

Operating The HPI Future SOC Lab - An Infrastructure Laboratory For Researchers - A Practitioners' Report

Rami Akkad, Henning Schmitz, André Deienno Pansani

Innovation Center
SAP AG
Konrad-Zuse-Ring 10
14469 Potsdam
rami.akkad@sap.com
henning.schmitz@sap.com
andre.pansani@sap.com

Abstract: The Hasso-Plattner Institute (HPI) Future Service Oriented Computing (SOC) Lab is a research institution of the HPI in Potsdam. It aims at enabling and fostering the exchange between industry and academia. Researchers are provided with free-of-charge access to state-of-the art software and hardware. The HPI Future SOC Lab hosts events for researchers and industry representatives to create a forum in which research ideas and results are discussed. The administration and operation of the lab requires processes for data life cycle management, privacy and security concepts, data backups and restores, provisioning techniques, and data center monitoring. In this report we will give insights into how the FSOC Lab is operated and what processes are in place.

1 Introduction

The Hasso-Plattner Institute (HPI) Future Service Oriented Computing (SOC) Lab is a research institution at the HPI in Potsdam¹. It was founded in 2010 as a collaboration of the HPI and four partners from industry; HP, EMC, Fujitsu, and SAP. The lab's overall goal is to enable and foster the exchange between industry and the research community. The HPI Future SOC Lab provides researchers with free-of-charge access to an infrastructure of state-of-the art software and hardware. This infrastructure is funded by the collaboration partners and operated at the HPI. Some highlights of the infrastructure are a 1000 core compute cluster, multi-core servers with up to 64 cores and up to 2TB RAM, GPU systems, a cloud computing infrastructure with 32 blades, and SAP software products. SAP sponsors three full-time employees and a student position at the lab, while the HPI funds one full-time employee, a student position, and the operational costs of the lab.

¹ HPI Future SOC Lab Homepage, <http://hpi.de/en/research/future-soc-lab.html> [last access 12.06.2014]

The Future SOC Lab also organizes and hosts events for researchers and representatives from industry in which research ideas and results can be discussed. These events include scientific workshops on questions regarding in-memory computing and cloud operations, regular HPI Future SOC Lab Days (see 2.2), and a workshop on Big Data Benchmarking². Main focus areas of the research activities at the HPI Future SOC Lab are defined by the partners of the lab and currently include in-memory computing, cloud computing, and GPU computing.

The activities at the HPI Future SOC Lab provide the partners from industry with insights into latest research trends, contacts to researchers, and a forum to position their own research questions in academia. Researchers get the infrastructure means to conduct their research while getting a chance to network with other researchers and potential partners from industry.

1.2 How The HPI Future SOC Lab Is Different To Other Infrastructure Labs

There a number of institutions worldwide that provide infrastructure services to researchers. Some of them are mentioned in the following with a focus on a highlighting one or more differences to the HPI Future SOC Lab.

The HPI Future SOC Lab's compute power can be considered small compared with what the Leibniz Supercomputing Centre offers with its SUPERMUC and other systems³. Hence it is better suited to support research in high performance computing than the HPI Future SOC Lab. The infrastructure of the HPI Future SOC Lab is located at the HPI in Potsdam and not distributed such as the infrastructure of the FutureGrid Project, which is a collaborative effort of many institutions⁴. The HPI Future SOC Lab is free-of-charge for researchers from any institution, which differentiates it from labs such as the GWDG in Göttingen, Germany⁵, or the Konrad-Zuse-Zentrum in Berlin, Germany⁶, which provide free services to certain institutions only. Providing datasets and workloads from industry is currently not in scope of the HPI Future SOC Lab either. By contrast, the Smart Data Innovation Lab with its focus on big data research intends to offer exactly such data sets⁷.

2 The HPI Future SOC Lab Periods

The work of the HPI Future SOC Lab is organized in periods of which two are conducted per year. The following chapter provides details about these periods, which comprise a proposal submission period, the HPI Future SOC Lab day as the main event of a Future SOC Lab period, and the project phase, in which projects are conducted.

