Data governance
Cyber protection
Data mobility

Nanosecond machine calls
High concurrency
Unbounded scale

Speed and revenue generation depend on
effortless, worry-free handoffs to the machine… Sounds Easy Enough.. BUT!

Gravity is Real!!

1) https://www.cio.com/article/3314742/trends-in-big-data-and-artificial-intelligence-data.html

We Want To React! What's the cost of opportunity?

It's Your Call, But!!!

Adjacency to Stream, Compute & Code

- HCI
- Containers
- Infrastructure as code
- Bare Metal
- GPU
- FPGA
- ASICs

1) https://www.cio.com/article/3314742/trends-in-big-data-and-artificial-intelligence-data.html

We Need To Learn! ! What's the cost?

Still your Call, But…..!!

Adjacency to Stream, Compute & Code Storage, Network

- HCI / ASICs/ GPU/FPGAs
- Containers
- Infrastructure as code
- Bare Metal
- Object vs File
- More Data Gravity
- Scaling
- Life cycle management
- Security
- Cost
- Complexity
- New Tools
- Where do you want to run your work!

1) https://www.cio.com/article/3314742/trends-in-big-data-and-artificial-intelligence-data.html

Defining the Cloud Ecosystem

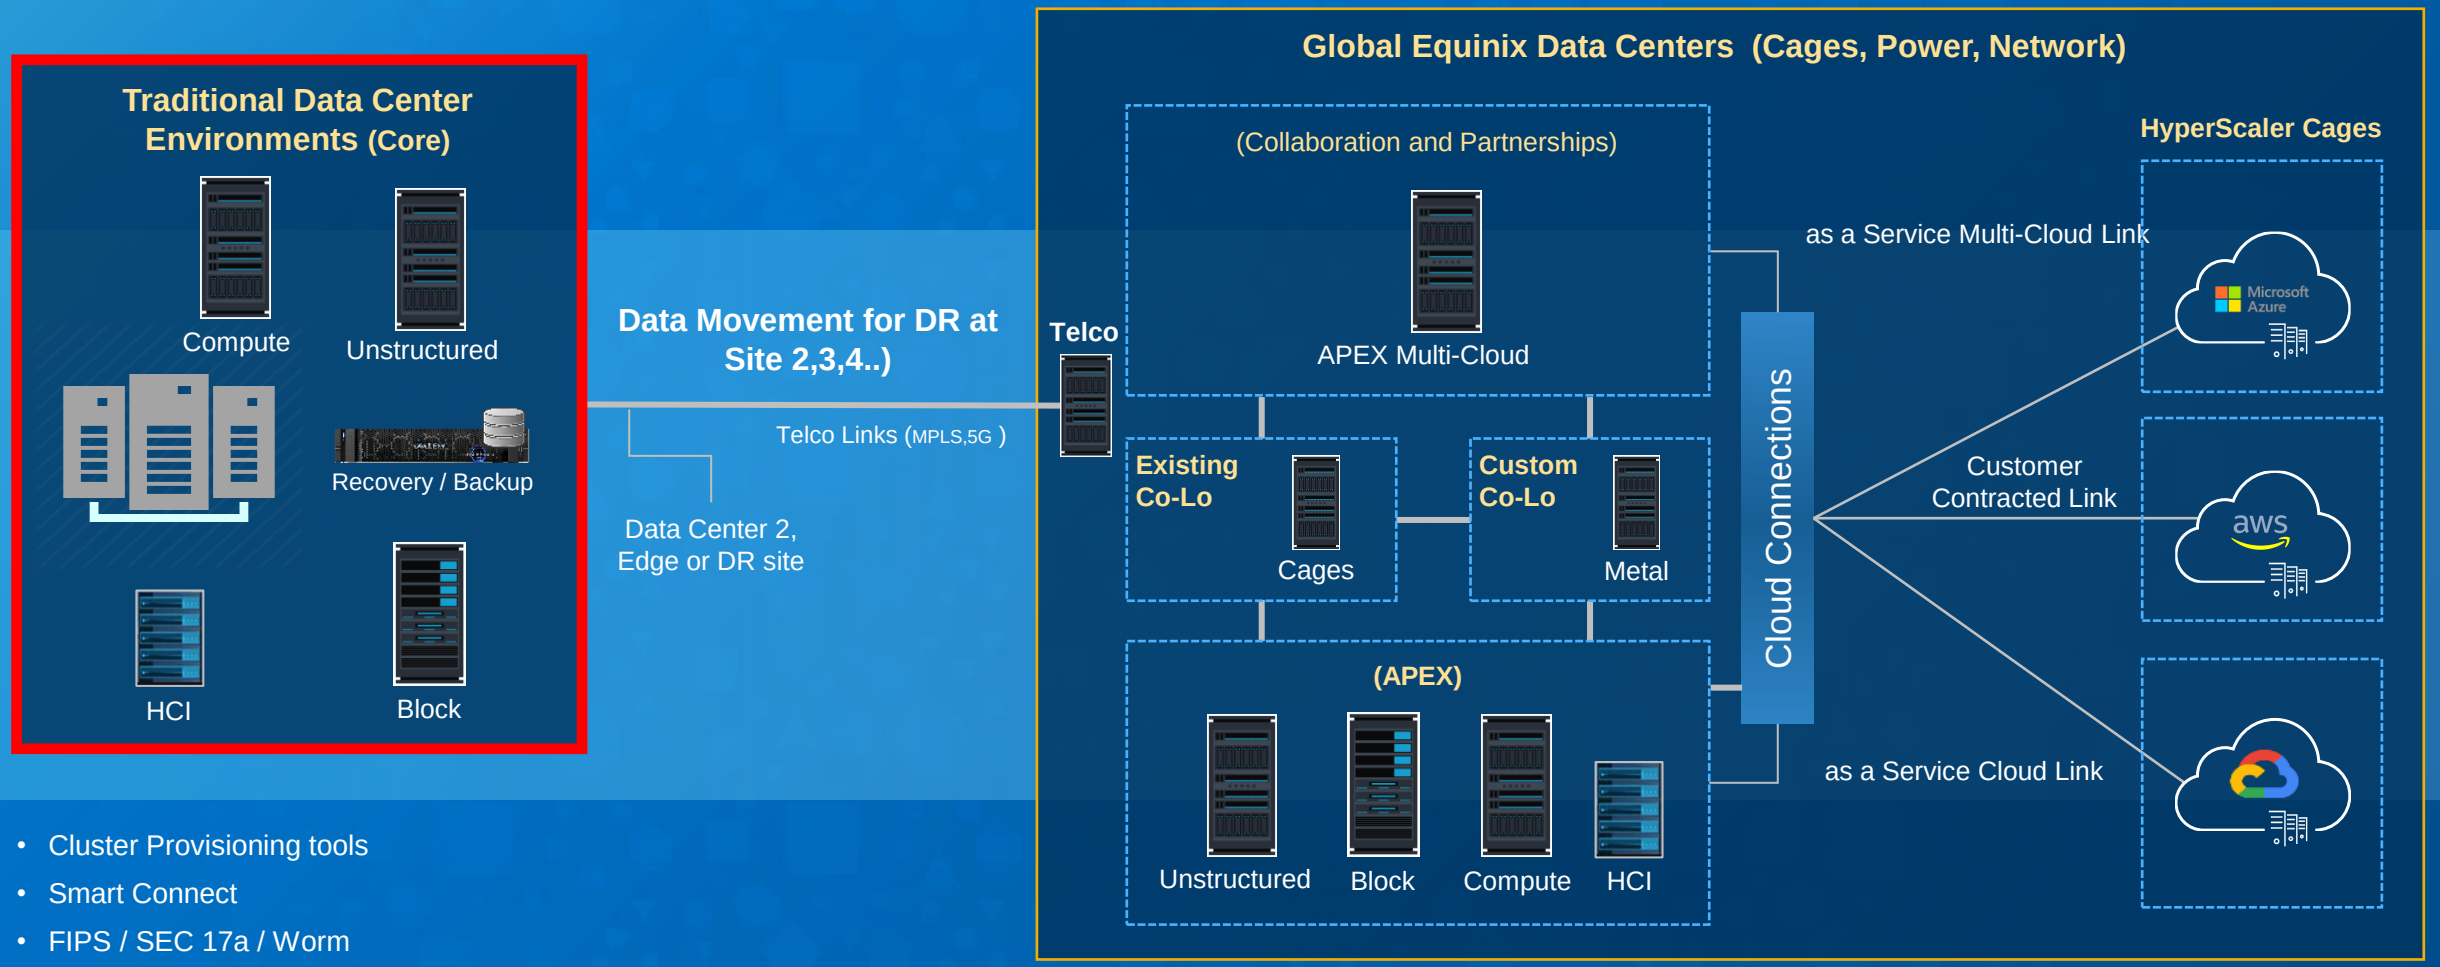

- Cluster Provisioning tools
- Smart Connect
- FIPS / SEC 17a / Worm
- Next Gen Access: RAN / CSI / API / S3 / V-Connect
- GPU direct and RDMA over NFS
- Certified HDFS qualification with Cloudera
- Advanced Security tools: Ransomware Defender, Easy Auditor & advanced Air Gap tools

