



Advancements in Storage QoS Management in National Data Storage

Darin Nikolow¹, Renata Słota¹, Stanisław
Polak¹ and Jacek Kitowski^{1,2}

¹ AGH University of Science and Technology, Faculty of Computer Science, Electronics and
Telecommunications, Department of Computer Science, Kraków, Poland

² AGH University of Science and Technology, ACC Cyfronet AGH, Kraków, Poland

emails: {darin,rena,polak,kito}@agh.edu.pl

Outline

- Introduction
- QoS and SLA for storage
- National Data Storage 2 project
- QoS management model in NDS2
- Test results
- Summary

Introduction

- Increasing role of storage systems
 - influence the performance and availability of applications and services
- More data – more challenges
 - Management, scalability, reliability, availability, performance, data endurance
 - Digital universe in 2012 – 2.8 ZB, growing exponentially
- Constantly growing users' requirements
- 4th paradigm of science
 - Scientific discoveries, data intensive applications
 - Many scientific applications process large amounts of data
- Applications executed in distributed environments

QoS and SLA for storage

- Scientific applications may have storage QoS requirements, e.g.,
 - minimal transfer rate
 - Writing irreproducible data coming from scientific instruments (HEP colliders, telescopes)
 - Recovering from backup
- SLA
 - Part of the contract between the resource provider and the user
 - Specify QoS requirements (SLO) and penalties
 - SLA parameters are used to form SLO,
 - Example: $\text{TransferRate} > 1 \text{ MB/s}$, in 99% of the transfers

SLA parameters and QoS metrics

- QoS metrics are low level QoS parameters usually obtained from a monitoring system or calculated from other QoS metrics
- SLA parameters are mapped to QoS metrics or calculated from more QoS metrics
- Types of SLA parameters
 - Performance parameters
 - Data protection parameters
 - Availability parameters

QoS metrics examples

- UserReadTransferRate
- SNAvailableReadTransferRate

$$SNRedundancyLevel = \begin{cases} NrOfCopies, & \text{dla HSM} \\ \frac{NrOfRedundantDisks}{NrOfAllDisk - NrOfRedundantDisks}, & \text{dla macierzy RAID} \end{cases}$$

$$SNLoad = \frac{SNCurrentReadTransferRate}{SNMaxReadTransferRate} + \frac{SNCurrentWriteTransferRate}{SNMaxWriteTransferRate}$$

SLA parameters examples

- ReadTransferRate
- DataSecurityLevel
- DataAvailabilityLevel

$$DataProtectionLevel = \sum_{i \in N} (1 + SNRedundancyLevel_i)$$

NDS2 project

- Provides backup, archiving and general data storage services
 - Successor of the NDS project
 - NDS software is used in PLATON-U4
 - Integration with PL-Grid is ongoing
- The stress (in NDS2) is put on security and data availability
 - Advanced encryption techniques
 - Multiple replicas stored in geographically different locations
- Data sharing
- SLA and QoS management

