David Minor

Chronopolis Program Manager
Director, Digital Preservation Initiatives
UCSD Library
San Diego Supercomputer Center

- SDSC Cloud now in production
- UCSD Library DAMS use of Cloud
- DuraCloud + SDSC Cloud
- DuraCloud + Chronopolis

Digital Preservation Network 

SDSC Cloud Storage

- Built on OpenStack ‘swift’ object storage software. Supports Rackspace/ Swift API and a subset of Amazon’s S3.
- Initial 5.5PB Raw Storage, scalable to over 100PB with equal performance scaling.
- Supported by redundant Arista Networks 7508 switches, providing 768 total 10 gigabit (Gb) Ethernet ports for more than 10Tbit/s of non-blocking, IP-based connectivity.
- Peak transfer rates of up to 8GB/sec.
- Direct connections to CENIC, ESNet, and XSEDE networks provides high-bandwidth wide area connectivity.
- Managed with Rocks cluster toolkit. Swift ’Roll’ available soon.
- AES 256-bit server side encryption for at rest data is available and deployed in a local HIPAA environment.

http://cloud.sdsc.edu

Service Architecture

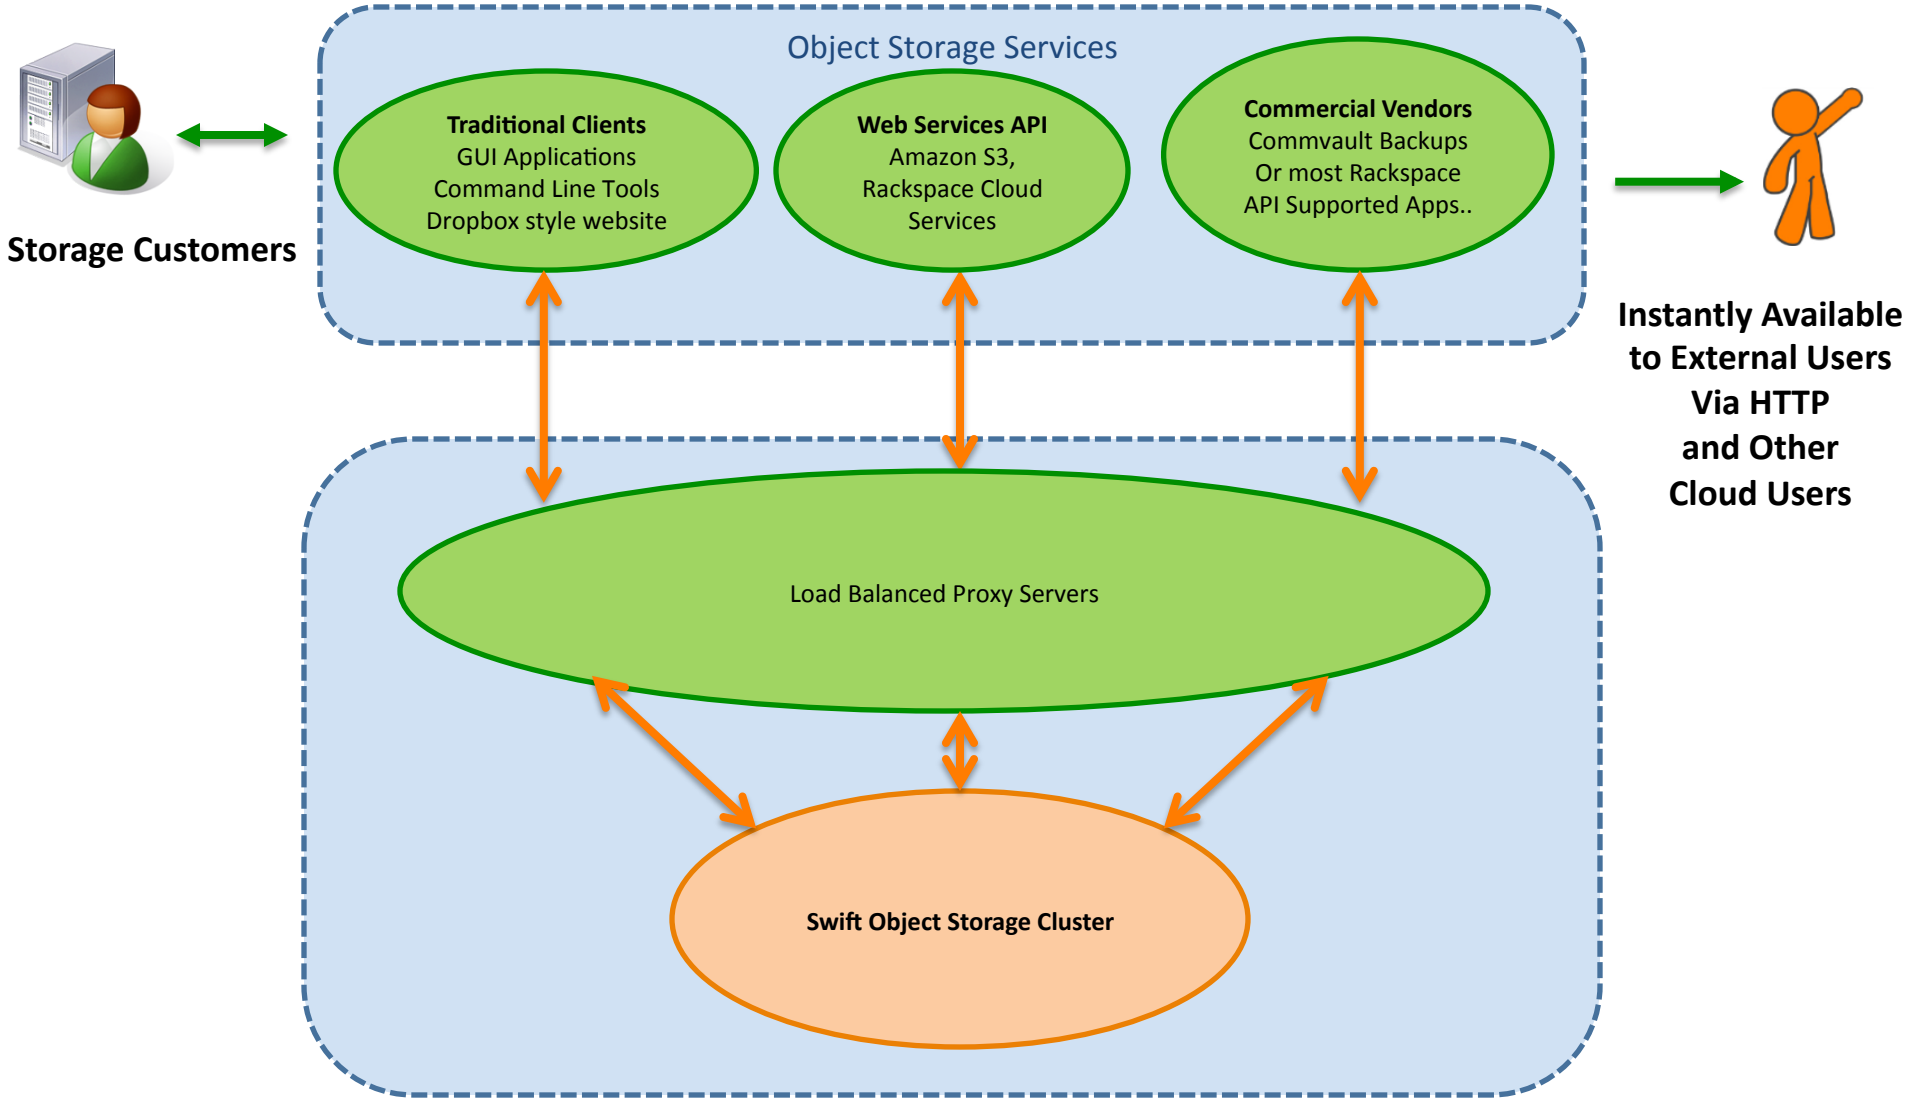

Goals for SDSC Cloud

http://cloud.sdsc.edu

- Support NSF Data Management Plans
  - Required Plan to describe how research results are shared.
- 99.5% system availability
  - File replication automated
  - Default 2 copies, able to keep additional offsite replications.
  - Automated checksum verification and error correction
- Scalable
  - Performance and capacity grows by incremental bricks.
