

Reliable Multicast for the Grid

Karl Jeacle & Jon Crowcroft


firstname.lastname@cl.cam.ac.uk

Key Results

TCP-XM: A modified implementation of the TCP protocol that allows reliable multicast data transport in addition to the existing unicast transport mechanism.

How does the work advance the state-of-the-art?

Allows Grid applications to transparently mix unicast and multicast data transmission thus increasing network efficiency.



Motivation (Problems addressed)


Large multi-site data transfers are commonplace in a Grid environment, but network resources can be wasted by not taking advantage of the efficiency provided by multicast transport mechanisms.

Multicast TCP

Applications use TCP for reliable unicast transmission. By modifying TCP to deliver data to more than one receiver, an application can transparently send data reliably to multiple recipients. By always attempting multicast transmission first, but falling back to unicast when necessary, existing TCP mechanisms ensure that data is always delivered reliably. We call this new protocol TCP-XM.

In order to test and deploy this new protocol on live networks, minimum requirements must be placed on end stations i.e. no root privileges or kernel changes.

Userspace applications can freely send and receive UDP packets; a small shim layer encapsulates our TCP engine's packets into UDP. There are performance implications with userspace, but the instant deployment potential, coupled with the scalability of multicast, means that these limitations are quite acceptable.





Globus and The Grid

Globus XIO is an eXtensible Input/Output library for the Globus Toolkit. It provides a POSIX-like API to swappable I/O implementations – essentially “I/O plugins” for Globus. We have implemented a multicast transport driver for Globus XIO that uses the TCP-XM protocol.

Experimental Results

Transfers using the new protocol are of equivalent speed to transfers using unicast alone. But with unicast, network load increases linearly with every new destination host; multicasting requires just one set of data segments to be transmitted regardless of the number of destination sites.



Summary

TCP-XM is a modified version of TCP that supports multicast. By running in userspace above UDP, we can test a new protocol in live networks, while delivering reliable multicast to the Grid community.