

Research Infrastructure



- Community-wide infrastructure support
 - HEPnet
 - HEPgrid (Grid Canada)
- Local infrastructure
 - MFA support for detector development
 - Computing infrastructure

HEPnet/Canada



Supports and funds wide area networking for SAP

- national links
 - SNO – Laurentian (50%)
- international links
 - now: CERN – Chicago (8% of N.A. cost)
 - in past: Canada – NSFnet, Canada – ESnet
- network monitoring and applications
 - high speed data transfer trials (CERN – Canada)
- dialup service



HEPnet/Canada people

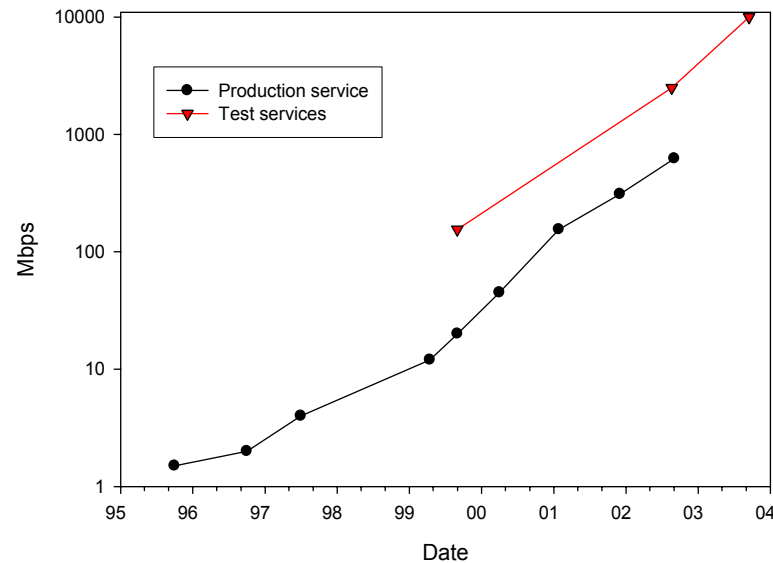


- Director – Dean Karlen
 - Manager – David Bickle
 - new hire: BSc. (CS), HPCVL & Bell Canada
 - previous manager, Wade Hong, now full time at Carleton University, remains active
 - Advisory committee
 - Richard Keeler
 - Renee Poutissou
 - John Martin
 - Affiliation with HEPgrid/Canada:
 - Randy Sobie & Bob Kowalewski
-

CERN – North America link



US - CERN link speed history



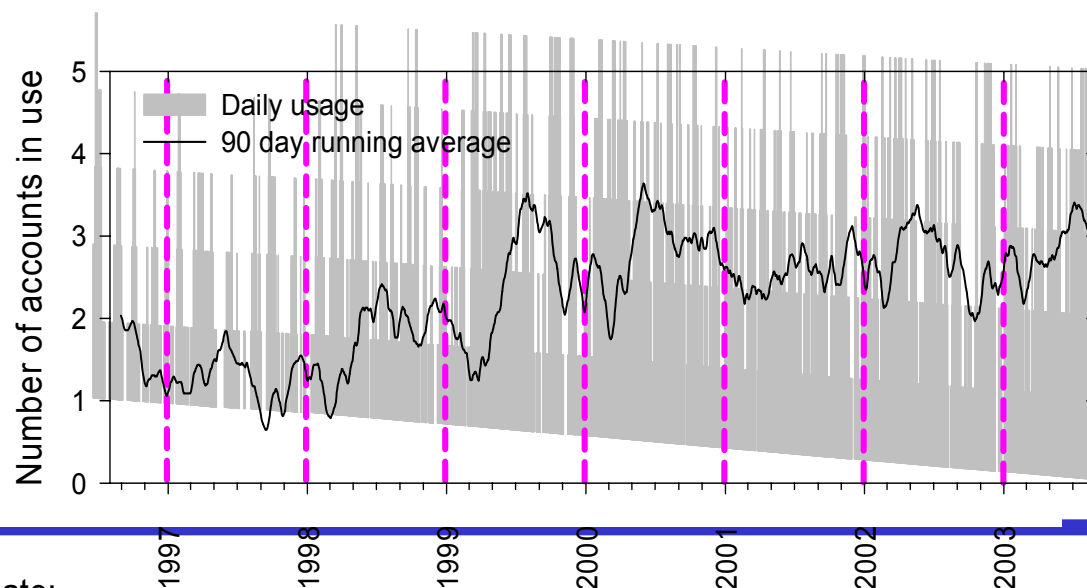
- “US line consortium” active since 1995
- frequent upgrades to reach with LHC requirements by 2007
- HEPnet/Canada support is 8% of N.A. costs

