



HAL
open science

DSL-Lab: a Platform to Experiment on Domestic Broadband Internet

Gilles Fedak, Jean-Patrick Gelas, Thomas Herault, Victor Iniesta, Derrick Kondo, Laurent Lefèvre, Paul Malecot, Lucas Nussbaum, Ala Rezmerita, Olivier Richard

► **To cite this version:**

Gilles Fedak, Jean-Patrick Gelas, Thomas Herault, Victor Iniesta, Derrick Kondo, et al.. DSL-Lab: a Platform to Experiment on Domestic Broadband Internet. [Research Report] RR-7024, INRIA. 2009. inria-00424936

HAL Id: inria-00424936

<https://inria.hal.science/inria-00424936v1>

Submitted on 19 Oct 2009

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.



INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

DSL-Lab: a Platform to Experiment on Domestic Broadband Internet

Gilles Fedak — Jean-Patrick Gelas — Thomas Herault — Victor Iniesta — Derrick Kondo
— Laurent Lefèvre — Paul Malécot — Lucas Nussbaum — Ala Rezmerita — Olivier
Richard

N° 7024

Août 2009

A large, light blue stylized 'R' logo is positioned to the left of the text. The text 'Rapport de recherche' is written in a serif font, with 'Rapport' on the top line and 'de recherche' on the bottom line. A horizontal line is drawn under the text.

*Rapport
de recherche*

DSL-Lab: a Platform to Experiment on Domestic Broadband Internet

Gilles Fedak^{*†}, Jean-Patrick Gelas[‡], Thomas Herault^{*}, Victor Iniesta^{*}, Derrick Kondo^{*§}, Laurent Lefèvre[‡], Paul Malécot^{*}, Lucas Nussbaum^{‡§}, Ala Rezmerita^{*}, Olivier Richard[§]

Thème : Calcul distribué et applications à très haute performance
Équipes-Projets Grand-Large, GRAAL, MESCAL

Rapport de recherche n° 7024 — Août 2009 — ?? pages

Abstract: This report presents the design and building of DSL-Lab, a platform for distributed computing and peer-to-peer experiments over the domestic broadband Internet. Experimental platforms such as PlanetLab and Grid'5000 are promising methodological approaches for studying distributed systems. However, both platforms focus on high-end services and network deployments on only a restricted part of the Internet, and as such, they do not provide experimental conditions of residential broadband networks. DSL-Lab is composed of 40 low-power and noiseless nodes, which are hosted by participants, using users' xDSL or cable access to the Internet. The objective is twofold: 1) to provide accurate and customized measures of availability, activity and performance in order to characterize and tune the models of such resources ; 2) to provide an experimental platform for new protocols, services and applications, as well as a validation tool for simulators and emulators targeting these systems. In this report, we report on the software infrastructure (security, resources allocation, power management) as well as on the first results and experiments achieved.

Key-words: DSL, Experimental platform, Desktop Grid, P2P

* INRIA Saclay-Île de France, Grand-Large, Orsay, F-91893 – LRI, Univ Paris-Sud, CNRS, Orsay, F-91405

† INRIA Rhône-Alpes, GRAAL, Montbonnot, F-38334 – LIP, ENS-Lyon, CNRS, UCB Lyon 1, Lyon, F-69364

‡ INRIA RESO, University of Lyon, Ecole Normale Supérieure de Lyon, France

§ INRIA Rhône-Alpes, MESCAL, Montbonnot, F-38334 – LIG, CNRS, Grenoble INP, INRIA, UJF, UPMF, Grenoble, F-38041

DSL-Lab : une Plate-Forme pour l'Expérimentation sur les Réseaux Domestiques à Haut Débit

Résumé : Ce rapport décrit la conception et la mise en oeuvre de DSL-Lab, une plate-forme d'expérimentation pour le calcul distribué et le Pair-à-Pair sur des réseaux Internet à haut débit domestiques. Si les plates-formes expérimentales comme PlanetLab et Grid'5000 offrent une approche méthodologique prometteuse pour étudier les systèmes distribués, ces deux plates-formes se concentrent sur les services à haute performance et les déploiements réseaux ne se font que sur une portion réduite d'Internet. Donc elles ne proposent pas les conditions expérimentales correspondant aux réseaux à haut débits résidentiels. DSL-Lab est composé de 40 noeuds silencieux et à basse consommation qui sont hébergés par des volontaires en utilisant leurs accès xDSL ou câble à Internet. L'objectif est double : 1) fournir des mesures précises et personnalisées de la disponibilité, de l'activité et des performances afin de caractériser et améliorer les modèles décrivant de telles ressources ; 2) fournir une plate-forme expérimentale pour de nouveaux protocoles, services et applications, ainsi qu'un outils de validation pour les simulateurs et émulateurs visant de tels systèmes. Dans ce rapport, nous décrivons les infrastructures logiciels (sécurité, allocation de ressources, gestion de l'énergie) ainsi que les premiers résultats et expérimentations effectuées.

