VirtLab: Virtual Laboratories in Federated Environments

José Quintino Rogado [1,2,3]

[1] CICANT Research Center, 376, Av. Campo Grande, 1749 - 024 Lisbon, Portugal.

[2] ECATI - Computer Engineering, Universidade Lusófona, 376, Av. Campo Grande, 1749 - 024 Lisbon, Portugal

Abstract: This document presents the ongoing VirtLab project which fosters an architectural model for modular and configurable virtual network laboratories, supporting selective access to different levels of resources, according to learning profiles, which can be securely utilized inside a campus or in the context of e-Leaning federations. This work references globally distributed testbeds such as PlanetLab or FEDERICA, and is aligned with projects developed in the European academia, which promote federated sharing of online educational content. The VirtLab Project is a relevant example of how state of the art information technology can be used to leverage successful teaching paradigms and make them accessible in the context of e-Learning initiatives.

Keywords: Virtual laboratories; federated eLearning; computer engineering.

1. Introduction

The level of success in learning results improves considerably when students have access to laboratorial environments from the very beginning of their curricula. This conclusion has been stated by different authors [1][2], and verified during several years as a lecturer of post and graduated courses in the area of computer engineering (operating systems, computer networks and distributed systems), for which a laboratorial environment was specifically created [3]. This laboratory, nicknamed Netlab (see **Fig. 1**), was designed and configured as a small replica of a real life network, with different network partitions and address ranges, created by switching devices and dual port hosts, which can be (re)configured during the classes.

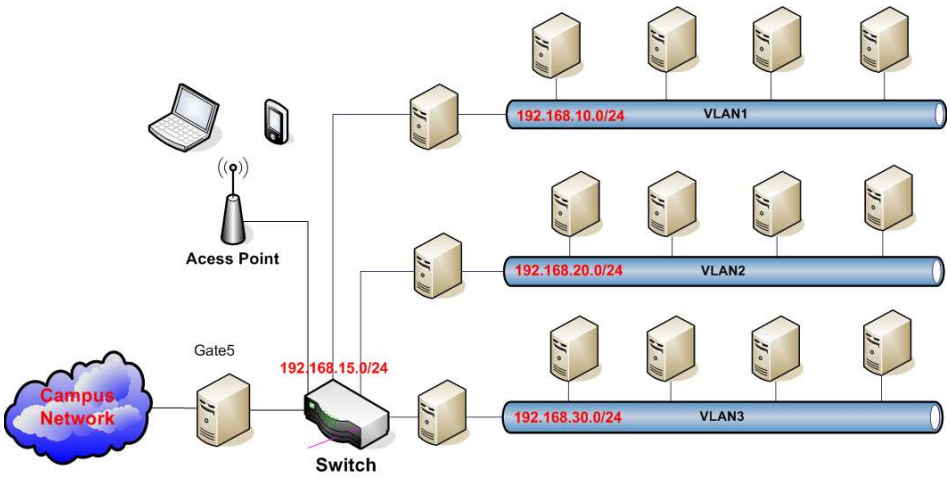

Fig. 1 *A simplified view of the Netlab Infrastructure.*

In post graduated courses, this infrastructure presents the advantage of being separated from the rest of the campus network, so that advanced distributed experimental environments, like VoIP networks or parallel computing clusters, can be deployed without disrupting the rest of the academic network.

Although quite successful from the educational point of view, the physical laboratory infrastructure does not easily support the current demanding learning requirements, which include e-Learning style remote access, a continuously increasing student population in computer engineering, and a wide range of concurrent configurations to support various teaching levels. Although an obvious solution for these problems would be to create physical replicas of the existing lab, space and cost considerations point to a different solution, which is to use virtualization technology instead.

[3] Email address: jose.rogado@ulusofona.pt

2. Moving towards Virtualization

In recent years, we have been observing more and more real life physical environments taking a leap towards virtualization, with medium to large enterprises and organizations progressively migrating their data centre resources to virtual infrastructures. This transition is being driven by economical and ecological reasons, and has been made possible by the enormous progress that virtualization techniques have undergone recently, both in its hardware [4] and software aspects [5][6].

Using these possibilities, virtual networks of machines and routing devices can be configured inside one or more virtualization hosts. Besides, different types of pre-configured virtual lab topologies can be activated on demand, based on the particular needs of a specific class or course, thus allowing simultaneous operation of different laboratory configurations.

