Recent and Upcoming Developments in the Tranche Data Repository

Computational Approaches and Proteomic Datasets

C552 C501-C553 HUPO 2009

Bryan Smith; James Hill; Mark Gjukich; Philip Andrews
University of Michigan, Ann Arbor, Michigan

Introduction

Ongoing development of the Tranche Distributed Repository (www.trancheproject.org) deals with issues of scalability, reliability, performance and security to address the growing community of users.

Tranche combines some of the strengths of a centralized, client-server model with a distributed architecture, but likewise must accommodate the challenges associated with both. Here we discuss these developments in the context of the related challenges.

What is Tranche?

The Tranche Distributed Repository is designed to meet the data sharing and public dissemination requirements of proteomics research.

Any number of Tranche networks can be created, offering separate maintenance options and security. The ProteomeCommons.org Tranche network is the first and primary Tranche network and is available at ProteomeCommons.org. (tranche.proteomecommons.org)

Registered users can upload data to the ProteomeCommons.org Tranche network, as well as manage and annotate their data sets through ProteomeCommons.org. (proteomecommons.org/signup.jsp) Data sets can be optionally encrypted. Anyone can download without registering.

More info at: trancheproject.org/about.jsp

How Does Tranche Work?

When a data set is uploaded to a Tranche network, each file is broken up into data chunks with a maximum size of 1MB. These data chunks in a file are described by a single metadata chunk, which is sufficient to locate and reassemble a file. All the files in a dataset are identified by a ProjectFile.

To download a dataset, the metadata for the ProjectFile is retrieved and used to download and assemble the ProjectFile. In turn, the ProjectFile is used to locate and download the metadata for each file.

The classic Tranche network involves separate servers, each housing a portion of data and metadata chunks. Since redundancy is essential for reliability in a distributed network, clients upload multiple copies to the appropriate servers. This requires that the client is connected to all (or at least most) servers in the network.

Additional server would result in more required client connections; however, there are operating system-specific limits to the number of connections that a Tranche client can make. A new network model is required to address this and related performance and security concerns.

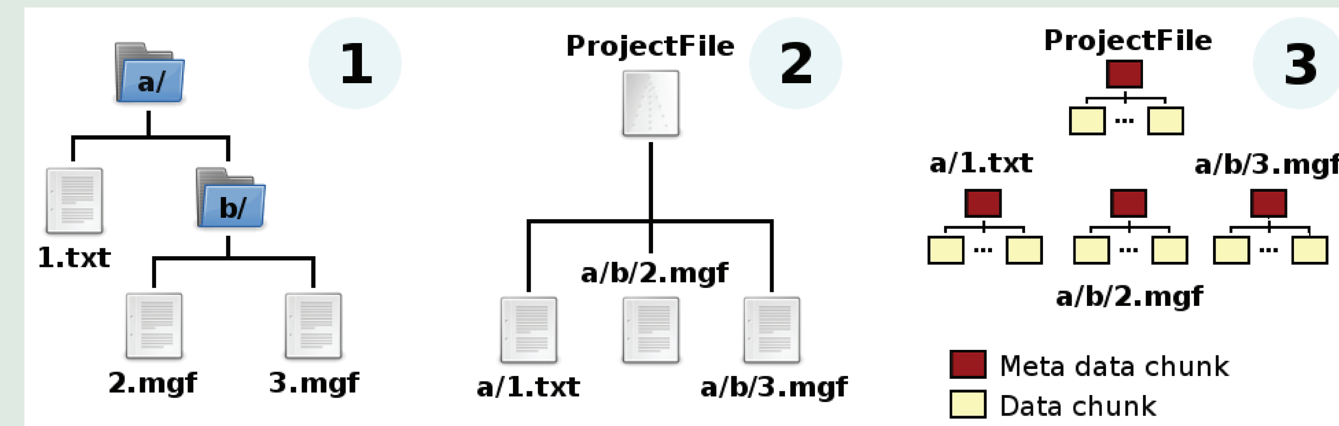

Figure 1. (1) A data set model is simply a connection of files. (2) When stored on a Tranche Network, a ProjectFile is created to describe the files and how to retrieve them. (3) Each file (as well as the ProjectFile) is made up of a meta data chunk and one or more data chunks.