Mots-clés : DSL, plate-forme expérimentale, Grille de PC, Pair-à-Pair

This report presents the design and building of DSL-Lab, a platform for distributed computing and peer-to-peer experiments over the domestic broadband Internet. Experimental platforms such as PlanetLab and Grid'5000 are promising methodological approaches for studying distributed systems. However, both platforms focus on high-end services and network deployments on only a restricted part of the Internet, and as such, they do not provide experimental conditions of residential broadband networks. DSL-Lab is composed of 40 low-power and noiseless nodes, which are hosted by participants, using users' xDSL or cable access to the Internet. The objective is twofold: 1) to provide accurate and customized measures of availability, activity and performance in order to characterize and tune the models of such resources ; 2) to provide an experimental platform for new protocols, services and applications, as well as a validation tool for simulators and emulators targeting these systems. In this article, we report on the software infrastructure (security, resources allocation, power management) as well as on the first results and experiments achieved.

1 Introduction

In the last decade, we have seen the development of distributed applications that allow the sharing of computing resources (CPU, storage or communication) at large-scales (GRID or P2P). Due to novel deployment platforms (institutional networks or the Internet) and applications (e-science, file sharing), many critical research issues around large-scale distributed systems have arisen.

In parallel with this evolution, high-speed Internet access has become common in homes; ADSL (Asymmetric Digital Subscriber Line) lines are widespread and fiber optic communication is now reaching the consumer market. The progress realized by these technologies allows Internet providers to offer their customers an Internet connection comparable, in terms of bandwidth, to 10 years old local network (up to 100Mb/s). However, the network architecture of home PCs interconnected by ADSL presents special characteristics:

- the physical characteristics of the network differ substantially from the LAN characteristics, already well-studied, because of the asymmetric communication performance (download/upload), the induced latency, and the internal ISP topologies ;
- within each family home, users share their Internet connection between several machines, using a wired and/or wifi local network as well as NAT and firewalls to protect their network ;
- new classes of network appliances, beside the regular PC join this network: Wifi phones, media centers and IPTV, Network Attached Storage, networked gaming consoles, etc. . . Furthermore, the network resource might be shared among several network-demanding applications (VoIP, P2P, gaming).

From the methodological point of view, evaluating the impact of broadband Internet characteristics over distributed applications presents several challenges. Building an experimental testbed is an approach that has proved to be essential as evident by successful projects such as PlanetLab [?] and Grid'5000 [?]. PlanetLab is a distributed platform for evaluating new classes of network services

such as content delivery networks, overlay networks, network measurement services and many more. PlanetLab interconnects dedicated servers, spread over Internet, usually made available by research laboratories. Grid'5000 is another experimental platform to study distributed systems composed of 5,000 CPUs distributed over 8 sites in France. Grid'5000 is designed as a cluster of clusters interconnected by 1 Gb/s to 10 Gb/s dedicated fiber links. The platform forms a VPN to launch experiments and the entire software stack is totally reconfigurable and customizable. For security issues, Grid'5000 is totally isolated from the Internet, in contrast with PlanetLab.

We think that there exists the need for a complementary approach to PlanetLab and Grid'5000 to experiment with distributed computing in an environment closer to how Internet appears, when applications are run on end-user PCs. Our approach consists of placing computing nodes at users' homes, which are called DSLnodes, available permanently for scientists to launch experiments. At the present stage, the platform deployed is composed of 40 nodes, covering the major French Internet Service Providers. It features a large variety of network connectivities (NAT, firewall) with several *DSL boxes*¹ technologies. The DSL-Lab platform appears to its users as a private virtual cluster, managed by a specialized version of OAR [?], a job scheduler which allows collaborators to launch experiments.

