



Estimation of Costs and Pay per Use on a Large-scale Shared Computer Science Testbed: the Grid'5000 Case

David Margery, Lucas Nussbaum

► To cite this version:

David Margery, Lucas Nussbaum. Estimation of Costs and Pay per Use on a Large-scale Shared Computer Science Testbed: the Grid'5000 Case. 2019. hal-02011425v2

HAL Id: hal-02011425

<https://inria.hal.science/hal-02011425v2>

Preprint submitted on 20 Feb 2019

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.

Estimation of Costs and Pay per Use on a Large-scale Shared Computer Science Testbed: the Grid'5000 Case

David Margery and Lucas Nussbaum

February 2019

Executive summary

Grid'5000 is a large-scale science testbed: multiple academic organizations contribute to some of the costs associated to operating and maintaining it, to the point many approximations must be made when attempting to build a view of the overall costs associated to it. However, such an approximated cost has been estimated, first in 2015, and more recently in 2017 and 2018 as part as a standardized infrastructure cost evaluation for all large scientific instruments on France's national research infrastructure roadmap. From this approximated cost, we can derive pricing.

However, this does not imply that we can make users pay for their usage of Grid'5000, as: (1) many users are affiliated in a way or another to the organizations that are contributing to supporting the costs of the infrastructure; (2) the way costs were approximated does not qualify Grid'5000 usage as an eligible cost in the context of funded projects as it is not auditable.

In practice, this only leaves as potential paying users: (1) foreign academics with no ongoing collaboration with their French colleagues, and (2) private sector users. Being a minority of the user base, we cannot expect a significant revenue stream from such categories of users. In theory, we could also hope to invoice organizations who benefit more than they contribute to costs, but in practice, such a scheme would be difficult to operate.

In conclusion, we do not believe that we can transition to a pay-for-use model for a majority of Grid'5000 users without a major increase in bureaucracy to associate usage to affiliation and a reduction to the number of organizations supporting the costs of the scientific instrument. This would certainly have a negative impact on usage, increasing the number of custom-built mini-testbeds and therefore the overall cost to the research ecosystem.

Contents

1	Introduction	2
2	Context: The Grid'5000 Testbed	2
3	Motivations	2
4	Challenges	3
4.1	Accurate Data Collection	3
4.2	Academic Culture of Invisible Costs	3
4.3	Other Incentives for not Using Pay-per-Use Infrastructures	3
4.4	Loss of Visibility for the Infrastructure	3
4.5	Quantifying Contributors' Contributions	3
4.6	Splitting of the Revenue Among Contributors	4
4.7	Legal Challenges Related to the Funding Sources	4
4.8	Complex Billing Scheme	4
5	The Grid'5000 experience	4
5.1	Evaluation of Costs	4
5.1.1	Data available	4
5.1.2	Initial cost estimation	5

5.1.3 Annual cost estimation	6
5.2 Definition of a Usage Cost for Pay-per-Use	6
6 Discussion and conclusions	6

1 Introduction

When discussing the costs of research infrastructures, it is often suggested to work on estimating costs in order to develop a pay-per-use model, that supposedly would be able to solve funding issues. This document discusses this issue in detail, detailing possible motivations, outlining some limitations and challenges, and illustrating this with the specific case of Grid'5000's attempts to estimate costs and establish a pay-per-use model.

This document is organized as follows. Section 2 provides a short description of the Grid'5000, in order to provide some context and some background for the assumptions made throughout the document. Section 3 discusses motivations for estimating costs and establishing various kinds of pay-per-use models. Section 4 then articulates limitations and challenges raised by those models. Finally, 5 describes Grid'5000's experience in that context, before the document is concluded in 6.

2 Context: The Grid'5000 Testbed

Grid'5000 is a distributed computer science testbed. It is composed of 8 sites, hosted by several institutions (nationwide research institutions, and several universities, sometimes in joint labs between research centers and universities). Seven sites are located in France, one is in Luxembourg. Each site is composed of several clusters (homogeneous sets of machines that are usually purchased at the same time). Different funding sources are used for the clusters: regional projects (sometimes combined with European funding), own resources of one of the institutions, own resources of one research team at one of those institutions, etc. The engineering workforce is composed of a mix of permanent staff, and staff on short-term temporary contracts. Engineers are employed by the institutions that also host sites.

