REGIONAL SDC 25
BY Developers FOR Developers

Regional SDC Denver
April 30, 2025

Storage for AI 101

An overview of AI Workloads from a Storage Perspective

Curtis Ballard
SNIA Technical Council, Co-Chair AI Task Force

SNIA

Speaker

Or ?

CURTIS BALLARD

SNIA Technical Council

Curtis Ballard

Hewlett Packard Enterprise

Bio available at: SNIA 2024/2025 Technical Council

REGIONAL SDC

What this presentation IS

High level
introduction to storage for AI

Foundation
for other AI presentations

Thinking Food
for planning storage for AI

REGIONAL SDC

What Kind of Storage?

- File storage?
- Object storage?
- Block storage?
- Direct attached storage?
- External (SAN/NAS) storage?

- **It Depends**
  - This presentation will cover general characteristics of AI workloads, from the perspective of the “storage”.
  - The specific storage implementation for any given workload has lots of choice points that need deeper analysis.

Storage, A Key Part of a Solid AI Foundation

77% of companies with AI in production plan to upgrade their storage! *

81% of companies evaluating AI plan to upgrade their storage. *

* IDC, Future Proofing Hardware for Artificial Intelligence, August 2022

REGIONAL SDC

Why is storage for AI different?

Requires huge quantities of data	• With different quantities in different phases
Extensive Data Manipulation	• With different types of data manipulation in different phases
It is a multi-phase workload	• Most traditional workloads, like databases, have predictable access patterns • AI has widely different workload patterns for different phases
Performance and capacity varies widely for different AI tasks	• Overview of these variations follows in this presentation
Highly Parallel Operations	• Usually multiple parallel operations with their own storage workloads

REGIONAL SDC

Food for AI

AI

- It runs on GPUs and CPUs

but

- Eats Data!

and that data

- Eats Storage!

