Provisional Thesis Outline

Version 1.0

29th July 2002

Abstract

Broad introduction into the subject of the thesis
Include concatenation of the thesis into a thesis statement – 'Existing networking technologies allow far greater performance than what it currently in use today'…? [for Grid performance etc]

Background - Requirement of Data Intensive Applications / Real-Time?

Promises of Real Time Application and Distributed Computing
Why we need to send data around quickly – include examples of projects such as astrogrid, tele-immersion etc.
Need details of the amounts of predicted data flows
Overview of the Grid and how it allows this
How the grid community (principally Europe? EDG and include US?) plan to use the Grid, include details of transfer requirements, processing requirements and raw data output. Both now and in the future
Focus on HEP application (LHC/CDF/SLAC)

- Why they send data around
- The way they send data – include protocols and programs
- How the Grid promises to make it easier

Overview of problems/possible improvements wrt Grid and HEP

- Replication issues?
- Existing networks and their limits

Performance Overview

What constitutes 'performance' - Define performance upon criteria of reliability, link utilization usage, use of available technologies (adaptability) and security.
Hierarchical Schematic of the components that affect network performance (ie the Grid)
Table/Section on performance issues

- Networks
- Hard disks
- Firewalls
- CPU
- Memory
- NICs

Fact that UKERNA network is underutilised – 10Gb! But only getting 200mb/s…
Do same test with higher performance machines… still get only about 700mbit/sec….
Identify that fast machines still only achieve a fraction of possible bandwidth
Problem statement of Thesis: 'existing infrastructure allows high speed transport, but is hampered by transport protocols. This thesis looks into the

state of the art technologies in data transport protocols for high speed Grid access'.

State of the Art - Analysis of Internet Performance

This section will give a broad overview of the work in progress in each field, followed by detailed personal work towards use and or implementation.
Maintain view of Problem Statement as an outline.
Problem today of HEP networks
Demonstration of Problem of existing networks – real transfer of real data with existing infrastructure
Identify areas

Network Interface Controllers

Detailed analysis of the performance of a range of gig nics – care of the network group and Nick
Define performance criteria for cards and repeat for all cards on the same test system – Dell machines
Latency vs. packet sizes for all protocols (ICMP, UDP, TCP)
Throughputs (UDP, TCP)
Vary the interrupt times of driver

Hard Disk (Host performance) and Transfer Protocols

Overview of hard disk technology
Analysis of disk speeds of computers around the world
Analysis of the affects of differing transfer protocols (ftp, bbftp, bbcp, gridftp, iperf)

TCP

Overview of TCP/IP stack
Describe what it is and does
Overview of the different flavours (history)

TCP Tuning

The effects of autotuning (theory and practical)
Monitoring links
Include memory/performance advantages of autotuning (esp for servers)

Multiple stream TCP

Performance advantages
Fairness argument
MulTCP

Test of Existing Infrastructure

Transfer of large datasets at high speed and long times
Requirements: eg large enough harddisk, large enough data set and OS (eg linux has 2gig limit on files)
Conduct tests to and from SLAC and CERN and Man using large files – gridftp if poss.
Demonstration of throughputs with tuned, untuned, multiple stream, single stream.
Real data! Using hard disks etc. Process data as well.. (but not concerned with)

ECN and Congestion Control

What it is
Work in field? Not TCP…

What are the requirements from the network
Mention overview of general mechanisms
Applications towards UDP and better TCPs

CPU and Transport Protocol Performance

Vary cpu loads and see result of protocol proformance in terms of loss, data/packet rates…
Icmp, udp and tcp
What can we do with results?

Summary

Identify optimal hardware and software using existing technology for use in HEP applications
Demonstration of how simple performance issues can be overcome with little knowledge
Network technicians can increase performance for particular situation/problem. - Need General solution to problems
Direction of thesis as result

Network Monitoring

Describe network monitoring and why network monitoring is important
Case study: How PingER has aided the HEP community
Introduce next generation of NM's: GridNM/IEPM-BW/GGF stuff (latter being architecture etc)
Use of infrastructures to analyse networks
Overview of existing network monitoring architectures
Continual analysis of the network vs short
Possibility of prediction engines and references
Identify the problem areas (no cooperation, no uniformity…)

GGF

Discussion of metrics and methodologies (GGF, IPPM)
Metric Schemas (with Warren)
How this all fits in with Grid and OGSA

Network Tools to take different metrics

Overview of most used (iperf, pathload, ping etc) and comparison of similar tools over same links
Focus on accuracy and intrusiveness of tools
Use of wrappers and web services to allow cross functionality and unified information

Network Monitoring Infrastructures – NG

Unification of programs using web services
Need of common schema for language – eg characteristics and observations
Mention Internet2 and GGF

Systematic Analysis of Network Performance

A set of problem scenarios and the causes and solutions
Analysis of Problem/Comparisons
Solution of Problem
Include stuff about different performance criteria: hard disk, memory, cpu…
With applications

Transport Protocol Comparison and Analysis

Solution of Problems using different protocols
Idealistic (ie like iperf) Performance Analysis of different transport protocols for network (fake) data transport
Demonstration of performance around UK, CERN and to SLAC

TCP

Overview of TCP technology and terminology
Include performance analysis of end hosts and LAN networks using existing tcps

High Speed TCP

Overview of differences to normal tcp and (dis)advantages
Adaptation of TCP for Sally Floyd's algorithm
Test results

Ravot TCP

Overview of differences
Test results

Fast TCP

Overview of differences
Test results

UDP

Benefits of UDP transport
Disadvantages of UDP transport (aka congestion control and fairness)

Tsumani

Overview of methods and application
Test results

Bill Allcock

Overview
Test results

UDP Blast

Overview
Test results

Multicast

What multicast is and what it requires from the network
Mention existing JANET and GEANT networks compatible and also ABILENE
Overview of the performance gains of multicast
Define performance analysis… not so easy….
Tests with people in SLAC and GEANT

Reliable transfer Protocol (java stuff – wp7)

PGM/CC

Application of Protocols to Real Life performance

Use of transport protocols for real data transport of files etc.
Try out and compare applicability
Mention GridFTP
Define schema for protocols? – defines when and how to use protocols
Small step may be to find out how many tcp streams to use…
Utility for choosing parameters without requiring expert knowledge
Prototype

QoS

How QoS fits into picture
Guaranteed bandwidth means that we should make use of it
Existing protocols do not allow this (udp throttles, tcp is too conservative)
MB-NG Stuff

Conclusion and Future Work

Appendix