

# Streaming and storing CineGrid data: *A study on optimization methods*

Sevickson.Kwidama  $\Rightarrow$  os3.nl

July 2, 2008



UNIVERSITY OF AMSTERDAM

## Mission statement

*“CineGrid’s mission is to build an interdisciplinary community focused on the research, development, and demonstration of networked collaborative tools, enabling the production, use and exchange of very high-quality digital media over high-speed photonic networks.”*



# Media Formats(1)



## Media Formats(2)

- **4K**

- $4096 \times 2160$
- $4.3 \times 1080p$  HD
- 36 bpp Color
- One frame = 40 MB
- Uncompressed 7.6 Gb/s

- **SHD<sup>1</sup>**

- $3840 \times 2160$
- $4 \times 1080p$  HD
- 24 bpp Color
- One frame = 25 MB
- Uncompressed 6 Gb/s
- Mostly used in CineGrid, because of hardware limitations

---

<sup>1</sup>Super High-Definition also known as Quad HD

# SAGE vs NTT jpeg2000 codec

- **SAGE<sup>2</sup>**

- Software implementation
- Performance may depend on hardware
- Used for tiled displays
- DXT compression technique
- Compression ratio 6:1

- **NTT jpeg2000 codec**

- Hardware implementation, 4 jpeg2000 codec boards
- Stream is not influenced by hardware
- jpeg2000 compression technique
- Compression ratio 12:1

# Research question 1

## Streaming

A streaming tool is an essential part in CineGrid.

## Streaming question

How do SAGE and NTT jpeg2000 codec compare against each other, regarding network streams?

## Test Material

---

### *7Bridges*

Length: 138 secs

Size: 4.3 GB

Rate: 30 fps

---



---

### *PragueTrain*

Length: 97 secs

Size: 2.5 GB

Rate: 24 fps

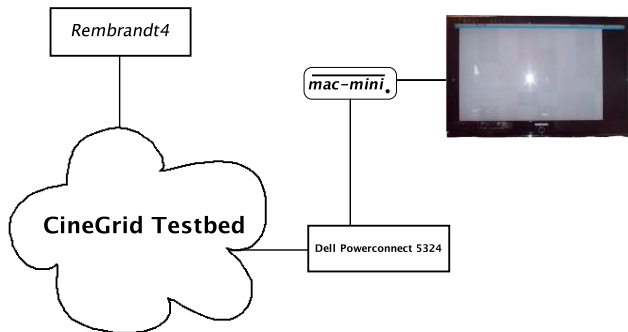
---



# Experiments

## Tools:

- top, measure CPU load
- tcpdump, capture packets send/received
- SAGE manager log, application layer measurements





SAGE test setup



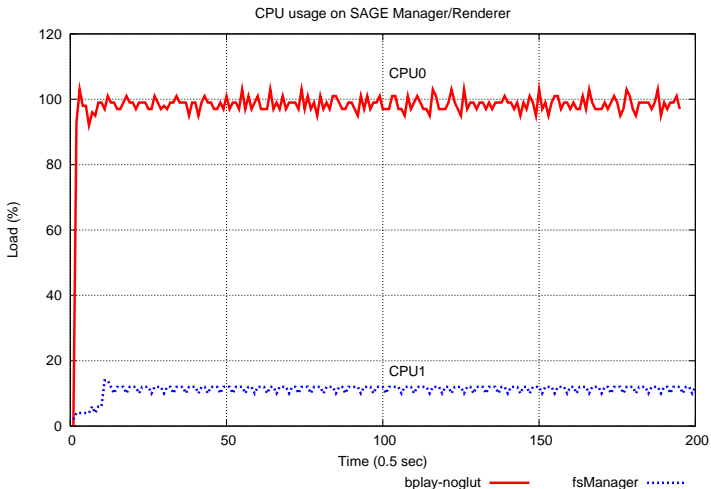
# Bandwidth Results

## SAGE Manager log



L: UDP stream 218.4Mb/s, R: TCP stream 208.1Mb/s



# CPU load SAGE Manager



fsManager: window manager SAGE | bplay-noglut: compresses video and sends

# CPU load SAGE Display

## CPU usage on SAGE display



L: UDP stream, R: TCP stream

## Research question 2

### Storage

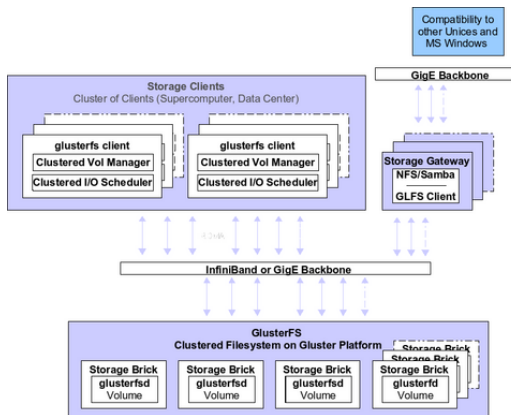
Local Filesystem does not meet demand storage space.  
NFS does not meet requirement of scalability and *read* speed.

### Storage question

Can GlusterFS improve the performance of the CineGrid storage?

# GlusterFS

- GlusterFS is a Cluster Filesystem
- Aggregates “storage bricks” → parallel network file system
- Several translators , i.e., clustering translators and performance translators




# Experiments

## Tools:

- dd, block level measurement
- iozone, file level measurement

## Methods:

- Read performance
- Striping translator
- Filesize 7GB



## Results(1)

4KB test results used for these percentages.

4KB was used because it is the standard block size in most Linux filesystems.

<b>Filesystem</b>	<b>Read speed</b>	<b>dd</b>
	<i>MB/s</i>	<i>%</i>
Local	85.9	
NFS	73.8	14.1%
GlusterFS 1 brick	56	34.8%
GlusterFS 2 bricks	56.3	34.5%
GlusterFS 3 bricks	104	21%

Absolute Read speed, percentage loss or gain in comparison with Local Filesystem

Bottleneck 1Gb/s interface.

Tests with more than three servers were not possible.

## Results(2)

Upgraded to 10 Gb/s connection to client.

Also used Read Ahead Translator and Threaded I/O Translator.

<b>Filesystem</b>	<b>Read speed</b>	<b>dd</b>
	<i>MB/s</i>	<i>%</i>
GlusterFS 1 brick	86.6	0.8%
GlusterFS 2 bricks	103	19.9%
GlusterFS 3 bricks	294	242.3%
GlusterFS 4 bricks	306	256.2%
GlusterFS 5 bricks	315	266.7%

Absolute Read speed, percentage loss or gain in comparison with Local Filesystem



## Future work

- NTT jpeg2000 codec measurements
- SHD instead of 1080p HD SAGE setup
- Test setup GlusterFS for long term stability
- Investigate Translators

# Conclusions

## Streaming

The SAGE performance was as expected.

The video player is CPU-bound.

Performance easily influenced by running processes.

## Storage

GlusterFS increases the read performance for streaming from CineGrid storage.