PDP Strategy

Introduction

The "Grid Computing" program at Nikhef is twelve years old as of this writing, if we take the start of the EU DataGrid project (January 2001) as the program start. The initial group activities were motivated by the needs of HEP computing, and the interest of the then-current group members. Later these activities were expanded to include middleware development as well as collaboration with non-HEP researchers, motivated by the successes of initial ad-hoc projects and collaborations, as well as (for the non-HEP collaborations) our desire to generate enough interest in a shared national e-infrastructure for science to make it a reality.

This strategy paid off handsomely with the approval of the BiG Grid project, which included the manpower and infrastructure costs necessary to build the Netherlands LHC Tier-1 computing center. We continued our investment in these collaborations with non-HEP researchers, partly due to the interesting research questions – particularly in the area of "security" – and partly to help ensure that the national infrastructure, including the Tier-1, would survive past the project stage, becoming a standard (and funded) part of the national scientific landscape.

The achievement of our goals, and the changing project and funding landscape for our group, make this a good moment to review our activities and define new objectives for the group. While it is quite difficult to make detailed predictions in the ICT world for more than a few years in advance, our aim was to document a high-level strategy that would serve us well for the coming ten years.

Strategic Directions

Scaling R & D

This direction concerns computing problems arising from some increase in scale. The specific dimension in which scaling is a problem is less relevant. This strategic direction exploits expertise present in our group. Given the historical scale increases observed in physics computing, and the likelihood that this trend continues, it is a safe choice.

In the last decade our scaling focus has been mostly on the grid infrastructure, with our group making significant contributions to the middleware used by wLCG. We envision a shift in this focus. wLCG has crystallized to a large extent; it is very difficult to make any changes, wLCG is mostly interested in service contributions from external providers. Furthermore while our group is a very strong one, it is small when compared to the whole of wLCG, resulting in a relatively low impact potential.

Targeting our efforts at local groups and other experiments has a much greater impact potential. Targeting local groups solves problems for our direct colleagues; at the same time our success for them generates useful publicity for uptake of our products outside of Nikhef. Targeting other experiments has a large overlap with "local groups" as we envision that the experiments we target (e.g. XENON, VIRGO, EUCLID) are experiments in which Nikhef (would like to) participate(s). In addition, most other experiments are much smaller than "wLCG" hence the relative impact of our effort is larger.

Participation in external projects, in the sense that it concerns something other than particle physics is more than welcome. However our contribution should be something that fits within other activities in the group, and for substantial effort, there should be associated funding (in some form).

Data management, GPU computing, and many-core computing are sub-areas that the group considers promising for future work in Scaling R&D.

Scalable Multi-Domain Security

This direction makes sense for our group for three main reasons:

1. History and reputation: Nikhef is well-known for this activity and we are recognized experts.
2. Experience has shown that since security touches most areas of infrastructure design, expertise in security gives our group a relatively large influence in major design decisions made by others.
3. Scalable multi-domain security is something needed by many researchers … of the software in the WLCG grid, the security (or better said the facilities for secure distributed collaboration) is the one thing consistently interesting to other (e.g. ESFRI) communities.

Identity federations, authentication/authorization infrastructure, and operational security are subareas in which we have expertise and for which the demand is relatively high. Our current products (glexec, LCAS, LCMAPS) are in maintenance phase, we will need to make some choices about which problem in SMDS we want to solve, and have some Nikhef "product" associated with that – this is necessary to maintain the reputation, expertise, and funding associated with the group.

Three focus areas for the security activity have been identified, areas for which we both have the expertise and where there is a relatively high demand from research communities and e-Infrastructure providers:

- Identity management (identifying users and resources and their properties and attributes) used as a basis for access control.

- Access Control and Provisioning: controlling access to (compute) resources for users beyond a single domain or community, addressing preservation of local control with global access, including the 'provisioning' across multiple

applications.

- Cross-domain operational security and incident response ("CSIRT"): addressing incidents beyond the network layer by designing methods for pro-active intervention based on foreknowledge of the community structure and behavior.

