



**PDC Center for
High Performance Computing**



**National Supercomputer Centre
at Linköping University**

NSC & PDC

Resource and Technology Providers for SeRC

Erwin Laure
Director PDC-HPC

Patrick Norman
Director NSC

NSC & PDC



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University



- The two leading HPC centers within Swedish National Infrastructure for Computing (SNIC)
- Founded in 1989 (NSC) and 1990 (PDC)

Resources and Technologies for eScience



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University

Frame of reference for NSC/PDC in SeRC

- eScience critically depend on the provision of excellent resources for
 - Networking (by SUNET)
 - Computing (by NSC & PDC)
 - Storage (by NSC & PDC)
 - Visualization (by NSC/PDC/LiU/KTH)
- To make efficient use of these resources, expert support and advanced technologies are needed
 - Advanced algorithms
 - Parallelization
 - Distributed computing
 - Data storage and management

New NSC Director: Patrick Norman



PDC Center for
High Performance Computing





National Supercomputer Centre
at Linköping University

- Appointed by the Rektor of LiU, contract until 2014/12
- 1998 PhD in Computational Physics
- 2010 Professor in Computational Physics
- Co-author:

DALTON <http://daltonprogram.org/>
quantum chemistry program package
SCF, DFT, MCSCF, MP2, CC
response theory, excited states, spectroscopies
2000 individual and 200 site licenses world-wide

DIRAC <http://diracprogram.org/>
4-component relativistic quantum chemistry program
heavy atom effects, spin-transition properties



Top-ranked in user satisfaction and science enabling



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University



$$I = \frac{G}{x + y} \sum_{i=1}^N f(i).$$

What is reasonable to ask for



PDC Center for
High Performance Computing

SNAC medium project: 200,000 hours/month

17 compute nodes



National Supercomputer Centre
at Linköping University

170 kkr per year

$\frac{1}{4}$ postdoc

What is a reasonable division of resources into the human vs. computer e-cloud



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University

SNIC cluster resource: 15-20 MSEK/year

For 10% one gets 2 AEs

What could we get for this investment?

HYPE code

The hydrological catchment model HYPE simulates water flow and substances on their way from precipitation through soil, river and lakes to the river outlet (Arheimer et al., 2008; Lindström et al., 2009).


<http://www.smhi.se/en/Research/Research-departments/Hydrology/hype-1.7994>



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University




2013-03-26 (from SMHI):

"I have made changes in the HYPE code after Jeff's report. [...] The code now runs faster (up to 50%) on some applications."

Flagship Systems

- Triolith @ NSC
 - 336 TFlop/s, 1200 HP SL230, Intel E5-2660, Mellanox FDR IB
- Lindgren @ PDC
 - 305 TFlops/s, Cray XE6, 3D Gemini Torus






PDC Center for
High Performance Computing



National Supercomputer Centre



Reorganization of PDC's Systems

- **Ferlin**

SNIC Throughput System – retired June 30, 2012

- **Ekman**

KAW-funded system for Climate and Flow research retired by Dec. 31, 2012

About half of the machine (~5,000 cores) remains available as **Ferlin** “new” to users from the Stockholm region using KTH funds


- **Povel**

Prototype system for PRACE; repurposed as pre- and post-processing to Lindgren

4320 cores (180 4x6core AMD nodes)

36 TF theoretical peak performance

5.76 TByte memory



- **Zorn**

GPU cluster co-funding by VIC-Stockholm and SNIC

40 GPUs




National Supercomputer Centre
at Linköping University

Major HPC systems at NSC

- **Krypton**
 - SMHI climate research and weather model development
 - 240 HP SL230, Mellanox FDR IB
- **Byvind**
 - SMHI weather forecasting
 - 140 HP SL6000
- **Kappa**
 - LiU material sciences and flow dynamics
 - 364 HP ProLiant DL170h, Voltaire QDR IB
- **Matter**
 - KAW-funded system for material sciences (LiU, KTH, UU)
 - 516 HP ProLiant SL2x170z, Voltaire QDR IB
- **Skywalker**
 - SAAB Aerospace research and development
 - 72 HP SL230, Mellanox FDR IB