Defining the Cloud Ecosystem

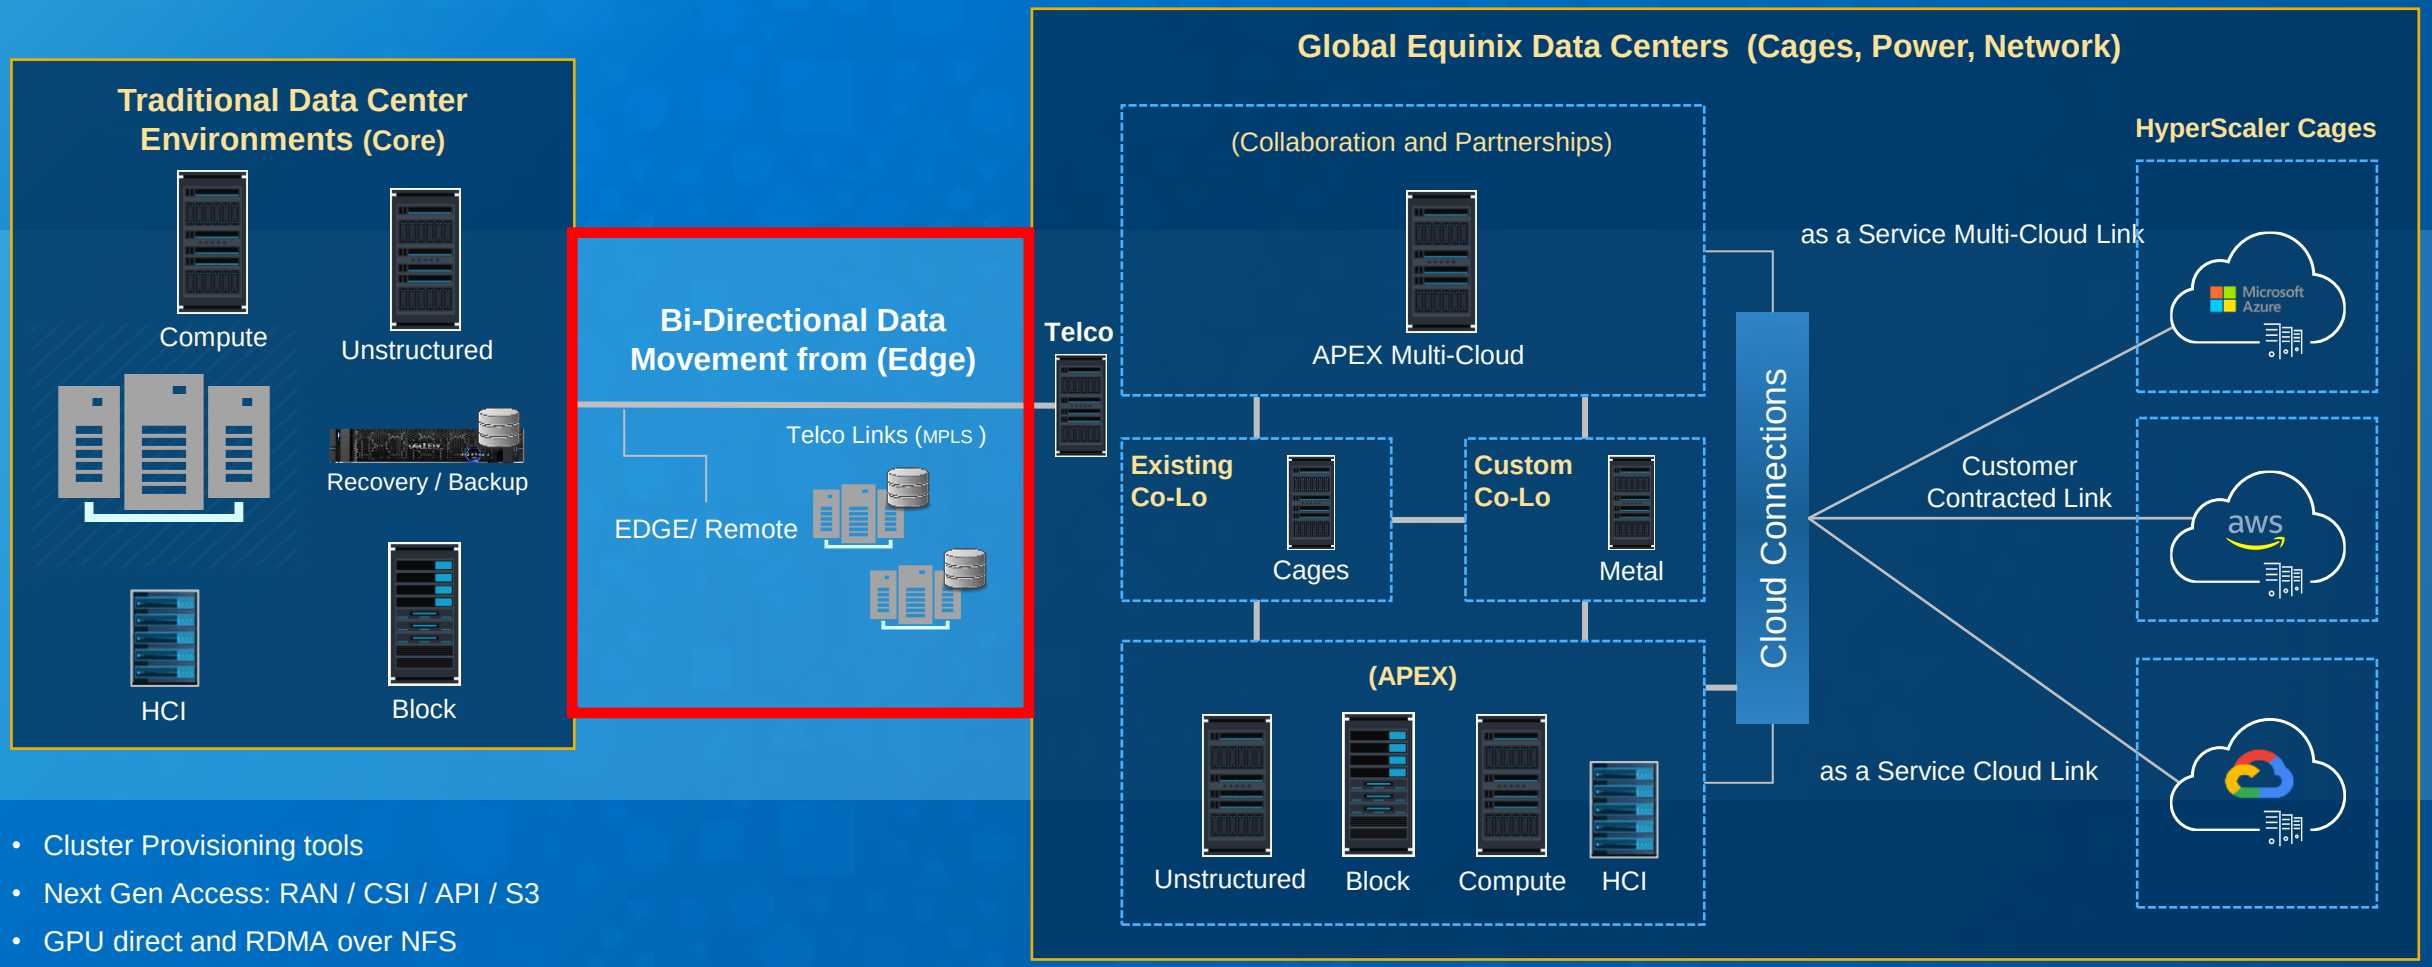

- Cluster Provisioning tools
- Next Gen Access: RAN / CSI / API / S3
- GPU direct and RDMA over NFS
- Certified HDFS qualification with Cloudera
- Advanced Security tools: Ransomware Defender, Easy Auditor & Air Gap
- Networking strategies
- Replication interface optimization

Defining the Cloud Ecosystem

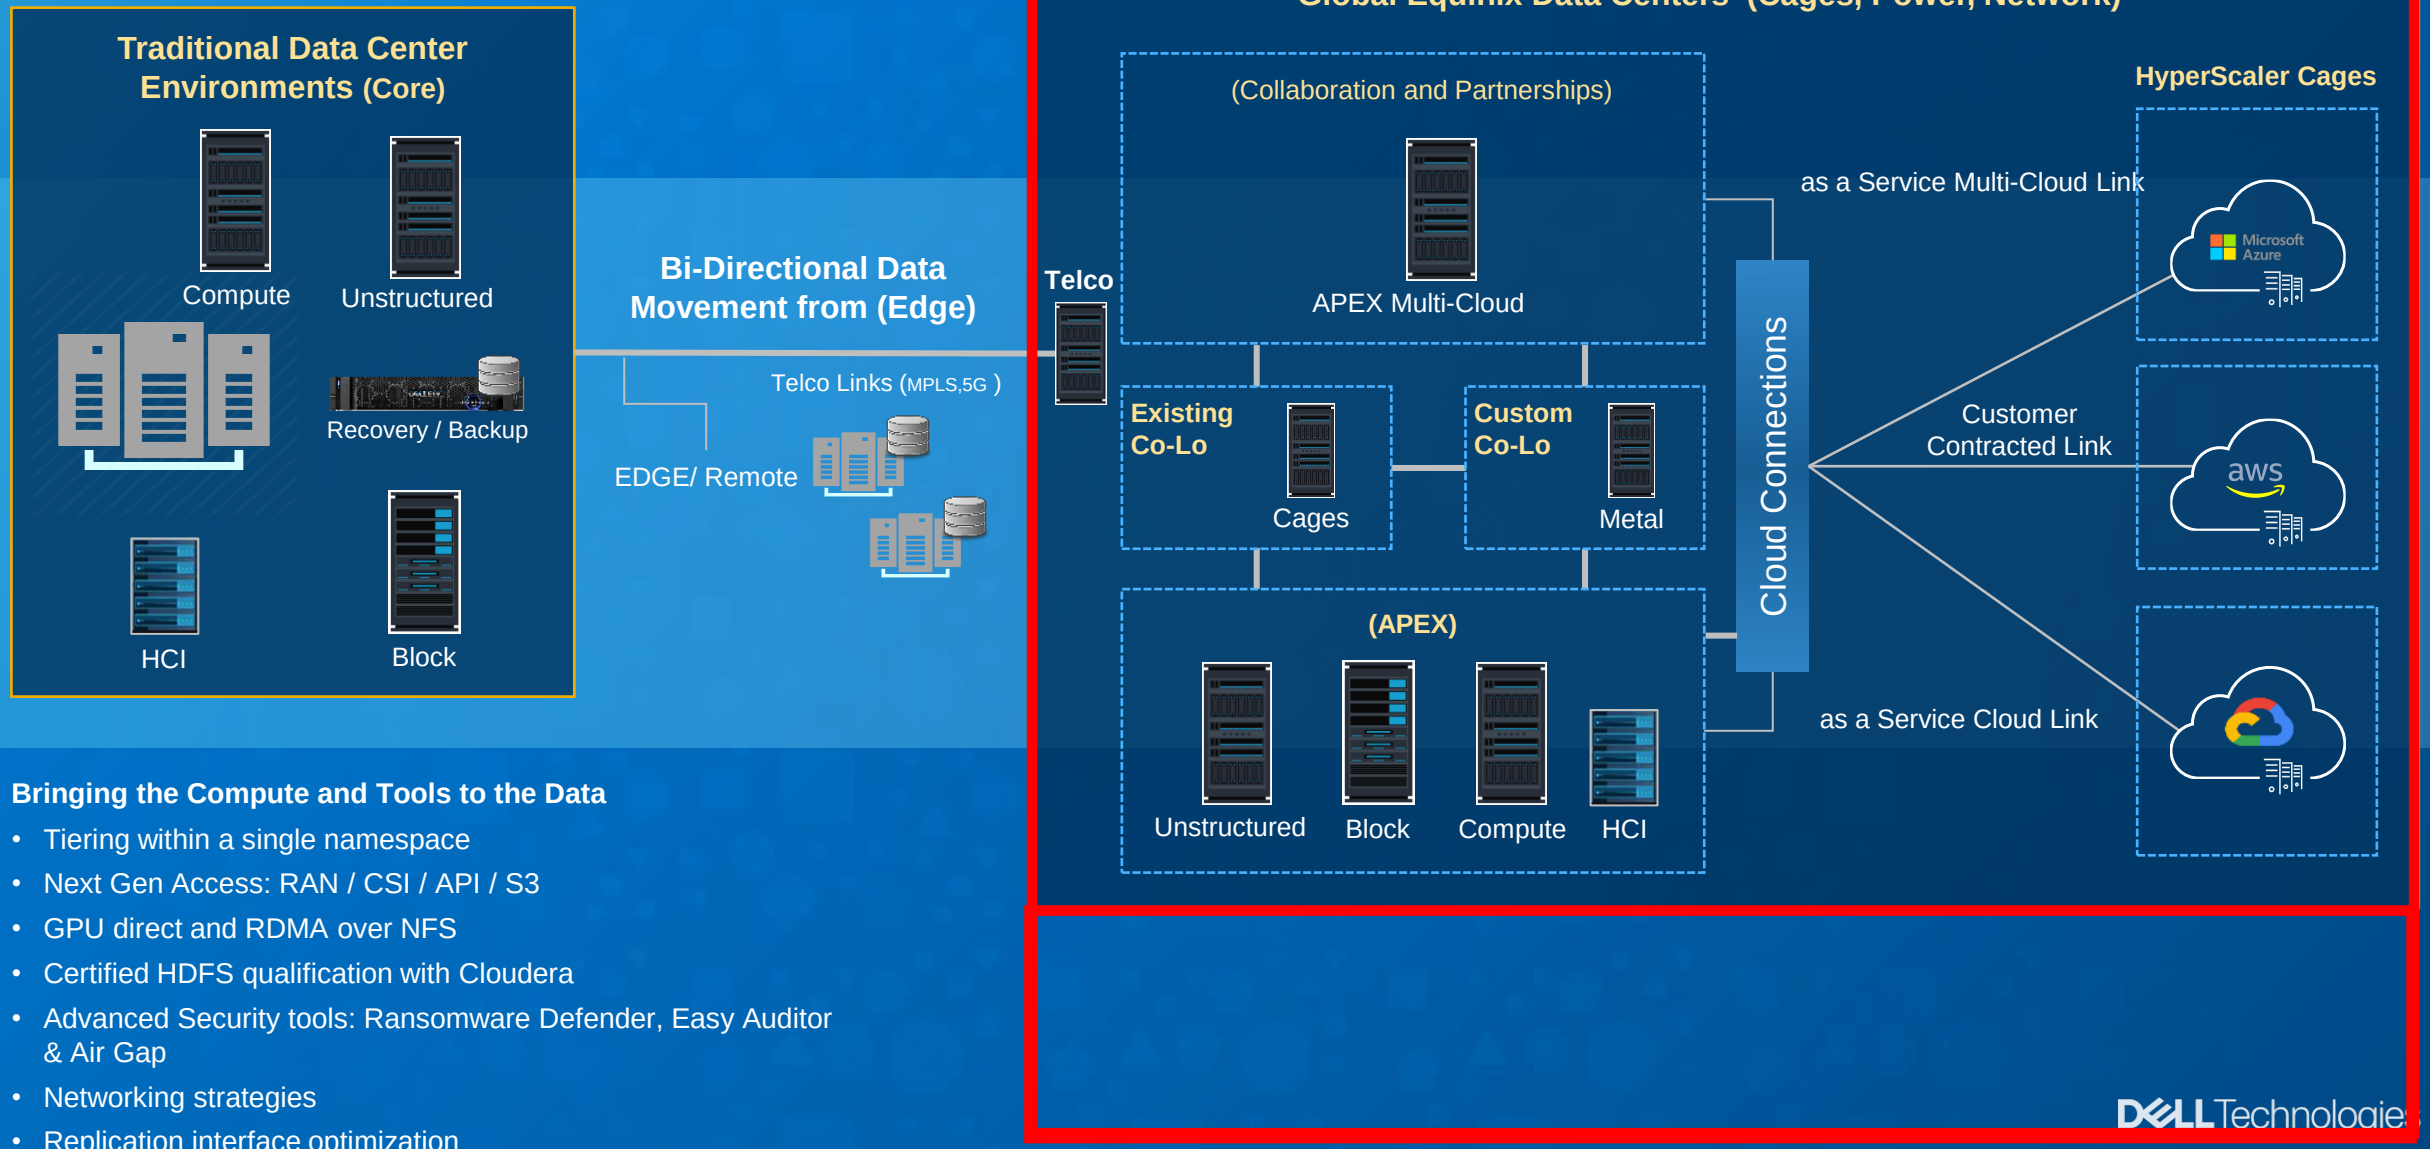

Bringing the Compute and Tools to the Data

- Tiering within a single namespace
- Next Gen Access: RAN / CSI / API / S3
- GPU direct and RDMA over NFS
- Certified HDFS qualification with Cloudera
- Advanced Security tools: Ransomware Defender, Easy Auditor & Air Gap
- Networking strategies
- Replication interface optimization

