/SX10X_User_Manual.pdf">http://www.mellanox.com/related-docs/user_manuals/SX10X_User_Manual.pdf</a>
Mellanox Ethernet adapter cards	<a href="http://www.mellanox.com/page/ethernet_cards_overview">http://www.mellanox.com/page/ethernet_cards_overview</a>

# 1 Overview

Deploying and maintaining a private or public cloud is a complex task – with various vendors developing tools to address the different aspects of the cloud infrastructure, management, automation, and security. These tools tend to be expensive and create integration challenges for customers when they combine parts from different vendors. Traditional offerings suggest deploying multiple network and storage adapters to run management, storage, services, and tenant networks. These also require multiple switches, cabling, and management infrastructure, which increases both up front and maintenance costs.

Other, more advanced offerings provide a unified adapter and first level ToR switch, but still run multiple and independent core fabrics. Such offerings tend to suffer from low throughput because they do not provide the aggregate capacity required at the edge or in the core; and because they deliver poor application performance due to network congestion and lack of proper traffic isolation.

Several open source “cloud operating system” initiatives have been introduced to the market, but none has gained sufficient momentum to succeed. Recently OpenStack has managed to establish itself as the leading open source cloud operating system, with wide support from major system vendors, OS vendors, and service providers. OpenStack allows central management and provisioning of compute, networking, and storage resources, with integration and adaptation layers allowing vendors and/or users to provide their own plug-ins and enhancements.

Mellanox Technologies offers seamless integration between its products and OpenStack layers and provides unique functionality that includes application and storage acceleration, network provisioning, automation, hardware-based security, and isolation. Furthermore, using Mellanox interconnect products allows cloud providers to save significant capital and operational expenses through network and I/O consolidation and by increasing the number of virtual machines (VMs) per server.


Mellanox provides a variety of network interface cards (NICs) supporting one or two ports of 10GbE, 40GbE, or 56Gb/s InfiniBand. These adapters simultaneously run management, network, storage, messaging, and clustering traffic. Furthermore, these adapters create virtual domains within the network that deliver hardware-based isolation and prevent cross-domain traffic interference.

In addition, Mellanox Virtual Protocol Interconnect (VPI) switches deliver the industry’s most cost-effective and highest capacity switches (supporting up to 36 ports of 56Gb/s). When deploying large-scale, high-density infrastructures, leveraging Mellanox converged network VPI solutions translates into fewer switching elements, far fewer optical cables, and simpler network design.

## **Mellanox integration with OpenStack provides the following benefits:**

- Cost-effective and scalable infrastructure that consolidates the network and storage to a highly efficient flat fabric, increases the VM density, commoditizes the storage infrastructure, and linearly scales to thousands of nodes
- Delivers the best application performance with hardware-based acceleration for messaging, network traffic, and storage
- Easy to manage via standard APIs. Native integration with OpenStack Quantum (network) and Cinder (storage) provisioning APIs
- Provides tenant and application security/isolation, end-to-end hardware-based traffic isolation, and security filtering


Figure 1: Mellanox OpenStack Architecture



## 2 Storage Acceleration

Data centers rely on communication between compute and storage nodes, as compute servers read and write data from the storage servers constantly. In order to maximize the server's application performance, communication between the compute and storage nodes must have the lowest possible latency, highest possible bandwidth, and lowest CPU utilization.

**Figure 2: OpenStack Based IaaS Cloud POD Deployment Example**



Storage applications that use iSCSI over TCP are processed by the CPU. This causes data center applications that rely heavily on storage communication to suffer from reduced CPU utilization, as the CPU is busy sending data to the storage servers. The data path for protocols such as TCP, UDPO, NFS, and iSCSI all must wait in line with the other applications and system processes for their turn using the CPU. This not only slows down the network, but also uses system resources that could otherwise have been used for executing applications faster.

Mellanox OpenStack solution extends the Cinder project by adding iSCSI running over RDMA (iSER). Leveraging RDMA Mellanox OpenStack delivers 5x better data throughput (for example, increasing from 1GB/s to 5GB/s) and requires up to 80% less CPU utilization (see Figure 3).