Storage (approx. numbers)



- Disk storage
 - 5.4 PB GPSF and Lustre storage @ NSC
 - 900 TB (raw) DDN Lustre storage @ PDC
- IBM tape storage (NSC and PDC)
 - 2900 slots, 3.0 PB
 - Accessible via HSM, TSM, dCache, and MARS
- Large users
 - WLCG (high-energy physics at CERN)
 - IceCube (neutrino observatory at the South Pole)
 - BILS/ELIXIR (Bioinformatics infrastructures)
 - ECDS (Environment Climate Data Sweden)
 - INCF (Allan Brain Atlas, Waxholm Space ...)
 - ODIN Space Mission
 - Human Proteome Project
- Swestore national storage

Future



PDC Center for
High Performance Computing



National Supercomputer Centre
at Linköping University

- Expansion of Triolith 2nd half of 2013
 - Additional 320+ compute server nodes
 - Installation in new computer room 4x1 MW
- Lindgren will retire 2nd half of 2014
 - Preparing for a replacement of at least 1 PF capacity
- MetCoOp cluster resource 1st half 2015
 - SE/NO weather forecasting
 - SMHI climate research and weather model development



International Dimension



PDC Center for
High Performance Computing

- Swedish bridgehead for European e-Infrastructures
 - Provide resources to European e-Infrastructures under the coordination of SNIC
 - Support Swedish users in gaining access to them



National Supercomputer Centre
at Linköping University



 www.swegrid.se
SweGrid



PRACE open calls

Tier-0 call for PRACE regular access – open every 6 months

get **CPU time** on one of the six PRACE **Tier-0** systems

Complemented by rolling **Preparatory access** calls, open all year round. No closing dates, but technical evaluations are every 3 months at defined cut-off dates.

- type A (*Code scalability tests*); maximum allocation period is two months
- type B (*Code development and optimisation*); maximum allocation period is six months
- type C (*Code development and optimisation with the support of PRACE experts*); maximum allocation period is six months including up to 6 months **expert help** scaling your code for **Tier-0** systems

DECI (Tier-1) call for applications – open every 6 months

get both **CPU time and** up to 6 months **expert help** on one of the many **Tier-1** systems across PRACE partner countries.



PRACE upcoming calls

Tier-0 Preparatory access (PA) call is constantly open, but **next cut-off** date will be 3rd of June!

It will feature **additional 150PMs** of expert help to scale your code for Tier-0. These efforts are leftovers reactivated from the first implementation phase of PRACE (1IP). They have to be spent until December 2013.

Extra-special for those **interested to test a MIC architecture**:

The 2IP Prototype system EURORA at CINECA (Intel Sandy Bridge, 16 GB DDR3, 1600 MHz) will be upgraded with Intel Xeon Phi cards, so that 64 of its nodes will have Xeon Phi cards inside (the other 64 have already Kepler GPU cards) and will be one of the PA systems available.

CINECA will run a Xeon Phi Summer School 8-13 of July.

Those interested should also check out the PRACE Intel Xeon Phi Best Practice Guide currently being improved.



PRACE 8th call: Tier-0 access March 2014 – March 2015

PRACE 8th call for applications will open in September 2013.

More information and application form can be found at:

<http://www.prace-ri.eu/Call-Announcements>

- Local experts can help you with the application form:
- PDC
 - Michaela Barth caela@kth.se
- NSC
 - Chandan Basu cbasu@nsc.liu.se
 - Soon-Heum "Jeff" Ko sko@nsc.liu.se
- HPC2N
 - Mikael Rännar mr@cs.umu.se
 - Jerry Eriksson jerry@cs.umu.se



DECI – Tier 1 access all over Europe

The amount of allocations received in DECI compares to what is available on the national level.