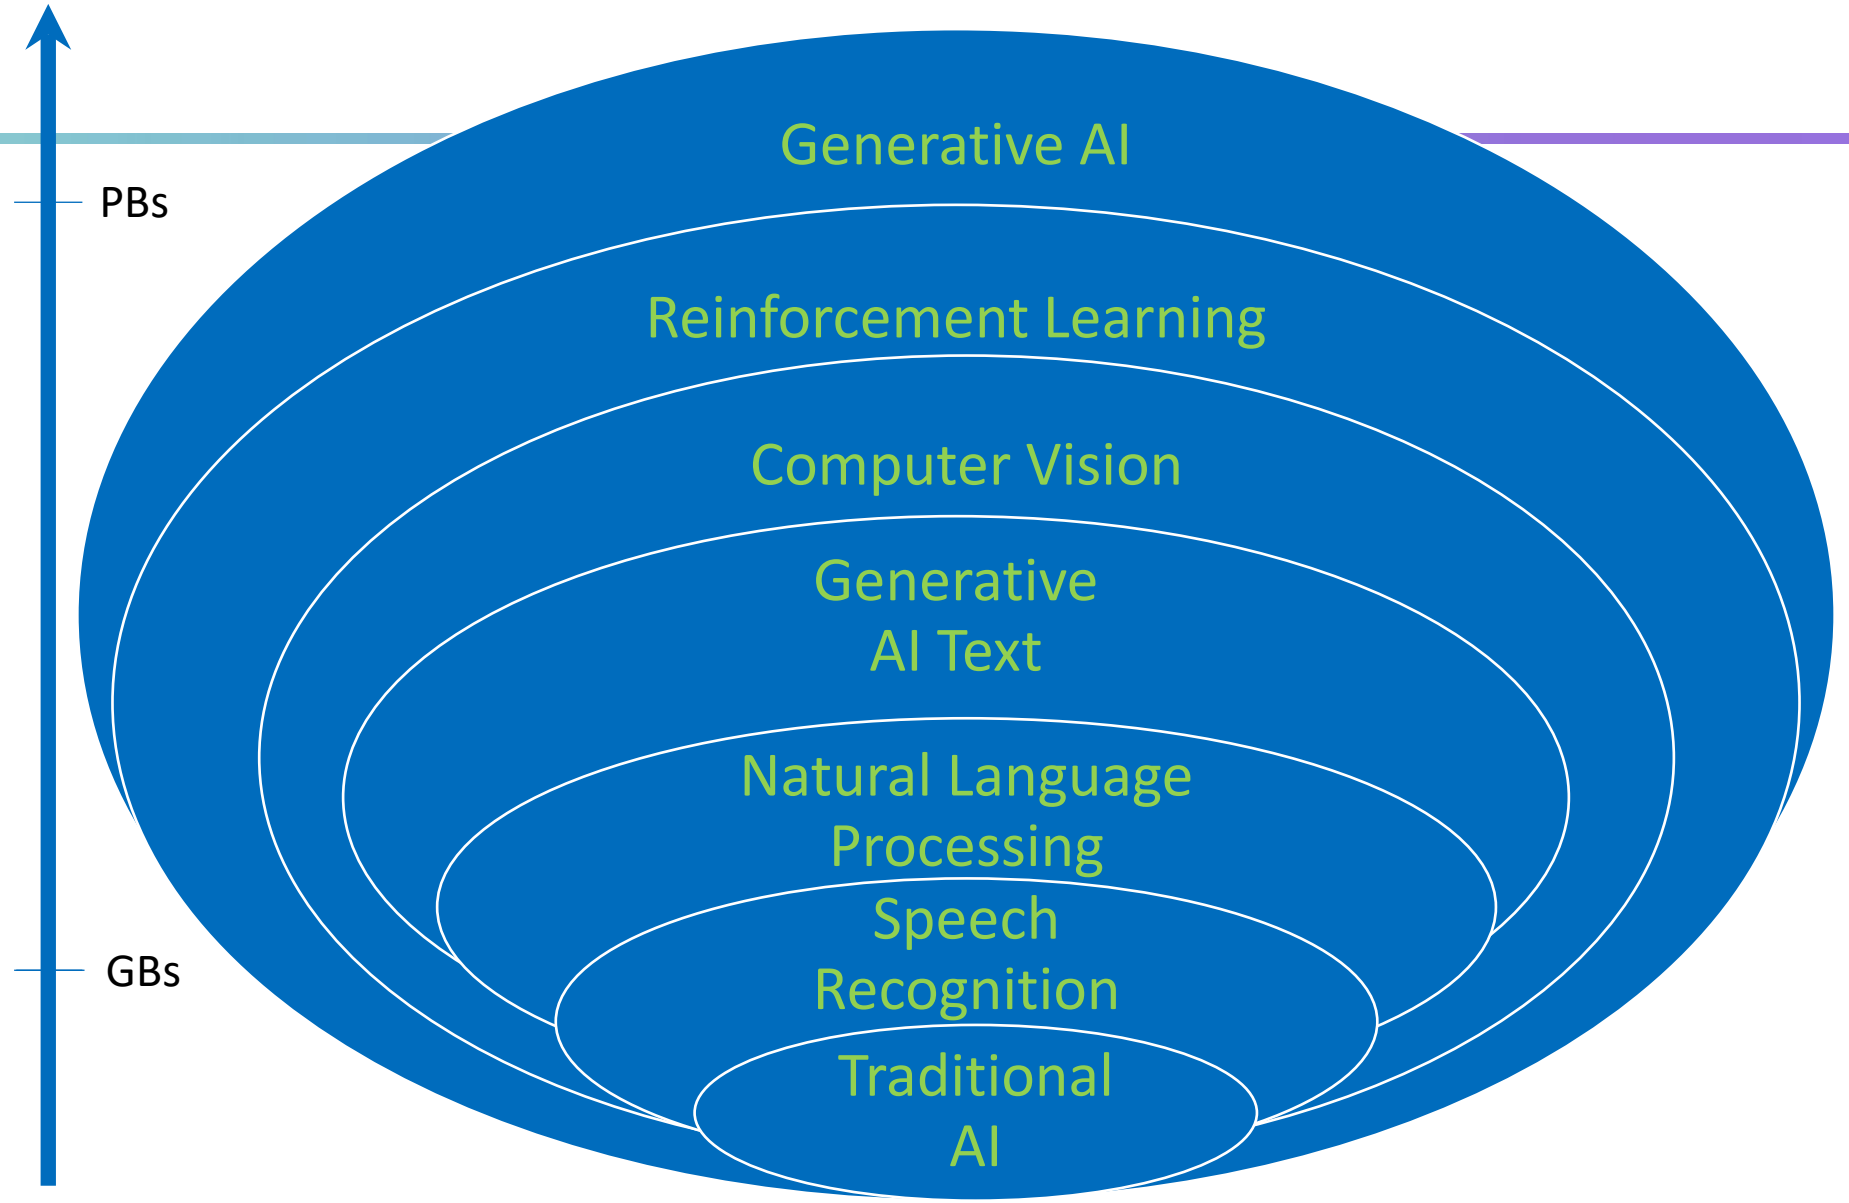

REGIONAL SDC

The Two Sides of AI: Training and Inferencing

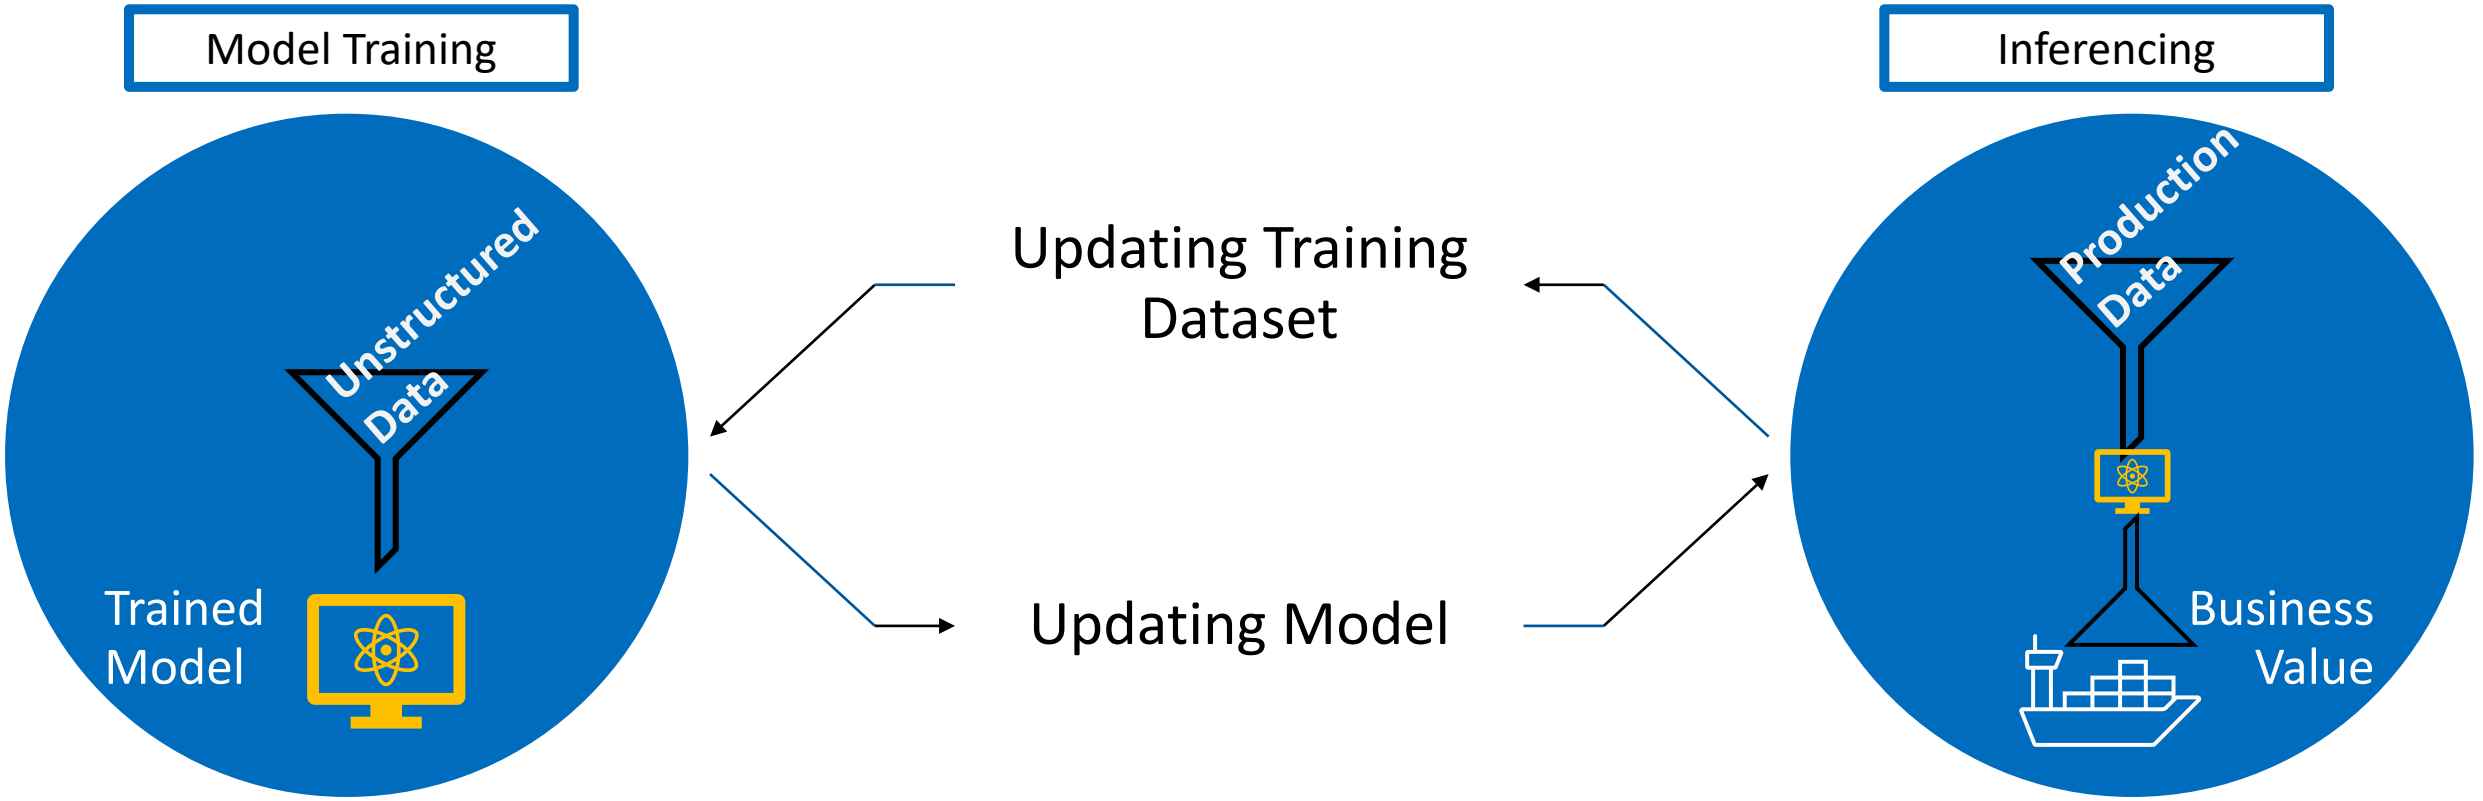

Several phases behind these

REGIONAL SDC

Storage Phases of AI one perspective

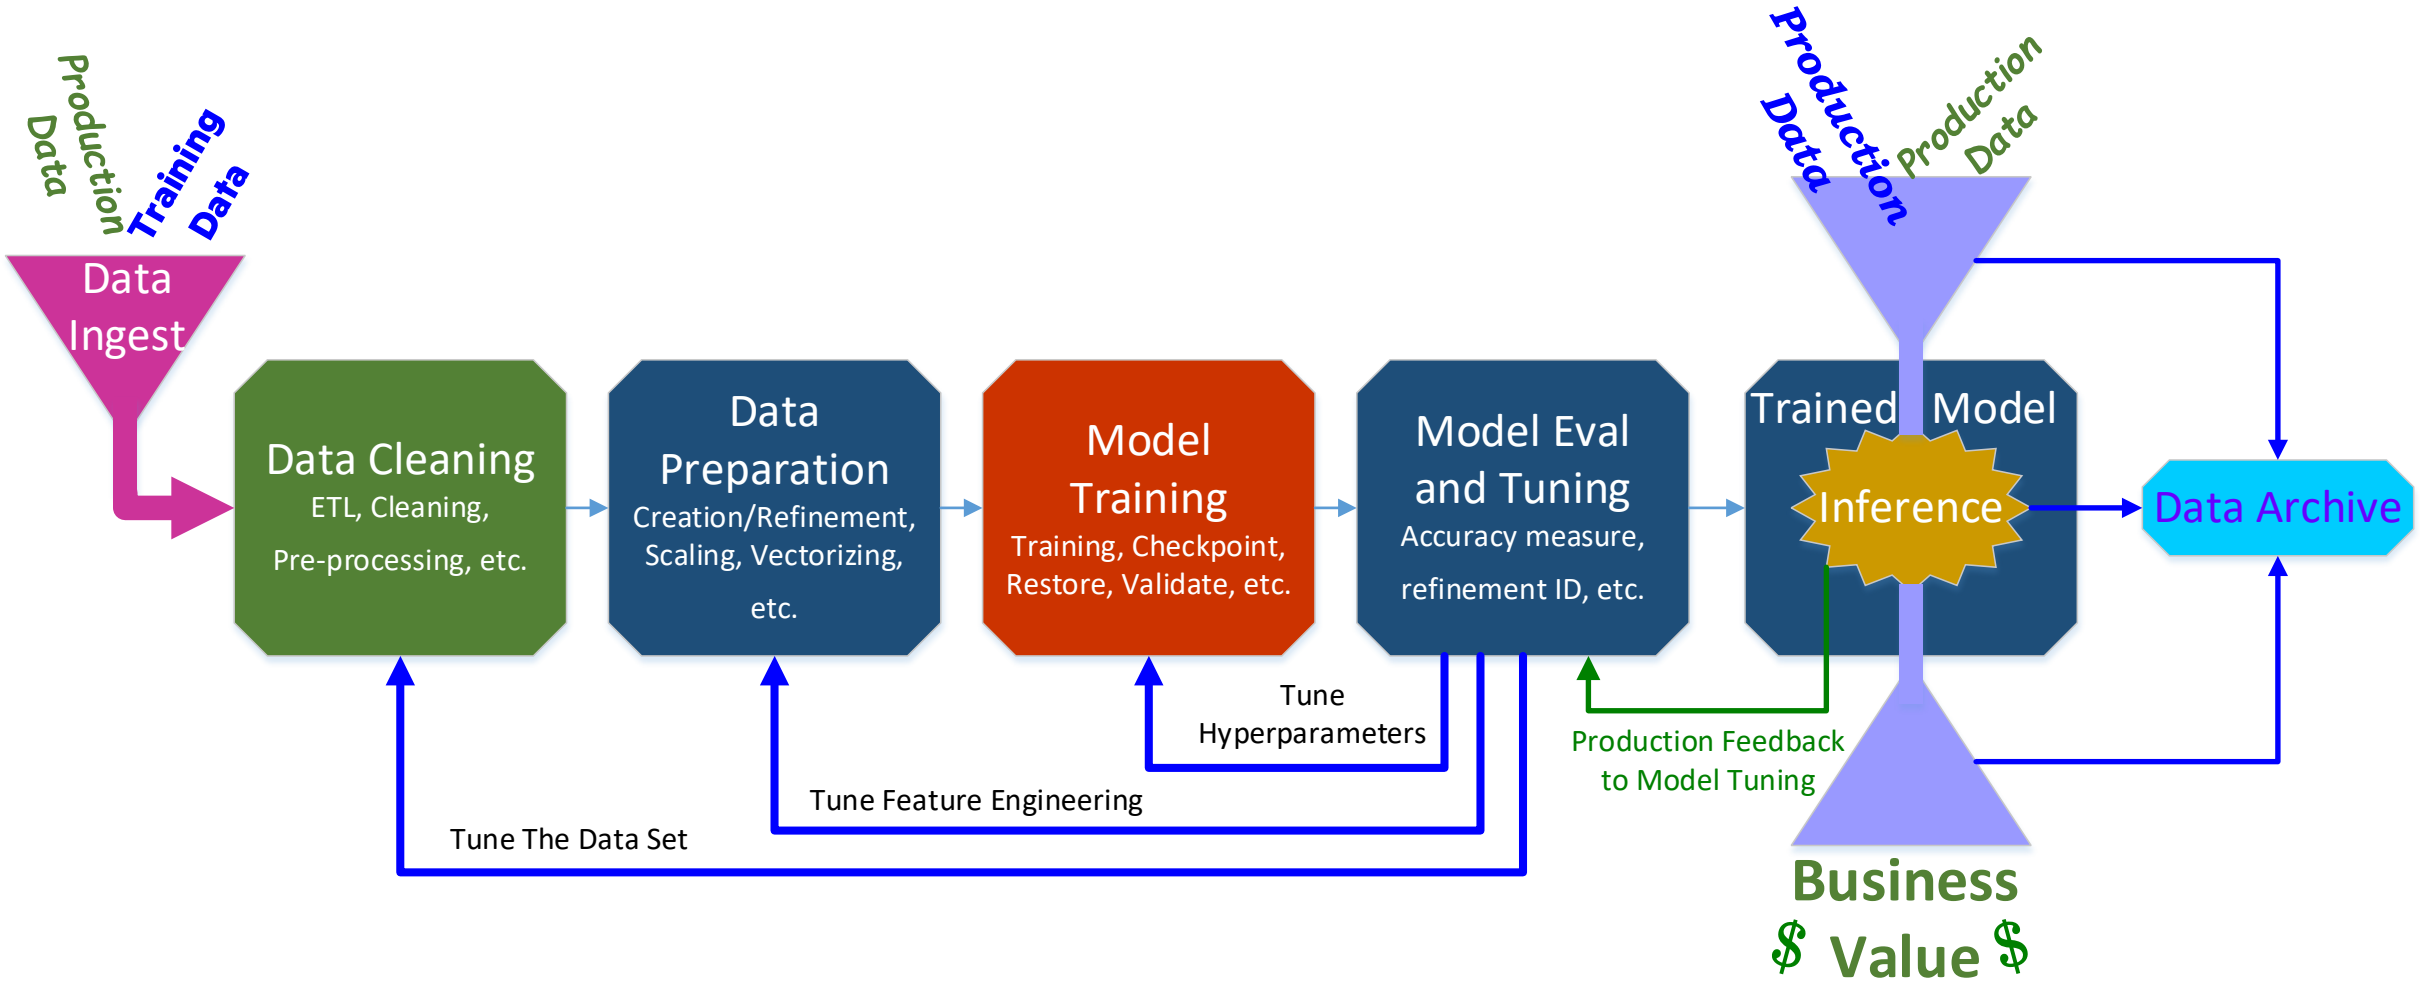

REGIONAL SDC

AI Affects All of Your Storage

Think about your needs for today and tomorrow

How does using AI change your storage requirements?

