UNIFIED INFRASTRUCTURE WITH RED HAT GLUSTER STORAGE

Marcel Hergaarden

Sr. Solution Architect

Red Hat Benelux

April 2017

Key Takeaways

- Red Hat Gluster Storage can be hyperconverged with the Red Hat product stack
- Benefits include infrastructure consolidation and operational efficiency
- Self healing and highly available architecture for Remote Office-Branch Office (ROBO) scenarios

The Need for Hyperconverged Infrastructure

Remote Office Branch Office (ROBO) use cases

What are the chief challenges?

- Compute and storage needs increasing, limited branch IT staff
- Space, cooling, talent (storage in particular) at a premium

What do they do today?

- Traditional compute + NAS or RAID array (expensive)
- Tape backup (error-prone)

What are the remedy options?

- Do nothing - continue to struggle with ever expanding gear, operating expenses & downtime
- Farm out - eat costs/delays of external IT management
- Deploy Hyperconverged Infrastructure

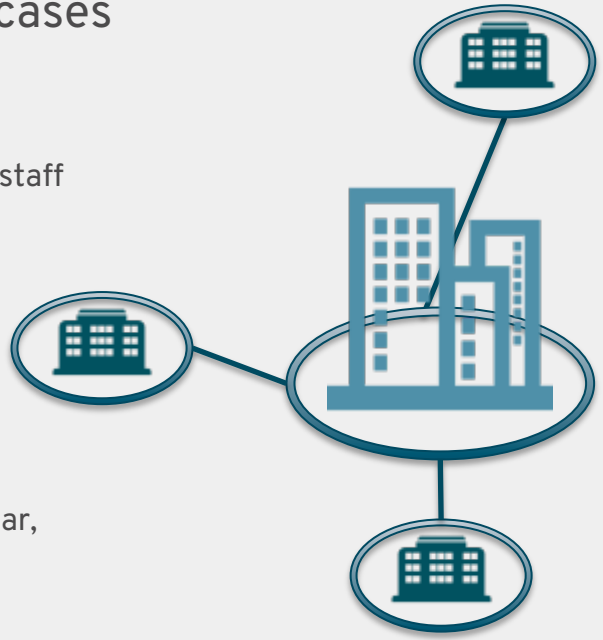

Red Hat Unified Infrastructure

Future ready with full integration with Red Hat stack

- Small footprint SDI for Remote Office Facilities
- Red Hat HCI Powered by:
  - Red Hat Virtualization SDI platform
  - Red Hat Gluster Storage
  - Red Hat's rich partner ecosystem
- Operational Efficiencies
  - Elimination of discrete storage tier
  - Reduced TCO

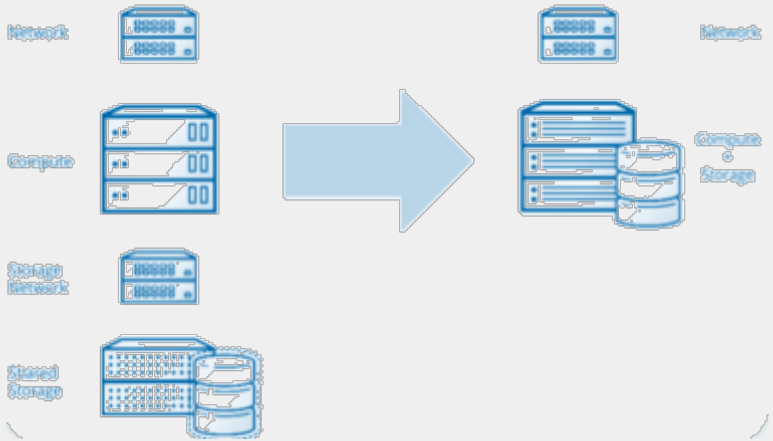

Architecture

Self healing and highly available

- 3-node base POD configuration
- RHGS with 3-way replica
- RHV-H with Self-Hosted Engine
- SSD LVMcache fronting spinning media
- Hardware and software monitoring and fault correlation
- Software solution with reference architectures
- Ansible based deployment tool & playbook

ANSIBLE
by Red Hat

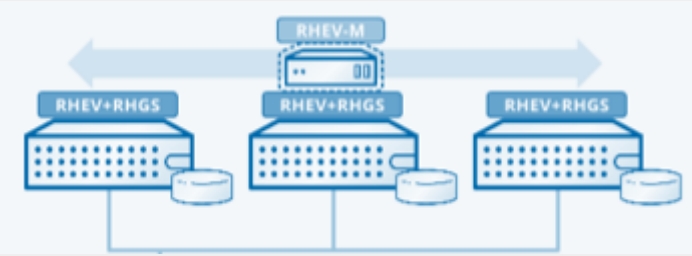

RED HAT VIRTUALIZATION

RED HAT GLUSTER STORAGE

Benefits of Red Hat Unified Infrastructure

Infrastructure Consolidation

- Eliminate storage as a discrete tier
- Virtualize business applications, maximizing resource utilization
- Single budget for compute & storage

Operational Efficiency

- Single team managing infrastructure
- Simplified planning and procurement
- Streamlined deployment & management

Operational Efficiency

Before and After Unified Infrastructure deployment

BEFORE

Ten steps across three teams

1. Purchase servers
2. Purchase storage (SAN and array)
3. Configure servers on network
4. Configure storage network (SAN)
5. Install management platform - storage
6. Configure array - pool(s), LUNs, masking
7. Install hypervisors
8. Configure storage (for VM's)
9. Install management platforms (virt)
10. Configure alerting servers, network, storage

AFTER

Six steps completed by a single team

1. Purchase servers (POD with single SKU)
2. Configure servers on network
3. Install hypervisors
4. Configure storage (for VM's)
5. Install management platform
6. Configure alerting servers and network

But What About Disaster Recovery?

Geo Replication

- Integrates geo-replication process with native RHV API
- Periodic synchronization based on checkpoints
- Target volume is *sharded*
- Shard size at the source & target may be different
- Crash-consistent DR images
- 1:1 source:target relationship

Learn More

redhat.

redhatstorage.redhat.com

red.ht/GlusterTestDrive

plus.google.com/+RedHatStorage

facebook.com/redhatstorage

linkedin.com/company/red-hat

twitter.com/RedHatStorage

youtube.com/user/RedHatStorage