

On the sustainability of large-scale computer science testbeds: the Grid'5000 case

David Margery, Frédéric Desprez

► **To cite this version:**

David Margery, Frédéric Desprez. On the sustainability of large-scale computer science testbeds: the Grid'5000 case. 2016. <hal-01273170>

HAL Id: hal-01273170

<https://hal.inria.fr/hal-01273170>

Submitted on 12 Feb 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Abstract

In this position paper, we look at the financial sustainability of Grid'5000. The duration of the project (over 12 years) owes more to successive investment decisions and continued support rather than from a well understood and operated business model generating enough revenue to covers costs and investments.

In this paper, we will give an overview of a typical cost structure for a large-scale testbed, we summarize with a few statements our views and develop pros and cons of a few funding sources. The way Grid'5000 is funded is detailed, before giving some data used to compute the unit cost for Grid'5000 resources.

On the sustainability of large-scale computer science testbeds: the Grid'5000 case

D. Margery, F. Desprez

February 12, 2016

Introduction

In this position paper, we look at the financial sustainability of Grid'5000[1]. With initial users using the platform in 2004, wide availability following the first winter school in Grenoble in March 2006, and hardware still being purchased in 2016, we can say that in a sense Grid'5000 has proven surprisingly sustainable.

But when one looks behind the scenes, the duration of the project owes more to successive investment decisions and continued support especially from Inria rather than from a well understood and operated business model generating enough revenue to covers costs and investments.

In this paper, we will give an overview of a typical cost structure for a large-scale testbed in section 1. In section 2, we summarize with a few statements our views. We develop in section 3 pros and cons of a few funding sources. In section 4, the way Grid'5000 is funded is detailed, before in section 5 giving some data used to compute the unit cost for Grid'5000 resources. Section 6 concludes with the next steps towards achieving sustainability for Grid'5000.

1 An overview of the cost structure

To understand this discussion, we will first enumerate and discuss the cost structure of running Grid'5000. The identified cost items are the following:

- Hardware costs. This includes
 - the server room, including cooling infrastructure
 - the servers to operate the infrastructure
 - the servers made available to users
 - the network fabric
 - the storage and backup infrastructure
- The maintenance costs. While for servers, maintenance costs can be paid upfront for their useful lifetime, it is often not the case for the rest, and there is an annual cost to pay to keep an active maintenance contract for them after an initial period. For some costs (fire protections), there might be a mandatory requirement.
- The utility costs, electricity in particular.
- Network connectivity costs, between sites and with Internet
- The costs for the personnel required to keep the instrument up and running while keeping the same functionality. This includes the costs incurred by doing software upgrades to avoid the accumulation of technical debt. We are mainly concerned here about direct costs, but indirect costs are also important when looking at the global picture.
- The costs to develop the instrument further, adding new services or exploring alternative approaches, to keep the instrument relevant to its target user community and to extend the applicability to other communities.
- The costs for travel to technical meetings and dissemination events.

Contrary to other sciences, the cost structure of a computer science testbed over time cannot be simplified as an initial upfront cost to buy and setup the scientific instrument, followed by annual operating costs. This is because a vast majority of the upfront costs is spend on servers that needs to be renewed 3 to 5 years after their initial purchase.

2 Opening Statements

1. Unless a stable cash flow is available to a computer science testbed to cover for hardware renewal, it will be on average operated with around no more than 2.5 years of visibility before end of life. This is very uncomfortable for prospective users when they consider the testbed for experiment-driven research for a significant research program. If you take into account that one year can pass between proposal writing phase and project start for funded research programs, this is clearly insufficient visibility.
2. As developed in the sustainability work of the BonFIRE and Fed4FIRE projects, a sustainable testbed will probably use a mix of funding sources **TODO: Citations needed, despite the fact the deliverable are not all public:**
 - (a) direct funding, where funding agencies, universities or research operators are convinced of the usefulness of the testbed to its users and will therefore fund it.
 - (b) indirect funding, where funding agencies, universities or research operators fund the users of the testbed, and users then in turn direct funding to the testbed.
 - (c) pay-for-use, where users pay to use resources on the testbed
 - (d) consulting revenue, where the testbed is paid to help or to completely operate an experiment at the request of a user.
3. The initial setup cost of a new testbed is much more important than the cost to renew its hardware, keep it running and adding new features. Or to say things differently: the initial investment is much more in bringing things together to build a testbed rather than in the cost of the initial hardware. This is badly understood by funding agencies and governing bodies.

