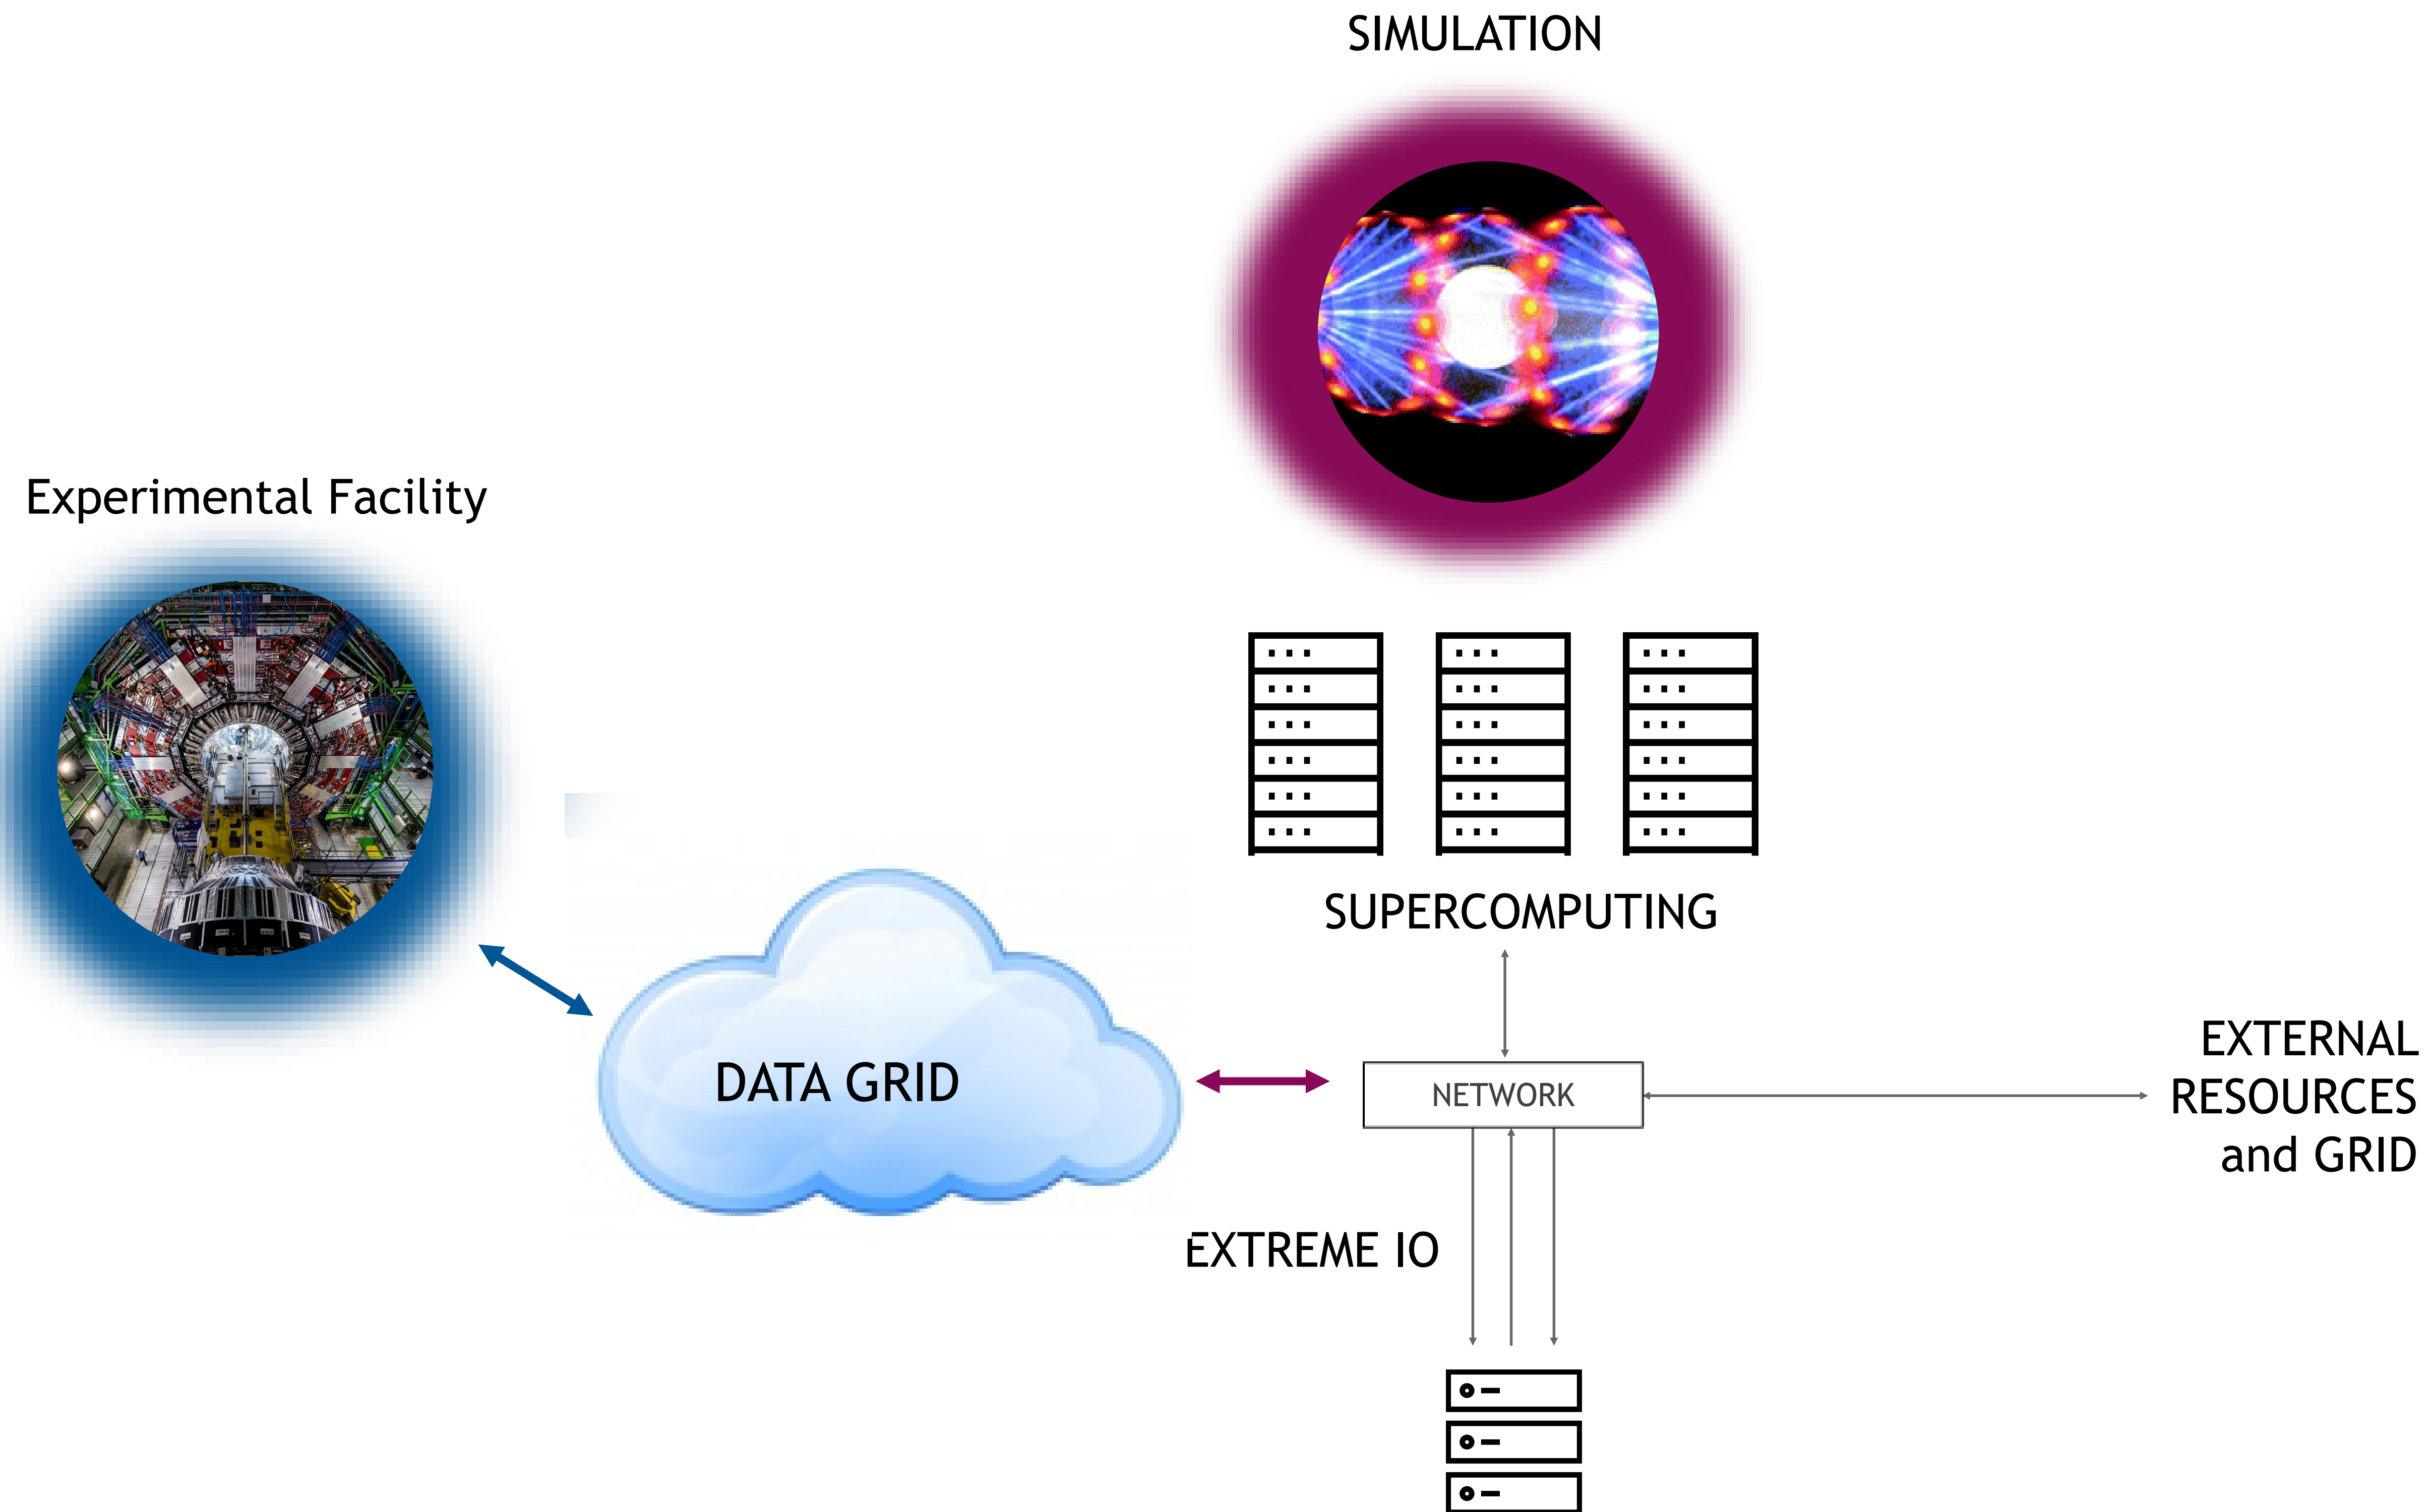





HPC BEYOND EXASCALE HARNESSING THE EDGE

TRADITIONAL HPC ECOSYSTEM




