TrustFS: An SGX-enabled Stackable File System Framework

Tânia Esteves[1], Ricardo Macedo[1], Alberto Faria[1], Bernardo Portela[2], João Paulo[1], José Pereira[1] and Danny Harnik[3]
[1]INESC TEC & U.Minho, [2]INESC TEC & U.Porto, [3]IBM Research – Haifa

Motivation

- Users demand **data confidentiality**.
- Service providers need to apply **content-aware functionalities,** for space reduction and query optimizations.
- How can we ensure data confidentiality and privacy while allowing content-aware computations?
- How can this be done without requiring a deep reimplementation of existing storage solutions?

Architecture & Prototype

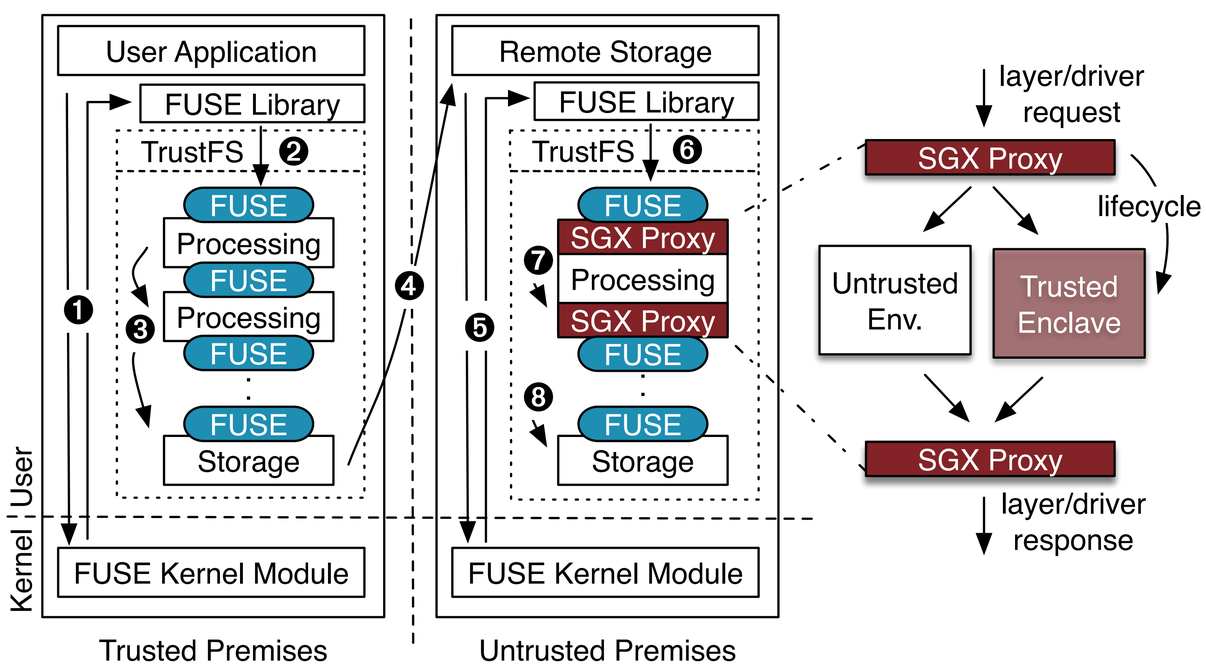

Fig. 1 - TrustFS Architecture

- **TrustFS**: a programmable and modular **stackable file system** framework for implementing secure content-aware storage functionalities resorting to **Intel SGX** enclaves.
- **SGX proxy** — a middleware component that can be used to transparently run layer and driver code in secure SGX enclaves.
- Prototype that enables **secure compression** over encrypted data.

Preliminary Evaluation & Conclusions

- The **integration** of FuseCompress as a TrustFS layer has a **small impact** in the performance of the different workloads, and requires modifying **less than 4% LoC** (230 of 5276).
- It is possible to provide secure compression while keeping the performance **overhead** between **6.5% to 31.3%**, with **less than 200 LoC**.
- FuseCompress (Vanilla setup) has a noticeable impact, in most workloads, when compared to the Native Setup.
- As data redundancy increases, write workloads performance can benefit from space reduction techniques.

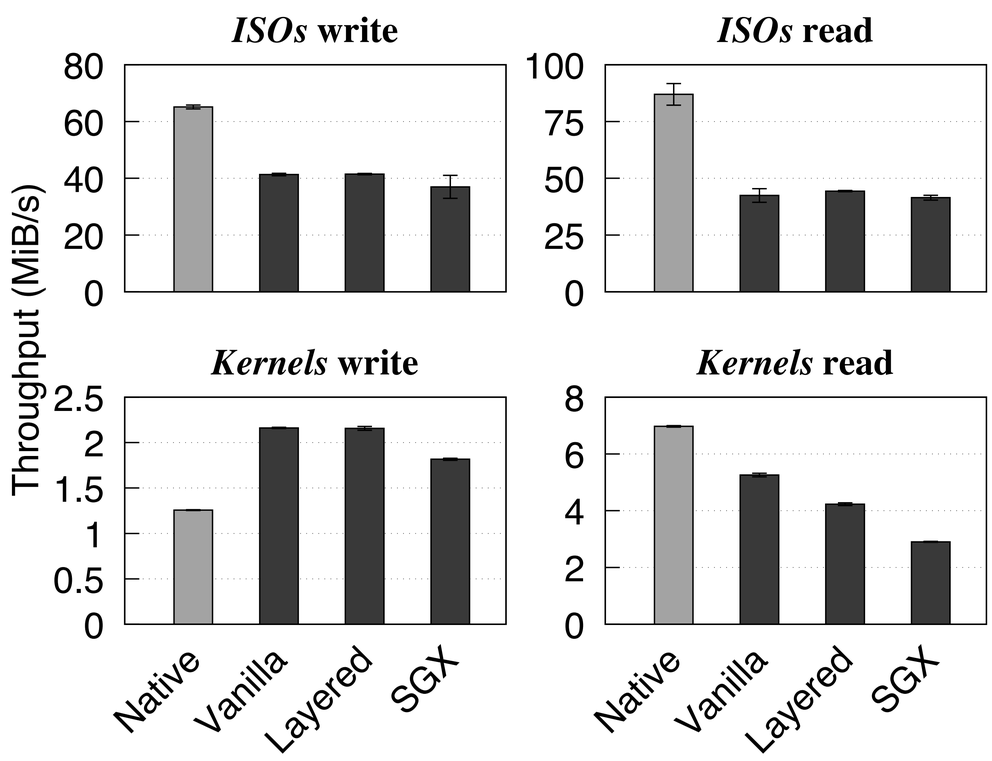

Fig. 2 - Performance evaluation of the compression prototype

Open Challenges & Future Work

- **Changing the storage layout** across layers can lead to significant performance penalty.
- TrustFS layers should implement proper mechanisms to detect **chunks splitting**.
- **Existing storage solutions** must be validated and evaluated before integration.
- A production-ready version should implement **cryptographic key exchange and management services**.
- The **number of LoC required** for implementing secure content-aware functionalities can be further reduced.

Acknowledgements

This work was supported by National Funds through the Portuguese funding agency, FCT – Fundação para a Ciência e a Tecnologia within the project UID/EEA/50014/2019, and by FCT/MCTES through project HADES (PTDC/CCI- INF/31698/2017).