TEXAS TECH UNIVERSITY
Edward E. Whitacre Jr.
College *of* Engineering
Computer Science

Data-Intensive Scalable Computing Laboratory (DISCL)

Technical Report

Department of Computer Science

Texas Tech University

Dynamic Active Storage for High Performance I/O

Chao Chen and Yong Chen

chao.chen@ttu.edu, yong.chen@ttu.edu

12/10/2011

Technical Report № TTU/DISCL-2011-04

http://discl.cs.ttu.edu

Dynamic Active Storage for High Performance I/O

Chao Chen and Yong Chen

Department of Computer Science, Texas Tech University, Lubbock, TX
chao.chen@ttu.edu, yong.chen@ttu.edu

Abstract

Many High-End Computing applications in critical areas of science and technology are becoming more and more data intensive. These applications transfer large amounts of data from storage nodes to compute nodes for processing, which is costly and bandwidth consuming and often dominates the applications run time. Active storage provides a promising solution for these applications by moving appropriate processing tasks from processing nodes to storage nodes. The prior research has achieved considerable progress and developed several active storage models. However, the existing studies have neglected that it is not optimal to offload every operation to storage nodes, especially when data dependence exists among operations. Such data dependence often exists in applications, such as the Geographic Information System and medical image processing. In such situation, it often wastes bandwidth for the conventional active storage architectures, because it is needed to transfer dependent data sets among storage servers. In order to resolve this problem, this paper presents a novel Dynamic Active Storage (DAS) system that analyzes the bandwidth requirement of an operation and then determines whether it is suitable for offloading to storage servers. Furthermore, based on the analysis of the bandwidth requirement, this paper proposes a data layout strategy for the DAS system to reduce the bandwidth requirement. Experimental tests have been conducted, and the results have confirmed that this new active storage architecture outperforms existing architectures. It significantly reduces the data movement caused by data

dependence and improves applications performance over existing strategies. It has a great potential for high performance I/O in high-end computing.

Keywords: dynamic active storage; active storage; high performance computing; data intensive computing; parallel I/O; parallel file systems