

# *Cloud computing, networks, data management and preservation*

*2026 IPP Annual General Meeting*

*Randall Sobie*

*IPP and University of Victoria*

## **Experiments and projects**

UA1, SLD, OPAL, BaBar, ATLAS, Belle II

## **Administrative responsibilities**

UVic (VISPA) HEP Research Centre and engagement with the MacDonald Institute and TRIUMF

Editor-in-Chief, Elsevier journal, SoftwareX

Director, HEPNET/Canada (national and international networks)

Researcher Council and Board of Directors, Digital Research Alliance of Canada (ended 2025)

## **Research activities**

Cloud computing for Belle II, ATLAS and DUNE

Belle II (CFI) Raw Data Centre

BaBar Long Term Data Access (LTDA) Facility







Terabit/sec network projects

CPU (and GPU) benchmarks for the WLCG

Software FPGAs and detector-simulations

# Experiments

UA1, SLD, OPAL, BaBar, ATLAS, Belle II



Pion-nuclear reaction studies (TRIUMF)

Jet physics (UA1/CERN)

Tau lepton physics (OPAL/CERN, BaBar/SLAC and Belle II/KEK)

Searches for new particles (OPAL/CERN)

Searches for non-standard CPV and rare tau decays (BaBar/SLAC)

Phenomenology papers

(Neutrino mass, isospin predictions for high multiplicity tau decays and the lineshape of the rho mesons)

# UVic VISPA Research Center

## Director of the UVic VISPA (particle physics) Research Center



University  
of Victoria

Victoria Subatomic Physics  
& Accelerator Research Centre

Manage administrative, detector support and computing support personnel

Secure funding from UVic, NSERC and other organizations

New engagement with the MacDonald Institute for detector development projects

Coordinate use of laboratory and assembly space

VISPA computing infrastructure

Gateway to Alliance and experimental computing resources

Liaise with the University research administration

Visits to TRIUMF, CERN, KEK, JPAC



## Senior, Editor-in-Chief of SoftwareX journal (Elsevier)

One of three Editors-in-Chief

~10 Associate editors, journal manager and 2 production managers

Authors must submit paper and software for review

Authors need to be reminded to cite software papers

Multidisciplinary journal

Some HEP papers, more NP - submission of software makes it challenging

Submissions to the journal have grown by 50% in 2025 and 2026

Expecting 1500 submissions in 2026 with ~50% acceptance

Many more papers using AI and LLMs



## HEPNET is responsible for national and international connectivity of the Canadian SAP community

Established in 1990

Funded with an NSERC MRS award until March 2029

## Activities

Connecting Canadian sites to the LHCOPN/LHCONE networks

Belle II and DUNE are part of the LHCONE network

Migration of sites from IPv4 to IPv6 (32 to 128-bit network addresses)

### Network monitoring

Managing Performance network systems for the Alliance and helping CANARIE

### Collaborating with Canadian and international network organizations

WLCG Network Group

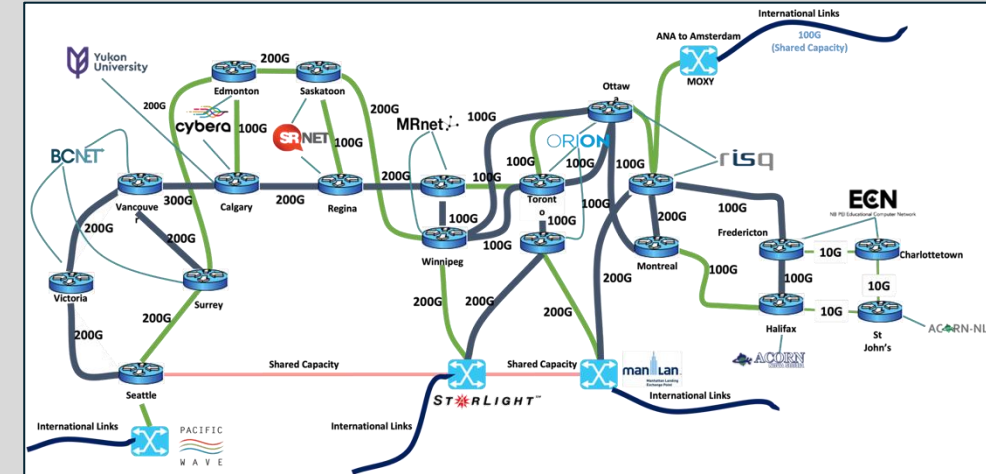
CANARIE, BCNET, Alliance, ESnet, Internet2, SciNet, Geant4, ..

