LEGS: A WSRF Service to Estimate Latency between Arbitrary Hosts on the Internet

R.Vijayprasanth[1], R. Kavithaa[2,3] and

Raj Kettimuthu[2,3]

[1]Coimbatore Institute of Technology, India

[2]The University of Chicago, USA

[3]Argonne National Laboratory, USA

Motivation

- Why do we need to estimate latency?
  - Enable distributed computing applications to find the closest server / replica of data needed for computation.
  - Internet content providers - often place data and server mirrors to improve access latency for clients - use this as a metric to find the appropriate mirror.
  - Useful for constructing topologically sensitive overlay networks.

WSRF

- A web service - Internet hosted application that is described via Web Services Description Language (WSDL) and is accessible via standard network protocols such as SOAP over HTTP.
- Web services must often provide their users with the ability to access and manipulate state
  - i.e., data values that persist across, and evolve as a result of, Web service interactions.

WSRF

- Grid service is a Web service that conforms to a set of conventions for purposes - lifetime management, inspection, and notification of service state changes.
- Grid services provide for the controlled management of distributed and often long-lived state that is commonly required in distributed applications.
- WSRF provides standardized specifications to perform the above-mentioned tasks.

LEGS

- LEGS - Latency Estimation Grid Service
- Grid service implemented on top of the Globus toolkit
- The Globus Toolkit is a software toolkit, developed by The Globus Alliance which can be used to program grid-based applications.

LEGS

- The Globus Toolkit 4 includes a complete implementation of the WSRF specification.
- We use the built-in security and the monitoring and discovery services provided by Globus.
- Web service - generic interface - enables users to develop clients easily

Latency Estimation Tools

- Ping, Synack - estimate end-to-end delay - need to be run on one of the two nodes between which the latency needs to be estimated
- King, IDMaps, GNP - calculate delay between arbitrary hosts on the internet
- IDMaps and GNP are better in terms of Speed of estimation and the estimation overhead, LEGS use King because of its ease of use.

King

- From any node on the Internet, measure latency between arbitrary hosts on Internet
- No additional infrastructure needed on end hosts
- Estimate latency between the domain name servers
- Claim ~75% of DNS servers support recursive queries from any host
- Assume name servers are located close to their hosts

Experimental Results

Src	Dest	RTT	L1	L2
ANL	OSU	10ms	5.2s	4.2s
ANL	USC	60ms	5.7s	4.2s
Maine	SDSC	105ms	6.0s	4.2s

Experimental Results

- A useful general service for the Internet - enable a host to quickly learn the distance between any two hosts.
- Such a service should provide an answer with a delay overhead less than the speed-up gained using the service.
- It takes a minimum of 4.2s to find the latency
  - Security exchanges and SOAP processing
  - May not work well for real-time applications

GridFTP

- Extends standard FTP protocol to provide a lot of important features
  - Striped data transfer (cluster to cluster)
  - Partial file transfer
  - Reliable and restartable data transfer
  - Data channel caching
  - Supports Grid Security Infrastructure
  - Setting of TCP buffer sizes

Optimizations

- TCP is the default transport protocol used by GridFTP
- Window-based flow control
- It is critical to use optimal socket buffer sizes to get maximum throughput
  - Bandwidth-delay product
- Optimal buffer size = 2*bandwidth*delay
- LEGS would help with estimating the delay here