



# Cloud Rationalization Next Year's Buzzword

March 6, 2025

Nathan Moore  
Principal Architect

# Our Story So Far...

## A Brief History of Data Centers

- ENIAC in 1945
  - University of Pennsylvania, required a dedicated building
- Computer Rooms
  - 1960's dedicated rooms in existing buildings to house transistorized computers
- Network Closet
  - 1980's Rise of client-server architecture, server housed in network closet and attached to office network
- Data Center
  - 1990's Internet connectivity and dedicated servers in centralized facility, back to a shared, dedicated building
- Cloud
  - 2000's Virtualized services delivered from shared physical servers
- Hyperscalers
  - 2010's Massive build-out of physical infrastructure to support expanded Cloud offerings
- Infrastructure-as-a-Service
  - 2020's The entire physical stack is now programmatically available with multiple providers at every level

**Growth has been driven by the quest to find economies of scale and economies of scope in Compute**

# A Brief History of Cloud

- Named in 1994
- VMWare virtualization in 1999
- AWS EC2 in 2006
- OpenStack in 2010
- Terraform in 2014
- Kubernetes in 2014
- \*-as-a-Service
  - API-driven services to be interconnected via common protocols

**Layers of abstraction mean physical infrastructure is well hidden**

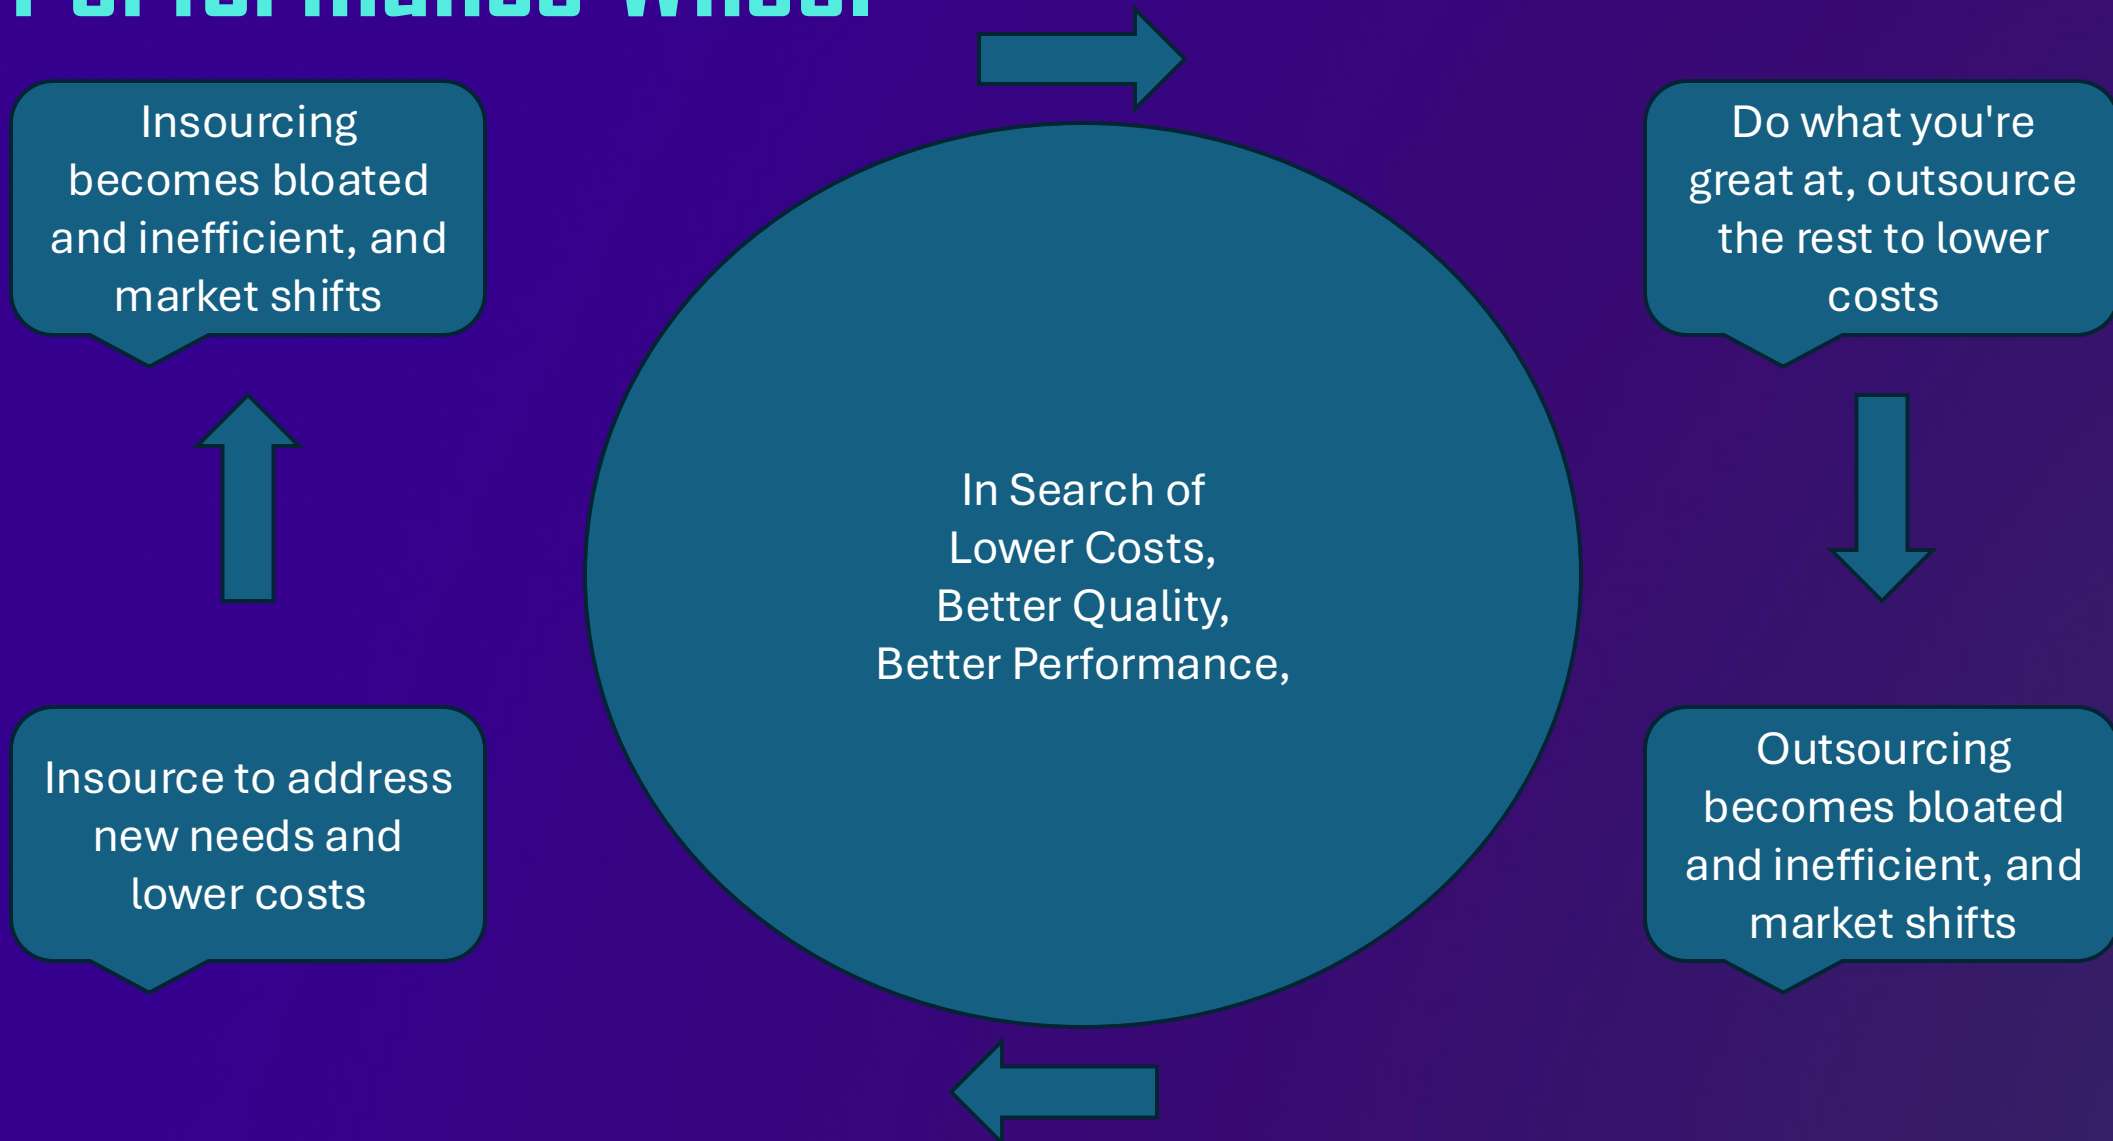
# How Did We Get Here

Cloud has been dominant for over a decade

- Cloud-only apps and services have been built and run and grown successfully for years
- There are plenty of companies with no expertise in anything behind the Cloud provider's API
