NSC News

National Supercomputer Centre at Linköpings universitet

2012:3

Grid computing and the Higgs boson

Photo provided by CERN. Used with permission.

Grid computing made Higgs discovery possible?

Dr. Oxana Smirnova describes how grid computing was an important technology in the observation of the Higgs boson.

Read further on page: 3

Triolith, Krypton and Skywalker

Mats Kronberg gives an update on the installation of NSC's three new systems.

Read further on page ?

Baton Received

In the latest editorial, Prof. Bengt Persson passed me the baton of directorship. I received it in the midst of the summer at vacation-slow times, so it is only now, a few months later, that I have grasped the extent to which NSC has transformed under his leadership. With offensive staff recruitment, NSC has acquired a broad spectrum of competences in high-performance computing (HPC). This development has been largely driven by the formation of the Swedish e-Science Research Centre (SeRC), and Bengt's involvement in this work leaves a footprint that affects the entire Swedish HPC landscape.

There was a time when research groups treasured their own dear and precious HPC resources, keeping them outside the reach of others and appointing a less fortunate soul in the unit to be responsible for the system management. Countless man-hours that could have, and should have, been dedicated to the advancement of science were spent in this way. Today, when visiting research groups around Europe, one can still see this pattern, whereas we in our country are fortunate enough to have national supercomputer centers to which we can send in applications that takes less than an hour to fill out and which gives us generous access to secure and well maintained HPC resources. I dare say that Swedish researchers in the HPC community have never before been blessed to spend so little time on issues directly related to the computing resources.

So, once again, thank you Bengt for, together with many others, having set the snowball in motion. As I see it, the main challenge that we now face is to define the interface between the research group and the centralized software service organization. This software interface is inherently more difficult to define than that of the hardware, which consists of standard compilers, mathematical libraries, resource schedulers, etc. At the software interface people with different expertise need to meet and share information and, as we all know, there are no clear and unambiguous protocols to control human interactions.

PATRICK NORMAN, NSC DIRECTOR

Apply for compute resources

A grand total of about 225 million core hours is available in the application round for Large Scale SNIC projects this fall. The call was issued on September 24 with a dead-line for submission on October 24 at 15:00. The new application system SUPR, SNIC User and Project Repository, https://supr.snic.se/, is used for the first time in a Fall Large Scale SNAC application round this time. It was introduced for the Spring 2012 Large Scale application round and is also used for Medium Scale applications since June this year.

New in the application round this fall regarding resources is that phase two of NSC's new main system Triolith is ready. Triolith is increased from 240 nodes in phase one to a total of 1200 nodes. The full system will be ready when allocations become effective on January 1, 2013. It is also worth to notice that the limit for Medium Scale allocations is higher on new systems than the traditional 80 000 core hours / month that it used to be on all systems. It is for example 200 000 core hours / month on Triolith.

Detailed instructions for applicants are available at: http://www.snic.se/apply-for-resources with specific instructions for Large Scale applications at: http://www.snic.se/apply-for-resources/large-scale-applications.

PETER MÜNGER, NSC

New staff member

I've had time for some exciting challenges. It began with moving from Västervik to Linköping to study Computer Science and Engineering at LiTH and continued with one year abroad at Nanyang Technological University in Singapore. Back in Sweden I signed up for a consultant mission at FOI, where we developed a game deployable on NSC's ex cluster Monolith and used for IT-security training. After that I moved to the west coast and an embedded Linux consultant job at Cybercom in Göteborg, where I participated in two global Linux projects within automotive and home electronics. Back in Linköping, I worked as UNIX/Linux system administrator at Ida Infront. Now at NSC I work as system developer with focus on security solutions.

ANDREAS LINDQVIST, NSC

Small Particle in Big Data

It was a warm late summer afternoon in a quiet town not far from Moscow, and I had to make up my mind on the subject of my master's project. The town hosted an acelerator facility and a branch of the Moscow State University where I was a student. The Large Electron-Positron collider (LEP) at CERN was turned on just two years ago, and a wise man suggested me to do a project with the DELPHI experiment at LEP: Russian physicists were in, and LEP was supposed to discover Higgs. Indeed, following discovery of W and Z bosons at CERN in 1983, the Standard Model of elementary particles was proven beyond reasonable doubt, and discovery of the only missing piece, the Higgs boson, was just a matter of time. It was more than 20 years ago.

Since then, LEP got disassembled to give way to the Large Hadron Collider (LHC). Superconducting Super Collider (SSC) in Texas didn't progress beyond a quarter of the tunnel before getting cancelled. Tevatron collider in Illinois got upgraded and kept on smashing protons and antiprotons until September 2011, but no Higgs boson was seen.