Defining the Cloud Ecosystem

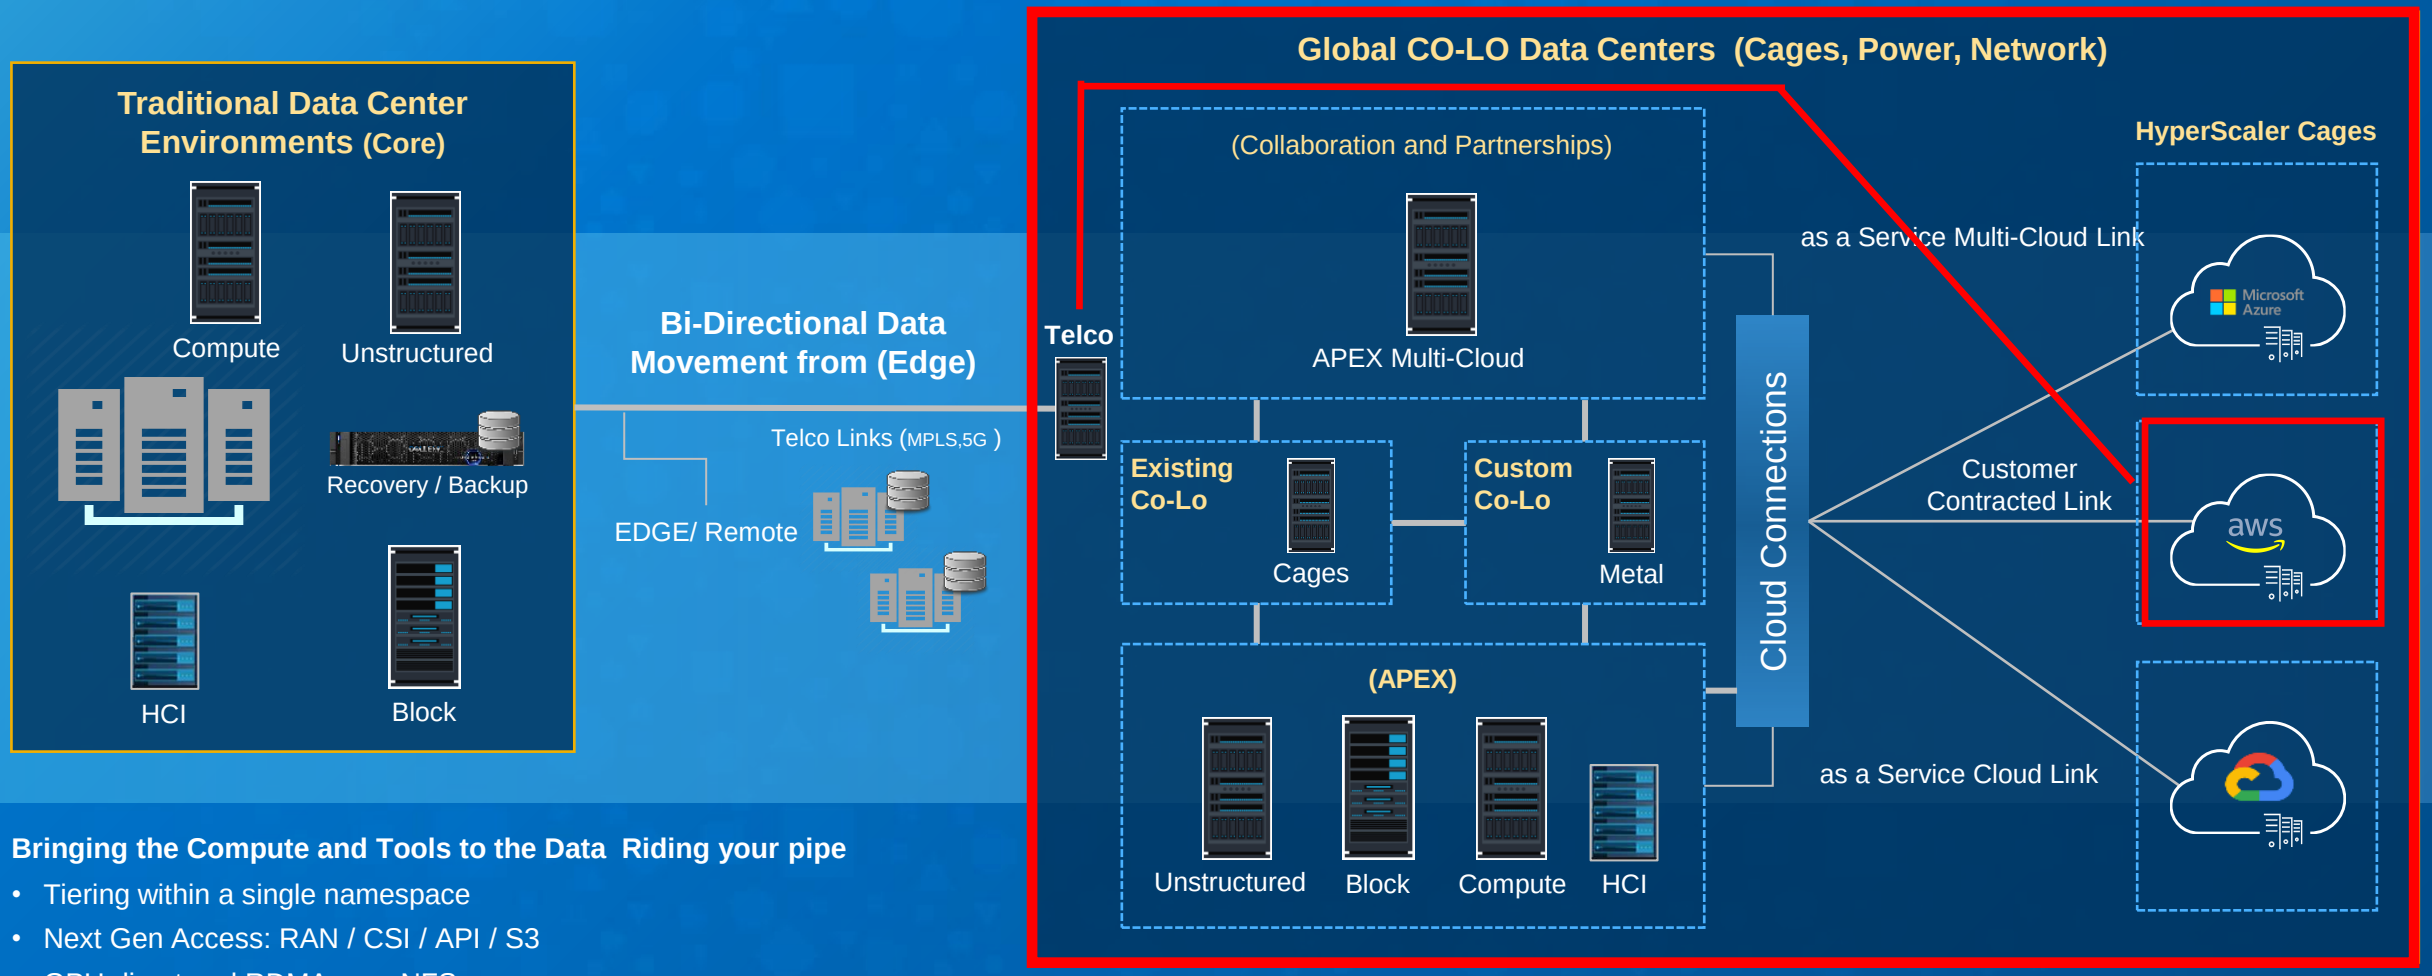

Bringing the Compute and Tools to the Data Riding your pipe

- Tiering within a single namespace
- Next Gen Access: RAN / CSI / API / S3
- GPU direct and RDMA over NFS
- Certified HDFS qualification with Cloudera
- Advanced Security tools: Ransomware Defender, Easy Auditor & Air Gap
- Networking strategies
- Replication interface optimization

Defining the Cloud Ecosystem

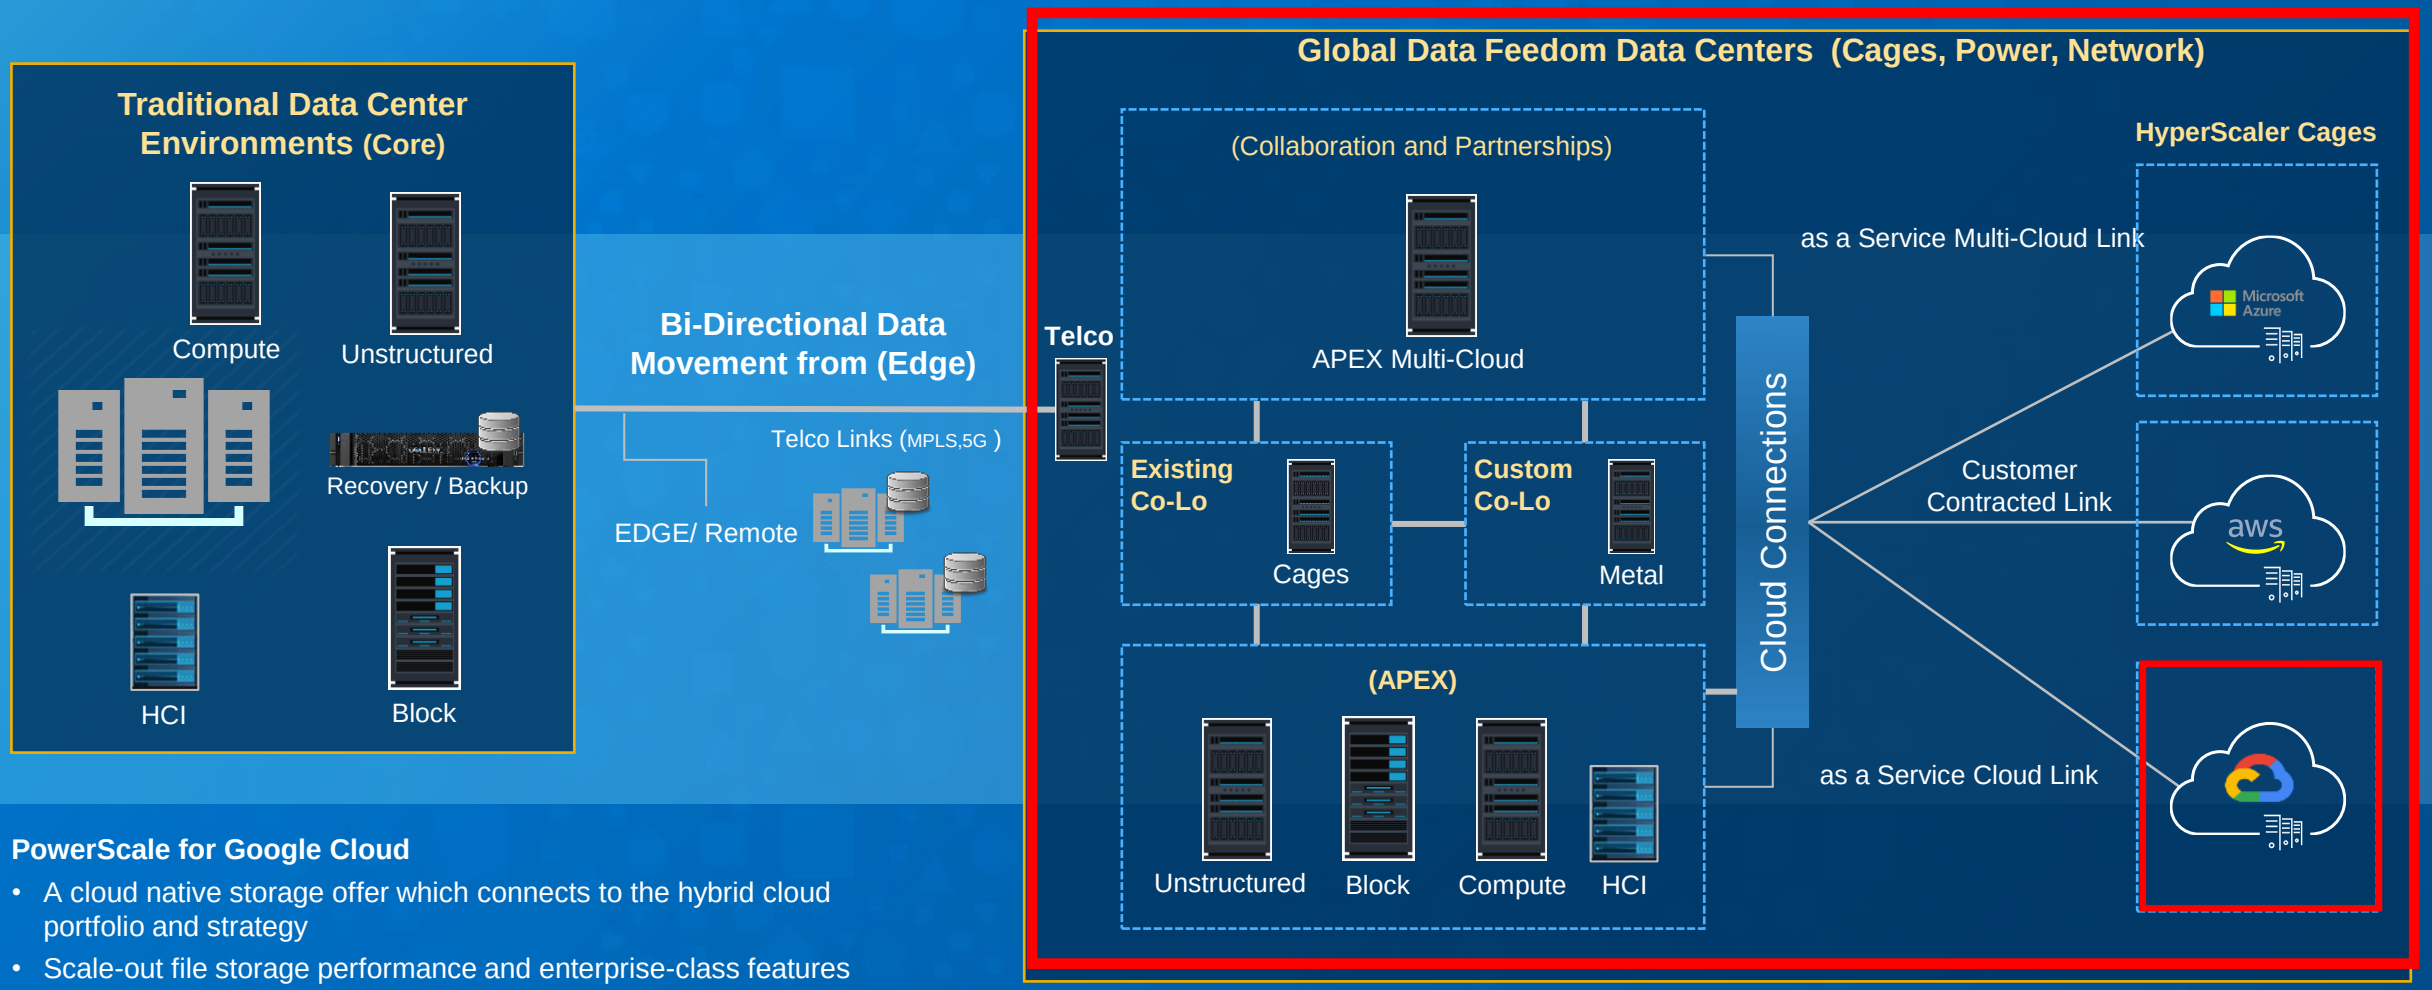

PowerScale for Google Cloud

- A cloud native storage offer which connects to the hybrid cloud portfolio and strategy
- Scale-out file storage performance and enterprise-class features
- Provisioning, configuration and management in Google Console
- Single bill and support from Google
- Predictable pricing and guaranteed performance
- Complete lifecycle management with SLAs from Dell Technologies

Challenges customers are facing

73% of company data goes unused for analytics and decision-making

49% of remaining 27% data set is poor quality data due to unclassified, unusable, legacy/old

>70% of data warehouses/lakes are siloed and disconnected across public / private clouds

Infrastructure complexities

People challenges

Data challenges & complexities

EVOLVING BUSINESS EXPECTATIONS

- ✓ Must meet rising customer expectations
- ✓ Varying sets of objectives
- ✓ Consume data in real time
- ✓ Use of AI & ML for decision making

22% of data management resources' time is spent on data innovation and monetization

80% of data created is unstructured and in decentralized locations. By 2025, 75% of data will be created at the edge

Why is Data Transformation important?

