DIGITAL REALTY™

The Role of Data Centres in Powering AI, Compute, Security & Connectivity

The meeting place for companies, technologies and data

DIGITAL REALTY

DIGITAL REALTY™

1) Who we are
2) History of Cloud
3) AI Architecture
4) Digital Realty Solution
5) Key take aways

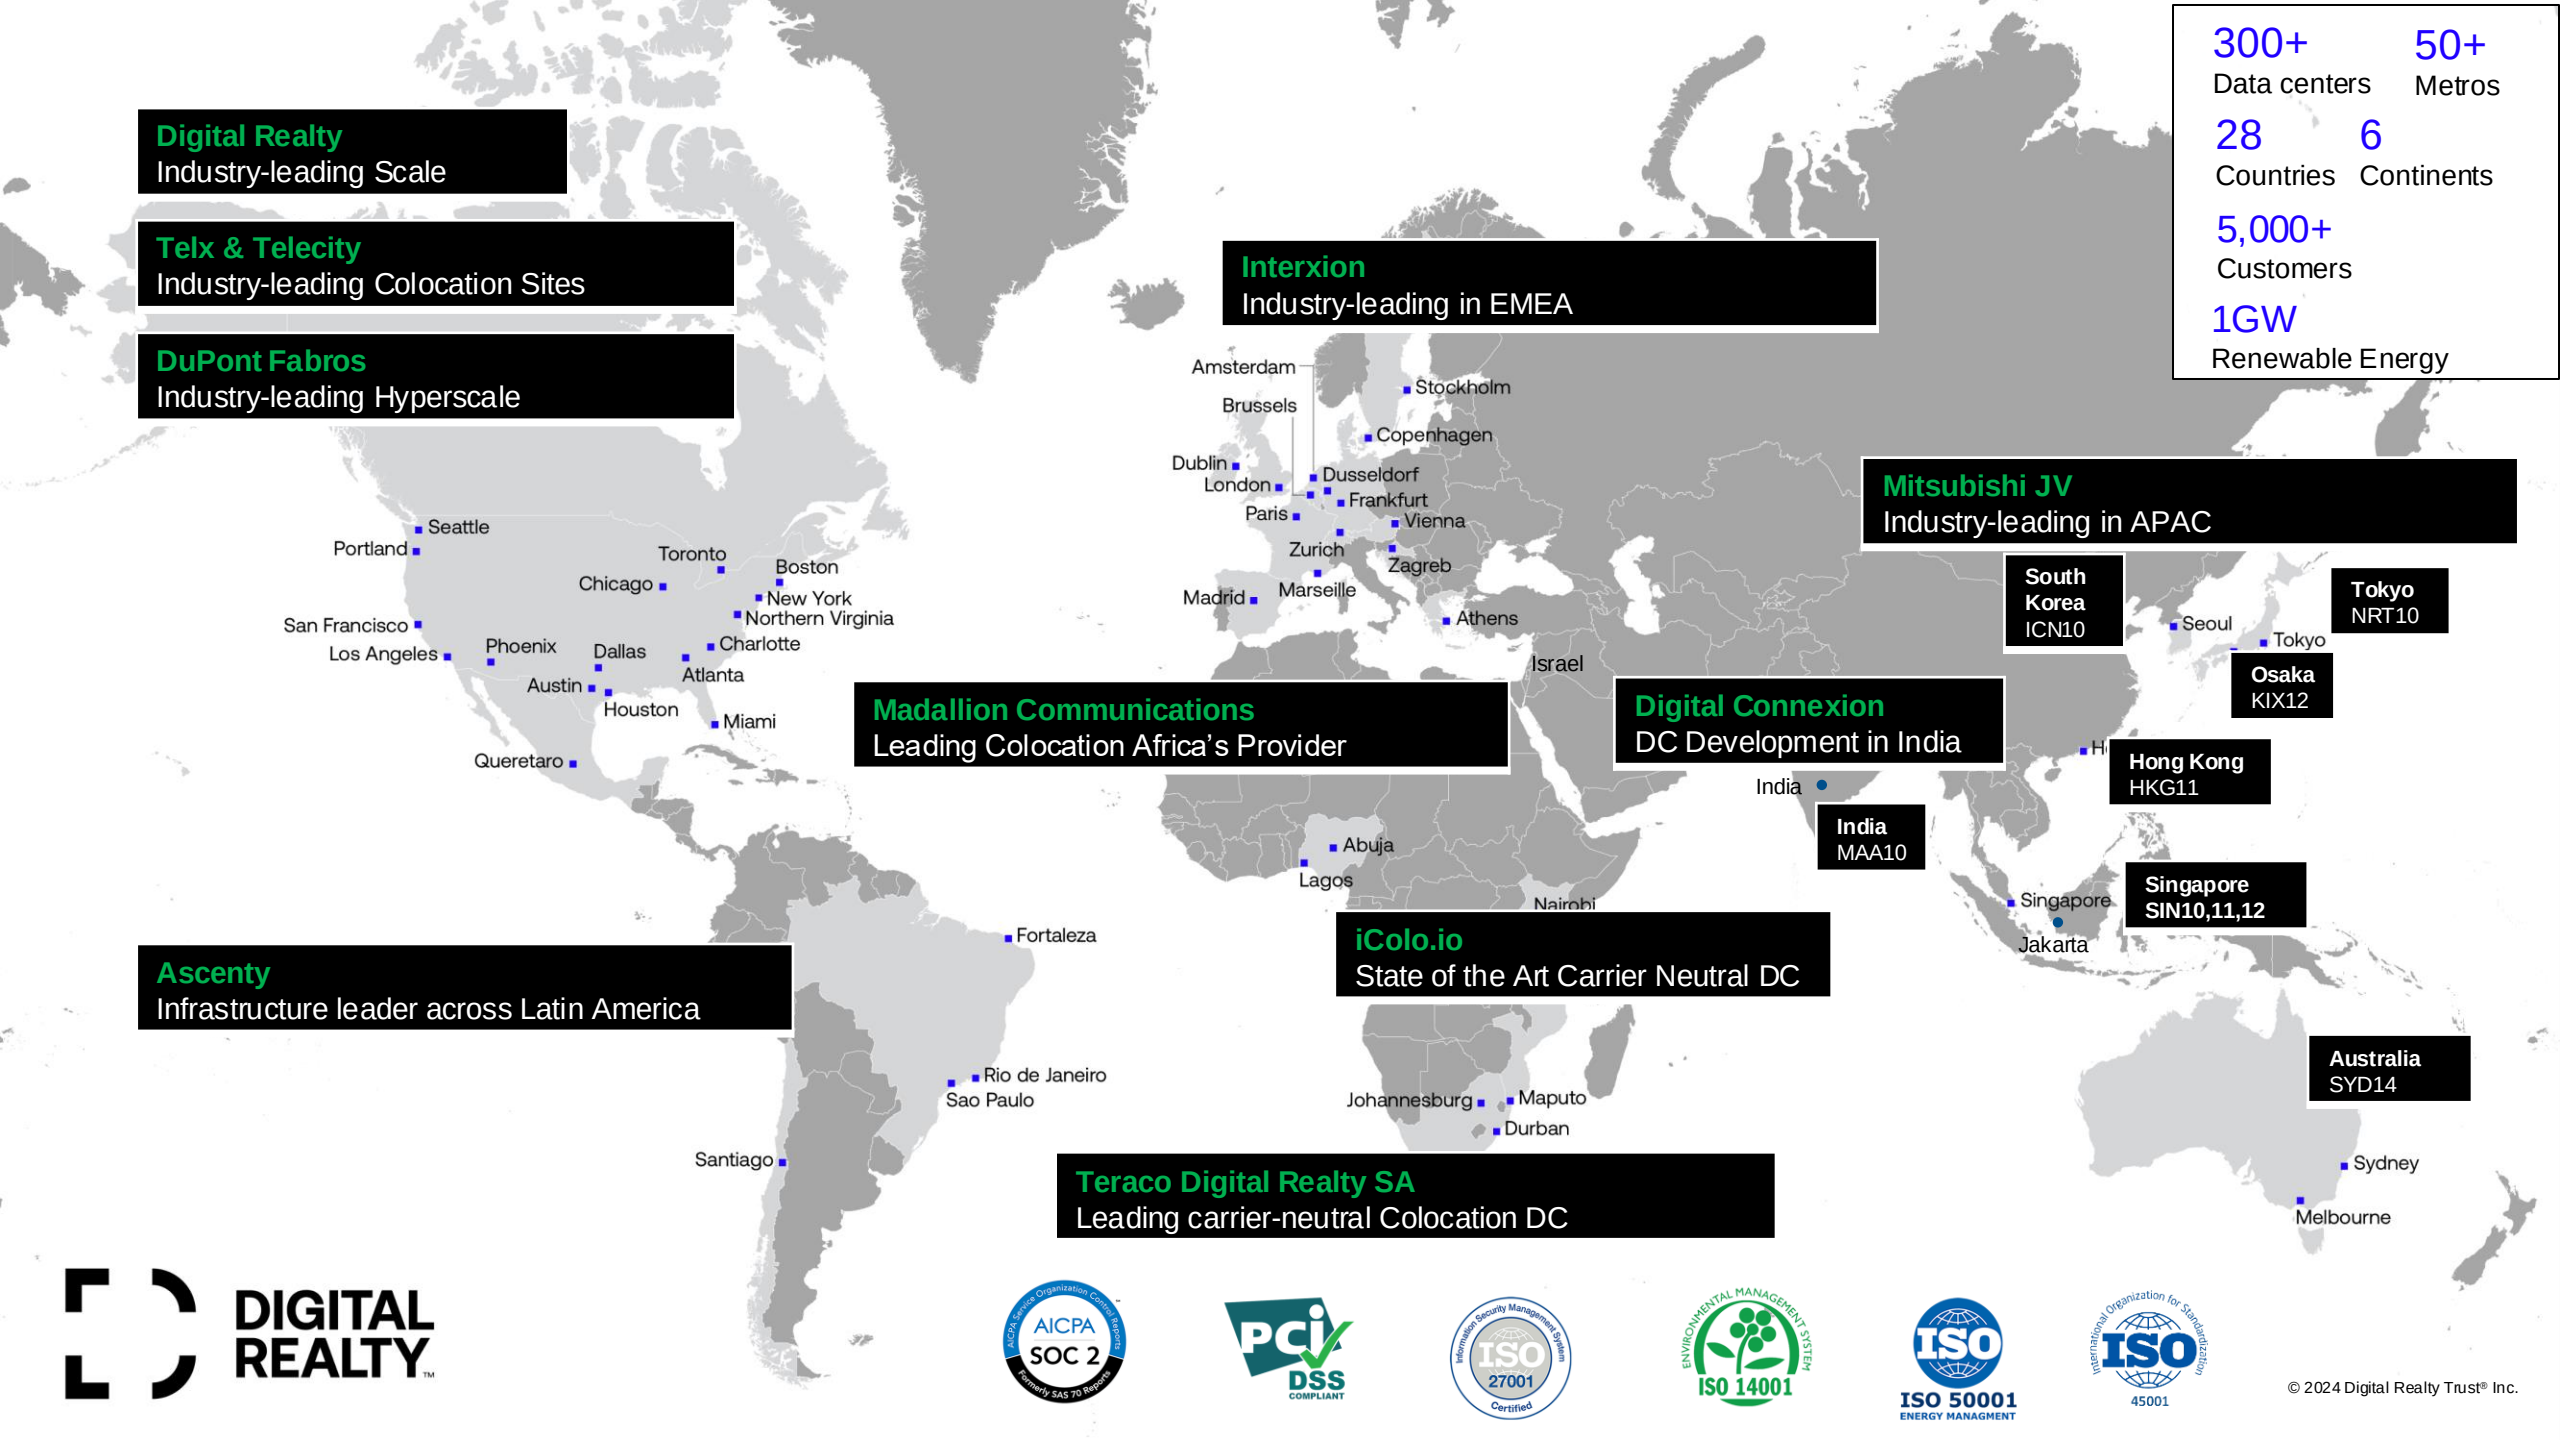

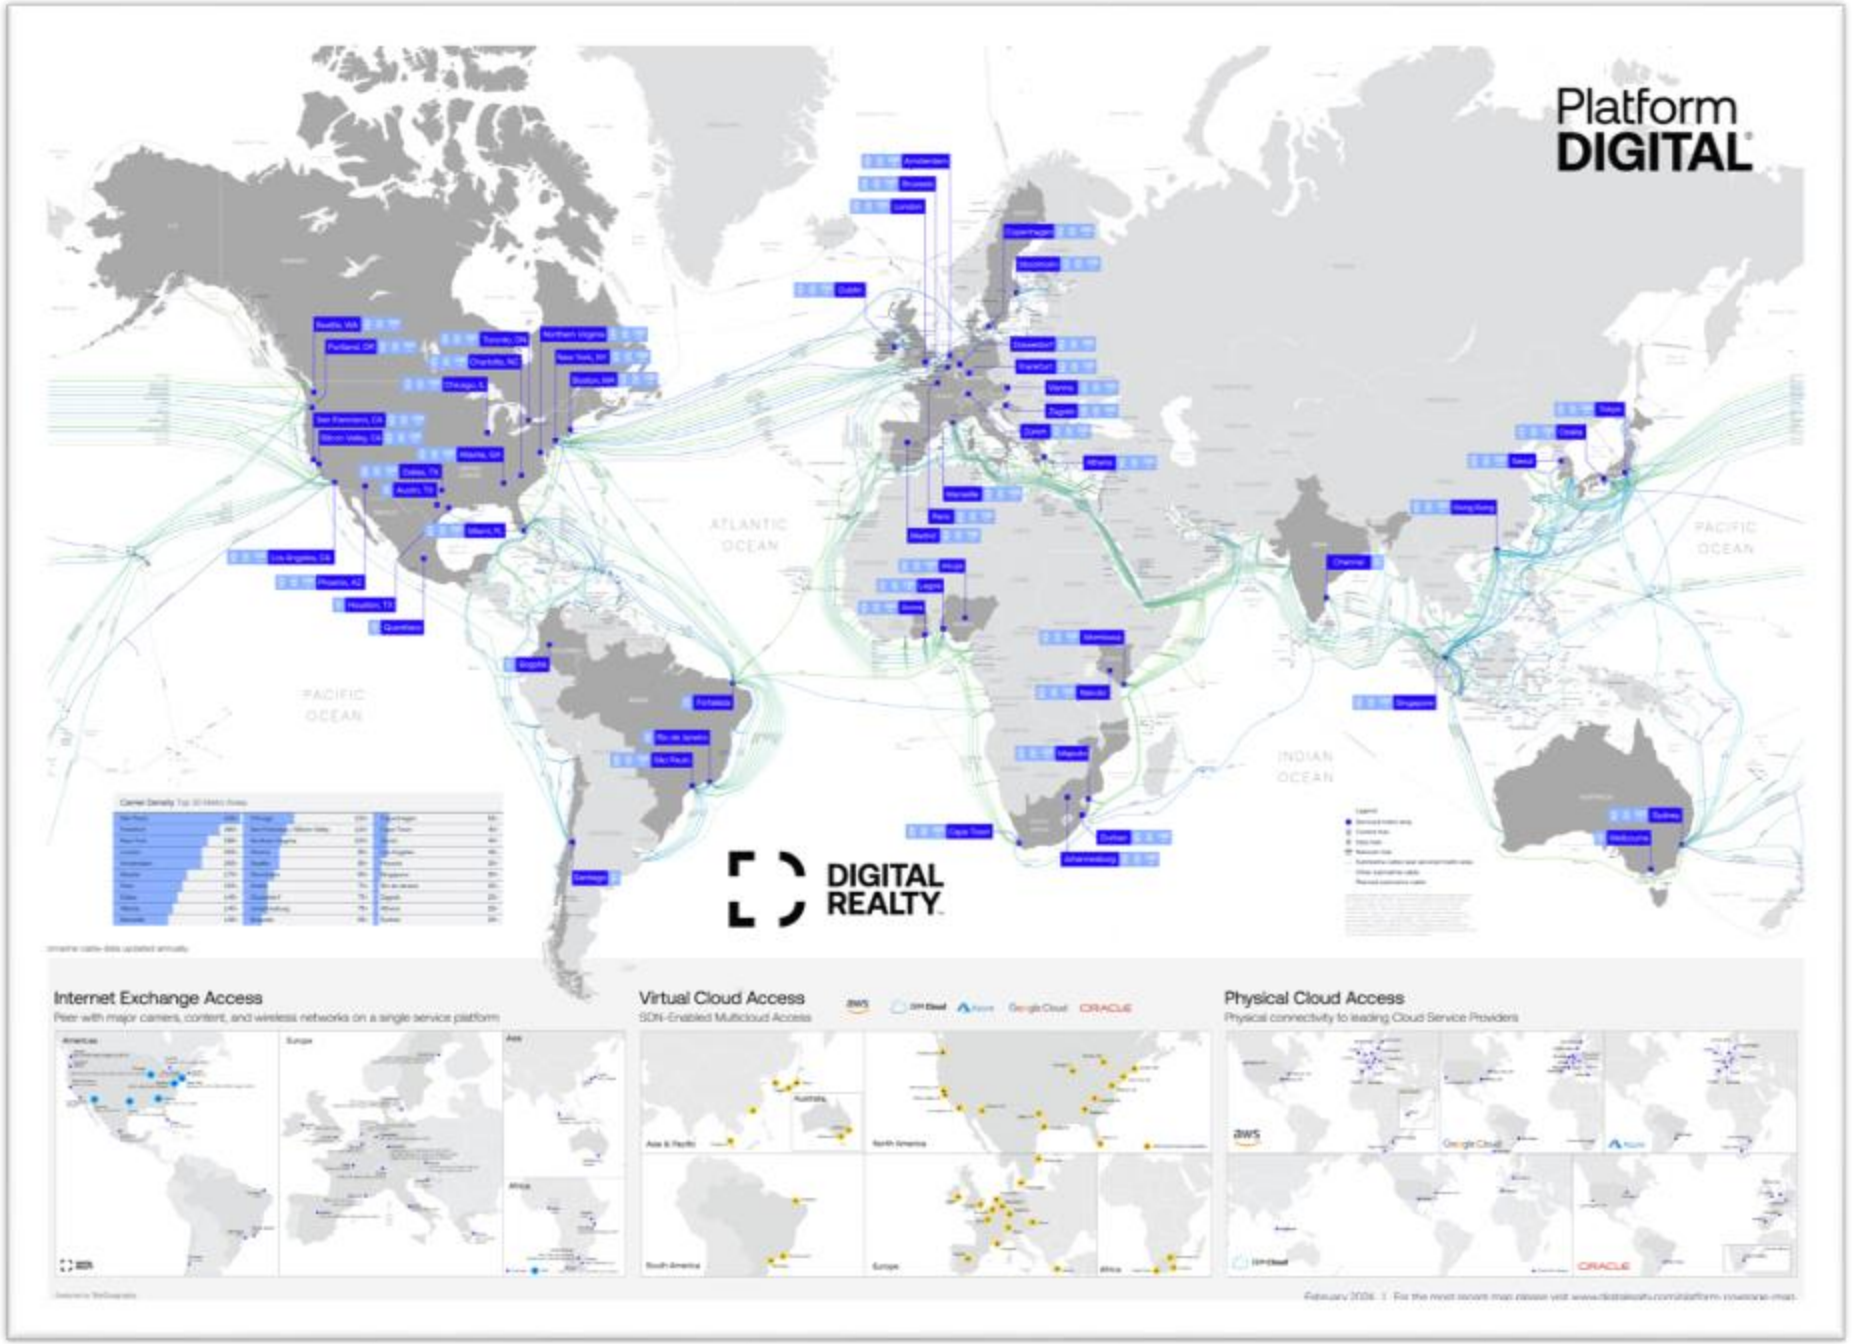

Building a Sustainable Future

- Digital Realty is the largest buyer of renewable energy among data center providers
- We cover 100% of our North American colocation business and European portfolios with renewable energy
- Digital Realty has certified more green buildings than any other data center provider.
- Digital Realty has more Energy Star-certified data centers than any other data center provider
- Digital Realty executed the data center industry's first green bond, a nearly half-billion dollar issuance earmarked exclusively for green projects.

68%

Global carbon reduction target

1GW

Renewable capacity

126

Data centers matched with 100% renewable electricity

12M

Square feet green building certifications

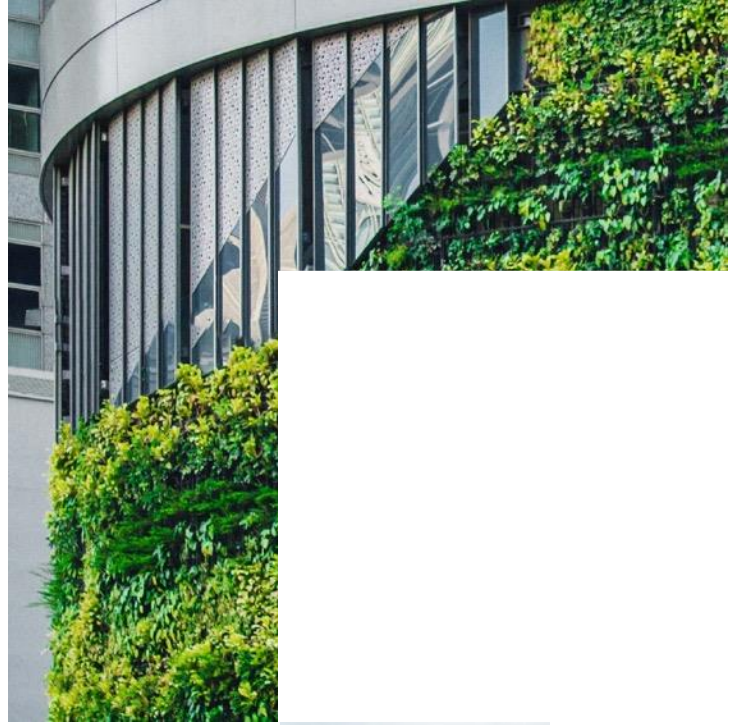

DIGITAL REALTY.

Continuing to Invest and Expand Across the Globe

January 2022

Digital Realty opens first carrier-neutral data center in **Seoul, South Korea.** Digital Seoul 1 (ICN10) will serve as a gateway to global expansion for enterprises in **Korea**.

June 2022

Digital Realty forms a joint venture with Mivne to develop and operate data centers in **Israel** to be operational in 2023**.**

August 2022

Digital Realty completes Acquisition of Teraco, a leading carrier-neutral data center and interconnection services provider in **South Africa**.

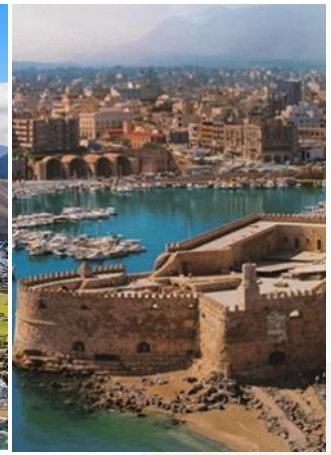

August 2022

Digital Realty announces plans to build its first data center in **Crete** to expand the presence in **Greece**, with a strong focus on subsea cables.

June 2022

Digital Realty expands Mediterranean presence with development of new colocation and connectivity hub in **Barcelona**.

February 2023

Digital Realty opens first data center in **Maputo**, **Mozambique** expanding the company's footprint in the eastern coast of **Africa**.

