

Towards Sustainable Software Infrastructures for Data-Intensive Systems

Guillaume Fieni
University of Lille / INRIA
Lille, France
guillaume.fieni@univ-lille.fr

Romain Rouvoy
University of Lille / INRIA / IUF
Lille, France
romain.rouvoy@univ-lille.fr

Lionel Seinturier
University of Lille / INRIA
Lille, France
lionel.seinturier@univ-lille.fr

ABSTRACT

Distributed environments are heavily used by nowadays data-intensive software systems to process continuous flows of data produced by a large number of devices (smartphones, IoT, cameras, etc.). While growing complexity, these distributed systems are required to mitigate two key concerns: security and energy. On the one hand, when offloading data processing to the cloud, the application may require security guarantees from the cloud providers in order to avoid privacy and security leaks. On the other hand, the cloud provider aims at continuously improving the energy efficiency of the cloud infrastructure by optimizing the resource usage. This PhD thesis therefore explores the security/energy trade-offs and optimizations that can be considered when deploying data-intensive software systems in distributed environments. In particular, this PhD thesis aims at exploiting all the possible levers that maximize the security of a data-intensive system while minimizing its power consumption.

KEYWORDS

Energy efficiency, Security & Privacy, Distributed systems, Cloud computing.

1 INTRODUCTION

Modern software systems are composed of complex components, which are deployed across a wide diversity of environments, from embedded sensors, to smartphones, to edge and cloud infrastructures. Part of these distributed systems are expected to process large volumes of raw data with *Quality-of-Service* (QoS) guarantees including not only timeliness or accuracy, but also privacy, security and energy efficiency. Yet, ensuring secure data processing while minimizing the energy footprint of the hosting infrastructure remains an open challenge as it requires to conciliate the potentially diverging objectives of applications and infrastructures.

This PhD thesis therefore intends to explore the trade-offs and optimizations that can be applied in the context of the development of data-intensive software systems, which are deployed in distributed environments.

2 CONTEXT AND CHALLENGES

In the context of this PhD thesis, we consider the case of data workflows that consume data produced by a multitude of devices and requires the execution of secured tasks (e.g., to anonymize some data attributes). Such application workflows are hosted in a cloud environment, which shares the resources across several applications and even organizations.

In this context, our research consists in identifying appropriate system-level mechanisms and heuristics to schedule the execution of secured or non-secured tasks in a multi-tenant environment, while minimizing the energy footprint, which requires to account the power consumption beyond the scale of nodes, and maximizing the availability of the hosting infrastructure. In particular, our security model builds on *Trusted Execution Environments* (TEE), like the *Intel Software Guard Extensions* (SGX), where the isolation mechanisms impact the performance thus increasing the execution time of the application, leading to an higher power consumption.

3 RELATED WORK

Several works in the literature propose different approaches to estimate and optimize the power consumption, especially for cloud infrastructures, where an intensive use of virtualization to maximize the use of resources can be observed. We can also observe the possibility to distrust the cloud provider while relying instead on hardware protection mechanisms guaranteeing the safe sharing of resource among different actors.

Power estimations. Several methods based on power models provides fine-grained power consumption estimations. For example, BITWATTS [3] delivers power estimations of virtual machines using *Hardware Performance Counters* (HPC) for the resource usage and an external power meter as machine-wide power consumption reference. However, such method requires a long calibration phase using synthetic workloads and external power meters, as power usage reference, which can be costly to deploy across a production infrastructure or not applicable in a heterogeneous environment. Moreover, external power meters monitor the power consumption at a coarse granularity, which is subject to noise from other components, such as disks and network interfaces, for example. Modern CPU support the RAPL feature providing power capping ability and power estimations of several components of the CPU itself including the RAM. The advantage of this method is its availability on all CPU since the "Sandy bridge" micro-architecture for Intel and the "Zen" micro-architecture for AMD. Hackenberg *et al.* [5] demonstrate the accuracy of RAPL power estimations for Intel's Haswell micro-architecture. However, this approach only provides socket-scale and per-domains power estimations and a power model needs to be created to reach finer granularity, like software containers.

Power usage optimization. While the above methods aim to profile the power consumption of systems, complementary work about the power optimization have been done by the community. We can cite POWERNAP [7] where the authors eliminate the IDLE power consumption by putting unused machines into a deep-sleep state, but such method requires specific-hardware and is non-trivial to deploy. vCAP [6] a power-capping technique where the resource

allocation decisions are optimized in order to maximize the *Quality-of-Service* (QoS) while meeting the power constraints of virtualized servers. DOCKERCAP [2] is a software-level power capping orchestrator based on RAPL power measurements, tuning the performance of Docker containers in order to respect a power consumption limit through resource management.