- Multifaceted accessibility
  - Web, API, Graphical and Command Line Clients
- Cost competitive
  - Operated as a recharge service
  - On par with current tape-based dual-copy costs of $0.0325/GB/Mo.

UCSD Library DAMS use of Cloud

(slides provided by Declan Fleming)

DAMS Storage Story

- Started with 4T on SRB 8 years ago
- Eventually moved everything to local SAMBA
  - Grown to 16T now
- UCSD research data pilots required 100T+

DAMS Implementation

- RackSpace CloudFiles Java Library
- RESTian
  - 13% faster for small files
  - 17% faster for large files
- SAMBA vs. cloud upload and download speeds quite similar
- Colocated server at SDSC produced download speeds sometimes 50% faster with OpenStack

Large Files with OpenStack

- Files larger than 10GB revealed a bug
- OpenStack requires breaking files larger than 5GB into segments
  - returns those segments in lexical sort order, rather than in creation order
    - So if the segments are named 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, they will be returned in the order 1, 10, 2, 3, 4, 5, 6, 7, 8, 9
- Padding the numbers "0001" to make sure they sort in numerical order fixed this problem
- Once dealt with, performance was independent of file size

OpenStack Storage Organization

- Containers vs Subdirectories
- How many containers should we use?
  - One per object?
  - One container for all objects?