Large amounts of **raw data**

Create **actionable insights**

Identify **strategic use cases**

Leading to **business value creation**

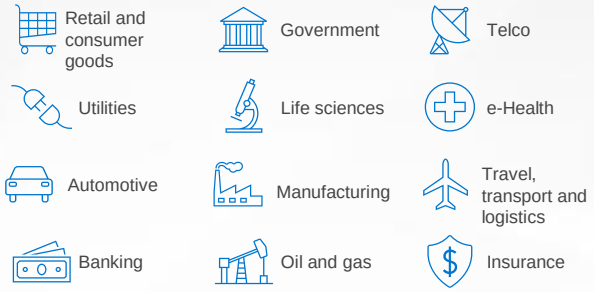

At the speed of opportunity

Customer Analytics

Drive Revenue
New revenue opportunities*

- Product Development
- Business Decision Making / Recommendations
- Customer Experience
- Personalized Healthcare
- Exploration & Discovery
- Billing Optimization
- Channel / mix optimization
- Pricing optimization

Operational Analytics

Reduce costs
Realized cost savings*

- Predictive maintenance
- Operations optimization
- Supply chain optimization
- Manufacturing operations
- Asset performance
- Process engineering
- Capacity planning
- Retail store operations

Risk Analytics

Avoid risks
Increased profitability*

- Cyber security
- Compliance / Audit
- Fraud detection
- Safety & Reliability
- Anti-money laundering
- Rogue trading
- Service Reliability / Resilience

Today's Challenges

10 years ago

Market Data is in
Gigs Daily
and put in memory for algorithms

Jobs isolated to Quant machines handful of jobs

5 years ago

Market Data is in
500+ Gigs Daily
Architectures for algorithms deviate (central vs node based)

Thousands of jobs on various compute boxes

Today

Market Data is in
5 to 25 TB Daily
Architectures for algorithms differ (central vs node based)