The SMDS direction also has excellent opportunities for outreach and societal visibility: the security community is a relatively small and closed niche, with at times widely different security topics closely coupled at the same conferences, in the same companies, or in the same individuals. This includes e.g. the combination of multi-domain incident response activities with cybercrime and law-enforcement; the use of some access control technologies for 'grid' are similarly to their use in the banking sector; and network security techniques are used not only to protect resource access in grids but also to connect secure remote workplaces for non-operational DoD personnel. It also leads naturally to other contacts, including some in the 'physical security' and other law-enforcement work. The strategic directions identified above will naturally lead to contacts in these directions, which are explicitly encouraged. Where specific work items emanate from these contacts, the work should be compatible with the global PDP aims (it should fit in one of the research areas, it should be fundable, and we should have the relevant expertise), and retain a link with 'security' in the broad sense of the word.

Where activities are deployed in this area, and where they use resources (people or hardware) from Nikhef or in our 'formerly-BiG Grid' infrastructure, our part of SURFsara S&D effort should be preferably considered to strengthen our 'embedding' with SURFsara S&D (and keep Nikhef visibility for these societally-relevant projects).

Collaboration and competition with other parties in the Netherlands needs special attention here. We expect that this work can be carried out in close collaboration with SURFsara and SURFnet. Especially for Dutch community support, SURFsara understandably positions itself as having all the answers here, and it is unclear whether NLeSC will choose collaboration or competition. The former is preferable and we should encourage this however possible. Our work on both identity management and for access control is well recognized at the European (and US) level, and as long as operational European funding is present we should be able to participate in these projects and in these areas. However, new European middleware development projects should be carefully considered to ensure the extra income merits the extra effort expended.

Activities in support of the Nikhef Mission

The Grid Computing program is a so-called "enabling physics program." These programs have in addition to their own research agenda, an important function in enabling advances in research to the physics-research programs. This function results in the following responsibilities for our group.

1. Operation of the Nikhef part of the NL LHC Tier-1 computing center, shared responsibility (with Sara) for planning & coordination of the Tier-1.
2. Computing support for local physics groups. A good example here is GPGPU work together with the Virgo and LHCb groups.
3. Targeted use of PDP group expertise for ICT infrastructure improvements at Nikhef. Here I refer to things like web sites, databases, mail servers and the like. The PDP involvement here is in the development phase; operational activities for these services fall under the CT group.
4. Partial support for the Stoomboot cluster. The PDP group is responsible for the batch scheduler configuration. We share responsibility for the large-scale storage cluster and network configuration (together with the core CT staff).

Where possible, these activities should be combined with one of the strategic directions. For example, the current pilot project concerning use of GPUs in the LHCb trigger (see 3 above) could be expanded to a full-blown "Scaling R&D" project for the PDP group if LHCb chooses to move this into production.

Supporting activities

This section lists activities we plan to carry out in order to support success of our stated strategic directions and mission-related activities.

Operation of technical infrastructure

We intend to continue to operate a significant computing "site" at Nikhef for several reasons:

1. having the infra here allows us to directly observe the interaction between the users and the stuff, as well as to know our users well.
2. Research on Scaling R&D requires having large-scale stuff on which to explore scaling properties
3. The facility serves an important PR function for both the group and for Nikhef. For example it is a popular stop at the Nikhef "Open Dag" and gives visitors a scale to which they can relate (they have computers at home too!)
4. It's a funding magnet: money sticks to stuff ….

Our "special ops" activity is an important part of this and should be continued as well. Members of the group are in close contact with computing, storage, and network vendors, and have pre-market access to new technologies. This privileged position has been earned by our demonstrated ability to discover the scalability limits of the hardware, and because what we ask for tends to be echoed in the larger HPC market a few years later, making Nikhef a valuable predictor of market trends. Our possession of, and expertise in, advanced hardware gives us an advantage in acquiring new projects.

Software Engineering

We intend to become more pragmatic (*i.e.*, less pedantic) in the area of software construction. The idea is to allow our own satisfaction with tools that solve 80% of the problem well, while maintaining high standards of portability, modularity, etc.