- Tested with 10000 TIFFS, 185G, from 1 to 10000 containers
- Sweet spot was around 100 containers
- Learned later that every container is a SQLite instance
- Also learned that some other services will limit the number of containers per user

Victory declared. Now starting process of moving dozens of terabytes of research data into cloud-hosted DAMS.

Declan Fleming - dfleming@ucsd.edu

DuraCloud and SDSC Cloud

DURACLOUD™

SDSC Cloud Storage Services

DuraCloud + SDSC Cloud

- SDSC Cloud now an option in DuraCloud
- Available in similar fashion as Amazon and Rackspace
- DuraCloud customers can now have a single service agreement and interface to access commercial and academic cloud providers

DuraCloud + SDSC Cloud

"SDSC is an important strategic partner for us in that they are the first academic-based production cloud offering and we envision together building an academic-based open cloud infrastructure with a layer of services provided by DuraCloud. We see this as the first step in this direction and are thrilled."

- Michele Kimpton, DuraSpace CEO

DuraCloud + Chronopolis integration

Chronopolis + DuraCloud

- Chronopolis will be a preservation option within DuraCloud offerings
- Rolling into production this fall
- On the backend, enabled through SDSC Cloud and related services

Process

- Data movement between systems happens in SDSC cloud
- Mediated by set of RESTful services
- Data handed off from DuraCloud with manifest, checksums, etc.

