

Selecting the Right Persistent Storage Options for Apps in Containers

Bipin Kunal & Niels de Vos, Red Hat

Who are Bipin and Niels?

Bipin:

Product lead for Gluster Support and Maintenance Engineer.

Niels:

Developer and Maintainer at the Gluster community and involved in related projects.

Agenda



- What is persistent storage
- Why do we need persistent storage
- Type of persistent volumes
- Access modes for your PVs
- Performance details between different Access Modes

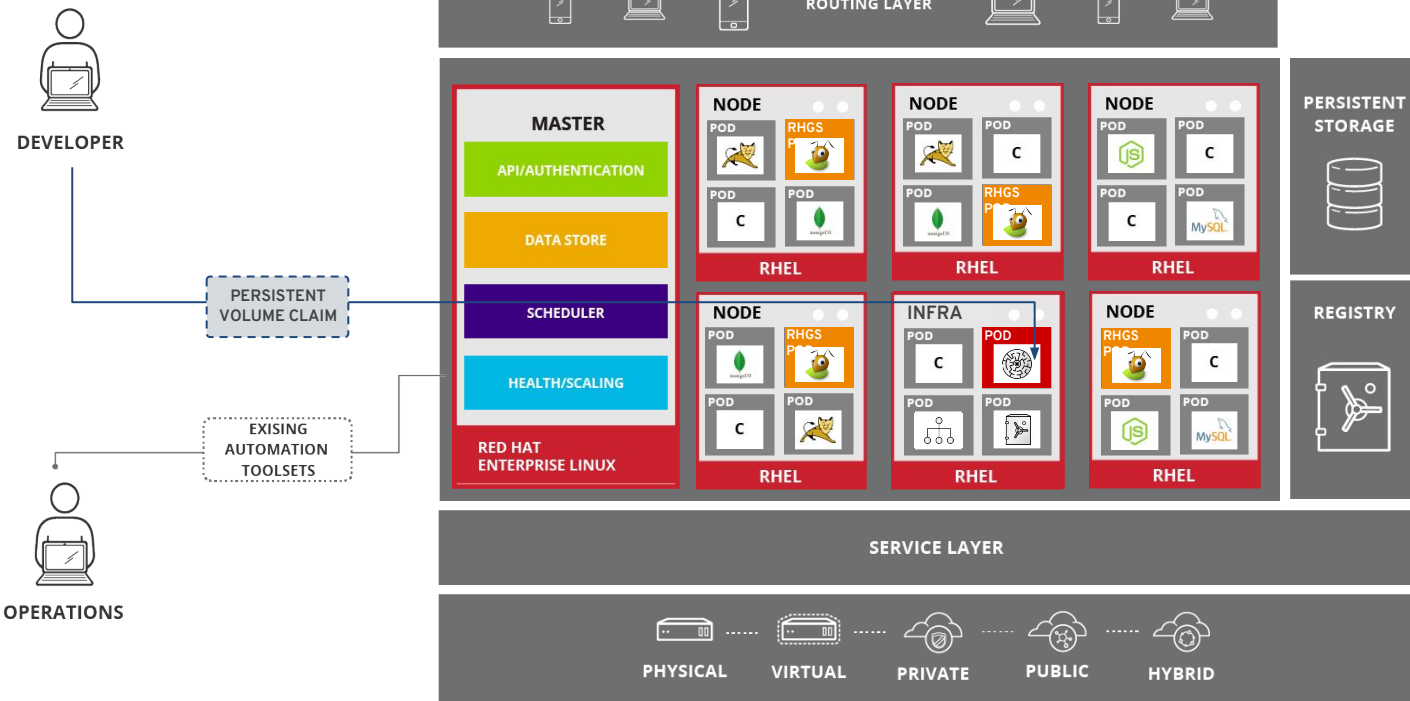
What is Persistent Storage?

- Persistent storage is any data storage device that retains data after power to that device is shut off.