### Network R&D projects

Network packet marking and terabit/second network demos

### Hosted LHC network meetings in Canada

Victoria 2024 and Montreal 2026



## Square Kilometer Array Observatory (SKAO)

Integrate SKAO into the LHCONE?

Many SKAO Centres are also LHC Centres  
Canadian SKAO Centre in Victoria

Canada does not have a trans-Pacific link  
(routed via Seattle and the US is not in SKAO)

# Chair, Alliance Researcher Council



**Digital Research  
Alliance** of Canada

## Roles over a six-year period

- |   |                   |
|---|-------------------|
| Selected to the first Researcher Council (March 2020) | - term ended 2025 |
| Chaired its first Governance Committee (2020)         |                   |
| First Chair of the Researcher Council (April 2021)    | - term ended 2024 |
| Board of Directors (September 2022)                   | - term ended 2024 |

## Key accomplishments

- Annual meeting with the Board of Directors
- RC Chair a member of the Board of Directors
- Editor of the RC document on Research Priorities documents
- Talks at public meetings and report at every Board meeting

*Firmly established the role of the Researcher Council in the Alliance*



## BaBar Long Term Data Access (LTDA) Facility

## Belle II Raw Data Center (CFI-funded project)

## Terabit/sec network demo projects

## CPU Benchmarks for the WLCG and HEP community

### **Cloud/grid computing infrastructure**

Belle II, ATLAS, DUNE production jobs using clouds in Canada, Europe, Australia and US  
Funded with a CFI Cyberinfrastructure award and CANARIE software programs  
Led the conversion of the ATLAS High Level Trigger (HLT) farm to an Openstack cloud

### **Belle II Analysis and Monte Carlo Facility**

We have a 4000-core allocation on the Alliance cloud (and managing clouds in Melbourne/Edinburgh)


### **Developing and testing the codes for network packet marking (WLCG network Group)**

Tested in the terabit/sec demo projects. Planned to be used for HL-LHC and upcoming data challenges

### **FPGA and detector simulation code development**

# Babar Long-Term Data Access Facility (LTDA)

## Status and Future of the BaBar Long Term Data Preservation and Computing Infrastructure



Marcus Ebert  
BaBar Computing Coordinator

on behalf of  
**BaBar**



CHEP 2026 Bangkok (May 2026)

### BaBar stopped recording e+e- data in 2008

Babar still has some of the largest upsilon data sets

### Cloud computing facility at UVic with data archive at Gridka (KIT)

Provides full analysis capabilities as well as collaboration tools.

### Still many active users and thousands of analysis jobs

25 BaBar conference papers in 2025

### However, it is a challenging infrastructure to maintain

Old operating system and software packages, security, knowledge, ..

# CERN COURIER

Reporting on international  
high-energy physics

Physics ▾ Technology ▾ Community ▾ Magazine

COMPUTING | FEATURE

## Hidden treasures

9 September 2025

As the LHC surpasses one exabyte of stored data – the largest scientific data set ever accumulated – Cristinel Diaconu and Ulrich Schwickerath call for new collaborations to join a global effort in data preservation, to allow future generations to unearth the hidden treasures.

