

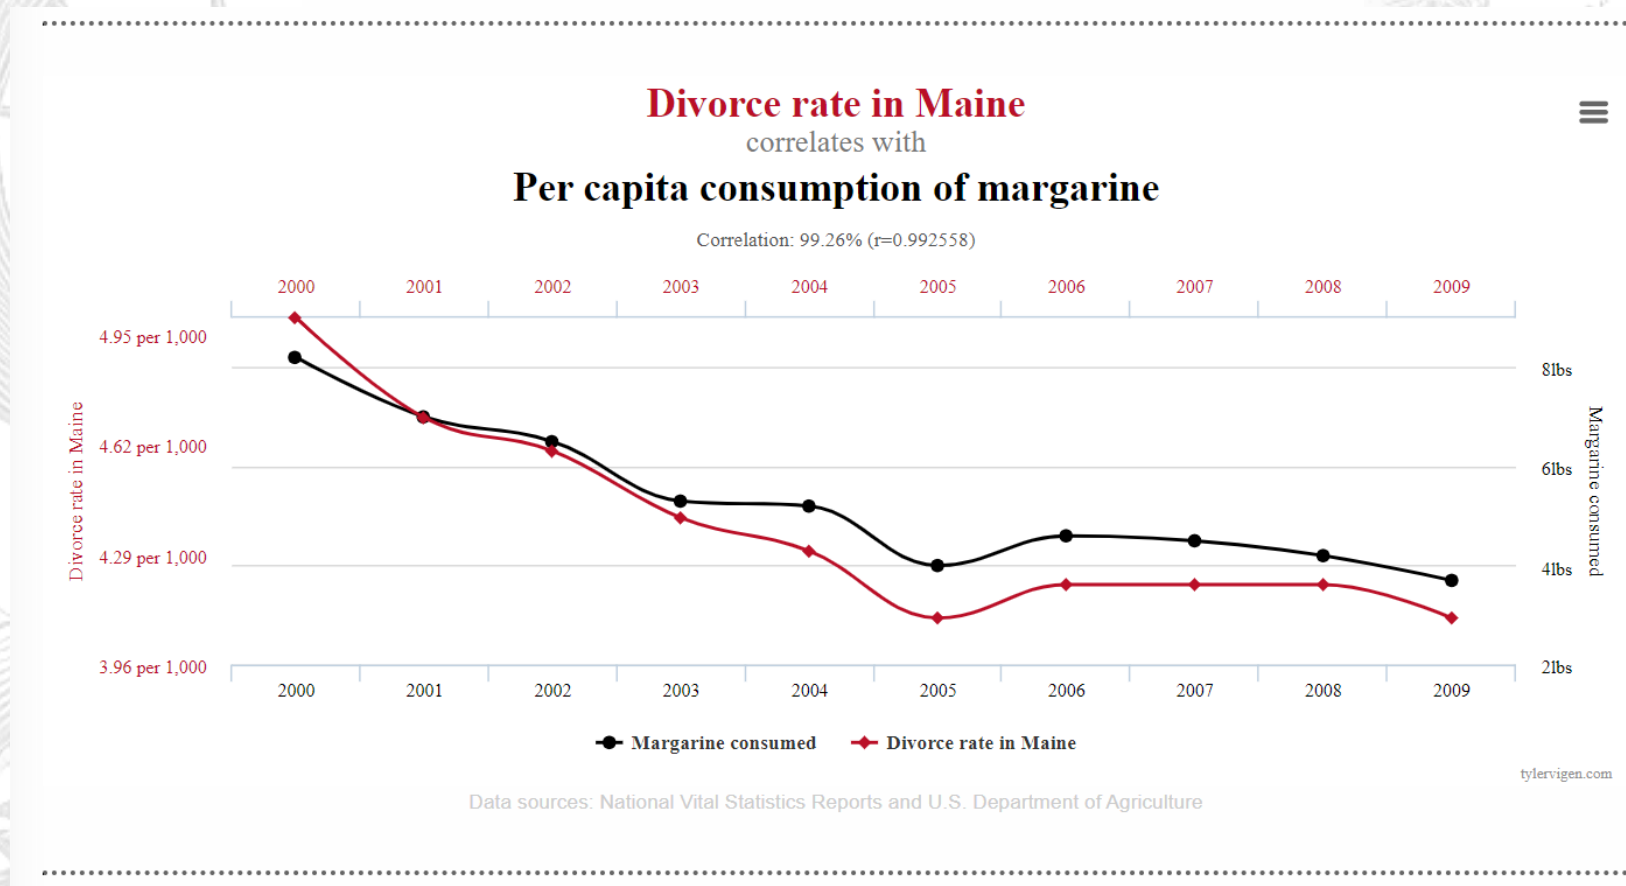


2023+ Data, Compute, Network, Power Consumption

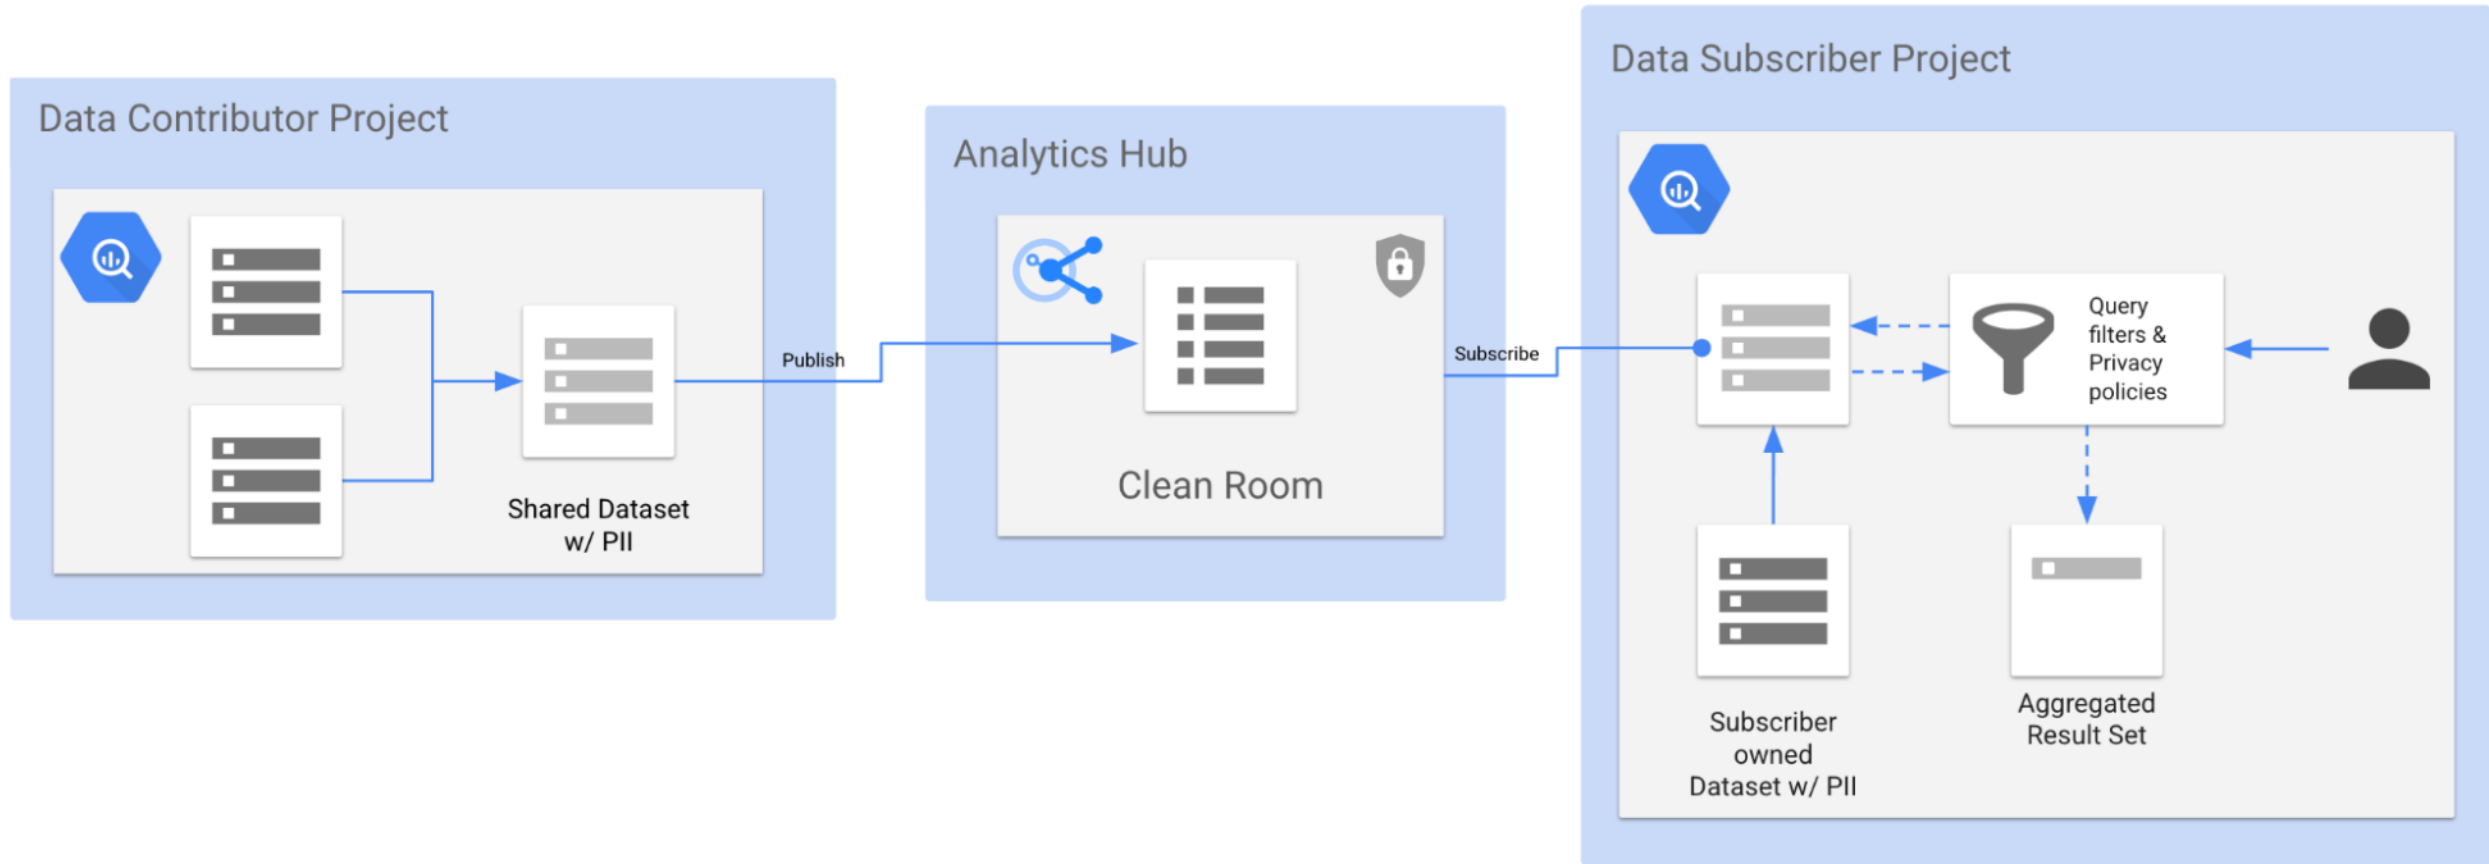
Alisson Sol

Relevant data can lead to irrelevant insights!

Spurious correlations



Data Clean Rooms



<https://cloud.google.com/bigquery/docs/data-clean-rooms>

AI is here... Let's talk with AI for a bit...

OP

You

Just the number: what is the maximum bandwidth for the Internet connection a company can buy today in the USA?



ChatGPT

The maximum bandwidth for an internet connection that a company can buy in the USA varies based on location, provider, and the type of connection (like fiber-optic, cable, DSL, etc.). As of my last update in April 2023, some of the highest bandwidth connections offered for commercial use were in the range of 10 Gbps (Gigabits per second) to 100 Gbps. These are typically offered through dedicated fiber-optic lines.

However, it's important to note that these figures can change rapidly as technology evolves and network infrastructure improves. For the most current and specific information, it's best to contact internet service providers (ISPs) directly, as they can offer tailored solutions based on a company's location and specific needs.



You

Just the number: What is the largest LLM publicly available for download in number of parameters?



ChatGPT

As of April 2023, the largest publicly available language model for download is Google's "PaLM" (Pathways Language Model) with 540 billion parameters.



You

Just the number: assuming 8 bytes per parameters, what is the size in TB of a model with 540 billion parameters?



ChatGPT

The size of a model with 540 billion parameters, assuming 8 bytes per parameter, is approximately 3.93 TB. [[↗](#)]

The background of the slide features a detailed illustration of various mechanical gears of different sizes and a bucket filled with popcorn, with a spoon resting inside. The style is a fine-lined, sketch-like drawing in a light gray color.