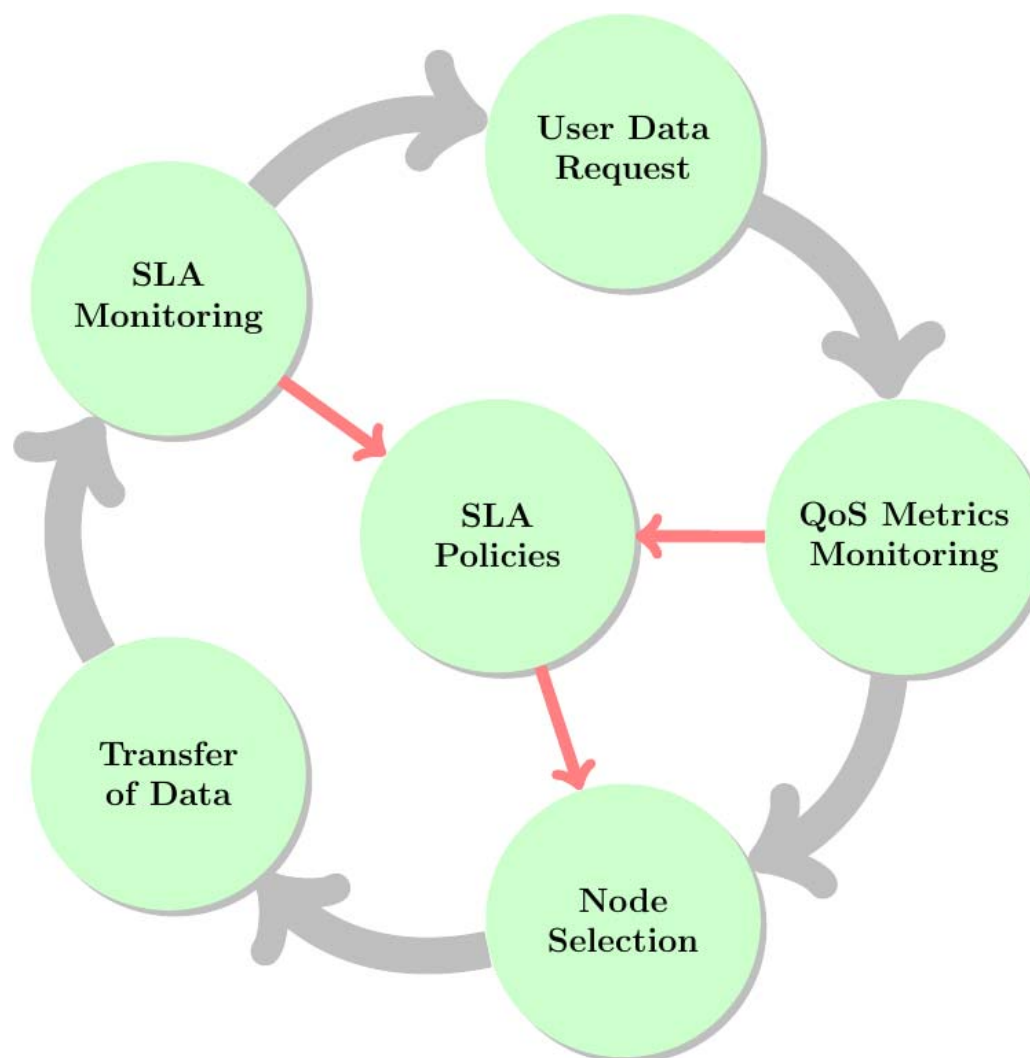
NDS2 components (physical layer)

- Storage systems
 - Disk array based
 - HSM (Tivoli Storage Manager)
- Storage Nodes (SN)
 - Allows access to storage systems
- Access Nodes (AN)
 - Responsible for serving client requests

NDS2 main software components

- Meta catalog (MC)
 - Holds the metadata
- Live Daemon (LD)
 - Monitors the status of other components
- Data Daemon (DD)
 - Provides access to the data via FUSE-based virtual filesystem
- Replica Daemon (RD)
 - Creates and remove replicas
- User Management System (UMS)
 - Provides access to user's accounts related data – profiles, passwords, etc.
- Quality Management System (QMS)
 - Provides the necessary SLA and QoS related functionality of NDS2

Model of storage QoS management



SLA Policies

- Define policies for storage resource management with respect to SLA
- Differentiate SN selection depending on SLA profiles
- Include a set of coefficients specifying importance of QoS metrics

SLA profiles

- SLA profiles define limits for a group of SLA parameters
- SLA profiles are assigned to users
- Three SLA profiles are defined in NDS2
 - Standard,
 - Fast access
 - data mostly on fast access storage, e.g. disk array,
 - High data protection
 - more replicas