50 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

50 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

50 Years of Microprocessor Trend Data



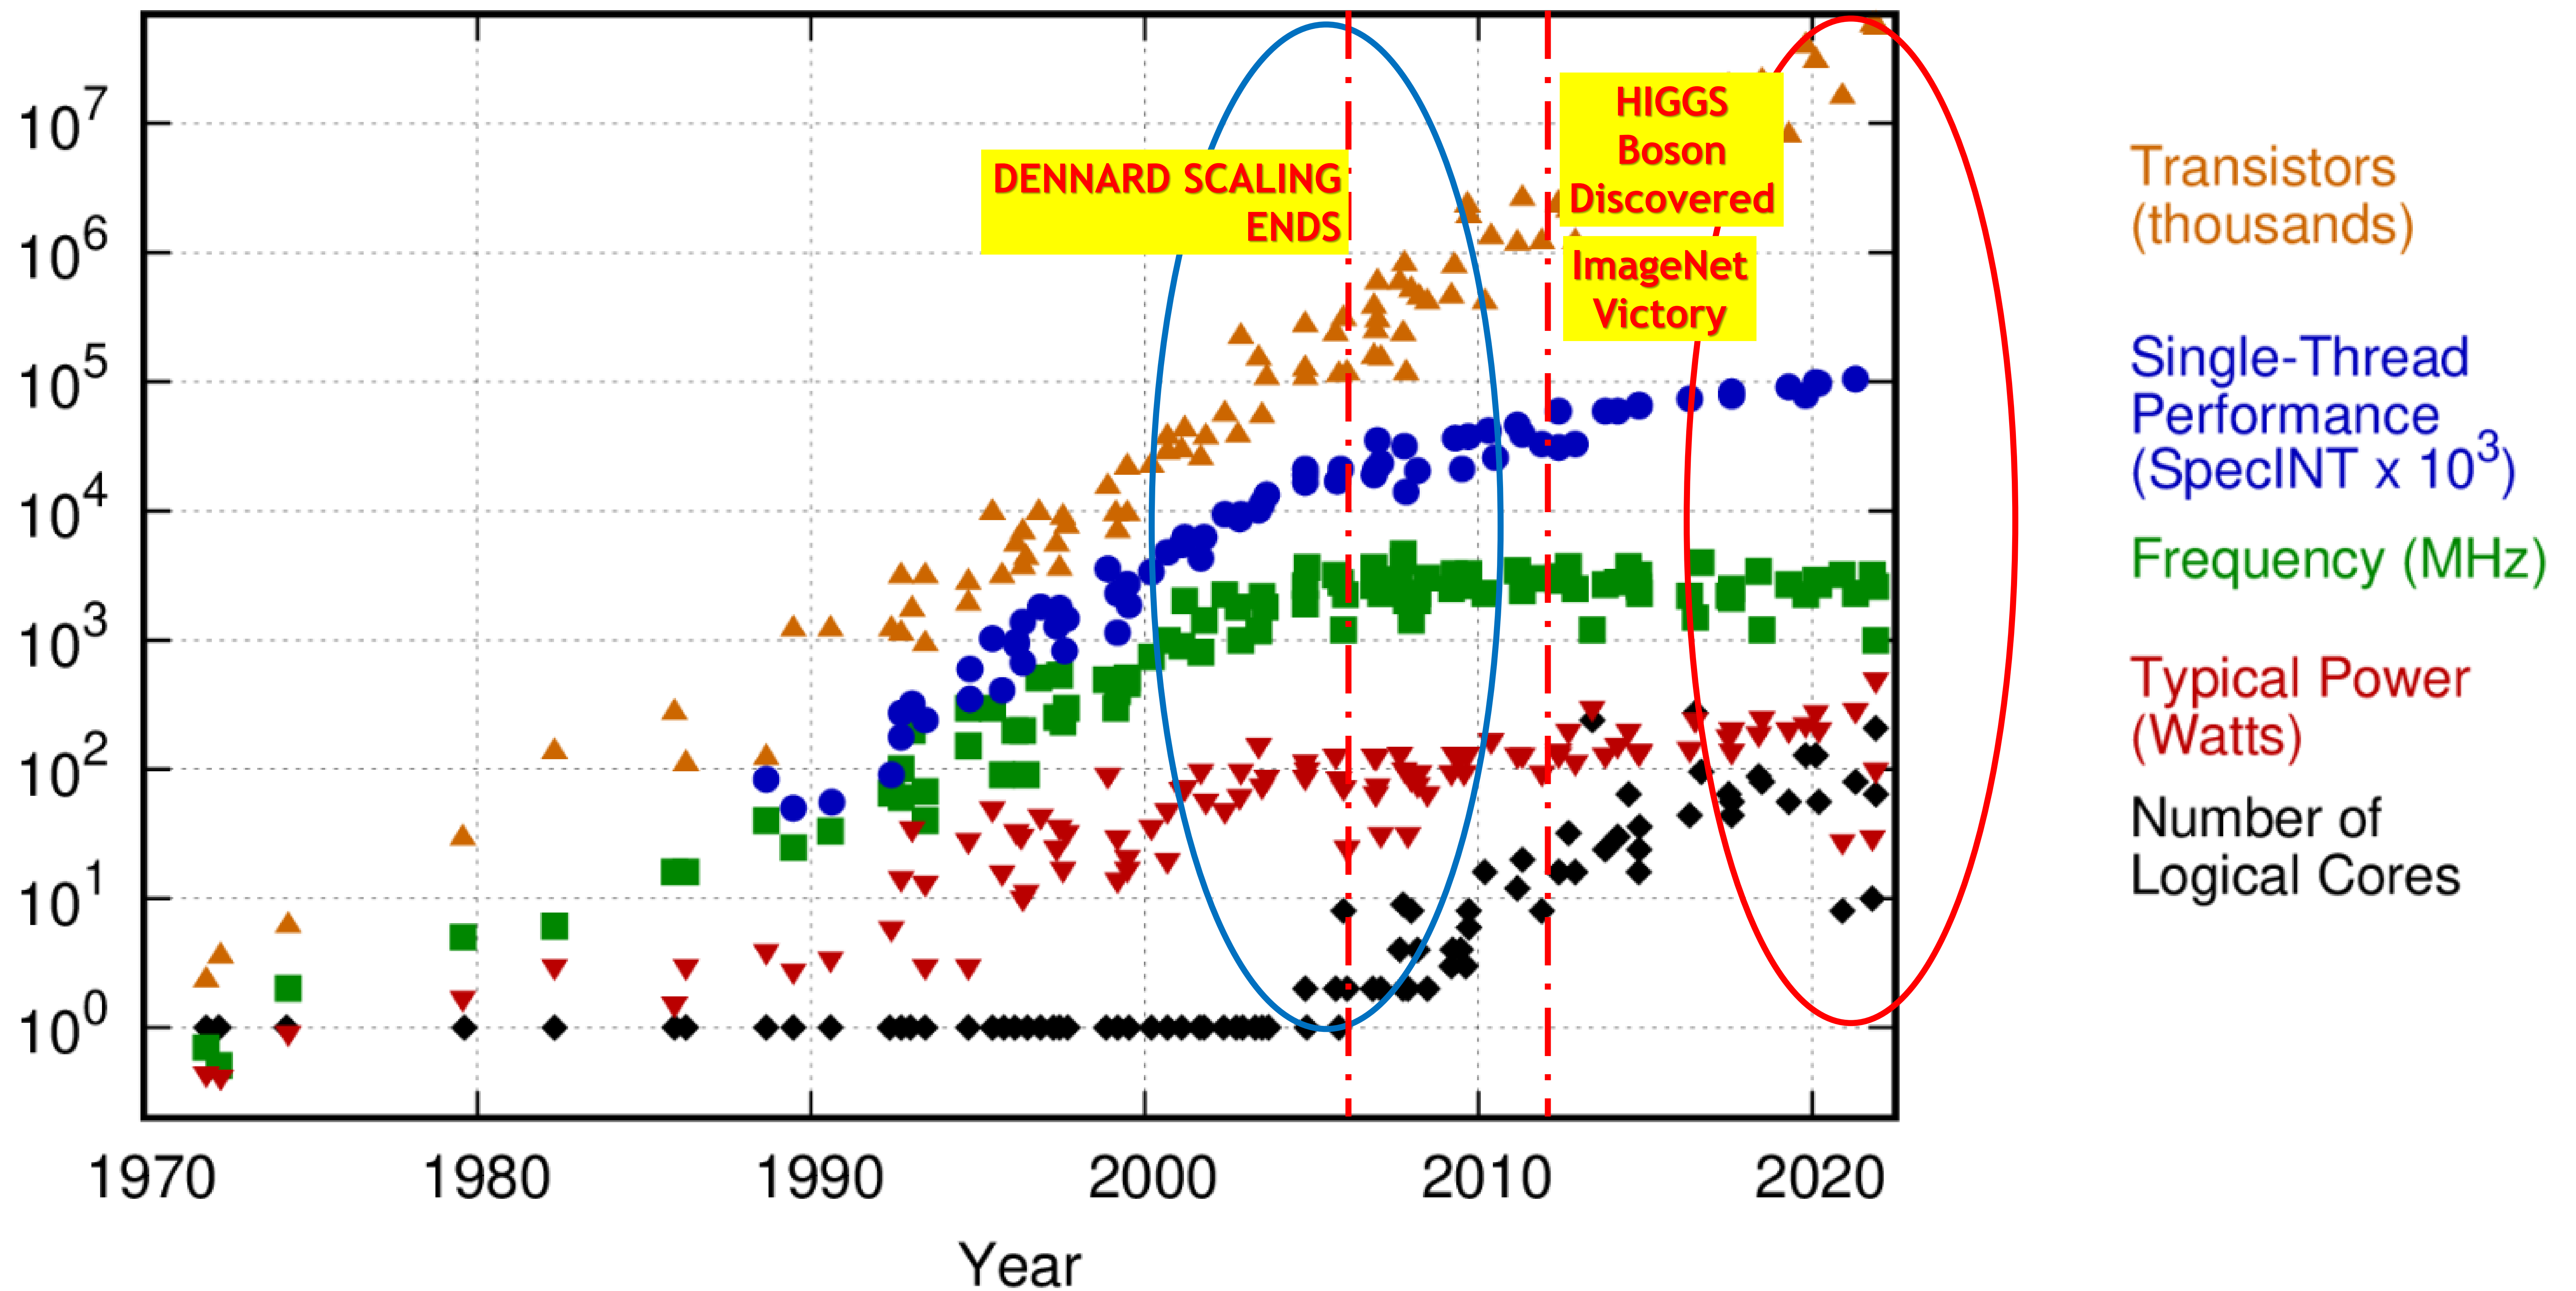
Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