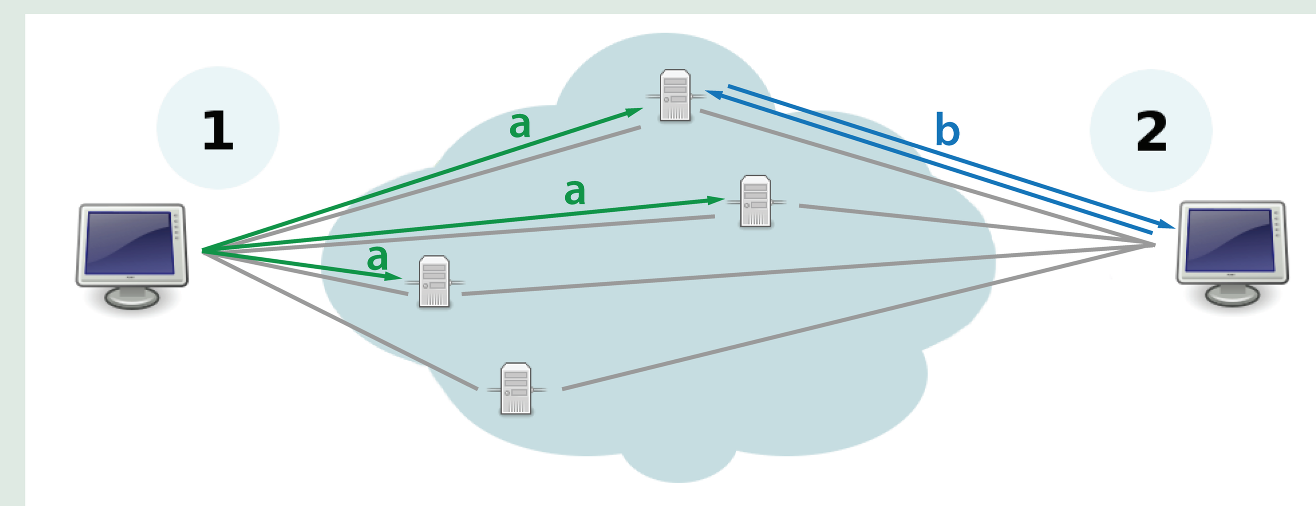

Figure 2. Old server model. Note that each client connects to every server (grey lines). Client uploads the required replications.

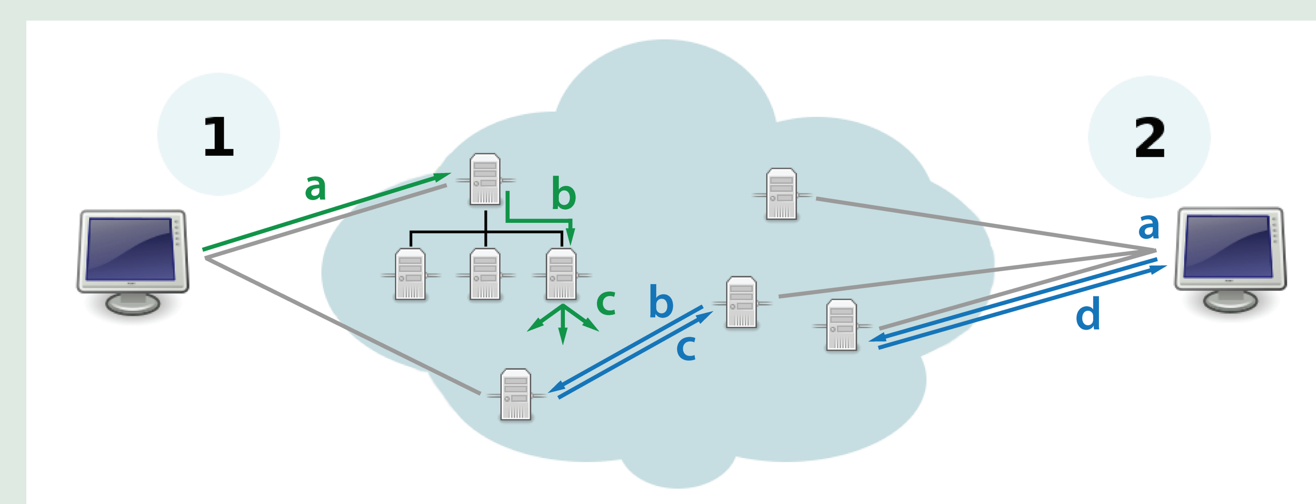

Figure 3. New server model. Note that each client connects to a limited number of servers (grey lines). Client 1 uploads chunk to routing server (a), which passes it to another server (b), which replicates it to the network (c). Client downloads a chunk (a-d).

Methods

Tranche is a highly-configurable, Java-based application currently used by ProteomeCommons.org (tranche.proteomecommons.org). Continuous development is directed by user feedback and well as testing (unit, stress and user), while major revisions are developed to address issues we can identify.

Challenges

- Scaling to petabyte capacities while minimizing connections
- Compensating for inevitable hardware failures
- Detecting and repairing corrupt data
- Offline servers, either temporarily (e.g., power outage) or indefinitely
- Managing security risks, even if compromised
- Network changes, such as new servers or adjusted server configurations

Results

Current statistics for the Tranche Distributed Data Repository as of September 17, 2009.

Servers
16 geographically-distributed servers

Registered Users
391 registered users

Data Sets
8,093 data sets

Total TB Online
10.7 TB of data (uncompressed)

Files
12,636,826 files

Recent Developments

- Servers devote their maintenance time to downloading new data and deleting old data (actively balancing), as well as checking for (and repairing) corrupted data and balancing local database b-tree
- Data is verified on clients and servers when uploaded and downloaded to check for integrity
- 'Shadow' Tranche network (with different trusted X.509 certificates) actively backup ProteomeCommons.org Tranche network
- Recent performance benefits from batching multiple requests as well as having servers 'compete' for a request

Upcoming Developments

- Multiple servers be managed by single 'routing' server
- Clients will connect to fewer servers
- Servers will have degrees of trust so can add data to each other
- Network will 'lazily' balance on demand by user requests
- Actions on network are logged and queried by servers coming online to check for deleted or updated content
- Server maintenance schedule more intelligent based on recent performance

Innovative Aspects

- Redundancy re-introduced (lazily and actively)
- Servers assist each other in repairing corrupted data using a degree of trust combined with autonomy
- Efficient logging of significant network activity
- Demonstrates data pedigree and integrity

Poster PDF and Additional Information:
The ProteomeCommons.org and Tranche Repository community resources are primarily sponsored by the NCRR grant #P41-RR018627 and the NCI subcontract #27XS115. Please contact markgj@umich.edu for any questions or comments regarding these projects.