Hundreds of thousands to millions of jobs

Dell’s Point Of View

Data Lifecycle Ecosystem

The complexity of process and technology integration

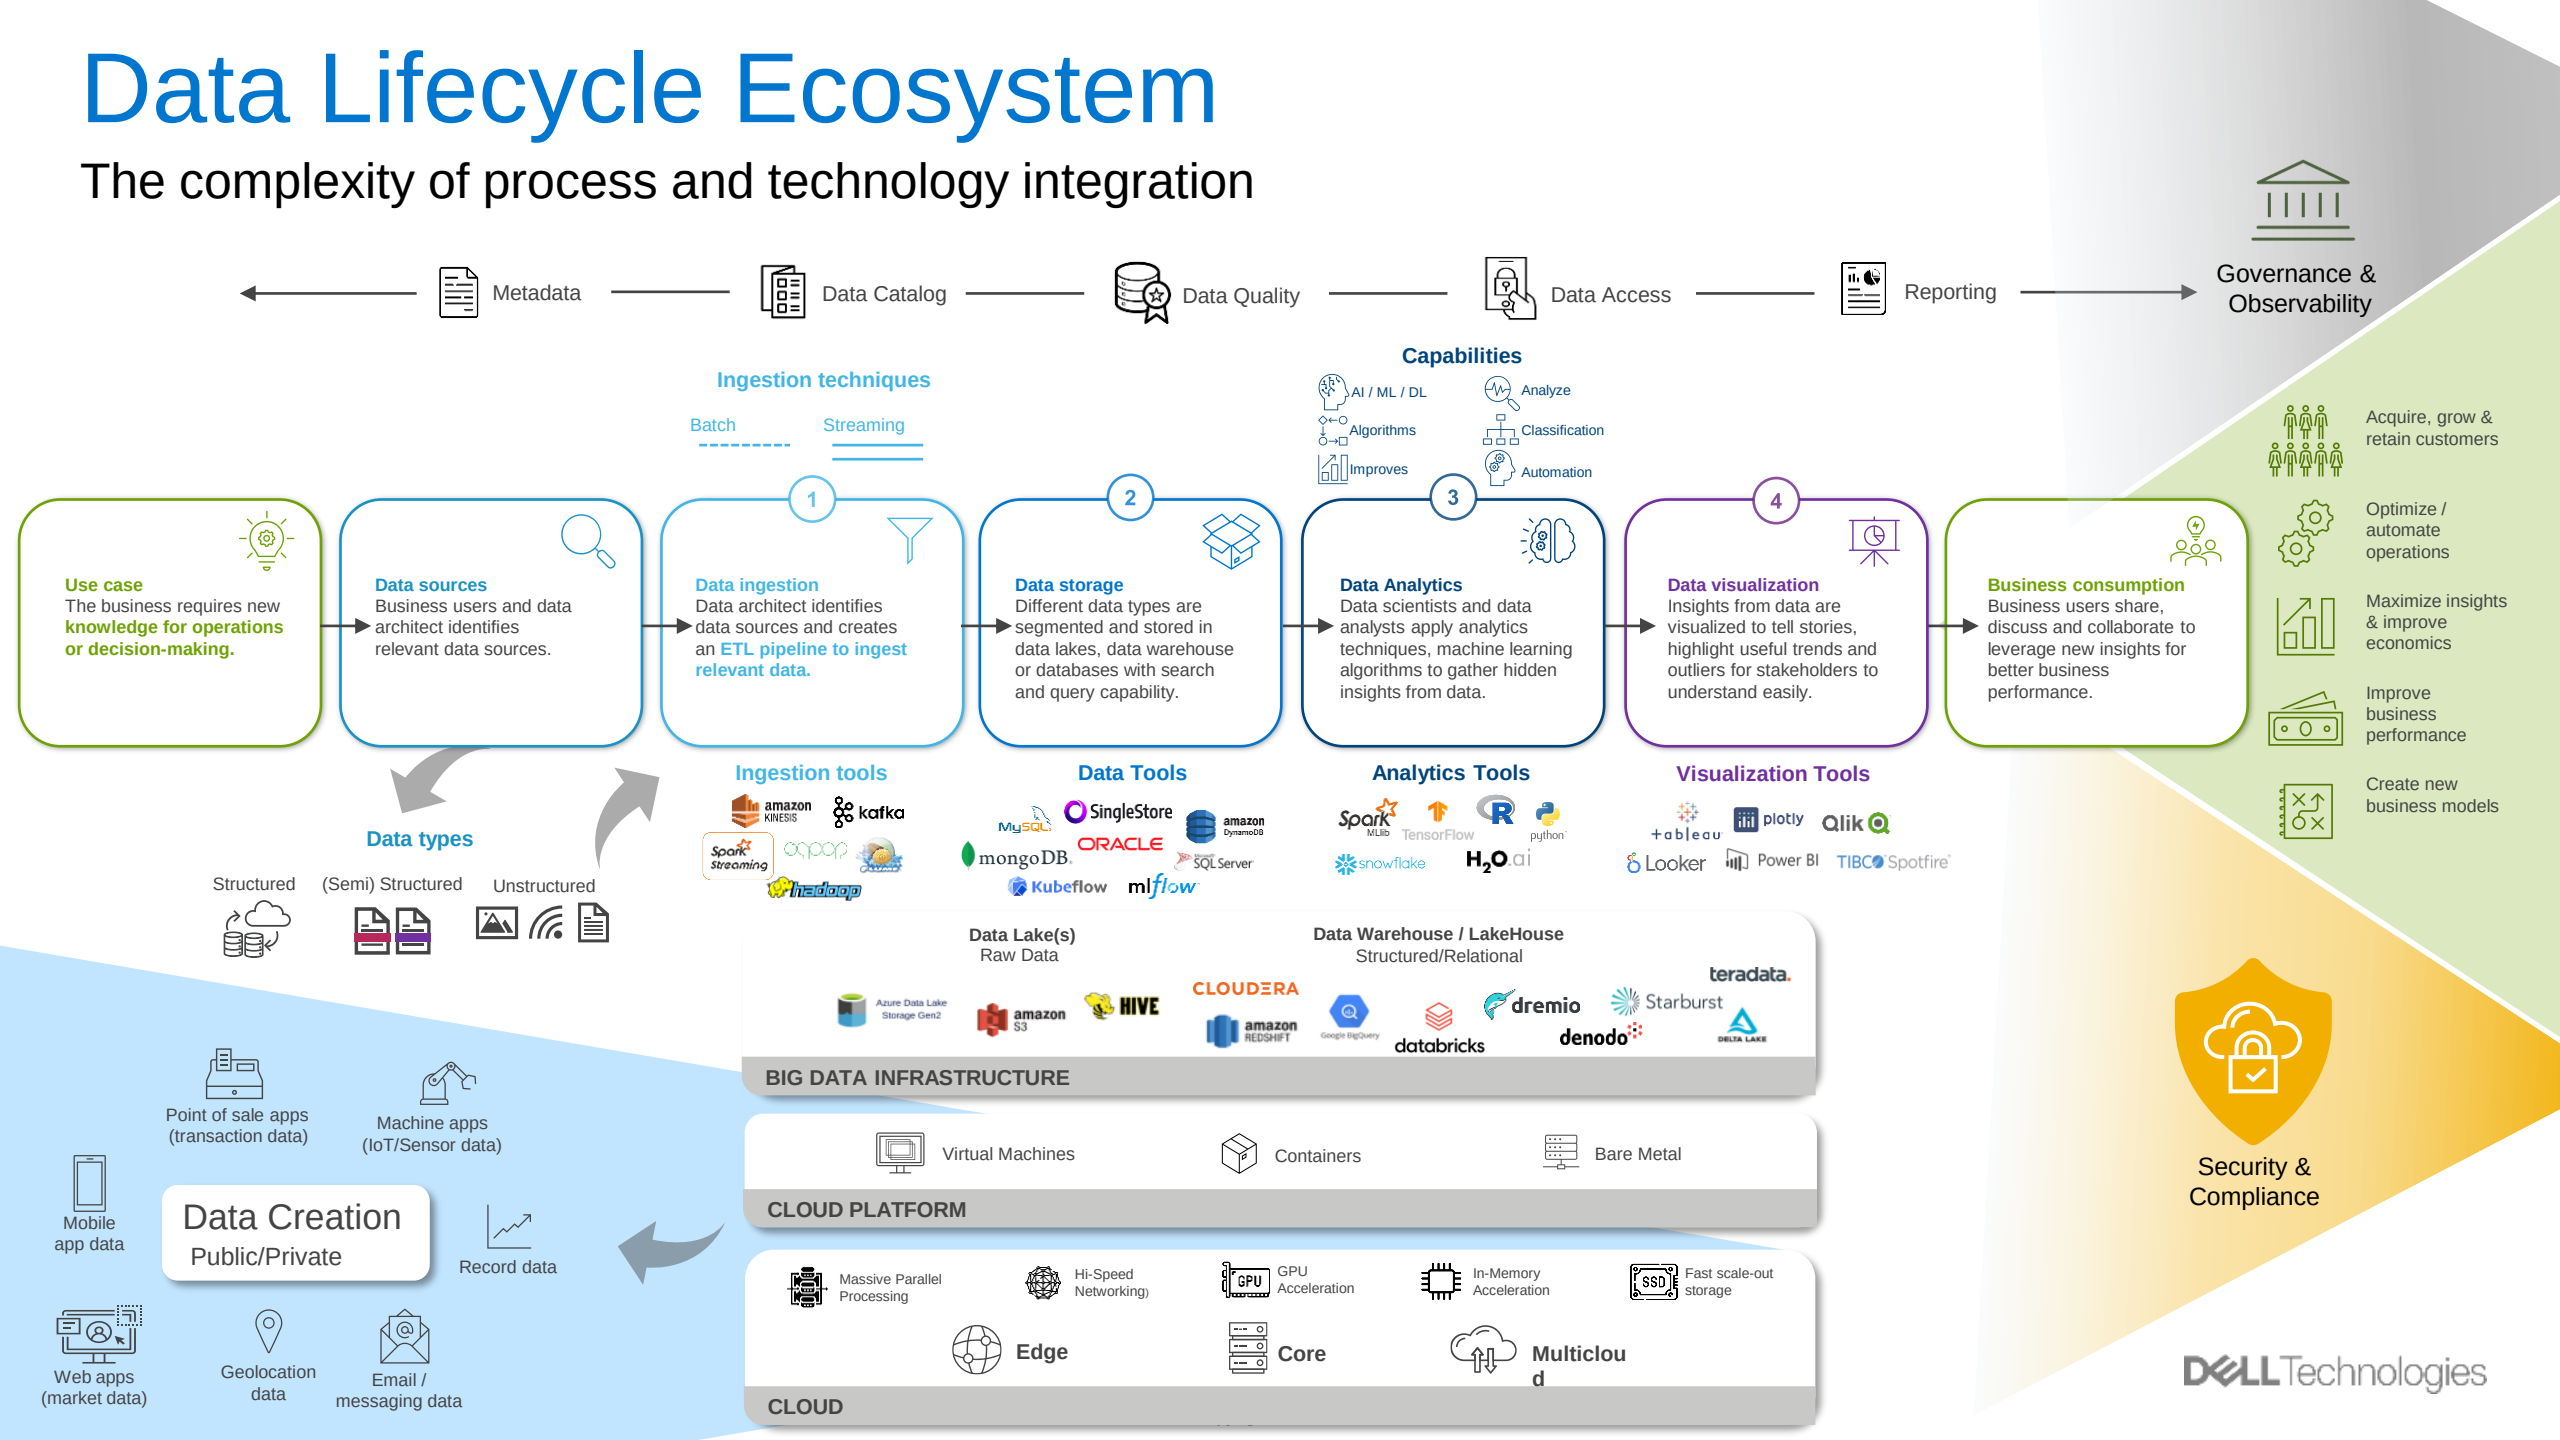

DELL Technologies

Data Management & Analytics Ecosystem

The complexity of process and technology integration

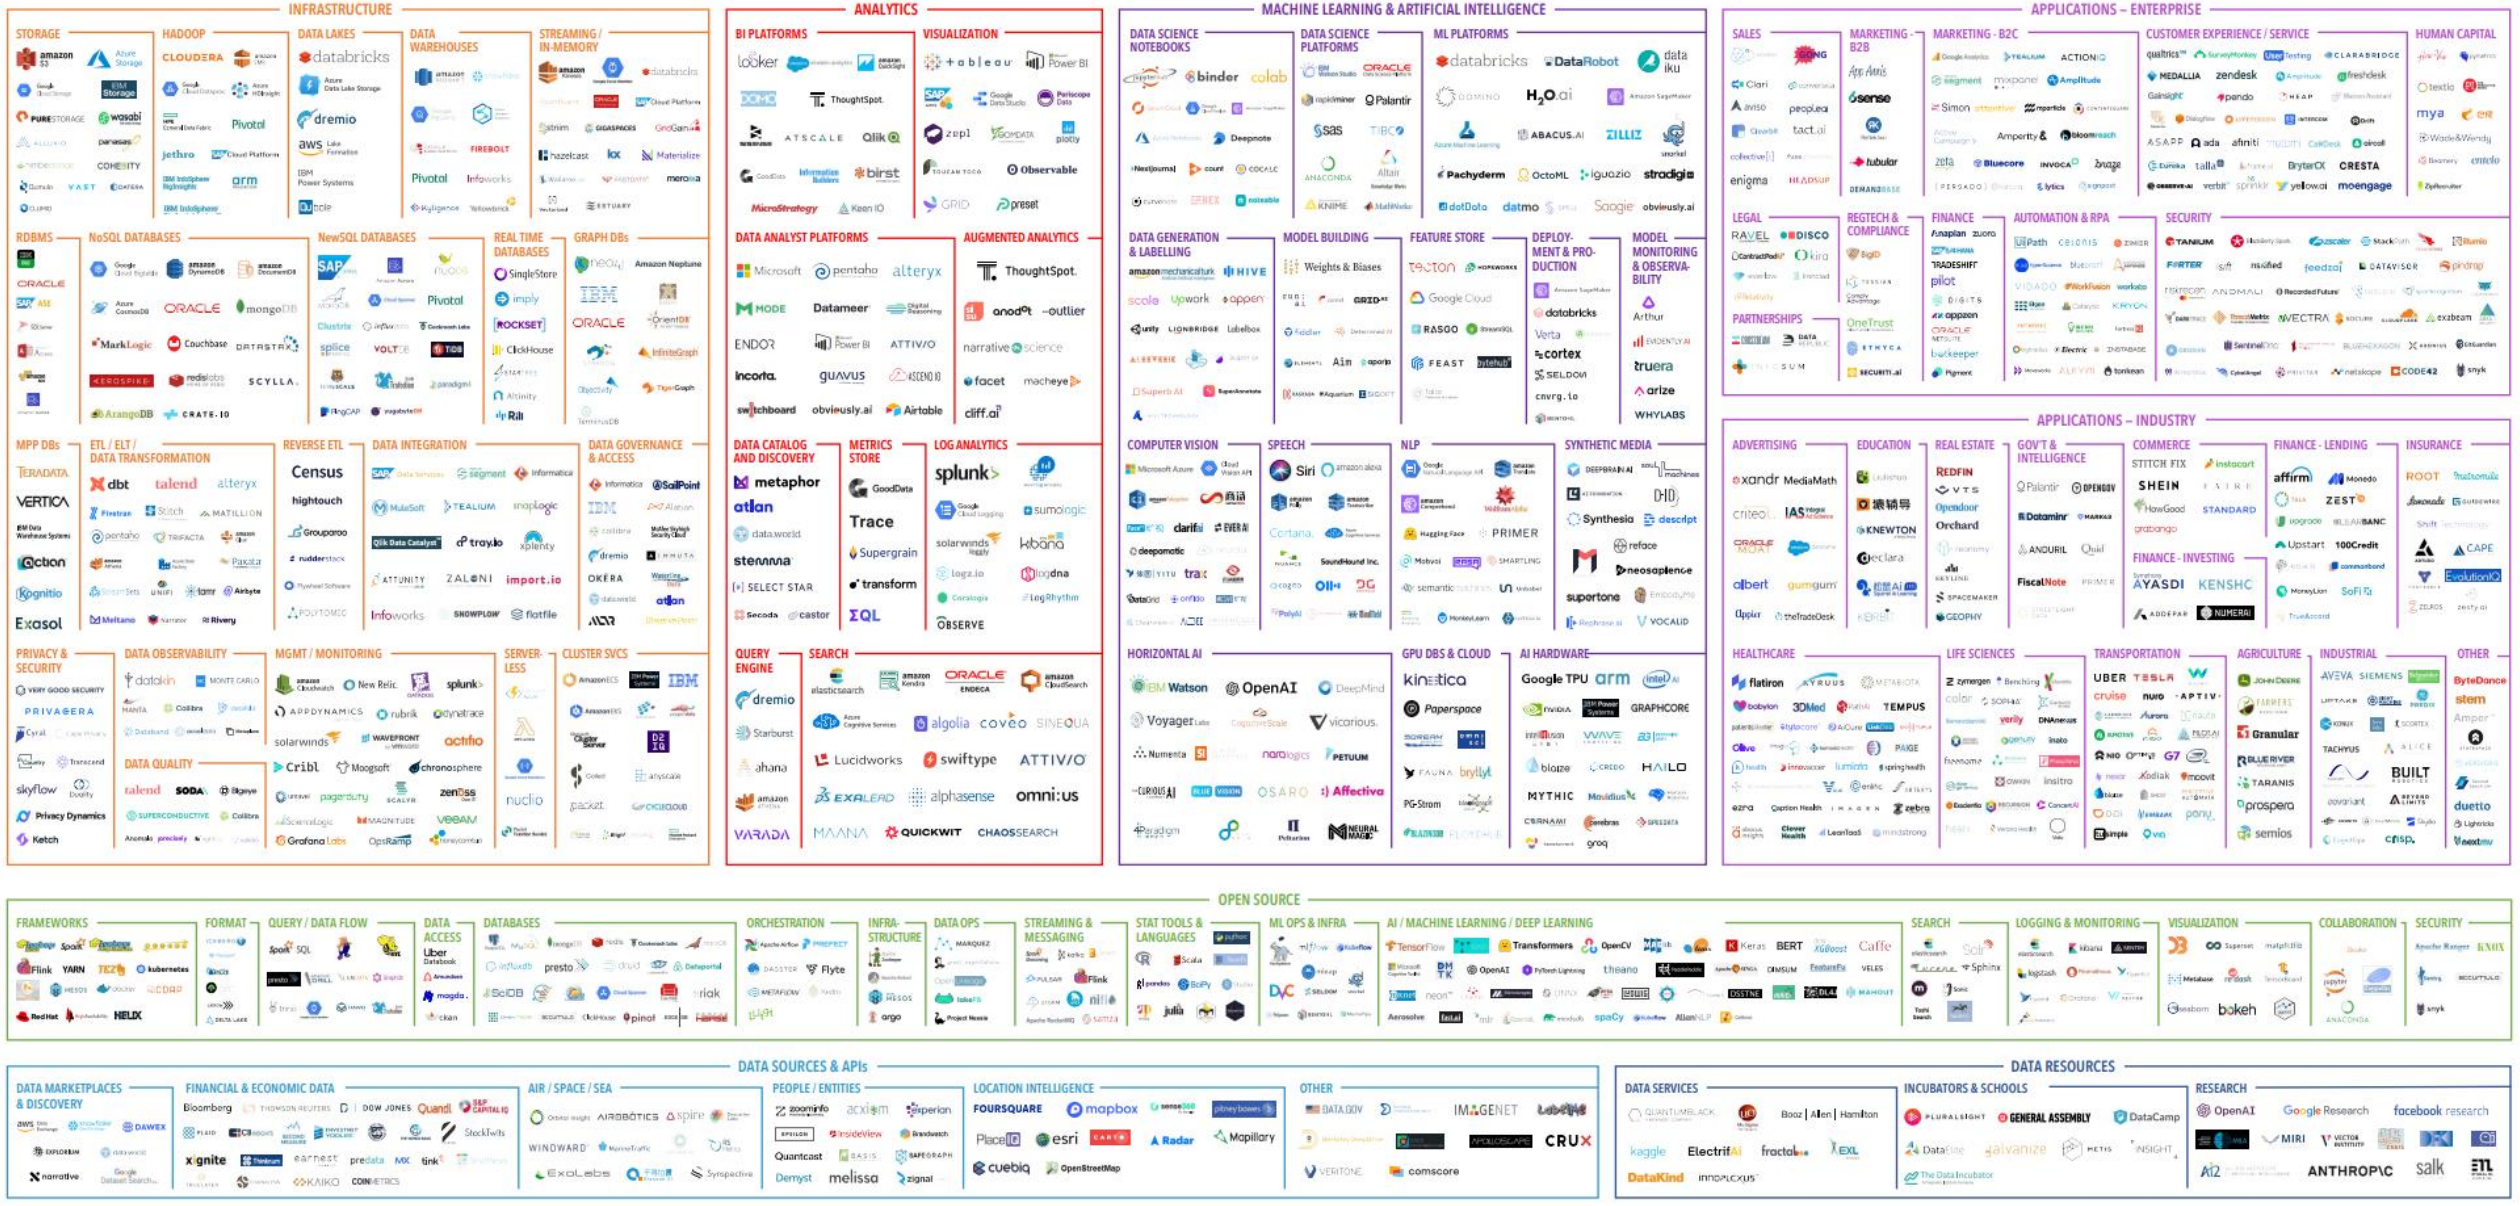

Data Management Vision

Strategic collaboration & integration for best of breed data lifecycle

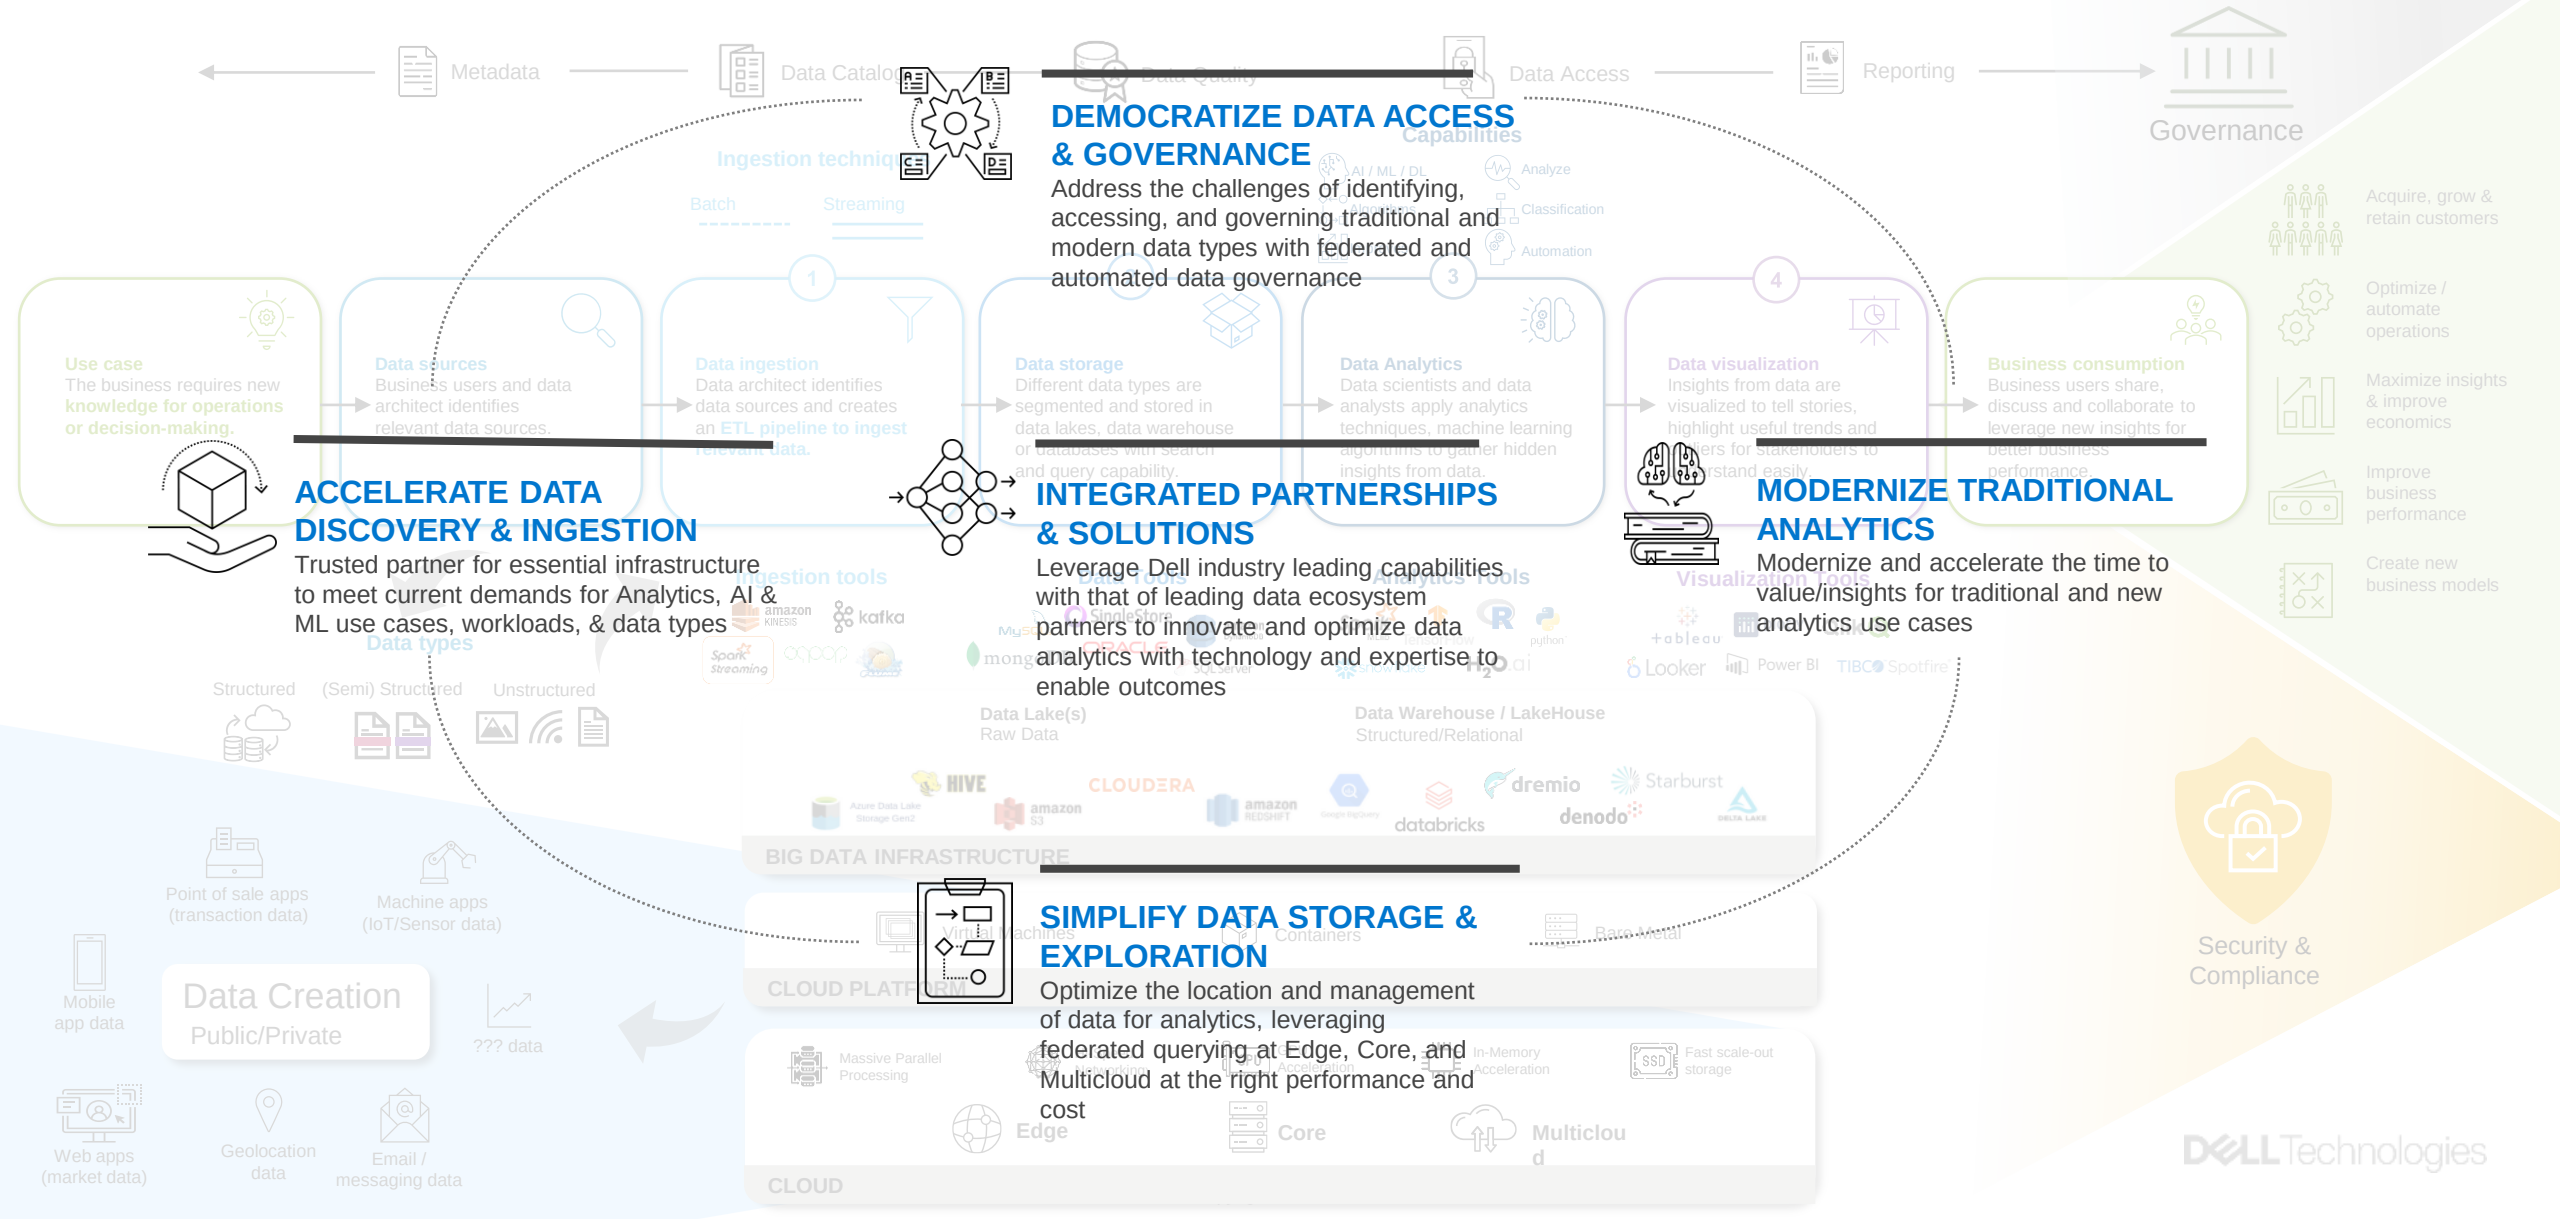

Apex Multi-Cloud Data Service & Data Analytics

Use the cloud you want, when you want

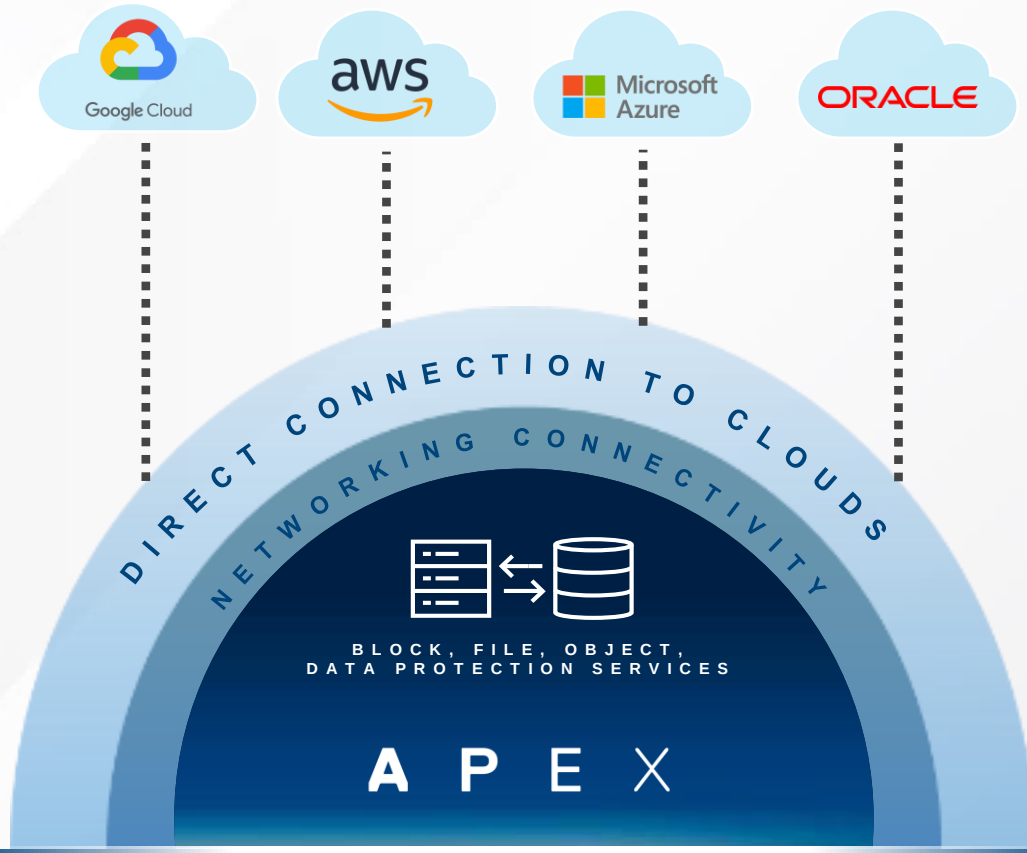

Multi-cloud agility with simultaneous public cloud access

Fully managed service available through a single console

No vendor lock-in with data independent of public cloud

Meet regulatory and compliance requirements when your data moves off-premises

No excessive egress fees