OP

You

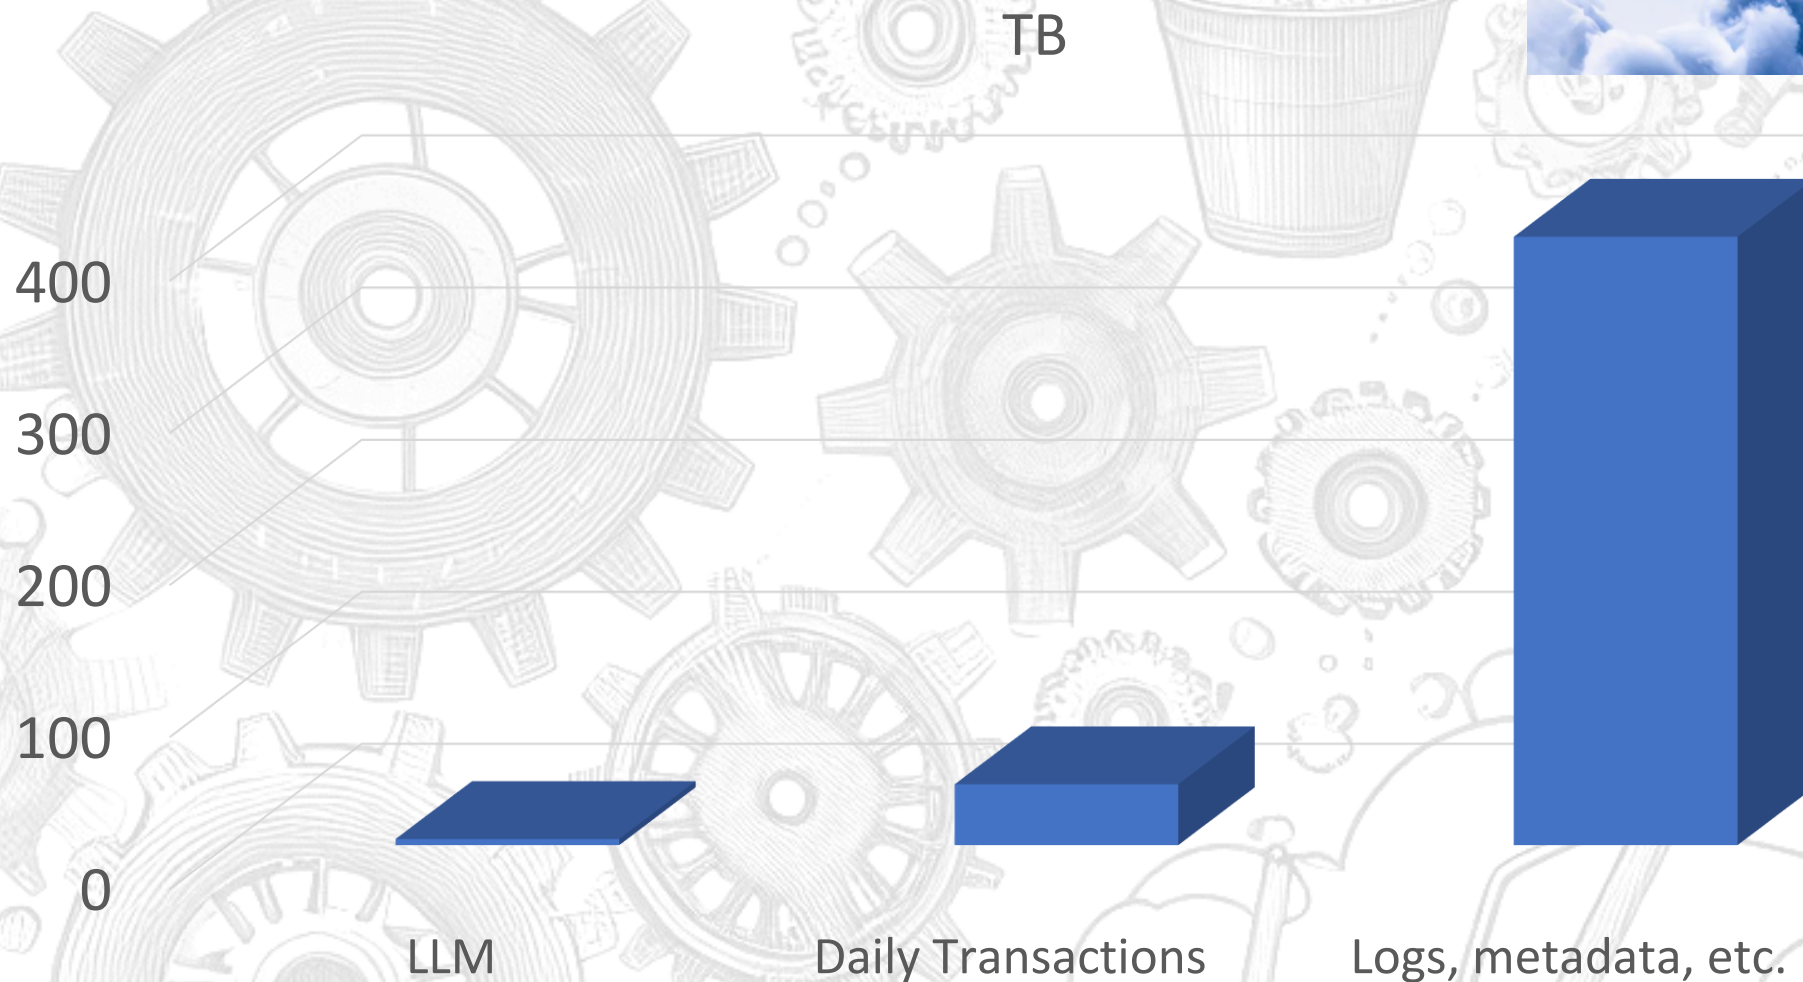
Assuming perfect performance, how long it would take to download 4 TB of data using an Internet connection of 100 Gbps?



ChatGPT

To download 4 TB of data using an internet connection of 100 Gbps, it would take approximately 328 seconds. This is equivalent to about 5 minutes and 28 seconds, assuming perfect performance and no interruptions or inefficiencies in the download process. [\[>-\]](#)

Houston, we have a problem?



Data Gravity

OP

You

What is "Data Gravity"?

