Unlocking Cross-Organizational Insights

Practical MPC for Cloud-Based Data Analytics

PEPR' 25 – June 9th 2025

resolve

Daniele Romanini

Agenda

01 Introduction: Cross-Organizational Analytics

02 A system for democratizing private analytics with SMPC

03 Carbyne Stack & MP-SPDZ

04 System enhancements

05 Conclusion

01 Introduction: cross-organizational analytics

resolve

resolve.tech

The privacy challenge in cross-organizational analytics

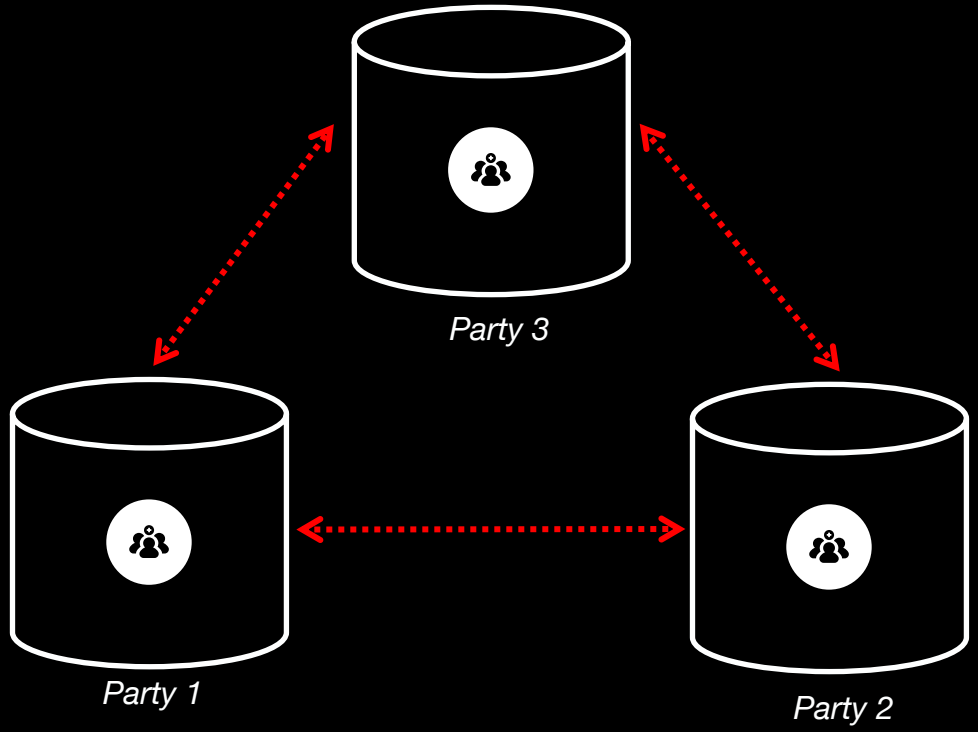