Example: Data Ingest

- Your business processes generate data today
- You already have storage for data ingest
  - *or do you?*

- Business data is already being captured, ***But:***
  - How does AI affect what you capture
  - How does AI affect how you store your business data
  - How does AI affect how you access your business data

REGIONAL SDC

One Example Company

What they had

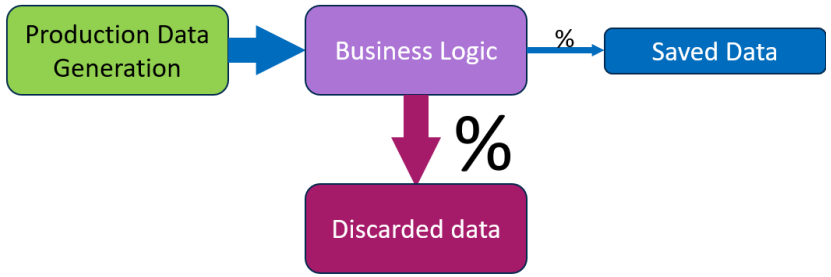

What they need

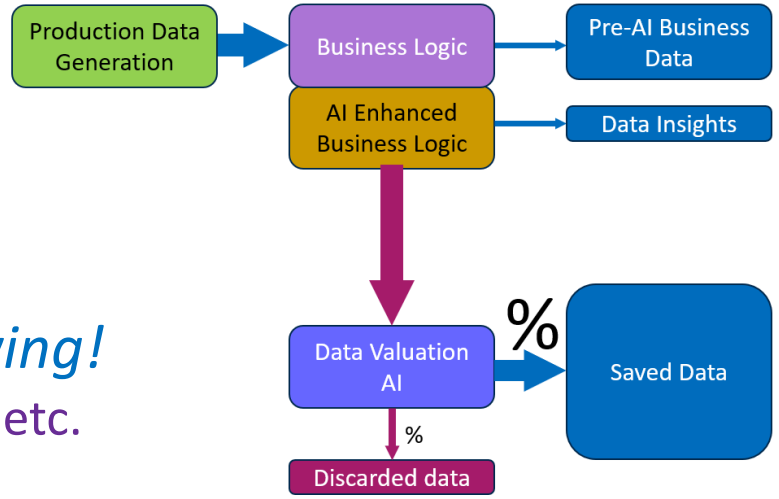

AI unlocked value in data that they weren't saving!

- Logs of data like mfg sensors, customer interactions, etc.
- Valuable business insights were hiding in sea of data

- AI PoC's demonstrated that they needed to save more data
- Their existing storage capacity was far too small!

