

Building a Reusable Privacy Substrate for Data Analytics Using Smart Storage Nodes

Overview of Work in Progress


Zsolt István (ITU Copenhagen)

and collaborators from UT Austin and UTCN, Cluj-Napoca



Smart Storage is Becoming Mainstream

- Compute pushdown available as cloud service, commodity, appliances...




Introducing
SmartSSD® CSD



Insights

Analytics,
ML Models


Distributed
Data Store



Priorities in Data Intensive Systems

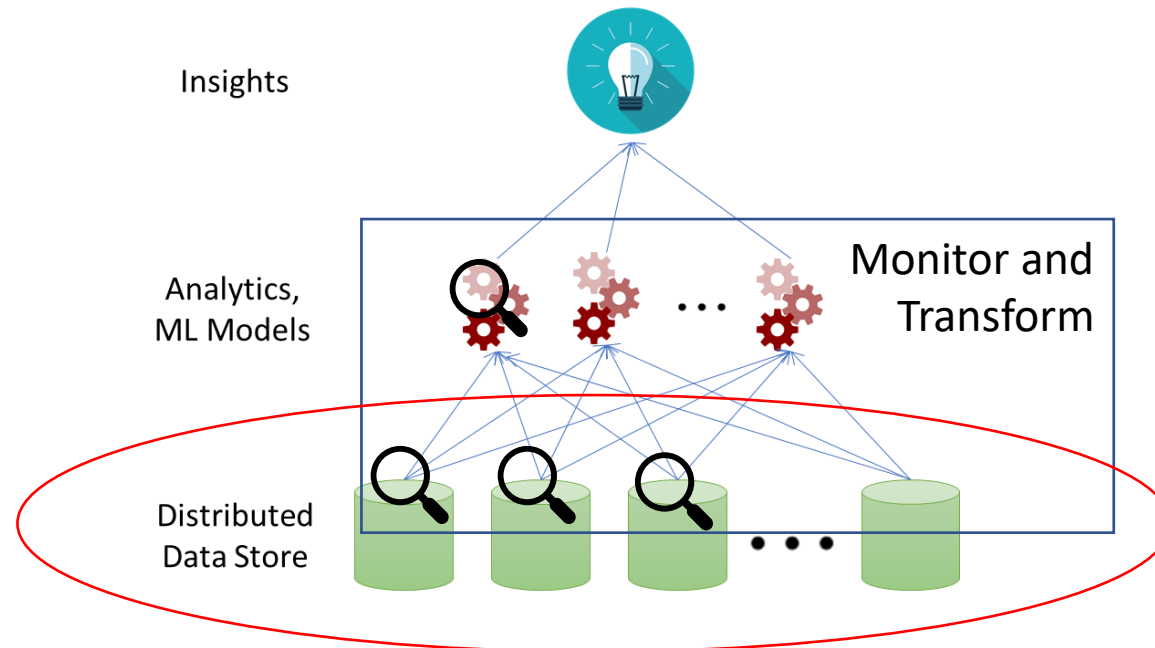
- Years of focus on performance
 - Speed up complex data analytics
 - Reduce data movement bottlenecks
 - Reduce coordination overhead
- Increasing privacy challenges
 - Risk of leaks, GDPR, etc.
 - Data flows between different teams

 Systems support for enforcing privacy rules



WiP: Systems Support for Enforcing Privacy

- Monitor and Transform data
- Avoid slowdown → Reimagine architectures
 - Specialized hardware in storage layer: low latency and predictability