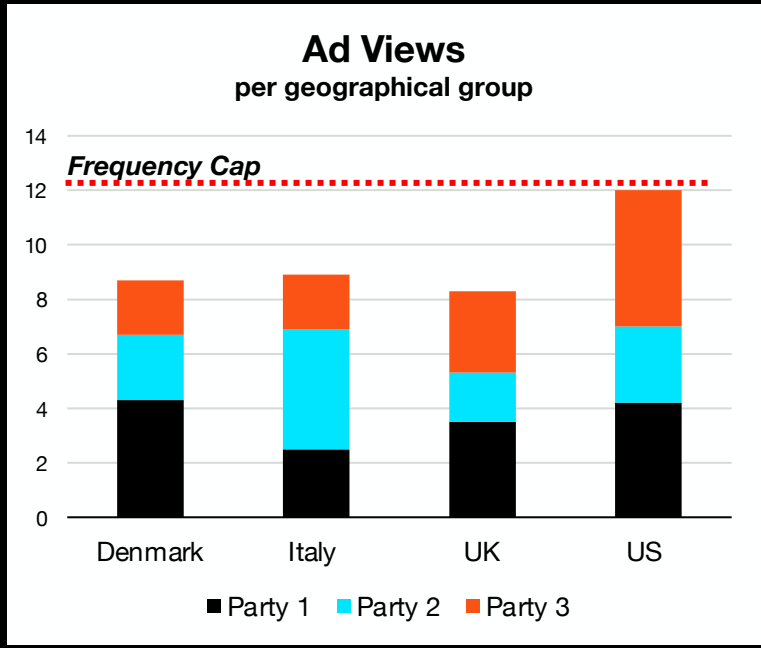

resolve.tech

Secure Multi-Party Computation to the rescue

01 Additive secret shares

02 SMPC as a blackbox

03 Usually not so user-friendly;
Usually not so fast

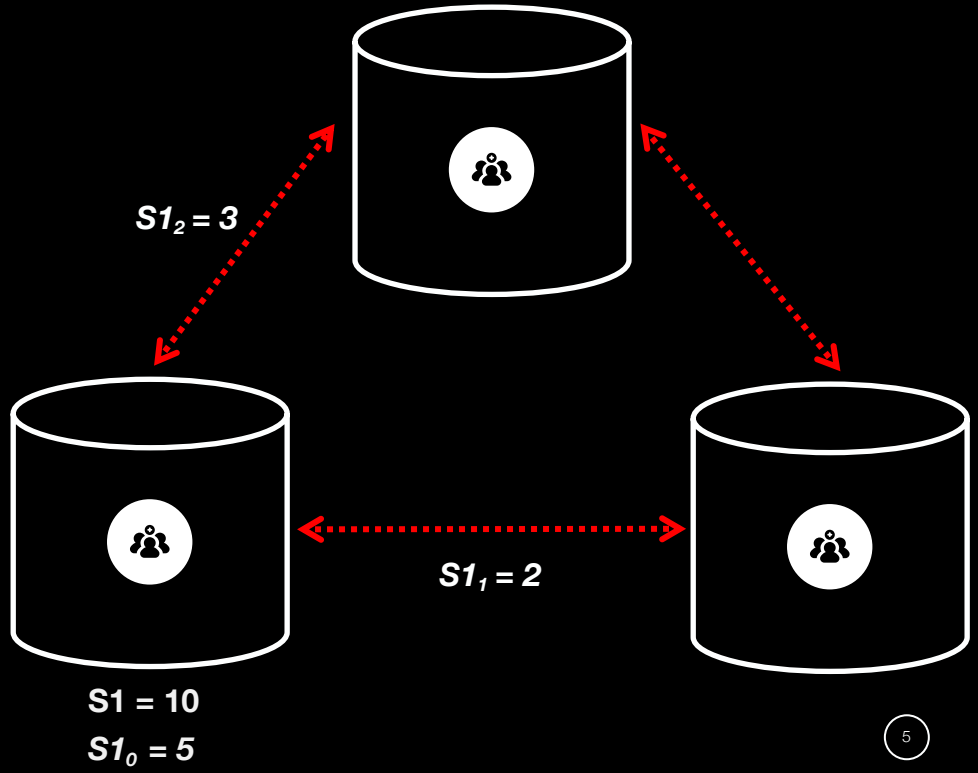

resolve.tech

Disclaimer:
input & output privacy

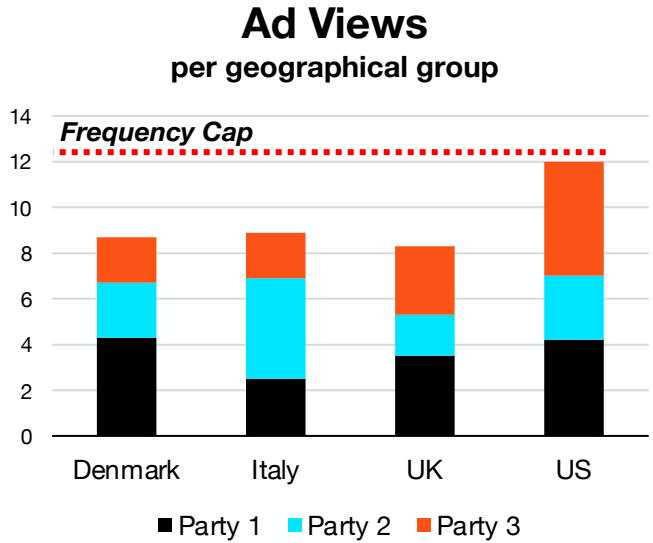

SMPC to protect inputs

Differential privacy to protect outputs
(not the focus of this talk)

Analytics as an example

More complex computations are possible
(e.g. training ML models)