- Infrastructure-as-a-Service has become "Service," and Infrastructure is a low class, dirty word Software does not speak

**There are lots of organizations which have either lost or never had expertise in systems and networks**

# Performance Wheel



# So What's the Problem

Cloud-native companies are discovering problems with

- Cost
- Control
- Performance
- Compliance
- Scope and Size of Service

**Cloud-native means Cloud-dependent**

# Cost

Cloud	Self Hosted	Versus																																								
<p>Costs are entirely variable, and increase linearly with usage. A second VM doubles the cost of a single VM.</p> <p>Costs vs. Number of VMs</p> <table border="1"><caption>Cloud Costs vs. Number of VMs</caption><thead><tr><th>Number of VMs</th><th>Total Cost</th></tr></thead><tbody><tr><td>0</td><td>\$0</td></tr><tr><td>1</td><td>\$100</td></tr><tr><td>2</td><td>\$200</td></tr></tbody></table>	Number of VMs	Total Cost	0	\$0	1	\$100	2	\$200	<p>Costs are fixed up to a given server. The added cost of running a second VM on top of an existing server is roughly \$0.</p> <p>Costs vs. Number of VMs</p> <table border="1"><caption>Self Hosted Costs vs. Number of VMs</caption><thead><tr><th>Number of VMs</th><th>Total Cost</th></tr></thead><tbody><tr><td>0</td><td>\$500</td></tr><tr><td>1</td><td>\$500</td></tr><tr><td>2</td><td>\$500</td></tr></tbody></table>	Number of VMs	Total Cost	0	\$500	1	\$500	2	\$500	<p>Cost lines cross at some point, and it becomes cheaper to self-host.