### CERN Courier article (“Hidden Treasures”) September 2025

*“The most striking example is BaBar, an electron–positron-collider experiment at SLAC that was designed to investigate the violation of charge-parity symmetry in the decays of B mesons, and which continues to publish using a preservation system now hosted outside the original experiment site.”*

### Members of the Data Preservation in HEP Working group

EU Strategy submission: <https://arxiv.org/abs/2503.23619>

### ICFA Panel on Data Lifecycle (IPP is a member of ICFA)

Highlights the importance of data preservation by our community

### Tri-Agency Research Data Management Policy

One must include plans for data preservation

# CPU Benchmarks

## CPU benchmark are a critical part of our global computing infrastructure (LHC, Belle II, DUNE, ..)

Countries make pledges to provide CPUs based on their benchmarks of their processors

The WLCG tracks the resources for the experiments (WLCG Accounting Group)

Facilities procure resources based on the performance using the benchmark

### WLCG Management Board created Task Force

#### HEPScore A new CPU benchmark for WLCG computing

Randall Sobie, University of Victoria

on behalf of the  
HEPiX Benchmarking Working Group

*Domenico Giordano and Michele Michelotto (co-chairs)*

Luca Atzori, Jean-Michel Barbet, Gonzalo Menéndez Borge, Christopher Hollowell, Ladislav Ondris, Andrea Sciaba, Emanuele Simili, Randall Sobie, David Southwick Tristan Sullivan, Natalia Szczepanek, Andrea Valassi

#### WLCG HEPscore Deployment Task Force

*Domenico Giordano and Randall Sobie (co-chairs)*

Julia Andreeva, Giuseppe Andronico, Gonzalo Merino Arevalo, Tommaso Boccali, Gonzalo Menendez Borge, Concezio Bozzi, Simone Campana, Ian Collier, Alessandro Di Girolamo, Michel Jouvin, Walter Lampl, James Robert Lettis, Zach Marshall, Helge Meinhard, Andrew Malone Melo, Bernd Panzer-Steindel, Stefano Piano, Danilo Piparo, Fazhi Qi, Matthias Jochen Schnepf, Oxana Smirnova, Jeff Templon, Andrea Valassi, Joshua Lee Willis, Tony Wong

CHEP 2023, Norfolk, May 2023

Computing and Software for Big Science (2021) 5:28  
<https://doi.org/10.1007/s41781-021-00074-y>

ORIGINAL ARTICLE



#### HEPiX Benchmarking Solution for WLCG Computing Resources

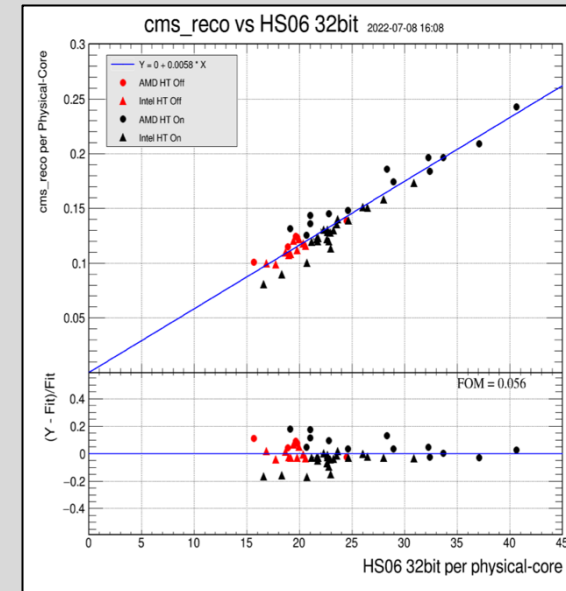
Domenico Giordano<sup>1</sup> · Manfred Ale<sup>2</sup> · Luca Atzori<sup>1</sup> · Jean-Michel Barbet<sup>3</sup> · Olga Datskova<sup>1</sup> · Maria Gironi<sup>1</sup> · Christopher Hollowell<sup>2</sup> · Martina Javurkova<sup>4</sup> · Riccardo Maganza<sup>1</sup> · Miguel F. Medeiros<sup>1</sup> · Michele Michelotto<sup>6</sup> · Lorenzo Rinaldi<sup>7</sup> · Andrea Sciabà<sup>1</sup> · Randall J. Sobie<sup>7</sup> · David Southwick<sup>1,2</sup> · Tristan Sullivan<sup>7</sup> · Andrea Valassi<sup>1</sup>