Selecting of SN

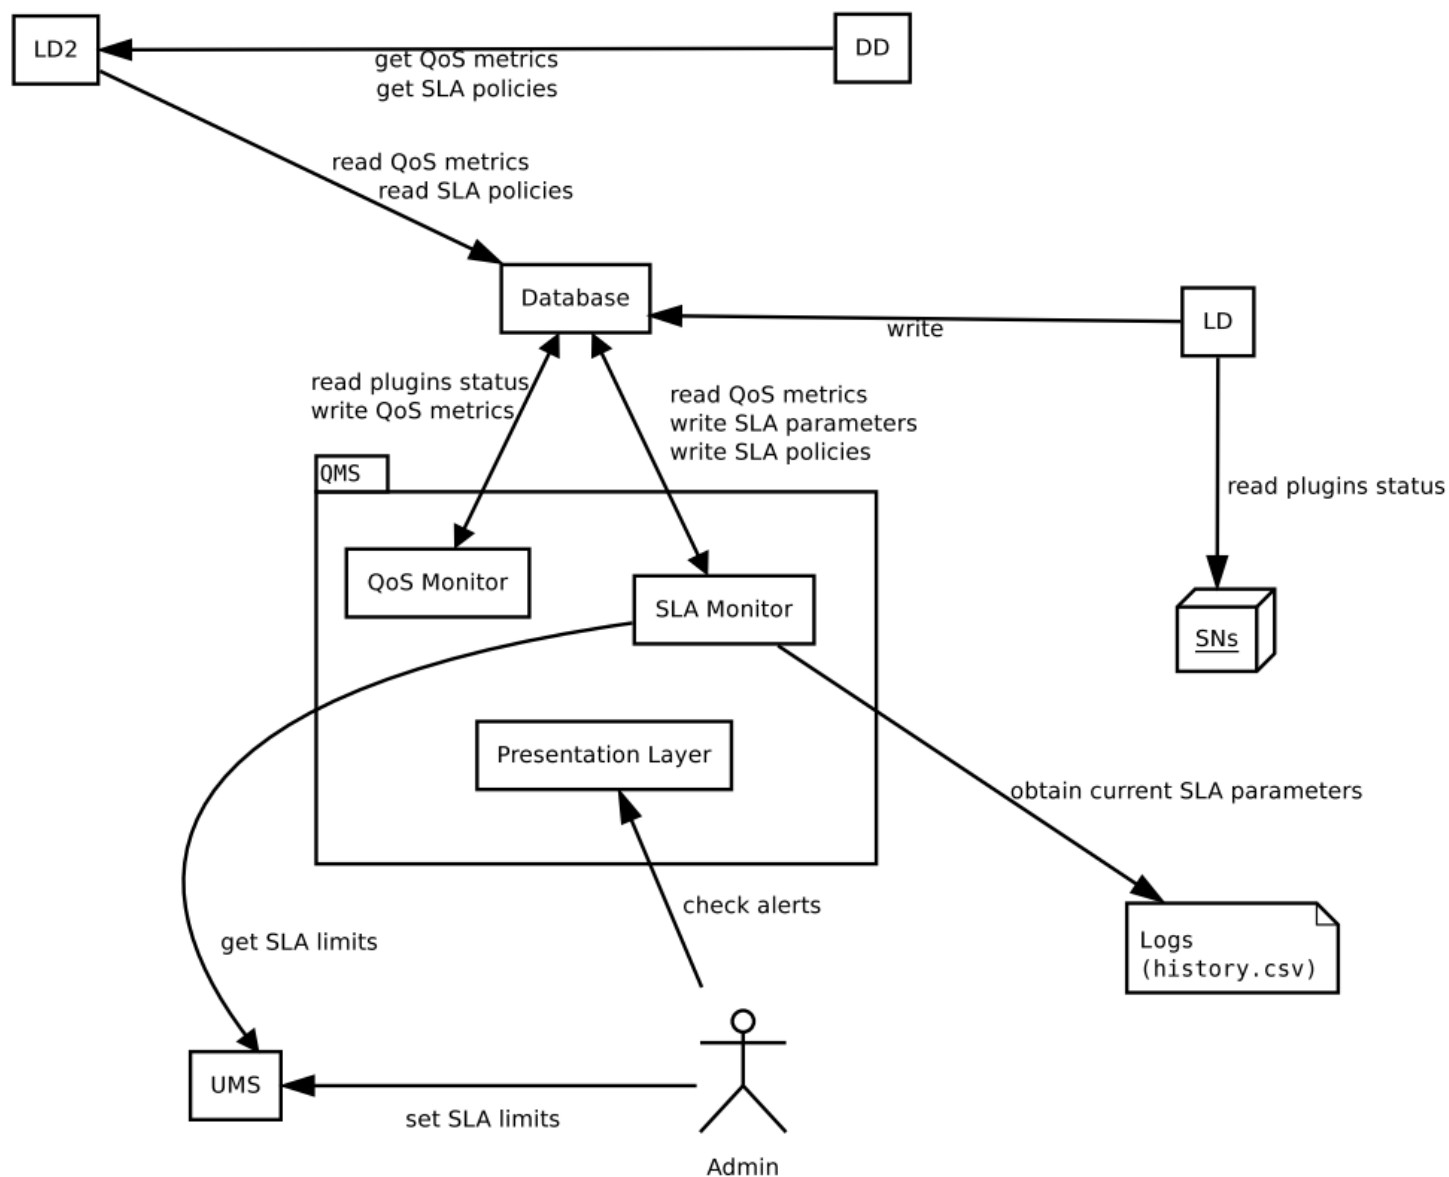
- Selecting of SN is required when a file is read (replica selection) or when a new replica is created