02 A system for democratizing private analytics with SMPC

resolve

resolve.tech

Enabling non-MPC expert devs to perform private analytics

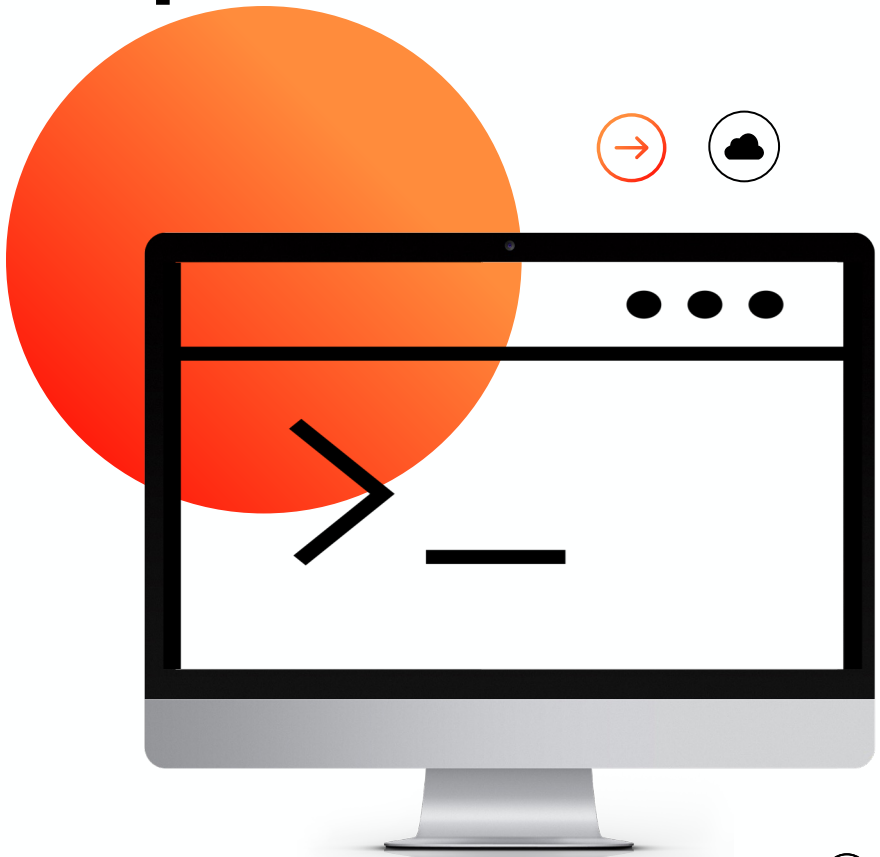

We were looking for...

- A framework usable by non-MPC expert developers
- With privacy guarantees underneath

resolve.tech

MP-SPDZ:
a friendly MPC framework

Protocols:
- **40+ protocols**
- Offline & Online phase
- **Malicious, Covert** & **Semi-honest** security

Library:
- **Open source**
- Well-maintained
- **Python-like** interface

Limitations:
- **No dictionaries** / hashmaps
- Memory allocated at **compile time** *(data structure cannot grow at run time)*

