Top 10 Rules for Building a Sustainable Grid

Wolfgang Gentzsch
RENCI, Duke & D-Grid, October 2007

After publishing a report on major grid initiatives on the RENCI webpage earlier this year, I am getting emails from colleagues who are currently setting up a grid infrastructure (or are planning to do so), be it in research or in industry enterprises, asking questions especially about a good strategy for making grids sustainable. Here, I am trying to answer by presenting "my" 10 important rules on how to build a sustainable grid (and I know there are more than 10). These rules are derived from mainly four sources: my research on major grid projects published on the RENCI website, the e-IRG Workshop on "A Sustainable Grid Infrastructure for Europe" in Heidelberg in spring, the 2nd International Workshop on Campus and Community Grids at OGF20 in Manchester, and my personal experience with coordinating the German D-Grid Initiative. The rules presented here are mainly non-technical, because I believe most of the challenges in building and operating a grid are in the form of cultural, legal and regulatory barriers.

Rule 1: Identify your specific benefits. Your first thought should be about your users and your organization. What's in it for them? Identify the benefits which fit best: transparent access to and better utilization of resources; almost infinite compute and storage capacity; flexibility, adaptability and automation through dynamic and concerted interoperation of networked resources; cost reduction through utility model; shorter time-to-market because of more simulations at the same time on the grid. Grid technology helps to adjust an enterprise's IT architecture to real business requirements (and not vice versa). For example, global companies will be able to decompose their highly complex processes into modular components of a workflow which can be distributed around the globe such that on-demand availability and access to suitable workforce and resources are assured, productivity increased, and cost reduced. Application of grid technology in these processes, guarantees seamless integration of and communication among all distributed components and provides transparent and secure access to sensitive company information and other proprietary assets, world-wide. Grid computing is especially of great benefit for those research and business groups which cannot afford expensive IT resources . It enables engineers to remotely access any IT resource as a utility, to simulate any process and any product (and product life cycle) before it is built, resulting in higher quality, increased functionality, and cost and risk reduction.

Rule 2: Evangelize your decision makers first. They give you the money and authority for your grid project. The more they know about the project and the more they believe in it (and in you) the more money and time you will get, and the easier becomes your task to lead and motivate your team and to get things done. Present a business case (current deficiencies, specific benefits of the grid (see Rule #1), how much will it cost and how much will it return, etc. They might also have to modify existing policies, top down, to make it easier for users (and providers) to cope with the challenges of and to accept and use the new services. For example, why would a researcher (or a department in an enterprise) stop buying computers when money continues to be allocated for buying it? This policy should be changed to support a utility model instead of an ownership model. If you are building a national grid, for example, convincing your government to modify its research funding model is a tough task.

Rule 3: Don't re-invent wheels. In the early grid days, many grid projects tried to develop the whole software stack themselves: from the middleware layer, to the software tools, to grid-enabling the applications, to the portal and Web layer…and got troubled by the next technology change. Today, so many grid technologies, products and projects exist that you want to start looking for similar projects, select your favorite (successful) ones which fit best

your users' needs, and 'copy' what they have built, and that will be your prototype. Then, you might still have some time and money left to optimize it so it fully matches the requirements of your users. Consider, however, that all grids are different. For example, research grids are mainly about sharing (e.g. sharing resources, knowledge, data), commercial enterprise grids are about cost and revenue (e.g. TCO, ROI, productivity). Therefore, if your community is academic, look for academic use cases, if it's commercial, look for commercial use cases in your respective business field.

Rule 4: KISS (Keep It Simple and Stupid). It took your users years to get acquainted with their current working environment and tools. Ideally, you won't change that. Try hard to stick with what they have and how they do things. Plan for an incremental approach and lots of time listening and talking. Social effects dominate in grids. Join forces with the system people to change/modify mainly the lower layers of the architecture. Your users are your customers, they are king. Differentiate between two groups of users: the end users who are designing and developing the products (or the research results) which account for all the earnings of your company (or reputation and therefore funding for your research institute), and the system experts who are eager to support the end users with the best possible services. Therefore, you can only succeed if you demonstrate a handful of clear benefits to these two user groups.

