SDC
STORAGE DEVELOPER CONFERENCE
SNIA ■ SANTA CLARA, 2014

Software Defined Storage:
Changing the Rules for Storage Architects

Ric Wheeler
Red Hat

Overview

- What is Software Defined Storage?
- What Changes for Storage Architects?
- Why Open Source Matters for SDS
- Impact on Users

What is Software Defined Storage?

What is Software Defined Storage?

- Another silly marketing term?
  - What storage does not have software in it?
  - Technology?
  - Deployment model?
- One of the more popular platforms for “Software Defined Storage”
  - Separates the hardware from the software
  - Significantly lower costs

History of SDS

- Fairly new term
  - Goes back a couple of years
  - ScaleIO & EMC ViPR both use this term
- Focus of this talk is on Linux based SDS
  - Gluster's scale out (primarily) file storage
  - Ceph scale out storage (block, object and file)
  - Linux based appliances with samba, NFS, iSCSI target, etc

A Variety of SDS Systems

- EMC Centera is a similar system
  - Linux base
  - Proprietary user space file system
  - Vendor specified & selected hardware system
- EMC ViPR is an SDS system that runs on Linux platforms
- Lots of non-open source code in all systems
  - Firmware in components like drives, HBA's, etc

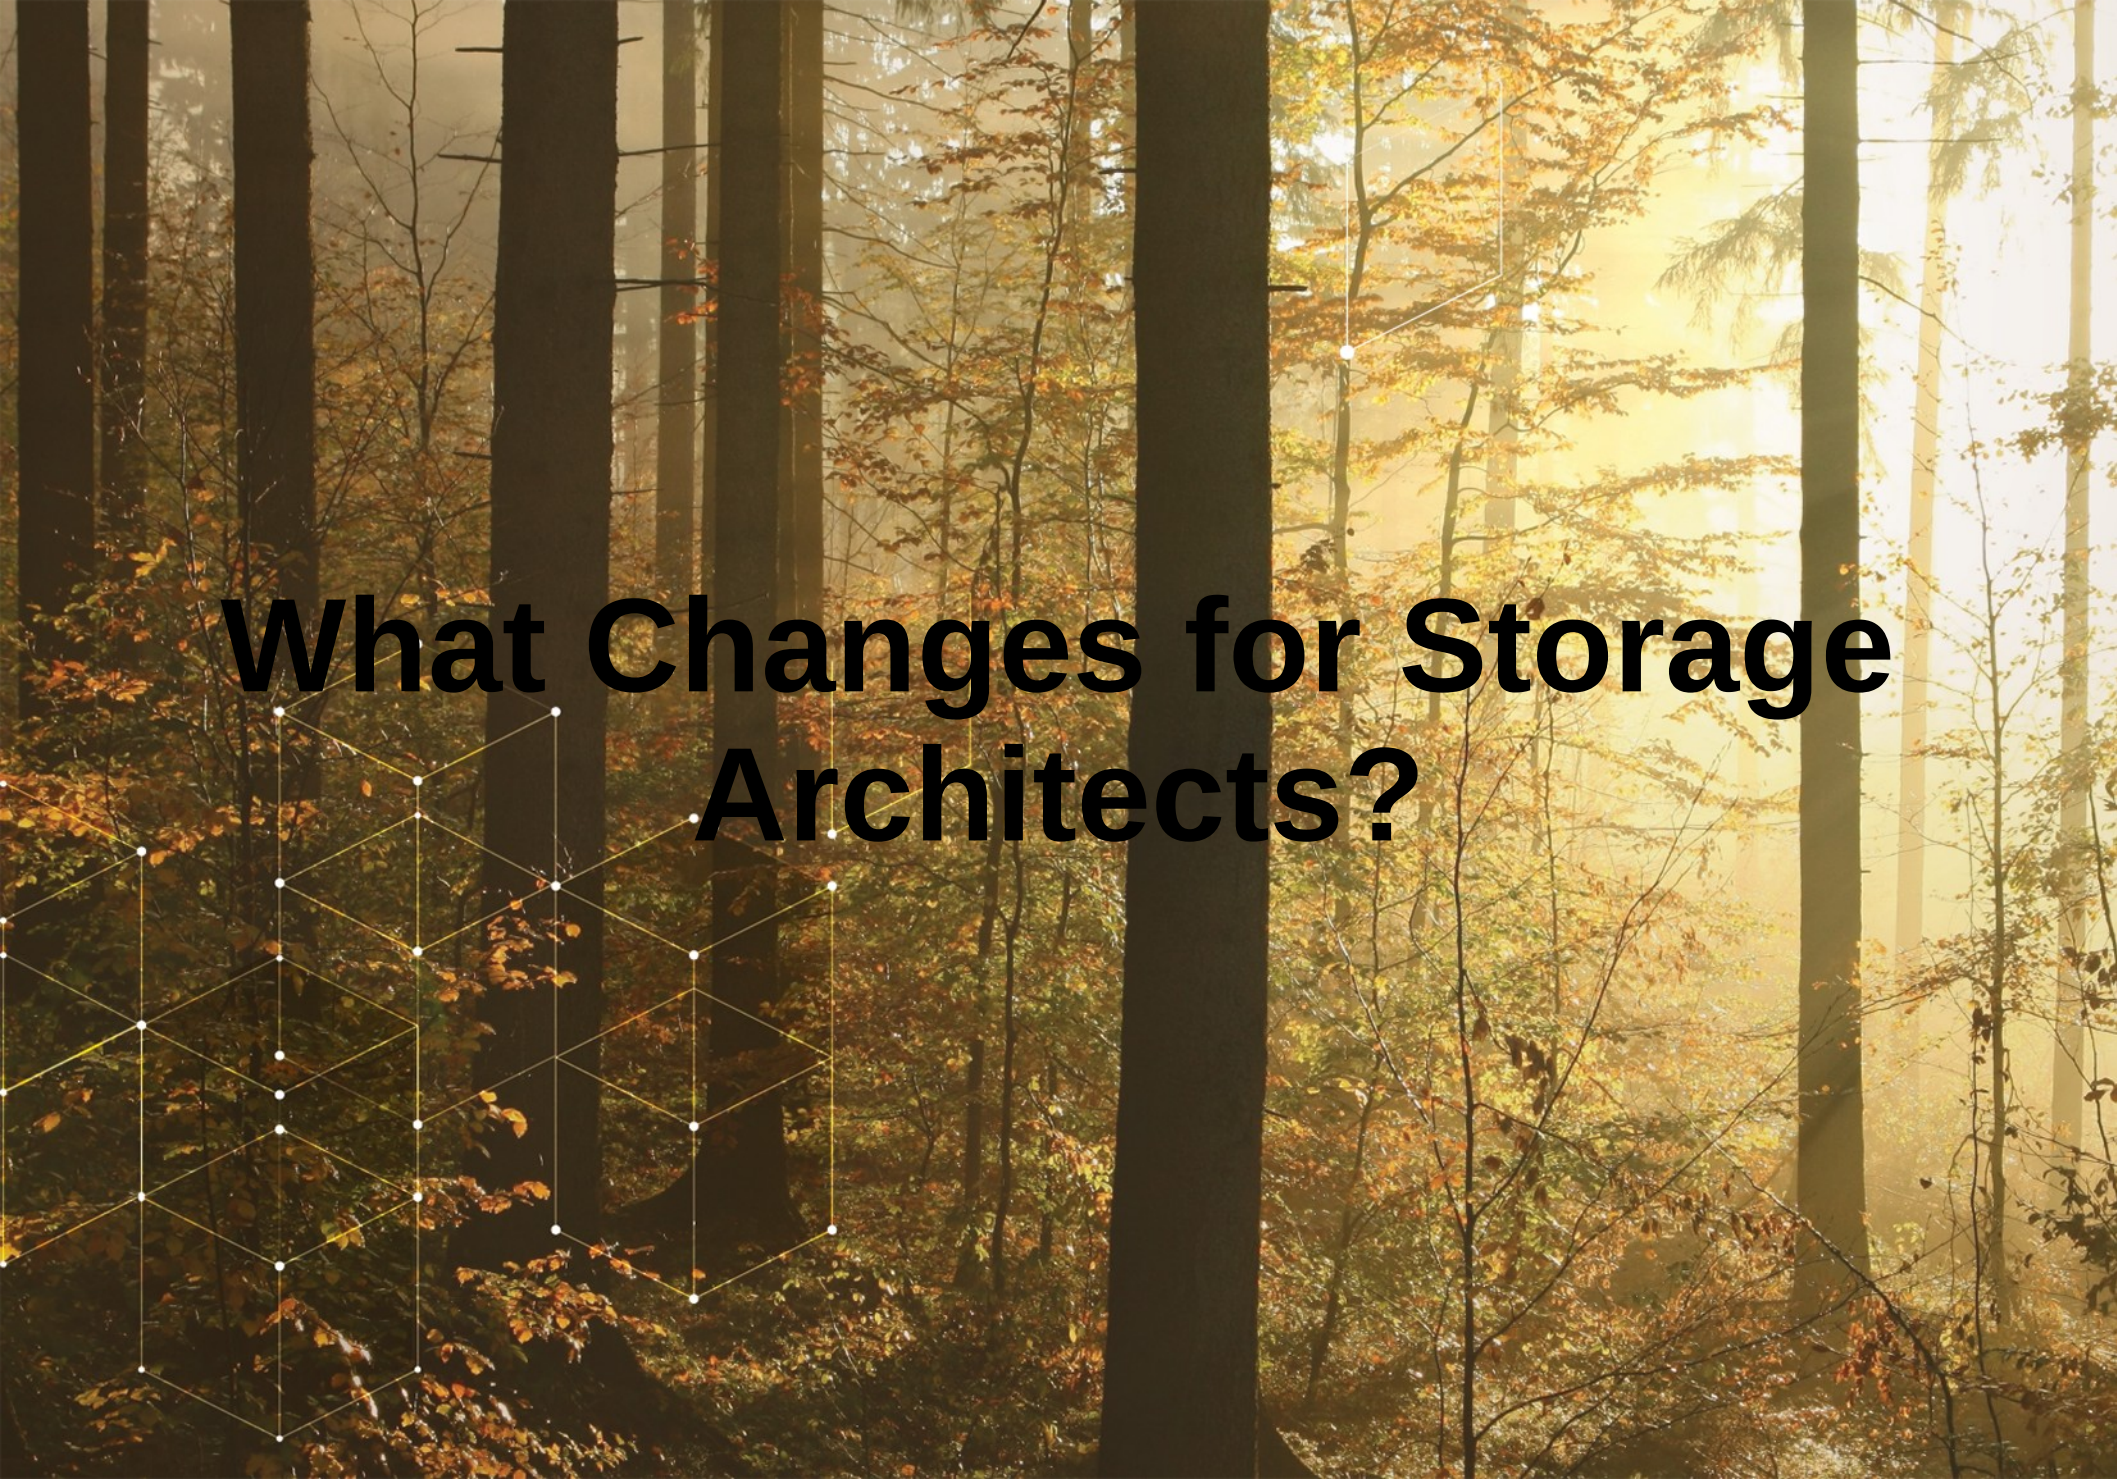
What Changes for Storage Architects?