And then it came, in December 2011: still tiny, but too noticeable to be ignored, a signal of a new particle with a mass quite consistent with what a Standard Model Higgs boson could have. The signal came from LHC. In July 2012, with increased accelerator collision rate, the discovery became obvious. With every new day of data taking, we are getting closer to the conclusion that we finally found the long-sought Higgs boson.

But why did it take so long? One obvious answer is that we didn't have such a great acelerator before, neither did we have such advanced experimental facilities. But have a look at the only slide shown by the CERN Director General on July 4 2012. It says, quote: "Results today only possible due to extraordinary performance of accelerators - experiments - Grid computing".

Grid computing made Higgs discovery possible? How is that?

Like most other elementary particles, Higgs won't leave traces in our detectors: it will decay well before reaching any sensor. Physicists can only record such decay products, and the challenge is to reconstruct what was the original particle. Long gone are the days when such reconstructions were made using a filmed snapshot of a collision, naked eye and a ruler. Modern detectors have spatial resolution of microns and geometrical size of a large house, meaning millions of measurements per one collision event. All these measurements are digitally recorded, constituting raw data. The magnificent LHC machine provides up to 600 million collisions per second. Assuming one bit of data per sensor hit, multiplied by number of hits per collision, and by number of collisions, will give a rough estimate of 100 Terabyte per second of data flow from one detector alone. No technology exists today to sustain such a load. But we don't really need to record all these data: after all, Higgs boson production is a very rare event, and during 12 months of LHC data taking in 2011-2012, it was expected to be observed only few hundred times. An obvious approach to data reduction would be to record only a subset of candidate events that carry a specific signature. This has been done very successfully: for example, all the data collected by single ATLAS experiment at LHC, raw, derived and simulated, fit just about 75 Petabyte, as of today (Fig. 1).

This is still a very significant data volume to process. And yet, the discovery announcement on July 4 was referring to data collected as recently as June 22! For comparison, similar analysis of Tevatron collider data was published 9 months after the machine was shut down. The overall data volumes collected by Tevatron and LHC so far are quite similar, the experimental setups are also very similar, and scientists are equally skilled. Where is the diference? Why could LHC deliver the results with such an amazing speed that never happened before?

Let's have a closer look at the Grid computing contribution to the LHC success. Collider experiments always relied on most advanced computing technologies, and accelerator facilities often hosted latest and greatest supercomputers. And even back in my student's days, a single supercomputer was not enough: data collected at CERN were written on tapes and shipped for processing to other centres. High bandwith, but high latency, too. The networks of today made tape shipings obsolete. Even if a single line to a small university is not

New staff member

I've recently joined NSC as a systems expert on parallel clusters. My primary task will be to help take care of SMHI production and research systems. For the past 12 years I've worked as a unix system administrator at Saab Aeronautics in Linköping. I've been part of a team maintaining the technical platform for flight simulators including real-time systems and image generators. It will be an exciting challenge to be part of the system expert team at NSC, contribute my experience to the group and become proficient in HPC technology.

FREDRIK NYSTRÖM, NSC

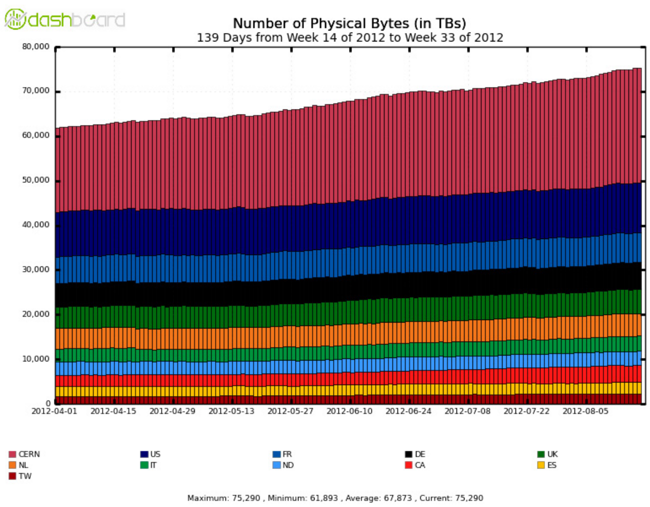

Figure 1: Aggregated ATLAS data on the Grid. Colors indicate regional computing and storage centres, ND representing the Nordic facility.