Several initiatives exist that use virtualization techniques to create experimental network environments. Probably one of the first and most well known is the PlanetLab project [7][8], that is being developed in the University of Princeton (USA) since 2002. The goal of the PlanetLab concept is to create a large scale virtual network spanning the whole planet, which can be used for experimenting new Internet protocols and application architectures. Basically, PlanetLab offers researchers and developers the possibility to share a large pool of network resources that can be divided in multiple partitions, on which several concurrent experiments can be undertaken simultaneously. This approach provides the necessary independence and isolation that are necessary for achieving innovative developments without disrupting the normal network traffic as it would be the case if new network architectures and protocols were to be tested in the real Internet infrastructure. To provide such isolation and accrued manageability, and "overcome the Internet impasse" [9], PlanetLab takes the option of implementing its overlay network layer using Virtual Machine Monitors (VMM). PlanetLab is currently in a mature phase, spanning about 1000 nodes worldwide, and currently supporting many research projects in the area of advanced network technology (see for instance the Grid Appliance project [10]).

Another project that is relevant for the present work is the European Union funded project FEDERIKA[11] which is part of the GEANT2 [12] initiative and started in January 2008. As the name implies, the goals of FEDERIKA are similar to those of the PlanetLab project, i.e.: to provide a virtual network infrastructure for fostering innovative projects in the area of network architectures. The FEDERIKA approach, however, federates the existing National Research and Education Networks (NREN) optical fiber infrastructures to provide 'slices' of this high speed and low latency network to research experiments. As part of its goals, FEDERIKA is creating points-of presence (PoPs) in almost all European countries, which provide access to network resources and implement access control policies based on (con)federating technologies developed in other related GEANT2 projects like eduroam [13] at the network level, and eduGAIN [14] at the service level. Interoperating with similar projects worldwide, like PlanetLab, is also part of FEDERIKA goals.

Although these projects develop and deploy extremely sophisticated virtualization techniques to create experimental network environments, none of them is directly suitable for solving the problems that Netlab faces, namely to provide an e-Learning configurable network environment for graduate and postgraduate students. In the case of PlanetLab, the complexity of the network overlay is of no utility in the Netlab case, and the FEDERIKA PoPs are not yet available in all the European countries and are currently reserved to advanced research projects. Interoperating with advanced environments like PlanetLab and FEDERIKA is part of the VirtLab project in a later phase, but in the meantime there are urgent needs for the creation of advanced virtual environments more specific to academic e-Learning purposes.

3. Customizing the Virtual Environment

As stated in Section 1, since NetLab is used by students of different levels, it needs to support concurrent configurations, which do not have the same degree of complexity. For instance, in Operating Systems courses, students need only one virtual machine they can control and configure, while in the first levels of network courses they need two machines per group to develop and test client/server applications and analyse the resulting network traffic. In more advanced network courses, there is often the need of having 3 or 4 hosts or two sub-networks per group to develop distributed applications or to study routing algorithms. With virtualization, this can be achieved by allowing students to access only those lab instances that are specifically associated to their course profiles, as shown in **Fig. 2**: section a) for operating systems, section b) for basic network and sections c) and d) for advanced network classes.

To implement this functionality, at least three conditions have to be fulfilled: 1) students should access their lab environment in authenticated sessions; 2) the student profile needs to contain explicit information about their course enrollment; 3) and the virtual environment has to be closely integrated with the campus authentication and authorization infrastructure (AAI). Today, in most universities the first two conditions are often met, since an e-Learning platform (like Moodle or equivalent) is usually present, on which the concepts of student profile

and course enrollment are already implemented. The third condition is harder to achieve, and to understand its implications, a brief overview of the generic architecture of VM platforms is necessary.

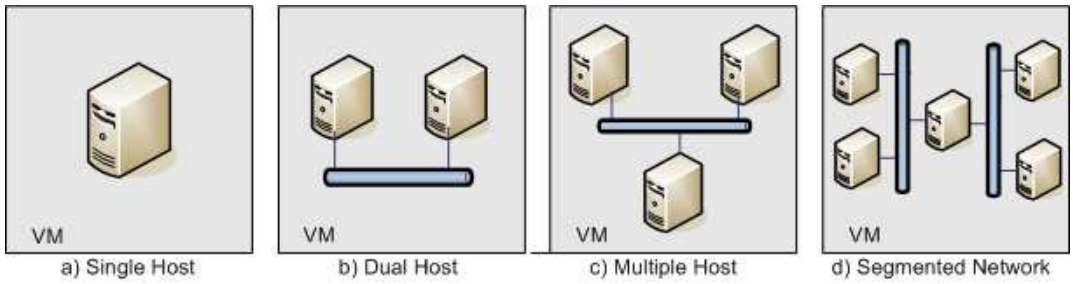

Fig. 2: *Different Virtual Laboratory configurations adapted to different course requirements.*