Where to Draw the Software/Hardware Line?

- Users can supply only the hardware
  - SDS architect controls operating system and above
  - Maximize the architect's control over the platform
  - Can insure better stability and performance
- Users supply the hardware and their own OS
  - Maybe just certify some versions of some OS'es?
  - SDS is just another application
  - Maximum deployment flexibility, minimum control for architects

Some Things Get Harder

- Storage architects lose control over supported hardware configurations
  - Users select drives, HBA's and server platforms
  - Users often do direct maintenance & software updates
- Servers can run other services
  - Co-locate virtual client on server with server software
  - Can cause irregular performance
- Management applications need to help users carry the extra burden

Some Things Get Easier

- Architects get to focus on the SDS application layer itself
  - No longer need to spend our resources on drivers or operating system stack design
  - Linux supports a wide choice of high performance, hardened software components
- Storage system design is open to a broader range of engineers
  - Lots of SDS code is user space
  - Easy enough to make a very reasonable storage array without having to write kernel level code

Good Fit for Openstack

- Openstack and SDS go well together
  - Lower cost is a goal for both
  - Deployment model both favor commodity components
  - “Agile” dev ops friendly management
- Performance and reliability goals not as critical
  - Reliability achieved by scale out design
  - Some environments have a looser consistency model

Why Open Source Matters for SDS

Linux and Open Source Ecosystem

- Enormous community of developers
  - Thousands of developers contribute to each kernel
  - Linux is a strategic platform for hardware vendors
  - Wide variety of supported hardware to choose from
- Same components and stack used widely
  - Dominance from embedded systems like Android up to super computing
  - Versions of same stack used with enterprise storage
- Allows storage designers (mostly!) to focus on the storage application itself

When Things Break

- Multiple vendors can provide technical support
  - Red Hat, SUSE, Canonical, Oracle, etc
  - Embedded Linux vendors
- Public support
  - Mailing lists like linux-scsi, linux-ext4, xfs, ...
  - IRC channels
  - Conferences and other public forums
- Still best practice to keep on staff kernel engineering
  - Code is open and available
  - Can investigate, tune and fix on your own terms

Linux Community Focus on Performance

- Multiqueue work
  - Multiqueue block layer work by Jens Axboe
  - Multiqueue SCSI work by Christoph Hellwig
    - See *Multiqueue Block* talk on Wednesday
- Support for advanced fabrics
  - See *Enhancements to the iSER iSCSI Protocol* or *NFS over 40Gbps iWARP RDMA* on Wednesday
  - Multiple talks this week
- Active work on caching, tiered storage, SMR drives and persistent memory


Impact on Users

Storage Users Become Partners

- Unlike traditional, enterprise systems, storage admins now are responsible for
  - Software updates
  - Servicing their storage servers and drives
  - Selection of the hardware itself
- Performance tuning can be painful
  - All SSD's are not normally an option
  - SDS vendor provides guidance, but ultimate choice and tuning hits the end administrator

Life Cycle Operations

- Populating new storage systems with data
  - Bring data in from traditional enterprise storage?
  - Load with new, user generated data?
- Migration off of old, end of lease hardware to new systems
  - Open source and open standards help prevent lock in of data in one system
  - Can be a slow and painful process

Red Hat's Software Defined Storage Systems

- Red Hat Storage Ceph and Gluster
  - Scale out SDS systems
- Red Hat Enterprise Linux based servers
  - Single node NFS and/or samba servers
  - Active/passive clusters with NFS and/or samba
- Red Hat Enterprise Linux as a base for third party SDS software


Questions?

Resources

- Multiple conferences each year
  - Linux Foundation Vault and LSF/MM this coming March in Boston
  - SNIA events like SDC!
  - USENIX FAST and other forums
- Mailing lists
  - Linux-scsi, xfs, linux-ext4, etc
  - IRC channels