July 2023

Digital Realty forms a three-way joint venture (JV) with Brookfield Infrastructure and Reliance Industries Ltd. in **India**. Operating as Digital Connexion: A Brookfield, Jio and Digital Realty Company.

September 2023

Digital Realty expands Mediterranean presence with development of new colocation and connectivity hub in **Rome**.

DIGITAL REALTY.

A Brief History of Cloud

2010

Cloud Computing

2015

Cloud Computing

Digital Transformation

2020

Cloud Computing

Digital Transformation

Artificial Intelligence

Q1 2023

A Brief History

2010

Cloud Computing

2015

Cloud Computing

Digital Transformation

2020

Cloud Computing

Digital Transformation

Artificial Intelligence

Q1 2023

Feb '23

Amazon launches multimodal CoT-model

Incorporates "chain-of-thought prompting and outperforms GPT-3.5 on several benchmarks

Feb '23

Microsoft launches Kosmos-1

Can respond to both image and audio prompts in addition to natural language

Mar '23

OpenAI launches GPT-4

Significantly improves accuracy - 40% improvements over GPT-3.5

Mar '23

Microsoft integrates O365

Brings GPT-4 into its software suite, enabling massive productivity gains

Mar '23

Bloomberg launches BloomberGPT

A trained model using financial data to suppor tasks in the financial sector

Feb '23

Meta launches LLaMA

Recognized as a smaller, more efficient model that still performas well in benchmarks

Mar '23

Salesforce launches Einstein GPT

The first GenAI technology for Customer Relationship Management (CRM)

Mar '23

Anthropic launches Claude

Uses "Constitutional AI" to reduce the likelihood of harmful outputs

Mar '23

Google launches Bard

Based on the LaMDA family of LLMs

Apr '23

Amazon launches Bedrock

The first managed service to make models available via API

What are the implications of Hybrid IT?

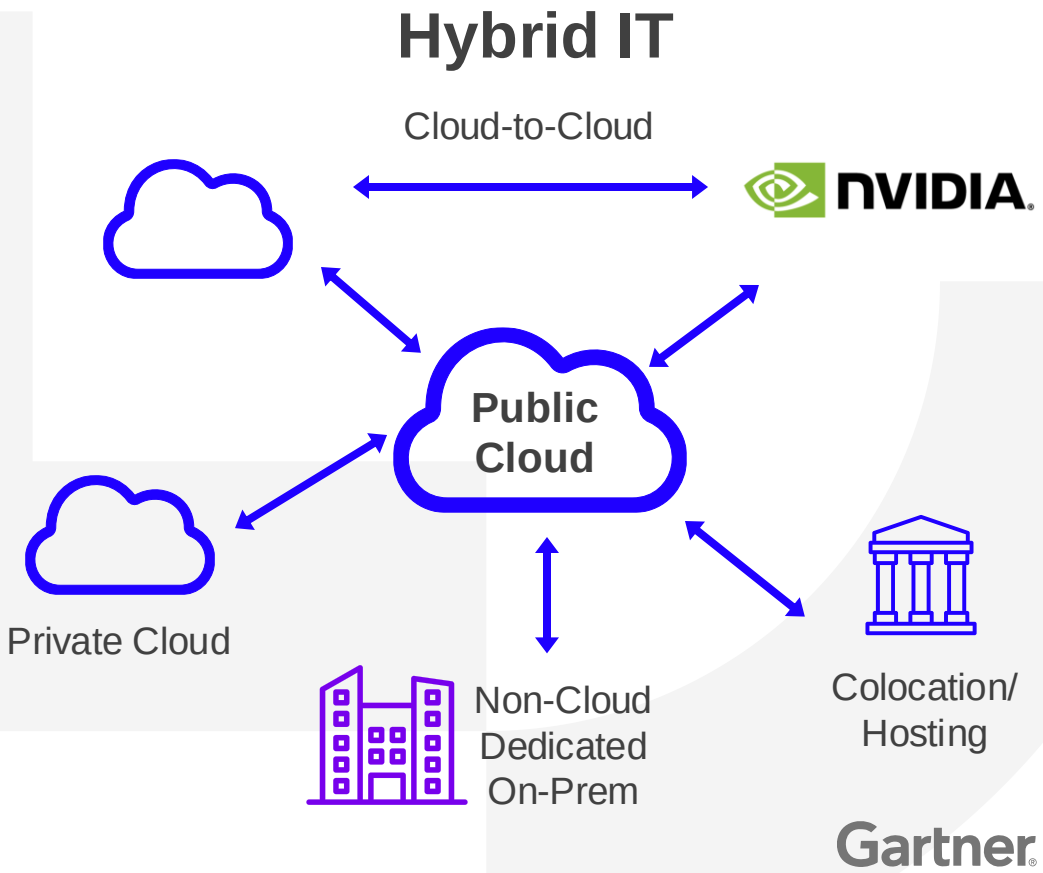

Hybrid IT is the evolution of Enterprise IT

- Clouds
- Locations
- Applications
- Data & Data Exchange
- Distributed Workflows
- Security Controls
- Complexity

Managing and orchestrating infrastructure and application workflow is increasingly complex and highly inconsistent.