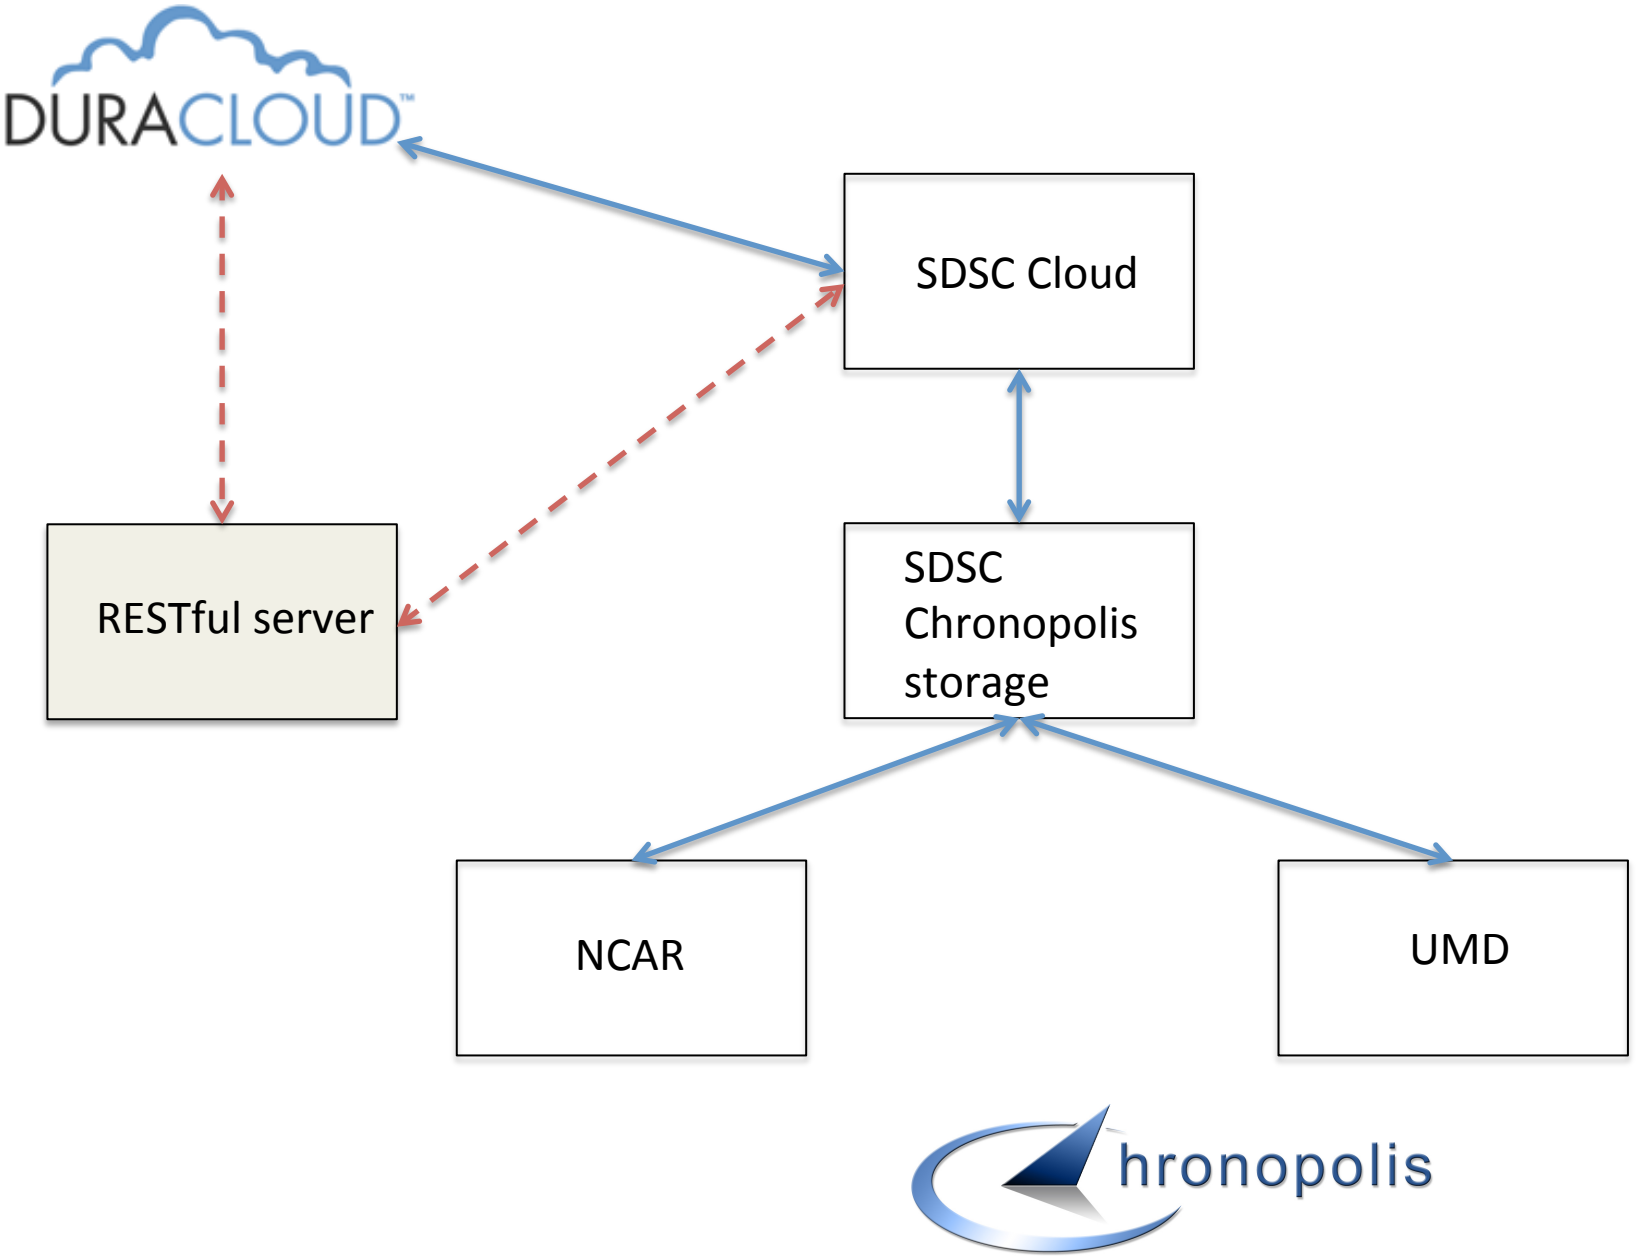
DURACLOUD
SDSC Cloud
RESTful server
SDSC
Chronopolis
storage
NCAR
UMD
hronopolis