REGIONAL SDC

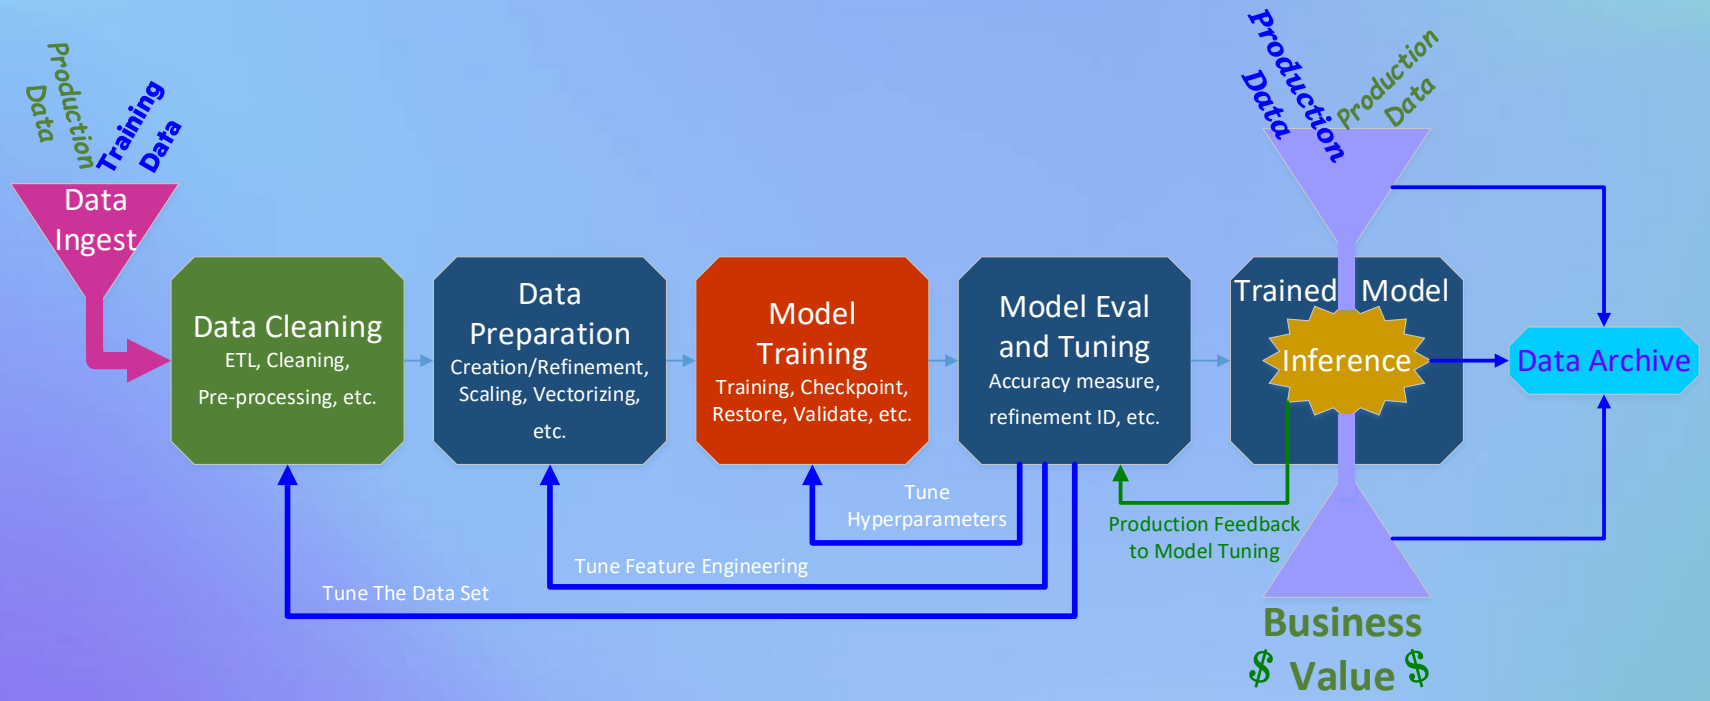

So what do the storage requirements look like for these Storage for AI phases?

Data Aggregation

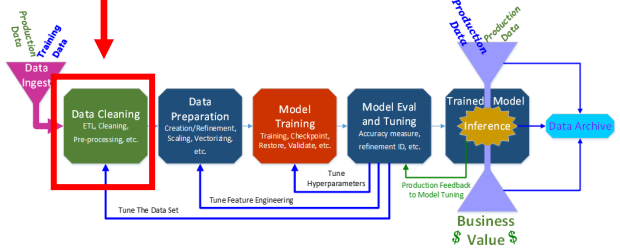

- Raw source data has to be prepared for use in AI
  - Logs, pictures, video, documents, etc.
- Data needs organized before becoming training data
  - Clean out noise
  - De-duplicate
  - Normalize
  - Privacy and Ethical processing, (anonymizing PII, removing bias, etc)
- Data is read from the ingest storage
- Cleaned data needs written to storage for data preparation
- Process may be able to be partially automated applying AI

Sequential Random

REGIONAL SDC

Data Preparation (e.g., feature engineering)

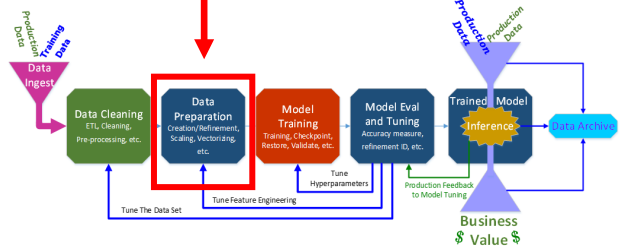

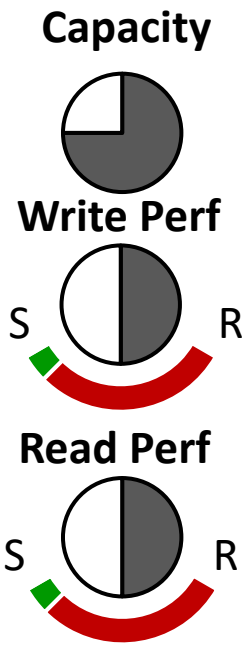

- Data Scientists serve as translators
  - Raw data → Food for AI (Numbers)
- Exploring the data – identifying patterns, outliers, relationships, etc.
- Splitting data for training and testing
- Feature extraction – converting key features into consumable nuggets
- Data transformation – converting data types (Vectorizing)
- Often highly parallel

Sequential Random

REGIONAL SDC

Model Training – Storage Interactions

- Loading memory for training
- Checkpointing
  - A point-in-time backup of the model copied off for recovery
  - GPU paused while model state is copied out of GPU Memory
  - Checkpointing may be synchronous or asynchronous

