
Hello


SOLIDFIRE
SOLIDFIRE
The All-Flash Array for the
Next Generation Data Center
IOPS

Agenda

- Storage and OpenStack
  - What do you mean when you say ‘storage’
  - Object vs. block
- What is Cinder?
- Considerations when thinking about a Cinder backend
- Configuring your storage for DBaaS
- OpenStack Trove and DBaaS
- The Future: Starring Trove and SolidFire

The IT Innovation Imperative

- Deploy new applications and services faster
- Provide more agile and scalable infrastructure
- Increase application performance and predictability
- Enable automation and end-user self-service
- Raise operational efficiency and reduce cost

OpenStack is driving the transformation of Enterprise IT globally

Quick Poll:

- How many of you contribute to OpenStack?
- How many of you are end-users of OpenStack?
- How many of you are OpenStack Operators?
- How many of you work for Vendor Organizations that contribute to OpenStack?
- How many are “all of the above”?

So…What do you mean when you say ‘storage’ ?

Storage. Storage. Storage. Storage. Hodor.

- **Ephemeral**
  - Non-persistent
  - Lifecycle coincides with a Nova instance
- **Object**
  - Manages data as.. well, an Object
  - Think unstructured data (Mp4)
  - Typically referred to “cheap and deep”
  - Usually runs on spinny drives
  - In OpenStack: Swift
- **Files System**
  - We all love NFS and CIFS…right!?

- **Block**
  - Foundation for the other types
  - Think raw disk
  - Typically higher performance
  - Cinder

Quick Poll:

	Block Storage	Object Storage
What does it do?	• Storage for running VM disk volumes on a host • Ideal for performance sensitive apps • Enables Amazon EBS-like service	• Ideal for cost effective, scale-out storage • Fully distributed, API-accessible • Well suited for backup, archiving, data retention • Enables Dropbox-like service
Use Cases	• Production Applications • Traditional IT Systems • Database Driven Apps • Messaging / Collaboration • Dev / Test Systems	• VM Templates • ISO Images • Disk Volume Snapshots • Backup / Archive • Image / Video Repository
Workloads	• High Change Content • Smaller, Random R/W • Higher / “Bursty” IO	• Typically More Static Content • Larger, Sequential R/W • Lower IOPS

Let’s talk Cinder!

Cinder Mission Statement

To implement services and libraries to provide on demand, self-service access to Block Storage resources. Provide Software Defined Block Storage via abstraction and automation on top of various traditional backend block storage devices.

Huh?

Very Simple.

Cinder is an API that allows you to dynamically create/attach/detach disks to your Nova instance.

Deep thoughts.

How it works…

- Plugin architecture, use whatever storage backend you want
- Consistent API, agnostic to the infrastructure
- Expose differentiating features via custom volume-types and extra-specs

Reference Implementation Included

- Includes code to provide a base implementation using LVM (just add disks)
- Great for POC and getting started
- Sometimes good enough
- Might be lacking for your performance, H/A and Scaling needs (it all depends)
- Can Scale by adding Nodes
- Cinder-Volume Node utilizes it’s local disks (allocate by creating an LVM VG)
- Cinder Volumes are LVM Logical Volumes, with an iSCSI target created for eac
- Typical max size recommendations per VG/Cinder-Volume backend ~ 5 TB
- No Redundancy (yet)

When LVM is not Enough…

Datera
fujitsu_eternus
Fusionio
hitachi-hbsd
hauwei
Nimble
Prophetstor
pure
Zfssa
netapp

coraid
emc-vmax
emc-xtremio
eqlx
glusterfc
hds
ibm-gpfs
ibm-xiv
Lvm
Nfs
nexenta

Ceph RBD
HP-3Par
HP-LeftHand
scality
sheepdog
smbfs
solidfire
vmware-vmdk
window-hyperv
zadara

Choosing a Vendor = the Hardest Part

- Ask yourself:

- Does it scale? How?
- Is it tested? Will it really work in OpenStack?
- How is it supported?
- How is performance and how does it handle noisy neighbors?
- Is the vendor active in the community?
- How easy is it to stand up?

