Future Generation Computer Systems 26 (2010) 622–632

Contents lists available at ScienceDirect

Future Generation Computer Systems

journal homepage: www.elsevier.com/locate/fgcs

A high performance suite of data services for grids

Alberto Sánchez [a], María S. Pérez [b,*], Jesús Montes [c], Toni Cortes [d]

[a] Universidad Rey Juan Carlos, Spain
[b] Facultad de Informática, Universidad Politécnica de Madrid, Spain
[c] CeSViMa, Universidad Politécnica de Madrid, Spain
[d] Universidad Politécnica de Cataluña and the Barcelona Supercomputing Center, Spain

ARTICLE INFO

Article history:
Received 24 September 2009
Received in revised form
20 November 2009
Accepted 29 November 2009
Available online 2 December 2009

Keywords:
Data access services
Data exchange
Data-intensive applications
Data grids
Data storage

ABSTRACT

Recent advances in computation have made feasible the resolution of very complex problems. Grid computing has matured to currently become one of the most successful initiatives for dealing with these challenging problems, which usually require tackling huge volumes of data. Although the main goal of the grid technology is sharing resources among several virtual organizations, few research works are oriented to increase the performance of grid solutions.

MAPFS-Grid is a complete suite of components for providing high performance access to huge volumes of data in a grid environment. MAPFS-Grid is composed of three different services: (i) Parallel Data Access Service, a WSRF-compliant grid service; (ii) MAPFS-DSI, a GridFTP-compliant service; and (iii) MAPFS-DAI, an OGSA-DAI-compliant service. Each one is suitable for different scenarios, which will be described in this article.

This paper describes and evaluates the performance of MAPFS-Grid, showing the main advantages of MAPFS-Grid vs other proposals.

1. Introduction

Advances in the computational field have made it possible to develop increasingly complex applications to face new and challenging problems. Most of these applications require the management and analysis of large volumes of data, which range from terabytes to petabytes. These data-intensive applications can be found in several domains, including Physics [1], climate modeling [2], Biology [3] or visualization [4]. The I/O phase usually constitutes the bottleneck of these applications.

Data grid [5–7] aims at developing suitable solutions to these kinds of applications by means of grid-based tools. Indeed, a data grid is specifically designed to store, manage, and provide reliable access to data.

Data grid services must provide a series of features in order to achieve maximum performance:

- Ability to search through numerous available datasets.
- Ability to select suitable computational resources to perform data analysis.
- Ability to manage access permissions.
- Intelligent resource allocation and scheduling.

Due to the basic grid principles, the environment is characterized by its heterogeneity. In the case of a data grid, this includes different storage systems, data access mechanisms, data access policies, and data formats. The data grid management infrastructure must act as an abstraction layer that provides a common, standard and efficient procedure to access the information stored.

Although data grid allows heterogeneous data resources to be shared, only few research works in the field of data grid are oriented to increase the performance of these solutions [8,9].

The aim of this work is to develop a complete suite of services for high performance access to huge volumes of data in a grid environment. Our approach, named MAPFS-Grid [10], is intended to provide a high performance access by means of the following kind of services:

(1) A generic WSRF-compliant data access service, which uses Simple Object Access Protocol (SOAP) and Web services technology. This proposal follows the OGSA guidelines [11], which propose Web services as basic technology for building grids.
(2) A performance-oriented data access service based on GridFTP and built within MAPFS-Grid. GridFTP [12] is a high-performance and reliable file transfer protocol, although it does not follow strictly the OGSA scheme.
(3) An OGSA-DAI-compliant service for providing a uniform access and better performance than OGSA-DAI [13]. OGSA-DAI provides a uniform way of querying, accessing, updating, and transforming different type of data resources by means of Web services.

* Corresponding author.
E-mail addresses: alberto.sanchez@urjc.es (A. Sánchez), mperez@fi.upm.es (M.S. Pérez), jmontes@cesvima.upm.es (J. Montes), toni.cortes@bsc.es (T. Cortes).

0167-739X/$ – see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.future.2009.11.006

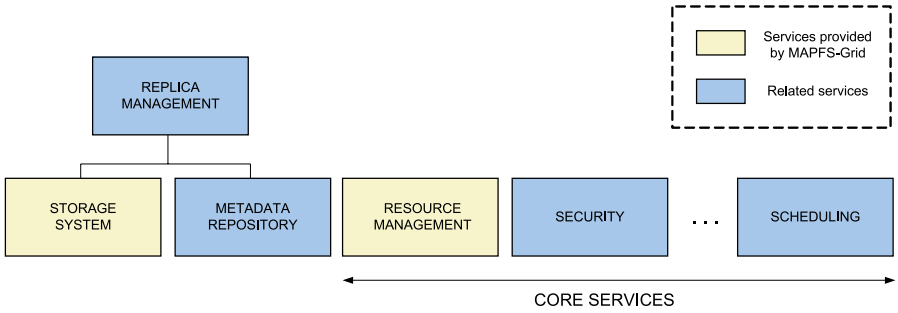

Fig. 1. Services provided by MAPFS-Grid and related.