Enable Data Analytics at the Edge and Core

New Object Services and multi-factor authentication, available in US, UK, Germany and Australia

SDP EDGE

MULTI-CLOUD DATA SERVICES
+
STREAMING DATA PLATFORM

SDP CORE

Maximize Your Data Capital on Your Terms

- Consistent lifecycle management of **endpoints**
  - Spin up/down endpoints
  - Grow/shrink capacity
  - Manage code upgrades & licenses
- Integrate with **key ecosystems** like K8s and the Cloud OS's
- Unified, consistent **API control plane**

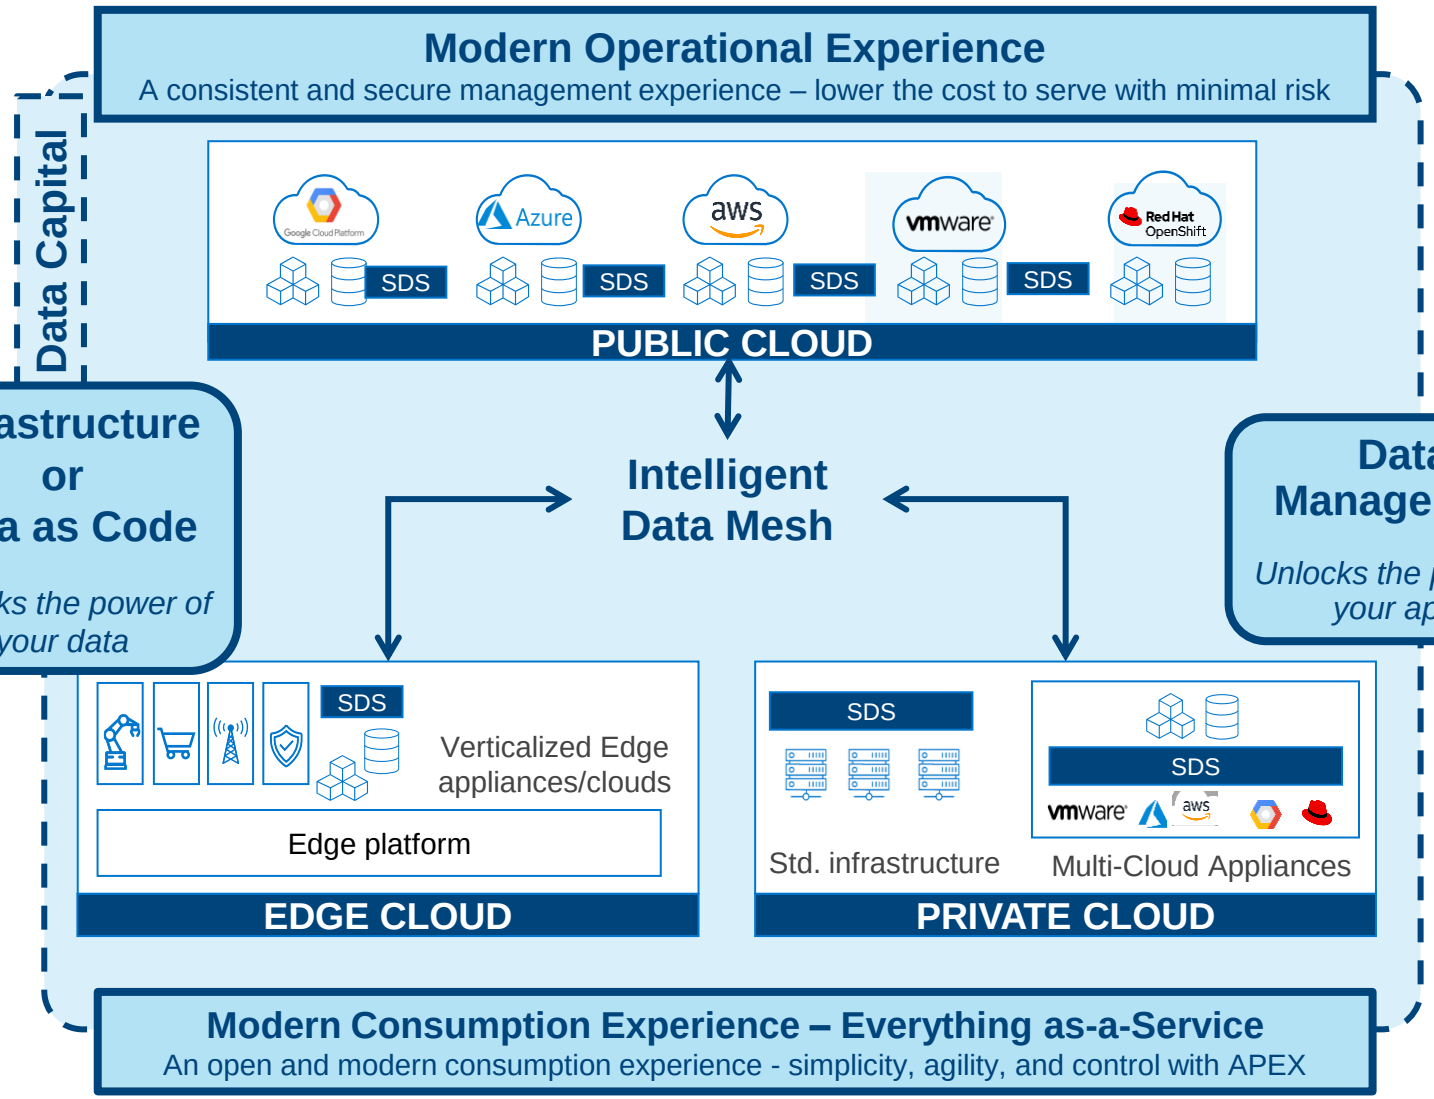

- **Uniform experience** between appliances, VE or cloud-native
- **Attribute persistence** even with movement
- Validated for **vertical specific use-cases**
- **Zero Trust** from the start

- Self-risk or **shared risk**
- **Consumption** based
- **QoS & chargeback** enabled
- **End-user portal**
- **Location independence**: edge; core; cloud