Trusted computing. Cloud infrastructures supporting secure enclaves are emerging, like Microsoft Azure confidential computing, allowing safe resources to be shared across multiple stakeholders. In this category, Intel SGX is a recent approach to safely isolate software processes sharing the same resources (CPU, RAM) of a given host. However, such security guarantees comes at a price and Gjerdrum *et al.* [4] evaluates the performance impact imposed by Intel SGX enclaves. Arnautov *et al.* [1] provides an approach to protect Docker containers from external tampering using Intel's SGX enclaves.

4 RESEARCH AGENDA

In this PhD thesis, we will assess the optimization problem of the resource usage of data processing on heterogeneous infrastructures under power usage and/or security constraint(s). For that, we choose to focus on suitable software isolation mechanism for this kind of requirements, specifically containers, virtual machines and secure enclaves. First, we will assess the problem of providing accurate and fine-grained power estimations of heterogeneous infrastructures when software isolation will be used. This will provide the baseline for evaluating the energy impact of secure isolation mechanisms. Finally, we will focus on the optimization of the resources usage (including power consumption) in order to produce flexible scheduling decisions according to the security and performance requirements.

5 PRELIMINARY RESULTS

Currently, we are working on a software-defined power meter based on the BITWATTS [3] middleware to estimate the per-container CPU and RAM power consumption. Based on RAPL power estimations and *Hardware Performance Counters* (HPC), this new software-defined power meter do not require to deploy hardware power-meters across the entire infrastructure while remaining precise and lightweight for the monitored hosts. We overcome RAPL's inability to produce fine-grained power estimations by building a power model at run-time using active-learning methods to detect and train a power model whenever an error threshold is exceeded. With this first contribution, we aim to provide a plug-and-play and self-adaptive power-meter for containerized environments. We plan to extend this contribution to estimate the power consumption of Intel SGX secure enclaves.

6 CONCLUSION

In this paper, we highlighted the importance of the efficiency of large scale data processing in distributed infrastructures under security and power consumption constraints. We introduced related works towards this goal along with their limitations, thus highlighting the challenge to take into account for such conflicting constraints.

We presented our research agenda covering this challenge onward our preliminary results on the power estimations of the CPU and RAM for software containers, by adopting an active-learning approach to provides self-adaptive power models using RAPL power measurements. In the future, we will work on delivering accurate power estimations of applications using secure enclaves and we will address energy optimizations for large-scale data processing systems deployed in distributed environments.

REFERENCES

- [1] Sergei Arnautov, Bohdan Trach, Franz Gregor, Thomas Knauth, André Martin, Christian Priebe, Joshua Lind, Divya Muthukumar, Dan O'Keeffe, Mark Stillwell, David Goltzsche, David M. Evers, Rüdiger Kapitza, Peter R. Pietzuch, and Christof Fetzer. 2016. SCONE: Secure Linux Containers with Intel SGX. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*. 689–703.
- [2] A. Asnaghi, Matteo Ferroni, and Marco D. Santambrogio. 2016. DockerCap: A Software-Level Power Capping Orchestrator for Docker Containers. In *IEEE Intl Conference on Computational Science and Engineering (CES) and IEEE Intl Conference on Embedded and Ubiquitous Computing (EUC) and 15th Intl Symposium on Distributed Computing and Applications for Business Engineering (DCABES)*.
- [3] Maxime Colmant, Mascha Kurpicz, Pascal Felber, Loïc Huertas, Romain Rouvoy, and Anita Sobe. 2015. Process-level Power Estimation in VM-based Systems. In *Proceedings of the 10th European Conference on Computer Systems (Eurosys)*.
- [4] Anders T. Gjerdrum, Robert Pettersen, Håvard D. Johansen, and Dag Johansen. 2017. Performance of Trusted Computing in Cloud Infrastructures with Intel SGX. In *7th International Conference on Cloud Computing and Services Science (CLOSER)*.
- [5] Daniel Hackenberg, Robert Schöne, Thomas Ilsche, Daniel Molka, Joseph Schuchart, and Robin Geyer. 2015. An Energy Efficiency Feature Survey of the Intel Haswell Processor. In *IEEE International Parallel and Distributed Processing Symposium Workshop (IPDPS)*. 896–904.
- [6] Can Hankendi, Sherief Reda, and Ayse K. Coskun. 2013. vCap: Adaptive Power Capping for Virtualized Servers. In *Proceedings of the International Symposium on Low Power Electronics and Design (ISLPED)*.
- [7] David Meisner, Brian T. Gold, and Thomas F. Wenisch. 2009. PowerNap: eliminating server idle power. In *14th International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS*.