Mellanox ConnectX®-3 adapters bypass the operating system and CPU by using RDMA, allowing much more efficient data movement paths. iSER capabilities are used to accelerate hypervisor traffic, including storage access, VM migration, and data and VM replication. The use of RDMA moves data to the Mellanox ConnectX-3 hardware, which provides zero-copy message transfers for SCSI packets to the application, producing significantly faster performance, lower network latency, lower access time, and lower CPU overhead. iSER can provide 6x faster performance than traditional TCP/IP based iSCSI. This also consolidates the efforts of both Ethernet and InfiniBand communities, and reduces the number of storage protocols a user must learn and maintain.


The RDMA bypass allows the data path to effectively skip to the front of the line. Data is provided directly to the application immediately upon receipt without being subject to various delays due to CPU load-dependent software queues. This has three effects:

- There is no waiting, which means that the latency of transactions is incredibly low.
- Because there is no contention for resources, the latency is consistent, which is essential for offering end users with a guaranteed SLA.
- By bypassing the OS, using RDMA results in significant savings in CPU cycles. With a more efficient system in place, those saved CPU cycles can be used to accelerate application performance.

In the following diagram, it is clear that by performing hardware offload of the data transfers using the iSER protocol, the full capacity of the link is utilized to the maximum of the PCIe limit.

To summarize, network performance is a significant element in the overall delivery of data center services. To produce the maximum performance for data center services requires fast interconnects. Unfortunately the high CPU overhead associated with traditional storage adapters prevents taking full advantage of high speed interconnects. Many more CPU cycles are needed to process TCP and iSCSI operations compared to that required by the RDMA (iSER) protocol performed by the network adapter. Hence, using RDMA-based fast interconnects significantly increases data center performance levels.

Figure 3: RDMA Acceleration





- Security: The eSwitch enables network isolation (using VLANs) and anti-MAC spoofing. In addition, by using OpenFlow ACLs, the eSwitch can be configured to filter undesired network flows.
- QoS: The eSwitch supports traffic class management, priority mapping, rate limiting, scheduling, and shaping configured via OpenFlow. In addition, DCBX control plane can set Priority Flow Control (PFC) and FC parameters on the physical port.
- Monitoring: Port counters are supported.


### 3.1 Performance

Many data center applications require lower latency network performance. Some applications require latency stability as well. Using regular TCP connectivity between VMs can create high latency and unpredictable delay behavior.

Figure 5 shows the dramatic difference (x20) when using para-virtualized vNIC running a TCP stream compared to SR-IOV connectivity running RDMA.

Due to the direct connection of the SR-IOV and the ConnectX-3 hardware capabilities, there is a significant reduction in software interference that adds unpredictable delay to the packet processing.

Figure 5: Latency Comparison



### 3.2 Quality of Service (QoS)

The impact of using QoS and network isolation is tremendous. The following example compares the various latency and bandwidth levels as a function of the QoS level.

The following test reveals the great advantage that can be achieved using the switch QoS capability:

Setup characteristics:

Streams:

In this test, two types of streams were injected:

- Blue (Storage stream): TCP stream in high volume. Latency is not crucial for such an application.
- Green (Messaging stream): Round robin (RR) TCP stream in low volume. Latency is crucial for such an application.

### QoS levels:

The following QoS levels were tested:

- Single Queue: both streams egress one queue.
- Dual Queues with no QoS: each stream egress a different queue while both queues have the same priority level.
- Dual Queues with QoS enabled: the green stream is prioritized over the blue stream.


**Figure 6: QoS, Setup Example**



The test results show the following:

- (1) When prioritizing a stream (green) and using dual queues, the low priority stream has a minor effect on the high priority stream (11.8 $\mu$ sec compared to 10.8 $\mu$ sec in Figure 7).
- (2) Bandwidth increases when prioritizing streams (9350GbE), as well as when increasing the number of queues (9187GbE), compared to regular non-QoS conditions (8934GbE).
- (3) The latency difference is dramatically reduced when using QoS (11.8 $\mu$ sec compared to 10548 $\mu$ sec).

**Figure 7: QoS, Test Results**



\* Results are based on 10GbE adapter card


### Conclusion:

The test results emphasize that consolidation is possible on the same physical port. Applications that require low latency will not suffer from bandwidth-consuming applications when using more than one queue and enabling QoS.

### 3.3 Seamless OpenStack Integration

The eSwitch configuration is transparent to the OpenStack administrator. The installed eSwitch daemon on the server is responsible for hiding the low-level configuration. The administrator will use the OpenStack dashboard APIs for the fabric management.

Figure 8: Network Virtualization



## 4 Setup and Installation

### 4.1 Basic Setup

The following setup is suggested for small scale applications.