- **Data discovery, catalog, & search**
- **Real-time** increments
- Custom **tagging**
- **Policy engine**, RBAC, ABAC

Apps Data

Move data across clouds with SmartSync – at the right place with the right SLAs

REPLICATE

Low RPO async copies for disaster recovery

Replication now; failover/DR in 2023

COPY

On-demand copy to file or object for burst, test/dev or analytics

Full copy now; incremental in 2023

BACKUP

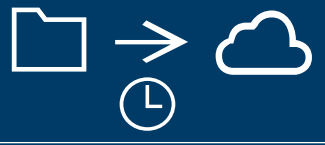

Scheduled copies of snaps to cloud for backup

Scheduled backups in 2023

TIER

Transparently archive data with option to recall

CloudPools now; replace in 2H 2023

ARCHIVE

Archive data and free up storage space

Delete from source in 2023

Data copied, tiered or archived to cloud will be directly accessible in cloud

DELL Technologies

Cloud Bursting

Quickly Serve Relevant Data to Compute Instances in the Cloud Collaboration

Key Features

- Automatically transfers from on-prem storage into the cloud
- No limit on the number of compute instances
- Minimizes data transfer costs with intelligent caching
- No manual data movement
- Provides each cloud compute instance with an easy access protocol
- Adheres to on-prem ACLs and security
- Only retrieve required data to reduce egress costs

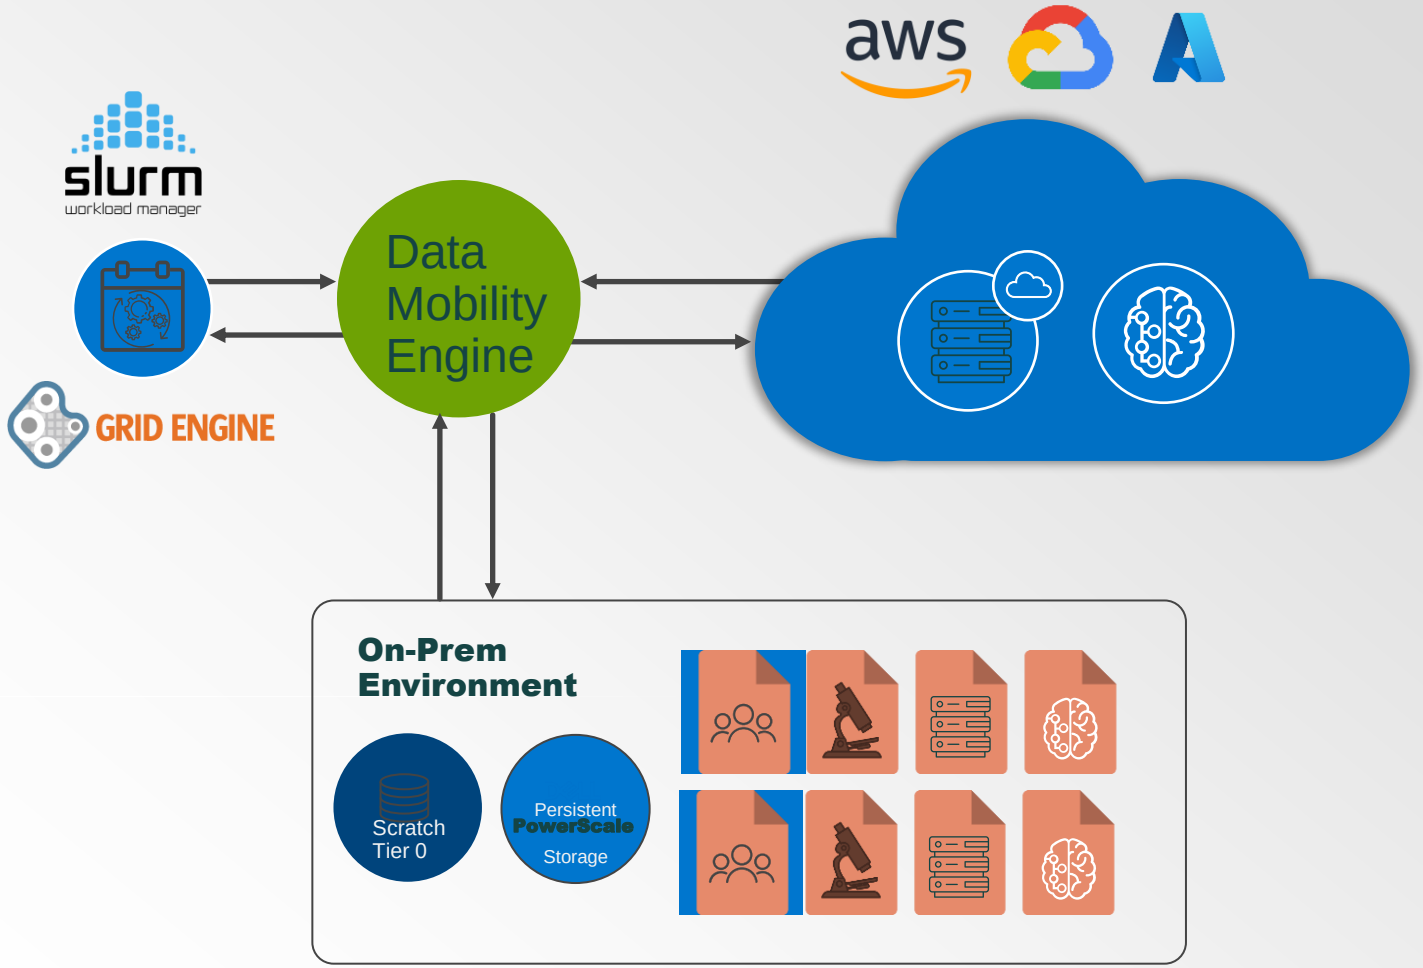

TEST ENVIRONMENT

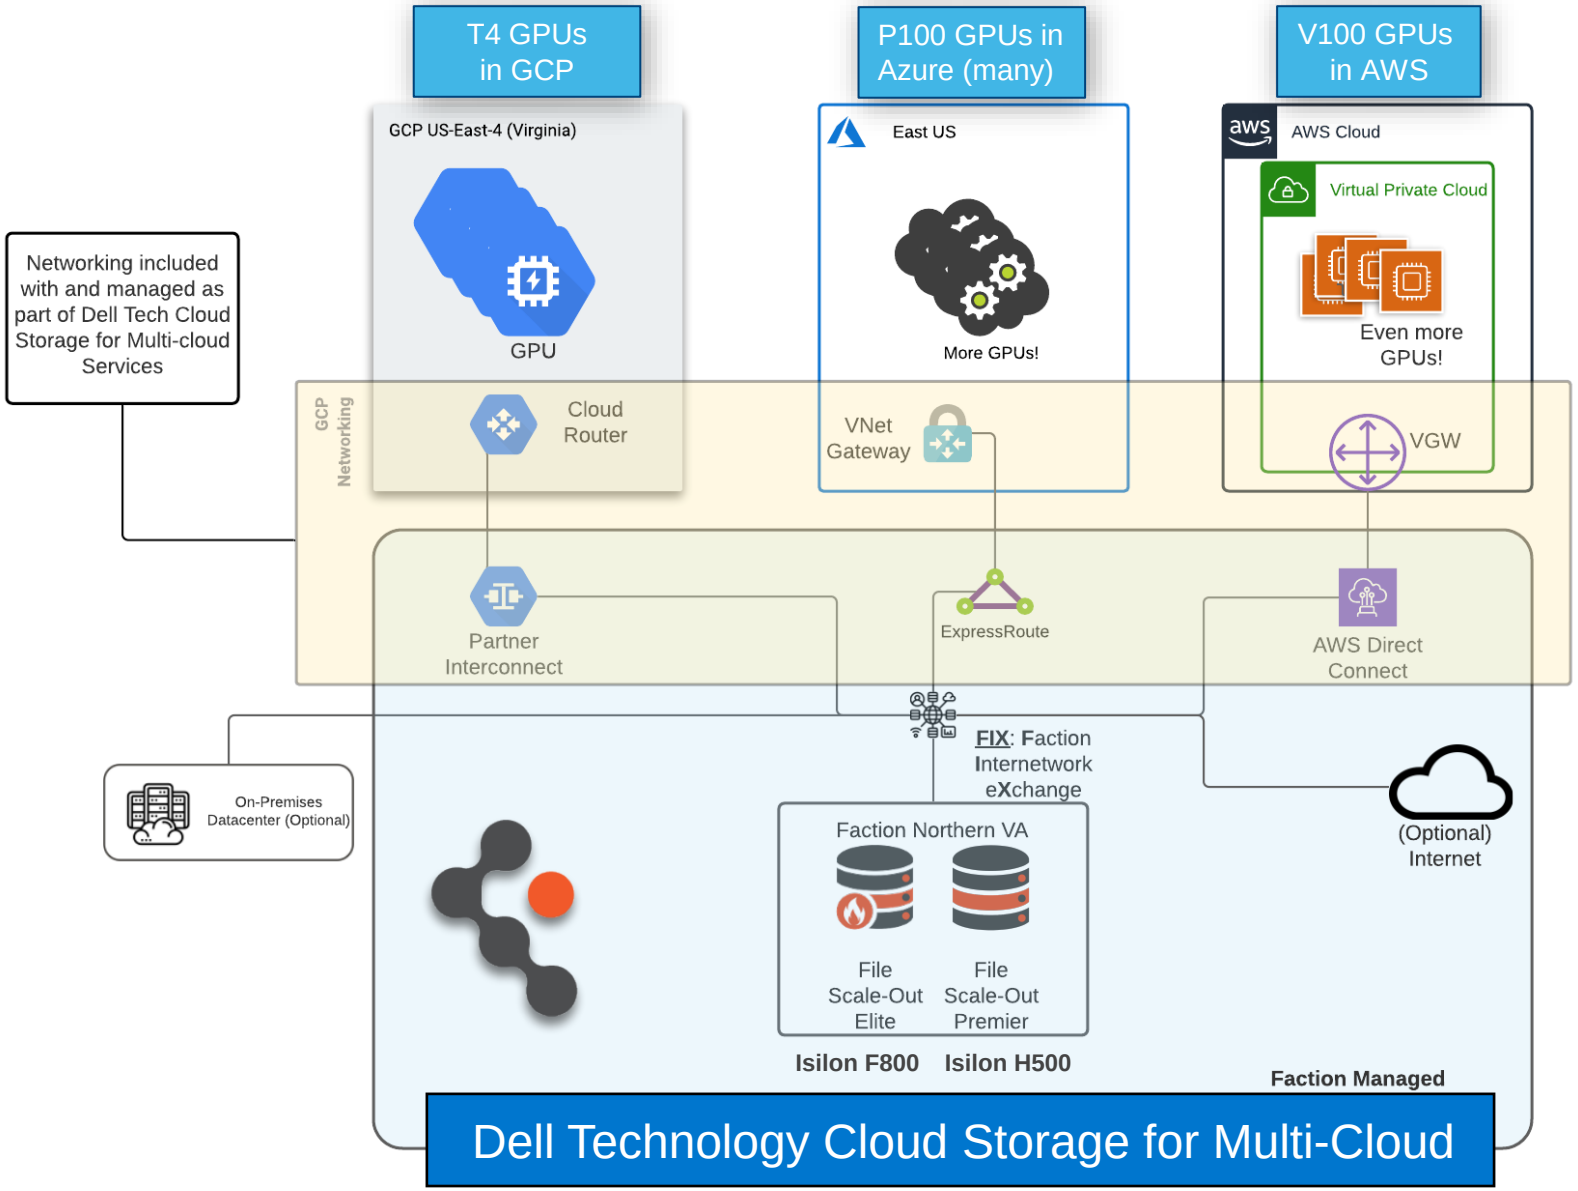

- Single Isilon-/PowerScale-powered Dell Tech Cloud Storage for Multi-Cloud data lake
- Connectivity to each cloud (included in service offering)
- Univa Grid Engine powered job execution engine (set up by Faction)
- Resources across GCP, AWS, Azure in Northern Virginia regions

THIS ISN'T THEORY

Azure

Google

aws

```
root@uge-master-new:~# qhost
HOSTNAME                ARCH         NCPU NSOC NCOR NTHR NLOAD  MEMTOT  MEMUSE  SWAPTO  SWAPUS
----------------------------------------------------------------------------------------------
global                  -               -    -    -    -     -       -       -       -       -
az-p100-w000004         lx-amd64       24    2   24   24  0.79  440.9G   45.0G     0.0     0.0
az-p100-w000005         lx-amd64       24    2   24   24  1.64  440.9G   44.4G     0.0     0.0
az-p100-w00000B         lx-amd64       24    2   24   24  1.89  440.9G  138.9G     0.0     0.0
az-p100-w00000E         lx-amd64       24    2   24   24  1.86  440.9G  271.2G     0.0     0.0
az-p100-w00000F         lx-amd64       24    2   24   24  2.04  440.9G   67.1G     0.0     0.0
az-p100-w00000G         lx-amd64       24    2   24   24  0.98  440.9G   48.0G     0.0     0.0
az-p100-w00000H         lx-amd64       24    2   24   24  1.67  440.9G   43.5G     0.0     0.0
az-p100-w00000J         lx-amd64       24    2   24   24  1.83  440.9G   47.7G     0.0     0.0
az-p100-w00000K         lx-amd64       24    2   24   24  1.48  440.9G   53.7G     0.0     0.0
az-p100-w00000L         lx-amd64       24    2   24   24  1.28  440.9G   26.1G     0.0     0.0
az-p100-w00000M         lx-amd64       24    2   24   24  0.64  440.9G   36.7G     0.0     0.0
az-p100-w00000N         lx-amd64       24    2   24   24  1.80  440.9G   43.4G     0.0     0.0
az-p100-w00000O         lx-amd64       24    2   24   24  2.33  440.9G  189.6G     0.0     0.0
az-p100-w00000P         lx-amd64       24    2   24   24  1.29  440.9G   29.3G     0.0     0.0
az-p100-w00000Q         lx-amd64       24    2   24   24  0.75  440.9G  236.3G     0.0     0.0
az-p100-w00000R         lx-amd64       24    2   24   24  2.19  440.9G  234.0G     0.0     0.0
az-p100-w00000T         lx-amd64       24    2   24   24  0.01  440.9G    1.8G     0.0     0.0
az-p100-w00000Y         lx-amd64       24    2   24   24  1.61  440.9G  234.0G     0.0     0.0
az-p100-w000014         lx-amd64       24    2   24   24  0.01  440.9G    1.8G     0.0     0.0
az-p100-w00001A         lx-amd64       24    2   24   24  2.06  440.9G  168.3G     0.0     0.0
az-p100-w00001B         lx-amd64       24    2   24   24  1.41  440.9G  227.6G     0.0     0.0
az-p100-w00001Y         lx-amd64       24    2   24   24  1.59  440.9G  184.8G     0.0     0.0
az-p100-w00001Z         lx-amd64       24    2   24   24  1.57  440.9G  115.9G     0.0     0.0
gcp-pb-workers-2nkl     lx-amd64       32    1   16   32  0.92  204.5G   32.9G     0.0     0.0
gcp-pb-workers-76tv     lx-amd64       32    1   16   32  1.06  204.5G   41.5G     0.0     0.0
gcp-pb-workers-frqr     lx-amd64       32    1   16   32  1.27  204.5G   46.8G     0.0     0.0
gcp-pb-workers-vjfl     lx-amd64       32    1   16   32  0.84  204.5G   24.7G     0.0     0.0
ip-172-31-32-118        lx-amd64       32    1   16   32  0.16  239.9G   29.6G     0.0     0.0
ip-172-31-39-100        lx-amd64       32    1   16   32  1.37  239.9G   36.6G     0.0     0.0
ip-172-31-44-41         lx-amd64       32    1   16   32  1.40  239.9G   41.0G     0.0     0.0
ip-172-31-45-227        lx-amd64       32    1   16   32  1.15  239.9G   67.0G     0.0     0.0
ip-172-31-66-202        lx-amd64       32    1   16   32  1.42  239.9G   40.6G     0.0     0.0
uge-master-new          lx-amd64        4    1    4    4  0.00   31.3G    1.3G     0.0     0.0
```