In this report, we describe the architecture of DSL-Lab, the design decisions and implementation concerning hardware, operating system, security, network connectivity, resources management, power management and experiment deployment. We also report on early experiments and preliminary results.

The rest of the document is organized as follows: in Section 2 we describe the architecture of DSL-Lab, in Section 3, we present the early results and ongoing experiments, in Section 4 we present Related Works, and we conclude in Section 5.

2 DSL-Lab Architecture

Because the major part of the platform is hosted at individual's homes, we had several requirements when designing the platform. We call *volunteers* our users who agreed to host DSLnodes and *experimenters* the researchers who are using the platform.

- DSLnodes have to be as unobtrusive as possible in terms of appearance, size, noise and power consumption;
- The platform should not compromise the security of volunteers by making their network reachable from the outside world;
- Node management should be done entirely remotely and should not require any on-site intervention from the volunteers.

¹DSL boxes such as Freebox, Neufbox, Livebox or Dartybox are the equipment lent by Internet providers to their users. These terminals usually feature more than a simple DSL modem and offer a large set of applications such as local network router, firewall, NAT, port forwarding, wifi radius, community chili-spots, NAS, telephony such as VoIP, SIP, media center, TV over IP, Video on Demand, etc...



Figure 1: Neo CI852A-4RN10 barebone used for the DSLnodes.

- For experimenters, the platform should appear as a traditional cluster in terms of providing all the tools they are familiar with.

To meet the aforementioned requirements, we had to select specialized hardware so that it would be powerful enough to conduct all our experiments, but low profile enough to avoid causing trouble for volunteers. We selected the Neo CI852A-4RN10 barebone, which belongs to the Mini-ITX class of PC, characterized by a small form factor, absolutely silent operation (due to the absence of a fan or moving parts), and a low-power processor. It is powered by an Intel Celeron M 1 GHz processor, 512 MB of RAM and storage is ensured by a 2 GB Compact Flash disk. The node also features 4 Ethernet ports so that it can be potentially used as a router to conduct networking experiments. An interesting feature is that this hardware does not allow volunteers to interact with the nodes, which prevents the node from being used for personal purpose other than scientific experiments.

The DSL-Lab central server is connected to the French academic backbone (Renater) through the University Paris-Sud network. This backbone is large and peers almost directly with all French ISPs, so network quality depends mostly on the ISP and the volunteer's last mile.

As of January 2009, the DSLnodes were distributed as follows: 32 were hosted on the French major DSL providers (Orange, Free, Neuf, Tele2), 1 on the French cable network (Numericable), 2 were hosted on the US cable network, 3 were distributed to Orsay, Lyon and Grenoble laboratories for development purposes and the remaining 2 were waiting for new volunteers. Some of our nodes may switch to fiber networks in 2009.

The DSL "network" is very heterogeneous and offers a wide range of different environments. There are 4 technologies allowed in France: ADSL, ADSL2, ADSL2+ and ReADSL. In addition, some volunteers have the ability to fine tune their ADSL link. For instance, some are using the "fastpath" which reduces the ping delay but is more prone to transmission errors. Also the phone line length (from 50 m to several kilometers for the set of DSLnodes) and the domestic electrical environment impacts the quality of the DSL link and available bandwidth and latency.

In this project, we have developed our own software stack. We now describe the features provided by our architecture.

2.1 Remote OS Deployment

One of the major concerns when designing the DSL-Lab platform was to avoid as much as possible volunteer intervention on the nodes. So, we needed to be able to re-install (in case an experimenter breaks the installed system by mistake) and upgrade (for security reasons, or to install additional software) the whole software stack, including the operating system installed on the nodes, without requiring volunteer intervention.

Our solution is to use two disk partitions, for two different systems:

- a small partition (5 MB) with a read-only, minimal system (we chose to use a customized TTYLinux), that is used for remote OS deployment and as a fallback in case of problems. As such, the experimenter is not allowed to modify it in any circumstance. This minimal system has been thoroughly tested to ensure maximum reliability;
- a larger partition, with the experimentation system (we use Debian GNU/Linux), that experimenters are allowed to modify.

The bootloader is configured to alternate booting on the first and the second partition.

The first partition will never be modified. It contains code that connects to the central DSL-Lab server, checks for an updated version of the experimentation operating system or a special request to re-install the DSLnode, and proceeds accordingly. After checking if re-installation is necessary, the DSLnode reboots the experimentation system, and awaits the experimenter's instructions.