→ DECI can be considered as an additional resource to what we already have within Sweden.

DECI offers a variety of different architectures across Europe. The amount of CPU time available per DECI call is about 5% (12m CPU hours) of Lindgren@PDC which is the Swedish PRACE resource. Over subscription is typically a factor 4.

DECI gives you the possibility to apply for up to 6 months PRACE expert enabling help on your applications' scalability.

Post Award obligations include a final report and acknowledgement of PRACE support.



Available user documentation

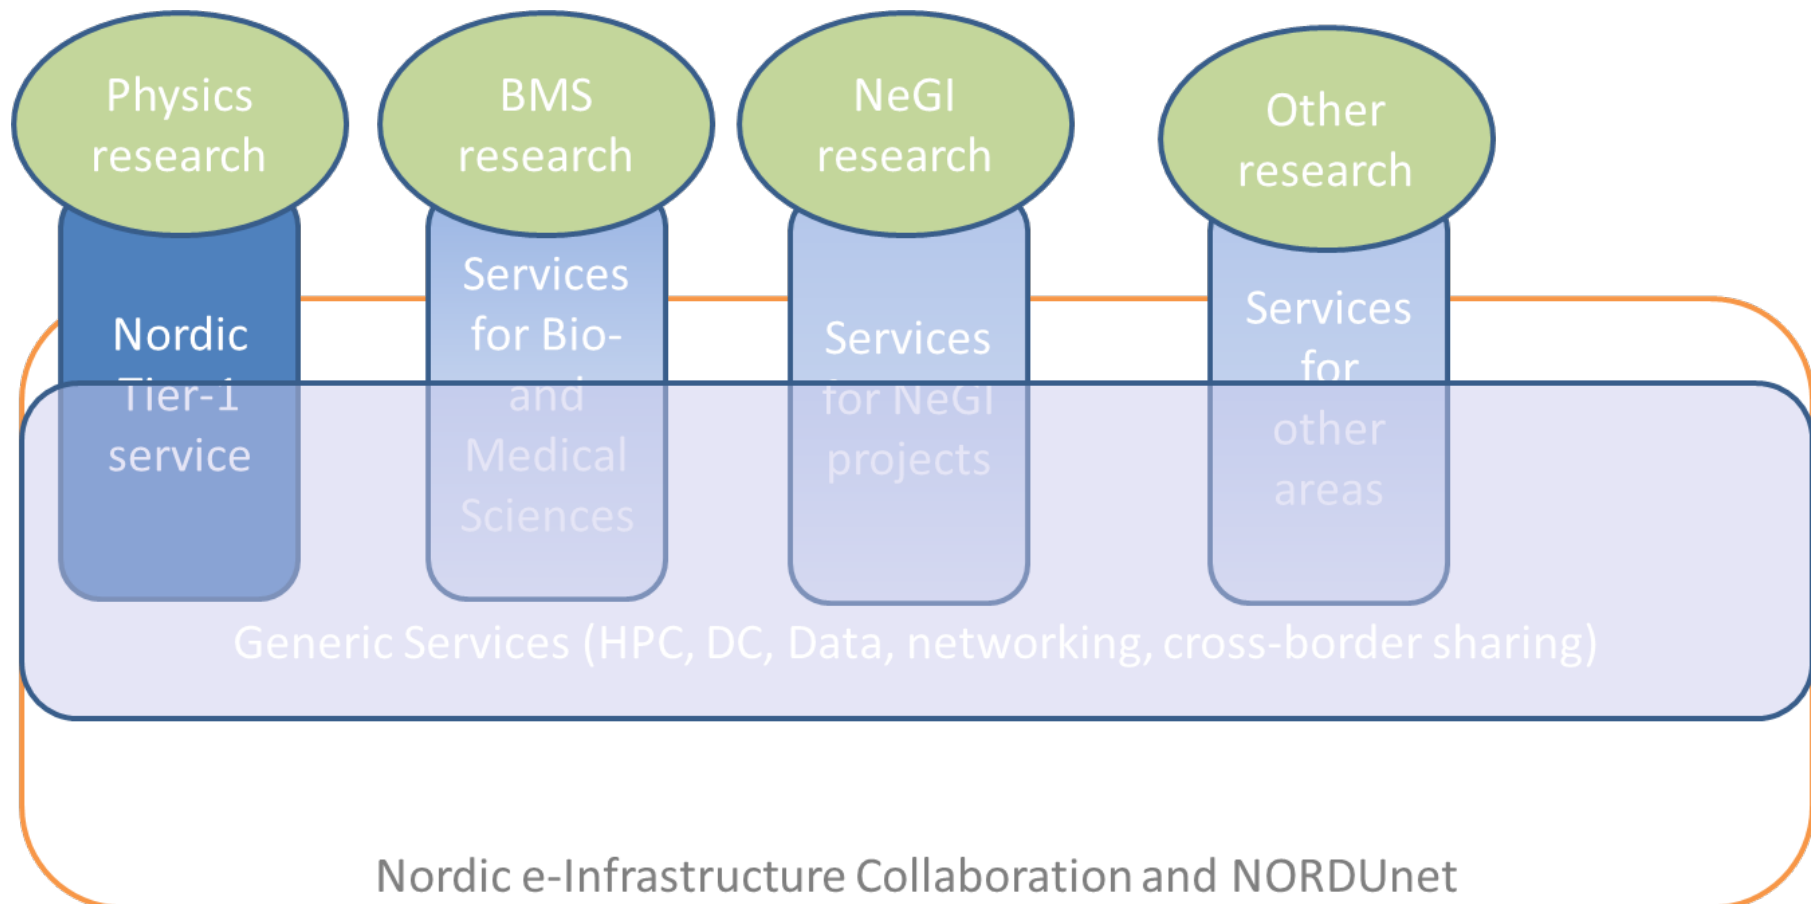
- PRACE Best Practice Guide:
<http://www.prace-ri.eu/Best-Practice-Guides>
- PRACE User Documentation:
<http://www.prace-ri.eu/User-Documentation>
- Training Portal <http://www.training.prace-ri.eu/>
with many hours of video training material available