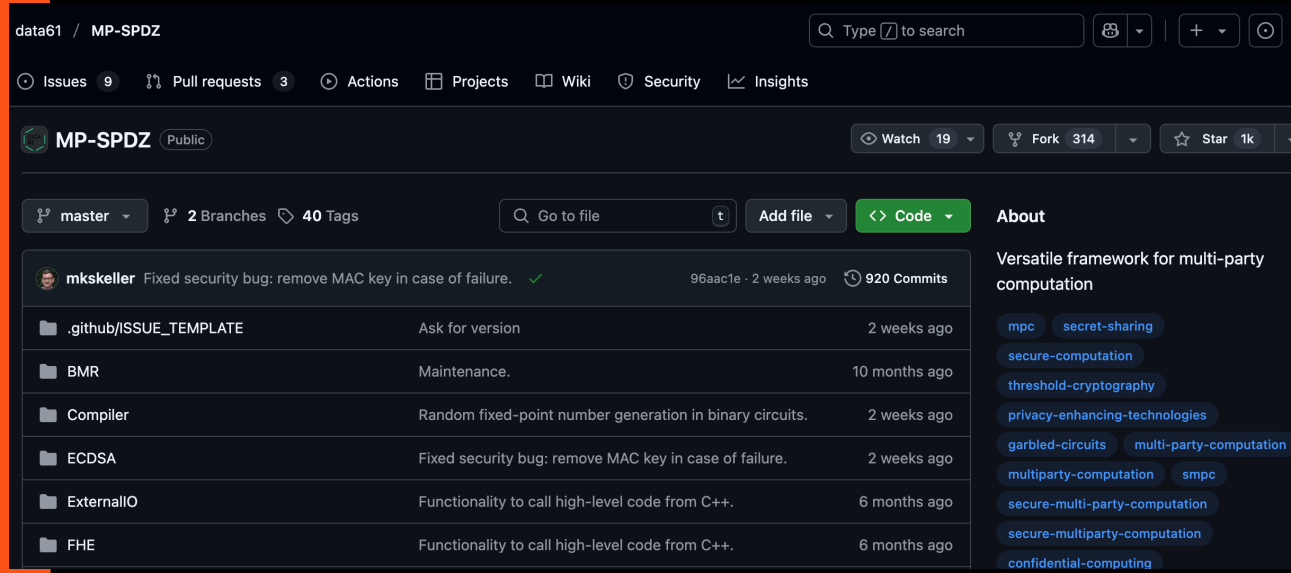

03 Carbyne Stack & MP-SPDZ

resolve

resolve.tech

Carbyne Stack donation

resolve.tech

Carbyne Stack: why?

01

Open-source: active development

Active community

02

Deploy and scale MP-SPDZ

But flexible and backend-independent

03

Client/Server style MPC

Data providers offload computation to computational parties*

**Damgård et al. – Confidential Benchmarking based on Multiparty Computation*

resolve.tech

Carbyne Stack: what?