Using that system, it is possible to safely and completely upgrade all DSLnodes remotely. It is possible, for example, to switch to another Linux distribution.

However, our current implementation is not hacker-proof. For example, a malicious experimenter could easily make the experimentation system unbootable by killing the software watchdog which is installed in the system. A solution could be to use a hardware Watchdog to solve this case, but the one provided on our DSLnodes failed to work. Another problem is that the minimal system is not write-protected. Hence, a malicious experimenter could also modify that system. PlanetLab, for example, uses a read-only medium (CDROM, write-protected USB key) to avoid this problem.

2.2 Connectivity and Security

The platform is managed in such a way that only identified experimenters have access to it. Experimenters first log into the server, which acts as a gateway and provides remote access to each DSLnode within a VPN through SSH.

To minimize the number of listening ports visible from the Internet, and entering in conflict with the volunteer security policy, there is no publicly accessible SSH daemon running on DSLnodes. Instead, the DSLnode establishes an SSH connection, including an SSH tunnel to its own SSH server port, to the DSL-Lab central server. That SSH TCP tunnel can later be used by authenticated experiments to connect to the DSLnode using SSH. (see Figure ??). This creates a VPN that connects the nodes together without modifying volunteers' security settings. To prevent the drop of the reverse connection caused by a

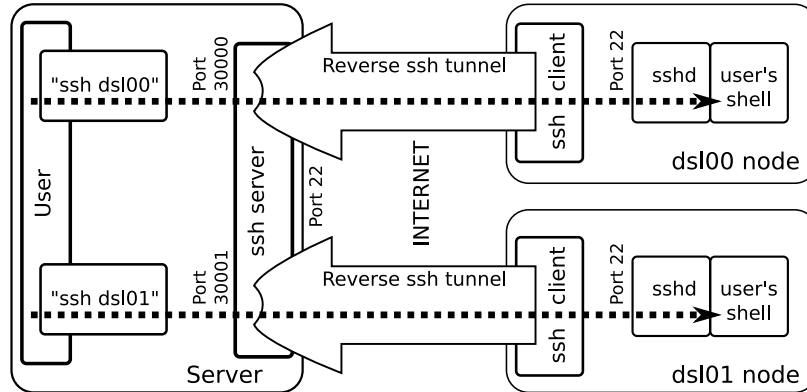


Figure 2: Establishment of connectivity layer within DSL-Lab using tunnel and reverse ssh.

TCP timeout in the case of no activity, we configured SSH to use a small keep-alive. In addition, the connection aliveness is checked every 10 minutes, and the ssh client from the DSLnode to the DSL-Lab server is restarted if necessary.

As the nodes are distributed to volunteer's home, we do not trust the software and private keys that they contain. One malicious volunteer (or the person who compromised their network) may retrieve the private key and try to connect to the SSH server with it. For this reason, the key only permits to create the SSH tunnel and cannot execute anything on the DSL-Lab central server.

On DSLnodes, only 4 ports, 2 TCP and 2 UDP, are inevitably used for experiments. Even if the experimental software behind those ports are not proven to be secure, they are running for short periods of time and with unprivileged accounts. To open direct connections between hosts, we are integrating a new solution called PVC (Private Virtual Cluster), which is described in the next section.

2.3 Resources and Power Management

Most of the DSL-Lab experimenters have are familiar with the Grid'5000 platform. To leverage their knowledge acquired on Grid'5000, we have adapted the Grid'5000 batch scheduler, called OAR, to the DSL-Lab platform so that: 1/ experimenters have a similar work environment and 2/ it would eventually facilitate the connection of both platforms. Thanks to OAR, several experimenters may reserve some nodes in advance and deploy their experiments simultaneously.

Because DSLnodes are hosted on a volunteer basis, a request of the volunteers is that the DSLnode does not waste power. Besides selecting thrifty hardware, the system ensures that the DSLnodes are powered-off when not used, thus reducing electricity consumption. To do so, DSLnodes periodically boot and check with the server if they have been reserved using OAR. The node then stays up (if the next reservation starts soon enough), or re-schedules its next wake-up time accordingly. Wake-ups are done using ACPI alarms, and have proven to be very reliable. Even if not reserved, the nodes wake up on a regular basis to check if new reservations have been created. Their wake-up