Dialup service



- for use by SAPHysicists when travelling
- local dialup access numbers in >1000 cities
- 660 periods reserved in past 7 years:

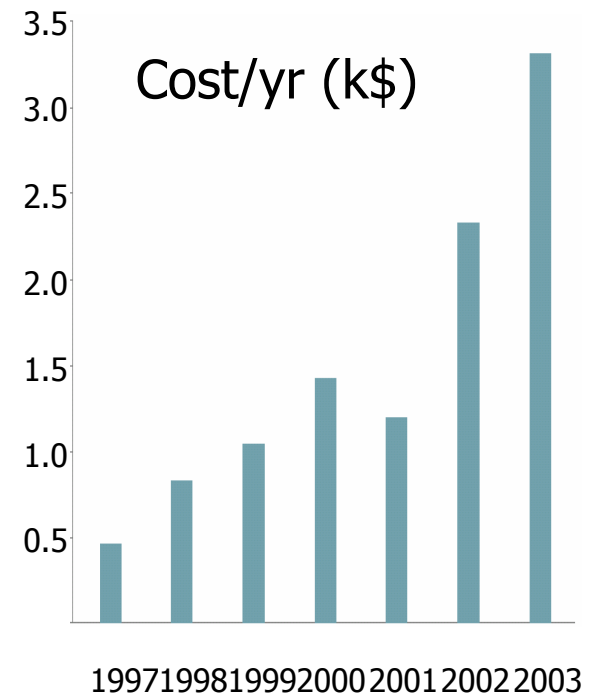

HEPnet/Canada Dialup Account Usage (Aug 1996 - Sep 2003)