² Future SOC Lab Event Overview, <http://hpi.de/en/research/future-soc-lab/events.html> [last access 12.06.2014]

³ Leibniz Supercomputing Centre Homepage, <http://www.lrz.de/services/compute> [last access 12.06.2014]

⁴ Futuregrid Homepage, <https://portal.futuregrid.org/about> [last access 12.06.2014]

⁵ GWDG Homepage, http://www.gwdg.de/index.php?id=wir_ueber_uns [last access 12.06.2014]

⁶ ZIB Homepage, <http://www.zib.de/de/institut/struktur-und-ziele-des-zib.html> [last access 12.06.2014]

⁷ Smart Data Innovation Lab Homepage, <http://www.sdil.de/en> [last access 12.06.2014]

2.1 Proposal Submission Period

Prior to the beginning of a period a call-for-proposals is published, which invites researchers to hand in project proposals to apply for infrastructure usage.

These proposals contain a description of the research purpose, some background about the researcher, the list of required infrastructure resources, the time for which access is needed, and whether this access is needed exclusively. The submission period lasts for roughly eight weeks. In exceptional cases project proposals can be submitted outside the official submission period. These are then handled in a separate process.

2.2 The HPI Future SOC Lab Days

At the end of the submission period a conference is hosted, the HPI Future SOC Lab day. This day marks the end of one and the beginning of a new Future SOC Lab period. Two Future SOC Lab Days are hosted annually, one in spring, usually at the beginning of April, and one in autumn, usually mid-October.

Representatives from the industrial partners, researchers, and guests meet to review and discuss the results of the projects on the day. Selected researchers present their results to the auditorium, while all researchers are asked to prepare posters and hand in reports summarizing their results. The reports are later published in HPI technical reports [Me10, Me11].

The HPI Future SOC Lab steering committee, which comprises representatives from the HPI and the partners from industry, approves or rejects the submitted project proposals during the Future SOC Lab Day as well.

2.3 The Future SOC Lab Period

The teams of accepted projects get access to the HPI Future SOC Lab for the following Future SOC Lab period, which is defined by the time between two subsequent Future SOC Lab Days and usually lasts six months. The principal investigator is required to accept the HPI terms of use, which includes the commitment to write a report and to create a poster by the end of the Future SOC Lab period. Afterwards the principal investigator provides the HPI Future SOC Lab team with the list of users.

For each project the required infrastructure is provisioned when needed by the project team. An average of 30 active projects and some 150 to 200 users access the infrastructure at any time. At the end of each Future SOC Lab period the users are locked. In some cases projects continue their work in subsequent Future SOC Lab periods. This requires the approval of a new project proposal but also operational consideration in terms of data and infrastructure handling, as a reuse shall be supported.

3 Technical Activities

The HPI Future SOC Lab comprises a highly heterogeneous infrastructure that is used for different purposes. To ensure a stable running environment and to simplify processes, several steps regarding resource sharing, security management, infrastructure provisioning, data- and user management, and monitoring have been implemented.

3.1 Resource Sharing

As projects require different access models to the resources the HPI Future SOC Lab offers, the lab implements two models:

1. Exclusive access to one or more requested servers / software services.

Due to high demand exclusive access to a server is usually granted for a limited amount of time, e.g. one week or one day, on a first come first served basis. Projects need to inform the lab two weeks in advance if they require exclusive access. Access to a software service, e.g. SAP HANA, is always exclusive and granted throughout one Future SOC Lab period.

2. Shared access to one or more requested servers.

Depending on the resource, shared access is realized through granting multiple projects access to one system without limitations (e.g. used in our GPU systems) or through limiting the host system's resources for a project with a control mechanism like Cgroups⁸ (e.g. used in our 1000 core compute cluster).

The lab uses active directory groups to automate the assignment of project teams to a specific server. Each server has one active directory group and all the server's group members are allowed to access the server.

3.2 Security And Data Management

After the Steering Committee accepted a project, all project related user accounts are created with an expiration date that matches the end of the current Future SOC Lab period. To access the lab's infrastructure a user needs to connect to a virtual private network (VPN). The lab avoids public IP addresses to avert public disclosure of our servers. To minimize the possibility of a brute-force attack, a user account is automatically disabled after five failed login attempts and automatically re-enabled after one hour.

As projects may deal with sensitive data the HPI Future SOC Lab has security guidelines to prohibit unauthorized access to files. However, the lab does neither follow ISO

⁸ kernel.org Cgroups specification, <https://www.kernel.org/doc/Documentation/cgroups/memory.txt> [last access 12.06.2014]

security policies nor does it guarantee data safety. It does not take responsibility in any case of data loss to third parties.

