Hewlett Packard
Enterprise

HPC Operational Data Analytics (ODA) for Digital Twins

Jeff Hanson

June 13, 2025

Digital Twins in HPC

- A digital twin is a set of virtual information constructs that
  - mimics the structure, context, and behavior of an individual/unique physical asset
  - is dynamically updated with data from its physical twin throughout its lifecycle, and
  - informs decisions that realize value.

AIAA Digital Engineering Integration Committee (2020)

- The ExaDigiT project aims to develop digital twins for HPC systems and data centers
- A critical challenge in this endeavor is bridging the gap between digital twin practitioners and the operational data needed to ground their models in reality.

ODA and Digital Twins

- ODA provides a comprehensive infrastructure for gathering and analyzing system and data center metrics end-to-end, supporting both real-time and long-term analysis.
- Digital twins provide models that are required to perform what-if analysis, enabling model driven optimizations of the system which is a basis of prescriptive analytics in ODA.

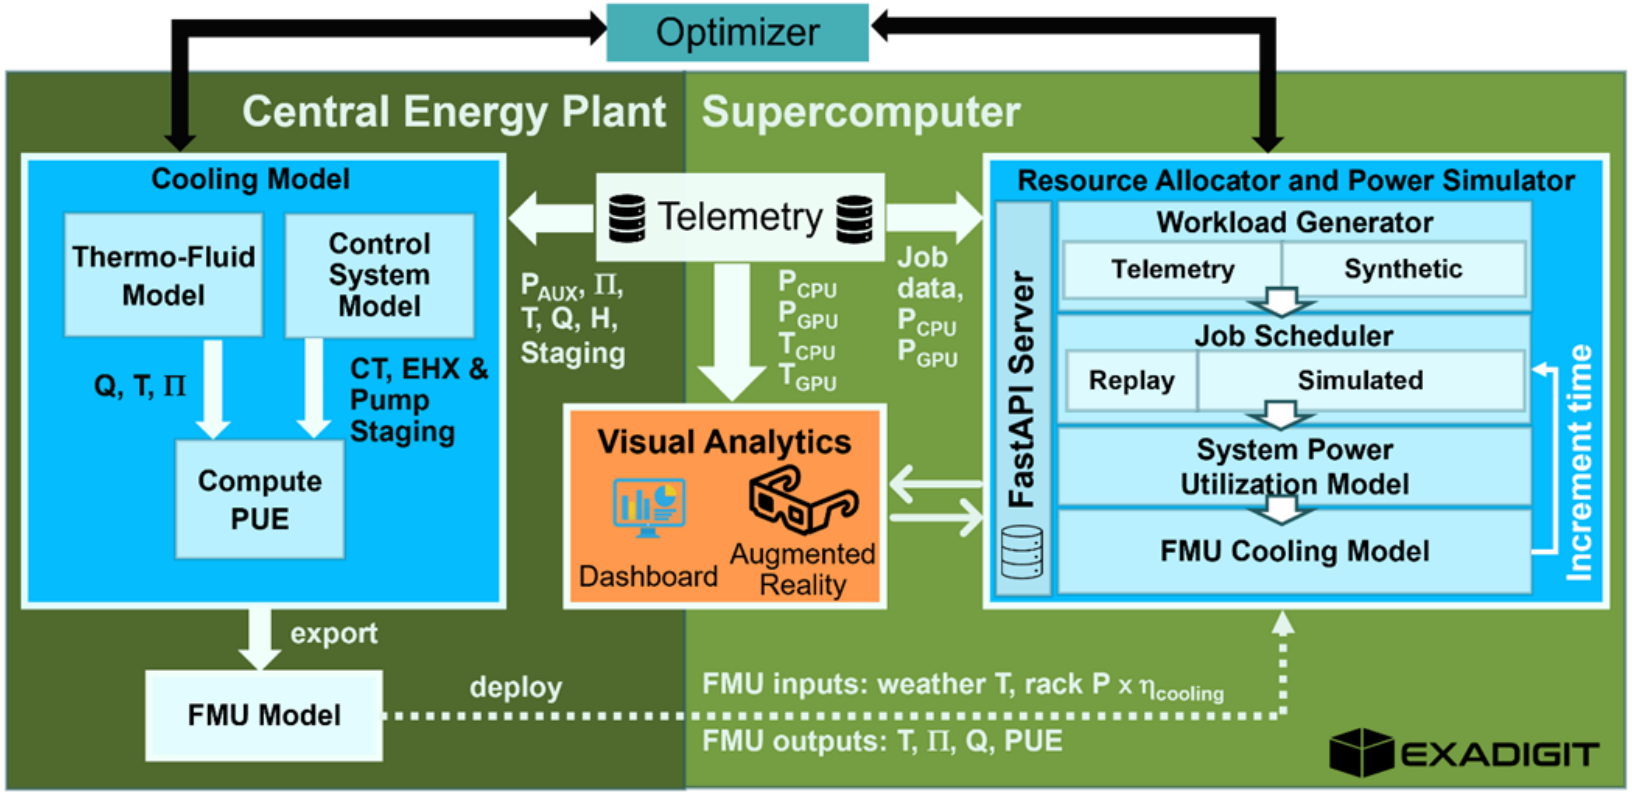

What should be done?

- I present an analysis of the trade-off between model accuracy and practical implementation, visualized on a dual-axis diagram comparing data volume against both model accuracy and practical usefulness.

Effective Digital Twins are limited by the data available requiring appropriate accuracy to be able to generate prescriptive insights