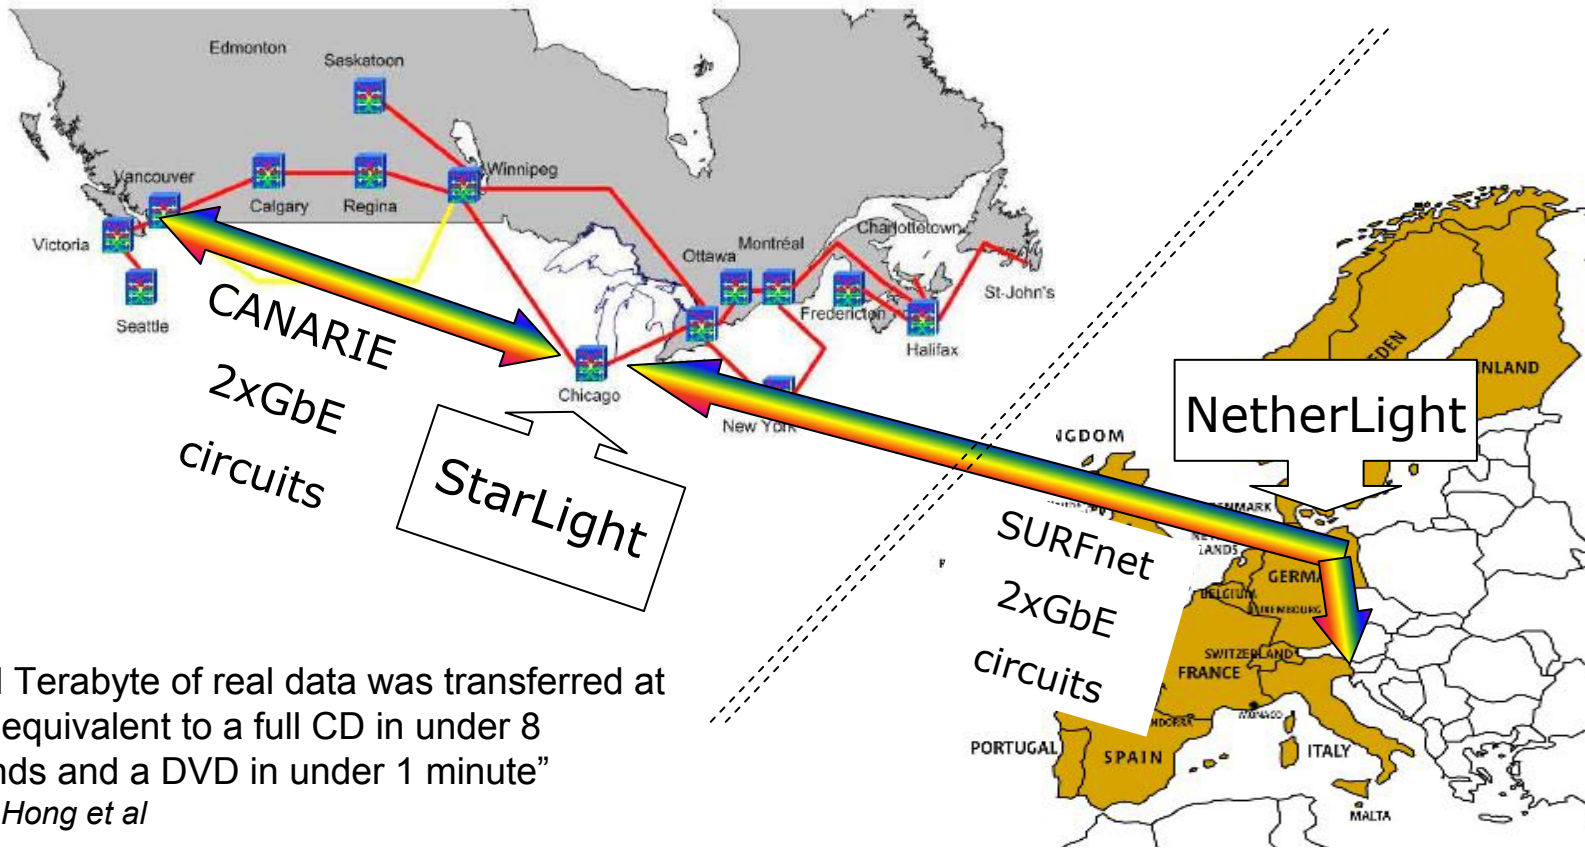
Dialup service



- usage is spread amongst several Universities
- usage fees are increasing...
- new system for requesting accounts in development



ATLAS Canada Lightpath trial TRIUMF CERN



“A full Terabyte of real data was transferred at rates equivalent to a full CD in under 8 seconds and a DVD in under 1 minute”

Wade Hong et al

High bandwidth tests



- A new demonstration/test is in preparation, in association with Telecom 2003 (Oct 12-18).
 - 10 Gbps line between CERN & Carleton U.
 - packet generators/ iPerf memory transfers / disk to disk transfers (up to several Gbps)
 - will be very useful tests of equipment/software to be used in the LHC era
 - A Canada-wide testbed for high bandwidth testing is being proposed, to be funded by Canarie
 - led by Wade Hong
 - primarily to benefit HEP
-

HEPnet Summary



- High performance networking is a very dynamic field
 - Our field will benefit from the close ties developed with CANARIE
 - Available manpower is increasing
 - support from Wade Hong
 - new full time manager, David Bickle
 - Expanded role:
 - merge network and grid support activities
 - single MFA for HEPnet and HEPgrid in the future...
-

Grid Computing



- The Grid is the solution for ATLAS computing
 - ATLAS-Canada model has TRIUMF as the link into the LHC Grid with the University sites forming a Canadian Grid
- The Grid has the potential to open up resources (manpower and hardware) not normally available to HEP
- UVictoria group active in Grid Computing since 2001
 - R.Sobie, R.Kowalewski, A.Agarwal + coop students
 - HEPgrid MFA 2 year funding for coop student support