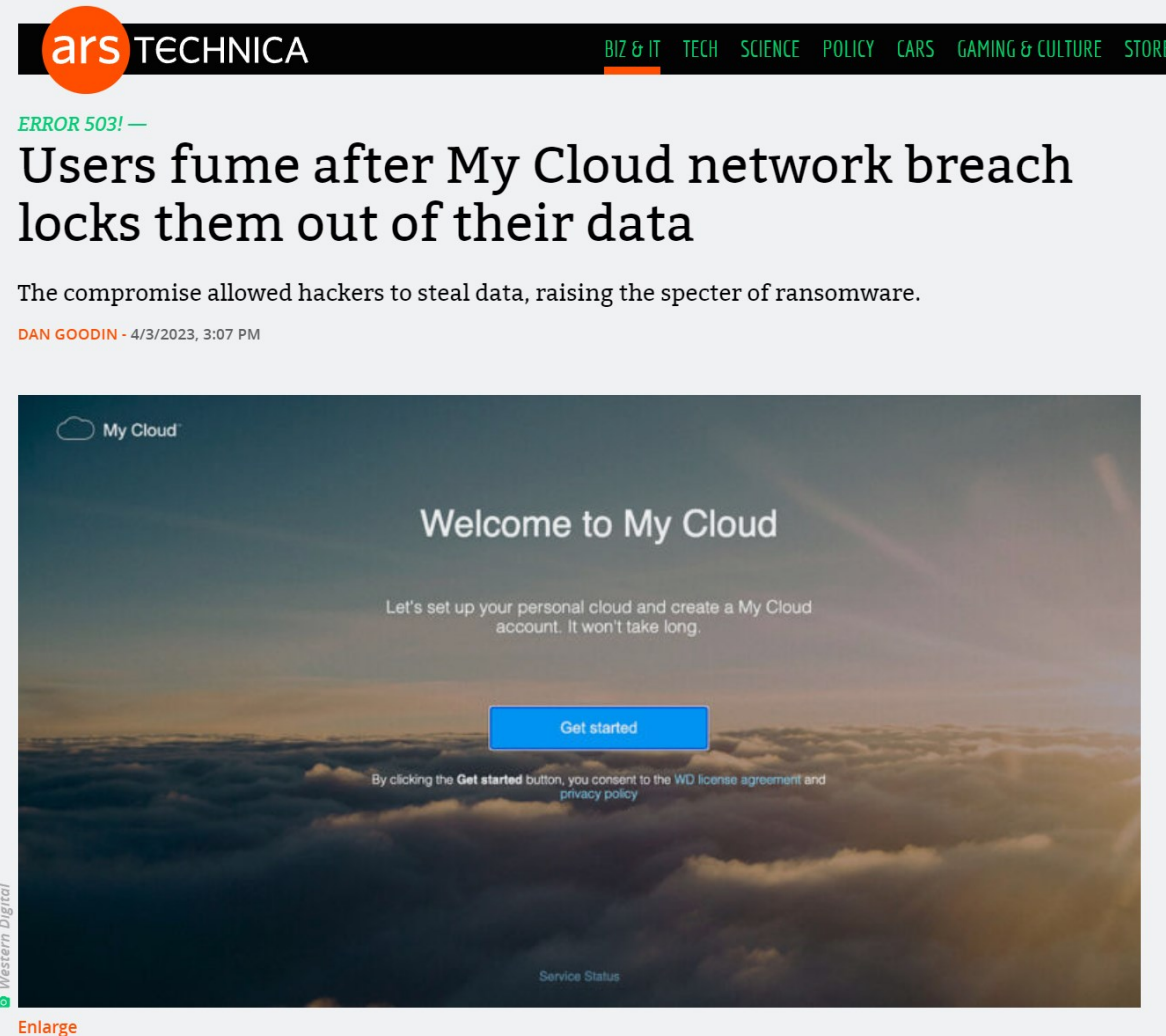


ChatGPT

"Data Gravity" is a concept in the field of data management and cloud computing. It was first coined by Dave McCrory in 2010. The idea behind Data Gravity is that as datasets grow larger and more complex, they become more difficult to move. This phenomenon is likened to the way a large object in space (like a planet) exerts a gravitational pull that attracts other objects to it.



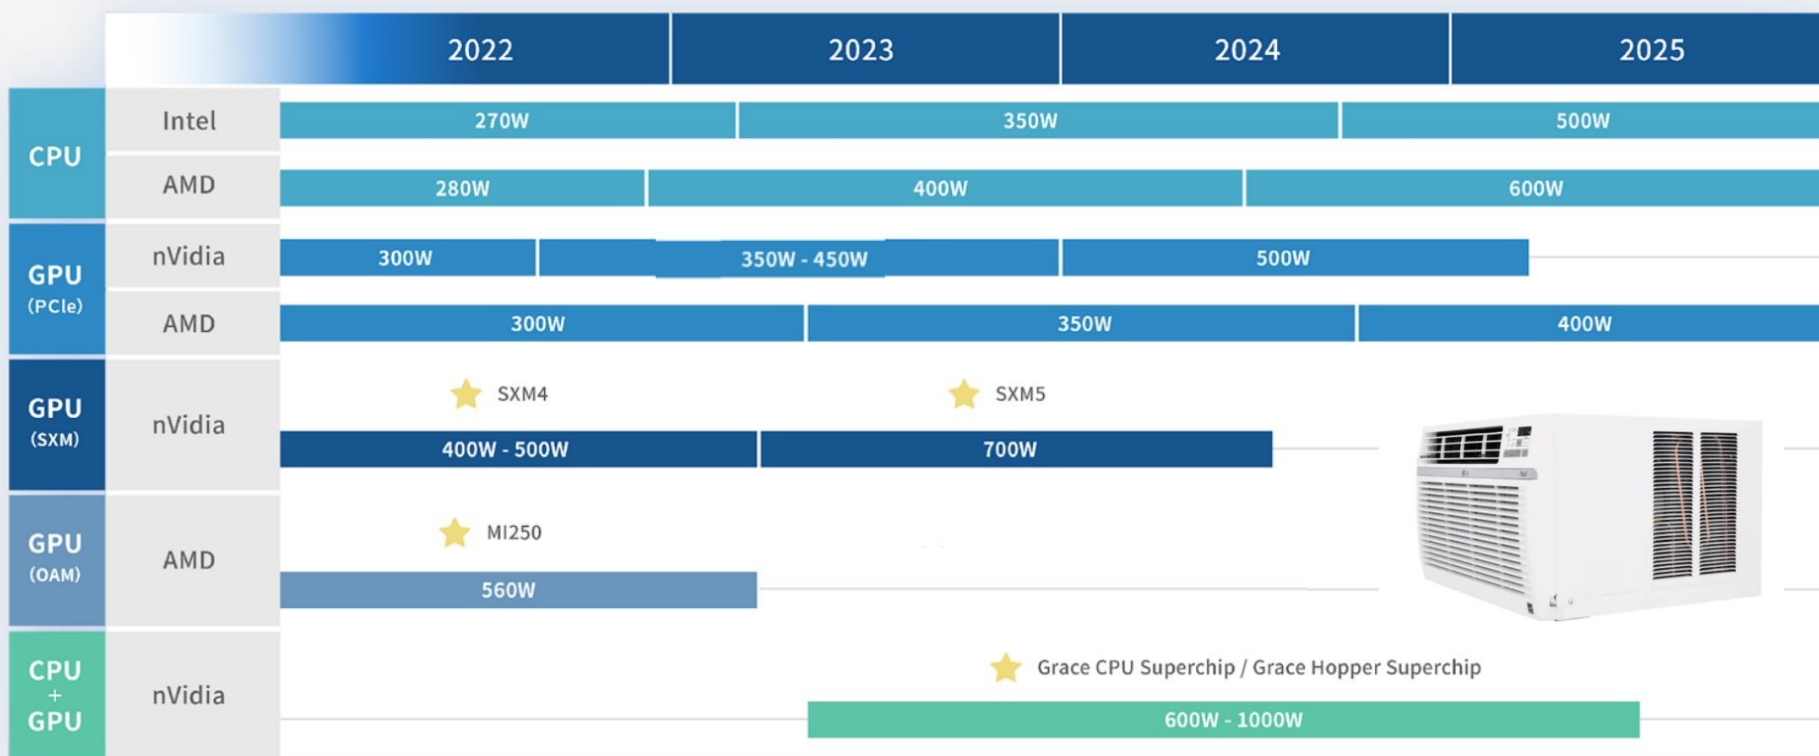
Cloud locked-in, yet locked out of your data!