AI Enablement: Applying Cloud Learnings to Accelerate AI

Cloud Architecture Evolution: We've Seen This Before

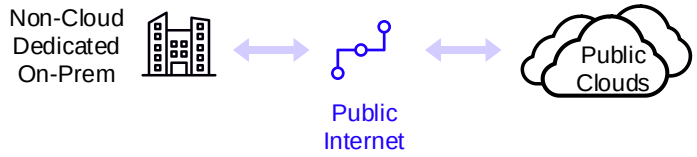

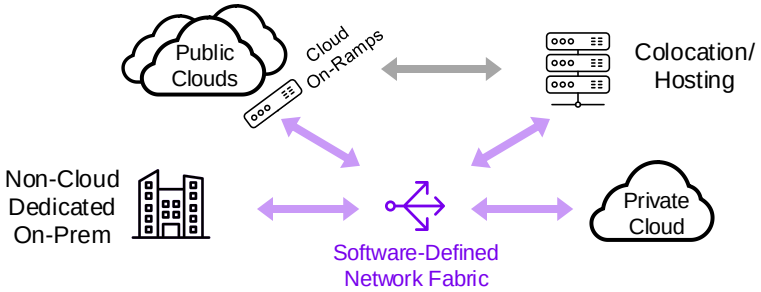

Early Cloud Challenges

- Limited, unsecure access options
- Initially built for single cloud
- Quickly evolving regulations
- Complex to deploy multi-cloud

Evolved Cloud Solutions

- Distributed, performant, secure access
- Mature service portfolio, including multi-cloud support
- Distributed AZs meet regulatory requirements
- Software-defined network fabrics remove complexity

AI Architecture Evolution: We See Where This Is Going

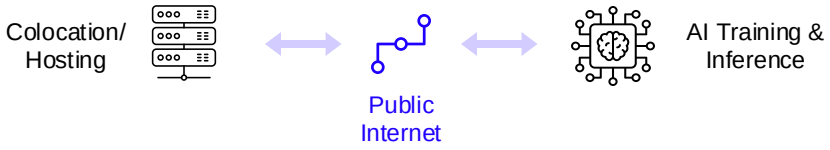

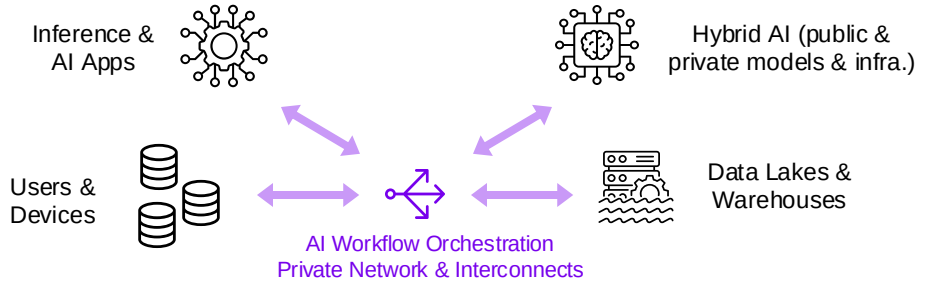

Early AI Challenges

- Limited, unsecure access options
- Initially built for single AI
- Quickly evolving regulations
- Immature use cases limit GPU utilization

Future AI Solutions

- Private network & interconnects enable performant, distributed access
- Workflow optimization increases infrastructure utilization
- Hybrid AI meets latency, security, & regulatory requirements
- Distributed inferencing unlocks enterprise data lakes & warehouses

Types of AI Architectures

Data Lake

Capacity Block Size	5-15 MW
Rack Density	Low to medium (5 – 15 kW)
Air Cooled	

Training

Capacity Block Size	5-100 MW
Rack Density	Very high (35 – 100 kW)
Liquid Cooling	

Inference

Capacity Block Size	1-5 MW
Rack Density	Medium to high (15 – 25 kW)
Location / Latency Sensitivity	

11

Generative AI Demand: Today

Primarily Demand for Training Deployments

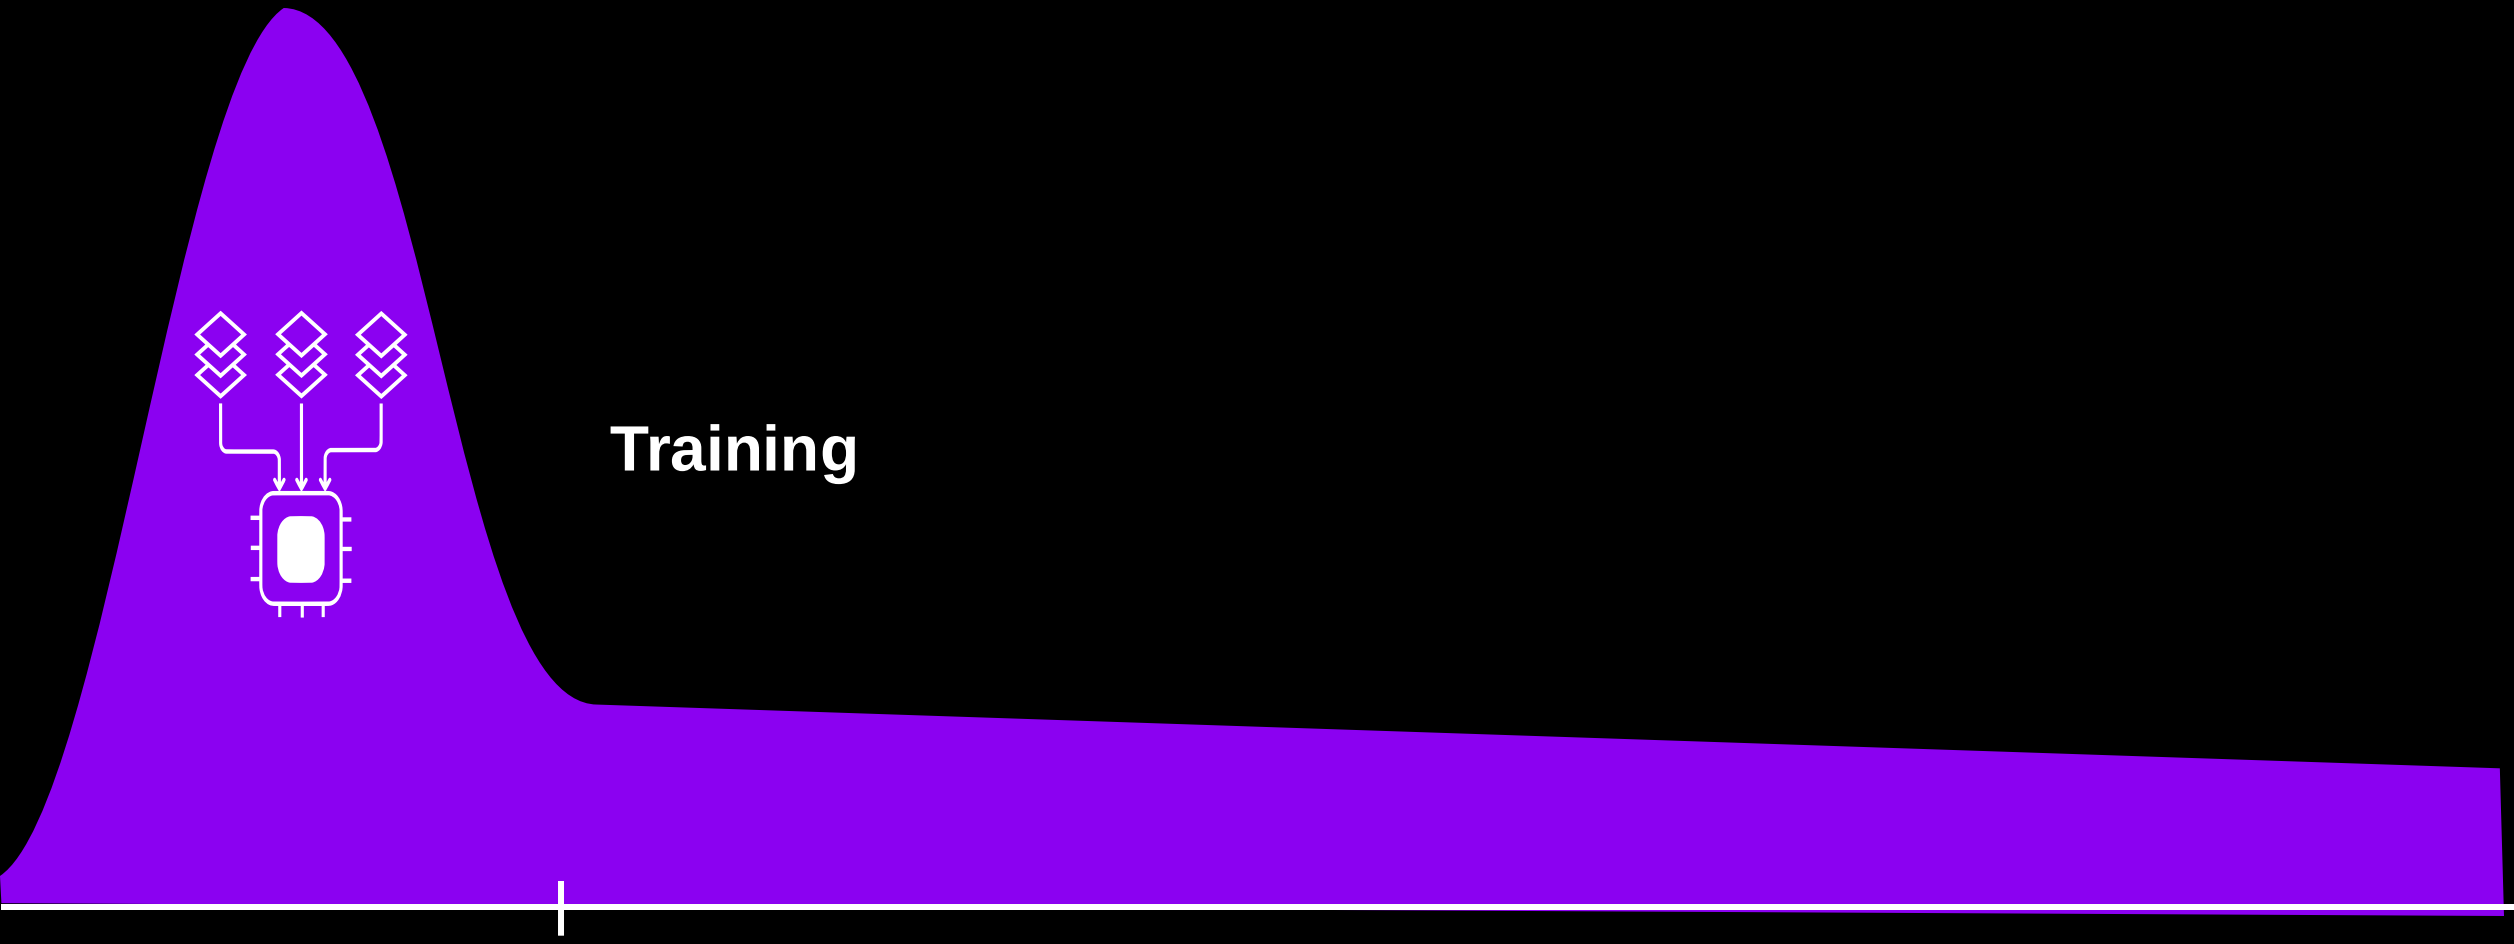

Generative AI Demand: 1-2 Years

Transition from Training to Inference as Products Mature

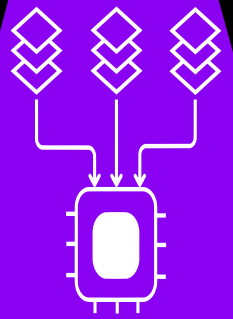

Training

Inference

Models become good enough, limiting incremental training demand. AI products still maturing with limited revenue, limiting inferencing demand