</p> <p>Costs vs. Number of VMs</p> <table border="1"><caption>Cost Comparison vs. Number of VMs</caption><thead><tr><th>Number of VMs</th><th>Cloud Costs</th><th>Self Hosted Costs</th></tr></thead><tbody><tr><td>0</td><td>\$0</td><td>\$500</td></tr><tr><td>1</td><td>\$100</td><td>\$500</td></tr><tr><td>2</td><td>\$200</td><td>\$500</td></tr><tr><td>3</td><td>\$300</td><td>\$500</td></tr><tr><td>4</td><td>\$400</td><td>\$500</td></tr><tr><td>5</td><td>\$500</td><td>\$500</td></tr><tr><td>6</td><td>\$600</td><td>\$500</td></tr></tbody></table>	Number of VMs	Cloud Costs	Self Hosted Costs	0	\$0	\$500	1	\$100	\$500	2	\$200	\$500	3	\$300	\$500	4	\$400	\$500	5	\$500	\$500	6	\$600	\$500
Number of VMs	Total Cost																																									
0	\$0																																									
1	\$100																																									
2	\$200																																									
Number of VMs	Total Cost																																									
0	\$500																																									
1	\$500																																									
2	\$500																																									
Number of VMs	Cloud Costs	Self Hosted Costs																																								
0	\$0	\$500																																								
1	\$100	\$500																																								
2	\$200	\$500																																								
3	\$300	\$500																																								
4	\$400	\$500																																								
5	\$500	\$500																																								
6	\$600	\$500																																								

# Control

## GCP auto-deleted the entire account of a \$125Bn pension fund

<https://www.theguardian.com/australia-news/article/2024/may/09/unisuper-google-cloud-issue-account-access>

- "UniSuper has approximately \$125bn in funds under management."