All these three scenarios cover the needs of most data grid-based applications. Moreover, services provided by MAPFS-Grid are incorporated within the generic architecture of a grid. Fig. 1 shows the scheme in which MAPFS-Grid is integrated. MAPFS-Grid uses other core grid services and supports data management and storage system, allowing other services (such as the replica manager) to offer more complex functionalities. All these other services are outside the scope of this paper, although they are integrated together with MAPFS-Grid.

The outline of this paper is as follows. Section 2 describes our proposal, MAPFS-Grid, as a suite of three different data services intended to provide high performance in different scenarios. Section 3 shows different experimental results, which evaluate the main characteristics of MAPFS-Grid. Section 4 enumerates different works related to our proposal. Finally, in Section 5 the conclusions of our work and the open research lines are described.

2. MAPFS-Grid, a suite of services for accessing large volumes of data

As we have seen in the previous section, different needs arise in data management and access in grids. We have noticed three important aspects related to these needs:

(1) An efficient data service that manages file resources following the OGSA architecture is required. The Web services technology is suitable for managing services and resources through Web Service Resource Framework (WSRF) [14]. As far as we know, there are not WSRF-based data services designed to increase the performance of the I/O operations.
(2) The most important drawback of the previous scenario is the low performance exhibited by Web services. In fact, the use of XML and SOAP as transfer protocol is not appropriate for performance-critical applications [15]. Although there are different proposals for dealing with this decrease of performance [16–18], none of them are suitable for scenarios demanding high throughput. In this context, GridFTP is an optimized protocol for transfer of large files, since it is not based on SOAP transfer. Additionally, GridFTP extends the basic FTP protocol to support data transfer among multiple servers (striping). Furthermore, GridFTP enables the use of multiple TCP streams in parallel from a source to a sink (parallelism) [19].
(3) Every grid project usually provides "ad hoc" solutions to the data management. A data service often offers a *native* interface, which does not provide interoperability with other I/O systems. OGSA-DAI [13] has emerged to provide a uniform access to data sources in a grid environment. However, OGSA-DAI is not focused on the performance of the I/O operations. Therefore, providing a bridge between the interoperability and the performance optimization is an important need of current data grid projects.

MAPFS-Grid tries to build a generic framework where all these problems can be solved, by means of the definition of different services, suitable for these three identified scenarios.

MAPFS-Grid makes use of MAPFS [20], a high-performance parallel file system for clusters of workstations. MAPFS (*Multi Agent Parallel File System*) has been developed in the Universidad Politécnica de Madrid since 2003. The main contribution of MAPFS is the conceptual use of agents to provide applications with new properties, with the aim of increasing their adaptation to dynamic and complex environments. MAPFS is based on a multiagent architecture that offers features such as data acquisition, caching, prefetching, and the use of hints.

The feasibility of the combination between MAPFS and MAPFS-Grid is due to the fact that grid environments are composed of different and heterogeneous resources, being clusters one of the most used because of its good relation power vs. cost. Thus, it is possible to improve the grid data operations through parallel accesses into the cluster resources. MAPFS distributes data stripes over all the nodes of a cluster. On the other hand, MAPFS-Grid allows heterogeneous servers connected by means of a wide-area network to be used as data repositories, by storing data in a parallel way through all the clusters and individual nodes which compose the grid.

The heterogeneity of grid environments makes the application of parallelism difficult. In fact, since every resource of the grid can be composed of several components (e.g., clusters of workstations), it is necessary to optimize the I/O performance of every resource before tackling the global I/O optimization. Therefore, MAPFS-Grid, as a complete suite of services, provides two levels of software parallelism in a grid:

(1) The high level provides parallelism among the grid storage elements, that is, *inter-storage element parallelism*.
(2) The low level provides parallelism among the set of nodes of each cluster, that is, *intra-cluster parallelism*. At this level, MAPFS is applied.

Both levels are integrated and cooperate with the aim of providing an enhanced I/O bandwidth. The inner level is achieved only if the storage element is a cluster of nodes. The outer level is provided among a set of storage elements, at grid level. Next subsections describe the three MAPFS-Grid services.

2.1. MAPFS-Grid Parallel Data Access Service (PDAS)

Our first proposal is to provide a grid-like interface to MAPFS. This WSRF-compliant service, named PDAS, allows parallel I/O operations to be made in a cluster environment. The idea of this service comes from Data Access and Integration Service (DAIS) [21]. PDAS is an adaptation of this concept from the performance and parallelism viewpoints.

The two levels of parallelism provided by PDAS are shown in Fig. 2. The level 1 parallelism is provided by several PDAS (in every storage element), which give support to a distributed data