
# LIGO_CIT

**A storage and compute cluster at the California Institute of Technology**

OSG Storage Forum
06/30/09

Robert Engel
engel_r@ligo.caltech.edu
California Institute of Technology

# What is LIGO_CIT?

A small storage and compute cluster build from commodity hardware part of the Integration Testbed, which is part of the Open Science Grid.

## Compute Element

- 104 cores
- 208 GB RAM
- Condor Queue
- GRAM
- GRAM-WS

## Storage Element

- 23TB Hadoop partition
- BeStMan-Gateway
- GridFTP

# Hardware

The cluster consists of 15 Dell PowerEdge 1950 1U units:

- ✔ 2 Intel Xeon 5300 Quad Core Processors
- ✔ 16GB FBD at 667MHz
- ✔ 2X 1TB SATAII hard disks
- ✔ 2X 1Gbit/s Ethernet

connected by a GB Ethernet network in channel bond mode:

- ✔ 2 Netgear 24Port GB switches

hosted at the 'Synchrotron' (Caltech) in its own rack.

# Network

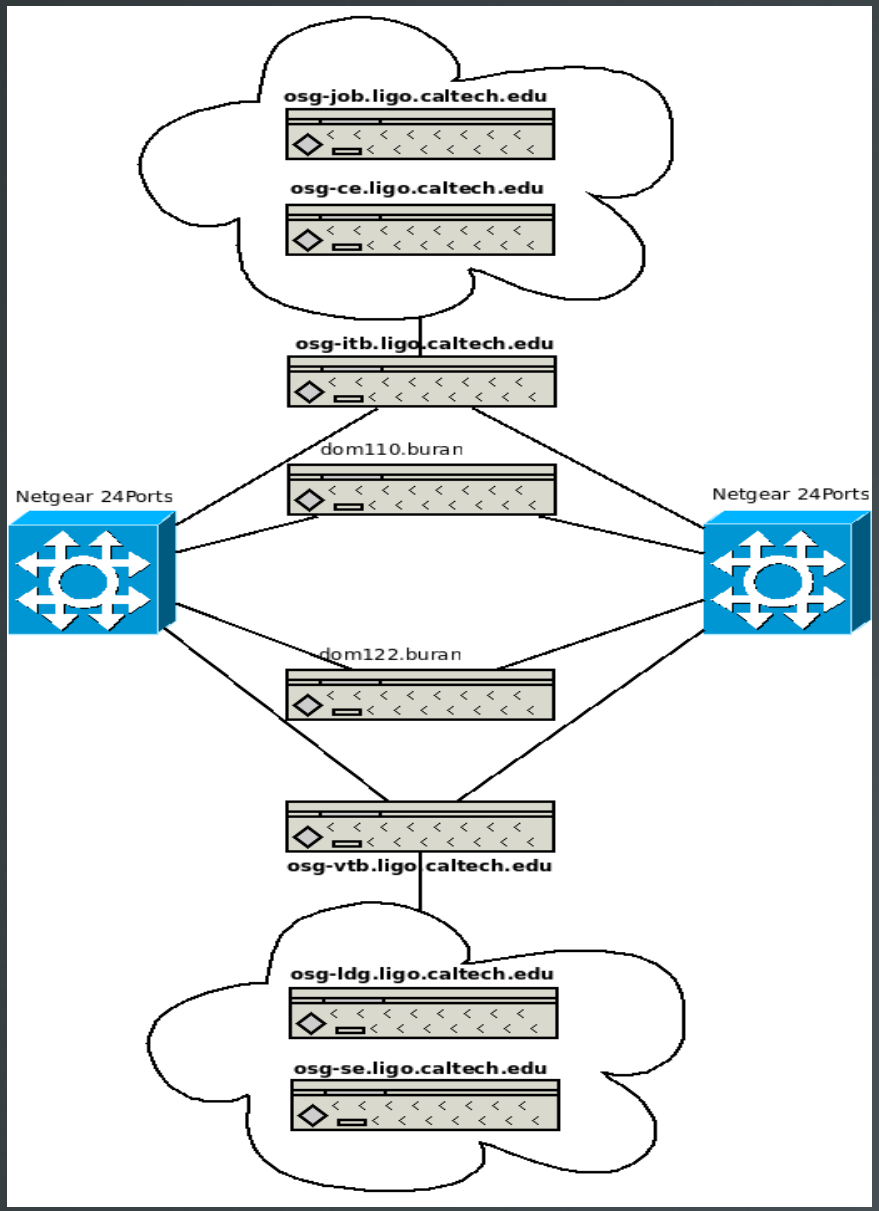


3 hardware frontends:
osg-itb.ligo.caltech.edu
osg-vtb.ligo.caltech.edu
osg-itb-se.ligo.caltech.edu*