<https://cloud.google.com/blog/products/infrastructure/details-of-google-cloud-gcve-incident>

- "... there was an inadvertent misconfiguration of the GCVE service by Google operators due to leaving a parameter blank. This had the unintended and then unknown consequence of defaulting the customer's GCVE Private Cloud to a fixed term, with automatic deletion at the end of that period."

**Bad things happen- who should bear the risk? At small scale, outsourcing risk is acceptable because the cost of a handful of angry customers is relatively small. At large scale, a formal SLA will bound risk to the company offering the service.**



# Performance

VM contention

Network contention

API access contention

- Rate limits

Data Gravity

- Egress is expensive, Ingress is free

# Compliance

PCI

CCPA


GDPR

SLA


**Even if your cloud vendor claims compliance, can you enforce it?**

# Size and Scope of Service

Optimizing for base load



Optimizing for individual projects

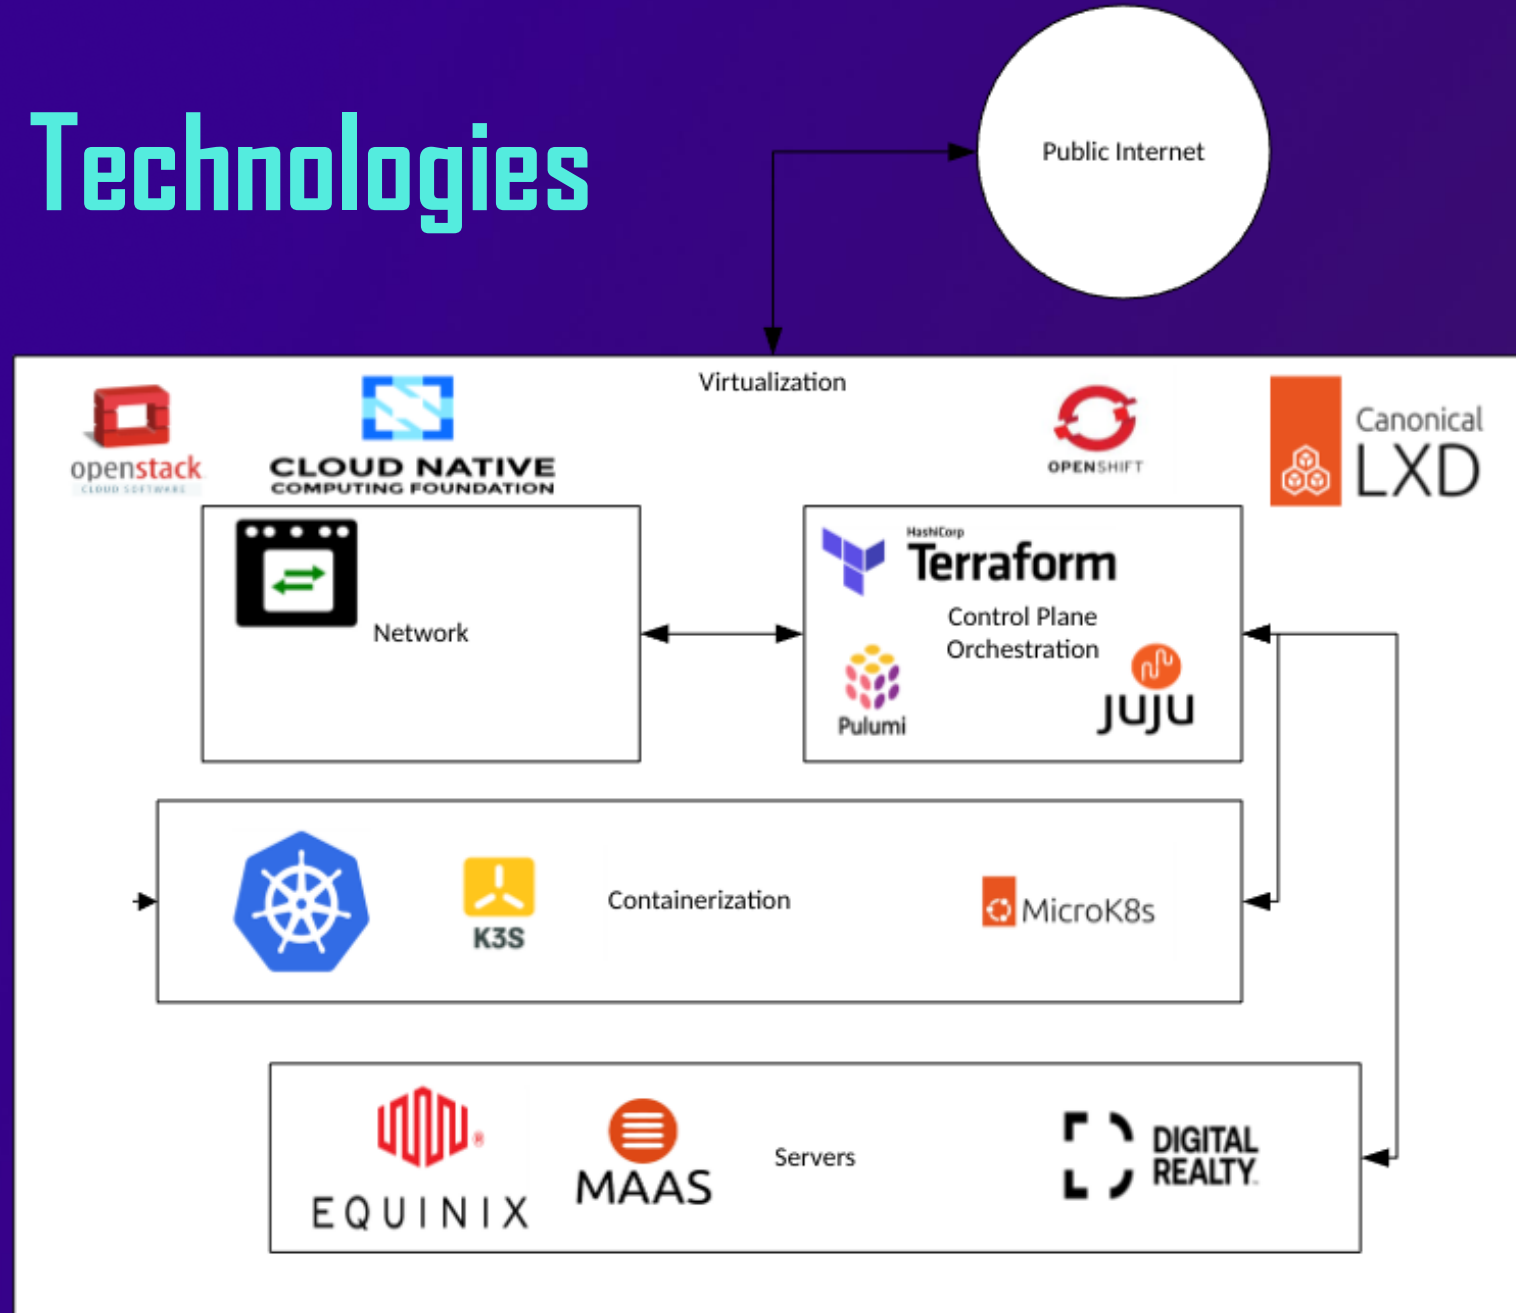


# Negative Externalities


- Cloud where it makes sense
  - New projects, dev work, scaling risk, business continuity
- Self hosting where it makes sense
  - Known large scale, large expected traffic, older projects where scope is known, batch processing
- Anything in-between
  - Shop specific custom knowledge