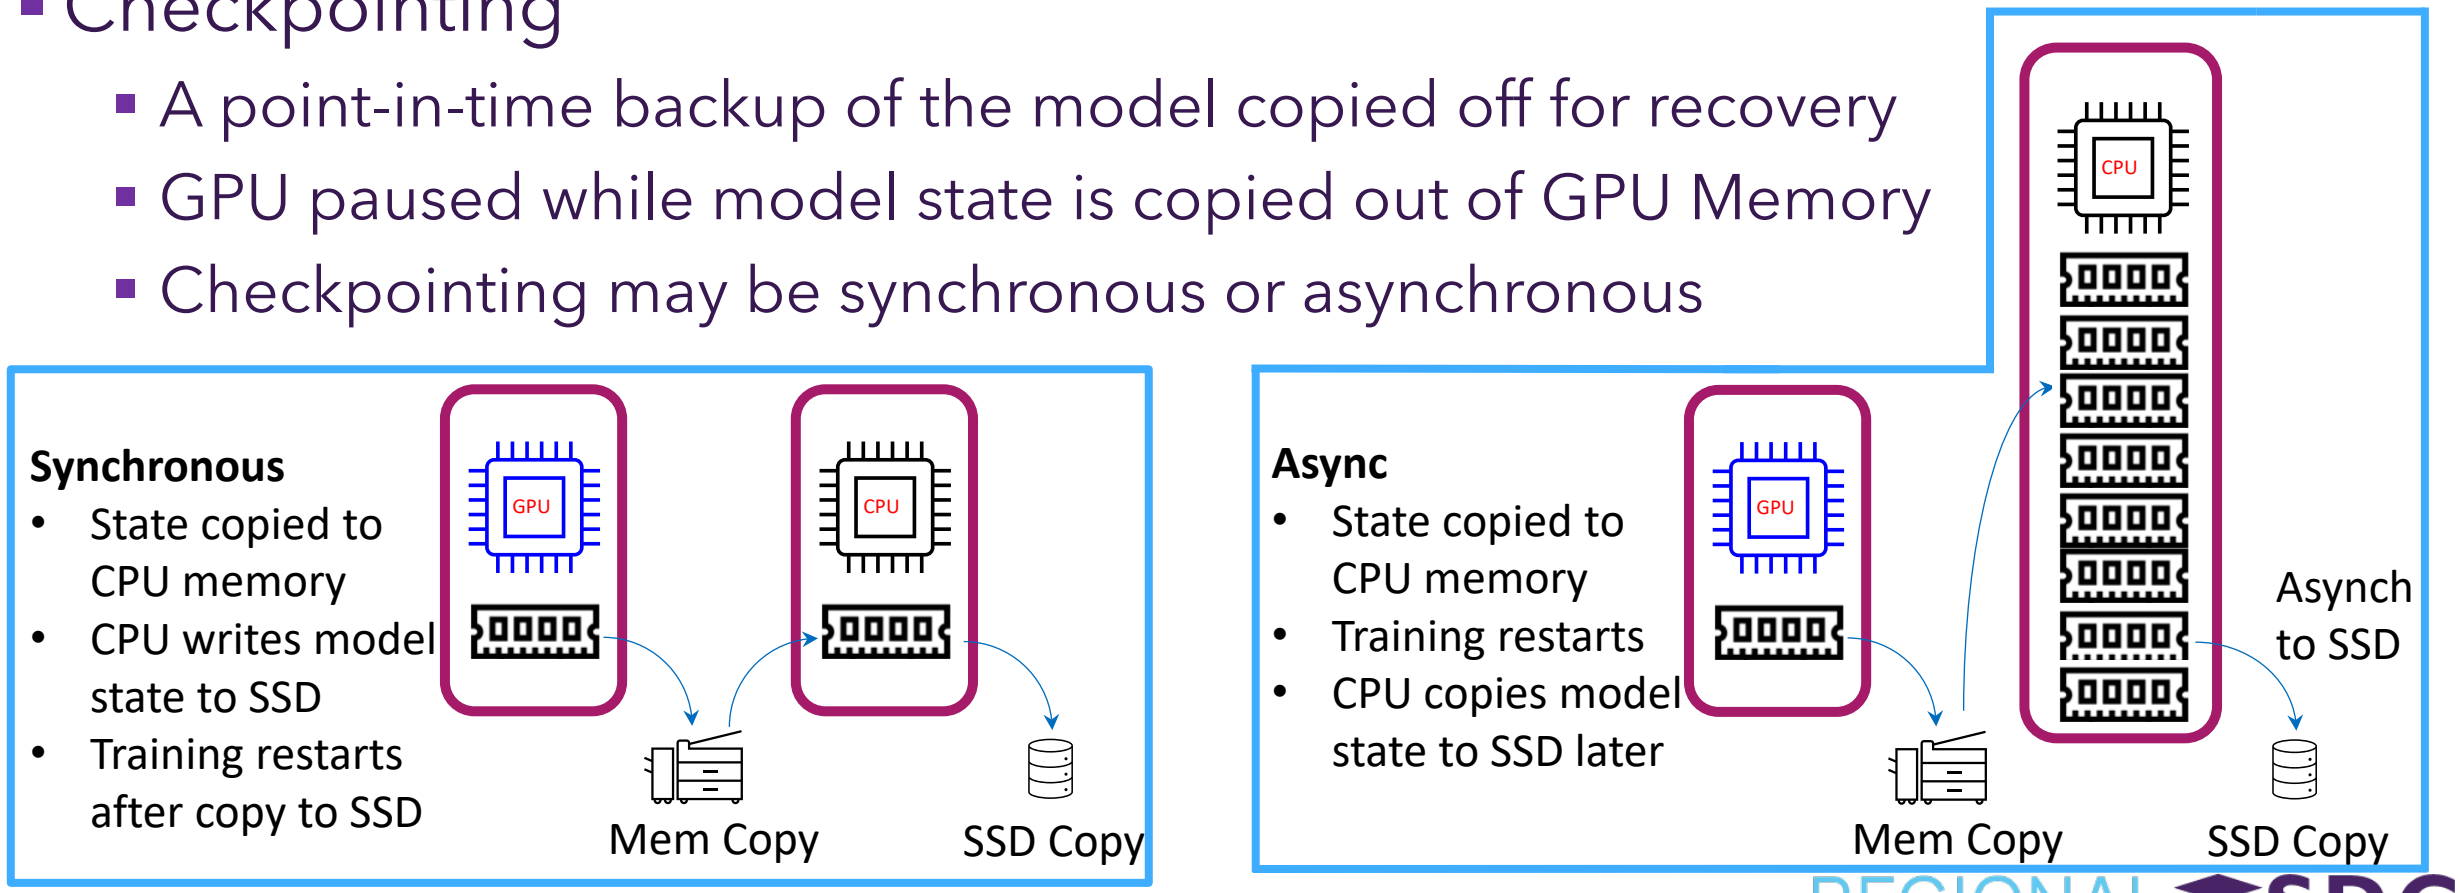

REGIONAL SDC

Fast Object for AI

- Object storage is convenient but often has higher latency than file
- File has been more popular than object for AI workloads
- Fast object has been gaining in popularity
  - SSD based object in most cases
  - RDMA accelerated object in some cases
- Stay tuned!