Rule 5: Evolution, not revolution. As the saying goes: "never change a running system".. We all hate changes in our daily lives, except when we are sure that things will drastically improving. Your users and their applications deeply depend on a reliable infrastructure. So, whenever you have to change especially the user layer, only change it in small steps and in large time cycles. And, start with enhancing existing service models moderately, and test suitable utility models first as pilots. And, very important, part of your business plan has to be an excellent training and communications strategy

Rule 6: Establish a governance structure. Define clear responsibilities and dependencies for specific tasks, duties and people during and after the project. An advisory board should include your representatives of your end-users as well as application and system experts. In case of more complex projects, e.g. consisting of an integration project and several application or community projects, an efficient management board should lead and steer coordination and collaboration among the projects and the working groups. The management board (Steering Committee) should consist of leaders of the sub-projects. Regular face-to-face meetings are very important.

Rule 7: Money, money, money. Don't have unrealistic expectations that grid computing will save you money initially.. In their early stage, grid projects need enough funding to get over the early-adopter phase into a mature state with a rock-solid grid infrastructure such that other user communities can join easily. In research grids, for example, we estimate this funding phase currently to be in the order of 3-5 years, with more funding in the beginning for the grid infrastructure, and later more funding for the application communities. In larger (e.g. global) research grids, funding must cover Teams or Centers of Excellence, for building, managing and operating the grid infrastructure, and for middleware tools, application support, and training. Also, today's funding models in research and education are often project based and thus not ready for a utilitarian approach where resource usage is based on a pay-as-you-go approach. Old funding models first have to be adjusted accordingly before a utility model can be introduced successfully. For example, today's existing government funding models are often counter-productive when establishing new and efficient forms of utility services (see Rule #2). In the long run, grid computing will save you money through a much more efficient, flexible and productive infrastructure.

Rule 8: Secure some funding for after the end of the project. Continuity especially for maintenance and support are extremely important for the sustainability of your grid

infrastructure. Make sure at the beginning of your project that additional funding will be available after the end of the project, to guarantee service and support and continuous improvement and adjustment of the infrastructure.

Rule 9: Try not to grid-enable your applications in the first place. Adjusting your application to changing technologies costs a lot of money and takes a lot of your precious time. Did you macro-assemble, vectorize, multitask, parallelize, or multithread your application yourself in the past? This is not what you or the user should do. Better to use the money to buy (lease, rent, subscribe to) software as a service.or to hire a few consultants who grid-enable your application and/or (even better) help you enable your grid architecture to dynamically cope with the applications and user requirements (instead vice versa). Today, in grids, we are looking more at chunks of independent jobs, (or chunks of transactions). And we let our schedulers and brokers decide how to distribute these chunks onto the best-suited and least-loaded servers in the grid, or let the servers decide themselves to share the chunks with their neighbors automatically whenever they become overloaded. .

Role 10: Adopt a 'human' business model. Don't invent new business models. This usually increases the risk for failure. Learn from the business models we have with other service infrastructures: water, gas, telephony, electricity, mass transportation, the Internet, and the World Wide Web. Despite this wide variety of areas, there is only a handful of successful business models: on one end of the spectrum, you pay the total price, and the whole thing is yours. Or you pay only a share of it, but pay the other share on a per usage basis. Or you rent everything, and pay chunks back on a regular basis, like a subscription fee or leasing. Or you pay just for what you use. Sometimes, however, there are 'hidden' or secondary applications. For example, electrical power alone doesn't help. It's only useful if it generates something, e.g. light, or heat, or cold, etc. And this infrastructure is what creates a whole new industry of new appliances: light bulbs, heaters, refrigerators, etc. Back to grids: providing the right (transparent) infrastructure (services) and the right (simple) business model will most certainly create a new set of services which most probably will improve our quality of life and make a few companies and some creative individuals very rich.

References:

2nd International Workshop on Campus and Community Grids (Draft), Eds: Laura McGinnis, David Wallom, Wolfgang Gentzsch, http://forge.gridforum.org/sf/go/doc14617?nav=1

Grid Initiatives: Lessons Learned and Recommendations, Report, RENCI Website: http://www.renci.org/publications/reports.php

A Sustainable Grid Infrastructure for Europe, Executive Summary of the e-IRG Open Workshop on e-Infrastructures, Heidelberg, Germany. e-IRG Website: http://www.e-irg.org/meetings/2007-DE/workshop.html

D-Grid Website: http://www.d-grid.de/index.php?id=1&L=1