Grid Canada

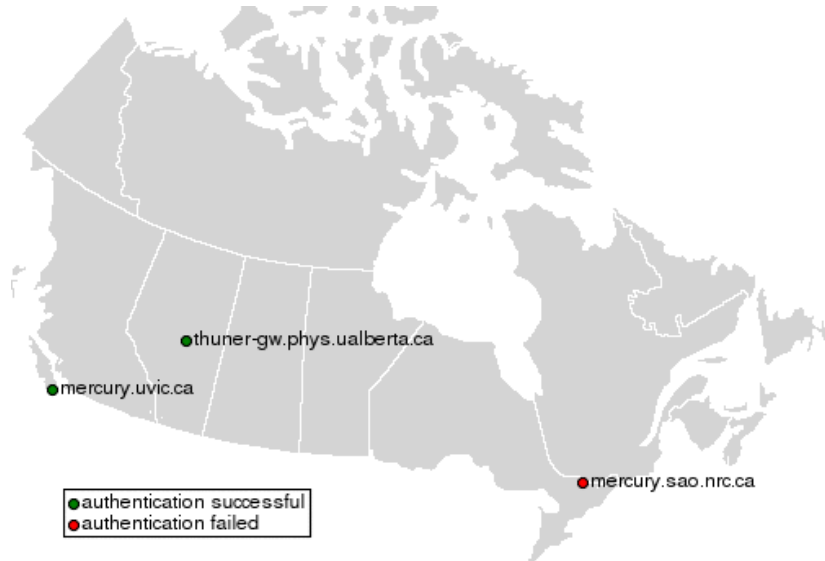


- Collaboration with NRC Institute for Information Technology (Ottawa), CANARIE, Alberta, Pacific Forestry Centre
- GC is responsible for Canadian Certificate Authority
- Established a Grid Testbed (2001-2003) using ~50 computers at 12 sites across Canada
 - Ran BaBar MC simulation and ATLAS DC
 - Results presented at the CHEP Conference (2003)

Canada-wide Grid



- Creating a Grid using “production” facilities at Victoria (Mercury), Alberta and NRC (~300 CPUs)



Operational Oct 2003

Large scale test using
ATLAS DC software
prior to CANARIE
meeting

GC Status



- Grid enabled 3 clusters (globus)
 - Involves sys admins who are not part of GC
 - Common accounts, gc-production queue
 - Site monitoring (ganglia)
 - “gcsb” batch submission and resource broker
- Short term goal
 - Working Grid by Oct 22 CANARIE WORKSHOP
 - Running ATLAS DC
- Longer term
 - Add additional sites
 - Apply to NSERC Strategic Grant (NRC, Eng, Forestry)

Training of HQP



- Tremendous interest in learning about the Grid from students
 - 40 applications per term from outstanding students in physics, engineering physics, computer and electrical engineering from SFU, UBC and Victoria
- Physics and Engineer departments at the University of Victoria are jointly supervising students on Grid projects
- 2003 IEEE Gold Medal – D.Vanderster
 - Computer Engineering
 - MSc on Grid Computing



UVic MFA



- MFA to the University of Victoria group
 - Support Detector Development
 - ATLAS and Linear Collider
 - Software Engineering Support
 - All experiments
 - Application-specific software and data installation and maintenance
-

Detector technologist: Poffenberger



- *Paul Poffenberger*
 - many years of experience in providing technical support to the experimental groups at the University of Victoria
 - FY2003 50/50 (ATLAS MIG), FY2004 100% on MFA
 - Current activities
 - ATLAS feedthrough project:
 - managed day to day activities in laboratory
 - successful completion of project on time & on budget
 - now overseeing installation of completed feedthroughs at CERN
-

Detector technologist: Poffenberger



- Current activities (cont.)
 - Linear collider TPC project:
 - detector modifications
 - data acquisition system development
 - setup and operations of TPC at TRIUMF and DESY
- Future activities
 - Linear collider TPC project:
 - micromegas detector and test beam preparations
 - JPARC-SK long baseline neutrino program:
 - “near detector” design studies
 - detector construction activities