Dissemination of Expertise

People in our group are "experts", "pioneers", and "bleeding edge" with ICT technology compared to most Nikhef staff. Transfer of this expertise from our group to the staff at large – in a form that they can easily digest – creates a win-win situation. Some example ideas:

- Consulting on realizing a "scientific workspace" for our researchers – "dutch dropbox", "dutch evernote", etc. They need something that will work straight out of the box (like the non-dutch versions) but do not have all the legal / copyright problems associated with cloud services hosted in the US.
- Seminars on the "app of the month"
- Computer courses for new users
- Tropical Lecture contributions
- Other "internal outreach" stuff, such as periodic seminar explaining new and useful technologies / applications to Nikhef staff

Dissemination of Results

Sometimes we have failed to be recognized for our contributions, stemming from a lack of (healthy) self-promotion. What we do needs to be publicized well, as future opportunities for our group are in large part generated by our (perceived) track record. Promotion used to be much, as we were long ago more or less the only "big fish" in the Netherlands in the area of distributed computing, and one of a handful in Europe. Now there are many more parties who are quite effective at positioning themselves as "the" party in NL to consult on topics (e.g. Surf, NLeSC and SARA), and even more in Europe.

Specific areas of improvement have been identified:

- Our activities should be published whenever possible
- Top layer of PDP website needs to be redesigned and repopulated with dissemination in mind.
- Give talks wherever relevant and possible.
- Document activities via tech reports (when a paper is not appropriate) and disseminate widely.

Dissemination of vision

Partly this has to do with the formulation of our strategy as contained in this document – a refined version should be made public and should be reflected on our website. There is another aspect: explanation of why we do what we do, why it is important, and why Nikhef is doing this. This fits into the questions raised in our discussions about why is PDP considered a scientific program at Nikhef.

Much of what we do, in an academic sense, justified by the apparent absence of academic programs like "experimental computer science" or "software engineering" in the Netherlands. While the Dutch academic CS programs with which we have significant contact are active in the area of distributed computing, most of their "experiments" in this area are either simulations, or use something like the DAS system, which falls short of realistic scale in several dimensions. Hence these programs are more correctly identified as "theoretical computer science". "Software engineering" refers to an understanding of program construction and large-scale deployment of software. Activities in NL are also largely theoretical in this area, with the exception of the OOTI program in Eindhoven (with whom we already have a collaboration).

Assuming that the above impressions are correct, a consistent story along these lines positioning us as filling the gap gives us

- an academic "position" in the netherlands
- an important position in CS or ICT research, as a place where realistic experiments can be carried out. This allows us to attract PhD students (both Physics and CS)
- a line along which university contacts can be established.

Before making strong statements along these lines, some effort needs to be expended in understanding the role of TUs in the Netherlands and to what extent they might fill this gap.

Short-term commitments

A few of our current activities *must* be continued in the coming years due to commitments we have made by accepting hardware to operate on behalf of NWO or Surf, due to the "software provider" role we occupy in certain international grid projects, and via contracts signed with the EU. We list them here.

1. Operation of the NDPF (Nikhef Data Processing Facility a.k.a. "the grid cluster"), including the Nikhef part of the NL LHC Tier-1 service. We do not plan to stop this operation after three years, however even if we did, we are committed to operate the facility *at least* until the current hardware (awarded to Nikhef to run on behalf of NWO and later Surf) reaches end-of-life (about three years from now).
2. LCMAPS / LCAS / glexec software maintenance. We plan to continue maintenance of these products as long as there are "important customers" (like Fermilab / OSG / wLCG).
3. Operational Security (funded by EGI until april 2014).

Appendices

In the appendices we include some examples of projects along our 'strategic directions' that might be carried out in the coming three to five years.

Appendix A: Scaling Research and Development

- Manycore computing.

  The days of speed increases for computing tasks are more or less over. The per-core HS06 rating has increased only about 25% over the last four years. Machines become "faster" by adding cores; clusters become faster by running more jobs in parallel. As the number of cores per box, bottlenecks are encountered like access to memory or disk – this comes due to an increasing number of processes accessing the same resources, with no central coordination. Also overhead increases – 17 copies of an ATLAS reconstruction program running on a worker node require 17 copies of the conditions database in memory, 17 copies of the geometry, *etc.* – this is expensive and limits the amount of computing capacity an experiment can collect.

  Most groups with a significant computing load are starting to explore manycore computing, in which a single program instance manages threads executing on many cores in parallel, allieviating some of the problems due to overhead and contention. While there are enough groups already working on this for the CERN experiments, our group could play a role here for some other experiment.