Nordic e-Infrastructure Collaboration – Mission and Goals

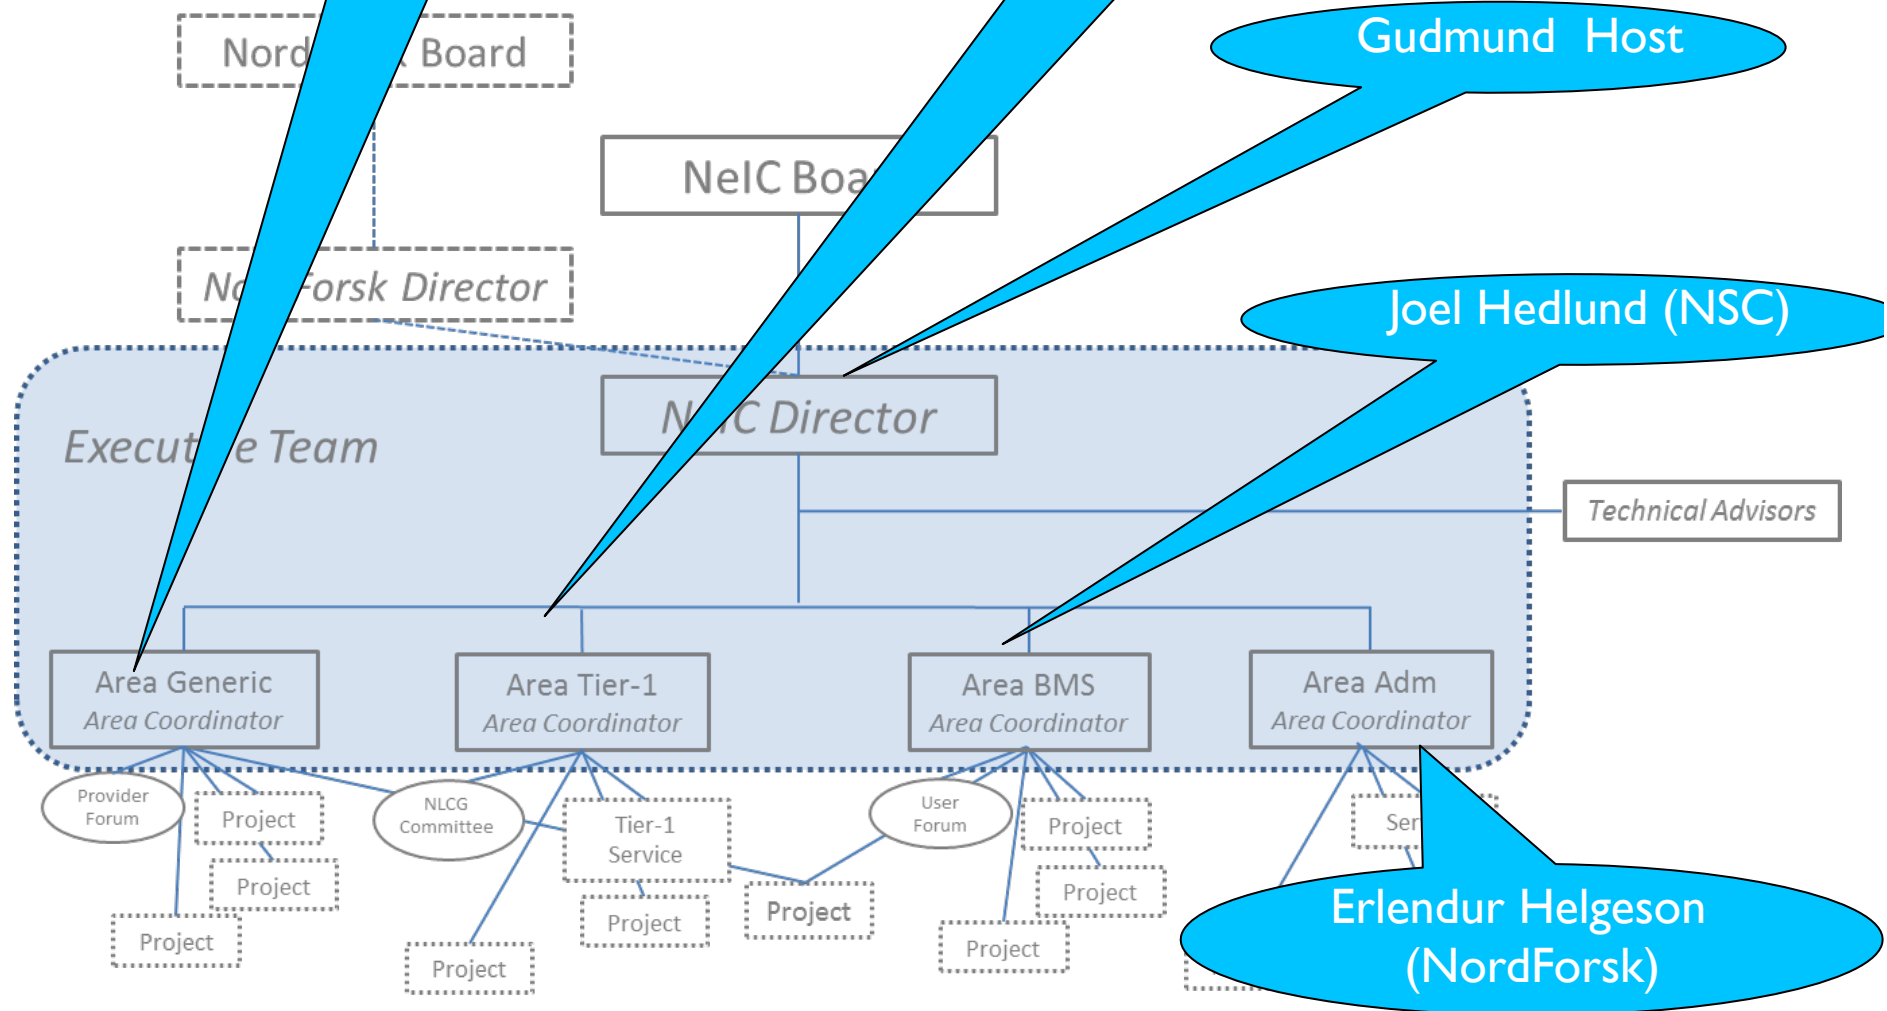
The mission of NeIC is to facilitate the development of high-quality e-Infrastructure solutions in areas of joint Nordic interest.