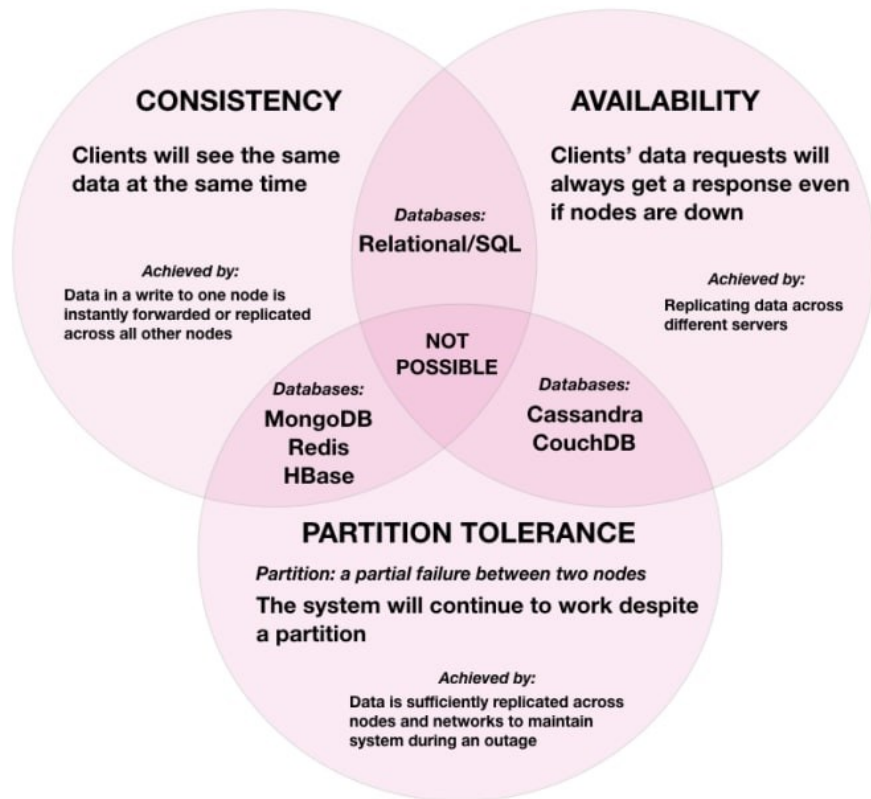
Power!

GIGABYTE™

CPU/ GPU Power Consumption



Distributed: Processing, Data, Networking

[Latest Client Download](#)[Terms](#)[Previous Versions](#)

Connection Health Check

Service Availability	US West (Oregon)	US East (N. Virginia)	South America (São Paulo)	Europe (Frankfurt)	Europe (Ireland)	Europe (London)
Workspace Services ⓘ	✓	✓	✓	✓	✓	✓
Round Trip Time (ms)* ⓘ	50	13	101	68	69	62
Speed Rating ⓘ	FAST	FAST	ACCEPTABLE	FAST	FAST	FAST
Recommended Region						

*Round trip time is the time to connect from your browser to the Amazon network, not the average latency for a given region.

Round trip time may vary due to network conditions. The region recommendation is based on service availability and latency. [Learn more about](#)