3 Discussion

3.1 Direct Funding

While at first this seems the most comfortable option for a testbed, and seems inevitable to establish a new testbed, this funding model comes with its own problems, as the testbed will always need to answer the question of whether it is still useful to fund it.

An important work is therefore required from the testbed operator to collect usage reports from its users (looking into the past) and future usage intentions (what research will be impacted if the testbed is no longer funded).

In a sense, users pay to use the testbed by providing relevance indicators in the form of reports or by declaring publications made using the testbed. These are summarized and put forward to convince funding agencies to support a testbed or future testbed.

We have found in Grid'5000 that it is very difficult to collect such relevance indicators when the number of users is high, as the link between users and the people operating or promoting the testbed can only become less and less direct as the number of users grows. Grouping users into projects might help scale a bit, but our experience is that even people that are very aware of the issue of collecting relevance indicators (members of some Grid'5000's committees) sometimes forget to check that usage is properly reported by their students when they publish their work. The issue only gets worse when federation is considered, as relevance indicators are rarely included in the federation mechanisms.

3.2 Indirect funding

In the indirect funding model, users are given a way to spend resources on any testbed they need to perform the work they are funded for. The exact cash-flow could be indirect (for example with testbeds collecting a virtual currency giving rights to receive direct funding, as explored by NIH (**TODO: track reference of this program**)). The key insight here is that funding agencies no longer need to assert the relevance of a testbed, as as long as users are willing to spend a limited resource to use the testbed, it must be because it is relevant.

Of course, this model also comes with issues, as testbeds are now at the same time competing for users and setting the price for usage. A newly funded testbed might be tempted to undercut other testbeds to prove its usefulness, or enter into other kinds of mutually destructive practices when setting the price.

We have also discovered that this model is not without administrative challenges:

- Setting the unit price is useless if it cannot be asserted as eligible costs for funding. If the rules vary between funding agencies, things could get tricky.
- the same is true about asserting usage. Can access rights, usage of access rights, CPU usage, data usage, network usage all be legitimately used as quantities to which a price is attached ? Is it acceptable to invoice the price of a complete package if it is only partially used ?
- If the testbed operator is a large research institute or university, it might be at the same time the funded user (or a member of the funded project), and the testbed operator sending invoices to users from a legal point of view, even if different departments are involved. From an accountant's or auditor's point of view, this is a legal entity asserting a cost for which it is setting the price and certifying the quantities. Unless acceptable rules are defined, this will seem suspect.

As it is expected that the majority of users of a scientific instrument will be academics, the chances of having to handle the rules and regulation hurdles described here are high. This cannot be done without forward thinking by funding agencies.

3.3 Pay-for-use

Should enough users be in a situation where they are able and willing to pay for usage of a testbed at the conditions set by the testbed operator, it may be possible for a testbed to come closer to sustainability.

It is only possible to envision such a future if these users are not able to find a similar service from a testbed with free access to resources, or even worse (but very common due to the Open Call mechanism in the Horizon 2020 framework program) a testbed that is willing to pay users to be used. Here in essence we have mature and established testbeds looking at pay-for-use business models that have to compete with funded testbeds trying to establish themselves and that might even be willing to fund users to attract them.

Suppose one moment that a testbed has unique selling points that make it attractive to users despite an unfavorable eco-system? To stay attractive, it will need to protect these unique selling points instead of sharing them with other testbeds. This is not good to promote experiment portability, experiment reproductibility and other fundamental challenges that need addressed on the road to open science.

Of course, we should expect funding agencies to **not** fund the development of testbeds that directly compete with existing testbeds. But as feature overlap is inevitable, or proposals wanting to replicate existing work but with an additional feature that is essential to a given scientific field, we expect some degree of competition for at least part of the user base of a given testbed.

We therefore conclude that pay-for-use revenue cannot be the main pillar towards sustainable scientific testbeds. It can nevertheless be part of the cash-flow mix of a given testbed.

3.4 Consulting

Designing and running good experiments is difficult. One not only needs access to a suitable testbed, but also to a lot of expertise. As understanding the performance and behavior of current and future computer systems is also a requirement for some businesses, it is expected that some form of consulting could generate revenue for testbed operators. We understand iMinds in Belgium have such services to offer.

If a testbed's staff is experienced enough to offer such services, and that they can work in the context of a consulting contract, this might also generate revenue for the testbed.