- GPGPU computing

  GPGPUs are the processors found on graphics cards in ordinary PCs. For a special class of computing problems, these processors can be much faster than general purpose CPUs. We already have some collaboration with the VIRGO people on this, we follow the work being done in ANTARES, and a new recent project has been started together with LHCb to investigate how their trigger would perform on GPGPU hardware. Any of these efforts, or related ones (AUGER?) might develop into a full-blown project.

- Data Management

  Most HEP experiments still have a problem dealing with distributed data. The system used in WLCG 'works' but is far from optimal, requiring a large manpower effort to keep all the various catalogues in a consistent state and to deal with disk crashes that occasionally permanently lose data. Nikhef has been involved in "data management next generation" discussions for quite some time. It is difficult to play a major role here since the LHC experiments would like to construct the new system themselves, however we follow

discussions in this area, looking for some area of interesting research in which our group could make a major impact. We expect that we could also play a major role here for some non-HEP experiment. EUCLID is a good example; if TARGET fails to deliver (their track record would assign a reasonable probability to this), our group could play a major role here.

Appendix B: Scalable Multi-Domain Security

- Identity management (identifying users and resources and their properties and attributes) used as a basis for access control.

  In this area we have ongoing work in both policy work and software implementation & support which gives us a strong starting position. The goal is ensuring a transparent 'market' for resource sharing across all research domains, without obstacles for users when switching between local and global resource access, or when organizing the mapping of their research community on the e-Infrastructure. At the same time this work aims to prevent a 'policy gap' opening up between wLCG and the developing e-Infrastructures (also beyond EGI).

  Existing opportunities in this area include collaboration with SURFnet on community structuring, with ESFRI projects (in particular CLARIN via MPI Nijmegen, but also via EUDAT) on identity management mechanisms, and with a large (and more amorphous) group of EGI, OSG, wLCG, PRACE, XSEDE, REFEDs, TERENA and others on global policy coordination and a 'single identity market'. The associated 'product' is in the form of documents, support and the recognition thereof. We do not have nor foresee large software products in this area beyond integration. For the Dutch community our most visible product ("jGridstart") ensures Nikhef visibility.

- Access Control and Provisioning: controlling access to (compute) resources for users beyond a single domain or community, addressing preservation of local control with global access, including the 'provisioning' across multiple applications.

  In this area we have invested significantly through our site access control suite (LCAS/LCMAPS/gLExec) which is extensively used but now moving to a maintenance phase. The goal is enabling structured community formation and enabling access to (global) resources for these communities whilst respecting both community structure and local access decisions. The emphasis is on compute and web-based services (mainly because raw data access is too broad to consider). With more 'general purpose' infrastructures emerging (including our own Dutch infrastructure), communities will be discouraged to build their own silos. So the need for generic access control is likely to stay. wLCG is an obvious case in point, but the work with MPI touches the same point.

  New opportunities in this area include building access control mechanisms for

new (ESFRI) communities – building on the MPI work – or even access controls to (virtual) networks for cloud services. Also taking ownership of (or building) base libraries in the security area for inclusion in higher-level products is a way to 'embed' the group in key products that have big impact and thereby ensure visibility and impact (trying to make the group 'unavoidable', akin to what was done with for dCache with jGlobus). For the time being the associated products are still LCAS/LCMAPS/gLExec, and the MPI OAuth service. New work should be identified in this area, preferably linked to federated identity management with SURFnet.

- Cross-domain operational security and incident response ("CSIRT"): addressing incidents beyond the network layer by designing methods for pro-active intervention based on foreknowledge of the community structure and behavior.

  In this area we have a prominent position both nationally and internationally. The work done in EGI coordination (we have the sole EGI Security Officer), the Security Service Challenges, and the participation in incident response (also beyond pure grid activities) is very visible. The goal is twofold: ensure pro-active protection of our own ICT infrastructure (grid and non-grid) by being engaged, and secondly foster security and availability of services for research globally. Given the high impact of security incidents (and their visibility), this is essential for all infrastructures we participate in (including wLCG and EGI).

  New opportunities in this area include building up explicit training services for resource owners (as part of a joint pan-European effort with our collaborators), joint operations with SURFnet on incident response, and taking on the CSIRT role on campus. This is also closely linked with the non-PDP ICT activities at Nikhef. We are not building concrete software product in this area, but instead create procedures and policies which have high visibility in the grid (e.g. EGI) and non-grid community (e.g. at FIRST conferences).