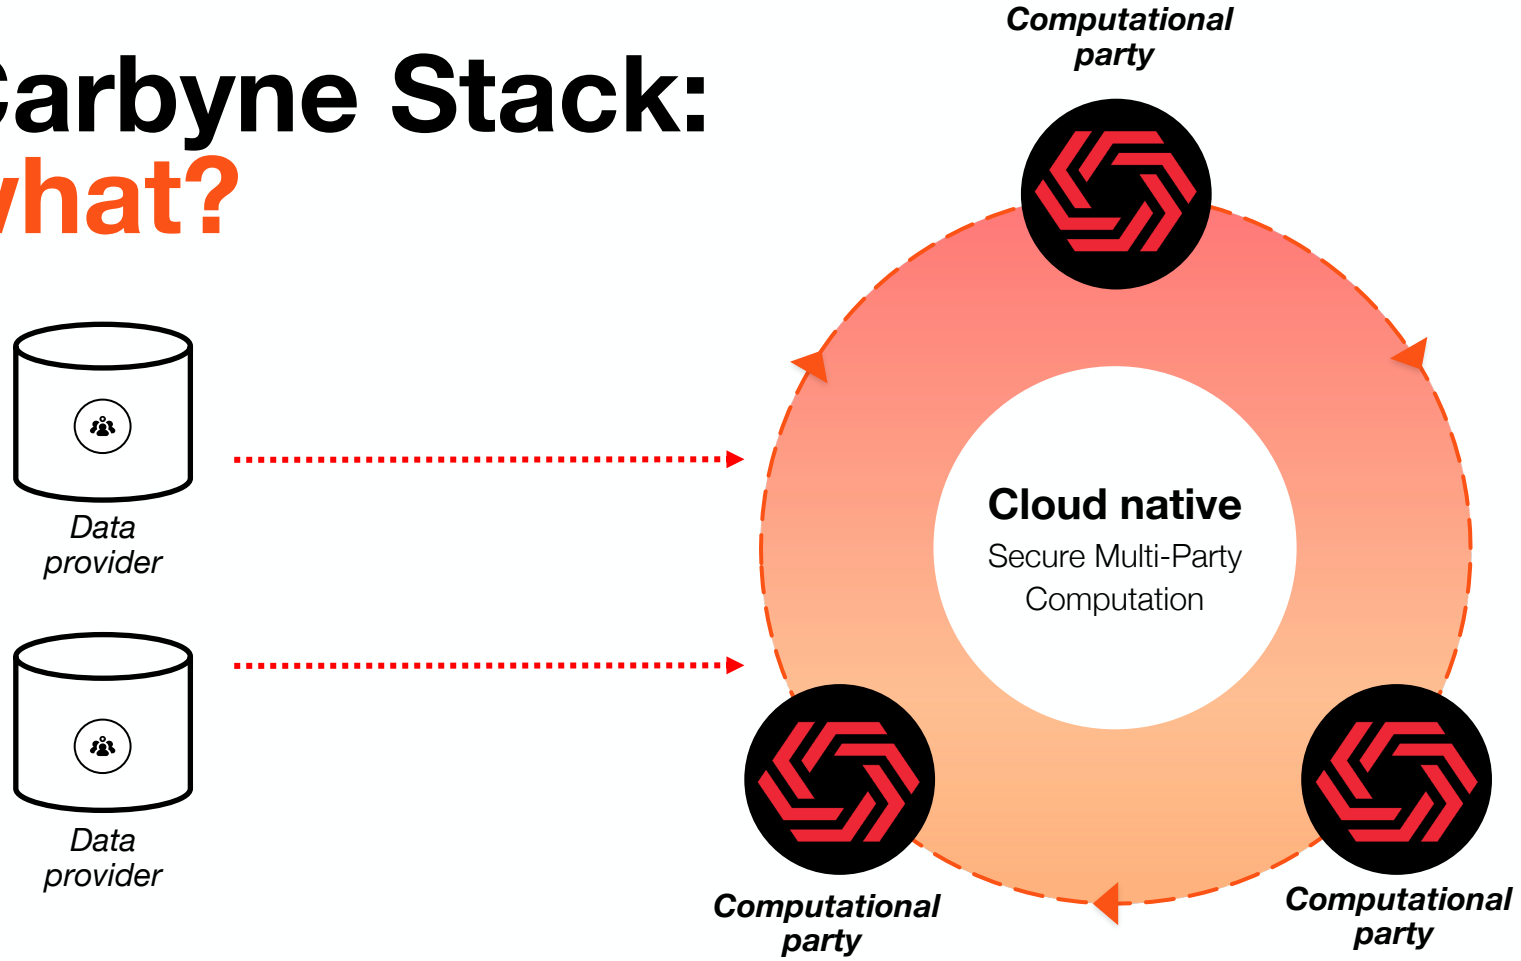

CARBYNE STACK

carbynestack.io
github.com/carbynestack

Carbyne Stack: what?

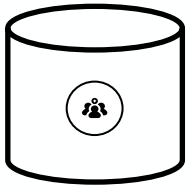

Data provider

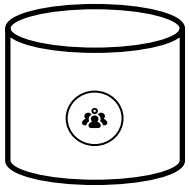

Data provider

- No party sees **raw data**
- Data are split into **secret shares**
- Secret shares uploaded to **secret store** of computational parties