Generative AI Demand: 3-5 Years and Beyond

Inference is Long Tail of Enterprise AI Demand

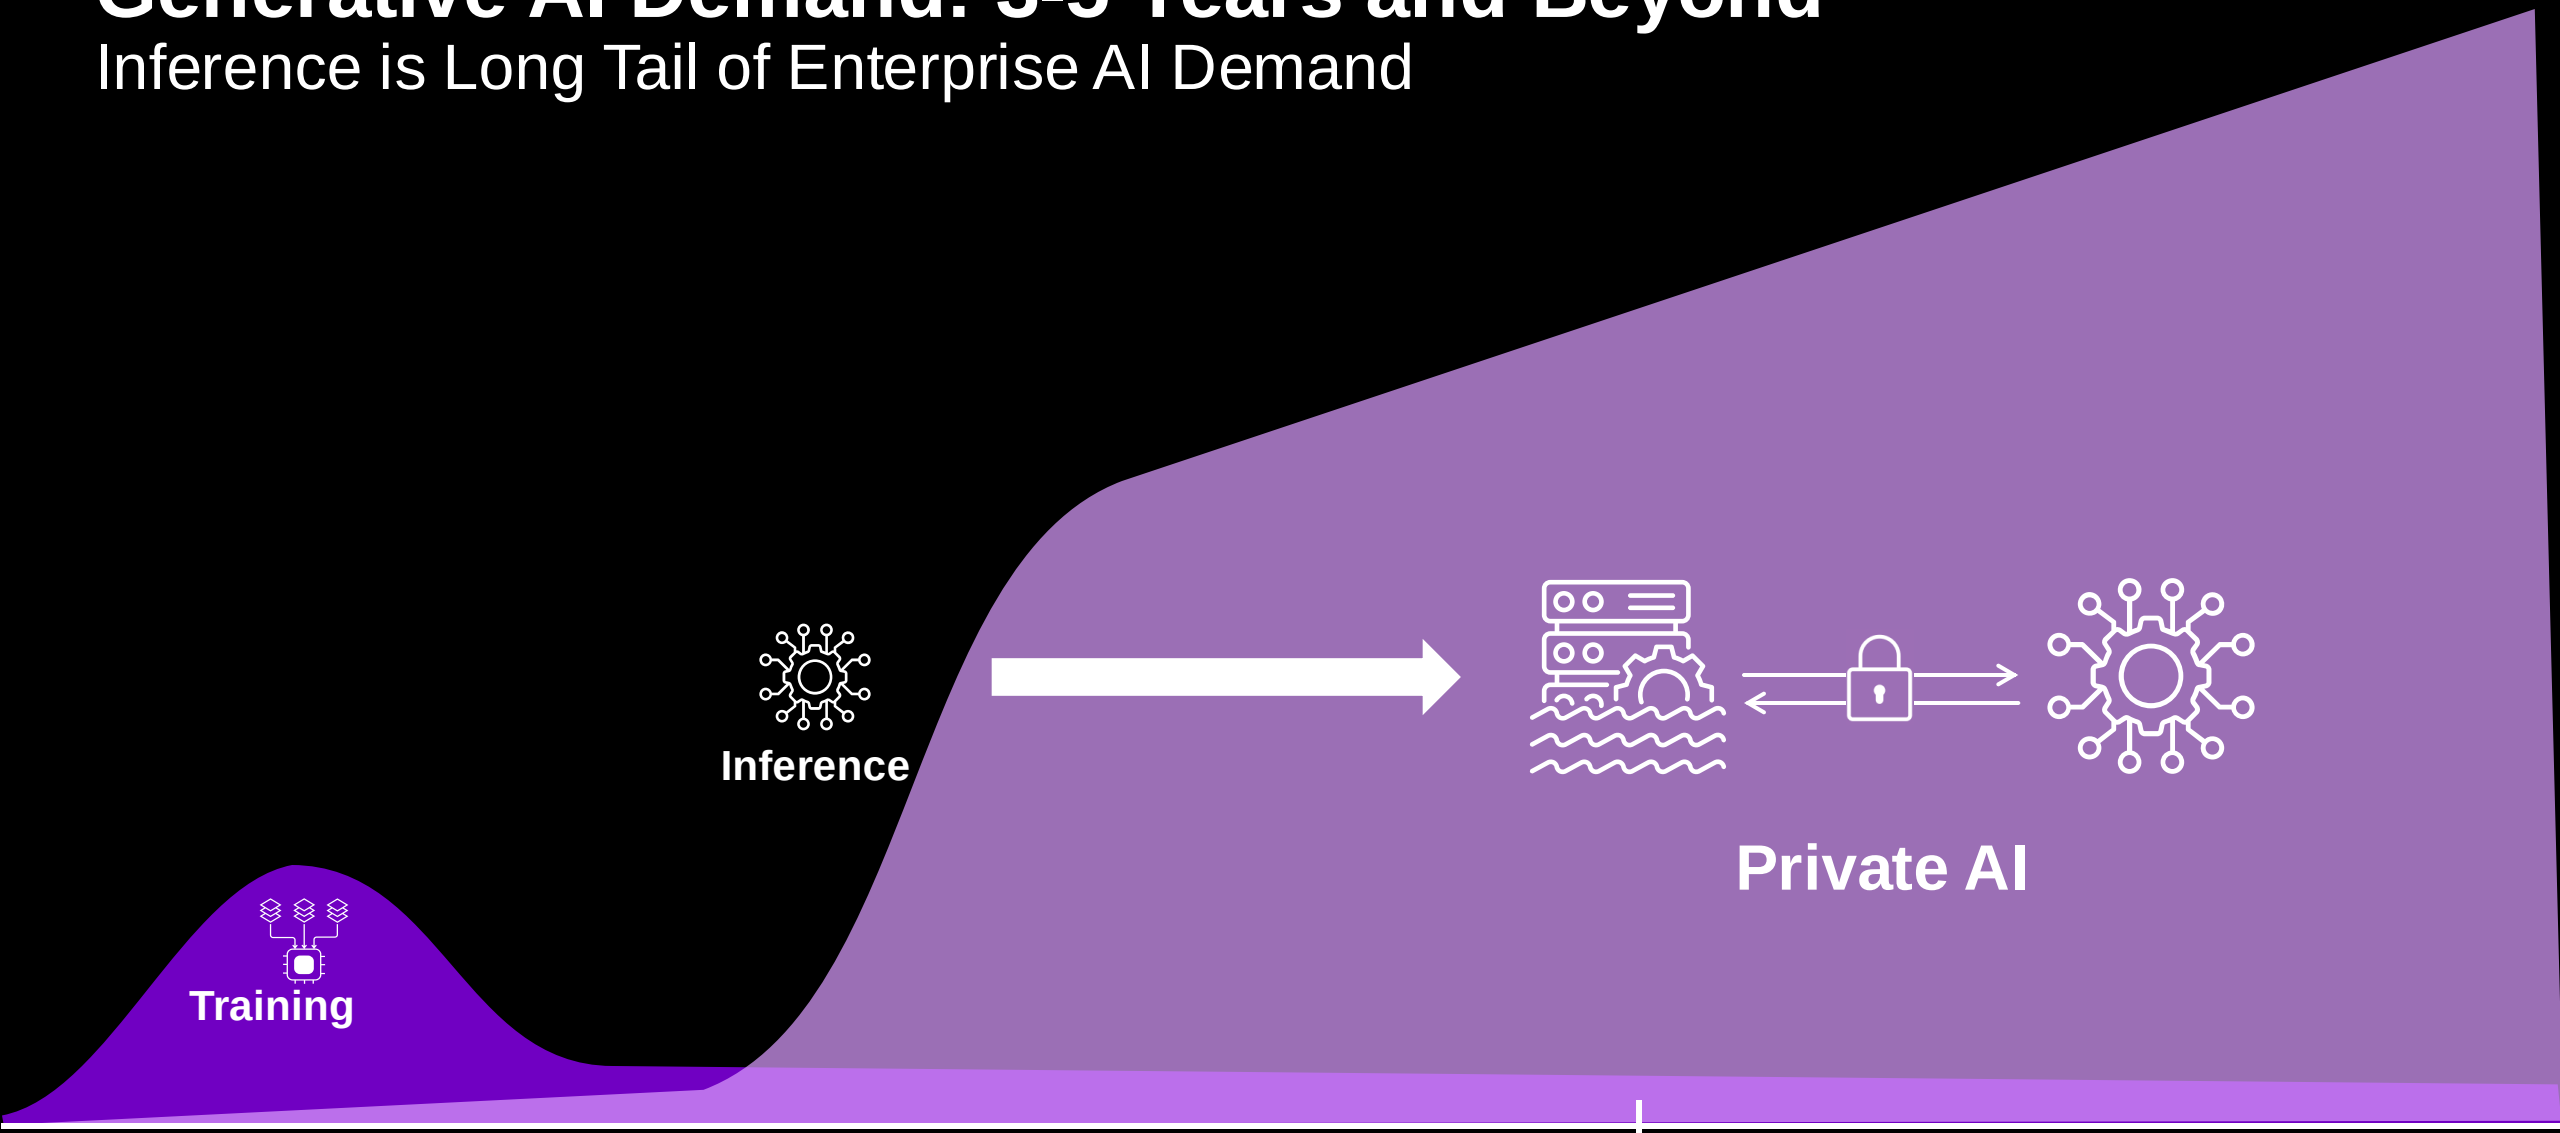

2nd and 3rd generation AI products reduce need for large-scale training. G2KEs deploy distributed inferencing to serve global users

