



Ovirt and Ceph : Perfect Combination ?

Thibaut Démaret – CTO
info@worteks.com





Services

Heterogeneous and complex infrastructures, cloud, mail, authentication, security

- **Studies, audit and consulting**
- **Technical expertise**
- **Technical support**
- **Training**
- **R&D**

Edition



Collaboration and application portal



Mutualized platform for development



Identity and Access Management

Partnership



READY

BUSINESS PARTNER



Spoiler Alert



Spoiler Alert





Ovirt : What is it ?



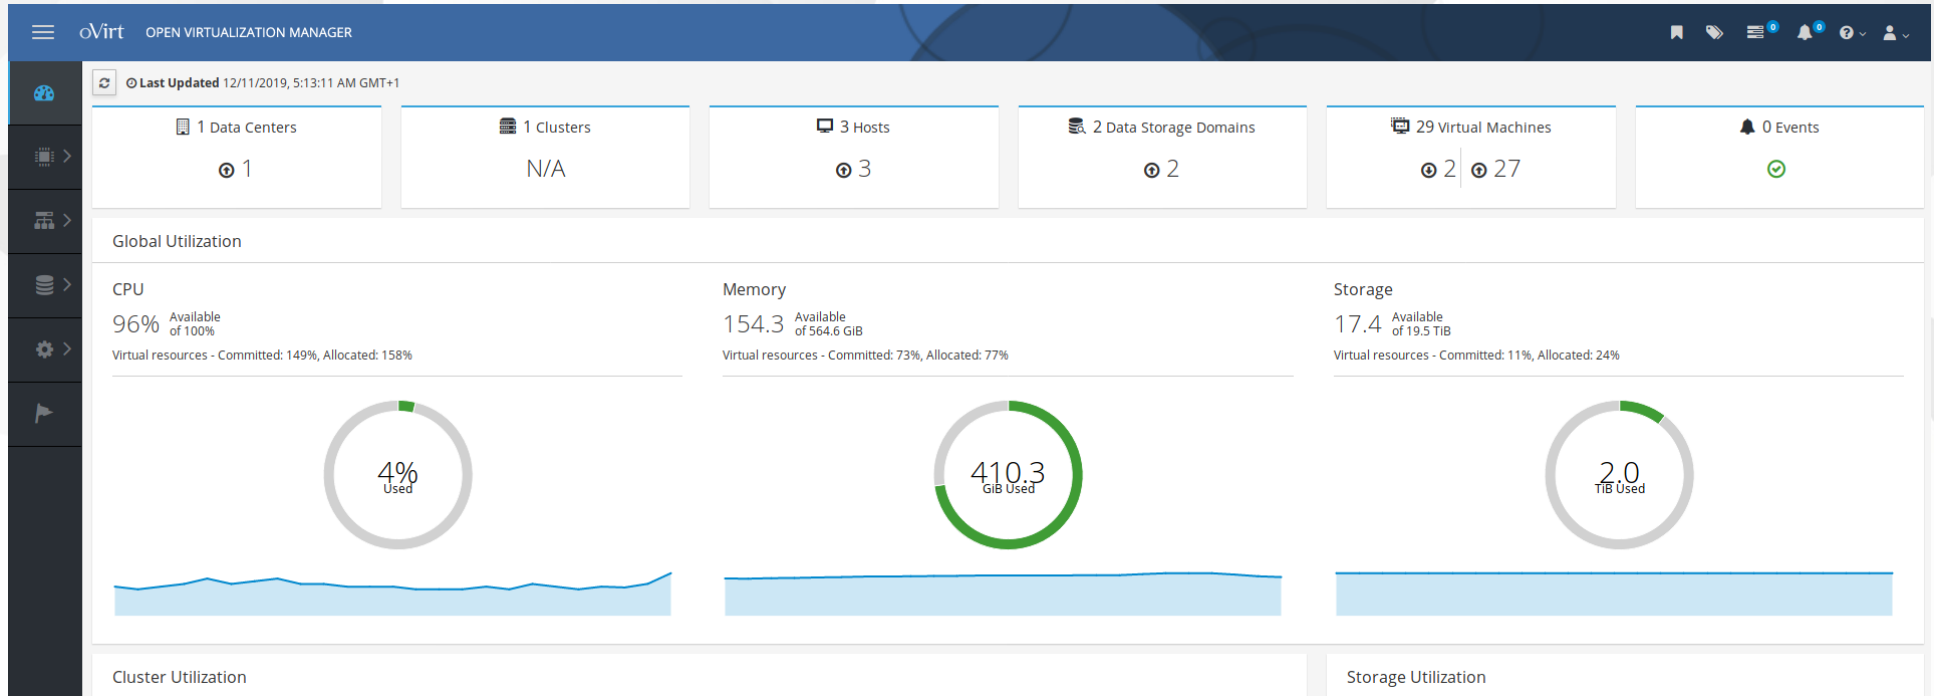
Ovirt - What is it ?