Received: 30 August 2021 / Accepted: 8 November 2021 / Published online: 14 December 2021  
© The Author(s) 2021

#### Abstract

The HEPiX Benchmarking Working Group has developed a framework to benchmark the performance of a computational server using the software applications of the High Energy Physics (HEP) community. This framework consists of two main components, named HEP-Workloads and HEPscore. HEP-Workloads is a collection of standalone production applications provided by a number of HEP experiments. HEPscore is designed to run HEP-Workloads and provide an overall measurement that is representative of the computing power of a system. HEPscore is able to measure the performance of systems with different processor architectures and accelerators. The framework is completed by the HEP Benchmark Suite that simplifies the process of executing HEPscore and other benchmarks such as HEP-SPEC06, SPEC CPU 2017, and DB12. This paper describes the motivation, the design choices, and the results achieved by the HEPiX Benchmarking Working group. A perspective on future plans is also presented.

**Keywords** CPU benchmark · GPU benchmark · High throughput computing · WLCG · LHC computing · HEP experiments · High-Energy Physics · Heterogeneous computing



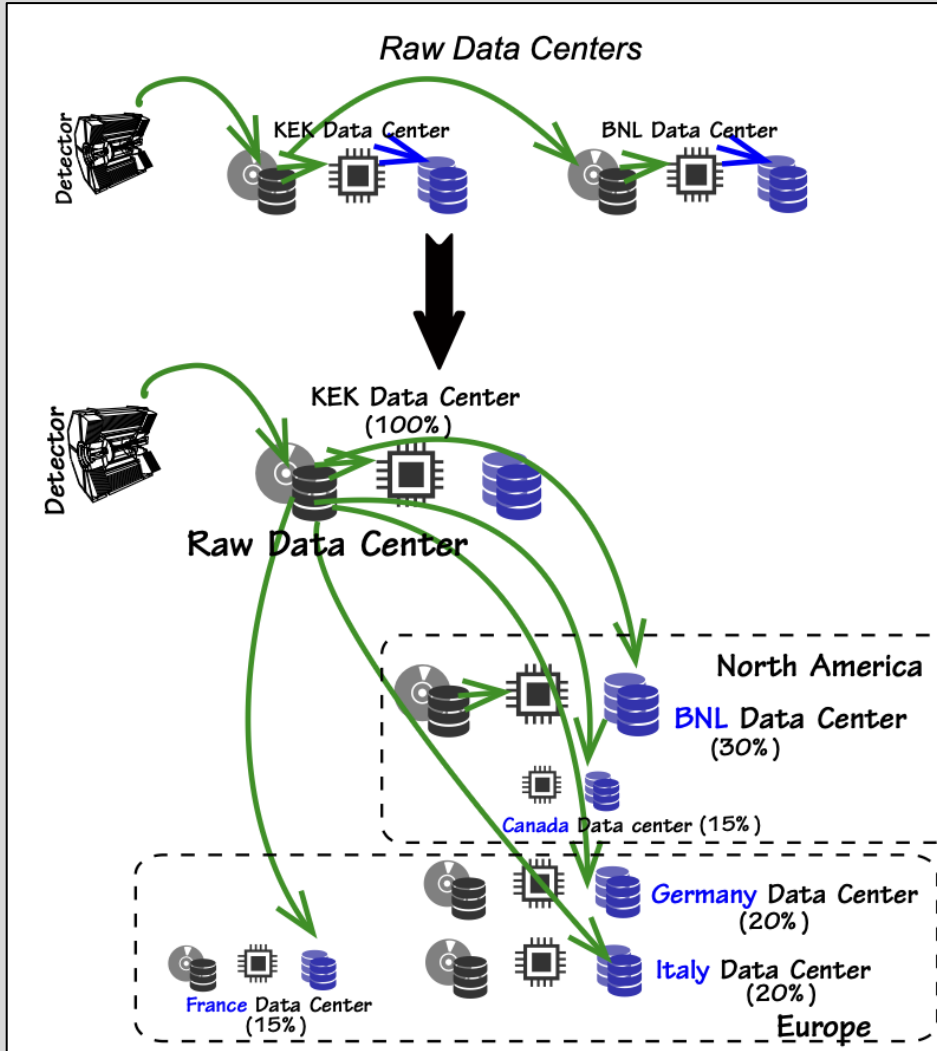
## Created a new CPU Benchmark (HEPScore) that uses experiment workloads