AI Training / Inference – Hybrid Cloud

AI Training / Inference Cloud Hosted, Data-On-Prem

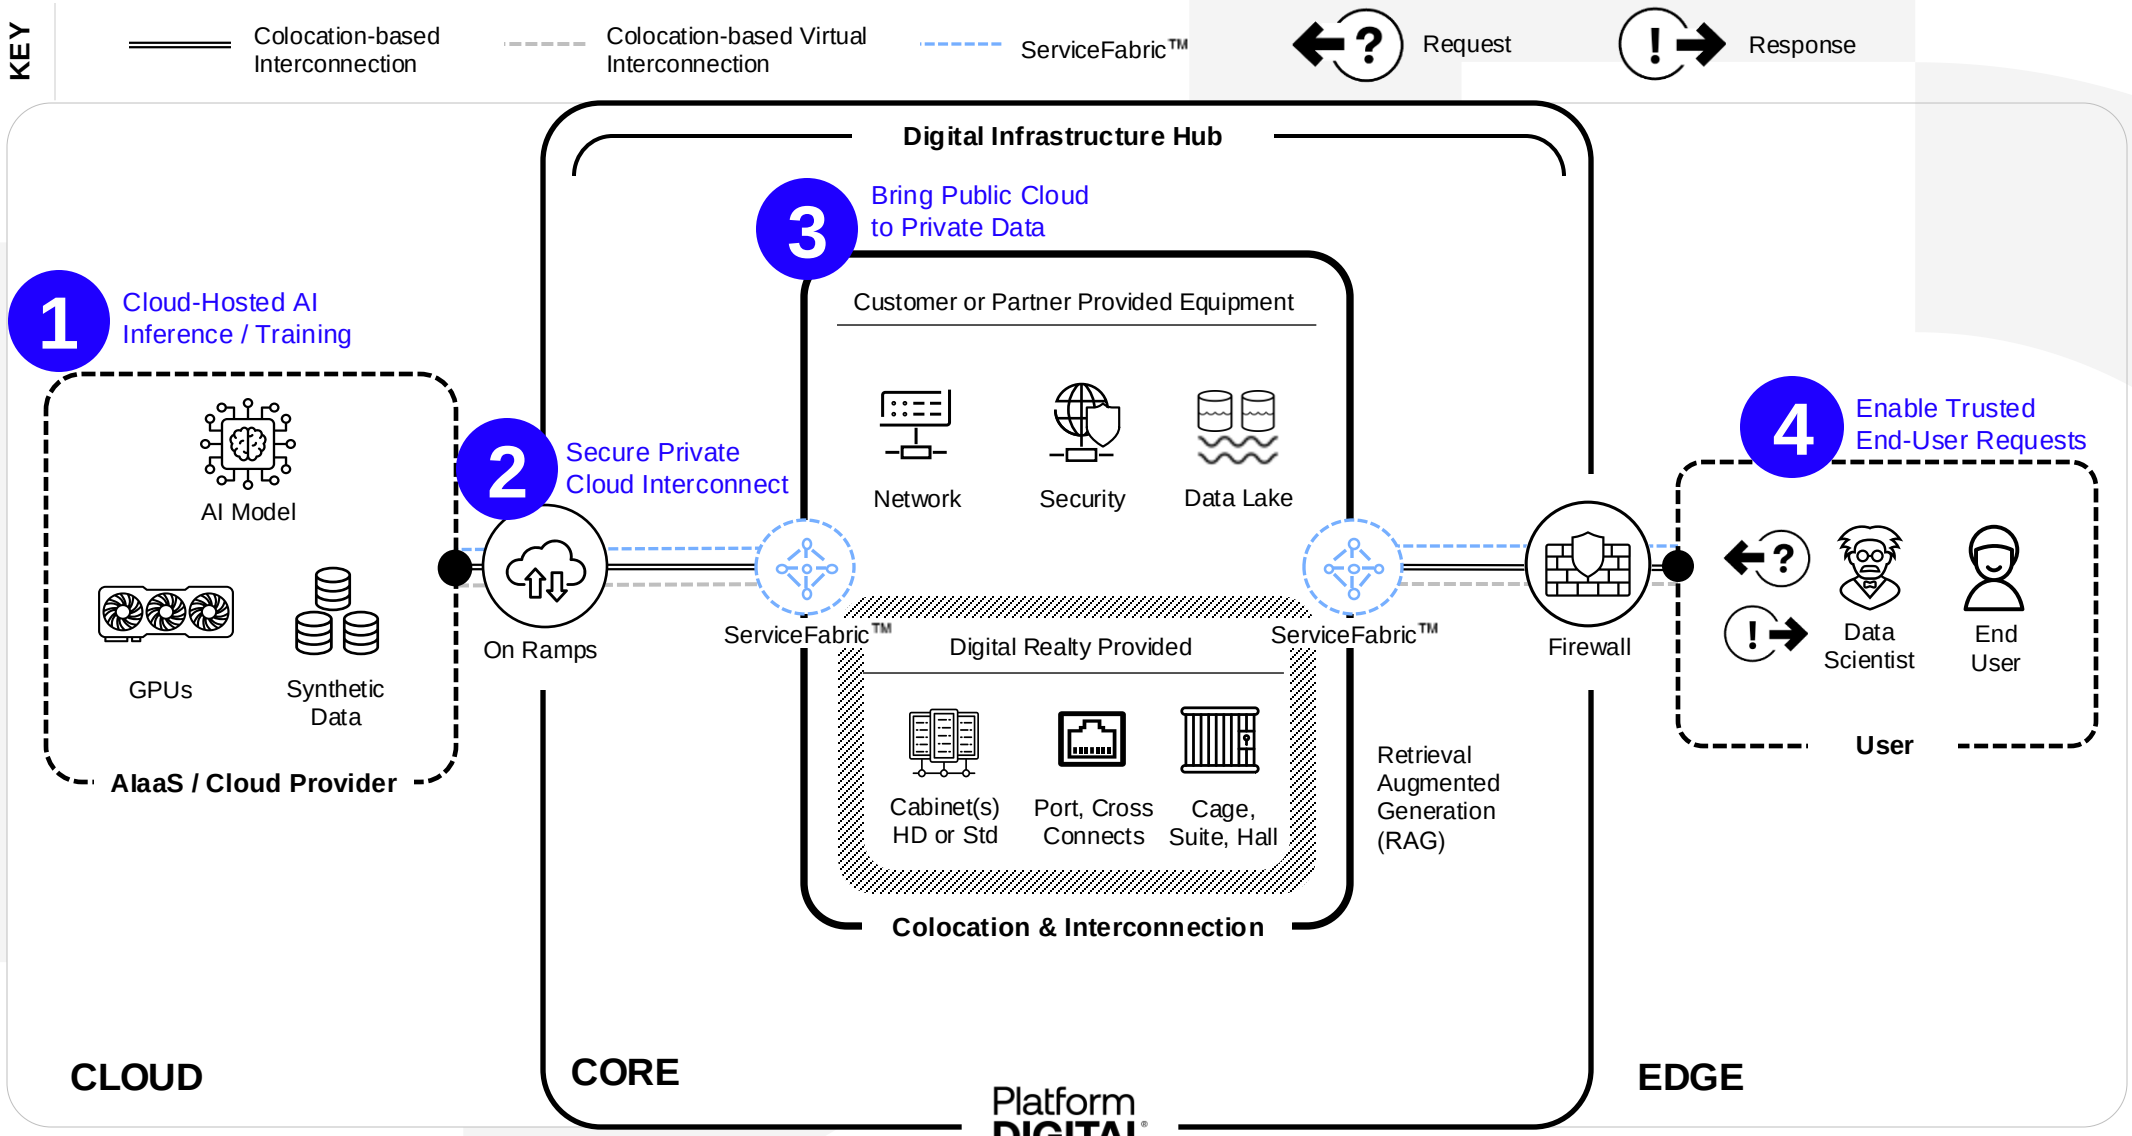

AI Training / Inference

Think AI?
Most people think of HPC

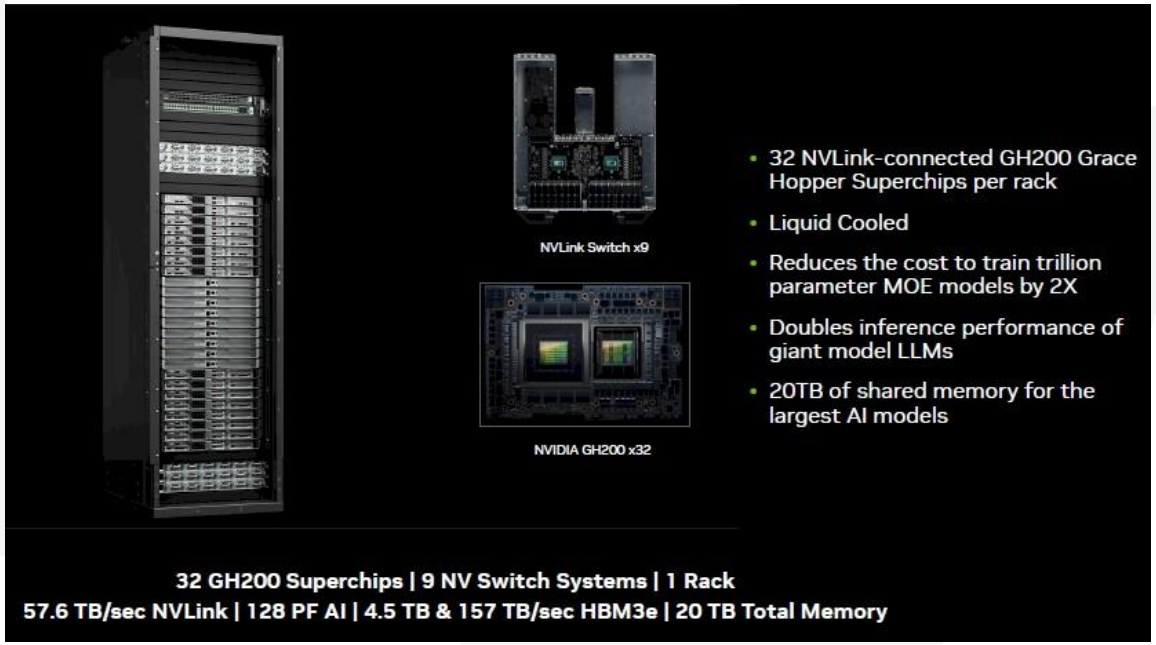

Actually, it can be done with this

Intel® Xeon® W5-3435X (up to 4.7 GHz with Intel® Turbo Boost Technology, 45 MB L3 cache, 16 cores, 32 threads)
Starting with 1125 W to 2250 W can be stack with 4x NVIDIA RTXTM A6000

Studio AI pre-build software

DIGITAL REALTY

PDx®: Data Loss Protection

Data Protection by Customer, Data from Public Clouds & Customer Facilities

PDx® Use Case
Logical View

Needs

- Provide Zero Trust Environment to End-Users
- Secure private connectivity to trusted data
- Performant access to data anywhere
- Scalable footprint to support growth

Challenges

- Cost-effective interconnected Metro sites
- Unpredictable scaling requirements
- Maintaining security and compliance of data
- End-to-end control of ingress/egress access points

Actions

- Establish direct connections to public cloud
- Connect to your private infrastructure
- Deploy your private infrastructure in colocation
- Operationalize Digital Infrastructure Hub

Benefits

- Establish Zero Trust Environment
- Cloud-adjacent colocation reduces latency
- Enables direct secure access to public cloud
- Security and compliance for data

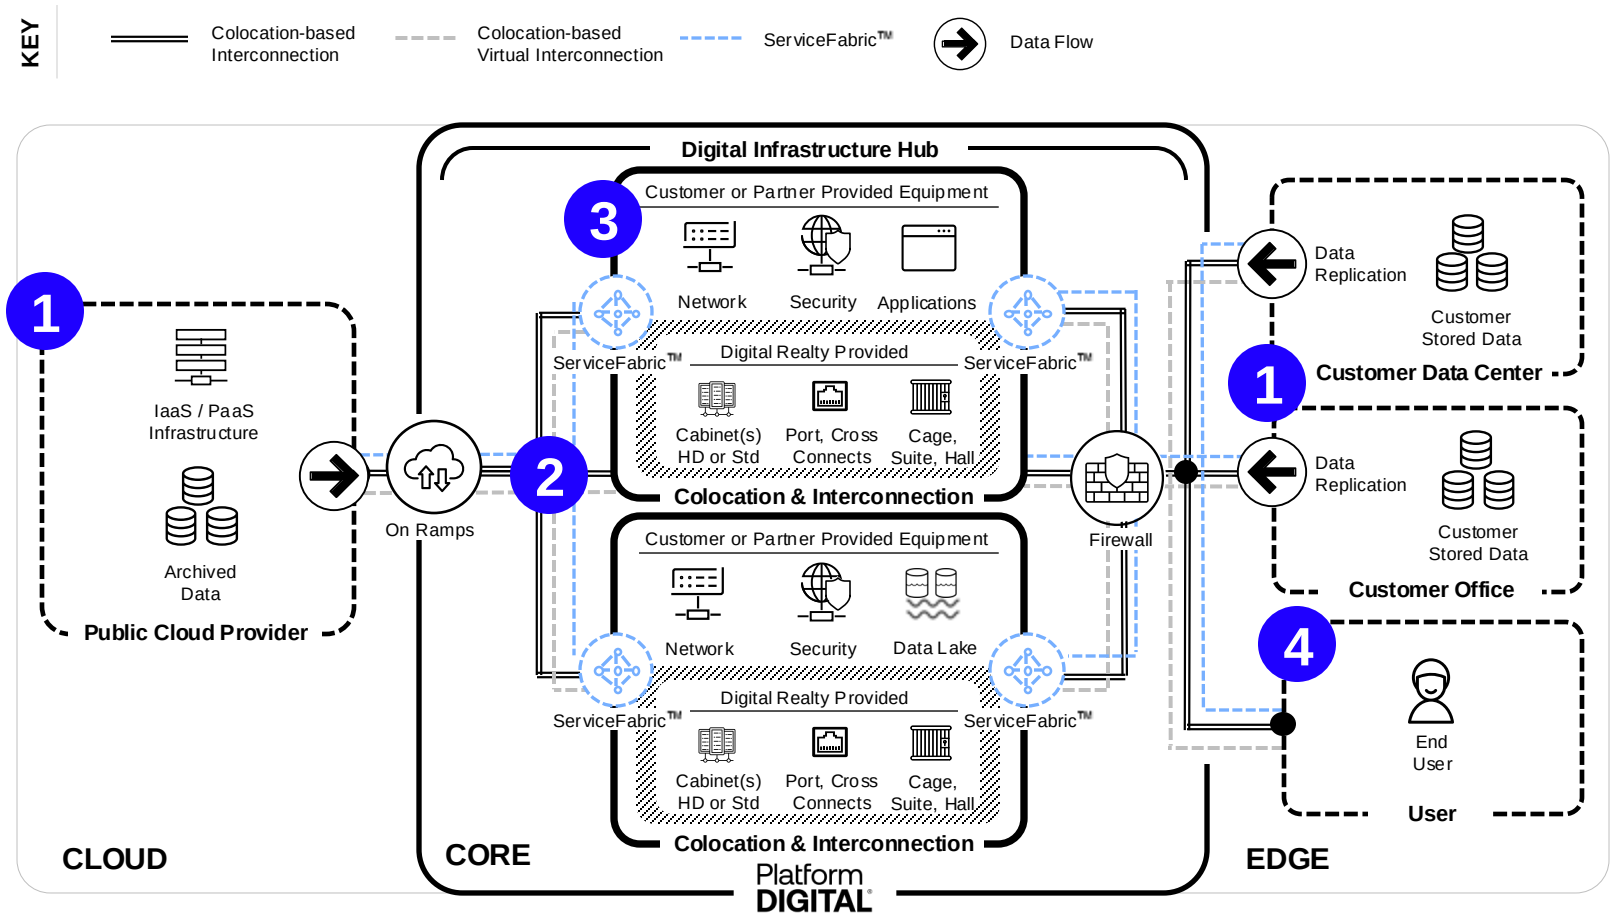

1. Bring Data to Data Exchange

2. Secure Private Cloud Interconnect

3. Quarantine Bad Traffic, Load Balance Clean Traffic

4. Zero Trust Environment for End-Users

NOTES

1. Platform DIGITAL®, the company's global data center platform, provides customers with a secure data meeting place and a proven Pervasive Datacenter Architecture (PDx®) solution methodology for powering innovation and efficiently managing Data Gravity challenges.
2. CLOUD - Massive compute and storage facilities provided by large hyperscale cloud service providers (CSPs) that provide IaaS, PaaS, and SaaS.
3. CORE - Digital Realty's Platform DIGITAL® Colocation featuring Digital Infrastructure Hub.
4. EDGE – Remote on-premises locations for Enterprises where data may be collected, processed, and managed.
5. IaaS / PaaS – Infrastructure-as-a-Service / Platform-as-a-Service
6. SDN – Software-Defined Network
7. SD-WAN – Software-Defined Wide Area Network
8. VPN – Virtual Private Network
9. HPC – High-Performance Compute
10. CPE – Customer Provided Equipment
11. PPE – Partner Provided Equipment
12. GPU – Graphics Processing Unit
13. HD or Std – High-Density or Standard Cabinet
14. BI – Business Intelligence Analytics Tools

HYBRID WORK	**H3**	
Enterprise End User		
Data Loss Protection		A **PDx®** Solution

DIGITAL REALTY™

London Metro Connect

6
Interconnected Campuses

13
Data Centers

75%
of Fortune 500

195+
Carriers

Alibaba Cloud · aws · Microsoft Azure · Google Cloud · IBM Cloud · ORACLE CLOUD · salesforce

