Chapter 6
Data Intensive Computing with Clustered Chirp Servers

Douglas Thain
University of Notre Dame, USA

Peter Bui
University of Notre Dame, USA

Michael Albrecht
University of Notre Dame, USA

Rory Carmichael
University of Notre Dame, USA

Hoang Bui
University of Notre Dame, USA

Scott Emrich
University of Notre Dame, USA

Patrick Flynn
University of Notre Dame, USA

ABSTRACT

Over the last few decades, computing performance, memory capacity, and disk storage have all increased by many orders of magnitude. However, I/O performance has not increased at nearly the same pace: a disk arm movement is still measured in milliseconds, and disk I/O throughput is still measured in megabytes per second. If one wishes to build computer systems that can store and process petabytes of data, they must have large numbers of disks and the corresponding I/O paths and memory capacity to support the desired data rate. A cost efficient way to accomplish this is by clustering large numbers of commodity machines together. This chapter presents Chirp as a building block for clustered data intensive scientific computing. Chirp was originally designed as a lightweight file server for grid computing and was used as a "personal" file server. The authors explore building systems with very high I/O capacity using commodity storage devices by tying together multiple Chirp servers. Several real-life applications such as the GRAND Data Analysis Grid, the Biometrics Research Grid, and the Biocompute Facility use Chirp as their fundamental building block, but provide different services and interfaces appropriate to their target communities.

DOI: 10.4018/978-1-61520-971-2.ch006

INTRODUCTION

It is not enough to have raw hardware for data intensive computing. System software is also needed to manage the system and make it accessible to users. If data is scattered all over the disks of a cluster, then it must be tracked, replicated, and periodically validated to ensure survival in the face of failures. Programs that execute on the system must be able to locate the relevant data and preferably execute on the same node where it is located. Users must have a reasonably simple interface for direct data access as well as for computation on the data itself.

Over the last five years, we have gained experience in designing, building, and operating data intensive clusters at the University of Notre Dame. Working closely with domain experts in physics, biometrics, and bioinformatics, we have created several novel systems that collect data captured by digital instruments, make it easy to search and access, and provide high level facilities for processing the data in ways that were not previously possible. This chapter outlines our experience with each of these repositories.

Chirp (Thain, Moretti, & Hemmes, 2009) is our building block for clustered data intensive scientific computing. Chirp was originally designed as a lightweight file server for grid computing. It was first used as a "personal" file server that could be easily deployed on a user's home machine or temporarily on a computing grid to allow a remotely executing application to access its data. However, we quickly discovered that the properties that made Chirp suitable for personal use – rapid deployment, flexible security, and resource management – also made it very effective for building storage clusters. By tying together multiple Chirp servers, we could easily build a system with very high I/O capacity using commodity storage devices.

Chirp is currently deployed on a testbed of over 300 commodity storage devices at the University of Notre Dame, in is in active use at other research institutions around the world. Using this testbed, we have constructed a number of scalable data storage and analysis systems. Each uses Chirp as its fundamental building block, but provides different services and interfaces appropriate to the target community:

- The **GRAND Data Analysis Grid** provides a scalable archive for data produced by the GRAND cosmic ray experiment at Notre Dame. It provides a conventional filesystem interface for direct access, a shell-like capability for parallel processing, and a web portal for high level data exploration.
- The **Biometrics Research Grid** archives all of the photographic and video data collected in the lab by the Computer Vision Research Lab at Notre Dame. It provides a combination database-filesystem interface for batch processing, several abstractions for high level experimental work, and a web portal for data exploration.
- The **Biocompute** facility provides a web interface to large scale parallel bioinformatics applications. Large input databases are obtained from national repositories and local experimental facilities, and replicated across opportunistic storage devices. Users may run large queries using standard tools which are decomposed and executed across the storage cluster.

In this chapter, we will explain the architecture of each of these systems, along with our experience in constructing, deploying, and using each one. We begin by describing the fundamental building block, the Chirp file server.

THE CHIRP FILE SERVER

Chirp is a practical global filesystem designed for cluster and grid computing. An overview of the

Related Content