Now the standard for LHC, Belle II and other experiments (replacing HEPspec)

## New functionality is being added to measure the power consumption and we are testing benchmarks that use CPUs+GPUs

Sustainable computing and provides ability to power up/down resources based on power availability and price

# Belle II Raw Data Centre



## Canada has an MOU with KEK to store Belle II raw data

The first copy of the Belle II raw data is stored at KEK

The second copy is distributed sites in North America and Europe

Canada is responsible for storage and computing of 15% of the second copy

## Belle II Canadian Raw Data Centre in Victoria

Funded with a \$2M CFI award

Co-managed by the local UVic/Alliance research computing group

Fully operational

# Network packet marking and Super Computing demo projects

## HEPNET has been active in network R&D projects with the WLCG network group for 20 years

CERN, Caltech, Michigan, CANARIE, Internet2, ESnet, GEANT2, ..

## UVic is the only university in Canada with a direct link to CANARIE (specifically for our demo projects)

In 2025, the link was 1.6 terabits/second and was fully utilized using 3 servers (two in UVic and one in SFU)

Technical goal of the demo was to see if we could use high-speed networks with the network packet marking software

In Focus • Network

Canadian Particle Physics team achieves first terabit-per-second data transfer from a Canadian University to the SC25 Supercomputing Conference

7 months ago



## R&D: SC26 Demo and mini Challenge




- **SC26 Goal:** Showcase Scitags for high-fidelity network visibility
- **Area 1: Global 1.2/2.8 Tbps Packet Marking**
  - 1.2 Tbps end-to-end (3x400G DTNs); potential 1.6 Tbps (2x800G) in addition
  - Sites: UVic, CERN, StarLight, FNAL, & SC26 booth
  - Tests: Memory/disk tests (iperf3, xRootd, dCache) and performance impact
- **Area 2: Production Demo of Fireflies (US/ATLAS & US/CMS)**
  - Scitags visibility in active HEP production (BNL & FNAL)
  - Dashboards at StarLight showing flows by scientific activity
  - Identify traffic owner/purpose for detailed visibility
- Connect to **ESnet high-touch** for packet-level drill-down

## HP will provide servers and switches for the 2026 demo project

Goal is 1.6 terabits/second

Disk-to-disk transfers with packet marking



- **(C) Community** identifier: "Who are you affiliated with?"
- **(A) Activity** identifier: "What are you doing within your community?"
- **(E) Entropy** bits sprinkled throughout

## Packet marking will be used to label the experiment, purpose, source and destination

Primary software development in Victoria (T. Sullivan)

Part of the next LHC data challenge (2027)

## Track record of training students and staff in physics/computing

Currently four long-term research staff

Graduate student authorship-qualification computing projects

Co-supervised two graduate students in Computer Engineering

One held a permanent CERN IT position

Coop students are a better match for our projects

Wide range of diverse student backgrounds

## Staff and student obtain valuable skills in cloud computing

Computing industry, government, research labs, medical, ..