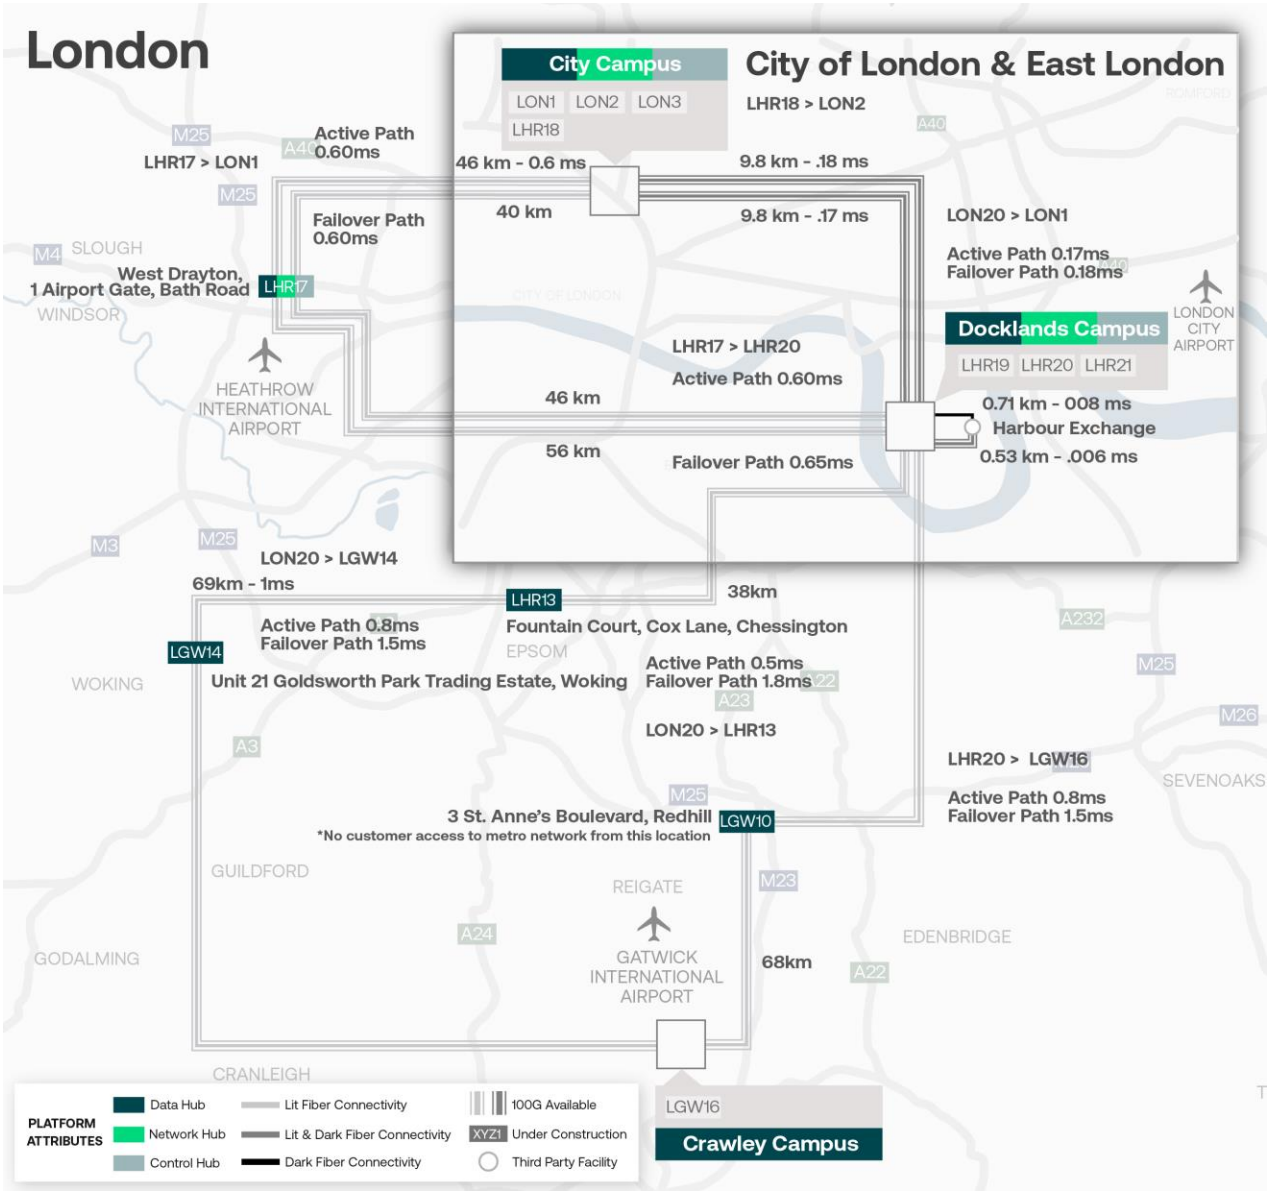

PlatformDIGITAL®

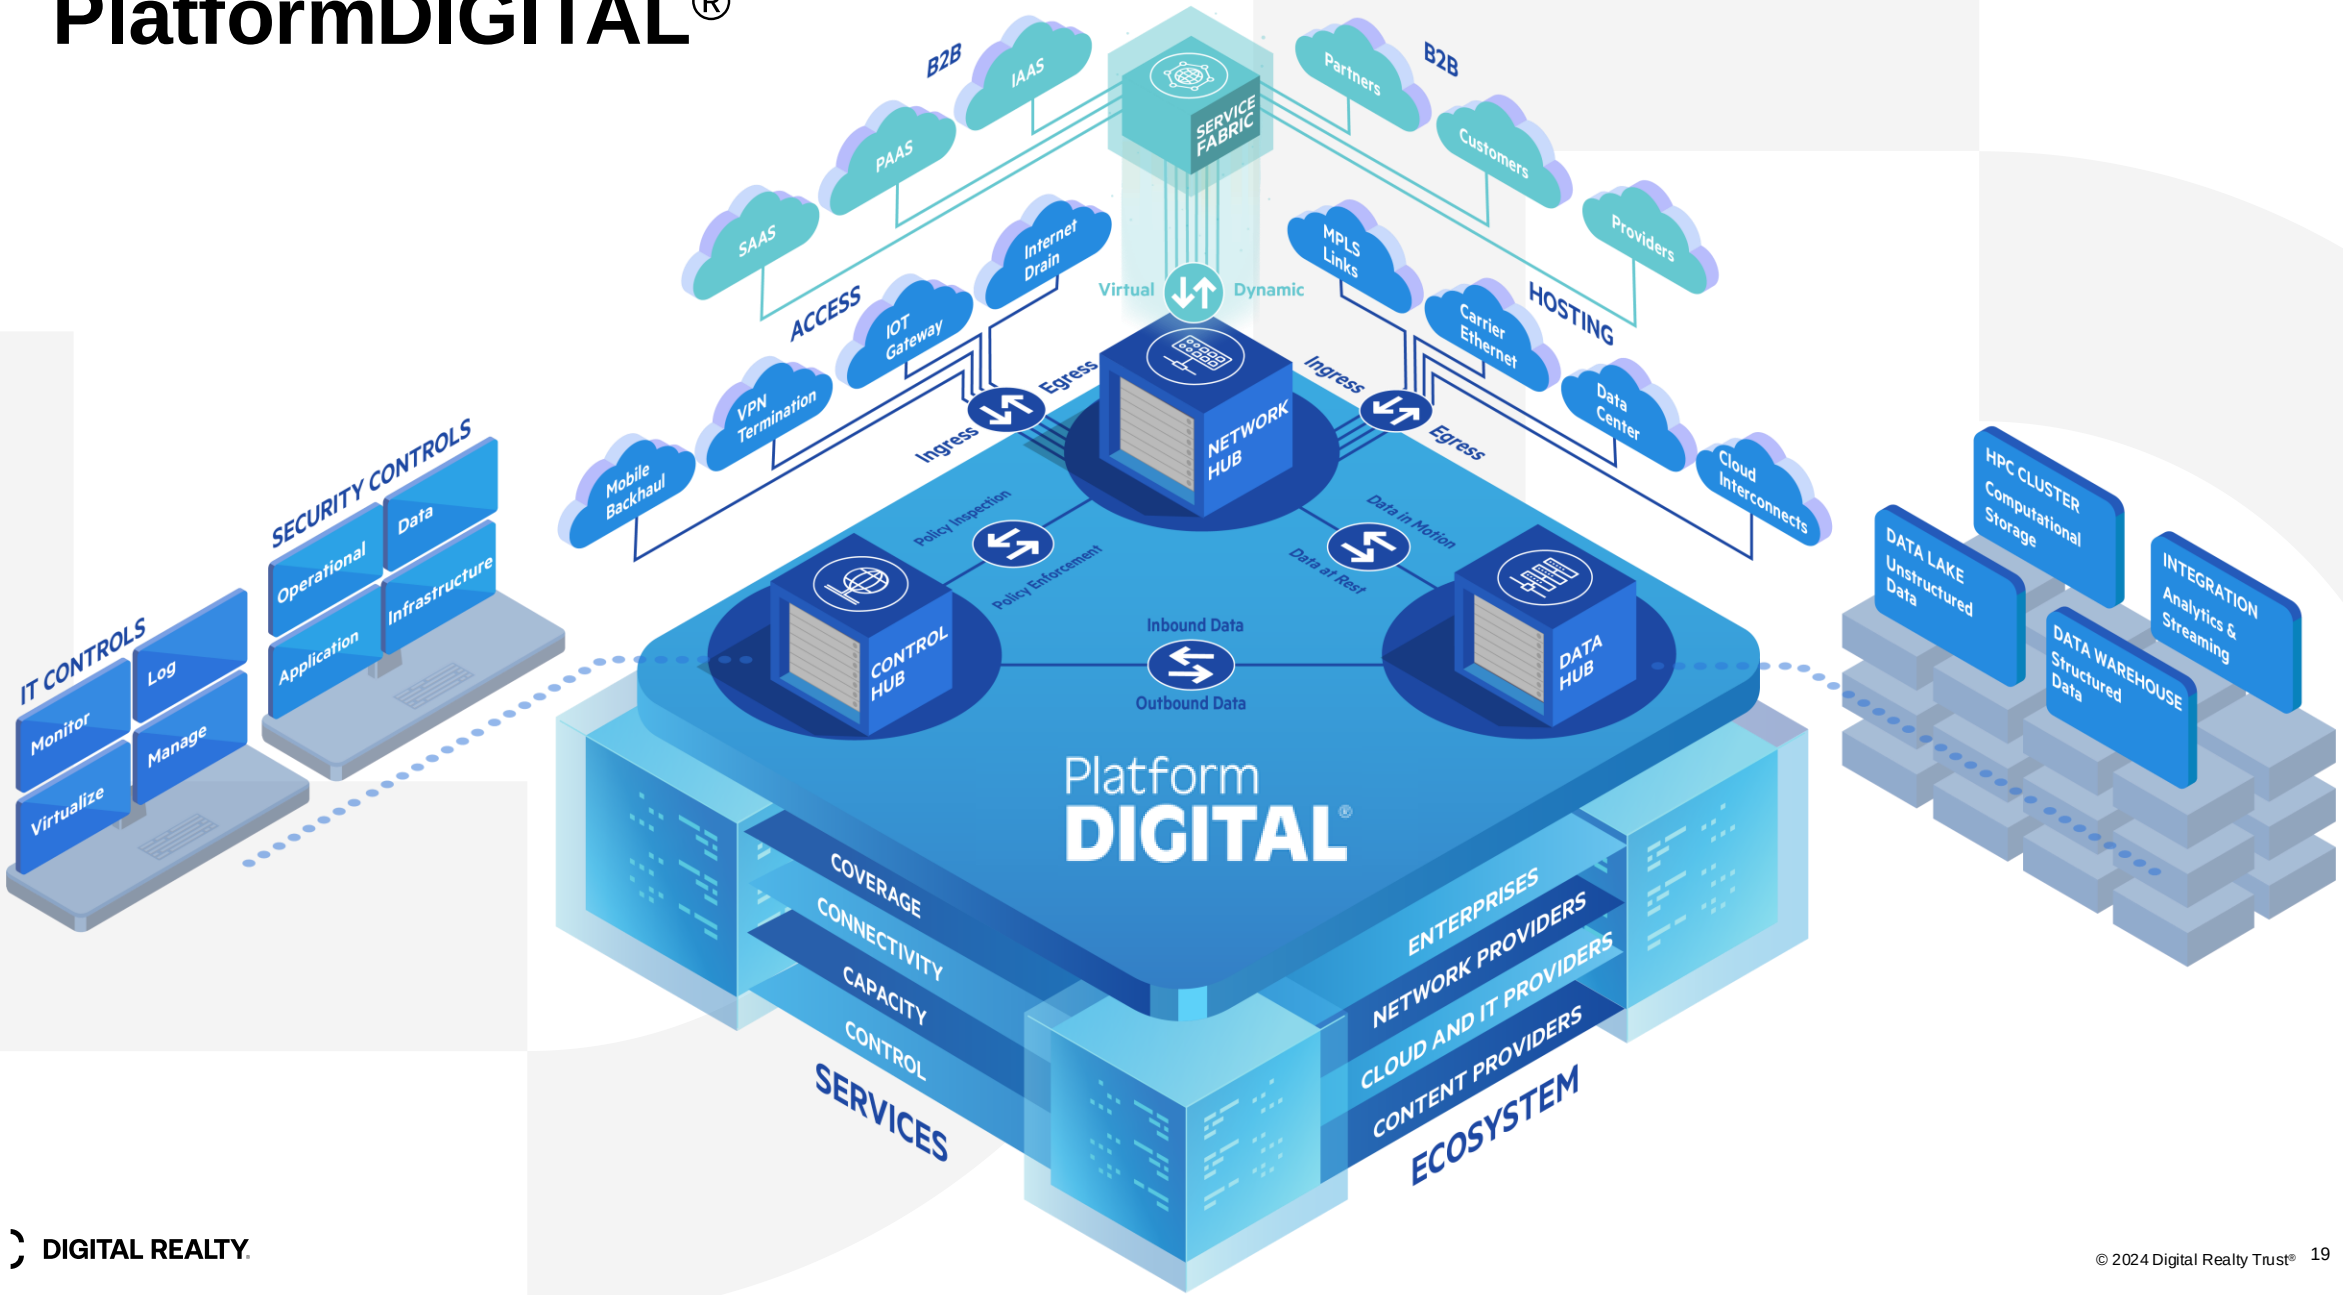

ServiceFabric™ Connect: Private Connections Made Easy

1. **Establish a Port** which supports multiple virtual private connections

2. **Create a Link** for direct, private connections to multiple Cloud Service Providers, Network Providers, SaaS Providers and other participants of the platform

3. **Add Services (optional)**
   Optimize and secure workflows

4. **Press** CREATE
   and the connection will be available within a couple of minutes

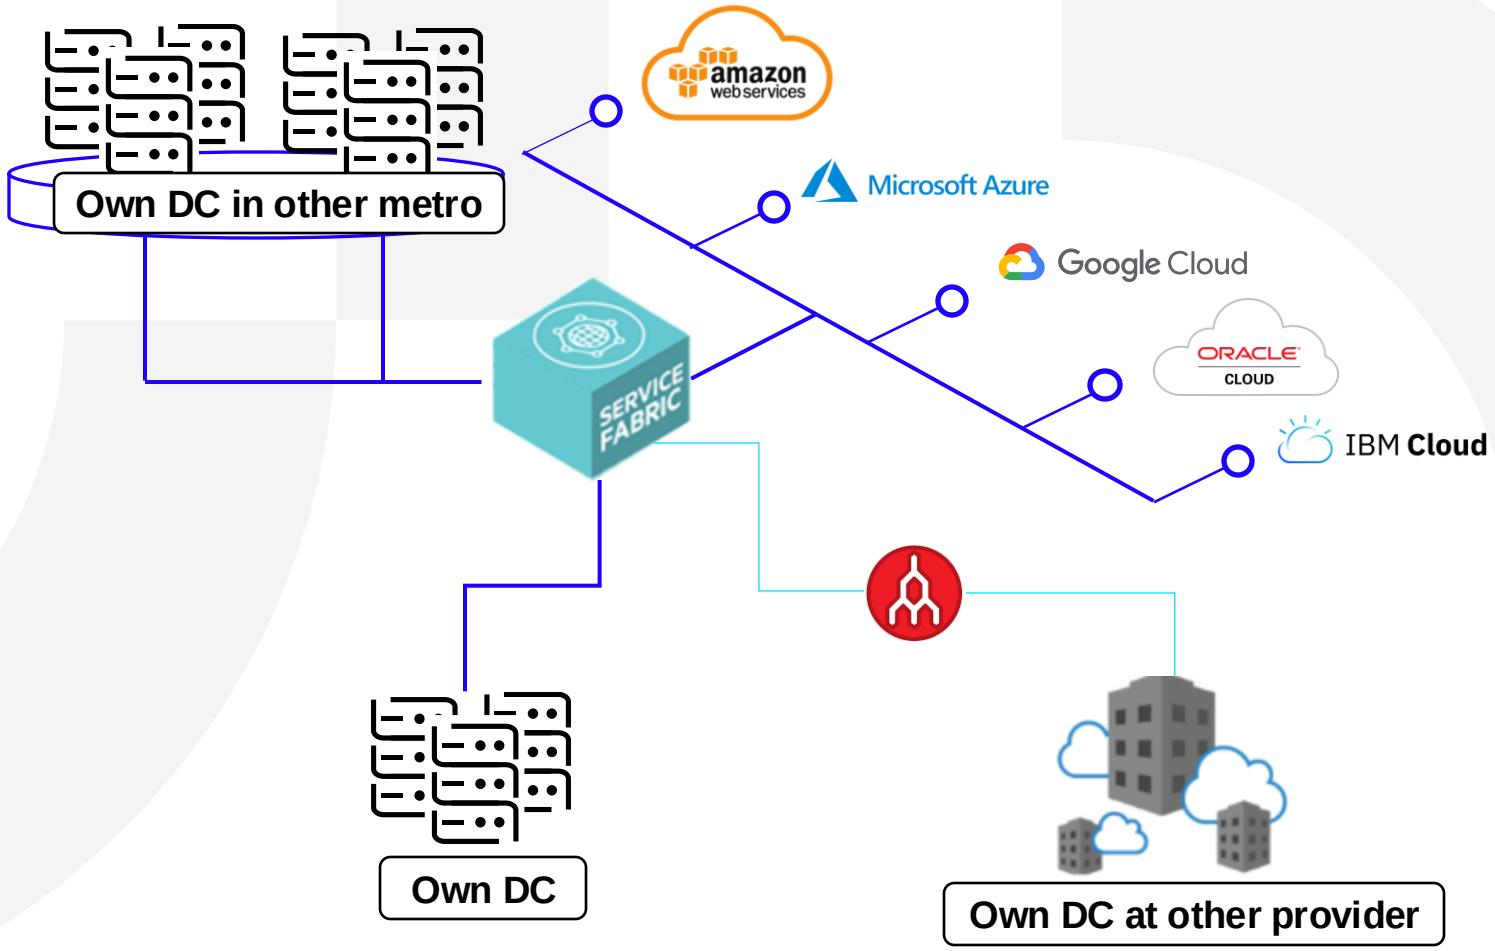

Gain First Mover Advantage by Enabling Private Interconnections

Gain first mover advantage by enabling private interconnections to the AI Ecosystem you (and your customers) need to access. Deploy AIPx solutions on PlatformDIGITAL.