In very brief terms, data center virtual machine platforms are basically composed of 3 components [15]: a low level Hypervisor, that directly manages the hardware and creates the virtual processor abstraction for the guest system, a Virtual Machine Process on which the guest OS runs, and a Host Process that manages the interaction between the guest OS and the peripherals. The VM platform operation is controlled by means of another component, the Remote Management Console (RMC). Usually, the communication between the RMC and the VM platform is authenticated, and uses public APIs that, on most VM platforms, are exposed and based on Web protocols like HTTP or SOAP.

In order to implement the customization described above, the interaction between the RMC and the VM Platform must be conditioned by the profile attributes, i.e.: the communication has to be intercepted by a specific component that we will call the Secure Access Gateway (SAG), which controls all accesses to the virtual infrastructure and takes its decisions based on enrollment profiles. Implementing this component and integrating it with a VM Platform is one of the primary goals of the VirtLab project.

4. Integrating with the Campus AAI

Integrating VirtLab with the Authentication and Authorization Infrastructure of the university provides a seamless access to the virtualization infrastructure from all computers in the campus. A mind map description of the participating entities and the relationships between them is given in
Fig. 3.

Fig. 3: *VirtLab - AAI integration: Entities and Relationships*

The process for selecting and accesing a virtual laboratory instance works as follows: when students pretend to use VirtLab to follow a practical course or perform a work assignment, they first login to their campus portal, providing their credentials to the Authentication Module. This module contacts the Identity Repository, retrieves the corresponding stored credentials, and validates them. In case they match, the authentication succeeds and the user is able to see and to select the lab icon on the corresponding e-Learning tool. At this moment the Secure Access Gateway intercepts the request, queries the Authorization Module to validate that the current user is allowed to access the virtual laboratory. If so, the corresponding student enrolment profile is returned. This profile is used to address the Laboratory Configuration Database, from which is extracted the set of virtual environments to which the student's profile has access. These environments are presented to the user, who is now able to select their option.

Upon selection of the desired lab instance, the Secure Access Gateway recognizes the resource as valid for the profile and allows the access. Consequently, the browser spawns a Remote Management Console that connects to the VM pre-configured environment selected, on which the user is allowed to perform some of the

possible VM configurations (adding network interface or changing its settings, etc.) and operations (start, stop, suspend, resume, etc.).

As seen in **Fig. 3**, the Secure Access Gateway is a central component in the construction of the VirtLab concept. It can be implemented using existing Open Software Single Sign-On components, but the choice of the technology is crucial to the extensibility of the VirtLab functionalities, given that remote e-Learning style access and interoperation with other laboratory platforms are also important requirements.

In this line of thought, the choice taken was to use federative technology at the very core of the system, so that the different components of the environment are tied together as pieces of an inter-campus federation. In that field, SAML [16][17] is the globally accepted standard for federative authentication and secure attribute propagation, and Shibboleth [18] is probably one of the most disseminated academic implementations of federative SSO based on SAML. These technologies have therefore been adopted as the foundations of the Secure Access Gateway functionality.

5. Federating Virtual Laboratories

Having that in mind, the VirtLab architecture is being built using the Shibboleth platform which is bascally composed of two main services [19]: The *Identity Provider* (IdP), which is responsible for validating identities assertions and providing attributes, and the *Service Provider* (SP), which protects the resources and establishes a secure dialog with an IdP to obtain identity assertions and associated attributes. The SP grants or denies access to the requested resources based on the validity and nature of these assertions. The choreography of the interactions between Shibboleth services and the VirtLab components are depicted in **Fig. 4**.