## Former Staff

Frank Berghaus

Kevin Casteels

Micheal Paterson

Ian Gable

Ron Desmarais

Chris Tooley

Ashok Agarwal

Kyle Fransham

Colin Leavett-Brown

Marcus Lee

Catherine Meng

Rolf Seuster

Matt Vliet

Duncan Penfold-Brown

Daniel Vanderster

## Former coop students

Valeria Gomez (Winter 2026)  
Sejion Lee (Fall 2025)  
Payne Jones (Winter 2025)  
Travis Lee (Fall 2024)  
Alyssa Blair (Fall 2023)  
Cindy Ng (Fall 2022)  
Victor Kamel (Fall 2021 and Summer 2022)  
Kathryn Dolby (Summer 2021)  
Garet Robertson (Winter 2020)  
Matthew Ens (Fall 2019)  
Shaelyn Tolkamp (Winter 2019)  
Tayha Weiss-Gibbons (Fall 2017 and Winter 2018)  
Jodie Weldon (Winter 2017 and Fall 2018)  
Trevor Lee (Fall 2016)  
Darryl Ring (Winter 2016)  
Matthew Murray (Summer 2015)  
Martin Conklin (Winter and Summer 2015)  
Colson Driemel (Summer 2014 and Winter 2015)  
Alex Lam (Winter 2014)  
Robert Prior (Fall 2013)  
Michael Chester (Winter 2013)  
Pranav Shrestha (Winter 2012)  
Andrey Polyakov (Fall 2011)  
Tubego Phamphang (Fall 2011)  
Robert Rusnak (Summer 2011)  
Tony Bashi (Winter 2011)  
Brennan McKinney (Winter 2011)  
Drew Harris (Fall 2010)  
Michael Anderson (Fall 2010)  
Holly Leavett-Brown (2010-2012)  
Matt Vliet (Winter 2009 and Summer 2010)  
Adam Bishop (Summer 2009 and Winter 2010)  
Michael Paterson (Summer and Fall 2008, Fall 2009, Winter 2010)  
Vikram Sandhu (Fall 2009)  
Chris Usher (Winter 2006, Summer 2009)  
Cole Uhlmann (Summer and Fall 2008)  
David Bartle (Winter 2008)  
David Grundy (Winter 2008)  
Sean Manning (Winter 2008)  
David Gong (Fall 2007)  
Sydney Schaffer (Summer 2007)  
Cameron Sobie (Summer 2007)  
Tristan Sullivan (Winter 2007)  
Duncan Penfold-Brown (Winter 2007)  
Sergey Popov (Fall 2006)  
Simon Ramage (Summer 2006)  
Patrick Armstrong (Summer 2006)  
Avery Berman (Winter 2006)  
Angela Norton (Fall 2005, Fall 2006)  
Ron Desmarais (Summer 2005)  
Marco Yuen (Winter 2005)  
Quinn Matthews (Fall 2004)  
Clayton Lindsay (Summer 2004)  
Alexandro Dimopoulos (Winter 2004)  
Lila Klektau (Fall 2003, Summer 2004)  
Manjinder Benning (Winter 2003)  
Jenny Allan (Winter 2003)  
Dan Vanderster (Fall 2002)  
Ian Zwiers (Summer 2002)  
Graeme Smetcher (Winter 2002)

# Impact with industry and public

Canadian Particle Physics team achieves first terabit-per-second data transfer from a Canadian University to the SC25 Supercomputing Conference

UVIC Open-source code tracks data's international travels

**Caltech**

About Caltech News & Events Research

11/21/2012

### High-Energy Physicists Smash Records for Network Data Transfer

New methods for efficient use of long-range networks will support cutting-edge science

**North Korea Times**

Thursday 15th December 2011

**Application Server Monitor**  
Pinpoint issues up to 90% faster See the video + Free Trial

**Researchers transmit record 186 GB data per second**  
North Korea Times  
Wednesday 14th December, 2011 (IANS)

Researchers have successfully transmitted 186 gigabits of data per second (GBps), accomplishing a new world record and helping usher in the next generation of high-speed network technology.