The major sources of cost for the infrastructure are:

- The purchase of hardware resources;
- The hosting of hardware resources (such as the operation of the server room itself, and electricity costs);
- The engineering workforce required to develop and maintain the infrastructure's software stack;
- The engineering workforce required to operate the testbed and its services.

3 Motivations

The motivations for estimating costs are described below.

Objective 1: Advertising costs to users to increase self-awareness

The use of the infrastructure would remain free from the point of view of users, but, in order to increase the self-awareness of users and their teams about the important costs generated by their use of the infrastructure, some kind of *usage report* or *virtual billing* would be done, with an estimation of the generated costs, without resulting in actual billing.

Objective 2: Billing some users of the infrastructure

The use of the infrastructure would be billed to some categories of users, such as:

(Objective 2.1) Users from the private sector

(Objective 2.2) Academic users affiliated to institutions that do not contribute to the infrastructure

(Objective 2.3) All users (including those affiliated to institutions that contribute to the infrastructure)

Objective 3: Using the estimated costs for reimbursement in the context of contracts

The costs generated by the use of the infrastructure could be considered in the context of contracts (industrial collaborations, collaborative projects funded by ANR or UE), in order to be compensated.

4 Challenges

4.1 Accurate Data Collection

(Affects all objectives, but especially Objective 3)

Collecting correct data about costs can be extremely difficult. Some costs are trivial to identify (for example, the initial purchase costs) but raise questions on how to amortize them over the lifetime of a resource. Some costs are general to the infrastructure (for costs related to the purchase and operation of the network infrastructure, or of the service machines) and must be shared by all resources. Other costs are hard to evaluate, such as electricity costs in a data center that hosts both machines for the testbed, and machines for other purposes. More examples of difficulties are provided in Section 5.

This might not be a severe issue for Objective 1, as long as the costs provide a good estimate. It might also not be a problem for Objective 2, as long as the costs are fairly accurate. However, it might be a severe problem when costs must be justified to a third party such as a funding agency.

4.2 Academic Culture of Invisible Costs

(Affects Objectives 2.2 and 2.3)

Most costs related to the life of research teams are generally hidden from the research teams themselves, and amortized over the whole research institution. For example, costs related to accessing publications, or to hosting people, are generally not billed to the teams.

Now let's suppose that the use of the testbed is billed to a researcher, who is a member of a French public research organization that does not contribute to funding the infrastructure. That researcher could be tempted to build her own testbed, hosting it in her lab, with local human resources. It is likely to be a less expensive solution for her, because many of the costs will be covered transparently by her lab. However, the overall cost for the French public research ecosystem would be higher due to the duplicate work and the lack of consolidation between testbeds.

4.3 Other Incentives for not Using Pay-per-Use Infrastructures

(Affects Objective 2.2 and 2.3)

There are other incentives for not using pay-per-use infrastructures. First, researchers have an incentive to build their own infrastructure, because the building of testbeds (even private testbeds that duplicate the functionality of other pre-existing testbeds) is usually seen favorably, for example when the researcher's lab is evaluated.

It might even be possible for the researcher to get paid to use other infrastructures. For example, many research projects on testbeds design have *Open Calls* where prospective users of those testbeds can be funded to use the testbed (and provide early feedback to the testbed operators).

4.4 Loss of Visibility for the Infrastructure

(Affects Objectives 2.1 and 2.2)

Providing free access to the testbed helps lower the barrier to entry, and increasing the visibility (especially the international visibility) of the testbed. Free access to the infrastructure can also act as a *honeypot* to attract collaborations with industry. One could even question the interest of asking industrial users for pay-per-use if it is possible, instead, to build mutually fruitful collaborations.

This could, of course, be mitigated by providing a *free preview* mode (enabling prospective users to explore the infrastructure capabilities by providing free access to the infrastructure for limited amounts of time and resources), but then care must be taken to ensure that users do not abuse the system by creating multiple accounts.

4.5 Quantifying Contributors' Contributions

(Affects Objectives 2.2 and 2.3)

The split between users from institutions contributing to the infrastructure, and users from institutions that do not contribute to the infrastructure, is not that simple.

First, there is a wide spectrum of contribution levels. If all contributions are considered the same, then this acts as an incentive for institutions to minimize their costs and avoid renewing hardware (if an old cluster is sufficient to be considered a contributing institution, what is the point of replacing it?). Therefore it is important to evaluate and differentiate contributions.