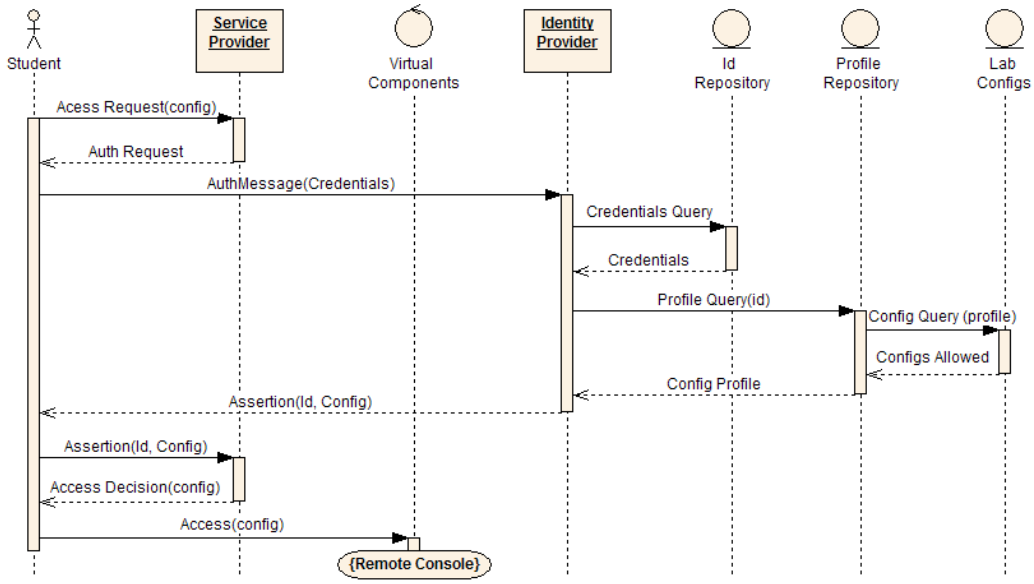

Fig. 4: *Interactions between Shibboleth modules and VirtLab components*

As we can see, the Access Control Gateway functionality is now split between the two Shibboleth IdP and SP modules. For intra-campus accesses, a local IdP is used, which acts as a front end for the local identity and profile repository. But any other Identity Provider can be used, provided that trust relationships are created, so that the local SP can securely accept assertions from it. This fact shows the advantage of using a federative platform to implement local authentication and authorization functionalities: by simply choosing other certified IdPs for providing identity and profile assertions to the VirtLab SP, students from other campuses or universities can be granted access to the virtual infrastructure. This approach can thus support seamless federated access from other academic institutions to the VirtLab components, provided that the necessary conventions, trust relationships, and profile attribute mapping are established between institutions.

In northern European countries there are many examples of well established academic federations, like in the United Kingdom [20] or in Nordic Countries [21]. In Portugal, this degree of collaboration is not yet achieved in the academia. To fill this gap, an ongoing initiative lead by FCCN (National Foundation for Scientific Computing - www.fccn.pt) called the AAI Pilot [22], aggregates some of the most relevant Portuguese universities, with a first short term goal of creating a federated access to Moodle e-Learning content. Once the pilot enters a stable state, other applications will be added to the federation, like access to the b-on Library [23]

and to academic VoIP network. The VirtLab platform, due to its federated infrastructure, constitutes another natural candidate for inclusion in this initiative.

6. Ongoing Work and Final Remarks

The VirtLab project is currently in an early phase where the architectural choices have been taken and preliminary proofs of concept are being made. One of these proofs consists in intercepting the message flow between a VM platform and its remote console management in order to introduce the authentication schema described in this paper. This task is made easier by the fact that these messages are in most cases sent though HTTP and/or SOAP, which are the protocols the Shibboleth services use. The result of this task will most probably also condition the choice of the virtualization infrastructure. The other important work now starting is to find a way for propagating the generic VM configurations available to different student profiles, possibly extending the Attribute Release Policy mechanism supported by Shibboleth. Both tasks are being achieved in the context of two Master's Thesis in Computer Engineering, which are in progress in our University. Some other less complex tasks, like the specification and configuration of the virtual network modules, are being achieved in the context of graduate studies final projects. The first working prototypes are expected next fall.

To conclude, and although there is a significant amount of work ahead, we strongly believe that this project will bring a significant contribution to the improvement of computer engineering student results by extending the range of intra-campus resources, thus augmenting their possibilities to use experimental environments. With respect to remote access, given the commoditization of broadband networks and the accrued possibilities for creating complex Web interfaces, e-Leaning platforms are enabling access to increasingly richer contents, to which the laboratorial environment described in this paper brings another degree of sophistication.