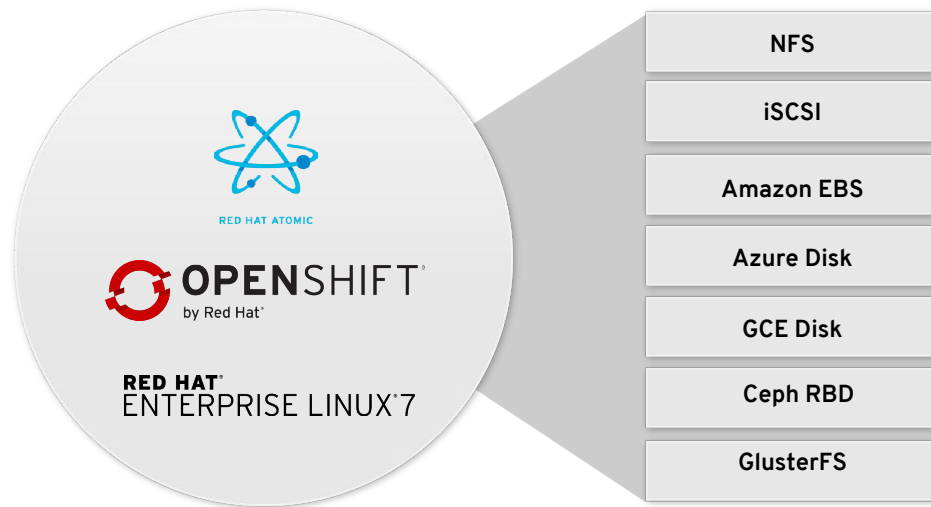
Why do we need persistent storage?

- Containers are mostly stateless, but produce or gather data for later consumption.
- Containers may not always be running or available, the data should be kept so that a next run can continue where it left off.