50 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

50 Years of Microprocessor Trend Data



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp


50 Years of Microprocessor Trend Data



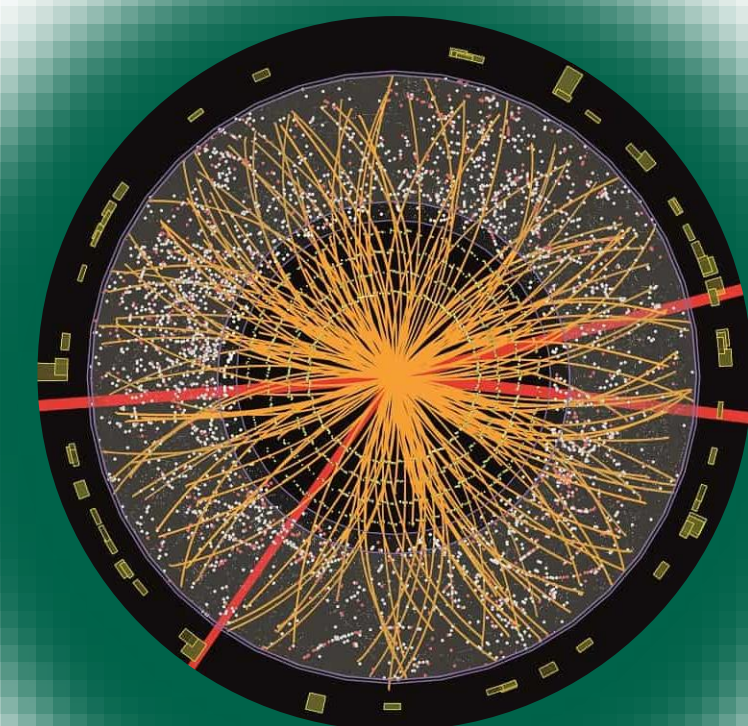
Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2021 by K. Rupp