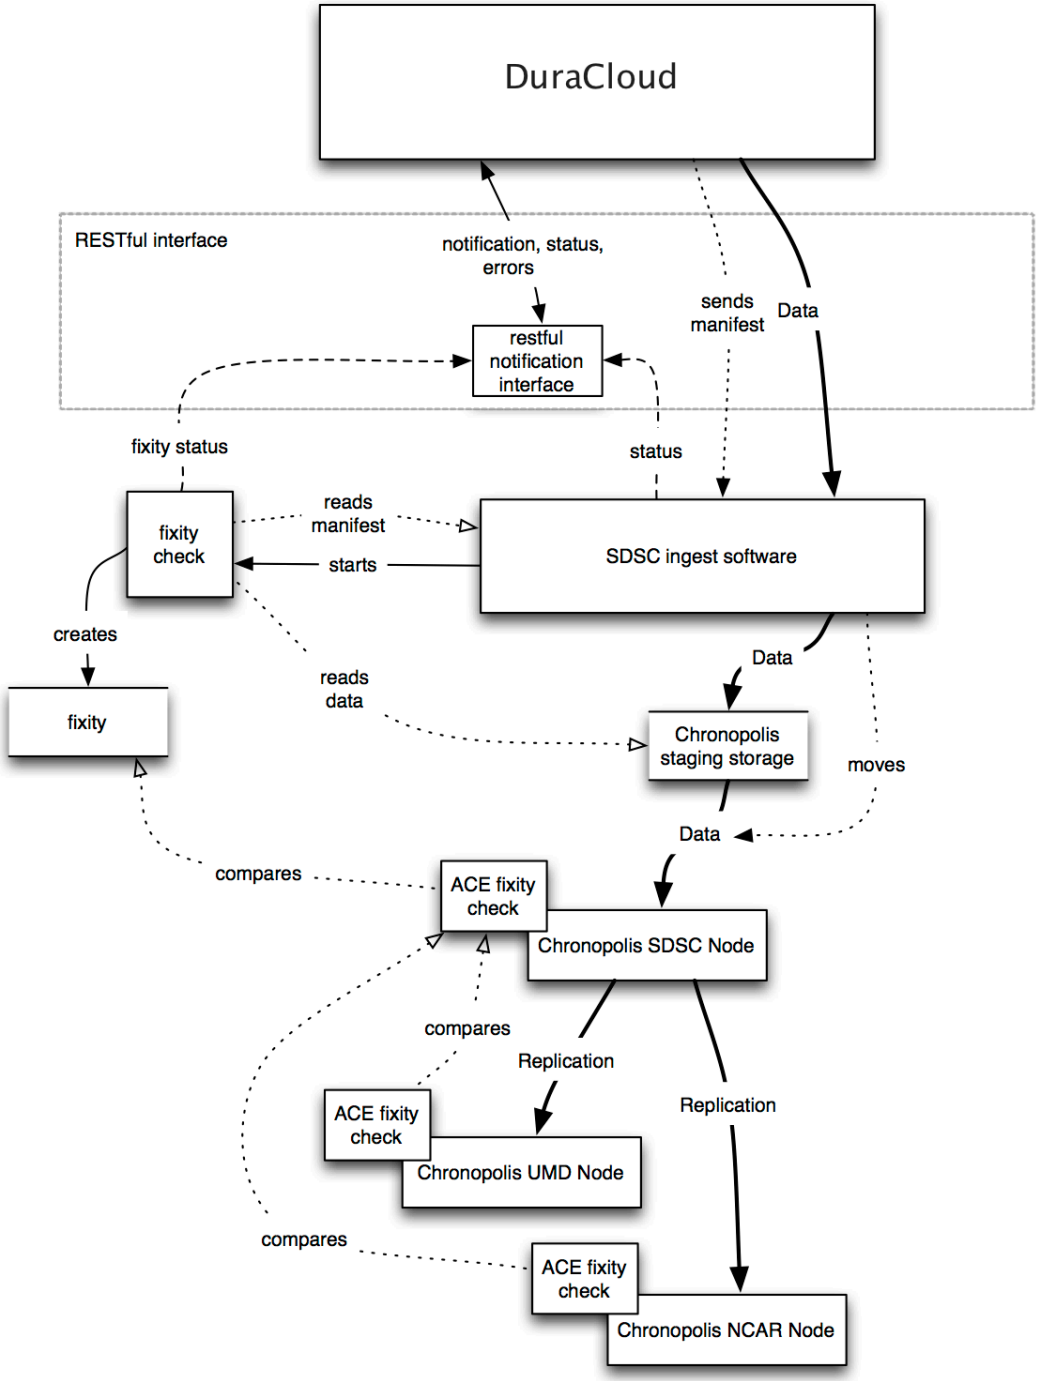
DuraCloud
RESTful interface
notification, status, errors
sends manifest
Data
restful notification interface
fixity status
status
fixity check
reads manifest
starts
SDSC ingest software
creates
fixity
reads data
Data
Chronopolis staging storage
moves
Data
compares
ACE fixity check
Chronopolis SDSC Node
compares
Replication
Replication
ACE fixity check
Chronopolis UMD Node
compares
ACE fixity check
Chronopolis NCAR Node

Outcome

- Successful integration of two NDIIPP-supported preservation projects
- Chronopolis offers a TRAC-certified archival backend to DuraCloud
- Chronopolis offers a low-cost snapshot of DuraCloud holdings
- DuraCloud offers a seamless option for orgs wanting data services and preservation
- Demonstration of uniting two systems via channels of communication and services

For more information…

- Andrew Woods, DuraCloud:
  awoods@duraspace.org

- David Minor, Mike Smorul, Chronopolis:
  minor@sdsc.edu
  msmorul@sesync.org

Also will be presented as a poster at iPRES