OpenShift: Container-Native Storage



Storage Options for OpenShift

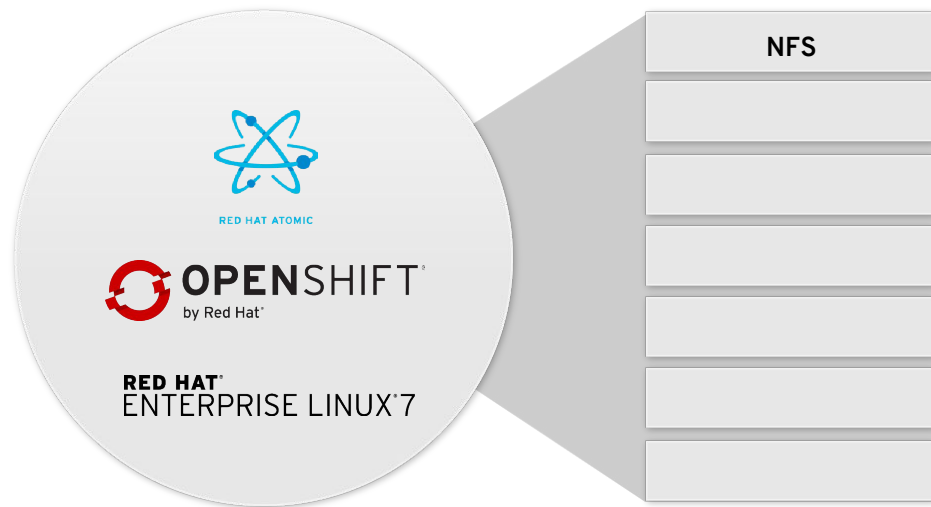


AUTOMATED CONFIGURATION

SINGLE CONTROL PANEL

CHOICE OF PERSISTENT STORAGE

Storage Options for OpenShift: NFS



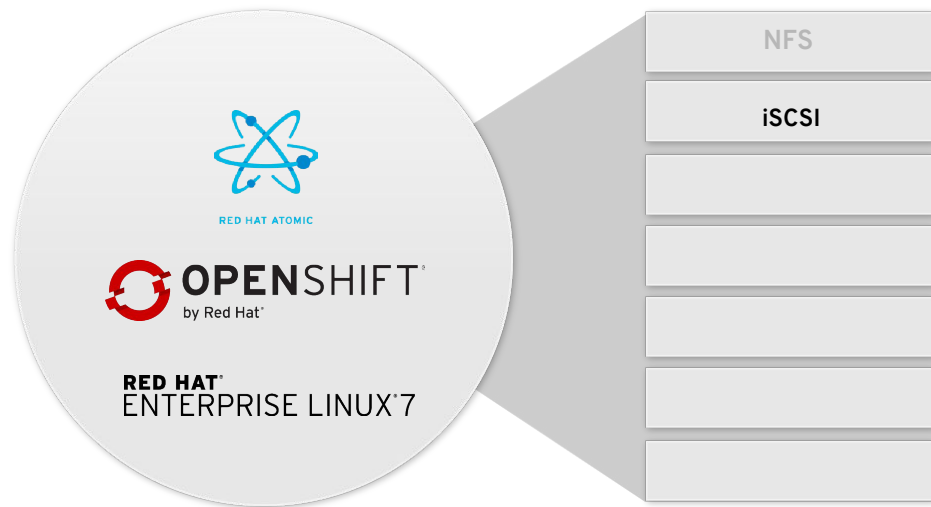
WELL UNDERSTOOD, EASY TO DEMO

PROVIDES RWX STORAGE

NO HA, NOT AVAILABLE IN THE CLOUD

REQUIRES PRE-PROVISIONING

Storage Options for OpenShift: iSCSI



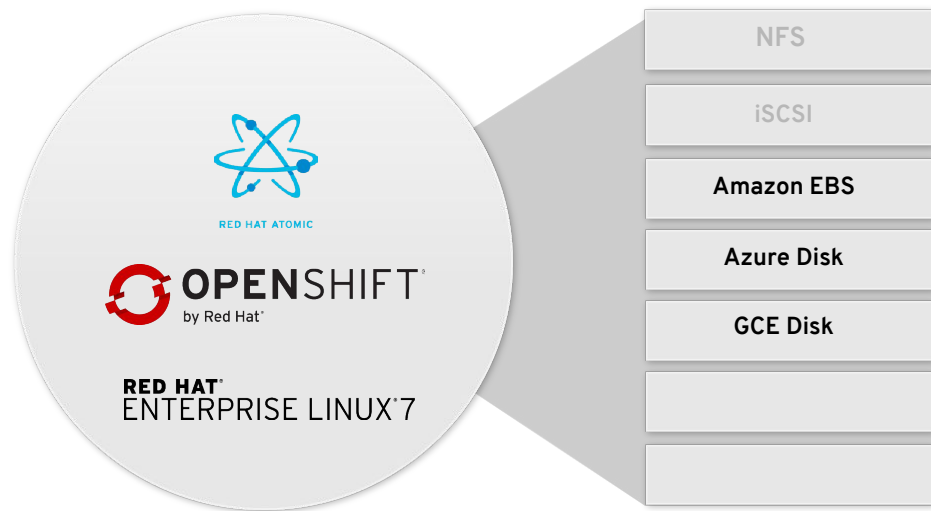
WELL UNDERSTOOD IN THE DC

BLOCK STORAGE, NO RWX SUPPORT

NOT AVAILABLE IN THE CLOUD

REQUIRES PRE-PROVISIONING

Storage Options for OpenShift: Cloud Provider Storage



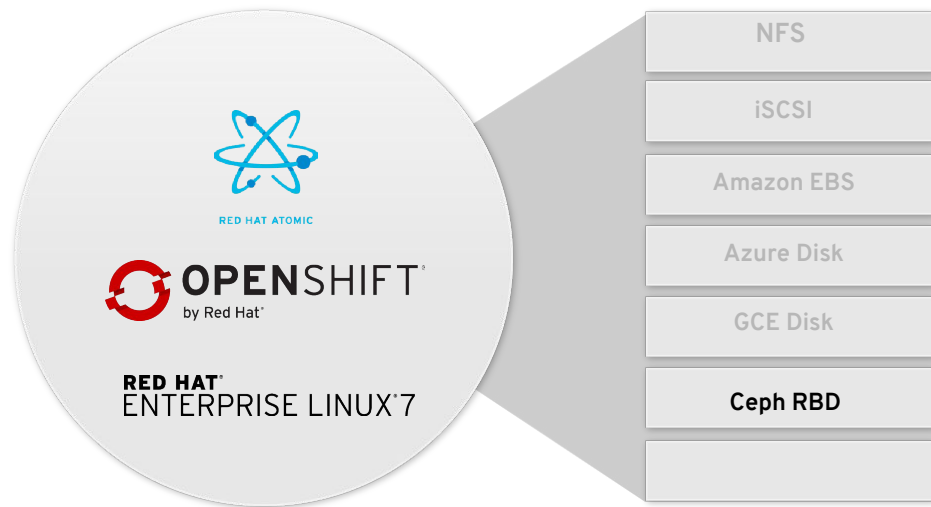
AVAILABLE IN THE CLOUD ;)