Software Engineer: Agarwal



- *Ashok Agarwal*
 - several years of software support to the Babar group
 - FY 2003 – FY 2004, 50 / 50 (Babar operating)
 - Current activities:
 - Babar experiment:
 - MC production, data management and software support
 - ATLAS experiment:
 - Grid Computing, Data Challenge, cluster software support for test beam and Monte Carlo studies
 - Linear Collider project:
 - Cluster software support for cosmic ray and MC studies
-

Local Resources



- Desktop
 - Linux/Windows workstations and laptops
 - Muse Linux Cluster
 - Commodity Intel processors and inexpensive disk
 - Mercury Linux Cluster (470th in TOP500)
 - IBM Intel processors connected to SCSI disks via FC
 - Storage Facility
 - 200 TB disk/tape system designed for moving data fast
-

Muse Cluster



- Commodity Intel cluster (Linux)
 - Funded by Startup Funds, CFI and NSERC Eqpt grants
 - Interactive users from all experiments
 - BaBar MC production
- Very successful - 3rd stage of “renewal”
 - Well managed by Jan van Uytven
- Manpower intensive and I/O limitations
 - Trading low-priced hardware for manpower
 - I/O limited - BaBar NTUPLE analysis



CFI Funded Resources



- UVictoria Research Computing Facility
 - 1999 \$2M CFI/BC for a parallel processor (N.Dimopoulos)
 - 2002 \$6M CFI/BC for storage facility (R.Sobie)
- Multi-discipline facility
 - Biology, Earth and Atmospheric Sciences, Engineering and Physical Sciences
 - NRC Herzberg Institute for Astrophysics and Canadian Forest Service Pacific Forestry Centre

CASS



- Managed by UVictoria Computing Centre (CASS)
 - UVictoria allocated \$1M of its CFI IOF to the project
- 2 CASS FTE's dedicated to Research Computing
 - Colin Leavett-Brown (25 years HPC experience)
 - Drew Leske (UVic BSc with 5 years experience)
 - Access to pool of experts in Networking, Storage Management, Facilities
- Strong University support
 - VP Research and CASS

IBM



- Excellent relationship with IBM Corporation
 - IBM has provided hardware and support for many years
- IBM has been generous in its support
 - 1999 \$840K IBM SUR grant (R.Sobie)
 - Mercury Linux Cluster was donated to UVictoria
- In response UVictoria has worked with IBM
 - IBM Beta Test for LTO2 Tape library (1/2 in Canada)
 - Working with their Storage Software team (Tivoli)

Year 1 Layout



1st of 4 Year Project

Fully Operational

- 12 TB disk
- 220 LTO2 tape library
- 8-way Storage Server
- Storage Area Network with 2 Gigabit/s Fibre Channel
- 108 2.4 GHz Xeon CPU's
- 128 CPU IBM SP
- 2 Gigabit/s Ethernet to BCNET/CANARIE

Current Status



- Hardware delivered January-April 2003
 - Operational May 2003 and Grid-enabled Sept 2003
 - Integrating large storage users
 - Particle Physics, NRC HIA and Pacific Forestry
 - 10 TB data on the system
 - Mercury Linux Cluster
 - 470th fastest computer in the world (100% utilized)
 - We are demanding users
 - Hardware is excellent; software is not meant for Research
-

Challenges



- HEP software no longer operating system, compiler or software independent
 - Impediment for making a Grid
 - IBM/Tivoli Storage Manager struggling with millions of files
 - Astronomers putting in data faster than the SM can cope
 - Inject some intelligence into the SM like tarring common files or keeping frequently used NTUPLES disk resident
 - Working with IBM to help speed it up (an Is of the Astronomy disk takes 9 hours)
-



Future plans



- Fall 2003
 - Disk from 12 TB to 24 TB
 - 28 dual Xeons using BaBar NSERC Eqpt Grant
 - Link Muse cluster into the Storage Facility (2 Gbit/s ETH)
 - Social Science CFI Project (TAPOR)
- 2004, 2005 and 2006
 - Double storage and processing capacity per year?
 - CFI for new parallel processor and extension to Storage
 - MFA for system support?