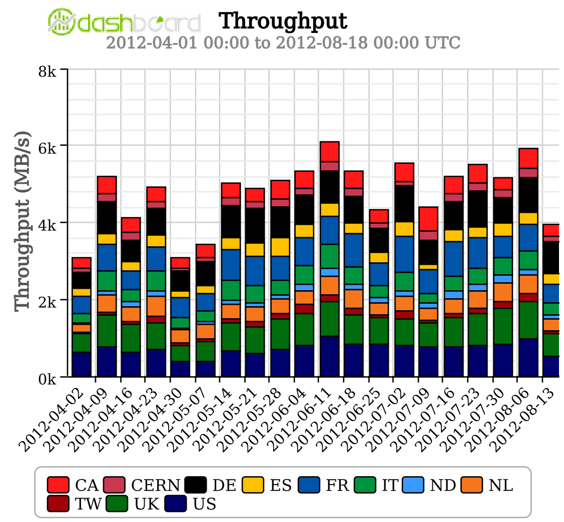

Figure 2: ATLAS Grid data throughput. Colors indicate regional computing and storage centres, ND representing the Nordic facility.

too impressive, combining together all possible destinations results in sustained throughput reaching several Gigabytes per second for one LHC experiment alone (Fig. 2). And LHC data are very easy to split, since collision events are independent from each other. This means we don't even need expensive supercomputers: a lot of data can be processed by regular Linux clusters, if we can get hold of many of them.

When you use few clusters for your analysis, you can handle it manually, or with a little help of few scripts. Having to handle hundreds of Petabytes of data, LHC uses more than 140 Linux clusters all around the planet. This is what makes it different from all previous collider experiments: it outsources computing and storage to previously unseen number of facilities, linked together into a huge coherent e-Science infrastructure. This is what allows LHC physicists to get the results almost instantly.

How does one create a coherent infrastructure from 140 disparate sites? This is where Grid technology comes in. While conventional computing service providers offer shares of their resources to customers, Grid is a technology that allows users to federate their conventional shares into one distributed resource, which they use collectively. This technology is implemented as specialised software, called middleware. Grid middleware provides common interfaces to very different systems, and helps supporting the hierarchial infrastructure for LHC computing. This hierarchy is centered at CERN (Tier 0), from where raw and derived data are transferred to 11 regional centres (Tier 1) for safe keeping, reprocessing and eventual analysis. From Tier 1 centres, data are distributed further to Tier 2 centres for data analysis and simulations. Within 48 hours of being recorded at CERN, raw and derived data are available to physicists all over the world, not in one central location, but wherever it is faster to process them.

Sweden was among first countries that established a Grid infrastructure already in 2003, in cooperation between the Swegrid project and NorduGrid ARC middleware developers. Swegrid consisted of 6 major scientfic computing centres, and NSC traditionally offered not just excellent hardware resources, but also outstanding human expertise. Since 2006, Swegrid contributes to the

SMHI personnel works tightly together with NSC

Hi, my name is Sébastien, I am employed by SMHI in the FoUp group as scientific programmer since June 2011. Since I maintain, optimize and administrate SMHI's Meteorological Archive and Retrieval System (MARS) hosted at NSC, I tightly collaborate with NSC experts and work one day per week there.

My scientific background is at the intersection of chemistry, physics and computer science. I obtained a PhD in Quantum Chemistry from Strasbourg University. I worked two years as a postdoc at Uppsala University, and two years at Linköping University. My expertise encompasses C++, FORTRAN and python programming, parallel and high performance computing, data design and data flow management.

SÉBASTIEN VILLAUME, SMHI

LHC computing both via the Nordic Tier 1 centre, and also by providing Tier 2 services. When CERN Director General says the discovery is impossible without Grid, he means it is impossible without NSC contribution to it as well. And this discovery is not the end of the LHC programme, but merely a beginning: the machine will be providing excellent data for many years to come, and we will have to match it with excellent computing resources. And this is not the end of Big Data, but just the beginning: data volumes in all sciences are set to increase dramatically, and trail blazes of LHC will be essential to sucess.

OXANA SMIRNOVA

Dr. Oxana Smirnova is an Associate Professor at the Particle Physics division in Lund University, Sweden. She is a member of the ATLAS collaboration at CERN, working with Grid computing. Oxana Smirnova is a coordinator of LHC-related computing at the Nordic e-Infrastructure Collaboration (formerly known as NDGF), and is one of the developers of the NorduGrid ARC middleware.

New systems in operation

The first phase, 240 nodes, of NSC's new large system Triolith is now in full production. Assembly and testing of the second phase, 960 more nodes, started in late September.

If the second phase hardware turns out to be as reliable as the first was, we should be able to begin full operation of the whole system during November.

As the second phase only contains more compute nodes of the same type as the first phase, there will be no regular pilot test. We plan to go directly from acceptance testing to normal production. However, during the acceptance testing of the second phase nodes and during the final full-scale testing of the system we will consider testing user-supplied codes on a case by case basis.

If you have codes that you would like to test on job sizes up to 19200 cores, please contact support@nsc.liu.se.

The first phase, 96 nodes, of SMHI's new system Krypton is now in full production, and acceptance testing of the remaining 144 nodes has started.