- Support new collaboration areas
- Strengthen collaboration on e-Infrastructures, all levels
- Increase visibility of e-Infrastructure collaboration
- A sustainable Nordic e-Infrastructure ecosystem
- Elicit new collaboration opportunities

Nordic e-Infrastructure supporting Nordic research



NeIC Organisation





norden

NordForsk



Nordic e-Infrastructure
Collaboration

Principles (Business Plan)

- Projects are organised in Activity Areas
- Projects are co-funded by partners
- Projects are led by NeIC Area Coordinator
- Activity in each Area is advised by User/Provider Forum
- Activity Areas are decided by NeIC Board



norden

NordForsk



Nordic e-Infrastructure
Collaboration

Principle of Co-funding

NeIC will provide base funding for projects

- In line with common Nordic (NordForsk) practice
- Co-funding from participating partner institutions
- 30-50% from NeIC, co-funding in-cash or in-kind
- Co-funding ensures
 - Quality and relevance
 - Risk-aversion
 - Variable geometry
 - Sustainability
 - Nordic coherence

Co-funds are quality-assessed funds



norden

NordForsk



Nordic e-Infrastructure
Collaboration

Nordic e-Infrastructure conference in Trondheim

13-16th of May:

Monday and Tuesday will be dedicated to workshops:

- Security workshop

- Center operations: discussing common challenges and best practices

- Data Services

- Science Gateways

- Bioinformatics..

Wednesday and Thursday will be an ordinary conference with presentations.

Generally: discussing ideas for NeIC activities and to actively explore opportunities for the future.