- Virtualization management platform
- Based on KVM and Libvirt
- Apache License 2.0
- Version 4.3.6
- Main Contributor : Redhat



oVirt

Ovirt - What is it ?



Ceph : What is it ?

Ceph - What is it ?

- Storage Platform
- Different usages :
 - Block Storage
 - Object Storage
 - FileSystem
- High Scalability
- No SPOF
- Version Mimic



Ceph - Options



- RBD : Block Storage
- RADOS : Object Storage
- CEPHFS : FileSystem
- ISCSI Gateway : ISCSI with additional nodes
- NFS Gateway : NFS with additional nodes





Ok, why this presentation ?



Ovirt – Default Storage Options



- You can use :
 - NFS
 - ISCSI
 - GlusterFS
 - POSIX Compliant
- Hyperconverged Mode :
 - Based on glusterFS



Ovirt – Why not use default options ?



- NFS :
 - Really Simple

- But :
 - Huge SPOF
 - Don't Use NFS 4.1 : pNFS



Ovirt – Why not use default options ?



- ISCSI :
 - High Availability
 - Multi Path
- But :
 - No SAN



Ovirt – Default Storage Options



- GlusterFS :
 - High Availability
 - Block Storage
 - FileSystem
- But :
 - Configuration Sync
 - Limit :
 - Max Host
 - Max Volume



Ovirt – Default Storage Options



- Hyperconverged Mode :
 - Based on glusterFS ...
 - Max 12 Hosts



Ovirt – Default Storage Options



- Best Option :
 - POSIX Compliant ??
- After the deployment :
 - Use external Cinder :
 - Use Cinder in Openstack
 - Use Cinder Standalone





High Level Design



Ovirt Cluster Size



- Mode : Hosted Engine
- Three Hosts Ovirt Minimum



- Maximum of Hosts :

Option : 1



- Use GlusterFS
 - Ovirt Hosted Engine
 - Ceph MON
 - Cinder Standalone
- Use Cinder with backend Ceph RBD
 - For all other VMs
- Number of Host : 11
- Number of SPOF : 1



Option : 2



- Use GlusterFS
 - Ovirt Hosted Engine
 - Ceph MON
- Use POSIX Compliant with CephFS
 - For all other Vms
- Number of Host : 10
- Number of SPOF : 0



Option : 3



- Use Ceph ISCSI
 - Ovirt Hosted Engine
- Use CephFS
 - For all other VMs
- Number of Host : 12
- Number of SPOF : 0



My opinion



- Option 2
- Hyperconverged :
 - Ovirt
 - GlusterFS
 - Ceph OS
- GlusterFS use :
 - Hosted Engine
 - 3 Ceph MON
- Number of Host : 7
- Number of SPOF : 0





Specification



Specification



- SNAPSHOT :
 - Don't use Storage native snapshot
- CephFS preparation :
 - Chown id 36
 - SELinux customization



Mount Options



Data Center	Default (V4) ▾	Name	cephfs
Domain Function	Data ▾	Description	
Storage Type	POSIX compliant FS ▾	Comment	
Host to Use	t[REDACTED]n ▾		

Path	[REDACTED]:/
VFS Type	ceph
Mount Options	rw,relatime,acl,name=admin,secret=[REDACTED]

Advanced Parameters

Warning Low Space Indicator (%)	10
(1532 GB)	
Critical Space Action Blocker (GB)	5
Warning Low Confirmed Space Indicator (%)	10
Format	V4 ▾

Wipe After Delete

Backup





After 2 years



Update



- Ovirt :
 - Minor update without Downtime
 - Update 4.2 to 4.3 : no issue
- Ceph :
 - Minor update without Downtime
 - Update Mimic to Nautilus : In progress



Incident



- Ovirt :
 - Memory Issue the first month
- Ceph :
 - Nothing special



Why Ovirt / Ceph ?



Why Ovirt / Ceph ?



- OpenShift Compatibility :
 - Ovirt Openshift Integration
 - Ceph RBD and CephFS integration





Questions ?





Thanks for your attention

More informations:



info@worteks.com



[@worteks_com](https://twitter.com/worteks_com)



[linkedin.com/company/worteks](https://www.linkedin.com/company/worteks)