Now…a word from our Sponsor

All-flash storage platform for the next generation data center.

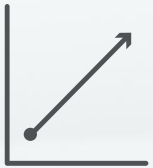

Scale-Out
Infrastructure Agility

Self Healing
High Availability

Guaranteed
Quality of Service

Complete
System Automation

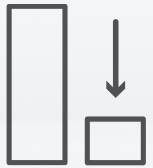

In-Line Data
Reduction

The Block Storage Choice for OpenStack

Deep OpenStack Integration

- SolidFire Cinder driver enables all OpenStack block storage features
- Ability to set and maintain true QoS levels on a per-volume basis
- Create, snap, clone and manage SolidFire volumes directly
- Run OpenStack instances on a SolidFire volume
- Eliminates arduous management layers between OpenStack and the storage system
- SolidFire driver fully integrated into OpenStack - no additional features / licenses required

Validated Interoperability

POWERED BY rackspace the open cloud company

redhat OPENSTACK CLOUD INFRASTRUCTURE PARTNER

MIRANTIS Pure Play OpenStack

Customer Success

ebay Comcast

iweb

PayPal

mercado libre

INTERNAP

BRINKSTER

QuoVadis

tableau SOFTWARE

SolidFire and Cinder

- SolidFire led the creation of Cinder (break out from Nova)
- Full SolidFire driver integration with every new OpenStack release
- Set and maintain true QoS levels on a per-volume basis
- Web-based API exposing all cluster functionality
- SolidFire integration with OpenStack Cinder can be configured in less than a minute
- Seamless scaling after initial configuration
- Full multi-tenant isolation

Configuring SolidFire in <60 Seconds

Edit the cinder.conf file

- *volume_driver=cinder.volume.solidfire.SolidFire*
- *san_ip=172.17.1.182*
- *san_login=openstack-admin*
- *san_password=superduperpassword*