$$\begin{aligned} SNReadPreference = & k_1 * SNMaxReadTransferRate * filesize - \\ & k_2 * SNCurrentReadTransferRate * filesize - \\ & (k_3 * \frac{filesize}{SNTapeReadRate} + k_4 * SNTapeLatency) * \\ & k_5 * SNisFileCached - \\ & k_6 * SNLoad - \\ & k_7 * SNIOps \end{aligned}$$

$$\begin{aligned} SNWritePreference = & k_1 * SNMaxWriteTransferRate * filesize - \\ & k_2 * SNCurrentWriteTransferRate * filesize - \\ & k_6 * SNLoad - \\ & k_7 * SNIOps \end{aligned}$$

- SNisFileCached = 0 – cached, 1 – on tape
- The coefficients k_1 - k_7 can be different for each SLA level and define the SLA policy
- k_1 - k_7 are set by QMD and are not expected to be changed often

QMS architecture



SLA Tests


- Test environment
 - 2x HP blade BL460c 12 cores (2x Intel Xeon CPU 2.53GHz)
 - VMware vSphere virtualization environment
 - Storage nodes
 - SN1 – 2 cores, 2GB RAM, EVA8100 disk array volume, 256GB
 - SN2 – 2 cores, 2GB RAM, EVA8000 disk array volume, 2TB
 - Management nodes
 - Access Node – 2 cores, 2GB RAM
 - Database node – 2 cores, 2GB RAM

SLA testing procedure

- A set of test files is created (10x 2GB)
- The filesystem buffers are flushed
- Additional storage load is generated locally to one of the SNs using the *fio* Linux filesystem benchmark tool
- The functionality of QMS (SLA) is turned off
- A number of read transfers are sequentially committed and performance data logged
- The filesystem buffers are flushed again
- The functionality of QMS (SLA) is turned on
- The same sequence of read transfers is committed again

Test results:

Performance of data access with SLA



Summary

- SLA parameters and QoS metrics for distributed storage systems are proposed
- A model of storage QoS management has been proposed
 - SLA monitoring
 - SLA policies
 - SN selection via heuristics
 - SLA-aware load balancing
- Integration is ongoing