**PDC Center for
High Performance Computing**



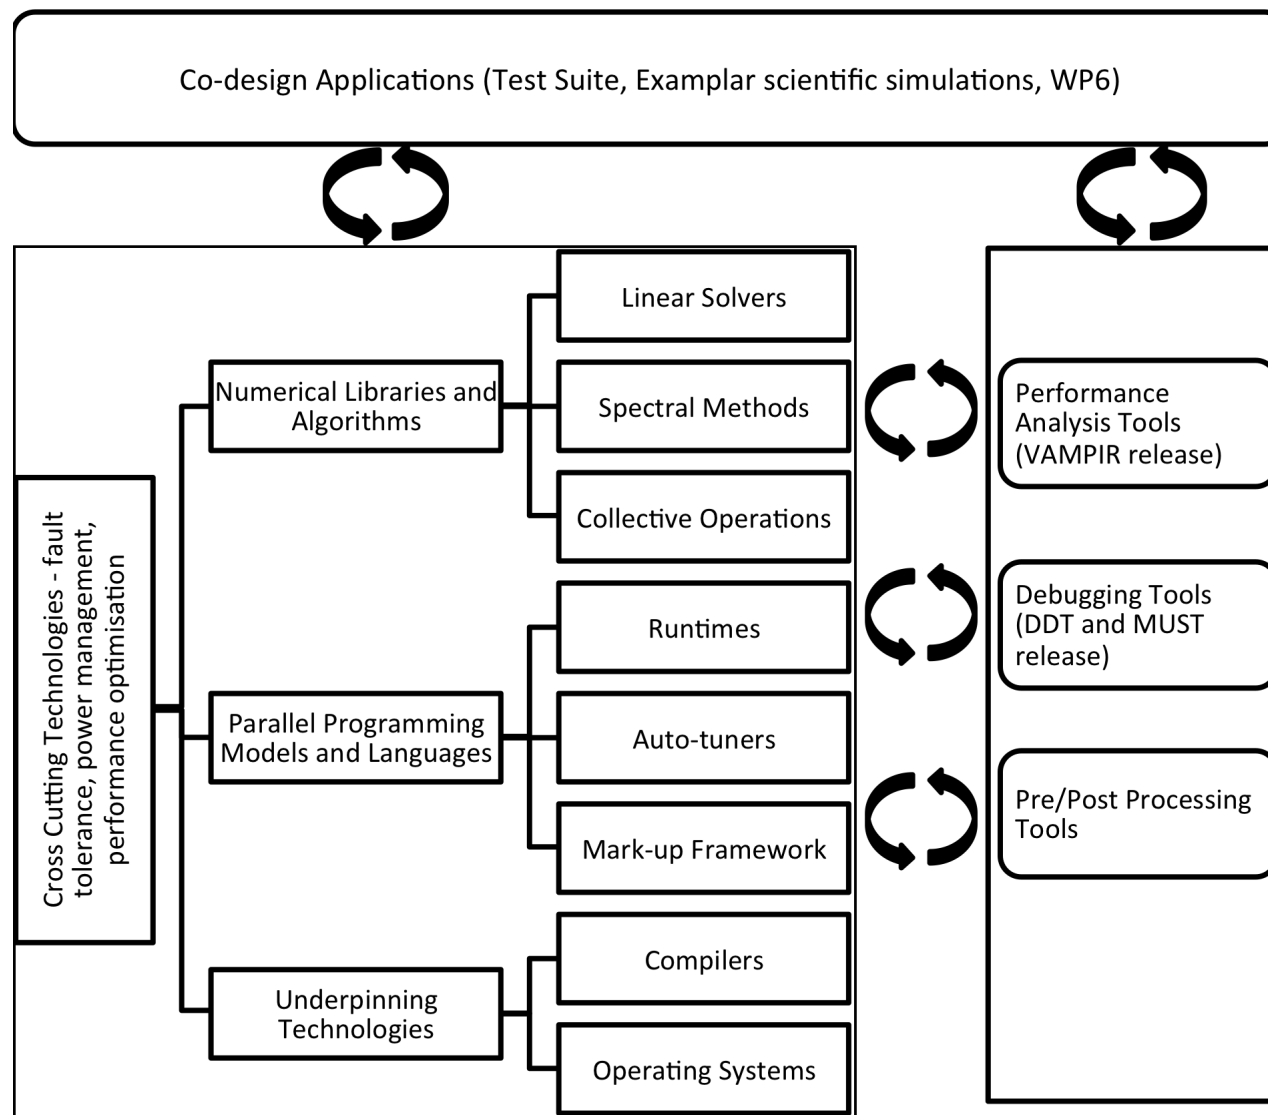
**National Supercomputer Centre
at Linköping University**

Selected New Activities

CRESTA objectives

- Main goal is to develop techniques and solutions which address the most difficult challenges that computing at the exascale can provide
- Success metrics by 2014
 - Co-design Applications
 - Co-design applications tested successfully on leading-edge petascale platforms and delivering previously unattainable simulations on those platforms
 - Roadmap to achieving application exploitation of exascale platforms
 - Systemware
 - Integrated CRESTA software stack successfully tested on petascale platforms
 - Co-design application simulations exploiting the CRESTA software stack to demonstrate massive and previously unattained scalability, reliability and usability
 - Roadmap to integrated CRESTA software stack on exascale platforms, reviewed by the Scientific Advisory Board