Eliminating Noisy Neighbor

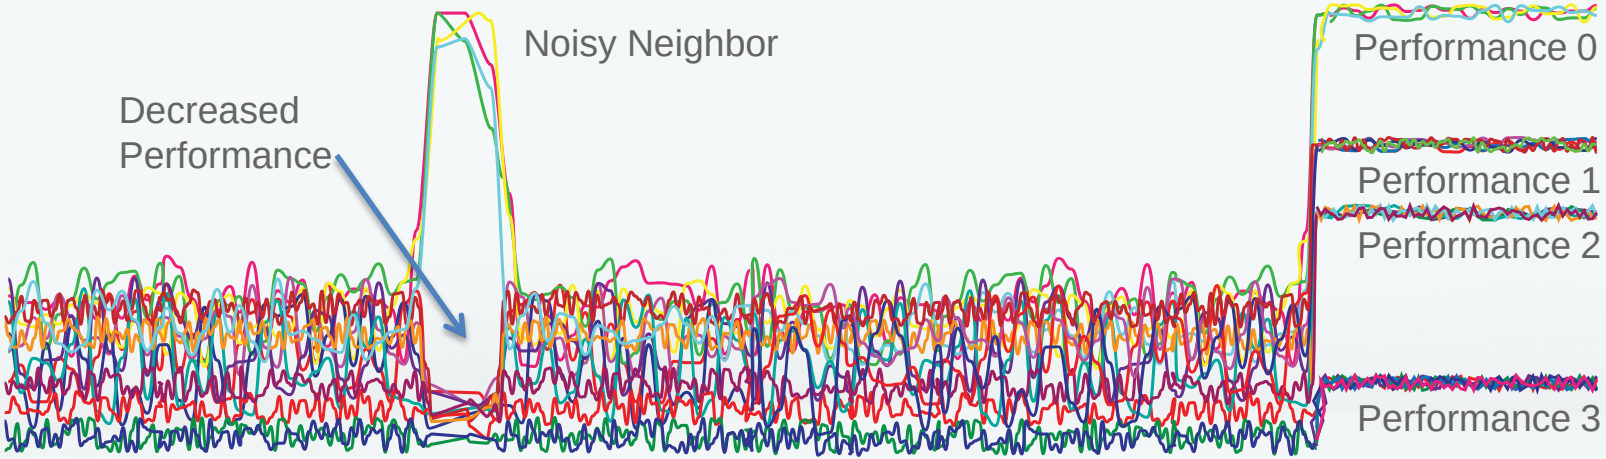

The Noisy Neighbor Effect

- Individual tenant impacts other applications
- Unsuitable for performance sensitive apps

SolidFire QoS in Practice

- Create fine-grained tiers of performance
- Application performance is isolated
- Performance SLAs enforced

Creating Volume-Types and Extra-Specs

```
griff@stack-1: cinder type create super
+--------------------------------------+-------+
|                  ID                  |  Name |
+--------------------------------------+-------+
| c506230f-eb08-4d4e-82e2-7a88eb779bda | super |
+--------------------------------------+-------+
griff@stack-1: cinder type create super-dooper
+--------------------------------------+--------------+
|                  ID                  |     Name     |
+--------------------------------------+--------------+
| 918cf343-1f3d-4508-bb69-cd0e668ae297 | super-dooper |
+--------------------------------------+--------------+
griff@stack-1: cinder type-key super set volume_backend_name=LVM_iSCSI
griff@stack-1: cinder type-key super-dooper set volume_backend_name=SolidFire \
qos:minIOPS=400 qos:maxIOPS=1000 qos:burstIOPS=2000
```

Configuring Storage for DBaaS

- Database as a Service
  - Provisioning simplification
  - Automation and SDS
  - Guarantee service levels
  - Protect against problem users
  - Advanced functionality: data set deployment

Configuring Storage for DBaaS

- IOPS and Bandwidth (QoS)
  - Min IOPS guarantees
  - Max IOPS ceilings
  - Burst IOPS credits
  - Alternatively gate/allow by bandwidth

Configuring Storage for DBaaS

- MySQL Workload Characteristics
  - 16K avg block size
  - Flushing can be bursty
  - Depends on various tunings (innodb_io_capacity)
  - Set Max IOPS and Burst accordingly
- Other DB Workload Characteristics
  - 4K to 64k avg block size (or larger)
  - Flushing can be *very* bursty
  - Sometimes controllable
  - Set Max IOPS and Burst to match DBMS profile

Boom. Time for Tesora.

OpenStack Trove: Database as a Service

- Simplifies use of databases in the enterprise
- Not another database
- Currently supports
  - MySQL
  - Percona
  - MariaDB
  - PostgreSQL
  - MongoDB
  - CouchDB
  - Couchbase
  - Redis
  - Vertica
  - DB2 – Express

The OpenStack Open Source Database as a Service Mission: To provide scalable and reliable Cloud Database as a Service provisioning functionality for both relational and non-relational database engines, and to continue to improve its fully-featured and extensible open source framework.

https://wiki.openstack.org/wiki/Trove

The load that databases place on infrastructure

- Compute
- Storage
  - Random Access
  - Periodicity
  - Write pattern

What this means to me (as an administrator)

What should I monitor?

