"Private-by-design operating system for smart homes"

Ph.D. thesis proposal in the Cloud and Large-Scale Systems group (CLSC) @ UCLouvain.

Advisor: **Pr. Etienne Rivière**, UCLouvain

This Ph.D. thesis proposal is offered in the context of the CyberExcellence project.

The CyberExcellence project started in January 2022. Its goal is to position Wallonia as a major actor in cybersecurity, covering many aspects from research to applications, and involving all stakeholders (industry, academia, institutions, etc.). The project aims at producing excellent science and train top-level experts in the topics of cybersecurity broadly construed. Among the different research objectives of the project is the security of novel and complex cloud infrastructure such as multi-vendor clouds, distributed clouds, and hierarchical clouds. This Ph.D. proposal will participate towards this goal by improving the security and ease of use and operation of complex, emerging cloud infrastructures, and will contribute to prototypes and proof-of-concept implementations in modern distributed cloud software.

Context

Smart environments are physical spaces equipped with a variety of sensors and actuators (IoT devices, but also larger smart equipment and appliances). Smart homes are gaining in interest with the advent of commercial offerings allowing one to automate and control their homes using, e.g., voice-controlled systems or pre-defined automation rules. An example of use of smart homes could be for elderly persons wishing to stay in the comfort of their homes rather than going to a care home, yet benefit from comfort functionalities and safety measures (e.g., communication with relatives and monitoring of the living environment on the one hand, and emergency services following the detection of a fall or inactivity on the other).

Problem

Smart home software is essentially split in two categories today[1].

On the one hand, commercial systems (Amazon Alexa, Samsung SmartThings, etc.) are *exclusively* cloud-based. This means that all sensor and control data from the smart home is sent to the cloud where it is stored and processed. These systems offer advanced features, are easy to use for non-experts, and can integrate easily with available online services (e.g., streaming services, delivery, emergency, etc.). The fact that these systems are proprietary and cloud based is, however, a major threat to users' privacy, as they require to blindly trust a commercial operator, often not based in the EU, with potentially highly sensitive data.

On the other hand, open-source, community-driven systems exist that do not rely at all on cloud hosting or online services. Examples include OpenHAB or Home Assistant. These systems are designed to run on a small, dedicated PC running in the smart home (**smart hub**) and connecting to a range of compatible devices through open communication standards. They allow users to define a wide range of automation using IFTT (If-This-Then-That) rules or even their own Python scripts, and do not pose issues with data privacy as everything remains local to the smart home. They are, however, complex to use and do not easily allow advanced features such as voice control or the interaction with online services. They are also ill-suited

[1] For more background information, check our recent systematization of knowledge paper in PETS2022: https://petsymposium.org/2022/files/papers/issue4/popets-2022-0097.pdf

to support applications provided by third-parties, as all code running on these platforms is considered trusted by default.

The objective of this thesis is to design and implement an operating system for smart homes that allows (1) privacy-by-design by running computation and storage primarily on a smart hub, yet allowing the use of external services and cloud-based backends when necessary and agreed by the user; (2) allows the installation of third-party applications on the smart hub and their execution in a safe environment, with strict control of the access and use of sensitive smart home and user data; and (3) provides support for security-enhancing features, raising awareness of the user about the use and spread of their data and offering tools to control this use and spread.

Approach

The PhD work will build over recent research effort performed by the Cloud and Large-Scale Computing and Software Security groups @ UCLouvain, in cooperation with a research center in Brussels, towards privacy-preserving smart homes for smart care.

The objective of this PhD thesis will be to design, implement, and evaluate operating system solutions for smart hubs that follow a private-by-design principle. The vision is that application developers, at ease with common development toolsets and languages for the cloud and for web application, should be able to write applications for this operating system with minimal adaptation effort. For this purpose, we propose to re-use recent technological advances such as WebAssembly (WASM) to offer near-native performance for demanding applications running on a smart hub, similarly to its use in web browsers.

System security will be the most important goal of the study. In contrast with existing, open-source solutions for smart hubs, the designed operating system services should:

- Provide support for strict control of the use and spread of data from smart home appliances and sensors to applications and to cloud-based services. This requires designing appropriate access control models, and the runtime services enforcing them.
- Ensure that malicious code provided by an untrustworthy application provider is unable to harm privacy and the well-functioning of other, legitimate applications. Techniques such as taint tracking and dynamic code analysis are interesting directions for this purpose.
- Allow application developers to provide user-understandable guarantees about the use and spread of smart home data, and verifying at runtime that these guarantees are enforced (e.g., by offering controlled communication/interaction services to the applications).
- Facilitate the deployment and installation of applications from multiple providers, possibly with help of decentralized trust-enabling technologies (blockchains).

Profile and skills

This thesis is particularly suited for a candidate with interests in Cyber-Physical systems / IoT, Cloud Computing, and Operating Systems Security. The research will focus on a *systems* approach, meaning that the conceptual contributions must be validated by the development and evaluation of proof-of-concept prototypes.

Proficiency with a low-level systems language such as C or Rust is required, as well as good knowledge of typical cloud backend languages such as node.js. The researcher should have a good command of GNU/Linux systems. The candidate must have followed graduate-level security courses and be familiar with cryptography and systems security concepts.

In addition to the technical skills, the candidate is expected to be able to work in collaboration with other junior and senior researchers. A team player attitude is necessary. Doing a Ph.D. in computer science and particularly in computer systems is a very rewarding but also a very challenging experience due to the need to build prototypes and to run evaluation campaigns on real hardware. The candidate should, therefore, be ready to invest significant conceptual and technical effort towards the realization of the Ph.D.

English will be the primary working language. Knowledge of French is a plus but not a requirement. A candidate willing to learn or improve her/his French will have access to training by the university.

Environment

The Ph.D. position is fully funded for a total of 4 years (salary and research expenses). The Ph.D. candidate will be hosted in UCLouvain's ICTEAM (Institute for Information and Communication Technologies, Electronics and Applied Mathematics). She or he will interact with members of other research groups in ICTEAM (Software Security and Performance of Networked Systems), as well as with researchers from research centers in Brussels and Wallonia. She or he will also participate to the training and dissemination activities of the CyberExcellence project.

ICTEAM is a top-level research institute with excellent facilities. The candidate will have access to several distributed clusters with high-end, recent servers and edge resources to test the contributions. UCLouvain is Belgium's largest French-speaking university and is ranked in the first 100 institutions for computer science and information systems by the QS World University Ranking. The position is in Louvain-la-Neuve, a modern city 30 kms south of Brussels with excellent quality of life and connectivity to Belgium, Europe, and the world.

Application

Interested students should send an application file (PDF only, in English) with the following information to the promoter:

- An up-to-date CV;
- Link to a copy of the publication(s), if any, and to Masters' thesis;
- If applicable, a link to a GitHub (or similar) profile with examples of personal software contributions;
- Masters' degree transcripts;
- The name and emails of up to three reference persons able to provide an assessment of the application. Reference letters sent by the applicant her/himself will be ignored;
- A short cover email detailing your motivation for pursuing a Ph.D. and your interest in this specific topic.

Application deadline: ASAP -- applications will be screened as they arrive. Feel free to send a declaration of intent or request for clarification before sending a formal application.

Suggested starting date: January 1st, 2023 (flexible)

Contact: etienne.riviere@uclouvain.be