CRESTA development model





- Goes back to Nordic cloud **project**, NEON (2010)
 - Practical evaluation [1], testing public vs private cloud for eScience users (bioinformatics)
- SNIC Cloud **project** (2011.6-2012.6)
 - Enabled cloud access (public and private) to SNIC users.
 - 14 (some recurring) users of SNIC Cloud for Amazon [2] (e.g. running Galaxy) and 54 on the private cloud (currently only PDC Cloud, partially from outside SNIC)
- SNIC Galaxy **project** (2013.3-2014.3)
 - The goal of the project is to deliver Galaxy as a service, using the Galaxy cloud management platform, Cloudman, on local cloud installations (private clouds).
- SNIC Cloud **Infrastructure** (long-term, ongoing proposal)
 - A (generic) IaaS on which communities/users can build their PaaS. Strong emphasize on user communities and their commitment.


[1] "Practical Cloud Evaluation from a Nordic eScience User Perspective", VTDC'11, ACM conference San Jose (2011) by Åke Edlund and Maarten Koopman, Zeeshan Ali Shah, Ilja Livenson, Frederik Orellana, Jukka Kommeri, Miika Tuisku, Pekka Lehtovuori, Klaus Marius Hansen, Helmut Neukirchen, Ebba Þóra Hvannberg

[2] Amazon usage during the SNIC Cloud project: USD 2,110 (June 2011 – June 2012), USD 6,715 (July 2012 – January 2013)


SNIC Cloud Infrastructure

A new proposal to SNIC following previous experiences


The overall idea is to create a **(generic) IaaS** on which communities/users can build their PaaS. This would be a natural division of tasks and responsibilities between SNIC and users. The requirement is then that one finds a few communities that are indeed willing to build such PaaS.

A purple callout box with a rounded bottom and a tail pointing towards the text "The requirement is then that one finds a few communities that are indeed willing to build such PaaS.".

#1 User input and
commitment on
proposal



SNIC Project



Autoscaling is **off**.
Turn **on**?

The goal of the project is to deliver Galaxy as a service, using the Galaxy cloud management platform, Cloudman, on local cloud installations (private clouds).

Adding Galaxy onto SNIC Cloud Infrastructure – a private cloud environment - will give our life science users a number of benefits, especially with respect to storage and usability. The Galaxy environment will later be of use for other communities as well.

In addition we will implement the Galaxy Cloudman service, enabling the user to elastically scale his/her applications with the resources needed. Through the Galaxy Cloudman it will also be easier for user groups to administrate the allocated resources between the users.

About the project:

- Running 2013.3-2014.3
- Budget: 2,5 MSEK (2,75 FTE, one year) – PDC-HPC (coordinating), UPPMAX, and UPPNEX
- Building on SNIC cloud and will contribute to the further development of SNIC Cloud Infrastructure (see previous slide)
- Early users: from ScilifeLab, UPPNEX, ScalaLife
- Ongoing interactions with ELIXIR (UK and Sweden (BILS) nodes) and NeIC (with CSC Finland)



New Department at KTH

HPCViz

Bring together competences and resources at CSC to address current and emerging challenges for efficient use of large-scale computational resources, efficient and varied manipulations of massive data sets, and method and model development, combining theory- and data-driven methods

- Laure (Head), Hoffman (Deputy), Thuresson (Head VIC)
- PDC (HPC research/resources), CTL (simulation/software), VIC (visualization/interaction/collaboration)
- 2 Full Professors, 2 assoc. professors, 3 Senior Researchers, 3 PostDocs, 10 PhD Students, 1 Research Engineer, 1-2 administrators
- Upcoming (short term): prof. in visualization, 2 ass. prof., PhDs and postdocs
- Currently 13 courses on undergraduate and graduate level



**PDC Center for
High Performance Computing**



**National Supercomputer Centre
at Linköping University**

NSC & PDC

Thank you for your attention!

Erwin Laure
Director PDC-HPC

Patrick Norman
Director NSC