4 virtualized frontends (xen):
osg-job.ligo.caltech.edu
osg-ce.ligo.caltech.edu
osg-ldg.ligo.caltech.edu
osg-se.ligo.caltech.edu

13 hardware nodes:
dom110 ... dom122

all: CentOS 5.3 x86_64

# Partitions

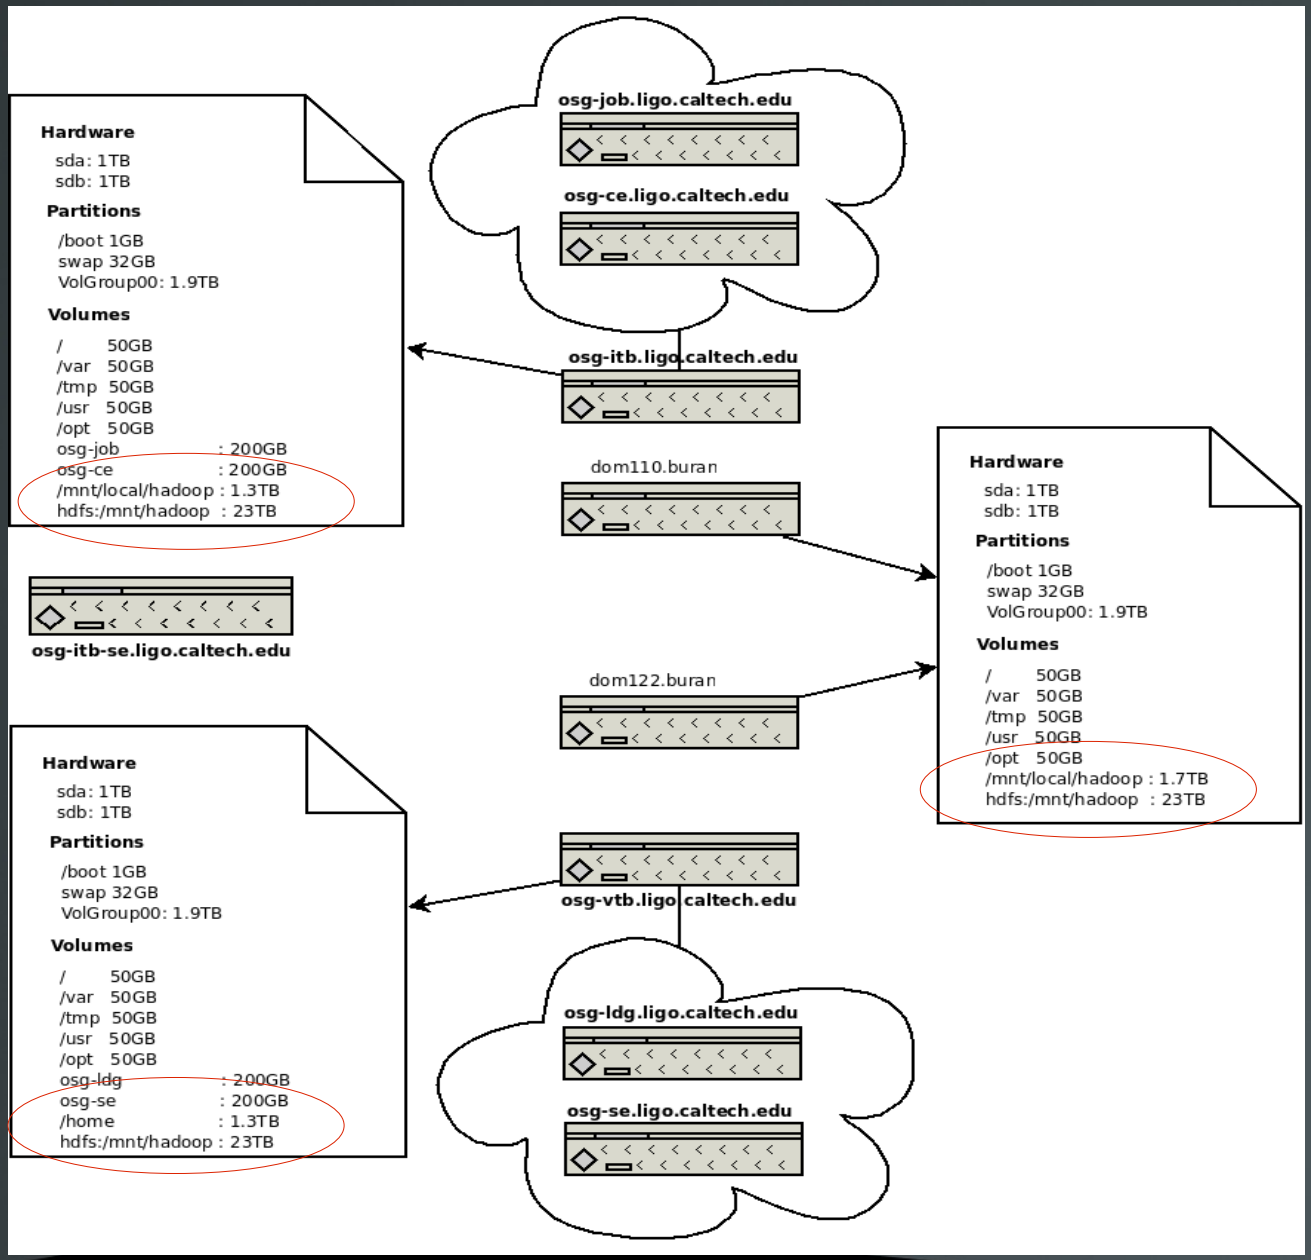


23TB /mnt/hadoop
1TB /home

80% of available storage space made available to applications.