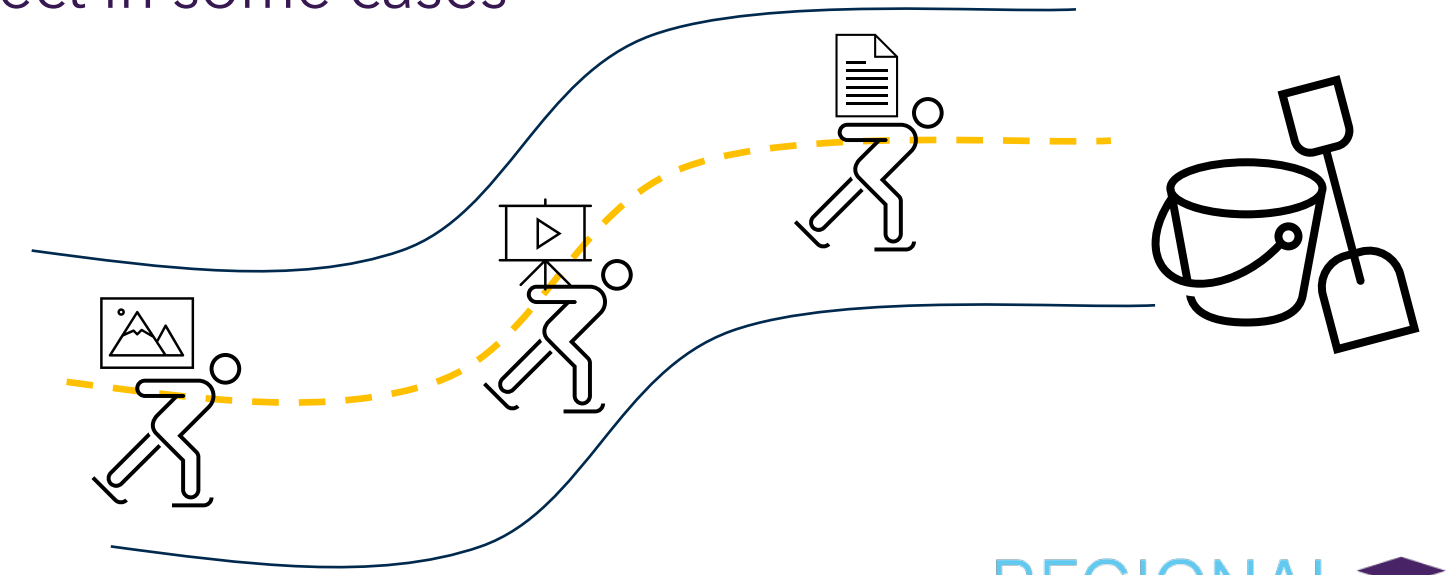

Model Training

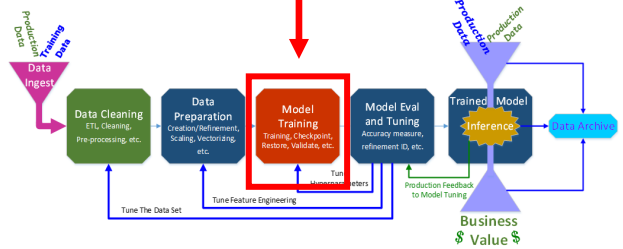

- Checkpointing – saving model weights and other state
  - Model weights are expensive when training takes a long time
  - Checkpointing saves state to allow restart after an error
- Checkpoint files are written sequentially
  - May be multiple sequential writes in parallel
- Checkpoint restoration is reversed
  - high sequential read, parallel reads to restore to multiple GPUs
- Training is paused, part of checkpoint, and all of restore time
- Storage performance determined be save/restore time goals

Sequential Random

REGIONAL SDC

Model Training – General Storage Planning

- GPUs drive the cost – maximizing GPU utilization optimizes investment
- Design for a balanced architecture
  - Balance storage performance with GPU requirements
- Consider data sources
  - May require both file and object access
- If known training workloads – match storage performance to workload *
  - AI GPU benchmarks can show peak performance for various models
  - MLCommons MLPerf Training benchmarks is a good source
  - Determine size of training examples
  - Multiply throughput and size to estimate required read bandwidth
- For general purpose training may need to support GPU max read speed
  - Can be up to 1GB/s per GPU for high end GPUs today, increasing regularly

* https://www.snia.org/educational-library/storage-requirements-ai-2024

Model Evaluation and Tuning

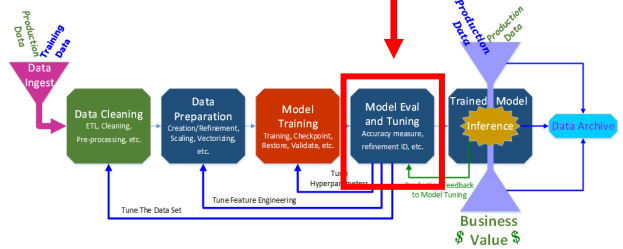

- Evaluation – measuring how well the results of the model match expectations
  - Accuracy – how often is it correct?
  - Precision/Recall – roughly a measure of how often wrong vs right
  - Measures such as F1 Score and AUC-ROC (area under the curve/receiver operating characteristics)
- Tuning – Adjusting hyperparameters to improve evaluation
- Produces a dataset containing the Model Parameters
  - Internal representation of the neural network
- Model Parameters size is constant, based on # of weights

Sequential Random

REGIONAL SDC

Inference

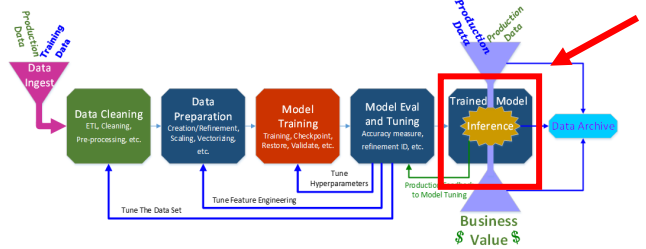

- Running production data through the finished model to generate business value
- Inference = Inferring information from the data
- Multiple types of Inference
  - Retrieval Augmented Generation from LLMs
  - Predictive analytics
  - Computer Vision
  - Anomaly detection (e.g., malware, fraud)
- Access pattern can vary some depending on type of inference
  - RAG can produce a random workload similar to databases