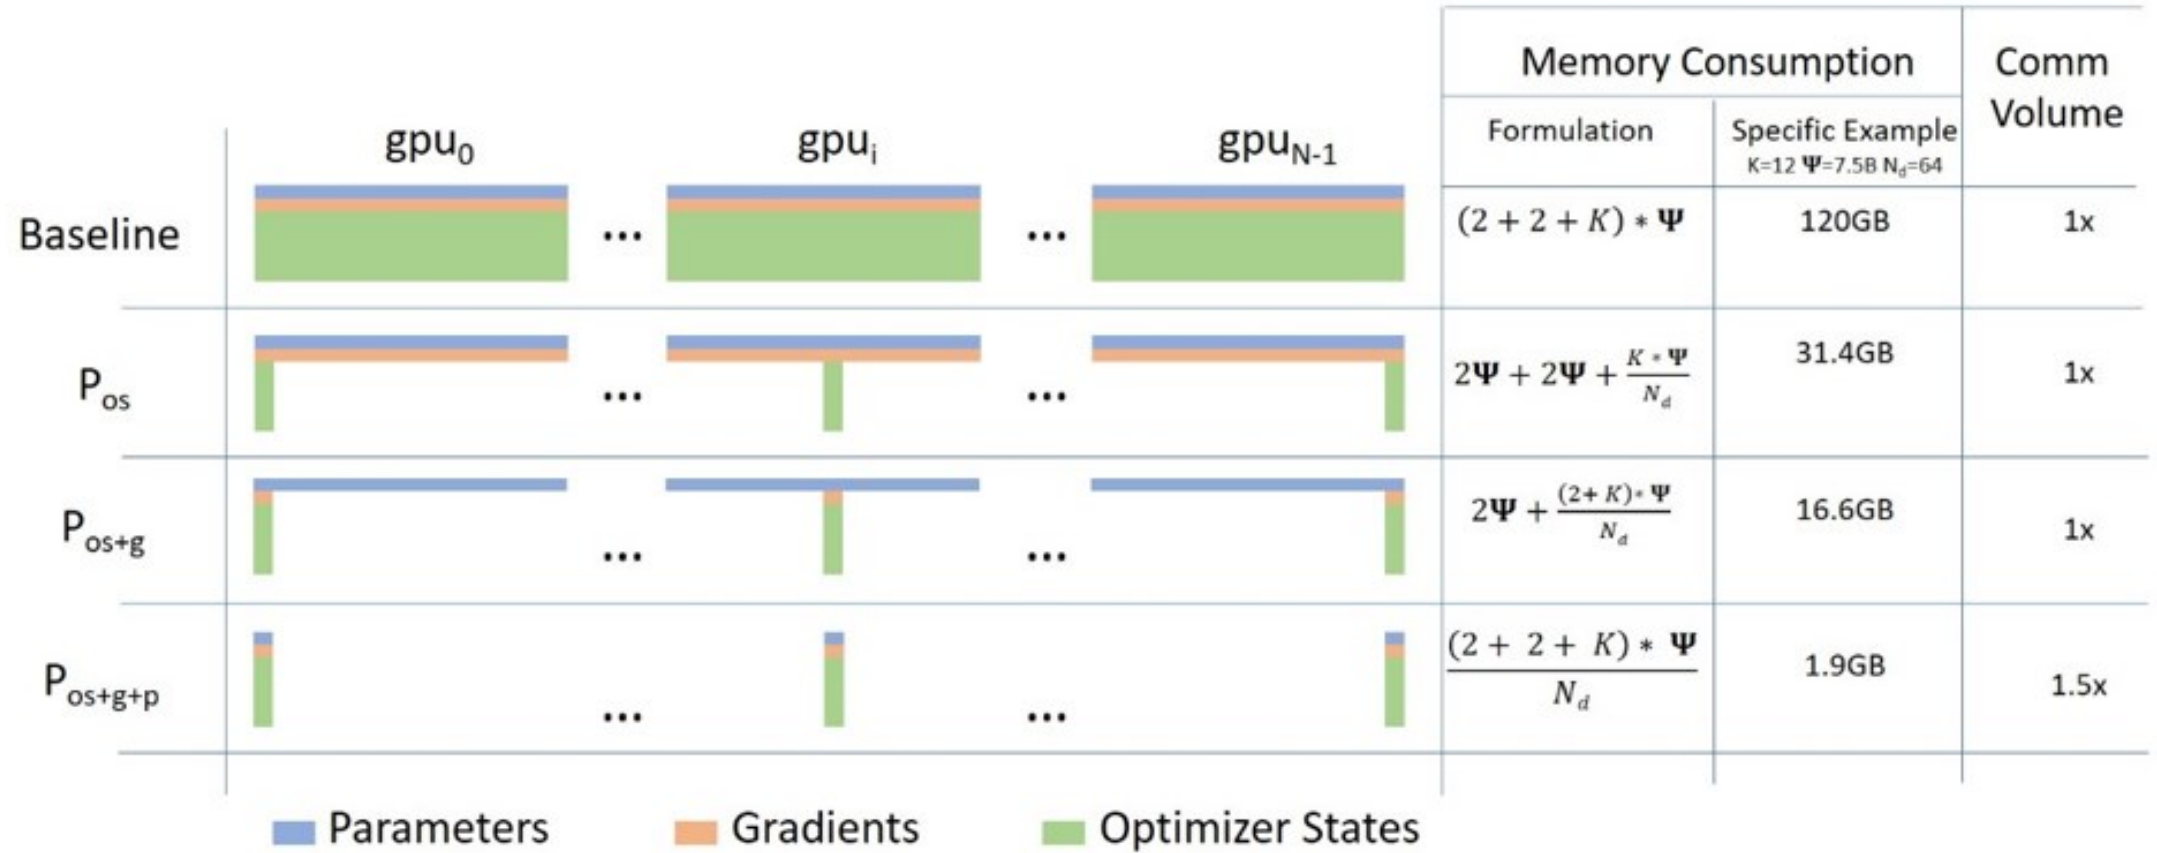
If this page is accessed from within a WorkSpace it will show results between your WorkSpace and the available regions.

[Check Again](#)

Cloud vendors are realistic already...

- Raw data in cold storage: cheaper
- Processed signals in hot storage
 - With provenance information
- Multi-cloud signal processing is available now
 - Right processing and data architecture: costs go down, effectiveness up

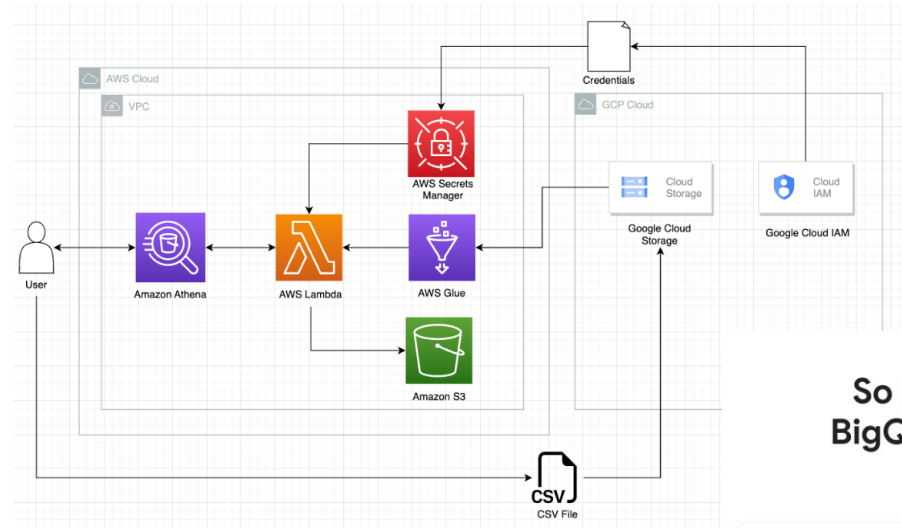
Even in your “local cloud”...



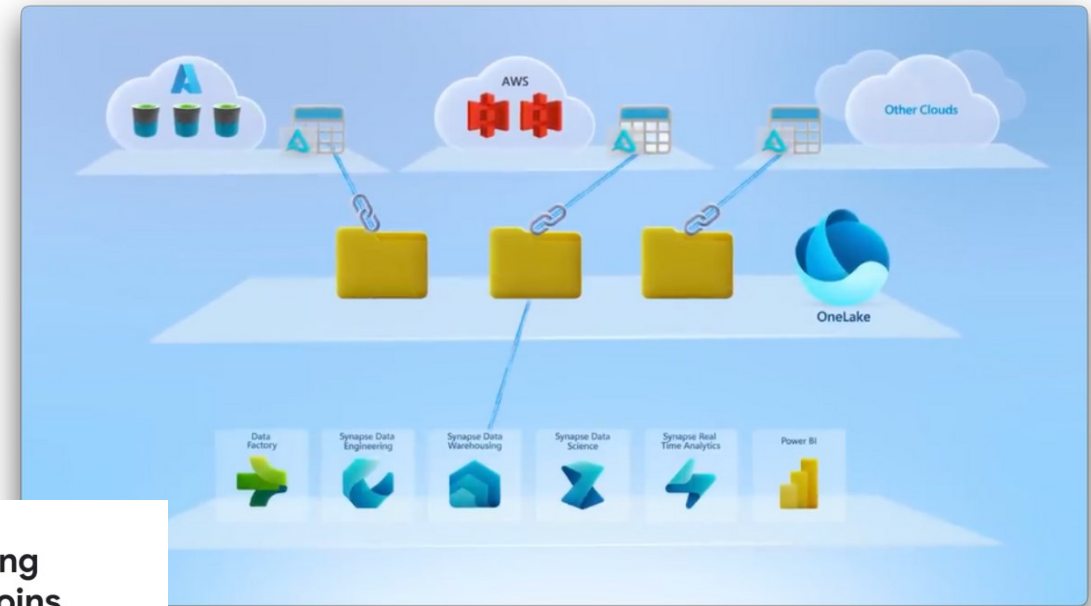
Cross-cloud data analysis

Solution overview

The following diagram illustrates the solution architecture.



