Chapter 9 Replica Management in Data Intensive Distributed Science Applications

Ann L. Chervenak University of Southern California, USA

Robert Schuler University of Southern California, USA

ABSTRACT

Management of the large data sets produced by data-intensive scientific applications is complicated by the fact that participating institutions are often geographically distributed and separated by distinct administrative domains. A key data management problem in these distributed collaborations has been the creation and maintenance of replicated data sets. This chapter provides an overview of replica management schemes used in large, data-intensive, distributed scientific collaborations. Early replica management strategies focused on the development of robust, highly scalable catalogs for maintaining replica locations. In recent years, more sophisticated, application-specific replica management systems have been developed to support the requirements of scientific Virtual Organizations. These systems have motivated interest in application-independent, policy-driven schemes for replica management that can be tailored to meet the performance and reliability requirements of a range of scientific collaborations. The authors discuss the data replication solutions to meet the challenges associated with increasingly large data sets and the requirement to run data analysis at geographically distributed sites.

DOI: 10.4018/978-1-61520-971-2.ch009



Figure 1. Distributed collaboration or virtual organization

INTRODUCTION

In the last decade, a large number of distributed scientific collaborations have generated an everincreasing amount of application data. These scientific collaborations include projects in high energy physics, gravitational wave physics, neuroscience, earthquake science, astronomy, and many others. Collaborations that span multiple institutions (see Figure 1) are typically called Virtual Organizations (Foster, Kesselman et al. 2001). Management of the large data sets produced by these data-intensive scientific applications is complicated by the fact that participating institutions are often geographically distributed and separated by distinct administrative domains. A key data management problem in these distributed collaborations has been the creation and maintenance of replicated data sets.

Data sets are replicated among the institutions of a scientific collaboration for a variety of reasons. One goal of replication is to provide high availability for data sets, so that if there are hardware failures at one site or if one region is affected by a network outage or a natural disaster, additional copies of the data sets may still be accessed at other locations. Another goal of replication is to improve performance by allowing a data set to be accessed at multiple locations. A client fetching the desired data can choose a particular site based on resource characteristics, such as the available bandwidth on storage systems or networks. Alternatively, clients can access portions of the file at multiple replica sites in parallel to increase the aggregate bandwidth available for downloading the data. Other reasons for replicating data may be the desire of scientists to have a copy of essential data sets at their home institutions. In addition, data sets may be replicated to improve the performance of scientific computations or workflows running on particular computational resources, such as a supercomputer or computation cluster. In this case, the data sets needed for the computation are frequently replicated or staged in to storage resources on or near the computational site.

Scientific data sets are frequently read-only, but in some cases, they may be updated to reflect corrections or recalibrations of data sets. When an original data set is updated, its replicas must also be updated using a consistency scheme agreed upon by the Virtual Organization. This may include a versioning scheme that keeps all versions of a data item and gives them increasing version numbers as data items are updated.

In this chapter, we trace the evolution of replica management techniques for distributed scientific collaborations, from an early focus on 16 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/replica-management-data-intensivedistributed/62828

Related Content

ACO Based Dynamic Scheduling Algorithm for Real-Time Multiprocessor Systems

Apurva Shahand Ketan Kotecha (2011). *International Journal of Grid and High Performance Computing* (pp. 20-30).

www.irma-international.org/article/aco-based-dynamic-scheduling-algorithm/56353

Prediction of Breast Cancer Using Extremely Randomized Clustering Forests (ERCF) Technique: Prediction of Breast Cancer

Akhil Gupta, Rohit Anand, Digvijay Pandey, Nidhi Sindhwani, Subodh Wairya, Binay Kumar Pandeyand Manvinder Sharma (2021). *International Journal of Distributed Systems and Technologies (pp. 1-15).* www.irma-international.org/article/prediction-of-breast-cancer-using-extremely-randomized-clustering-forests-ercf-technique/287859

Analyzing the Robustness of HPC Applications Using a Fine-Grained Soft Error Fault Injection Tool

Qiang Guan, Nathan DeBardeleben, Sean Blanchard, Song Fu, Claude H. Davis IVand William M. Jones (2016). *Innovative Research and Applications in Next-Generation High Performance Computing (pp. 277-305).*

www.irma-international.org/chapter/analyzing-the-robustness-of-hpc-applications-using-a-fine-grained-soft-error-faultinjection-tool/159049

Design and Implementation of Hybrid Time (HT) Group Communication Protocol for Homogeneous Broadcast Groups

Isamu Tsuneizumi, Ailixier Aikebaier, Makoto Ikeda, Tomoya Enokidoand Makoto Takizawa (2011). International Journal of Distributed Systems and Technologies (pp. 37-48). www.irma-international.org/article/design-implementation-hybrid-time-group/55420

PheGee@Home: A Grid-Based Tool for Comparative Genomics

Bertil Schmidt, Chen Chen, Weiguo Liuand Wayne P. Mitchell (2012). *Grid and Cloud Computing: Concepts, Methodologies, Tools and Applications (pp. 1885-1903).* www.irma-international.org/chapter/phegee-home-grid-based-tool/64572