

Pure Storage (including Portworx) Fast Facts



Growth

\$2.75 Billion

FY23 Annual Revenue



Customers

12,000+

Global Customers



Leadership



**Gartner Magic
Quadrant**

10 x leader!

\$1.3 Billion

Q3 FY24 Subscription Annual Recurring Revenue (ARR)

+\$60 Billion TAM

353

New Customers
in Q3 FY24

~60%

Of Fortune
500 companies

82

**Net Promoter
Score**

Highest in the industry



Data Storage
as-a-Service



Data
management



Cloud mobility



Kubernetes native
data



Database as-
a-Service



What you see and hear about AI



ChatGPT



Mid Journey



Llama2
(Meta)



Data selection, curation, cleaning

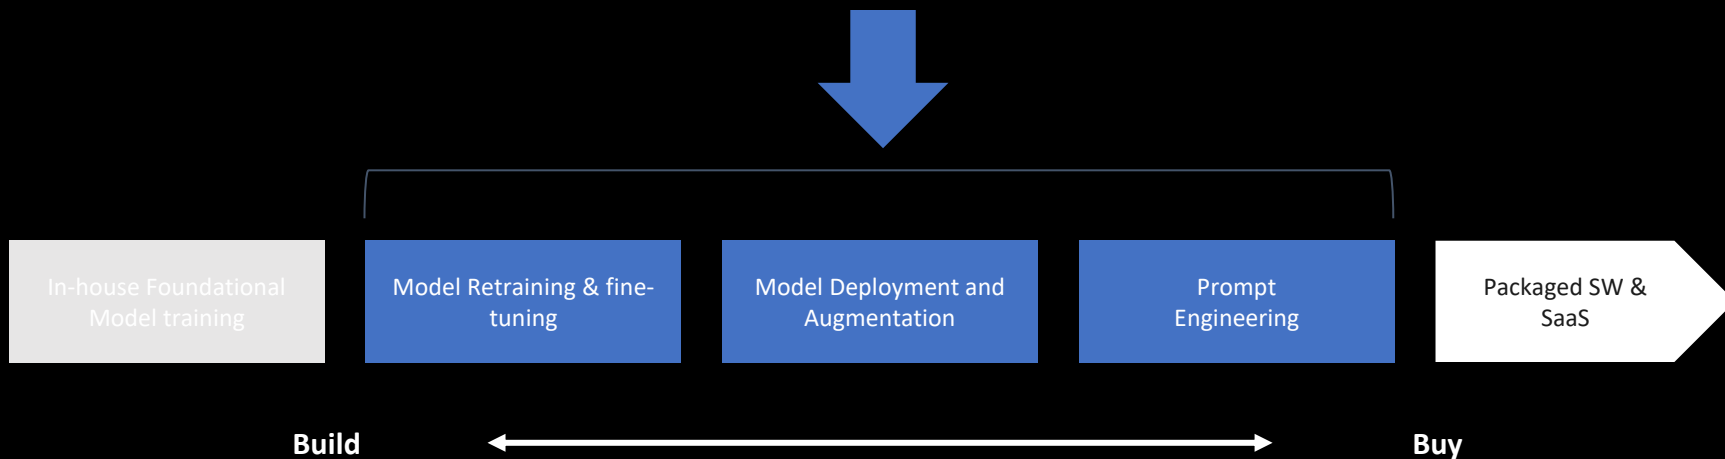
Data security and governance

Data integration

Data infrastructure

What we don't
talk (enough)
about...

The spectrum of AI usage



Powering Meta's Research SuperCluster



*"Meta chose Pure as it needed a storage partner that can deliver **robust and scalable storage capabilities** to power the **Research SuperCluster**."*






*"With Pure, RSC will have **unparalleled performance** to rapidly analyze both structured and unstructured data, underpinned by Pure's foundation of **simplicity, reliability, and sustainability**."*

Infrastructure impact

What did we learn?

- **Power efficiency is a key metric – Every kW matters**
- **GPU servers now require > 10kW**
- **Even the most modern datacentres are built for 20-25kW per rack, this means a lot of empty space in DCs**
- **High performance networks and storage are required to ensure the GPUs are not idling**
- **Petabyte scale datasets are common in AI use cases, some tending to Exabytes**


Wide range of needs in AI pipelines

	 kafka INGEST	 CLEAN & TRANSFORM	 EXPLORE	 TRAIN	 INFER
ACCESS PATTERN	SEQUENTIAL	SEQUENTIAL OR RANDOM	RANDOM	SEQUENTIAL OR RANDOM	SEQUENTIAL
ACCESS TYPE	WRITE	READ & WRITE	READ	READ & WRITE	READ & WRITE
OBJECT SIZE	SMALL TO LARGE	SMALL TO LARGE	SMALL TO LARGE	SMALL TO LARGE	SMALL TO LARGE
CONCURRENCY	DEPENDS ON # SOURCES	HIGH	LOW	HIGH	HIGH