EXPANDING UNIVERSE OF HPC


EDGE




HPC * AI



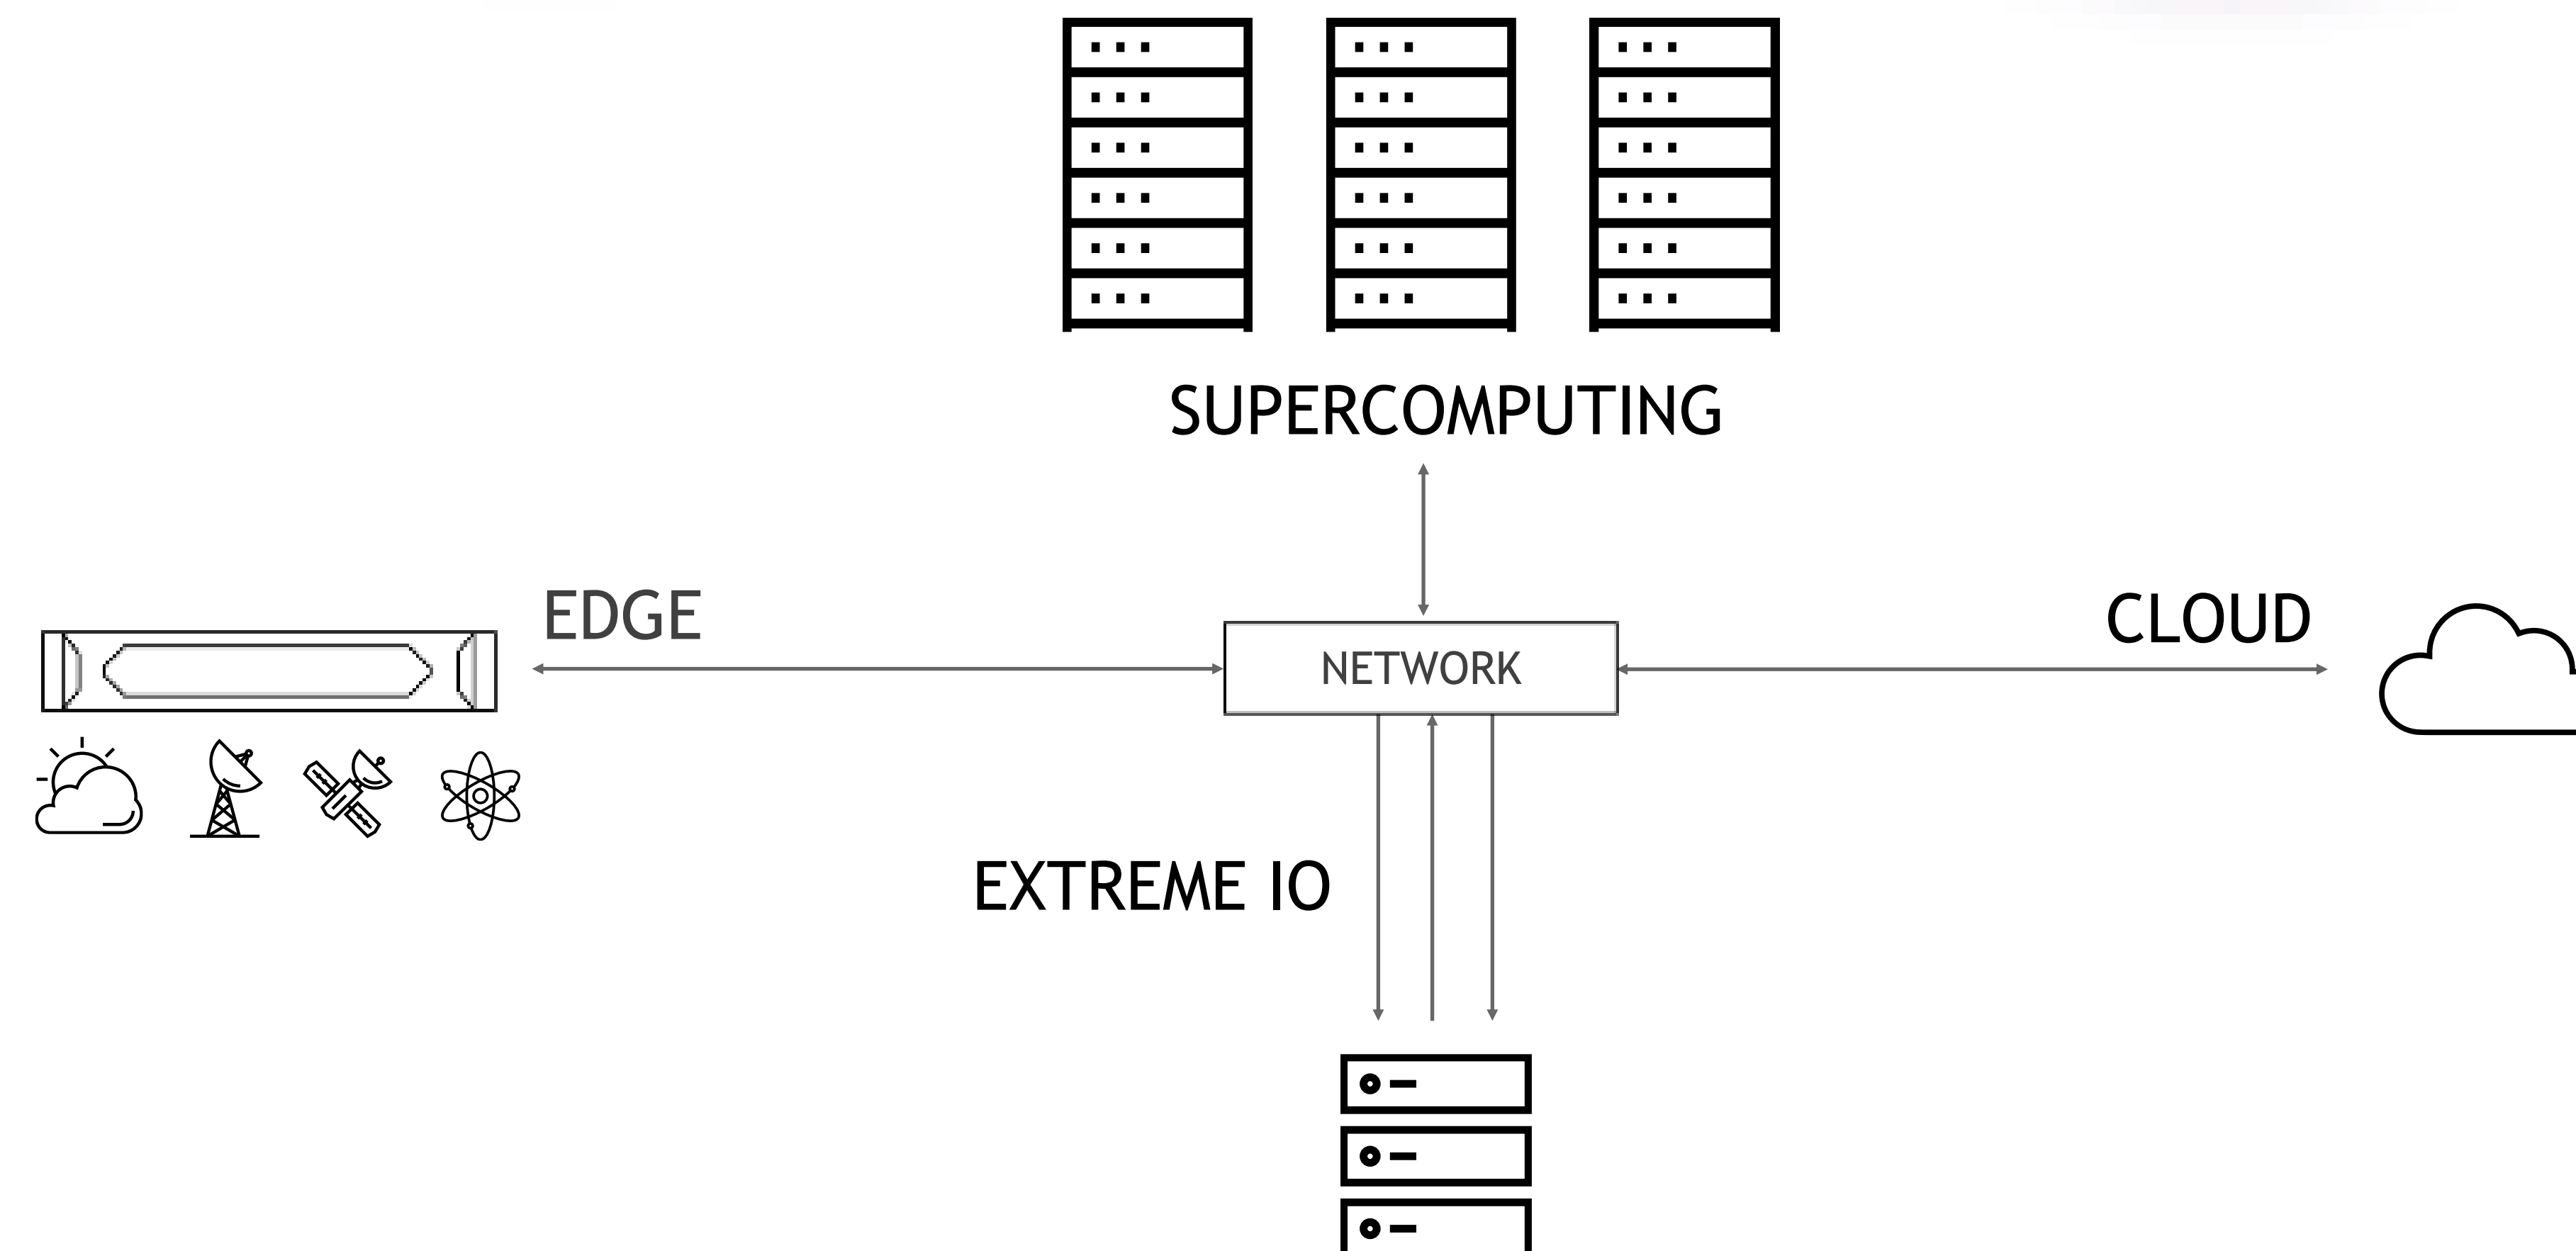

SIMULATION