**Retionalization does not mean no more cloud, instead it means evolving into an intelligent Hybrid**

# Enabling Technologies



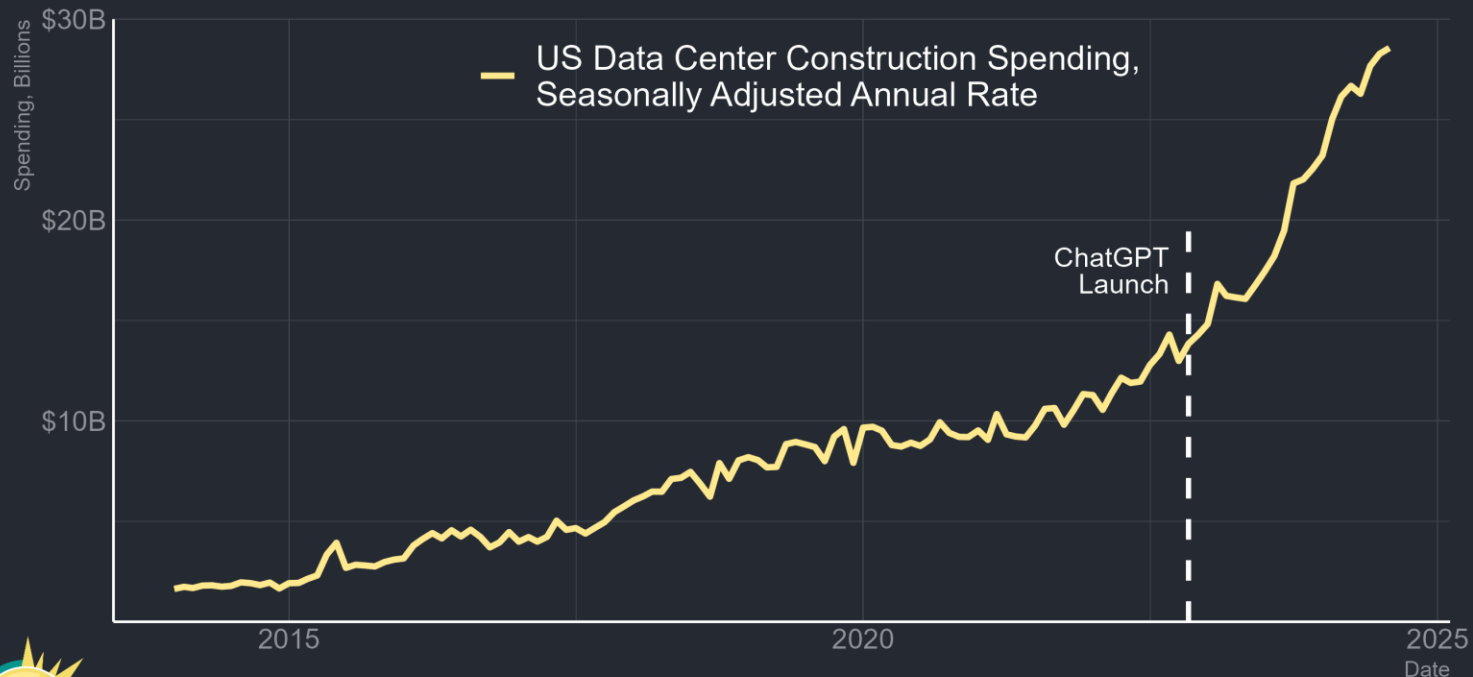
# The Latest Spin of the Wheel



# Wherever You Go, There You Are

## Data Center Construction at Record Highs

Data Center Construction Spending is Skyrocketing Amidst the AI Boom



Graph created by @JosephPolitano using Census data



# Build the Stack You Want

- There are multiple projects which replicate a professional cloud
- Pick-and-choose the parts which work for you, and ignore what doesn't
- Customize based on your shop's requirements
- Optimize to deliver an ideal production environment

**If you like your cloud, you can keep it**



# Vendors for Every Part of the Stack

- What you want to do for yourself
  - Control your own Cloud
  - Control your own network circuits and bandwidth
- What you want to outsource
  - Utilize a third party Metal-as-a-Service
    - Third party is responsible for physical hardware only

**Getting the risk-reward balance right**

# Questions?

✉ [hello@inertialabs.io](mailto:hello@inertialabs.io)

Nathan Moore  
Inertia Labs

TM