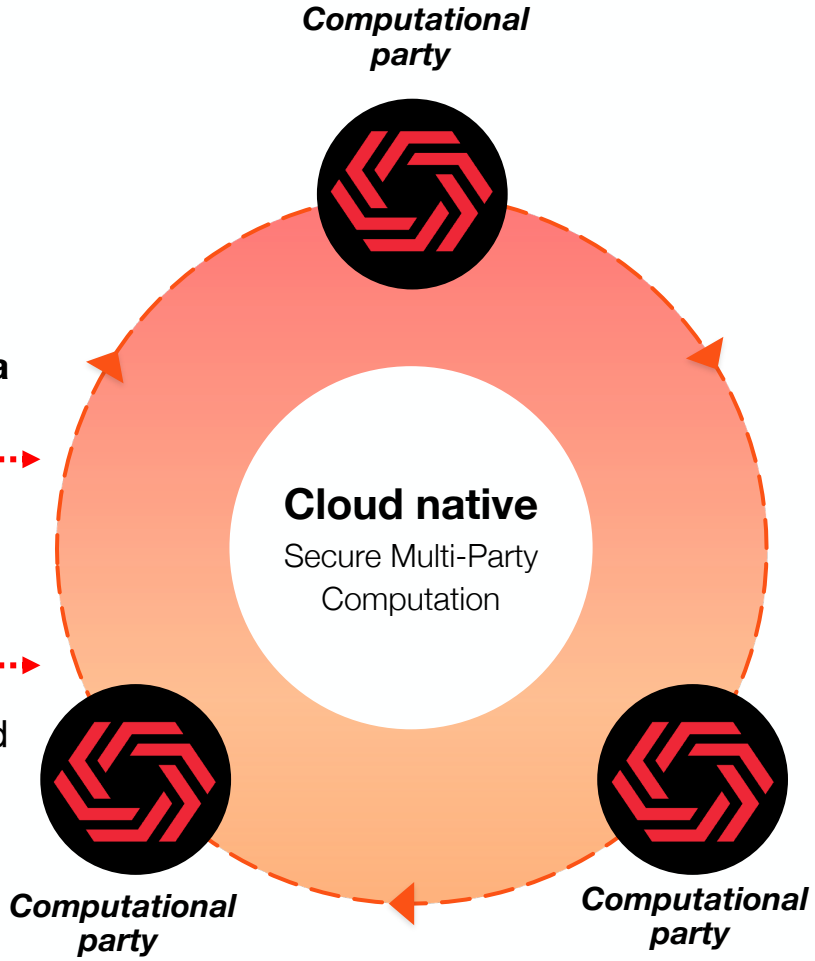

Carbyne Stack: components

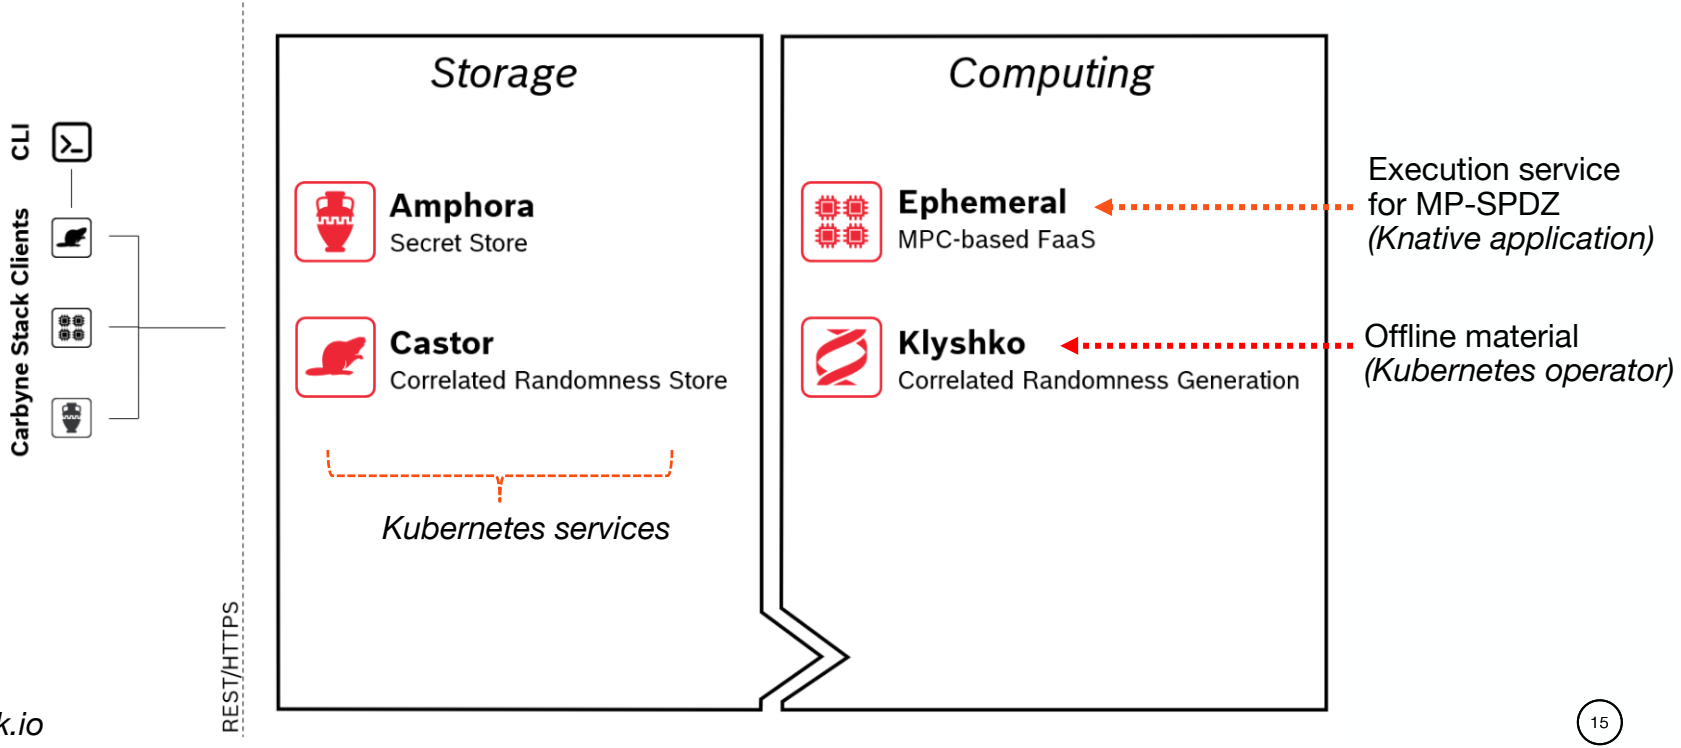

Image from carbynestack.io

Carbyne Stack:
possible extensions

Covert security, dishonest majority only

Semi-honest?

SMPC Protocols are network-bandwidth heavy

Infrastructure optimization?

SMPC protocols: sequential operations

Improve parallelism?