The rate is equivalent to moving two million gigabytes per day, fast enough to transfer nearly 100,000 full Blu-ray discs - each with a complete movie and all the extras - in a day.

Randall Sobie  
Institute of Particle Physics  
Research Scientist  
University of Victoria

Blue tubes contain the two beam pipes and register at 1.8 degrees below

openstack summit

Keynote: Clouds in High Energy Physics

Research scientist uses Microsoft Cloud technology for high-energy physics computing

1,127 views

SanDisk Case Study: SanDisk Helps Accelerate High Energy Physics

Ian Gable  
HEPnet Canada Technical Manager  
University of Victoria

**BBC NEWS TECHNOLOGY**

14 December 2011 Last updated at 10:16 ET

## Scientists break world record for data transfer speeds

SUCCESS STORY

**BROCADE**

University of Victoria

The Fast Path to International Research Collaboration

The University of Victoria (UVIC) located in British Columbia, Canada, is a globally recognized, research-intensive university offering innovative programs for more than 20,000 students. Approximately 1,000 research faculty teach and conduct work in a wide range of areas, including oceans and climate, genomics and proteomics, physics, astronomy, chemistry, engineering, and computational modeling. UVIC research teams connect to colleagues, cutting-edge projects, and powerful resources around the world through a number of networks.

**SanDisk Helps Accelerate High Energy Physics Research for the Origin of the Universe**

Summary

Through the use of Fusion ioMemory™ PCIe solutions, the University of Michigan and University of Victoria are implementing a multi-site supercomputing project to help them share massive volumes of data among the CERN Laboratory in Geneva, Switzerland and 100 computing centers around the world. The solution has the potential to accelerate access to physics data-allowing teams of physicists to expedite the potential discovery of new particles and forces that will help explain the nature of the Universe.

canarie

heprc.phys.uvic.ca

**HEP Data Intensive Distributed Cloud Computing System**  
Uncovering the secrets of the universe

Articles

Open-source code tracks data's international travels [UVIC News February 2008](#)

Canadian Particle Physics Team Achieves First Terabit-Per-Second Data Transfer from a Canadian University to the SC25 Supercomputing Conference [BCNET News Item November 20, 2012](#)

GEANT Connect December 2012

CANARIE Network Quadruples Capacity for Research and Education Data to 397 Gigabits per second [SANSUR News Item January 28, 2013](#)

\$16.4M reboot elevates cloud computing performance and energy efficiency [UVIC News June 5, 2014](#)

Rebuilding a plane while it's still flying: UVIC supercomputer to get \$16.4M overhaul [Times Colonist June 5, 2014](#)

UVIC receives \$16.4M for supercomputer infrastructure upgrade [Victoria News June 6, 2014](#)

The Alliance awards \$500k to Simon Fraser University and the University of Victoria to renew advanced research computing infrastructure [Digital Research Alliance of Canada News June 3, 2014](#)

Alliance 2020-2021 Annual Report [Researcher Profiles](#)

NDRI's Researcher Council Elects New Chair, Dr. Randall Sobie [SORDI News March 2011](#)

CFI Award for the Belle II Raw Data Centre [UVIC News March 2011](#)

Canada's New Digital Research Infrastructure Organization Names Inaugural Researcher Council [NDRI 2010](#)

A data system for the world's largest physics experiment [UVic Ring 2010](#) and [BC Technology 2010](#)

The Fast Path to International Research Collaboration, [Brocade 2015](#)

HEP Data-intensive Distributed Cloud Computing System [CANARIE 2015](#)