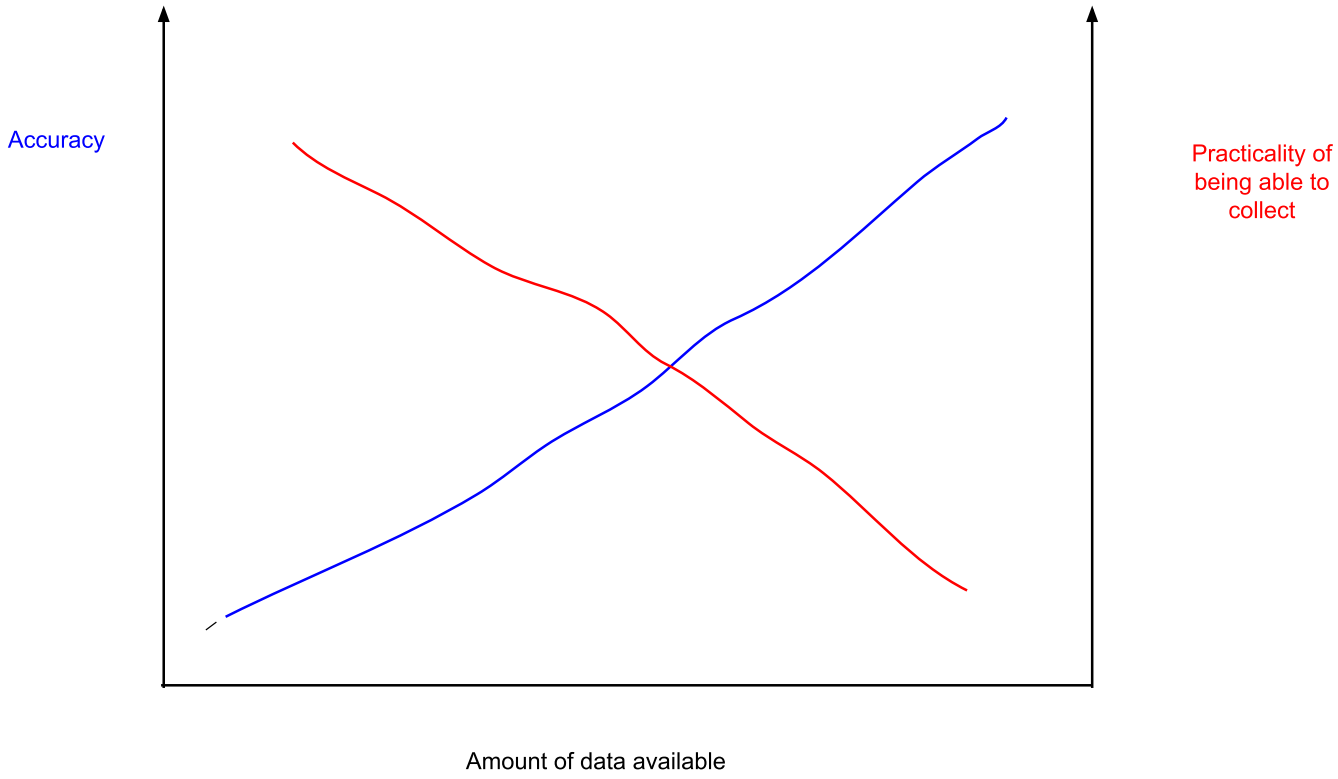

What should be done?

- As data volume increases, model accuracy improves, but practical implementation becomes more challenging due to resource and budget constraints. To find the optimal balance, we propose three key efforts:
  - Use case-based standardization to guide data curation and model development
  - A system of metadata collection and analysis for telemetry
  - Determination of appropriate sampling rates and time domains for data collection across a wide range of interfaces.

Use case-based standardization

- Attempts to make **universal** standards for data are doomed.
  - Application Fingerprinting Working Group in ExaDigiT has worked for a year with a group of public datasets to build a single presentation layer with limited progress.
  - EEHPCWG ODA has a published conceptual model but no standard for data itself
- A different approach is to come from the use case side and build software bridges between disparate data sources
  - EEHPCWG ODA Dashboard team has created a method to make a data source plugin from a template
  - ExaDigiT ODA WG expects to be able to connect to data sources
- Use cases that express the what will lead to a how
  - As a researcher in sustainable systems I want to track energy used at component level
  - As a workload programmer I want to observe the effect of different algorithms on GPU utilization.
  - Et cetera

System of metadata collection

- Metadata systems enrich telemetry databases by providing details about sensors, events, node location, system attributes and more
  - Metadata system records min, max values, reading type, location on board, criticality of the sensor, …
  - This makes sensors understandable to Digitial Twin use cases.
- Example from HPE CrayEx235a
  - Sensor is Accelerator0VoltageRegulator0OutputVoltage which at a timestamp has value 47.914
  - An expert in the design of the hardware could identify what this is and why it might matter but it is obscure for most
  - The metadata includes

Geoname for system with sensor is x1000c2s0b0n0

Sensor name is PDB INRUSH ACCEL0

Sensor is on node 0 which is the only node in the chassis

Sensor is for 48V part

Part is on an accelerator

It is the 1 of 4 such accelerators on the board where counting is 0 based

Part is a voltage regulator

Sensor units are volts

Sensor thresholds are defined for caution, critical and fatal for both upper and lower

Part is of medium importance in a scale of high, medium, low

Alerts should be done on lower critical when seen for more than 30 seconds

Metadata service provides query access for metadata collections

- Dashboards can use generic names like cpu temperature and metadata maps to actual sensor
- Digital Twin models can use generic names like gpu power and metadata maps to the right data source
- System users/analysts can browse available metrics using a search UI