NOT AVAILABLE ACROSS AZs

PERFORMANCE DEPENDS ON SIZE

BLOCK STORAGE, NO RWX SUPPORT

Storage Options for OpenShift: Ceph

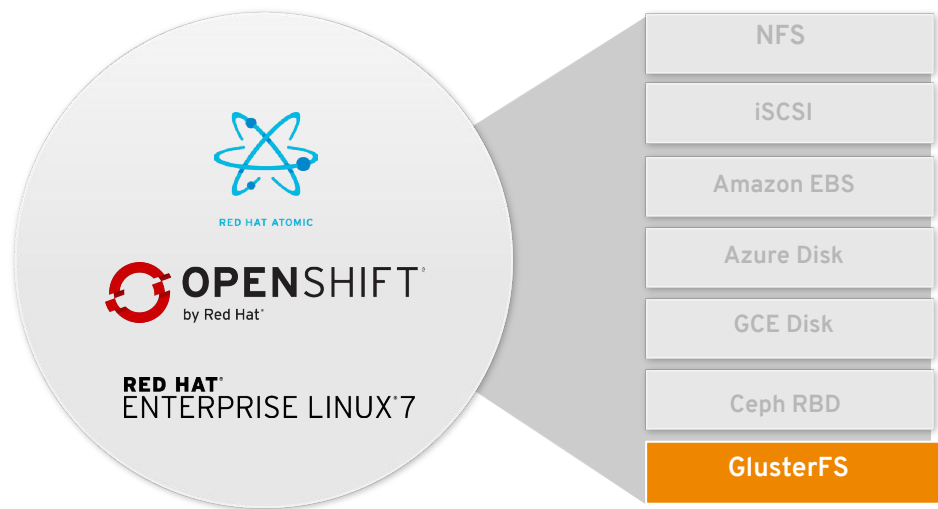


SCALABLE, HA

BLOCK STORAGE, NO RWX SUPPORT

CINDER PREFERRED FOR OCP on OSP

Storage Options for OpenShift: Gluster



RUNS ON TOP OF OPENS SHIFT

AVAILABLE EVERYWHERE

SCALABLE FILE STORAGE, RWX SUPPORT

AVAILABLE OUT OF THE BOX

GEO-REPLICATION, SNAPSHOTS

Access modes for your PV's

- ReadWriteOnce (RWO)
 - Single node is allowed to read/write
- ReadOnlyMany (ROX)
 - Multiple nodes can read at the same time
- ReadWriteMany (RWX)
 - Multiple nodes can read and write at the same time

ReadWriteOnce (RWO)

- Applications that do not need to share a filesystem among instances
- Optimizations possible, block-storage mounted in the container
- Fewer filesystem operations, only read/write data

ReadOnlyMany (ROX)

- Static content that was generated
- Snapshot of RWO/X promoted to ROX
- Read-only mounted NFS-export

ReadWriteMany (RWX)

- Dropbox like applications, Wordpress with uploading images, Jenkins Artifacts
- Performs best with limited create/write, many read
- Mounted NFS-export

Example: ReadWriteMany (RWX)

- Mounted NFS-export
- 'git clone /srv/glusterfs.git' into the mountpoint
- Repository after clone ~80 MB

50 seconds, 17133 NFSv4 operations

Example: ReadWriteOnce (RWO)

- Disk-image stored on NFS-export
- Mounted NFS-export on the host
- Mounted (loopback) disk-image on the host

< 1 second, ~1 NFSv4 operation

8 seconds, 195 NFSv4 operations (+ sync)

Thank you!



- Red Hat is still not in shipping business!
- Visit the Red Hat stand register for a workshop
- Visit for cool pictures



OPEN
EUROPE

SOURCE SUMMIT

THE LINUX FOUNDATION