- *Compute:*
  - *Utilization, context switches, steals*
- *Storage*
  - *Disk Utilization*
  - *I/O Queue Length*
  - *Average I/O wait*

Trove Architecture

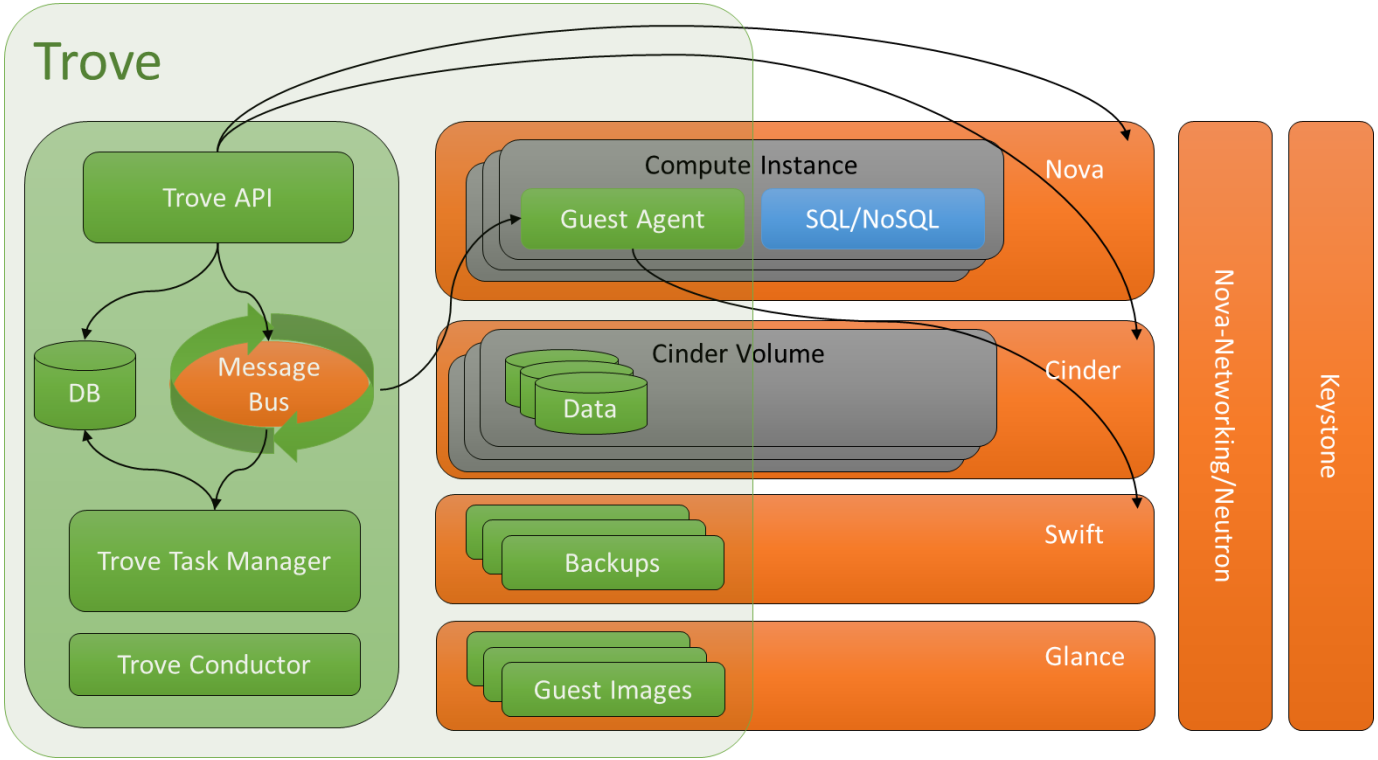

Storage Performance

- Why it matters
- Things databases do to improve storage performance

- What's the ideal storage for a database?

Contact

- Amrith Kumar
  - Founder & CTO, Tesora
  - @amrithkumar
  - IRC: amrith @freenode

- Tesora
  - www.tesora.com
  - @tesoracorp
  - Sign up for Short Stack, our informative weekly newsletter about OpenStack

What's next for DBaaS/Storage?

- Tighter integration between Cinder and Trove
- Up-level Cinder API call in Trove
- Surface more granular storage options
- Enable per-database storage customization

DO EPIC STUFF

Go forth and be a force of the awesome

Thank You