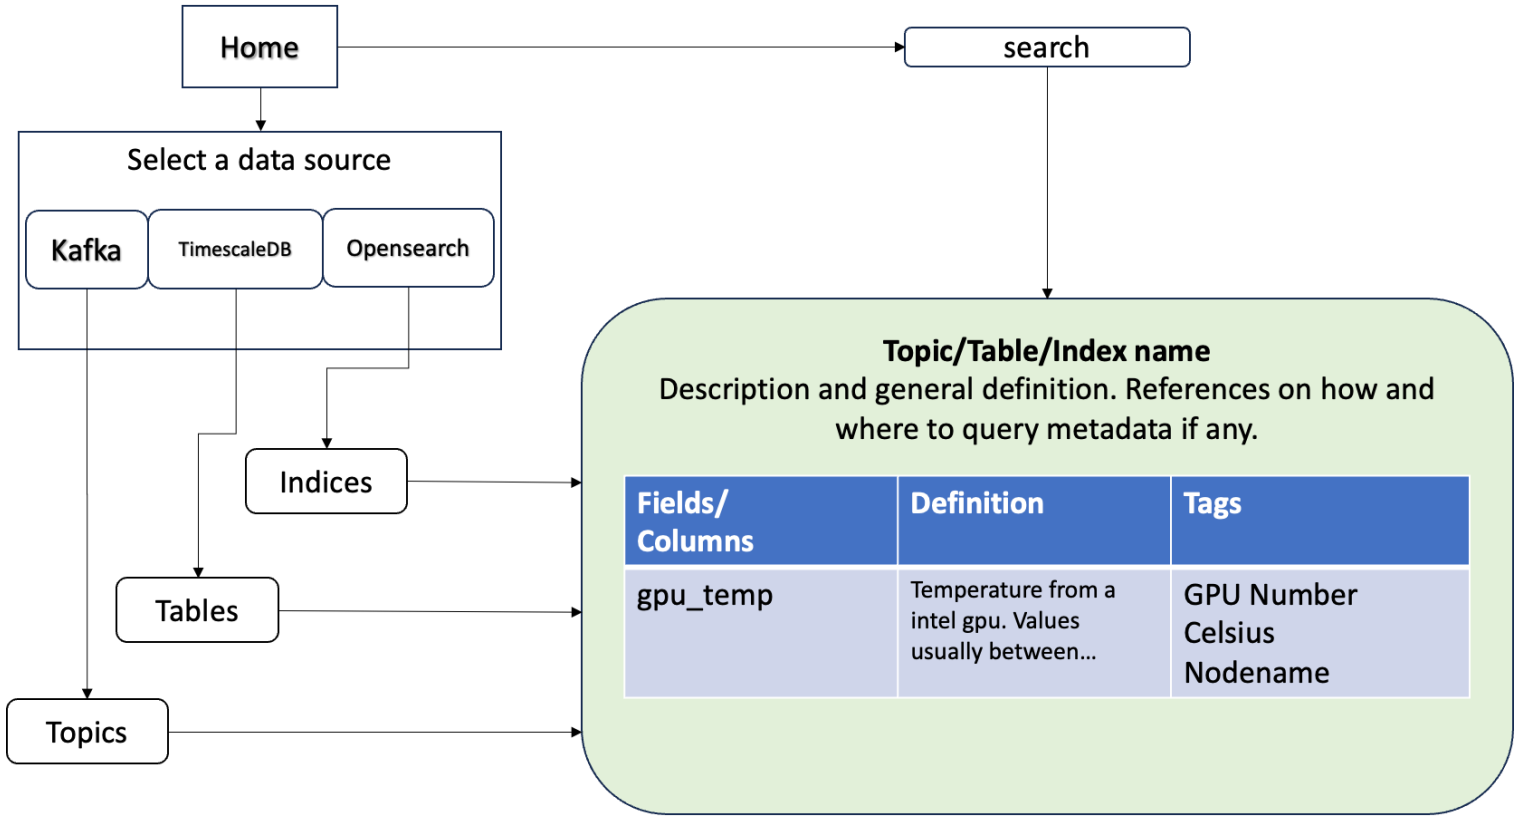

Determination of appropriate sampling rates

- Effective Digital Twins are limited by the data available requiring appropriate accuracy to be able to generate prescriptive insights.
  - Example if the total IO over the life of a multi hour jobs is known and not the pattern of IO the scheduling model could try to run several at once because they fit inside the IO profile only to find out that the initial startup IO causes all the jobs to stall.
  - Hernandez and Elwasif - **Fine-Grained Application Energy and Power Measurements on the Frontier Exascale System** CUG 2025 with 1ms sampling
- Solution involves the same concepts from the prior slides
  - Use metadata system when generating data to give monitoring system hints on what to retain at what time resolution
  - Use metadata system to annotate the existing data sets for time resolution

