

## **A SPDEcon Cloud Preview**

Mark Carlson, Oracle Michael T Elliott, Dell Gene Nagle, Nirvanix Sage Weil, InkTank





- Storage Plumbing and Data Engineering Conference
  - By experts, for experts (and those who want to be...)
  - June 10-12, 2013, Hyatt Regency Santa Clara, Santa Clara, CA
- Register at http://snia.org/spdecon



□ A Review of the Cloud Sessions at SPDEcon

- Mark Carlson, Oracle
- Michael T Elliott, Dell
- Gene Nagle, Nirvanix
- Sage Weil, InkTank



# Interoperable Cloud Storage with the CDMI Standard

Mark Carlson, Principal Member of Technical Staff, Oracle

The Cloud Data Management Interface (CDMI) is an industry standard with ISO ratification. There is now an open source reference implementation available from SNIA as well. Storage vendors and Cloud providers have started announcing their implementations of the CDMI standard, demonstrating the reality of interoperable cloud storage. This talk will help you understand how to keep from getting locked into any given vendor by using the standard. Real world examples will help you understand how to apply this to your own situation.



#### Hybrid Clouds in the Data Center - The End State

Michael Elliott, Enterprise Cloud Evangelist, Dell

- In this presentation, I will define how to build clouds in a heterogeneous, open, and secure environment to take advantage of the benefits that hybrid cloud provides. I will cover the concepts of the cloud and detail:
  - Data Protection and Archive In the Cloud
  - Components of building a Private Cloud
  - Integration of Private and Public to form Hybrid Clouds
- The presentation will include case studies to highlight how fortune 500 and global companies are utilizing cloud infrastructure to gain agility and efficiency in their datacenter

## What are your needs?

Storage Plumbing and Data Engineering Conference

# SP⊒DEcon



#### **CXO** and **Execs**

- Reduce or eliminate IT complexities
- Reduce costs
- Better agility
- To gain a competitive
- advantage



#### **App developers**

- Quickly and easily obtain resources
- Effectively test, develop or run our applications



#### **IT** operations

- Provide deep cloud functionality to the organization
- Limit 'rogue' cloud use and spending outside our domain

# **Cloud Imperatives**

Storage Plumbing and Data Engineering Conference



- I. Cloud is here (and Hybrid is the end state)
- 2. IT must lead the discussion on Cloud
- 3. IT culture must change from one that controls physical or managed assets to one that delivers agile services that aligns with the business strategy

## **Adopting Cloud - Case Studies**

Storage Plumbing and Data Engineering Conference

# SP**⇒**DE con





### Combining SNIA Cloud, Tape and Container Format Technologies for the Long Term Retention of Big Data

Gene Nagle, Senior Systems Engineer, Nirvanix

Generating and collecting very large data sets is becoming a necessity in many domains that also need to keep that data for long periods. While this presents significant opportunities, a key challenge is providing economically scalable storage systems to efficiently store and preserve the data, as well as to enable search, access, and analytics on that data in the far future. Both cloud and tape technologies are viable alternatives for storage of big data and SNIA supports their standardization. This tutorial will present advantages and challenges in long term retention of big data, as well as initial work on how to combine SIRF with LTFS and SIRF with CDMI to address some of those challenges.

# Outline

Storage Plumbing and Data Engineering Conference



SNIA Long Term Retention technology

- Self-contained Information Retention Format (SIRF)
- Other SNIA technologies in development
  - Cloud Data Management Interface (CDMI)
  - Linear Tape File System (LTFS)
- Combining SNIA technologies
  - SIRF Serialization for CDMI
  - SIRF Serialization for LTFS
- EU Enabling kNowledge, Sustainability, Usability and Recovery for Economic Value (ENSURE)

# The Need for Digital Preservation of Big Data

Storage Plumbing and Data Engineering Conference

SP<u>⇒</u>DE con

#### Regulatory compliance and legal issues

- Sarbanes-Oxley, HIPAA, FRCP, intellectual property litigation
- Emerging web services and applications
  - Email, photo sharing, web site archives, social networks, blogs
- Many other fixed-content repositories
  - Scientific data, intelligence, libraries, movies, music
- Domains that have Big Data require preservation





### Ceph: A Unified Distributed Storage System

Sage Weil, Founder & CTO, Inktank

Ceph is a massively scalable, open source, distributed storage system. It is comprised of an object store, block store, and a POSIX- compatible distributed file system. The platform is capable of auto-scaling to the exabyte level and beyond, it runs on commodity hardware, is self-healing and selfmanaging, and has no single point of failure. Ceph is in the Linux Kernel and is integrated with the OpenStack and CloudStack cloud platforms. This talk will provide an intro into the Ceph architecture and how it unifies storage for the cloud.

# **Ceph Storage Cluster**

SP⊒DEcon



# **The Ceph Ecosystem**

SP⊒DEcon



SP=DECON

# **The Ceph Value Proposition**





#### Q&A

### Attending SPDEcon

- June 10-12, 2013, Hyatt Regency Santa Clara, Santa Clara, CA
- Register at http://snia.org/spdecon