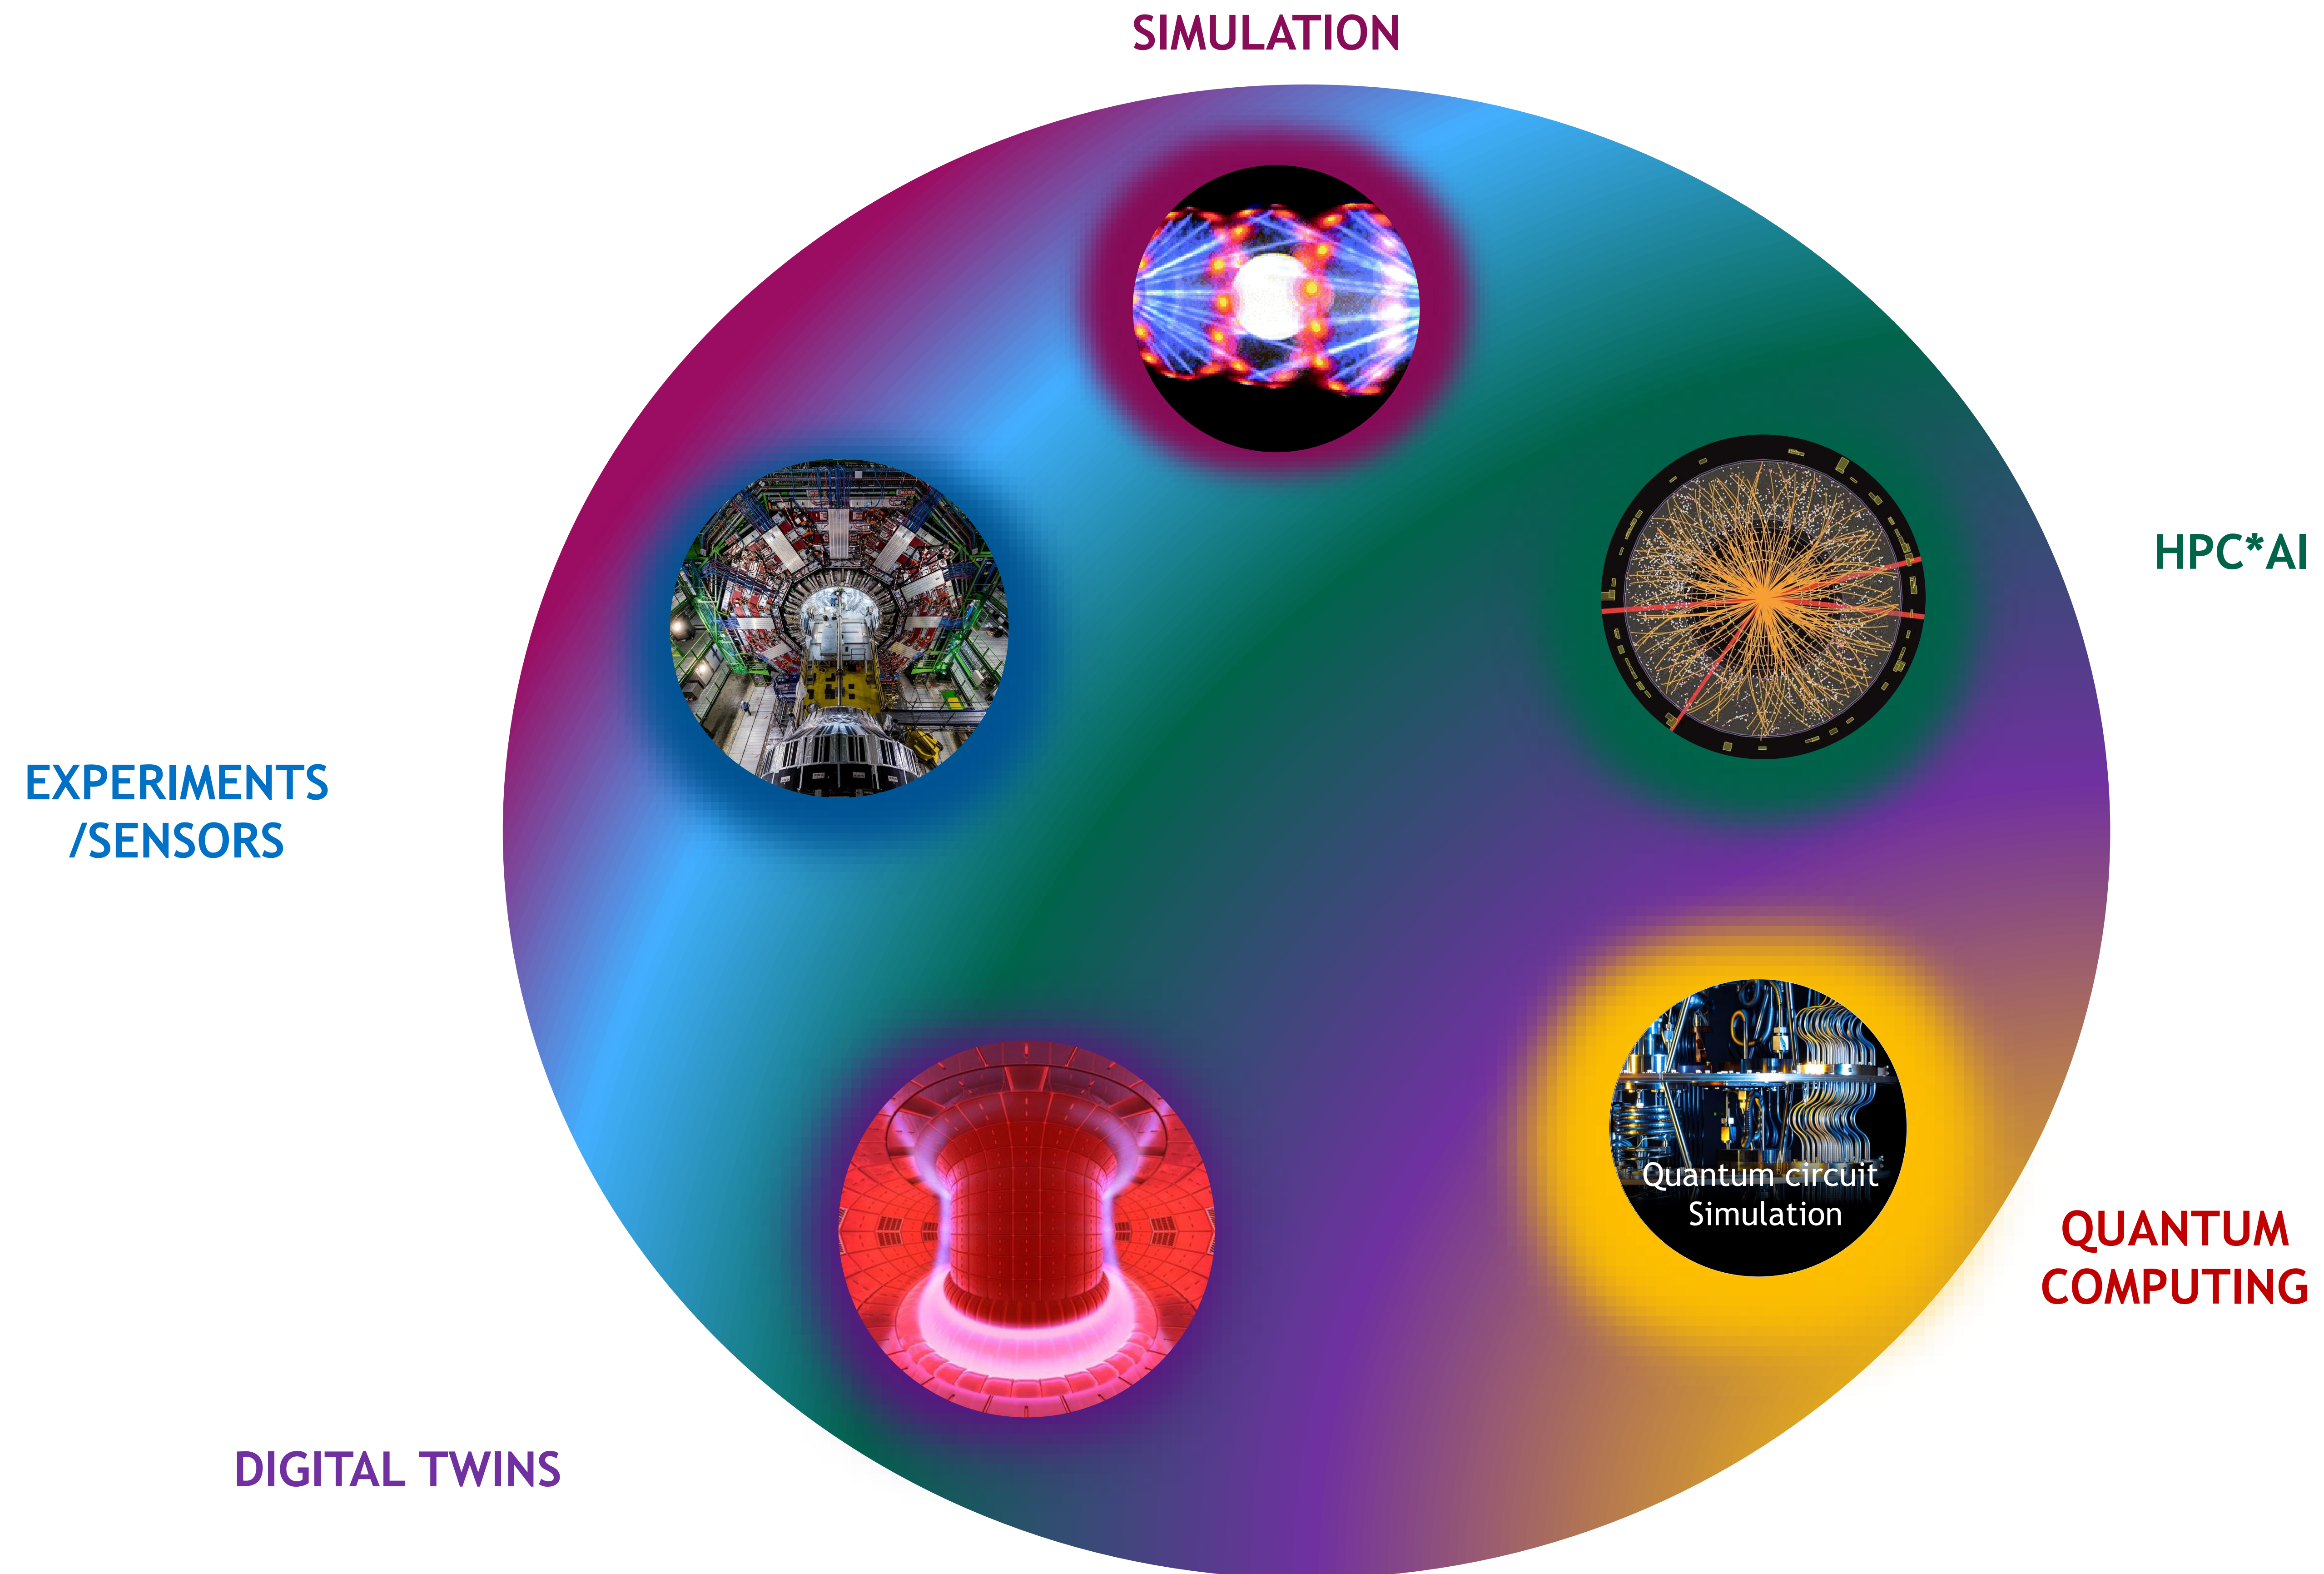
DIGITAL TWIN



