



# PERFORMANCE AND SIZING WITH SOFTWARE DEFINED STORAGE

MARCEL HERGAARDEN

Sr. Solution Architect, Red Hat Benelux


# Key Takeaways

- Gluster architectures should be tailored for your workloads
- SDS without the right server config is like one hand clapping
- Red Hat offers:
  - High performance, scale-out reference architectures,
  - Pre-configured, workload optimized SKUs via partners
  - Centers of Excellence to build and test enterprise grade solutions

# Red Hat Gluster Storage Scales Up AND Out

Serving a large spectrum of workloads

**Scale UP**  
Add perf per node



**Scale OUT**  
Add nodes



JPEG Web Image Files



DVD Movie Files




CCTV Surveillance  
Recording Cameras

# A six node cluster can process...



JPEG Web  
Image Files  
(32KB)



# Or...



DVD  
Movie Files  
(4GB)

72x 7.2K HDD

**1** DVD  
per second

or

Optimized  
72x 7.2K HDD

**2** DVDs  
per second

or

72x SSD

**4** DVDs  
per second

Or...



High-Def  
CCTV Camera  
Recording Streams

72x 7.2K HDD

**200** CCTV streams  
within latency threshold

or

Optimized  
72x 7.2K HDD

**500** CCTV streams  
within latency threshold

or

72x SSD

**?** CCTV streams  
within latency threshold



**OPTIMAL PERFORMANCE  
FOR YOUR WORKLOAD**

# Getting Optimal Performance

Customized to target workloads




## Pre-Configured SKUs

### Quanta Cloud Technologies

- QxStor RGT-200(SF)/(LF)
- QxStor RGT-400(LF)
- QxStor RGC-400

### Supermicro

- SSG-6028R-GL096/P
- SSG-6028R-GL192/P



## Do-It-Yourself

Based on detailed performance and sizing guides

x86 storage server from favorite vendor



# CLUSTER DESIGN CONSIDERATIONS

Select data  
protection method

Qualify need for  
scale-out storage

Determine fault-domain  
risk tolerance

Design for target  
workload IO profile(s)

Identify Capacity

Choose storage  
access method(s)



# Qualify Need for Scale Out Storage



**Elastic  
Provisioning**




**Commodity  
Hardware**



**Highly  
Available  
Storage  
Pool**



**Connect  
Storage  
Islands**



**Incremental,  
Non-forklift  
Upgrades**

<http://red.ht/2ntr5f3>



# PARTNER ECOSYSTEM

# Jump Start Technical Design Discussions

Reference Architectures and Centers of Excellence

IOPS  
OPTIMIZED

THROUGHPUT  
OPTIMIZED

COST/ CAPACITY  
OPTIMIZED

Use Case: MySQL

Use Case: Rich Media

Use Case: Active Archives



SanDisk



SAMSUNG



General SDS Hardware  
Performance & Sizing

Container Storage  
Center of Excellence

# Joint Innovation With Partners

## Performance and Sizing Guides

<http://red.ht/2mg9kQ5>



### INTRODUCTION

A significant portion of financial data analysis is performed by stitching together large volumes of big data. However, enterprise storage continues to evolve as well as the ways the data is used and rapidly expands in size from underlying core systems. Penguin Computing and Red Hat Storage have built a cost-effective and scalable solution to enable financial institutions to make more insight for their business decisions.

### THE PROBLEM

Financial institutions, traders, and investors closely track a number of asset classes across an emerging world. These may include securities, equities, futures, interest rates, foreign exchange, and commodities, such as oil, gas, or foreign bank data.

Also, operators demand investment – represented by a tick – over a period of time to maintain value to run analysis, operational procedures, trading trading and market making strategies, and sophisticated user needs.

This data needs to be analyzed, volume and many other dimensions such as volume and price, size or size class, order flow, trade time, or average order size for the vast range of products.

In a single trade, the current transaction occurs at a price that the previous transaction, the bid/ask spread, the current transaction occurs at a lower price than the previous transaction. Consequently, a series of orders to a trade stream from multiple parties and complete historical data may be used with a variety of analytical applications to increase profitability, including:

- Low-frequency correlations.
- Non-linear and event-driven correlation.
- Market impact, structure, and order flow simulation.

Active order book trading requires the data to be collected through a number of data sources including an electronic order book, a market data feed, and a data feed. Each file is usually written in a single sequential stream of bytes.

The granularity of data stored within each trade file depends on the client, the structure of the data, the frequency and what kind of analysis must be performed on the data. The higher the resolution of the data, the larger the volume of data. Thus, the amount of storage capacity required.

Storage capacity for big data could grow quickly from hundreds of gigabytes to terabytes or even petabytes. Efficient big data storage, therefore, is a critical business challenge for most financial institutions.



Red Hat Gluster Storage  
Penguin Computing

© 2015 Red Hat

<http://red.ht/2k6uYcg>



# Learn More



[redhatstorage.redhat.com](http://redhatstorage.redhat.com)

[red.ht/GlusterTestDrive](http://red.ht/GlusterTestDrive)



[plus.google.com/+RedHatStorage](https://plus.google.com/+RedHatStorage)



[facebook.com/redhatstorage](https://facebook.com/redhatstorage)



[linkedin.com/company/red-hat](https://linkedin.com/company/red-hat)



[twitter.com/RedHatStorage](https://twitter.com/RedHatStorage)



[youtube.com/user/RedHatStorage](https://youtube.com/user/RedHatStorage)