The OpenStack environment should be installed according to the OpenStack installation guide.

In addition, the following installation changes should be applied:

- A Quantum server should be installed with the Mellanox Quantum plugin.
- A Cinder patch should be applied to the storage servers (for iSER support).
- Mellanox Quantum agent, eSwitch daemon, and Nova patches should be installed on the compute nodes.

### 4.2 Hardware Requirements

- Mellanox ConnectX-3 adapter cards
- 10GbE or 40GbE Ethernet switches
- Cables required for the ConnectX-3 card (typically using SFP+ connectors for 10GbE or QSFP connectors for 40GbE)
- Server nodes should comply with OpenStack requirements.
- Compute nodes should have SR-IOV capability (BIOS and OS support).

There are many options in terms of adapters, cables, and switches. See [www.mellanox.com](http://www.mellanox.com) for additional options.

**Figure 9: Mellanox MCX314A-BCBT, ConnectX-3 40GbE Adapter**



**Figure 10: Mellanox SX1036, 36x40GbE**



**Figure 11: Mellanox 40GbE, QSFP Copper Cable**

### 4.3 Software Requirements

- Supported OS
  - RHEL 6.3 or higher
- Mellanox OFED 2.0 (SR-IOV support)
- KVM hypervisor – complying with OpenStack requirements

### 4.4 Prerequisites

- (1) The basic setup is physically connected.
  - In order to reduce the number of ports in the network, two different networks can be mapped to the same physical interface on two different VLANs.
- (2) Mellanox OFED 2.0 (SR-IOV enabled) is installed on each of the network adapters.
  - See Mellanox Community – Cloud developer zone for verification options and adaptation.  
<http://community.mellanox.com/community/develop/cloud-developers/blog/2013/04/07/mellanox-ofed-driver-installation-with-sr-iov>
- (3) The OpenStack package is installed on all network elements.

### 4.5 Software installation

For Mellanox OpenStack installation, follow the Mellanox OpenStack wiki pages:

- Quantum: <https://wiki.openstack.org/wiki/Mellanox-Quantum>
- Cinder: <https://wiki.openstack.org/wiki/Mellanox-Cinder>

For the eSwitch daemon installation, follow the OpenStack wiki pages (part of Mellanox Quantum):

- <https://wiki.openstack.org/wiki/Mellanox-Quantum>

## 5 Setting Up the Network

### 5.1 Configuration Examples

Once the installation is completed, it is time to set up the network.

Setting up a network consists of the following steps:

- (1) Creating a network
- (2) Creating a VM instance. Two types of instances can be created:
  - a. Para-virtualized vNIC
  - b. SR-IOV direct path connection
- (3) Creating a disk volume
- (4) Binding the disk volume to the instance that was just created

#### 5.1.1 Creating a Network

Use the `quantum net-create` and `quantum subnet-create` commands to create a new network and a subnet ('net3' in the example).

Figure 12: Quantum net-create/subnet-create Commands

```

root@xena016-2:~# quantum net-create net3
Created a new network:
+-----+-----+
| Field                | Value                                     |
+-----+-----+
| admin_state_up      | True                                     |
| id                   | 97127c33-0095-4776-abb5-20420987009a   |
| name                 | net3                                     |
| provider:network_type | vlan                                     |
| provider:physical_network | mlx1                                     |
| provider:segmentation_id | 3                                       |
| router:external      | False                                    |
| shared               | False                                    |
| status               | ACTIVE                                   |
| subnets             |                                           |
| tenant_id           | 9511bf012f00467481022e3235b5786f     |
+-----+-----+
root@xena016-2:~# quantum subnet-create --name usbnet_net3 net3 192.168.203.0/24
Created a new subnet:
+-----+-----+
| Field                | Value                                     |
+-----+-----+
| allocation_pools     | {"start": "192.168.203.2", "end": "192.168.203.254"} |
| cidr                 | 192.168.203.0/24                       |
| dns_nameservers      |                                           |
| enable_dhcp          | True                                     |
| gateway_ip           | 192.168.203.1                           |
| host_routes          |                                           |
| id                   | e0896b22-3668-4fce-bf8c-a843001858d5   |
| ip_version           | 4                                       |
| name                 | usbnet_net3                             |
| network_id           | 97127c33-0095-4776-abb5-20420987009a   |
| tenant_id           | 9511bf012f00467481022e3235b5786f     |
+-----+-----+

```

## 5.1.2 Creating an Instance (Para-Virtualized vNIC)


- (1) Using the OpenStack Dashboard, launch an instance (VM) using the Launch Instance button.
- (2) Insert all the required parameters and click Launch.