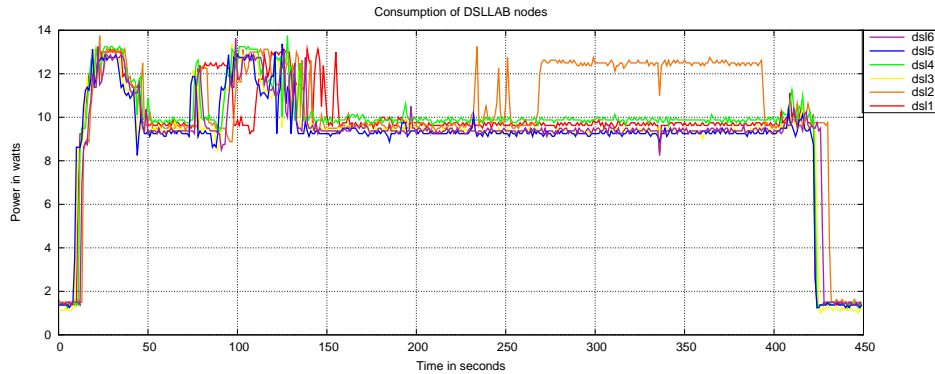


Figure 3: Power consumption of 6 DSLnodes during the boot process, low or full usage (during CPU intensive usage), and shutdown.

times are distributed within the work day, so that experiments are guaranteed to be able to access a node in less than 30 minutes, even without reserving a node previously.

3 Firsts Experiments and Results

3.1 Power Consumption of the platform

First, we have evaluated the power consumption of DSLnodes, in order to price DSLnode hosting as our volunteers are not refunded for the electricity consumed. Linux's frequency scaling is enabled using the on-demand governor and nodes are shutdown when unused by experimenters. Figure ?? presents real power usage as measured with a power sensor [?]. A node consumes around 1.5W when turned off (due to wake-on-lan and software power switch) ; 9-10W when the CPU is idle and 13-14W at 100% CPU load. According to the French regulated electricity price, the cost per hour is 0.00182€ if the DSLnode is reserved and 0.000195€ otherwise. If we assume usage of the platform to be 8 hours a day, 5 days a week, 11 months a year, it costs 3.57€/year for a volunteer to host a DSLnode.

3.2 Bandwidth

In our this experiment, we characterized our network by building a communication matrix between nodes. Figure ?? shows the DSLnode to DSLnode communication bandwidth² ISPs advertise that the ADSL download bandwidth is up to 28Mbit/s ATM and 1Mbit/s ATM for upload. But these bandwidths are not always observed in practice because of phone line length and attenuation. As a consequence, we observe that performance varies from 16Mb/s (if one end is located in a laboratory) to 256Kb/s (TCP/IP).

²At the time of this experiment, we were only able to conduct the measurement on a restricted set of DSLnodes where volunteers have agreed to modify their firewall settings in order to allow this measurement. In the future, thanks to PVC, we will be able to conduct similar experiments without requiring users to open specific ports in their network.

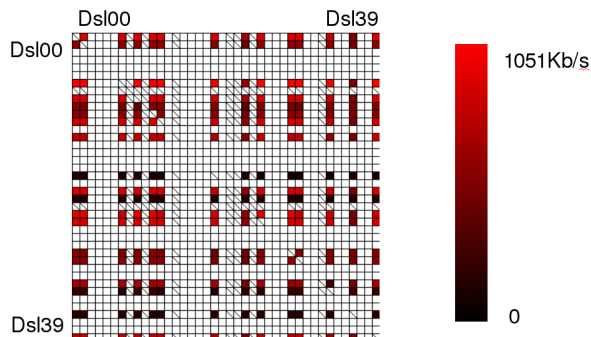


Figure 4: Communication matrix between all nodes of DSL-Lab. The color shows the TCP/IP bandwidth averaged over 8 measures.

3.3 Private Virtual Cluster (PVC)

The objective of Private Virtual Cluster (PVC) [?] is to provide, in a transparent way, an execution environment for existing cluster applications over nodes distributed on the Internet. Thanks to its capability to dynamically connect firewall-protected machines, without any intervention of domain administrators, without modification of the existing security policy, and with minimal performance overhead, PVC can serve as a middleware in order to test the network performance, with standard tools.

In order to validate PVC implementation, two types of experiments were performed on the DSL-Lab platform. The first one evaluated the overhead of the system for the connection establishment using a specific test suite. The second one demonstrated PVC's capabilities to execute unmodified applications (NAS benchmarks, MPIPOV, DOT) deployed over a set of machines behind firewalls blocking incoming connections, connected to the Internet by ADSL connections.