7. References

[1] Rossiter, J.A., "New Games for use in Lectures to improve Students learning", International Conference on Engineering and Education (ICEE): Coimbra, September 2007, icee2007.dei.uc.pt/proceedings/papers/56.pdf.

[2] Gardner, H., "Multiple Intelligences: The Theory in Practice", Ed. Basic Books, New York, April 1993.

[3] Rogado, J. Q., "A Prática Laboratorial no Ensino da Engenharia Informática", Caleidoscópio Review, Xth Edition on Computer Engineering, Ed. Lusófona, Lisboa, Portugal, December 2008.

[4] Leung, F. et al, "Intel® Virtualization Technology: Hardware Support for Efficient Processor Virtualization" Intel Technology Journal, August 2006, www.intel.com/technology/itj/2006/v10i3.

[5] Menon, A., et al, "Optimizing Network Virtualization in Xen", Pp. 15–28 of the proceedings of the Usenix Annual Technical Conference: www.usenix.org/events/usenix06/tech/menon.html, Boston USA, June 2006.

[6] Yang, S., "Extending KVM with new Intel® Virtualization technology", Kernel Virtual Machine Forum, Napa California, USA, June 2008.

[7] Peterson, L. et al, "A Blueprint for Introducing Disruptive Technology into the Internet", Proceedings of HotNets I, Princeton, New Jersey, USA, October 2002, www.acm.org/sigcomm/HotNets-I.

[8] Peterson, L. et al, "Experiences Building PlanetLab", Proceedings of the 7th USENIX Symposium on Operating System Design and Implementation *(OSDI '06)* , Seattle, WA, November 2006, nsg.cs.princeton.edu/publication/experiences_osdi_06.pdf.

[9] Anderson, T. et al, "Overcoming the Internet Impasse Through Virtualization", in IEEE Computer Review, Volume 38, Issue 4, pages 34-41, April 2005, www.arl.wustl.edu/~jst/pubs/hotnets04.pdf.

[10] Wolinsky, D., et al "Design of Virtual Machine Sandboxes for Distributed Computing in Wide Area Overlays of Virtual Workstations'. First Workshop on Virtualization Technologies in Distributed Computing, Tampa, Florida, November 2006, byron.acis.ufl.edu/papers/vtdc06.pdf.

[11] Maglaris, V., "Next Generation Networking in Europe: GÉANT3 and FEDERICA", e-IRG Workshop, Lisbon, October 2007, www.fp7-federica.eu/documents.php.

[12] GEANT2: The pan-European Research and Education Network, www.geant2.net.

[13] Eduroam Document "Eduroam Service Definition and Implementation Plan", Eduroam Deliverable DS5.1.1 January 2008, www.eduroam.org.

[14] López, D., "eduGAIN: Federation Interoperation by Design", TERENA Networking Conference 2006, Catania (Italy), May 2006, www.terena.org/events/tnc2006.

[15] Xen VM Document ,"Xen Architecture Overview", February 2008, wiki.xensource.com/xenwiki/XenArchitecture.

[16] Saldhana, A., "Oasis SAML and XACML", Exploring Identity Management Landscape Workshop, Interoperability Week at NIST, Gaithersburg, Maryland, April 2008.

[17] Lockhart, H. et al., "Defining and maintaining a standard, XML-based framework for creating and exchanging security information between online partners", OASIS Security Services Technical Committee.

[18] Morgan, R. L. et al., "Federated Security: The Shibboleth Approach", EDUCAUSE Quarterly, Volume 27, 2004, http://connect.educause.edu/Library/EDUCAUSE+Quarterly/FederatedSecurityTheShibb/39889.

[19] Carmody, S., "Introduction to Shibboleth and Phases for Deployment", EDUCAUSE Seminar on Flexible Web-Based Authentication and Authorization, June 2007 Portland, Oregon, http://net.educause.edu/Proceedings/12960.

[20] Tysom, M, The UK Federation, 2007, www.ukfederation.org.uk/content/Documents/FederationPresentations.

[21] Linden, M, ”The Nordic Middleware Identity Federation”, 23rd NORDUnet Conference, Gothenburg, Sweden, 2006, www.nordu.net/conference2006/presentations/We14.pdf.
[22] Piloto AAI - FCCN, http://www.fccn.pt/index.php?module=gps&proj_id=9&MMN_position=262:254.
[23] b-on: Portuguese Online Knowledge Library, www.b-on.pt.