

# Monitoring of a Grid Storage Virtualization Service

*Jacques Jorda, Université Paul Sabatier, Toulouse, France*

*Aurélien Ortiz, Université Paul Sabatier, Toulouse, France*

*Abdelaziz M'zoughi, Université Paul Sabatier, Toulouse, France*

*Salam Traboulsi, Université Paul Sabatier, Toulouse, France*

---

## ABSTRACT

*Grid computing is commonly used for large scale application requiring huge computation capabilities. In such distributed architectures, the data storage on the distributed storage resources must be handled by a dedicated storage system to ensure the required quality of service. In order to simplify the data placement on nodes and to increase the performance of applications, a storage virtualization layer can be used. This layer can be a single parallel filesystem (like GPFS) or a more complex middleware. The latter is preferred as it allows the data placement on the nodes to be tuned to increase both the reliability and the performance of data access. Thus, in such a middleware, a dedicated monitoring system must be used to ensure optimal performance. In this paper, the authors briefly introduce the Visage middleware – a middleware for storage virtualization. They present the most broadly used grid monitoring systems, and explain why they are not adequate for virtualized storage monitoring. The authors then present the architecture of their monitoring system dedicated to storage virtualization. We introduce the workload prediction model used to define the best node for data placement, and show on a simple experiment its accuracy.*

*Keywords: Data Storage, Distributed Architectures, Grid Computing, Monitoring Systems, Storage Virtualization*

---

## INTRODUCTION

The grid concept defines the aggregation of heterogeneous computing and storage nodes across administrative domains to scale the performance of individual computer nodes (Foster, 1999). There are two paradigms in grid computing: large-scale grids, involving at least thousands of nodes from standard com-

puters, and meta-cluster type grids involving a maximum of hundreds of nodes from high performance computers – typically servers or supercomputers. We will focus on the latter case, mostly used for high performance computing. In such systems, a wide range of distributed physical resources can be used to data processing activity, especially for scientific applications, such as biology research and simulation for

DOI: 10.4018/jghpc.2013010104

weather forecasts. These scientific applications require huge computations in order to achieve their goals, generating access to large amounts of data handled by the deployed grid storage resources.

Before launching such application, data are transferred to the dedicated nodes using tools like GridFTP; the application then runs on these nodes, and the results are merged to a user repository. However, in domains like weather forecasting, the granularity may be refined to produce more accurate results. Thus, these large scale applications involve an increasing number of nodes handling larger and larger amount of data. Consequently, a new problem arises, which is more about data storage than processing unit allocation. Indeed, whatever scientific computation is considered, more accurate results involve more data to store and handle. Then, manual strategies using tools like GridFTP are more difficult to implement, and the performance gain is not linear with the number of nodes: the data distribution on the nodes for HPC programs is complex, and impacts directly the sustained performance. Consequently, the need is to improve data storage on grids to overcome the current limitations.

The most advanced techniques in improving data storage on distributed systems deal with virtualization. Virtualization entails grouping physical resources in virtual ones, in order to mask the hardware complexity. Therefore, like grid middleware, storage virtualization systems handle the heterogeneity and dynamicity of nodes and network workload, while providing a transparent and uniform data access interface to the users, making data access as easy on the distributed environment as on a single machine. The work presented in this paper is part of the ViSaGe project. ViSaGe is a middleware designed to provide the set of functionalities needed for storage virtualization: transparent reliable remote access to data and distributed data management. This storage virtualization is of prime interest when working on grids. In fact, when working either with a meta cluster of a few supercomputers or a larger grid of a few hundreds of servers, the transparent ac-

cess of data is needed to easily distribute jobs among compute resources. Moreover, efficient mechanism must be employed to guarantee a high level of performance.

Using ViSaGe middleware, data access is easier, and complex data management like data reorganization, data redundancy, or replication can be used to improve performance. However, such data management techniques require an efficient mechanism to guarantee the required quality of service and to ensure maximal performance. This task is even more complex because the characteristics of physical storage resources evolve over time. Therefore the various resources need to be monitored in order to react to any significant changes of observed criteria, which require an adjustment to storage parameters.

A dedicated monitoring system must then be used to trace physical resources usage and to feed back the administration tool to update storage parameters adequately. These parameters are mainly the placement of redundant / spitted data on nodes to improve both data reliability and access performance. In order to give ViSaGe an appropriate Administration and Monitoring component, we first studied the major existing grid monitoring tools to assess their capabilities to be used in the context of storage virtualization. This study will be summarized in the second section, and we will explain why these tools cannot be used in ViSaGe. In the following section, we will introduce ViSaGe components, and present the architecture of Admon (the administration and monitoring component). We will then present in the next section the prediction model used, and show on a basic experiment that it accurately predicts the I/O workload variation. Finally, we will conclude the paper and present some future work.

## BACKGROUND

Monitoring large scale systems, like grids, implies facing a number of challenges. First, the monitoring system must be scalable. In fact,

15 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the product's webpage:

[www.igi-global.com/article/monitoring-of-a-grid-storage-virtualization-service/78735?camid=4v1](http://www.igi-global.com/article/monitoring-of-a-grid-storage-virtualization-service/78735?camid=4v1)

This title is available in InfoSci-Journals, InfoSci-Journal Disciplines Computer Science, Security, and Information Technology. Recommend this product to your librarian:

[www.igi-global.com/e-resources/library-recommendation/?id=2](http://www.igi-global.com/e-resources/library-recommendation/?id=2)

## Related Content

---

### On Construction of Cluster and Grid Computing Platforms for Parallel Bioinformatics Applications

Chao-Tung Yang and Wen-Chung Shih (2011). *International Journal of Grid and High Performance Computing* (pp. 69-88).

[www.igi-global.com/article/construction-cluster-grid-computing-platforms/52587?camid=4v1a](http://www.igi-global.com/article/construction-cluster-grid-computing-platforms/52587?camid=4v1a)

### Dynamic Overlay Networks for Robust and Scalable Routing

Bart De Vleeschauwer, Filip De Turck, Bart Dhoedt and Piet Demeester (2010). *Handbook of Research on P2P and Grid Systems for Service-Oriented Computing: Models, Methodologies and Applications* (pp. 543-565).

[www.igi-global.com/chapter/dynamic-overlay-networks-robust-scalable/40817?camid=4v1a](http://www.igi-global.com/chapter/dynamic-overlay-networks-robust-scalable/40817?camid=4v1a)

### Monitoring and Controlling Large Scale Systems

Valentin Cristea, Ciprian Dobre, Corina Stratan and Florin Pop (2010). *Large-Scale Distributed Computing and Applications: Models and Trends* (pp. 141-167).

[www.igi-global.com/chapter/monitoring-controlling-large-scale-systems/43106?camid=4v1a](http://www.igi-global.com/chapter/monitoring-controlling-large-scale-systems/43106?camid=4v1a)

## Keys for Administration of Reconfigurable NoC: Self-Adaptive Network Interface Case Study

Rachid Dafali and Jean-Philippe Diguët (2010). *Dynamic Reconfigurable Network-on-Chip Design: Innovations for Computational Processing and Communication* (pp. 67-83).

[www.igi-global.com/chapter/keys-administration-reconfigurable-noc/44221?camid=4v1a](http://www.igi-global.com/chapter/keys-administration-reconfigurable-noc/44221?camid=4v1a)