3.4 Evaluation of Desktop Grid Platforms

DSL-Lab is the platform of choice to evaluate experimental Desktop Grid middleware. This approach complements Grid'5000 where the experimental conditions are controlled but unlike residential broadband-network conditions. BitDew is a middleware dedicated to data management on Desktop Grid. In addition to traditional micro-benchmarks performed on Grid'5000, we have run several real-life scenarios, including multi-source file distribution, fault-tolerance, collective file distribution with replication and wide-area file transfers featuring P2P protocols (Bittorrent) as well as Amazon S3 and IBP protocols. Thanks to DSL-Lab, we have shown that our middleware was able to cope with complex scenarios in situations close to real-world conditions.

In partnership with C. Cerin from Université Paris XIII, we are now evaluating a new Desktop Grid protocol called BonjourGrid [?], which is dedicated to service discovery. In addition to Desktop Grid middleware, we are also working with A. Imanitchi from the University of Florida to evaluate a prototype of a middleware called MobiP2P, dedicated to data management in the field of mobile computing.

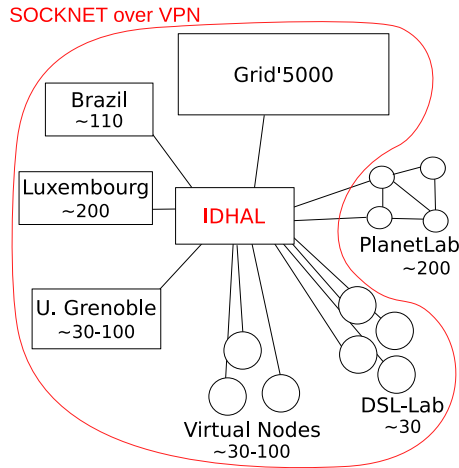


Figure 5: IDHAL topology: several platforms were linked together using either a VPN tunnel to a central node, or using TakTuk’s communication layer.

3.5 Integrating DSL-Lab with other platforms: the IDHAL experiment

In the IDHAL experiment, we investigated the possibility of connecting various experimental platforms together, namely DSL-Lab, Grid’5000, PlanetLab, desktop computing nodes located at a university in Grenoble, clusters in Luxembourg and Porto Alegre (Brazil), and virtual machines. Two challenges arose: being able to connect those machines together despite the differences in the way the platforms are connected together, and using all the platforms together to do a distributed computation. To connect all the platforms together, we combined two approaches (Figure ??). On most of the nodes, we set up a virtual private network (VPN) tunnel to a well-connected single host using *OpenVPN*. (We had the possibility to use several hosts instead of one, but this didn’t prove necessary.) On the PlanetLab nodes, where setting up a tunnel was not possible, we used TakTuk [?] to set up an overlay communication network.

After connecting all the nodes together, we focused on using them all for a single computation. We chose to use the KAAPI library [?], and instances of the N-Queens problem. Since the nodes have very different resources (both compute power and network), it doesn’t make sense to simply split the computation equally on all nodes. Instead, the work-stealing capabilities of KAAPI allowed it to efficiently share the computation on all nodes. While the goal of this experiment was mainly to prove the feasibility of using so many different resources together, we still managed to solve several instances of the N-Queens problem, for $n = 19, 20, 21$, using different combinations of nodes. During these experiments, we found several limitations, and gained perspective for future improvements, in the tools we used, in particular in the area of fault tolerance.

4 Related Works

Prior to DSL-Lab, some platforms were already dedicated to scientific research on large-scale Internet distributed systems. In this last section, we describe some of them and how DSL-Lab differs by providing a complementary experimental setup.

Grid'5000 [?] is a French academic cluster of clusters spread over 9 French locations. It gathers about 4800 cores of 4300 processors that are part of 1600 nodes. Due to its main purpose for experimentation and in contrast to traditional grids, it has been designed to be easily re-configured and managed. All experiments are isolated from each other and from the Internet. Grid'5000 features a wide variety of high-speed networks and high performance computers. Other similar platforms like DAS-3 [?] from the Netherlands achieve similar goals.