Second, not all affiliations are the same. Many users have multiple affiliations due to, for example, joint teams between several institutions. Also, there are nationwide research institutions with several sites (locations): if some sites contribute to an infrastructure, should users from other sites be identified as users from contributing institutions?

Third, not all institutions have the same user base. How should contributions be matched with the usage volume?

This could be solved by measuring the value of contributions and providing some kind of virtual billing even for users of contributing institutions, to ensure that the costs resulting from the usage of the infrastructure do not exceed the contributions (or else, the usage would be billed as for users of non-contributing organizations). But this is a complex scheme, which raises additional questions, such as which users should be billed when the institution's quota is exceeded.

4.6 Splitting of the Revenue Among Contributors

(Affects Objectives 2 and 3)

Another problem is the splitting of revenue among institutions contributing to the infrastructure. Should it be split according to the value of contributions, or to the actual resources used? The latter is harder to implement (especially since some elements are shared over all resources, such as development costs), but can act as an incentive to provide more interesting resources.

4.7 Legal Challenges Related to the Funding Sources

(Affects Objectives 2 and 3)

Some sources of funding that could have been used to purchase parts of the infrastructure have additional requirements, such as: forbidding to generate revenue from the resources, or forbidding commercial use. Special care would have to be taken to respect those additional rules, for example when allocating usage to resources.

4.8 Complex Billing Scheme

(Affects Objective 3)

When attempting to justify costs in order to claim their reimbursement from a funding agency, there are two different institutions at play: (A) the institution that provides the infrastructure (assuming, for simplification, that a single institution is in charge of the billing on behalf of all institutions providing the infrastructure); (B) the institution of the user of the infrastructure.

If (A) and (B) are different, and if the user also has free access to the infrastructure, there is no real incentive for the user and (B) to declare their use of the infrastructure to the funding agency (and thus pay for the use of the infrastructure), because it would not benefit (B), but (A). (B) also has an incentive not to claim costs for the use of the infrastructure, because projects generally operate on a budget fixed in advance, so reimbursement going to (A) is funding lost for (B).

If (A) and (B) are the same institution, it requires the funding agency to accept the evaluation of the costs made by (A/B) as correct, because (A) cannot invoice itself to present an invoice to the funding agency. Such a scheme might be an issue for some funding agencies, even if the French ANR seems to accept such scheme¹.

Finally, the only case where this works well is if (A) and (B) are different institutions, and (B) is always paying for its use of the infrastructure. But in that case, the fact that (B) gets reimbursed by the funding agency is invisible from the point of view of (A).

5 The Grid'5000 experience

This section describes the Grid'5000 experience of evaluating its costs (Section 5.1), and then for implementing a pay-per-use model for a company (Section 5.2).

5.1 Evaluation of Costs

5.1.1 Data available

The Grid'5000 testbed has a database with data that are crucial to estimate the costs of running Grid'5000.

¹See ANR – Règlement financier, section 3.1.3

- Total of core.hours used for each day, with complete data going back to 2007 and partial data for years 2005 and 2006.
- Price and date of purchase for all equipment since the start of the project. For the initial clusters, the price of purchase is estimated using the amount of budget available for the purchase. Errors should be minimal.
- For every engineer having worked on the project, dates and proportion of time devoted to the project, and whether he or she worked as an admin, a developer, a manager or for a project connected to Grid'5000.

This dataset is, however, missing some crucial information:

- The cost of the network interconnection between Grid'5000 sites, and between Grid'5000 and the rest of the internet. It has been guesstimated at 100k€ a year for each 10G lambda in 2015.
- A trace of the involvement of the different scientists in the management of the project. For all that follows, this cost has been neglected.
- The salaries and overhead costs for all the staff having worked on the project. There are too many organizations involved to be able to track that information, rather confidential by nature, in the project's stats database.
- The annual costs of hosting a Grid'5000 site: cooling infrastructure, maintenance of the cooling infrastructure, electricity, network link to the Renater's Point of Presence). No research institution or university hosting Grid'5000 is organized to estimate hosting costs at the granularity of a rack or to monitor electricity consumed. In the best case, we have a one-time measure of the proportion of electricity consumed by Grid'5000 over a one-month period on the overall consumption of the buildings.

Even if all that information was available, some difficult choices need to be made to compute a cost.