This operation will create a macvtap interface on top of a Virtual Function (VF).

Figure 13: OpenStack Dashboard, Instances




Figure 14: OpenStack Dashboard, Launch Instance



(3) Select the desired network for the vNIC ('net3' in the example).

Figure 15: OpenStack Dashboard, Launch Interface - Select Network



### 5.1.3 Creating an Instance (SR-IOV)

(1) Use the `quantum port-create` command for the selected network ('net3' in the example) to create a port with `'vnic_type=hostdev'`.

Figure 16: Quantum port-create Command

```
root@xena016-2:~# quantum port-create net3 --binding:profile type=dict vnic_type=hostdev
Created a new port:
```

Field	Value
admin_state_up	True
binding:capabilities	{"port_filter": false}
binding:profile	{"physical_network": "mlx1"}
binding:vif_type	hostdev
device_id	
device_owner	
fixed_ips	{"subnet_id": "e0896b22-3668-4fce-bf8c-a843001858d5", "ip_address": "192.168.203.5"}
id	099a1db7-8b0a-4a7e-8045-2fd4cff4c17f
mac_address	fa:16:3e:b2:f1:6f
name	
network_id	97127c33-0095-4776-abb5-20420987009a
status	ACTIVE
tenant_id	9511bf012f00467481022e3235b5786f

(2) Use the `nova boot` command to launch an instance with the created port attached.

Figure 17: Using the nova boot Command

```
root@xena016-2:~# nova boot --flavor m1.tiny --image rh6.3 vm3 --nic port-id=099a1db7-8b0a-4a7e-8045-2fd4cff4c17f
```

Property	Value
OS-EXT-STS:task_state	scheduling
image	rh6.3
OS-EXT-STS:vm_state	building
OS-EXT-SRV-ATTR:instance_name	instance-00000029
flavor	m1.tiny
id	f2b0d3ec-64b8-473b-866a-b29264f13219
security_groups	[{"name": "default"}]
user_id	2573e2ffdef241c29c932cc76f01e583
OS-DCF:diskConfig	MANUAL
accessIPv4	
accessIPv6	
progress	0
OS-EXT-STS:power_state	0
OS-EXT-AZ:availability_zone	None
config_drive	
status	BUILD
updated	2013-04-04T11:23:15Z
hostId	
OS-EXT-SRV-ATTR:host	None
key_name	None
OS-EXT-SRV-ATTR:hypervisor_hostname	None
name	vm3
adminPass	Uz3RPE3TvnJf
tenant_id	9511bf012f00467481022e3235b5786f
created	2013-04-04T11:23:15Z
metadata	{}

### 5.1.4 Creating a Volume

Create a volume using the Volumes tab on the OpenStack dashboard. Click the Create Volume button.

Figure 18: OpenStack Dashboard, Volumes




Figure 19: OpenStack Dashboard, Create Volumes

Figure 20: OpenStack Dashboard, Volumes

Volumes

Info: Creating volume "vol1"

+ Create Volume    Delete Volumes

<input type="checkbox"/>	Name	Description	Size	Status	Type	Attached To	Actions
<input type="checkbox"/>	vol1		5GB	Available	-		Edit Attachments    More

Displaying 1 item

### 5.1.5 Binding a Volume

Bind a volume to the desired instance.


Figure 21: OpenStack Dashboard, Manage Volume Attachments

## 5.2 Verification Examples

### 5.2.1 Instances Overview

Use the OpenStack Dashboard to view all configured instances.

Figure 22: VM Overview




### 5.2.2 Connectivity Check

There are many options for checking connectivity between different instances, one of which is simply to open a remote console and ping the required host.

To launch a remote console for a specific instance, select the Console tab and launch the console.

Figure 23: Remote Console Connectivity



### 5.2.3 Volume Check


To verify that the created volume is attached to a specific instance, click the Volumes tab.

Figure 24: OpenStack Dashboard, Volumes



Additionally, run the `fdisk` command from the instance console to see the volume details.

Figure 25: OpenStack Dashboard, Console



## Appendix A: Scale-out Architecture

A scale-out option is an important issue for cloud providers. Compute and storage networks should have the ability to scale easily and effectively with high availability options.

Figure 26: Scale-out Architecture