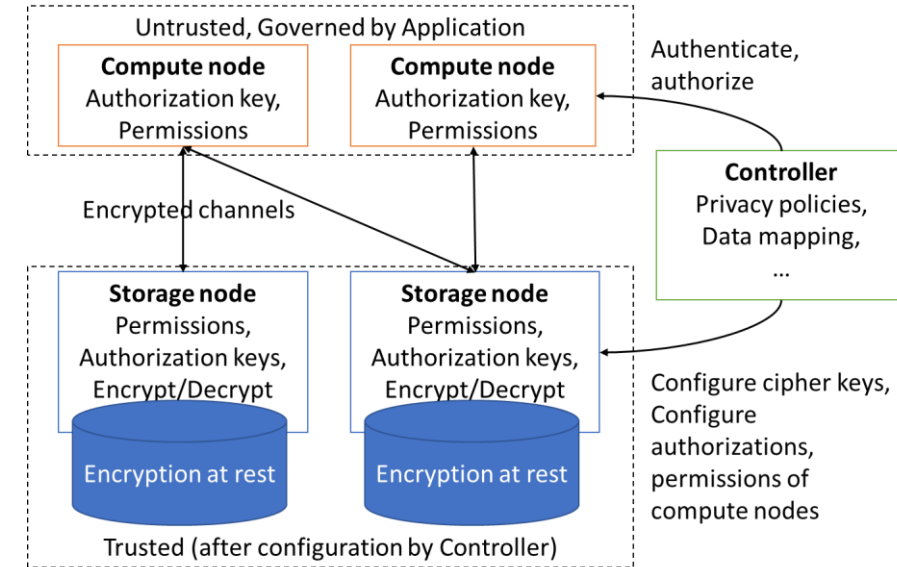


Software-defined Data Protection



Soujanya Ponnappalli, Vijay Chidambaram

- Policy compliance can be costly!
 - GDPR: >30% of data protection articles affect storage
 - **Software-Defined Data Protection**
 - SDS but Access and Protection instead of Performance
 - Decoupling enforcement from decisions increases performance
- Simplified logic in nodes → (Almost) achievable today!




Challenges: custom TEEs and policy interpretation

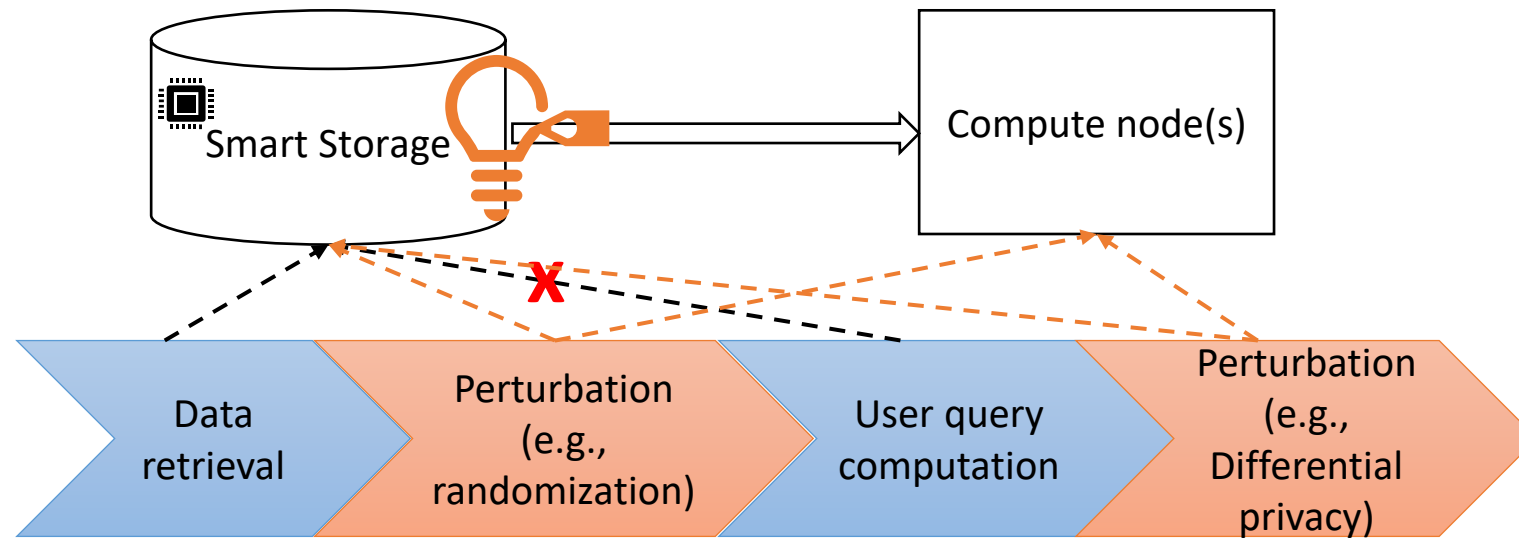
[2008.04936] Towards Software-Defined Data Protection: GDPR Compliance at the Storage Layer is Within Reach (arxiv.org)



Smart Storage with Privacy Operations



Smart Storage with Privacy Operations





Smart Storage with Privacy-preserving Perturbations



A. Hangan, G. Sebestyen, C. Mihali, A. Tosa

- Maintain some utility of data while protecting personal information
- 1) Perturb data before processing: Training classifiers on randomized data
 - Guarantee line-rate



- 2) Perturb data after processing: Differentially private histograms and group-by inside the storage
 - Integrate with control frameworks



We are hiring at ITU in these topics (PhD and Postdoc)!

Reach out for more information: zsis@itu.dk

WiP: Systems Support for Enforcing Privacy

- Monitor and Transform data
- Avoid slowdown → Reimagine architectures
 - Specialized hardware in storage layer: low latency and predictability