- Choosing the units of costs. For this document, the unit chosen is the core.hour. This makes it free to store data on Grid'5000, and does not differentiate between cores based on their performance or when they give access to a large amount of memory, accelerator cards, specific local storage performance, or high performance network interconnect. Should the cost be passed on to users, optimization based on this simple unit of cost would surely lead to a lot of undesirable side effects.
- Choosing how to depreciate the value of the hardware and of the engineering invested to build the infrastructure. When staff is only seen as an operating cost, the figures fail to capture the value of all the processes, software infrastructure and know-how required to build a research testbed. Value of depreciated but still used hardware must also be captured.

5.1.2 Initial cost estimation

First in 2015, and then at the start of 2016, a cost for Grid'5000 was computed based on the Grid'5000 history since 2005, without taking into account inflation. This has led to an initial estimation of a cost of 0,1 € per core.hour, using the following data extracted on February 11, 2016 from our statistics database:

- 264 333 207 core.hours had been used and registered on Grid'5000 since 2005. 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 person.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 10 Gb/s uplink to the core network, and Grid'5000 has its own uplinks (two) to the core network.
- 457 420 € had 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 € per core.hour.

5.1.3 Annual cost estimation

For the last 2 years, the French Ministry of research has made an attempt at standardizing the way the annual costs of all research infrastructures on the national roadmap. For this reason, we have an total annual cost for Grid'5000 since 2016 (published in 2017) and 2017 (published in 2018). This cost takes inflation into account but not networking costs as the National Research and Education Network (RENATER) also presents its costs, and that networking costs for Grid'5000 are included in RENATER's costs.

For 2016, the reported annual cost for Grid'5000 was 1.904 M€ (in 2015 €). In the same year, 47 993 263 core.hours were used on the infrastructure, putting the cost of each core.hour used that year at 0.0396 €. Even if networking was factored in, we would not reach the 0.1 € for each core.hour computed using the previous method.

For 2017, the reported annual cost for Grid'5000 was 1.845 M€ (in 2017 €). In the same year, 41 903 708 core.hours were used on the infrastructure, putting the cost of each core.hour used that year at 0,0440 €, when network is not factored in.²

5.2 Definition of a Usage Cost for Pay-per-Use

Using the data from the initial cost estimation, Grid'5000 governing body, the *Conseil de Groupement du GIS Grid'5000*, has formally approved selling access to Grid'5000 at the price of 0.1€ for each core-hour used.

Because of that decision, Inria, hosting the governing body, is able to invoice users of Grid'5000 provided they have agreed to use Grid'5000 on a Pay-per-Use basis.

Formalizing that agreement is the difficult part. A contract has been written to do so, and signed once with a private company. The contract allowed the company to buy the right to use a pre-defined number of core-hours. Grid'5000 would be allowed to block usage should they exceed that amount, but only real usage would be invoiced.

For history, it should be noted that because of the expertise required to run proper experiments, they started financing a PhD in a collaboration with a university involved in Grid'5000. They therefore acquired the right to use Grid'5000 freely before the first invoice.

We therefore have a contract to sell access to Grid'5000 on a Pay-per-Use basis, but no identified clients.

6 Discussion and conclusions

This document outlines the challenges of estimating costs and building a pay-per-use model. It is worth stressing that moving forward in that direction would require a lot of expertise that is not currently found in the current Grid'5000 workforce: collaboration with services responsible for building management, and with legal and financial services to understand all the legal consequences. Also, if significant revenue is expected from paying customers, it would also require the help of commercial staff to advertise the infrastructure and extend the customers base. Finally, administrative staff would be required to operate the scheme on a daily basis.

If continued, the current scheme (free use) could nevertheless be improved. Specifically, we could develop schemes that encourage institutions to contribute to Grid'5000, by providing higher priorities to users from contributing institutions. Such a scheme was outlined in [1], but has not been implemented yet, mainly due to lack of development resources in order to implement the required changes in the Grid'5000 services.

References

- [1] Lucas Nussbaum. Usages et utilisateurs de Grid'5000: stratégie pour l'accès aux ressources. Available online at <https://hal.inria.fr/hal-01294910>, March 2016.

²As a data point, an Amazon EC2 instance of type `m5.metal` (with 96 cores) in the Paris AWS region costs 5.375 USD per hour, or 0.056 USD per core.hour. That does not include data transfer or remote storage (EBS) costs.