DSL-Lab tries to retain all the features that have made Grid'5000 a success by adapting them to residential Internet nodes. We use the same OAR scheduler [?] but, as detailed previously, added power saving features and improved disconnected host management. Another example is the deployment software. For this, Grid'5000 includes Kadeploy [?] but this optimizes communications only for clusters and requires server management cards (IPMI). We had to rewrite a similar tool able to deal with Internet communications. The second difference is that our software is not able to take deployed OS control back due to lack of hardware support on our nodes. (PC's watchdog methods have been implemented but surprisingly not found on our nodes.)

Apart from the isolated grids world, there are platforms that include some real Internet links between their nodes. Experiments would be feasible with Volunteer Desktop Grids, but even if they offer a great number of nodes, the experimenters have very limited control of the volunteer's resources. In particular, there is no guarantee that volunteered desktop would provide full CPU power and availability 7 days a week. [?] So, in comparison, dedicated nodes provide a more flexible experimental setup.

PlanetLab [?] is a worldwide platform gathering about 900 nodes from 472 laboratories locations. It was originally a centralized network managed by the Princeton University, but now federates various administrative domains and projects. Nodes are shared among experimenters using virtualization (Linux VServer). In contrast, DSL-Lab allows, optionally, exclusive full access to nodes in order to prevent interactions with others experiments (required on low bandwidth links). The fact that there is no virtualization also allows DSL-Lab nodes to require less power and to consume less hardware resources than PlanetLab nodes. (For example, the absence of virtualization allows us to use up to 2 Gb of storage and 512 Mb of RAM, which is the full system capacity.) The lack of virtualization also suppresses the additional network latency during network time measurements, thus providing more precise measurements.

One of the commonly noticed limitations of the original PlanetLab project is that nodes offer only a small heterogeneity as they are only labs' nodes. We will now describe projects added to the PlanetLab's federation to deal with that issue.

The EU project OneLab, the European branch of PlanetLab, have made some studies that cover radio links such as 802.16, UTMS and WiFi. They plan

to works on the other types such as Bluetooth. They also provides network monitoring services.

SatelliteLab [?] is an extension of PlanetLab which includes nodes on the edges of Internet. The testbed consists of *planets* nodes formed by PlanetLab nodes and *satellites* nodes which are desktops, laptops and PDA connected to the Internet through DSL, Wifi, Bluetooth and cellular links. The requirements of PlanetLab hosts are very high. Server class machines running the PlanetLab OS must be configured with a static and public IP address, and this requirement excludes most of the nodes on the edges of Internet. With SatelliteLab, the application code of an experiment runs on the planets, whereas the satellites only forward network traffic. In contrast, the DSL-Lab software stack has been designed to run on a machine with very low requirements (less than 2 GB of storage) hosted on a domestic DSL, which allows experimenters to run experiments directly on each node of the platform.

Besides PlanetLab, Emulab, originating from Utah University, is a software and a platform of 450 nodes. Experiments are driven using a similar interface to the Network Simulator. It also offers access to some wireless links, sensor networks, access to PlanetLab nodes and can use emulation. Experiments, even if they get exclusive access to nodes, can be swapped out and may share local networks with other nodes. DSL-Lab has only one node per broadband line and thus does not suffer from this limitation. EmuLab is still being improved with projects such as FlexLab that try to restart experiments automatically as long as unrealistic conditions are detected by its network observations.

5 Conclusion

In this document, we have presented the experiments and results of the DSL-Lab project. To summarize our main achievements:

- We have built a new experimental platform, which allows scientists to experiment on the broadband DSL Internet. After a careful selection of the hardware, we have designed and developed a complete software stack which features OS deployment, resource management, security and connectivity. The platform is deployed over 40 nodes and is now opened to several collaborators from other universities.
- DSL-Lab has lead to several advances in evaluating Desktop Grids (for example, BitDew, PVC, BonjourGrid). To the best of our knowledge, none of the existing Desktop Grid middleware have presented experimental evaluation in the context of the DSL broadband Internet. We are now opening the platform to other areas of distributed computing research (self-stabilization, mobile computing etc...).
- We have successfully connected DSL-Lab with two other main experimental platforms, namely PlanetLab and Grid'5000. While these early results are promising, we plan to focus our next work on bridging Grid'5000 and DSL-Lab in order to make DSL-Lab available for the entire Grid'5000 community.

Contents



Centre de recherche INRIA Grenoble – Rhône-Alpes
655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex
Centre de recherche INRIA Sophia Antipolis – Méditerranée : 2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
<http://www.inria.fr>
ISSN 0249-6399