As described in 3.1, a project will get either exclusive or shared access to a resource. Depending on the operating system or offered service, the project is able to access the resource after successfully connecting to the lab's VPN through a protocol (e.g. SSH / RDP). If a project needs to execute root commands on a shared resource with no root access it needs to inform a lab admin who will carefully audition whether an execution of the command has any impacts on the security of the lab and who will perform the necessary operation.

To transfer huge amount of data the HPI Future SOC Lab runs an SFTP server that has a public IP address. Access to this server can be requested by projects at any time. All project related data will be deleted after the end of the Future SOC Lab period.

3.3 Infrastructure Provisioning

The lab uses different techniques to provision services, servers and virtual machines to meet the diverse requirements of projects running at the HPI Future SOC Lab.

Physical servers and virtual machines are deployed using Hewlett-Packard's private cloud solution Converged Cloud, in which pre-defined templates can be used to rapidly deploy operating systems and networks.

Most of the offered services comprise newest SAP software, such as SAP HANA. SAP HANA is currently installed on 1 TB RAM servers. Projects receive full control of their SAP HANA system. They have sufficient authorizations to perform customization and to develop, run, and debug their programs.

3.4 Backups

The HPI Future SOC Lab conducts backups on a voluntary basis and thereby does neither implement a standardized backup strategy nor does it guarantee the availability of fallback systems in case of data loss. All projects are responsible to create backups on their own. The lab does provide support in transferring data for backup reasons (e.g. when approaching the end of a Future SOC Lab period).

3.5 Monitoring

The HPI Future SOC Lab uses Icinga⁹ to monitor its infrastructure to ensure that servers and services run in a stable mode and to react on hardware failures.

The lab uses built-in Nagios plugins such as `check_ssh` and `check_root` to perform monitoring operations. Additionally the lab adapted Icinga to monitor SAP HANA by

⁹ Icinga Homepage, <https://www.icinga.org> [last access 12.06.2014]

writing a custom plugin that periodically checks that vital processes are running and that sufficient disk space is available on the relevant disks.

Icinga is also used as an internal reporting tool to generate statistics regarding load of servers and uptime of services. This helps the lab to understand problems and optimize daily operations.

4 Outlook And Conclusion

The HPI Future SOC Lab is a research institution at the HPI in Potsdam with the overall goal of enabling and fostering the exchange between industry and the research community. The HPI Future SOC Lab successfully provided over 200 [Bo11, LT11, SS14] projects from more than 40 institutes with free-of-charge access to an infrastructure of state-of-the-art software and hardware since it was founded in 2010. It hosted 12 events such as the HPI Future SOC Lab to create a hub in which industry and academia can connect.

The HPI Future SOC Lab will continue its work throughout the next years. To keep its infrastructure up to date and to attract and support cutting-edge research, the lab has to continuously update and improve its offerings and processes with the help of its partners.

This refers to both infrastructure and internal processes, particularly in regard to data and user management as well as automatic physical and virtual machine provisioning to enhance productivity and reliability.

References

- [Me10] C. Meinel, A. Polze, A. Zeier, G. Oswald, D. Herzog, V. Smid, D. D’Errico, and Z. Hussain: Proceedings of the Fall 2010 Future SOC Lab Day, Technische Berichte Nr. 42 des Hasso-Plattner-Instituts für Softwaresystemtechnik an der Universität Potsdam, Potsdam, 2011.
- [Me11] C. Meinel, A. Polze, G. Oswald, R. Strotman, U. Seibold, and D. D’Errico: HPI Future SOC Lab – Proceedings 2011, Technische Berichte Nr. 70 des Hasso-Plattner-Instituts für Softwaresystemtechnik an der Universität Potsdam, Potsdam, 2011.
- [Bo11] M. Boehm, P. Lehmann, P. Volk, and W. Lehner: Query Processing on Prefix Trees, Dresden, 2011
- [LT11] M. von Löwis, and P. Tröger: Safe and Scalable Build Automation through Virtualization, Tagungsband der INFORMATIK 2011, Berlin, 2011
- [SS14] J. Spillner, and A. Schill: Towards Dispersed Cloud Computing, 2nd IEEE International Black Sea Conference on Communications and Networking, Chişinău, 2014