04 System enhancements

resolve

04 System enhancements:

04.1 Security model

Security model: semi-honest

More efficient
in real-world environment

Parties follow protocols
(and are bound by legal contracts)

Added tuples (*offline material*) support

Covert: *tuple* + *MAC*
(no MAC in semihonest)

04 System enhancements:

04.2 Infrastructure

resolve.tech

Zone Affinity

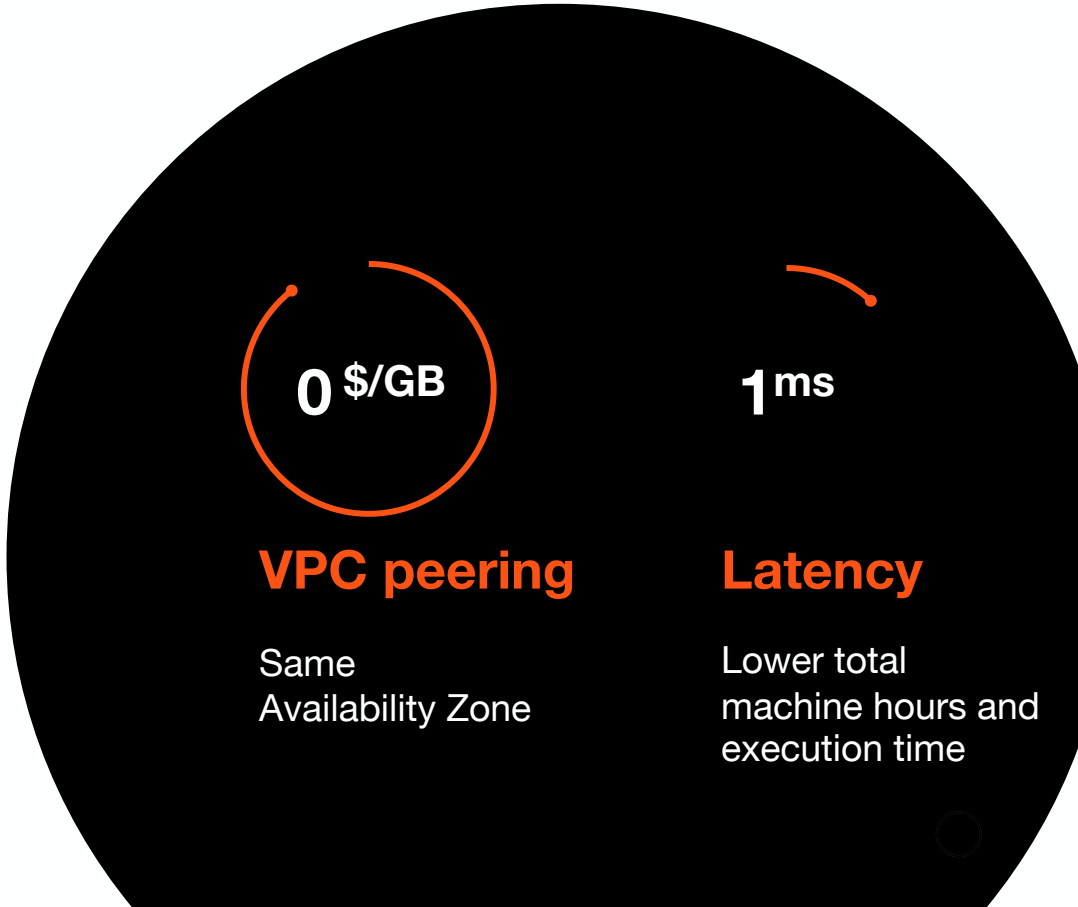

Disclaimer: same cloud provider

Threat modelling assumption:
data were already on the cloud

resolve.tech

Running on spot nodes:

90% cost reduction

Spot nodes are non-reserved nodes

Significantly cheaper

Downside:
more frequent crash scenarios

- Orchestrator for execution service
- Acceptable to take slightly longer

04 System enhancements:

04.3 Algorithmic

Fundamental function: Group by

We wanted something *general*:

groupby().agg()
as a primitive to developers

Sort + accumulation
Added *daBits* and *edaBits* support in Carbyne Stack
for efficient binary ops (e.g. comparisons, conditional logic…)