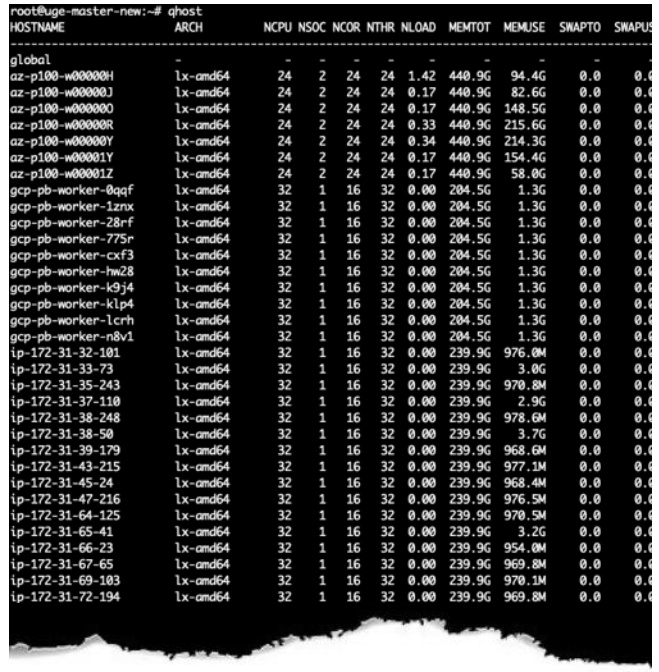

Peak cluster size > 1 Million CUDA cores

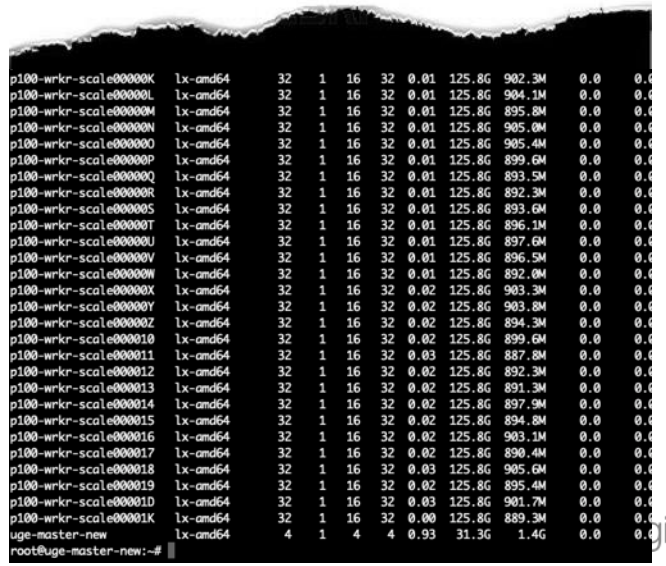

SCALE – GET TO 2Tb/S+ ACROSS MULTIPLE CLOUDS

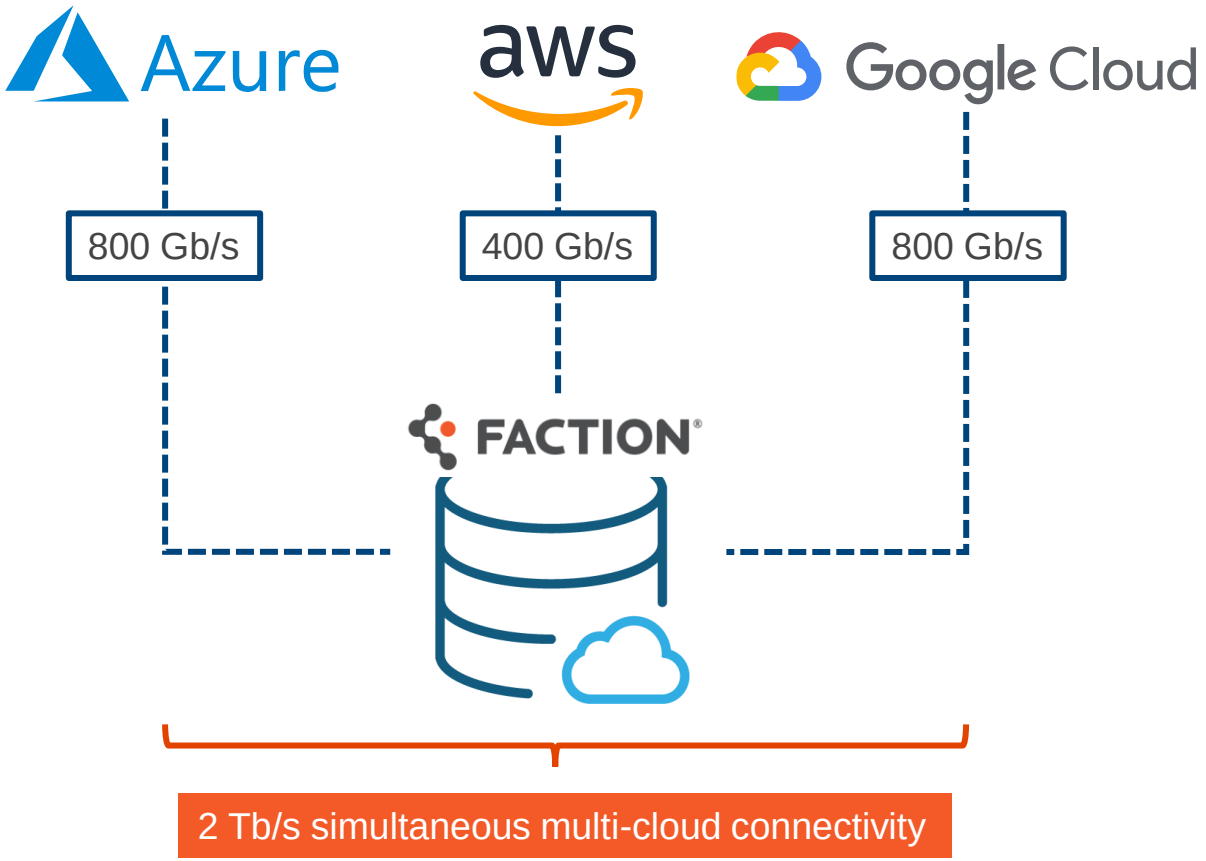

Dell & Kalray ngenea®

ngenea® Unifies Dell Storage Tiers

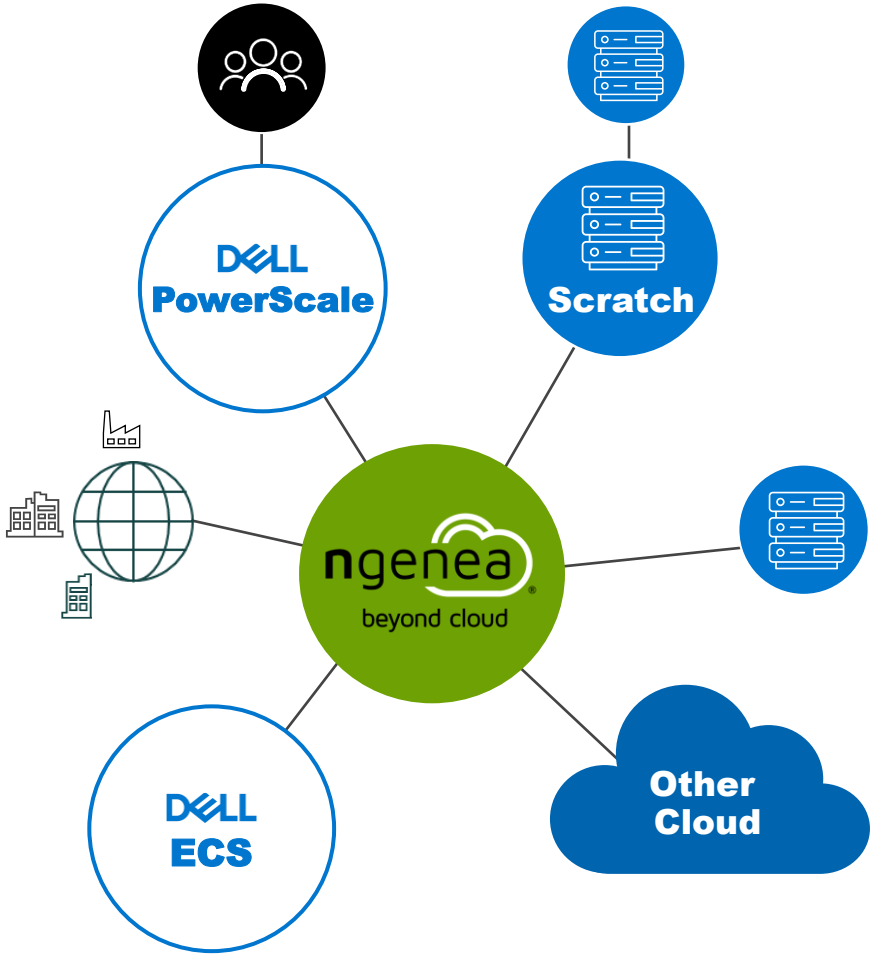

ngenea® is a Software-Defined Data Management Solution That Uses Qualified Building Blocks

- Dell & Kalray pixstor™ Tier-0
  - PowerEdge NVMe Servers
- Dell & Kalray pixstor™ Tier-1 Plus
  - PowerVault Storage Arrays
  - PowerScale
- Archive
  - ECS Enterprise
- Third-Party and Public Cloud

Dell’s all-flash platforms for tick data analytics

Supported by STAC benchmark results

Highest Scale and Performance

For the entire data set

F900

- F900
  - KDB220506 Antuco + Kanaga Jun. 11, 2022
  - KDB210929 Antuco Oct. 18, 2021
- F800
  - KDB190430 Antuco + Kanaga Apr. 30, 2019
- F200
  - KDB200914 Antuco Sep. 20, 2020

High Performance Edge

For data subsets

F200

Demonstrating real-time performance on smaller data sets (<10TB) and near real-time performance on large data sets (>10 TB) at high concurrency (100s into the millions).

DELL Technologies

"The Lab is staffed by engineers with advanced degrees and many years of industry experience in domains such as mechanical engineering and bioinformatics. We also have engineers with computer science backgrounds, providing expertise in file systems, interconnects and HPC management tools."

—Onur Celebioglu, Sr. Distinguished Engineer, ISG CTIO, HPC and Emerging Workloads

HPC & AI Innovation Lab

Dedicated to designing solutions while staying on the leading edge of new and emerging technologies

You've got the power

High performance computing (HPC) gives you the power to break new ground, make important discoveries, and solve some of the most important challenges of our time. But there are always bigger questions — and bigger data sets — on the horizon, requiring HPC solutions to keep pace with the speed of innovation.

That's why Dell Technologies is committed to enabling more organizations in industry, research and government to use HPC solutions for more innovations and discoveries than any other HPC systems vendor in the world. This passion for innovation has helped make Dell Technologies an industry leader in HPC clusters, storage, networking and software. We've built a nexus of collaboration in the industry, exemplified by the Dell Technologies HPC & AI Innovation Lab.

Working with the HPC community to go further, faster

The Dell Technologies HPC & AI Innovation Lab encompasses a 13,000-square-foot data center devoted to high-performance computing and artificial intelligence (AI). It houses thousands of servers, a wide range of storage and network systems.

But the Lab is more than world-class infrastructure. Bringing together HPC operational excellence and expertise, it is staffed by a dedicated group of computer scientists, engineers and subject matter experts who actively partner and collaborate with customers and other members of the HPC community. The team gets and provides early access to new technologies, integrates and tunes clusters, benchmarks applications, develops best practices, and publishes their results.

When you engage with the Lab, you work directly with these experts to design a solution for your unique HPC workloads. The opportunity to develop and test your configuration with an expert team prior to deployment reduces risk, and because your HPC system is tuned for optimized performance from day one, your team can get to results faster. And that means your organization can recognize a better return on investment.

"The HPC & AI Innovation Lab gives our customers access to cutting-edge technology from Dell, Intel, AMD, NVIDIA, Bright Computing and more. Customers can bring us their workloads and we can help them tune a solution before the technology is readily available."

—Garima Kochhar, Distinguished Engineer

DELL Technologies