New data access needs



No cold data anymore



No silos

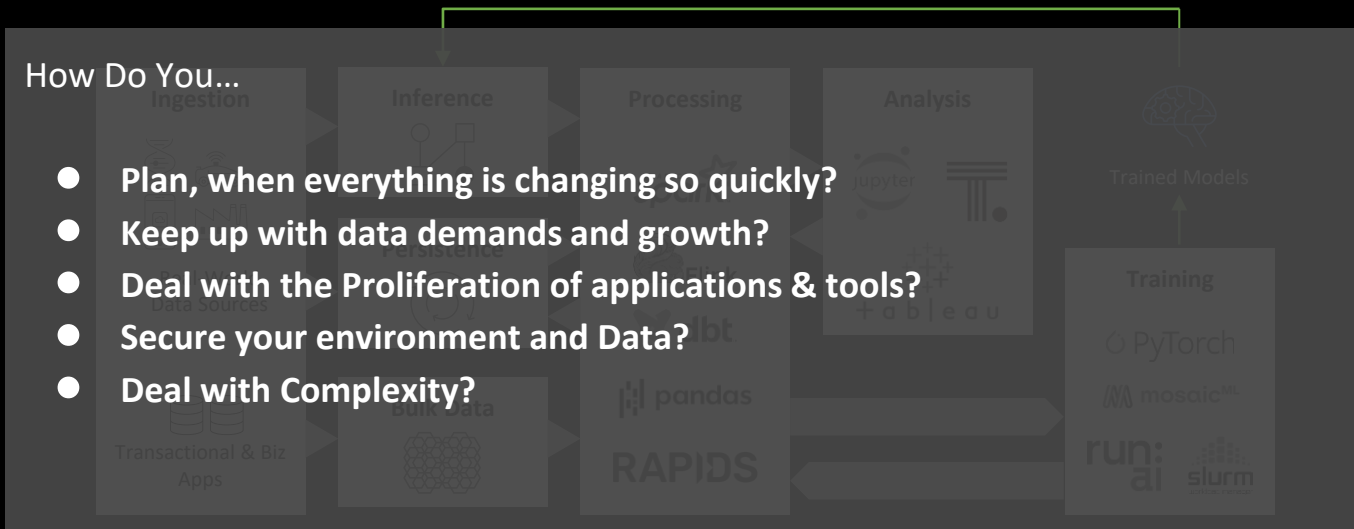


Variable, simultaneous
data access patterns




Checkpointing support

Consistent Challenges

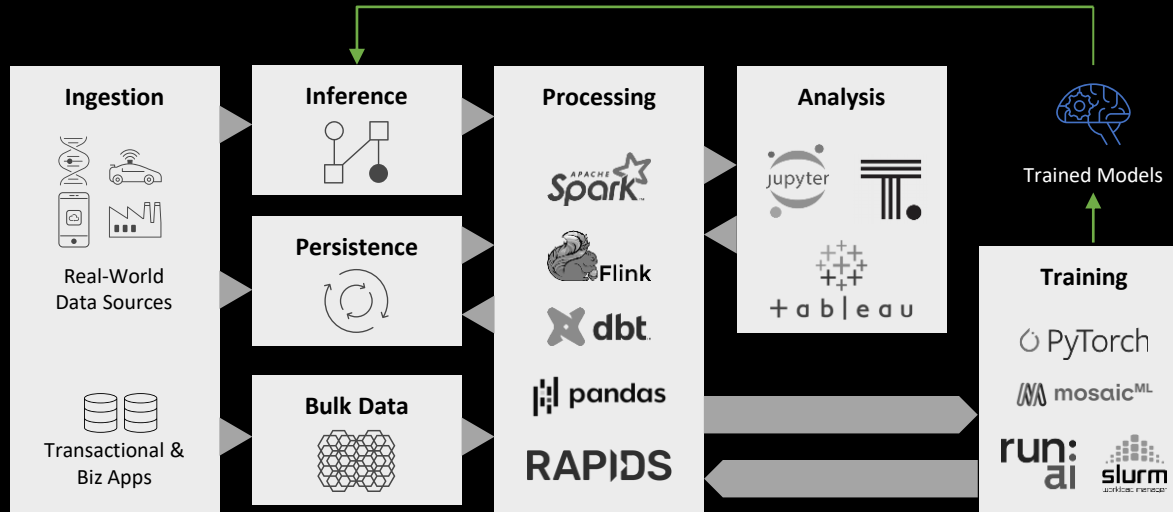


Anti-Pattern: Assembling a collection of specialized solutions

Time to science



Larger AI Journey




Training is just a piece of a larger & more complex picture

Serving data scientists

Data Science as-a-Service should enable:



Faster access to resources
Tools, environment and data



Faster results
Explore, train and evaluate



Self-Service
On-demand and programmatically

Best practices

1

Recognise the new data access requirements

2

Design for agility and flexibility

3

Incorporate flexible consumption models

4

Treat Data Scientists like developers and give them their own as-a-Service platforms


5

Optimise power usage everywhere you can. Efficient data systems mean more power for compute tasks



Pure Can Help You Move Quickly, Today

One Consistent offering for all AI Needs



Enterprise Container Storage
& Self-service DBs



Storage abstraction & orchestration
Policy-based management



Transactional &
High-Performance



Content Stores, Bulk
and Archive



Cloud Storage
Automation

On-demand consumption
SLA-based



Cloud service experience
Fully flexible

Move with **Speed** and
Agility at **Scale**