SanDisk Helps Accelerate HEP Research for the Origin of the Universe [SanDisk Media release 2014](#)

Canadian physicists achieve 100G transatlantic transmission [CANARIE 2014](#)

CANARIE invests in 9 innovative software projects [CANARIE 2014](#)

Team demonstrates intercontinental OpenFlow network at Super Computing (BCNET 2013)

Enterprises launch their own private clouds [CNN Money 2013](#)

Physicists link OpenStack, Nimbus clouds around the world to share resources [IT World 2013](#)

Moving at the speed of data [UVIC Internal 2013](#)

High-Energy Physicists Smash Records for Network Data Transfer 2012 [Caltech News](#)

Particle physicists use fast network to find the Higgs boson [CANARIE 2012](#)

Data goes faster than ever [CERN 2012](#)

High-Energy Physicists Set Record for Network Data Transfer [BBC](#) and [NSB](#) and [Ciena](#)

Preserving research data to enable long-term investigation and discovery, [CANARIE 2011](#)

Canada Explores New Frontiers in Astroinformatics, [UVIC in the Cloud 2011](#)

UVIC research data transfer = 132 million books, [UVic Ring December 2010](#)

Cloud computing at UVIC, [CANARIE 2010](#)

UVIC establishes compute cloud for particle physics, [UVIC Vice-President Research Newsletter Preserving the Data Harvest, Summary Magazine, December 2009](#)

Researchers Collaborate to Sustain Canada's Forests, [BCNET 2009](#)

\$1.55 Million Funds Data Preservation, [UVIC Media Release 2009](#)

CANARIE Invests Millions More in Cutting-edge Canadian Research, [CANARIE 2009](#)

Technology Brings Scientists Together [Oak Bay and Samich News 2008](#)

WestGrid linked to Nobel Prize in Physics, [WestGrid News, 2008](#)

BCNET Faculty Recognition Award, [BCNET 2008](#)

Building a high-speed network for High Energy Physics, [UVic 2008](#)

CANARIE Announces Network Enabled Platform (NEP) Projects, [CANARIE 2008](#)

BCs Big Science Projects get on Board to Use New Supercharged ROADM Network, [BCNET 2007](#)

ATLAS: Higgs Data Needs Huge Bandwidth, [BCNET 2007](#)

The Power of Grid Computing, [UVic 2006](#)

Canadian grid computing project finds place in ATLAS, [iBusiness 2005](#)

Creating the World's Largest Computer Grid [BCNET 2005](#)

GridX project joins LHC Computing Grid, [CERN Courier 2005](#)

Grid project probes our origins, [IT World Canada 2005](#)

UVic Helps Form World's Largest International Computer Grid, [Grid Today article 2005](#)

Dan Vandenberg : UVic graduate student wins BCNET coolest application contest [BCNET 2004](#)

Dan Vandenberg : IEEE Gold Medal in Computer Engineering [UVIC 2003](#)

Girding for the Grid, [UVIC 2003](#)

# Summary

## **Contributions to Canadian and international HEP community**

UVic (VISPA) HEP Research Centre and growing links with the MacDonal Institute and TRIUMF  
HEPNet/Canada  
Editor-in-Chief, Elsevier journal, SoftwareX  
Digital Research Alliance of Canada  
Co-Chair of the WLCG CPU Benchmark Task Force

## **Specialized computing and networking projects**

Compute clouding for HEP using national and international resources (Belle II, ATLAS and DUNE)  
Belle II Raw Data Centre  
BaBar Long Term Data Access System  
Network packet marking with the WLHC Network Group  
Terabit/second network demo projects with CANARIE, WLHC Network Group and industry  
FPGA, GPU-benchmarks and detector simulation software development projects

## **Funding, training of HQP**

Three major CFI awards, multiple CANARIE research software contracts, SLAC/BaBar LTDA contract, HEPNET  
Unique environment for training HQP in particle physics computing and networks

## **Active program and funding for the coming years**