FUTURE-PROOFED

HD Colo accelerates 70 kW/rack deployments in 34 markets across the global

ORCHESTRATION

ServiceFabric enables intelligent AI workflow orchestration via service chaining

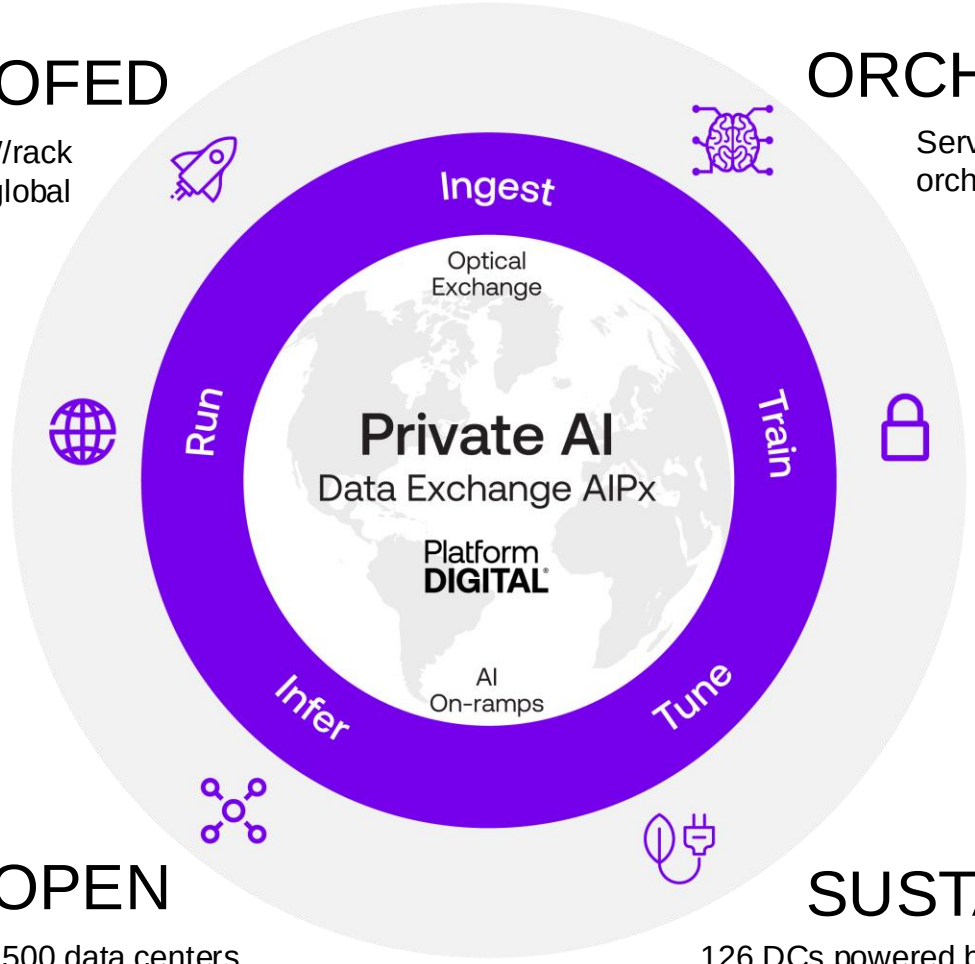

GLOBAL

Deployed in 50+ metros across 6 continents to meet your customers where they are

SECURE

Private network, private model access & tuning, distributed in key markets

OPEN

Access to 5,000 customers, 500 data centers, and 135 services from 80 service providers

SUSTAINABLE

126 DCs powered by renewable energy, 1.3 GW of wind and solar capacity globally

DIGITAL REALTY

ServiceFabric™ Connect

Private AI Data Exchange (AIPx) leverages ServiceFabric,™ enabling partners to extend their AI Infrastructure and Services over a single port (private, dedicated, AI On-ramp). The Service Directory, Service Key and Click to Connect features simplify and accelerate service distribution, consumption, and time to revenue to participating Enterprise and Service Provider customers and partners:

- **Distribution & aggregation ports** (AI On-ramps)
- **Service Directory** – Advertise services & manage distribution ports to the AI Ecosystem
- **Service Key** – Enable connections via discrete assignable connection keys
- **Click to Connect** - Designate and manage your click to connect ports, automatically accepting connections

Digital Realty is committed to making frictionless connections across workflow participants, applications, multiple clouds, networks and ecosystems.