Bring data from Azure storage accounts & other clouds into one unified data lake.

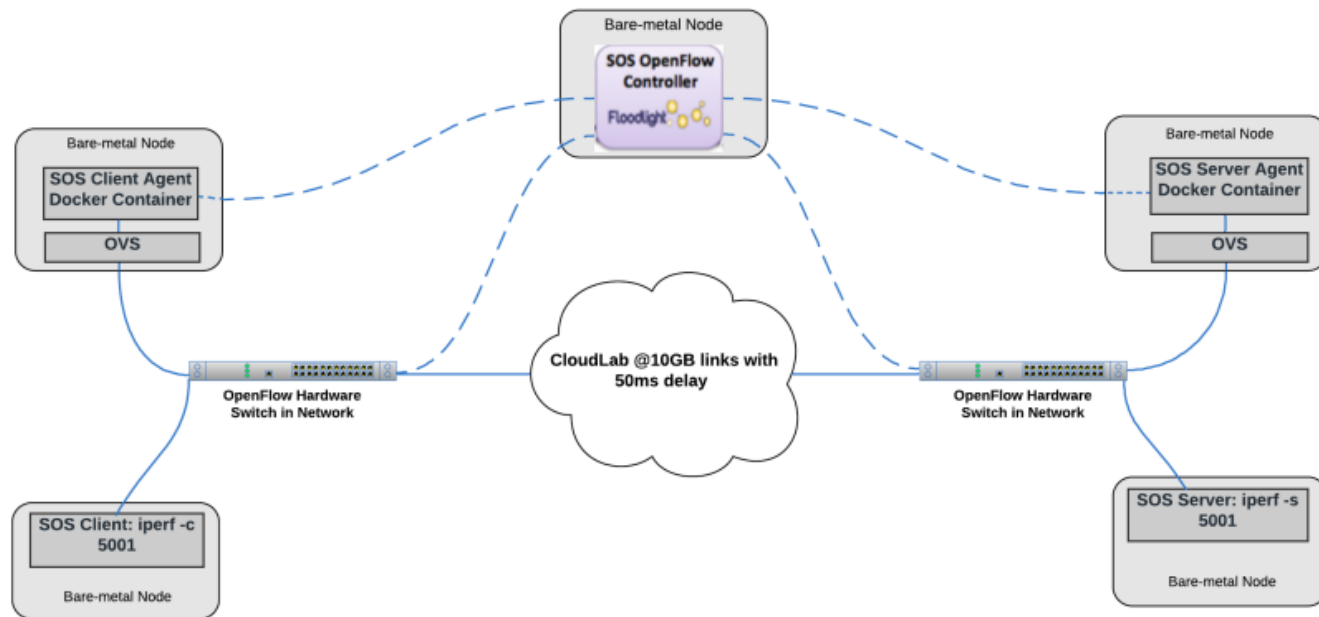


So long data silos: Announcing BigQuery Omni cross-cloud joins

September 26, 2023

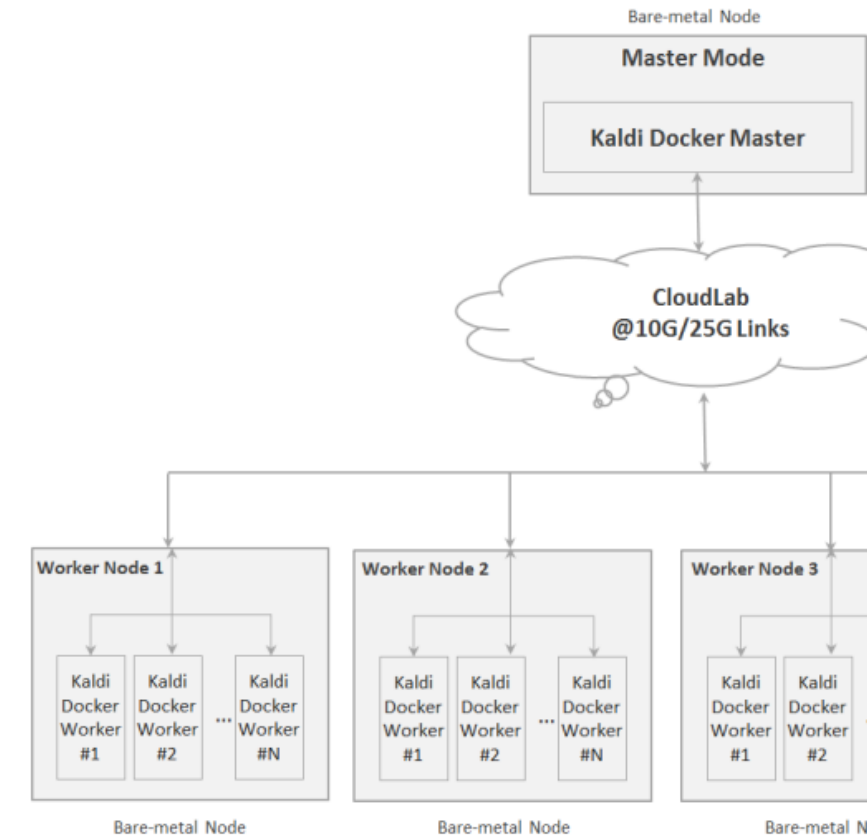


Nascent academic and practical field



Note:
Docker version: 19.03.11
Operating system: Ubuntu 16.04.1 LTS

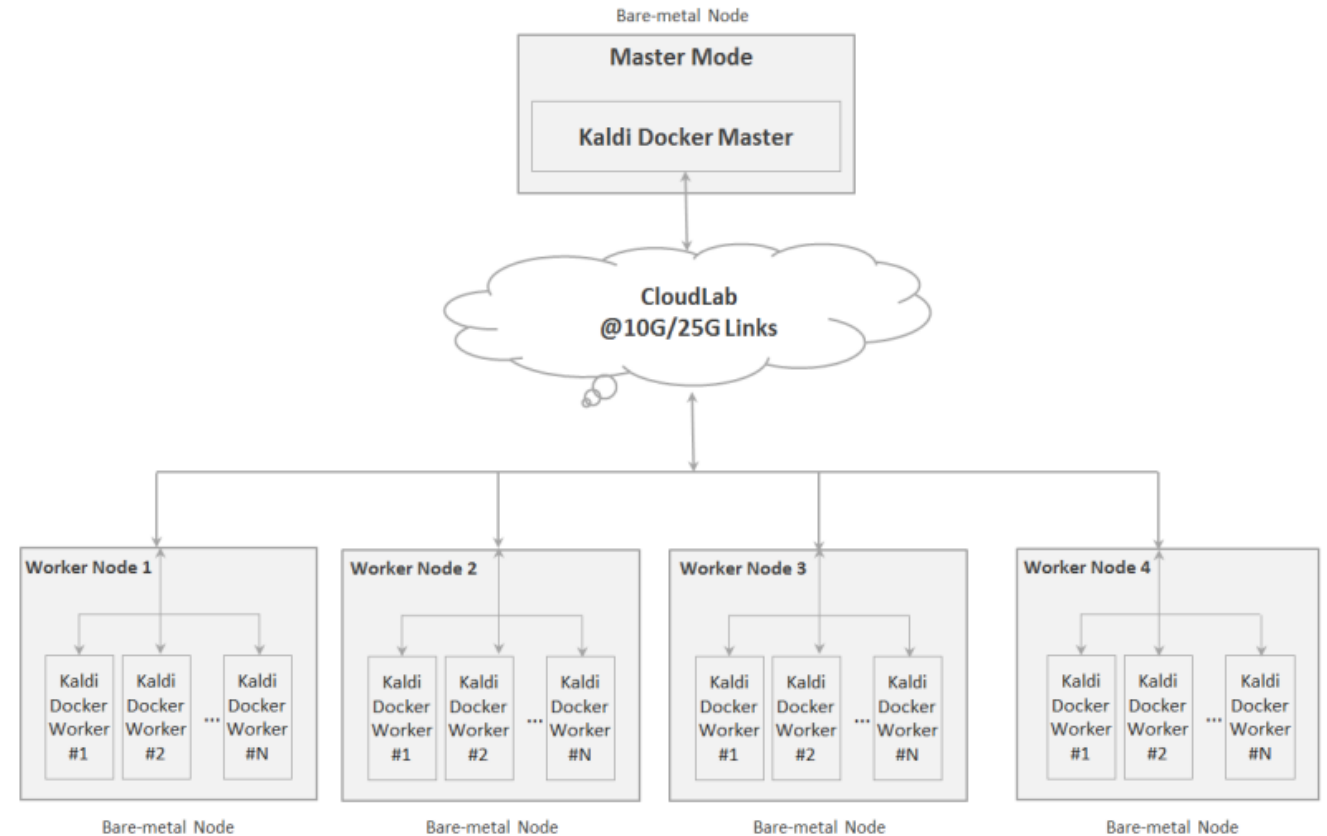
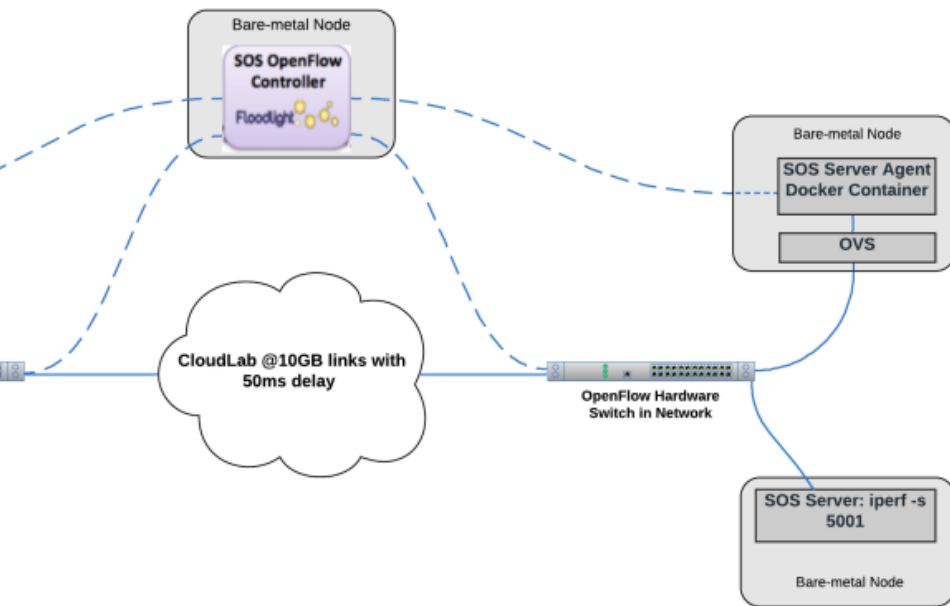
Figure 3: Containerized SOS deployment topology in the CloudLab testbed



Note:
Docker version: 19.03.11
Operating system: Ubuntu 16.04.1 LTS

Figure 4: Containerized Kaldi deployment topology

Nascent academic and practical field



Note:

Docker version: 19.03.11
Operating system: Ubuntu 16.04.1 LTS

ized SOS deployment topology in the CloudLab testbed

Figure 4: Containerized Kaldi deployment topology in the CloudLab testbed

What is cloud attestation?



ChatGPT

Cloud attestation is a security process in cloud computing that assures the integrity and trustworthiness of cloud services and infrastructure. It involves verifying that the cloud environment, including hardware, software, and firmware, remains secure and uncompromised. This is crucial in scenarios where sensitive data or critical applications are involved.

A key component of cloud attestation is remote attestation, where a cloud server demonstrates to a client or user that it runs specific, unaltered software in a secure configuration. This process often includes integrity checks, ensuring the cloud infrastructure, especially during boot-up, hasn't been tampered with.

The concept is closely linked to secure boot processes and Trusted Execution Environments (TEEs), which provide protected areas within processors to maintain the confidentiality and integrity of code and data. Cloud attestation is particularly significant in multi-tenant cloud environments, offering each client assurance about the security and isolation of their resources.

In summary, cloud attestation forms a vital part of cloud security, helping to establish trust in cloud services by confirming that they operate securely and as intended, which is essential in meeting compliance standards in regulated industries.

Message Summary

- The cloud distributed processing is real
- It demands investment in infrastructure
 - Test your infrastructure before your "processing layer"
 - Investigating a "Python timeout" or "Too many requests" inside a notebook is super hard in a distributed environment