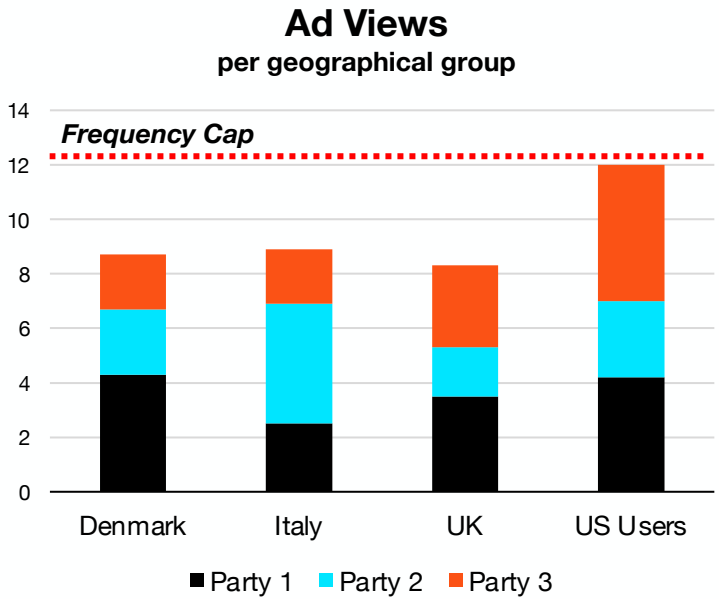

resolve.tech

Parallel sorting for complex analytics

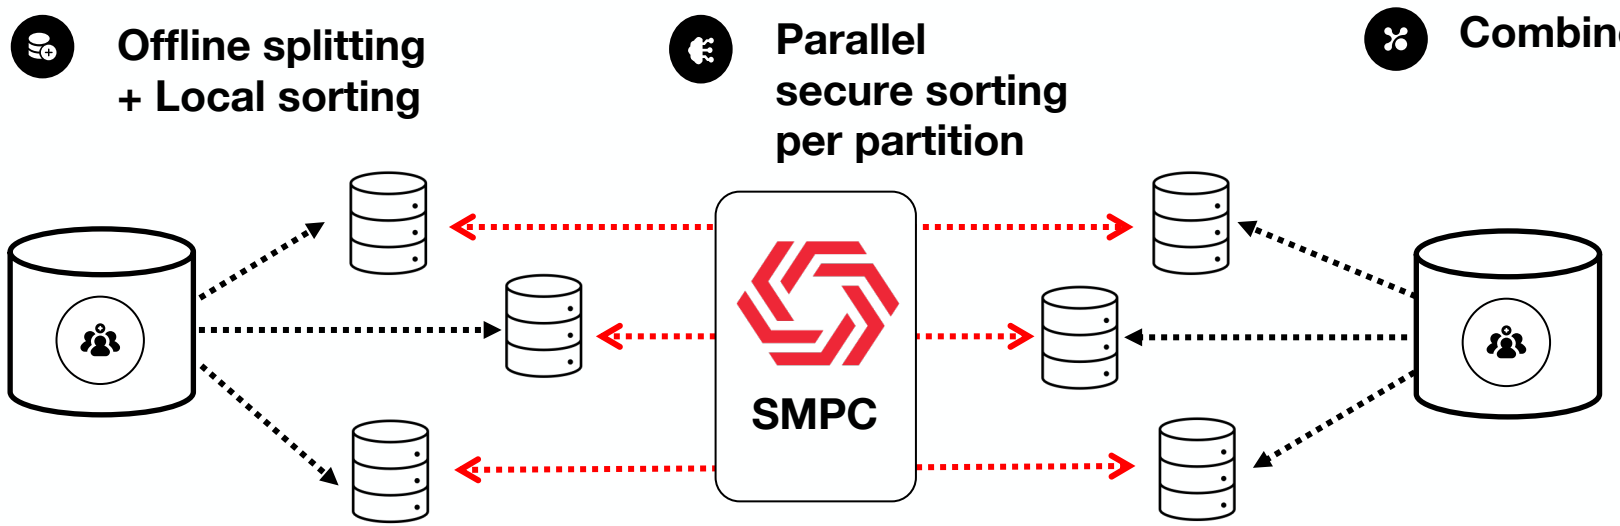

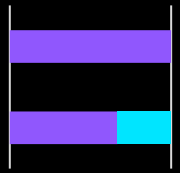

Fixed-size data structures

Padding to match size

05 Conclusion

resolve

resolve.tech

Practical MPC is possible

MPC as a tool to assure privacy
and available to non-MPC experts

Open-source projects
e.g. Carbyne Stack and MP-SPDZ

System enhancements
security models; algorithmic level

Infrastructure optimization
for real-world deployment


Thank You
For your time
resolve
resolve.tech

resolve.tech

The Team

Adrián Vaca

Goran Stipcich

Gerardo González

Daniele Romanini

Sergio Reinosa