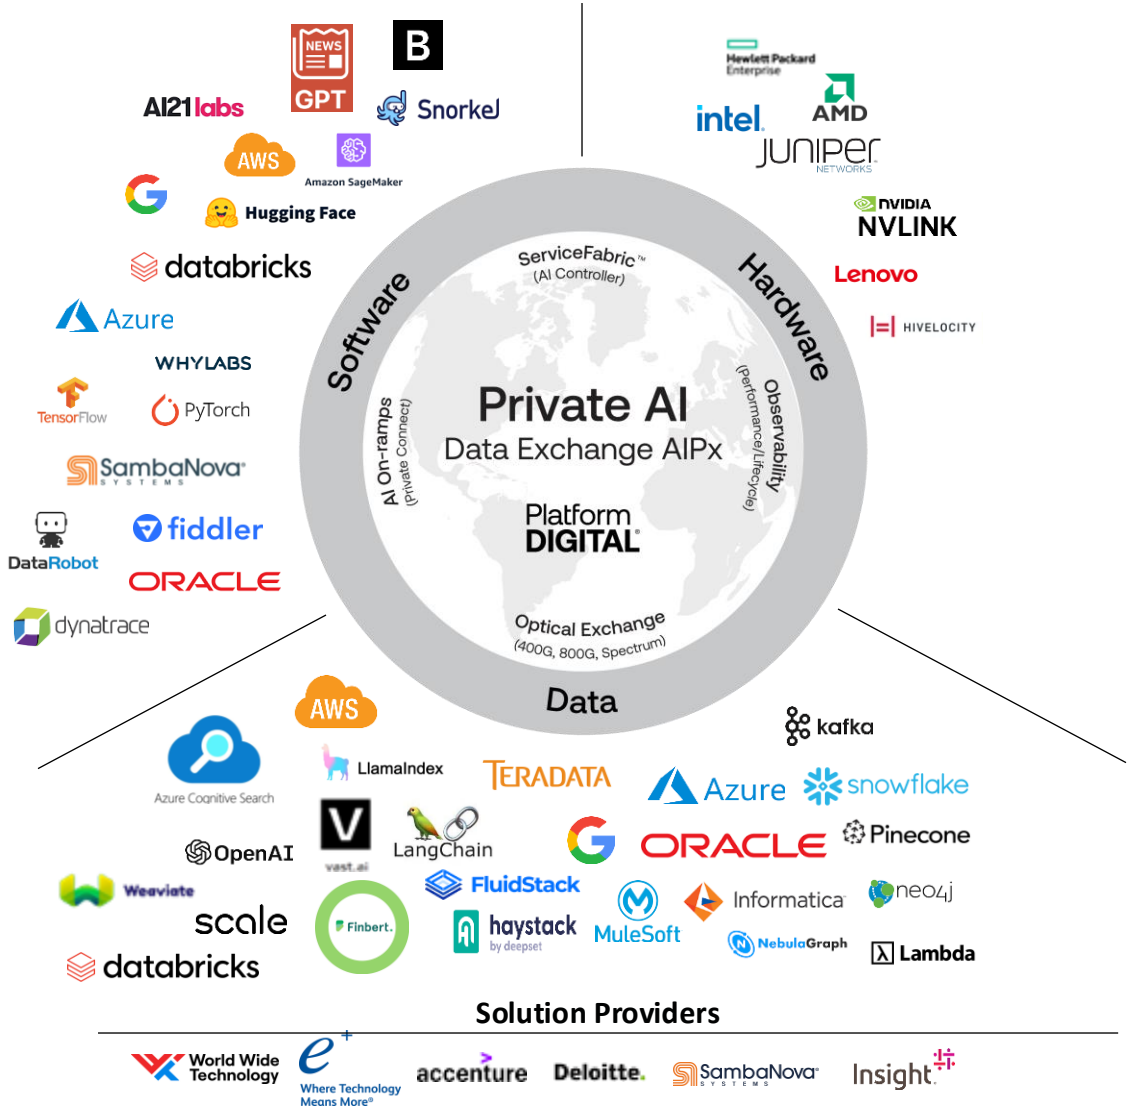

All logos are illustrative only

PDx Visualizer: A Virtual Tour of the Ecosystem

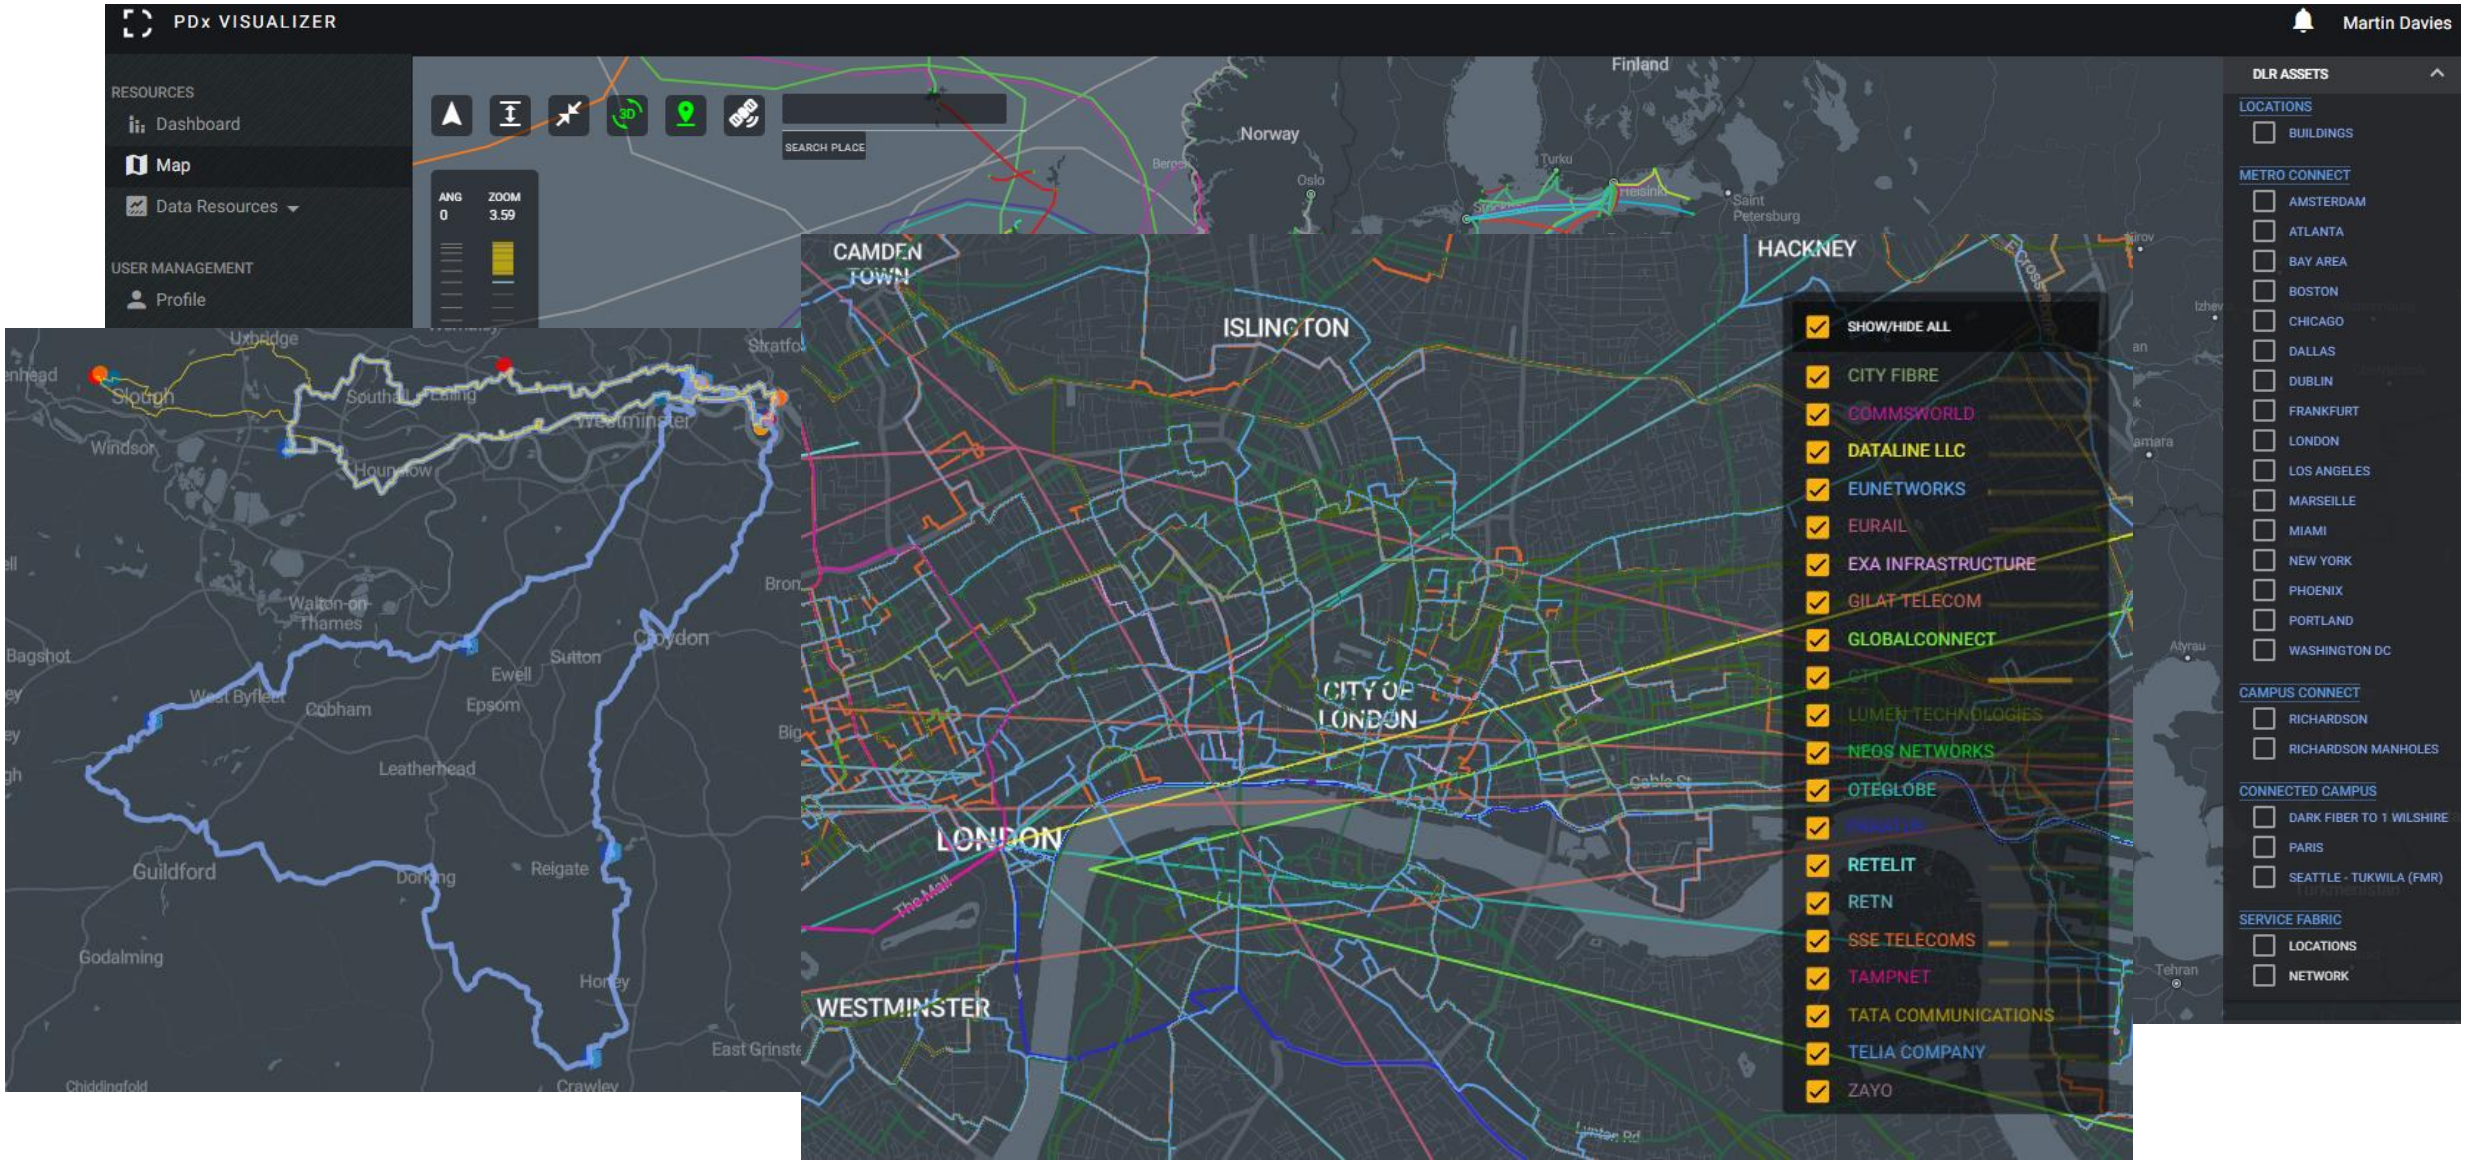

PDx Visualizer: A Virtual Tour of the Ecosystem

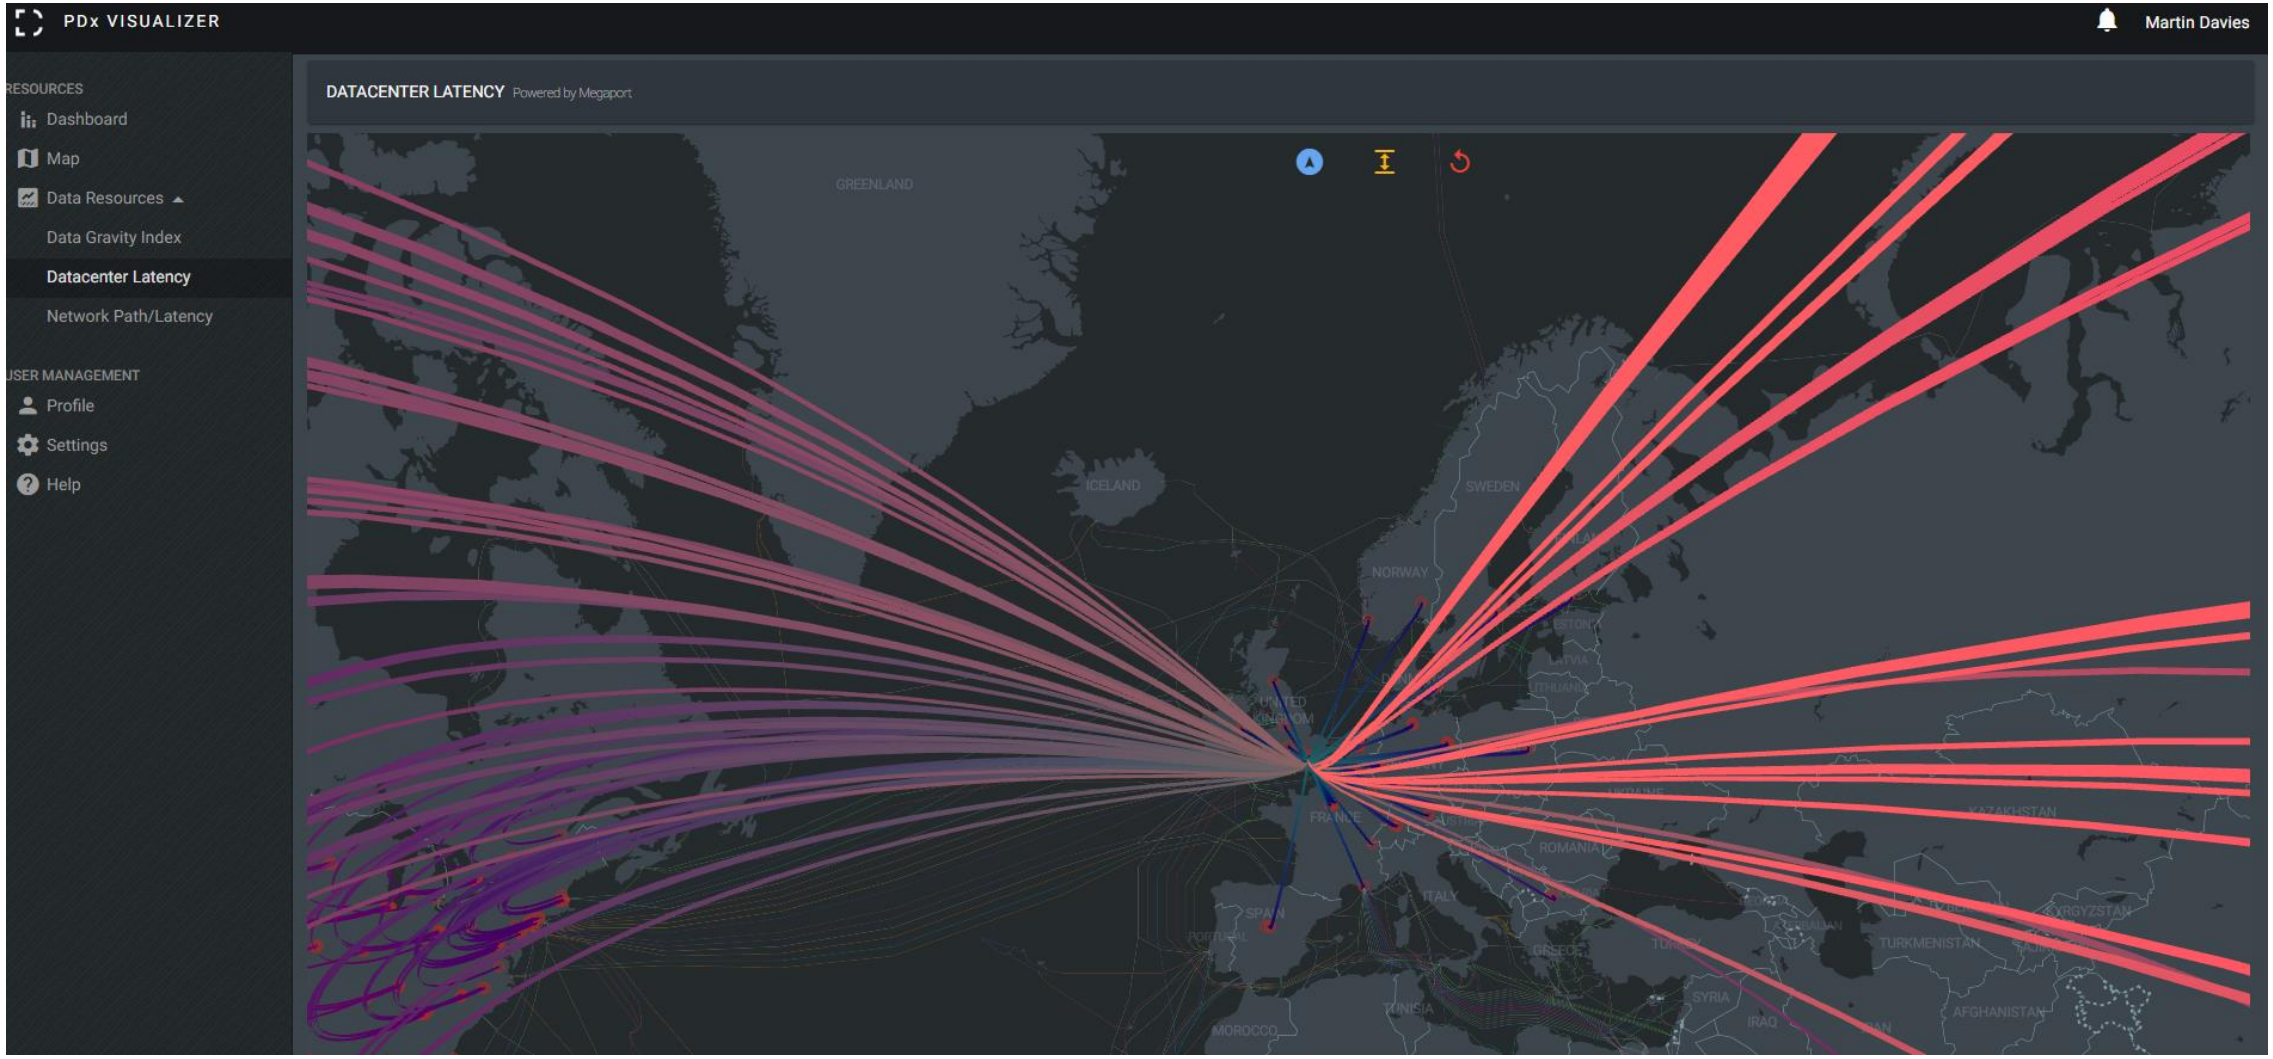

PDx Modeler: Model Customer Future Deployments

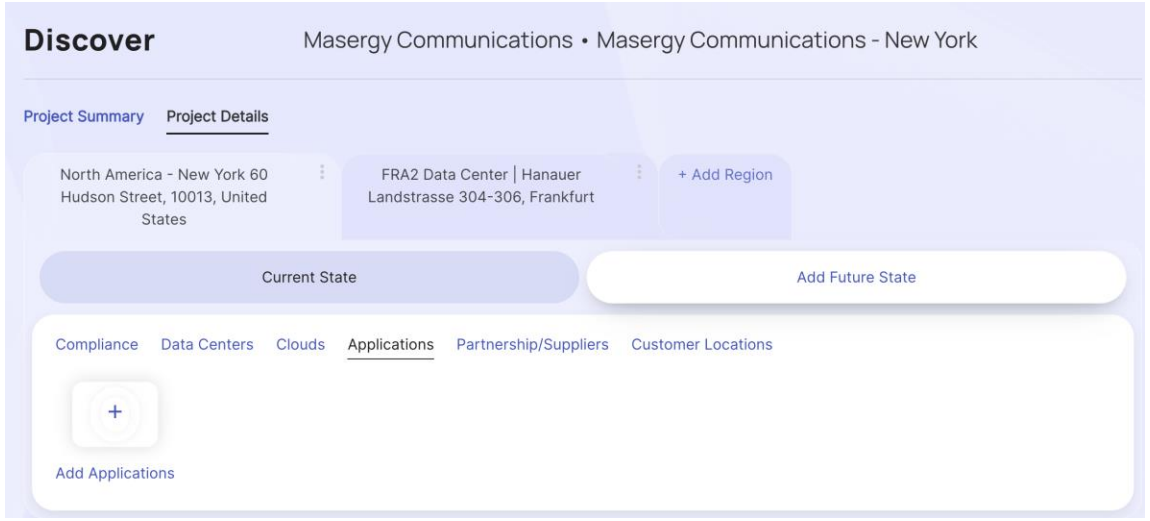

Choose Application

ADP

asana

Atlassian

AutoCAD

Choose Network Service Provider

Airtel

Altice Business

Arelion

Aryaka

AT&T

CATO Networks

Charter

Cogent Communications

Choose Cloud Provider

Alibaba

AWS

AWS Region

Azure

KEY TAKE AWAY

1. Network, increased ability to connect service providers together

2. Data centres, tackling constraint in high density compute

3. Connect to anyone, anywhere, anytime

Thank you.

Connect with Digital Realty

Darren Chan

Darren Chan
Sr. Solution Architect

darrenchan@digitalrealty.com