Sequential Random

REGIONAL SDC

RAG: Retrieval Augmented Generation

LLM Text Generation

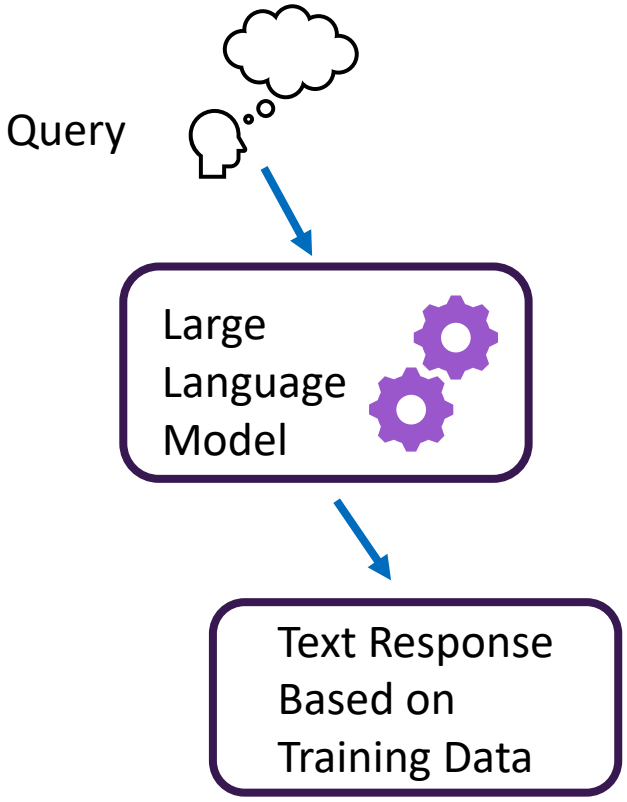

Text Generation with RAG

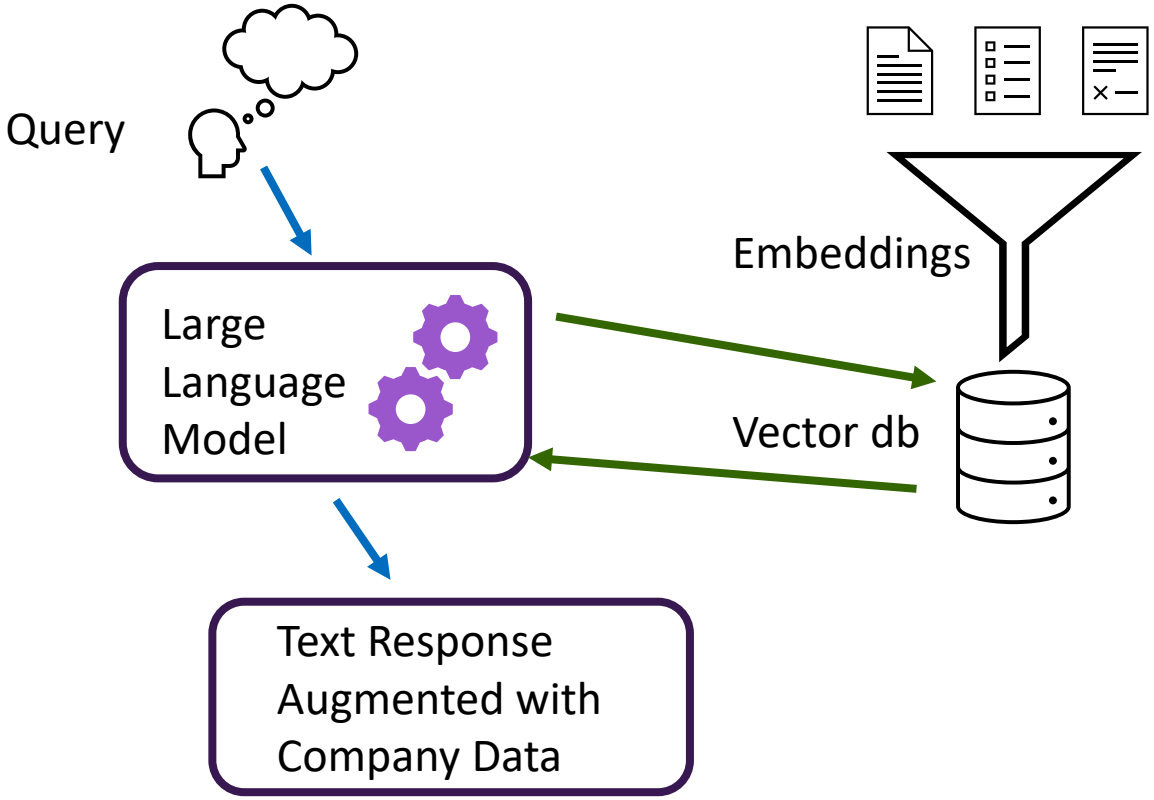

REGIONAL SDC

Agentic AI

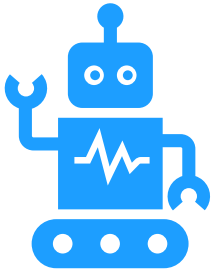

The next wave of AI digital transformation

An application of AI to perform tasks on behalf of users

Learn and act autonomously in complex environments

- Still in the early phases but generating a lot of excitement and activity
- May use RAG (Agentic RAG)
- May use LLMs or small domain specific models

And Don’t Forget! - Archive

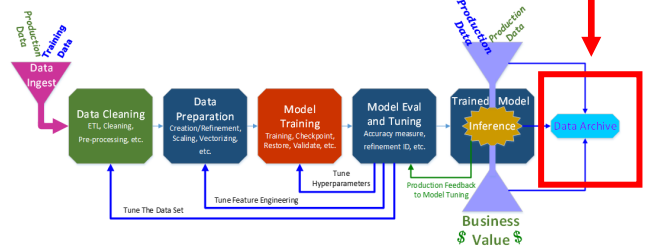

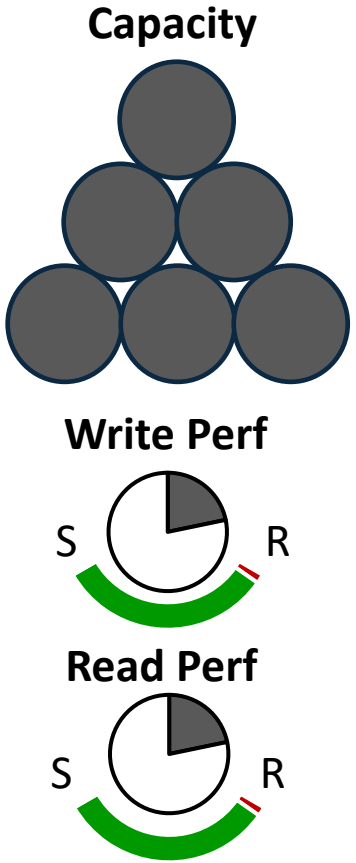

- Often overlooked, not core AI, but important AI storage
- Mandated by regulations for some AI applications
- Similar, but not traditional “archive”
  - Archived data may be brought back for training or new insights
- Performance needs vary but “just fast enough”
- No accepted terminology, maybe “Cold Storage”
- Continually growing data set
- Requires low cost and low carbon footprint storage
  - Opportunity for zero power storage such as DNA and Optical

REGIONAL SDC