The third new system Skywalker has completed testing and has been delivered to Saab where it will be integrated into their computing environment.

Finally, some bad news: In order to supply power and cooling to the new system Triolith, we are retiring the old system Neolith a few months earlier than originally planned. On September 16th, we powered off 40% of the Neolith compute nodes, and later on we will have to power off even more nodes to be able to properly test Triolith. Any data that you have on Neolith will continue to be available on Triolith, Kappa and Matter.

MATS KRONBERG, NSC

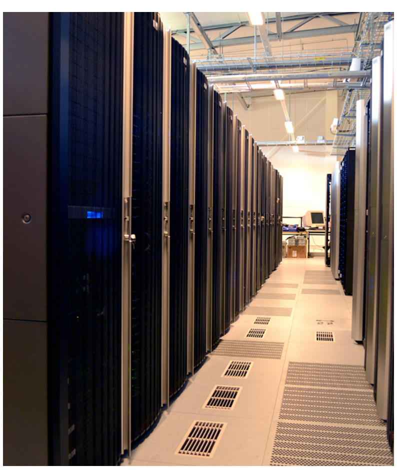

20076192

Linköping
Porto betalt

Training announcement!

The Intel HPC training at NSC continues. This time we invite you to learn about how to optimize serial and parallel performance using the tool VTune Amplifier XE during a two-day in-depth training class. The training will also include how to detect memory and threading errors using the tool Inspector XE.

This advanced training class is primarily intended for application/system experts at the various SNIC centers, as well as HPC users with a particular interest in performance analysis and optimization. The training will consist of both presentations and hands-on exercises that will run on the new NSC cluster, Triolith.

The training class is free of charge and includes lunch and coffee. Hence, participants will only need to cover travel and accommodation costs.

The training will take place on October 22–23, 2012 and will be given by the Intel expert Beverly Bachmayer, who is working with the Software and Solutions Group at Intel.

Early registration is an advantage, but not a guarantee for a seat in the class. Due to the limited number of seats, we reserve the right to select the participants that we think will benefit the most from the training.

Beverly Bachmayer, Bev, is a Technical Consulting Engineer at Intel Munich, she has worked in diverse software engineering, engineering and program management positions in the US and Europe during her 29 years at Intel. She is currently working in Developer Products Division and is Intel's lead on the EU PEPPHER project. Her key area of interest is performance analysis and optimization of software on new computer architectures. She develops and teaches courses on Intel Architecture, Optimization and Intel Tools, including Intel Parallel Studio XE 2013. Additionally, Bev supports increasing the number of professional females entering computer science/engineering programs worldwide through multiple local projects. Bev Bachmayer holds a Bachelor's degree in Computer Science from the University of Oregon (1983) and an MBA from Portland State University (1992).

More information and a registration form is available on: http://www.nsc.liu.se/intel-hpc-training.

TORBEN RASMUSSEN, NSC

User Outreach in Gothenburg

The application experts' user outreach initiative continues, this time in Gothenburg. Application experts from all major Swedish HPC centra were represented in a session lasting two hours where users met the application experts and discussed projects and usage issues in high performance computing.

The aim of the user outreach is to increase the HPC users' access to application experts and to provide a convenient opportunity for users for in-depth consultations and collaborative project discussions normally a bit outside the scope of the ususal support channel. However, any user issues are welcome.

The means of the outreach has been open face to face meeting sessions where anyone can approach application and system experts with any issue, large or small. The major NSC user sites have been visited so far, Uppsala, Stockholm, Linköping and Gothenburg. The intention is to visit all major university cities in Sweden in a round-robin fashion at a rate of one or two sites per term. In these meetings, system experts are also represented to allow comprehensive technical HPC support to complement the application specific support.

Available expertise has varied on the individual and application specific level between meetings, but will always be comprehensive in terms of scientific areas, so as to reflect and enable support to the diversity of Swedish academic HPC users. Represented disciplines are Molecular Dynamics, Electronic Structure Theory (Materials Science and Quantum Chemistry), Parallel Programming, Bioinformatics, Neuroinformatics and Computational Fluid Dynamics with Climatology soon to be added. Specific application expertise is typically available for the most popular software applications in the respective scientific areas.

The outreach initiative will now continue to the next site by the end of the fall term. It will be announced by mail and on the NSC web pages and is open to everyone interested.

JOHAN RABER, NSC

Linköping University

National Supercomputer Centre, Linköpings universitet, 581 83 Linköping, Sweden
Tel: +46 (0) 13-28 26 18, fax: +46 (0) 13-28 25 35, e-mail: info@nsc.liu.se
www.nsc.liu.se

Editor: Peter Münger
Use NSC Express, www.nsc.liu.se/express, or send an e-mail to support@nsc.liu.se to subscribe for a paper or electronic version.