4 Grid'5000's situation

Grid'5000 is sustained the following way:

- Hardware renewal costs are not covered by a regular budget. They are currently partially covered by a 900 000 grant by Inria, complemented by regional funds in Rennes, Nancy and Lille. A distinct operation paid by regional funds pays for the development of the site in Nantes.
- Maintenance costs are not covered. All hardware is bought including a 5 year maintenance contract. After those five years, clusters slowly die. We are at risk on some sites as the network equipment is no longer covered by a maintenance contract.
- The utility costs, electricity and maintenance of the cooling equipment, are covered by the hosting entities. The real cost is unknown to Grid'5000 management.

- Network connectivity costs, between sites and with Internet are covered by RENATER.
- The personnel cost to keep the platform running is covered by Inria, Université de Rennes 1 and CNRS (through LIP lab in Lyon) for the 1.3 full-time equivalent (FTE) in permanent positions, and by Inria for the 3 other engineers and 2 apprentices. We submit a request for the non-permanent personnel regularly through the InriaHub call for projects (previously known as ADT). One engineer is also financed through an European project.
- The development cost to extend the platform is covered by Inria, paying the technical director (0.8 FTE) and 2 engineers on temporary positions. Another project is submitted to the InriaHub call for projects for the new developments, with a description of the developments aimed for.
- The costs for travel, technical workshops and dissemination is paid by the 7 members of the Grid'5000 GIS (a formally defined scientific interest group), each paying 5000 a year. Minor hardware purchases are sometimes also paid on this budget.

Currently,

- Grid'5000 receives a small amount of funding for the BonFIRE and Fed4FIRE projects, used by Grid'5000 to understand and integrate in the European eco-system. This has had marginal benefit to the overall funding of Grid'5000, as most activities carried in these project have not had direct impact on Grid'5000. Three of the engineers working for these European projects have also been able to contribute occasionally to the work of the technical team.
- No indirect funding mechanism has been found workable for Grid'5000, but we are taking part in a proposal that will explore this topic further in an Horizon 2020 project built on the foundations of Fed4FIRE and due April 2016.
- Inria has signed, for Grid'5000, a first contract with a SME that is willing to use Grid'5000 in a pay-for use model.
- Lucas Nussbaum is consulting for the SME cited above, but with no arrangement in place to extract revenue for Grid'5000 through that consulting contract signed by University of Lorraine.

5 Calculating a cost for Grid'5000

In this section, we detail how a cost of 0,1 per core hour was established. On February 11, 2016, our statistics database reports that over the duration of the project:

- 264 333 207 core.hours have been used on Grid'5000. Note that these are not normalized to account for different processor performance, and data for older jobs is not complete.
- 6 093 870 was spent on personnel, outside European projects. This is an estimated value, with an average cost of 5000 per man month.
- 7 738 911 was spent on hardware. For older purchases, the funded amount was used instead of the final purchase price.
- 12 000 000 spent by RENATER on network. This is the less precise estimate. Each site has a 10Gb/s uplink to the core network, and Grid'5000 has recently its own uplinks (2) to the core network.
- 457 420 has been spent on travel and small equipment. Exact figures are used between 2009 and 2014, and the average of these figures is used for the other years taken into consideration.

When all those costs are added, and divided by the number of core.hours, you get a cost of 0.0995

6 Future work

On the topic of sustainability, the next steps envisioned are the following

- Check for the existence of rules in the context of Horizon 2020 to set an auditable unit price for the use of a scientific testbed
- Send a first invoice for resource usage by the SME that is using Grid'5000 in a pay-for-use mode. Check Inria can receive the money without a formal decision by its Conseil d'administration
- Understand how any cash generated by the pay-for-use customer using Grid'5000 must be spent. In particular, whether it must be spent on the year it has been received.
- Contribute to the successor project of Fed4FIRE to promote a sustainable eco-system for testbeds in Europe

References

- [1] Franck Cappello, Eddy Caron, Michel J. Daydé, Frédéric Desprez, Yvon Jégou, Pascale Vicat-Blanc Primet, Emmanuel Jeannot, Stéphane Lanteri, Julien Leduc, Nouredine Melab, Guillaume Mornet, Raymond Namyst, Benjamin Quétier, and Olivier Richard. Grid'5000: a large scale and highly reconfigurable grid experimental testbed. In *GRID*, pages 99–106. IEEE, 2005.