

# PROGRESS REPORT

**DOE Award number:** DE-SC0012115TDD

**Name of recipient:** University of Massachusetts, Lowell

**Project Title:** PROPER: Parallel Resource-Optimized Provisioning of End-to-End Requests

**Principal investigator:** Vinod Vokkarane

**Date of Report:** Sep 15th, 2014

**Period covered by the report:** June 15, 2014 - September 15, 2014

## INTRODUCTION

We intend to implement a Parallel Resource-Optimized Provisioning of End-to-End Requests (PROPER) Framework that uses the full potential of parallel architectures for reliable and flexible Big Data storage and retrieval. In the PROPER project, specific problems to be addressed include 1) Parallel Data Transfers over Multipath Circuits, 2) Reverse Multicast Data Retrieval, 3) Anycast, Multicast, and Reverse Multicast Survivability, 4) Enhancements to the Berkeley Storage Manager (BeStMan), and 5) What-If Co-Scheduler development. In what follows, we outline the progress in this quarter for categories 1), 2), 3), 4), and 5) (Year 1 deliverables).

## ACTIVITIES

The PROPER kick-off meeting was held via teleconference on the 11th of August, 2014, involving Vinod Vokkarane (PI), Chin Gouk (ESnet), Alex Sim (Lawrence Berkeley National Laboratory), Jeremy Plante and Dylan Davis (graduate research assistants). Discussion focused on work completed by the research assistants, future goals and deliverables, and access to documentation on BeStMan to aid students working on PROPER.

## GOALS

In this section we describe the progress and accomplishments in each of the categories (labeled A, B, C, D, E) as outlined in the project proposal:

- **A: Parallel Data Transfers over Multipath Circuits**

A multipath transfer service allows applications to exceed the maximum bandwidth along a single network path for both parallel storage and retrieval across the network. Development work of a multipath reservation service has begun with the creation of a client application that is mostly independent from the core of ESnet's On-Demand Secure Circuit Advance Reservation Systems (OSCARS). The client groups several independent unicast subrequests together logically for unit operations. Tasks within this category include:

- 1: Enhance the initial Multipath service to allow for heterogenous transmission rates along disjoint paths.
- 2: Support querying for bandwidth availability along (possibly) disjoint paths to allow Multipath requests to better utilize resources across the network.
- 3: Combine Multipath support with existing Anycast and Multicast OSCARS implementations.

**Work Performed:** A Multipath client for homogenous parallel transfers has been implemented to interface with OSCARS for both unicast and anycast scenarios. The work completed so far has been written to mimic the default OSCARS API as closely as possible. Item 3 from this category's tasks has been partially completed, with work remaining on Items 1 and 2.

- B: Reverse Multicast Data Retrieval

Parallelism will be provided for data retrieval through new retrieval paradigms supported at the network level, allowing a single sight to retrieve separate files from multiple data centers in parallel. We propose a reverse multicast data retrieval approach to handle this parallelism. Tasks include:

- 1: Investigate various algorithms and heuristic approaches to support the reverse anycast retrieval paradigm.
- 2: Investigate various algorithms and heuristic approaches to support the reverse multicast retrieval paradigm.
- 3: Demonstrate an end-to-end reverse anycast retrieval service on OSCARS.
- 4: Incorporate reverse multicast with the proposed multipath OSCARS client.

**Work Performed:** No implementations have been attempted in this category, focus has been on investigation for now. The groundwork for Item 4 has been completed with the implementation of the proposed multipath OSCARS client.

- C: Survivability

One of the fundamental features essential to reliable data transfer on a high-speed network is survivability against both link-failures and node-failures. The PROPER framework will support survivability for both storage and retrieval, and will be integrated with our proposed network services. This approach will provide not only flexible data transfers and efficient use of parallelism in the network, but also increase resiliency against network faults. Tasks include:

- 1: Provide anycast protection in OSCARS.
- 2: Provide multicast protection in OSCARS.
- 3: Provide reverse multicast protection in OSCARS.

**Work Performed:** Item 1 has been completed with the implementation of a Multipath client that interfaces with OSCARS. providing survivability for both anycast and unicast scenarios through the use of link-disjoint backup paths. Items 2 and 3 remain in the preliminary stages of research and implementation.

- D: BeStMan

BeStMan will be used storage management in the PROPER framework, and its functionality will be integrated over its current SRM web services interfaces. BeStMan relies on the underlying storage system and file system for storage contents, and performance in browsing the storage contents in its managed storage space depends on the performance of that underlying system. We plan to extend the browsing capability of the BeStMan through the completion of the following tasks:

- 1: Extend the browsing capability of the BeStMan for its managed storage space.
- 2: Implement dynamic cataloging of file-system contents including database design and caching implementation.
- 3: Develop queuing algorithms for meta-data queries.

**Work Performed:** Work on these tasks are in the beginning stages. Students working on PROPER have received documentation detailing the functionality of BeStMan so that research can begin in earnest.

- E: What-If Co-Scheduler

Due to pre-existing requests and a lack of available bandwidth on certain links, users intending to provision network resources to carry their data and satisfy their requests may not be able to successfully provision a circuit for the desired period of time or within time constraints. We have designed and constructed the What-If extension to OSCARS, which is an offline option-based provisioning tool that can reduce the number of provisioning attempts a user may have to make. This extension can be fully incorporated into the PROPER framework, not just to find available resources given time/bandwidth constraints, but may also be generalized to find combinations of suitable available resources and services based on user/application requirements. Tasks for incorporating What-If as a service include:

- 1: Develop interfaces for the OSCARS What-If scheduler to BeStMan.
- 2: Develop and implement anycast and multicast service negotiation within What-If.
- 3: Develop and implement a reverse-multicast service negotiation within What-If.

**Work Performed:** As one of the longer range goals of PROPER, this work has not progressed beyond initial research.

#### COST STATUS & UNEXPECTED FUNDS

See attached document

#### NEXT QUARTER DELIVERABLES

- Task A-1. Completed implementation of heterogenous transmission rates along disjoint paths extension to the Multipath OSCARS front-end client.
- Task B-1. Completion of investigation into reverse anycast retrieval approaches.
- Task C-2. Completion of multicast protection extension to provide multicast survivability in OSCARS<sub>i</sub>