Working Group Kickoff Operational Data Analytics

Goal: "Bridge realities"

- **Data standards & Data dictionaries:** Tools and frameworks for preparing data and utilizing data for DT
- **Tools and systems** for DT & ODA integration - Reference design and implementations
  - Validators, loaders, transformers, synthetic data ge nerators, adaptors, connectors

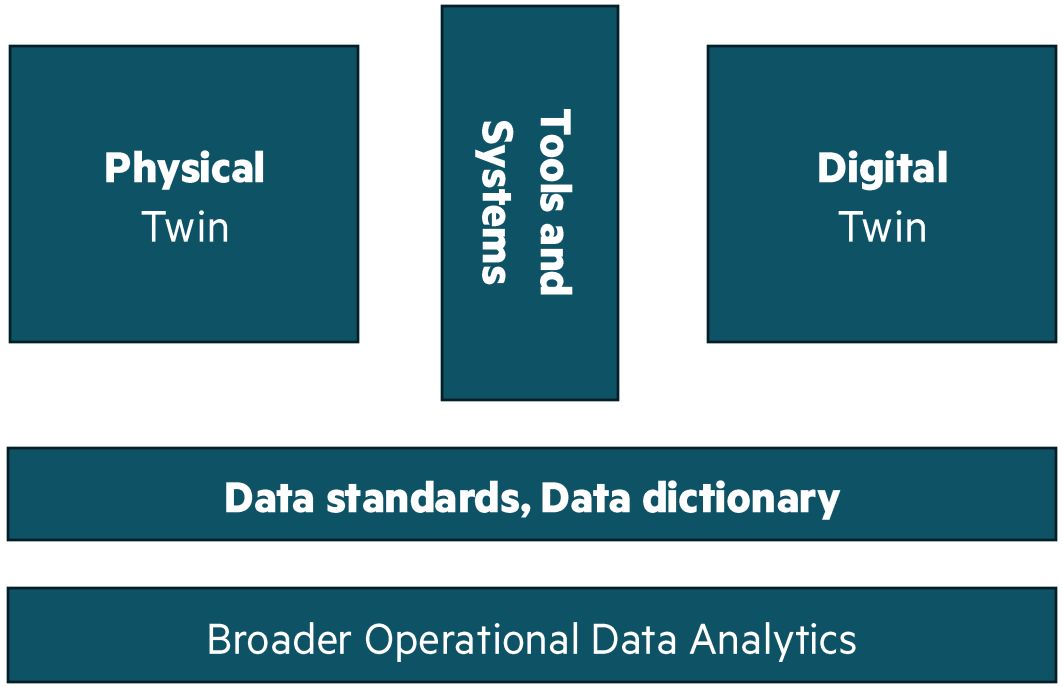

1st kick-off meeting was in May, 2025!!

Meeting Details

Woong Shin (ORNL) – shinw@ornl.gov

Jeff Hanson (HPE) – jeff.hanson@hpe.com

Thank you

ExaDigiT ODA working group (Email jeff.hanson@hpe.com or shinw@ornl.gov)
ExaDigiT Application Fingerprinting working group (Email trjones@ornl.gov)
Join the groups!

The author acknowledge the invaluable assistance to this presentation and the history of the ideas
- Dr. Woong Shin, ORNL
- Dr. Torsten Wilde, HPE