# Services

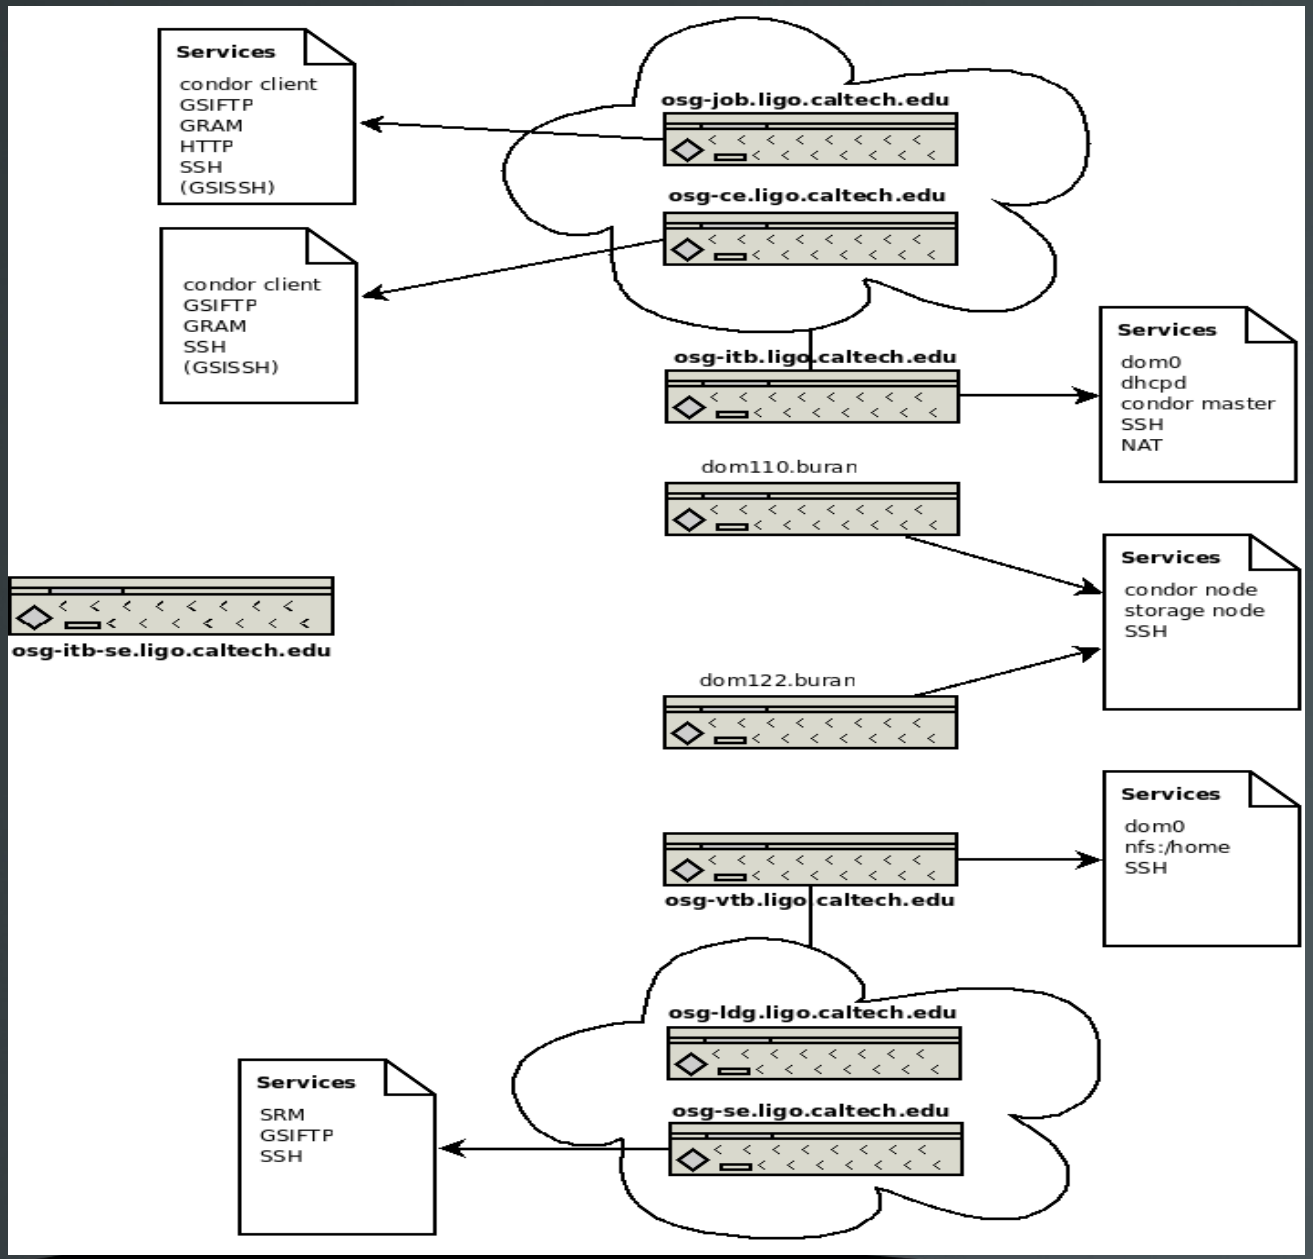


GRAM/Condor ..... 2119
GridFTP .....2811
SRM .....10443


# Why Hadoop?

- ✔ is well documented.
- ✔ has an active and helpful mailing list.
- ✔ is open source backed by Yahoo!
- ✔ is comparable simple to install and configure.
- ✔ is stable as of version 0.20.
- ✔ can be mounted using fuse.
- ✔ can be optimized for IO or redundancy.
- ✗ is not 100% Posix compatible!


# Storage Details

13 storage nodes each read and write a 4GB file.

- ✔ Hadoop service on hdfs://osg-itb.ligo.caltech.edu:9000
- ✔ read: 64MByte/s, write: 48MByte/s, combined: 26 MByte/s

Hadoop is serving a 23TB /mnt/hadoop partition to all nodes

- ✔ /mnt/hadoop/osg/app
- ✔ /mnt/hadoop/osg/data
- ✔ /mnt/hadoop/ligo

BeStMan-Gateway is used as SRM service using Hadoop as backend

- ✔ srm://osg-se.ligo.caltech.edu:10443/srm/v2/server?SFN=/mnt/hadoop/osg/


# LIGO

## Storage Usage and Requirements

- Raw data from detectors is split up into 16MB chunks ( 120s each ).
- Files are hosted by LIGO on several storage resources ( SRM + GridFTP ).
- Globus *Replica Location Service* lists location of files and replica copies.

- Good statistics requires analysis over extended period of times ( 1 wk, 80GB ).
- This amount of data needs to be staged-in before it can be analyzed.
- It can be reused for several data analysis runs.

LIGO requires storage resources that are close to computational resources.
Storage provided externally by SRM could be available internally through mounts.

In this case we need a mapping from SRM:

srm://osg-se.ligo.caltech.edu:10443/srm/v2/server?SFN=/mnt/hadoop/osg/

To internal mount points:

/mnt/hadoop/osg