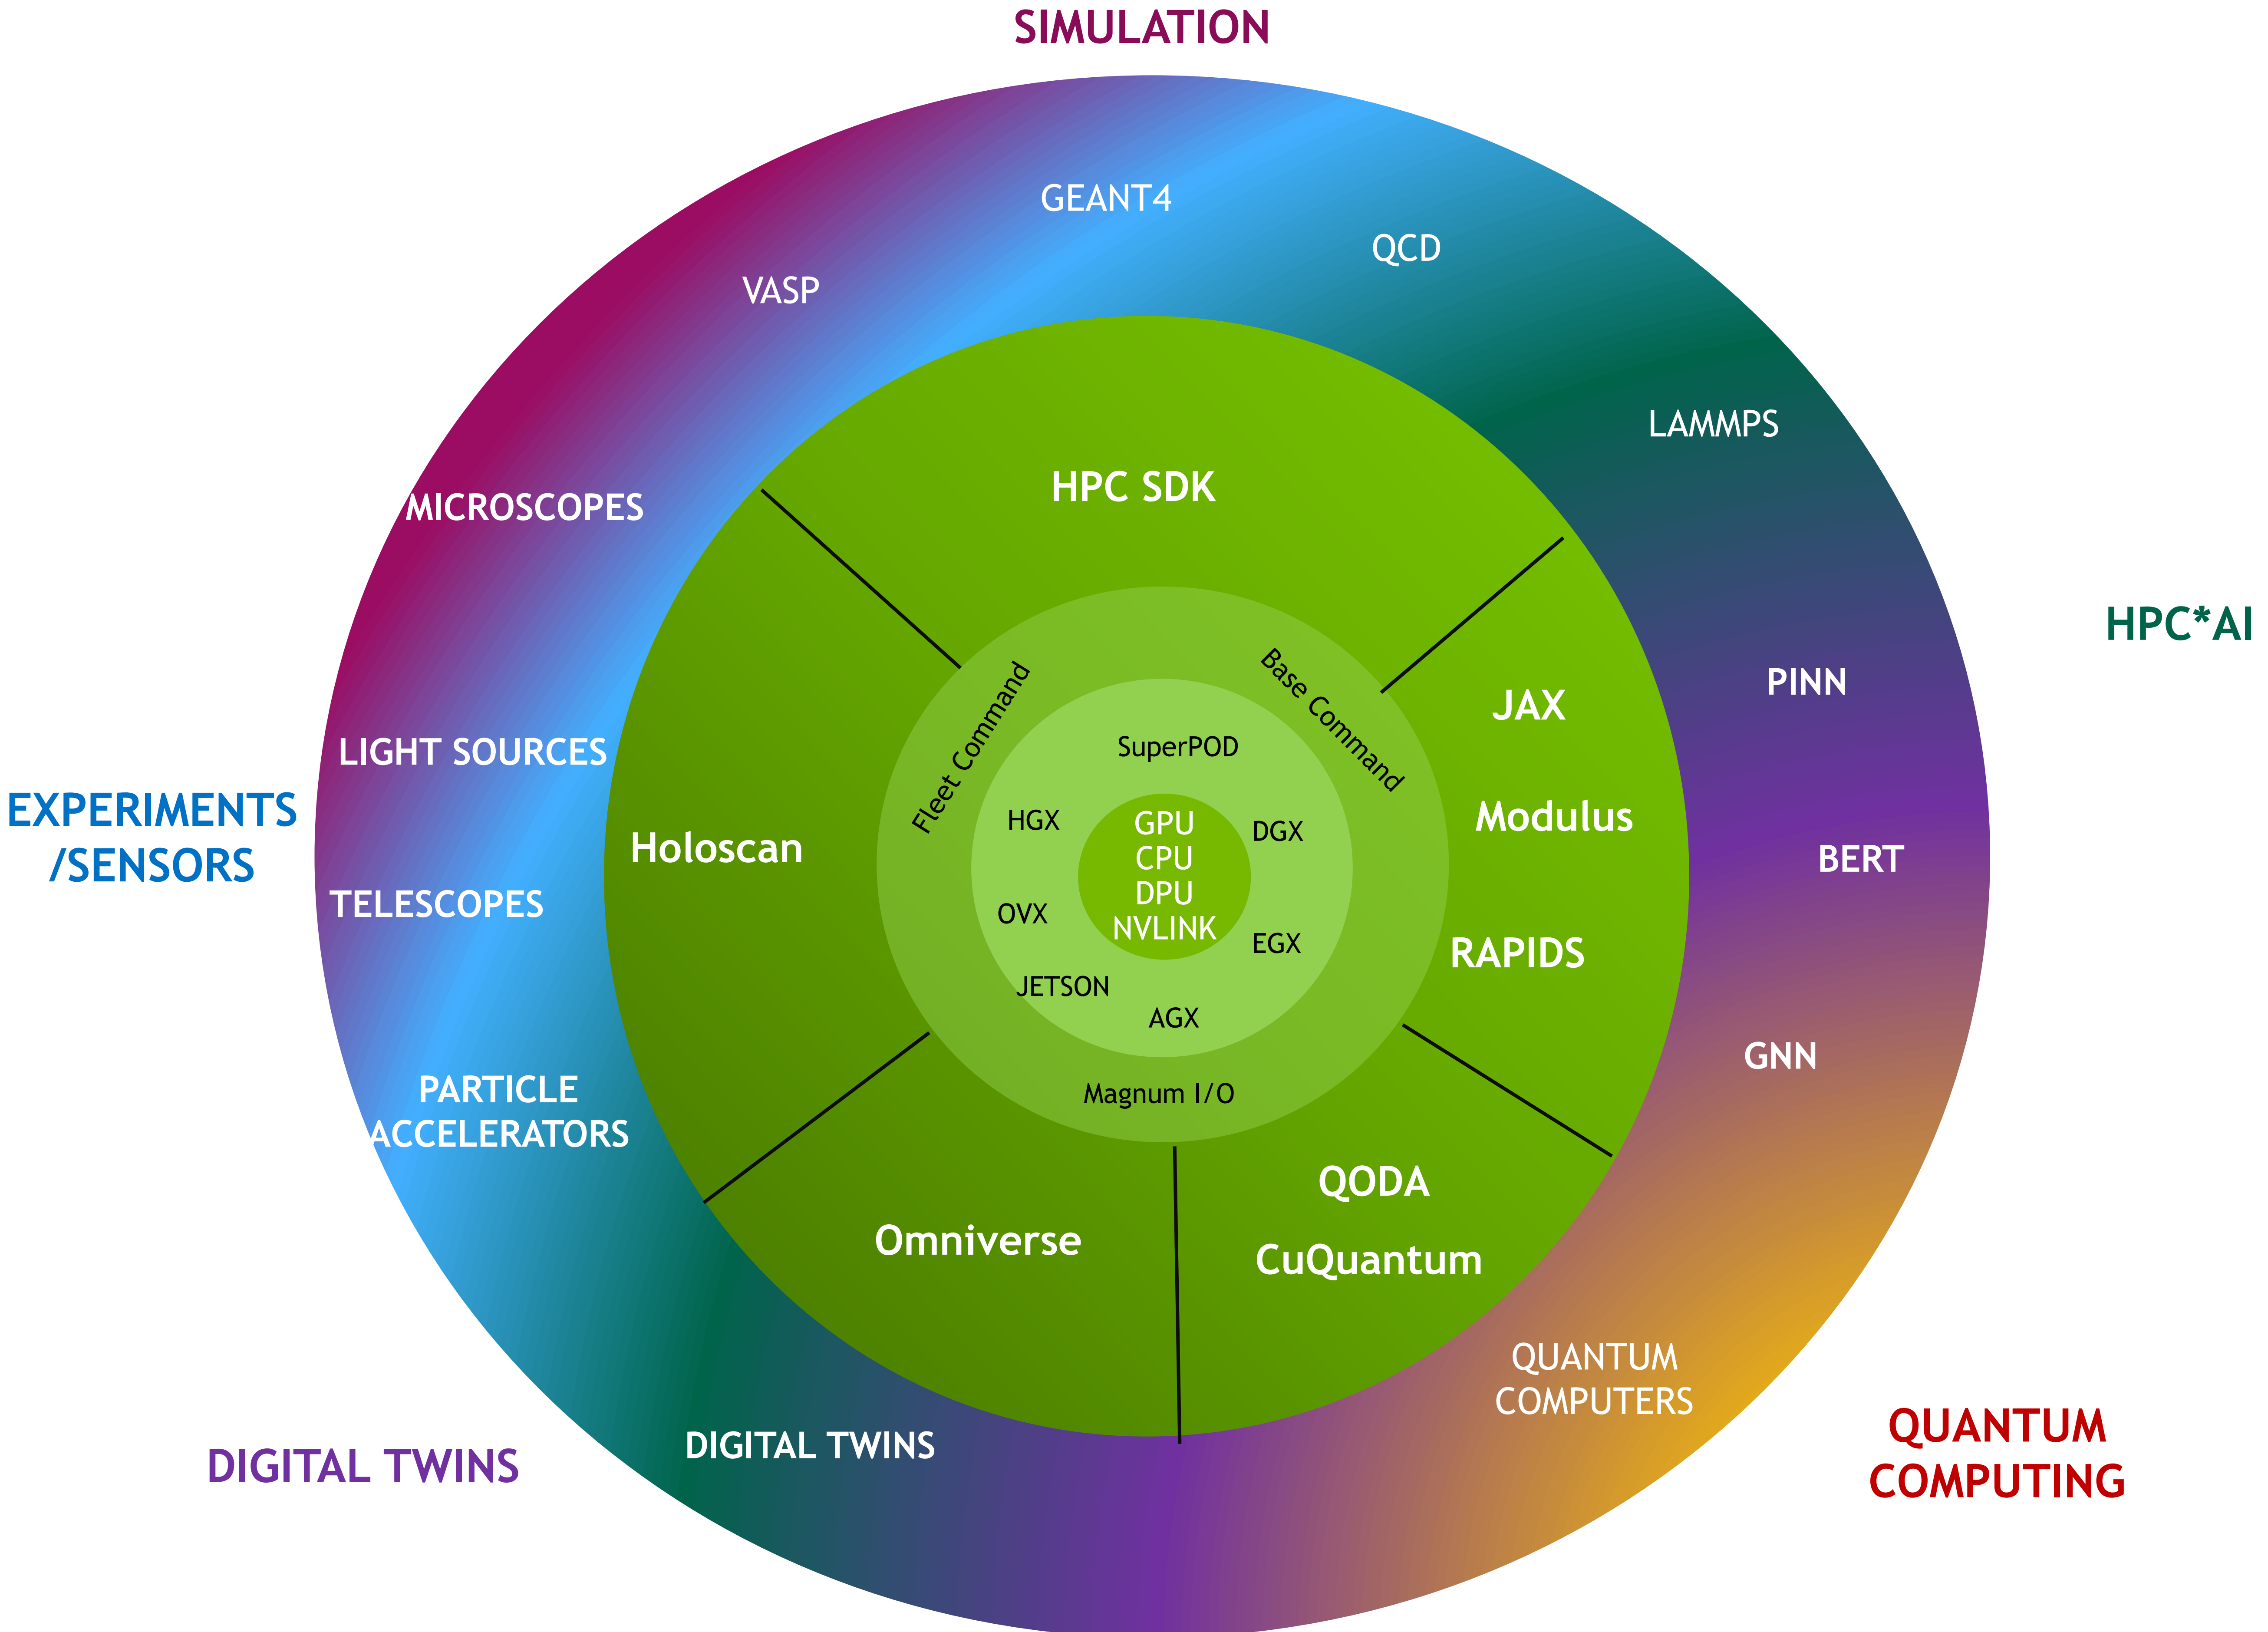
QUANTUM COMPUTING



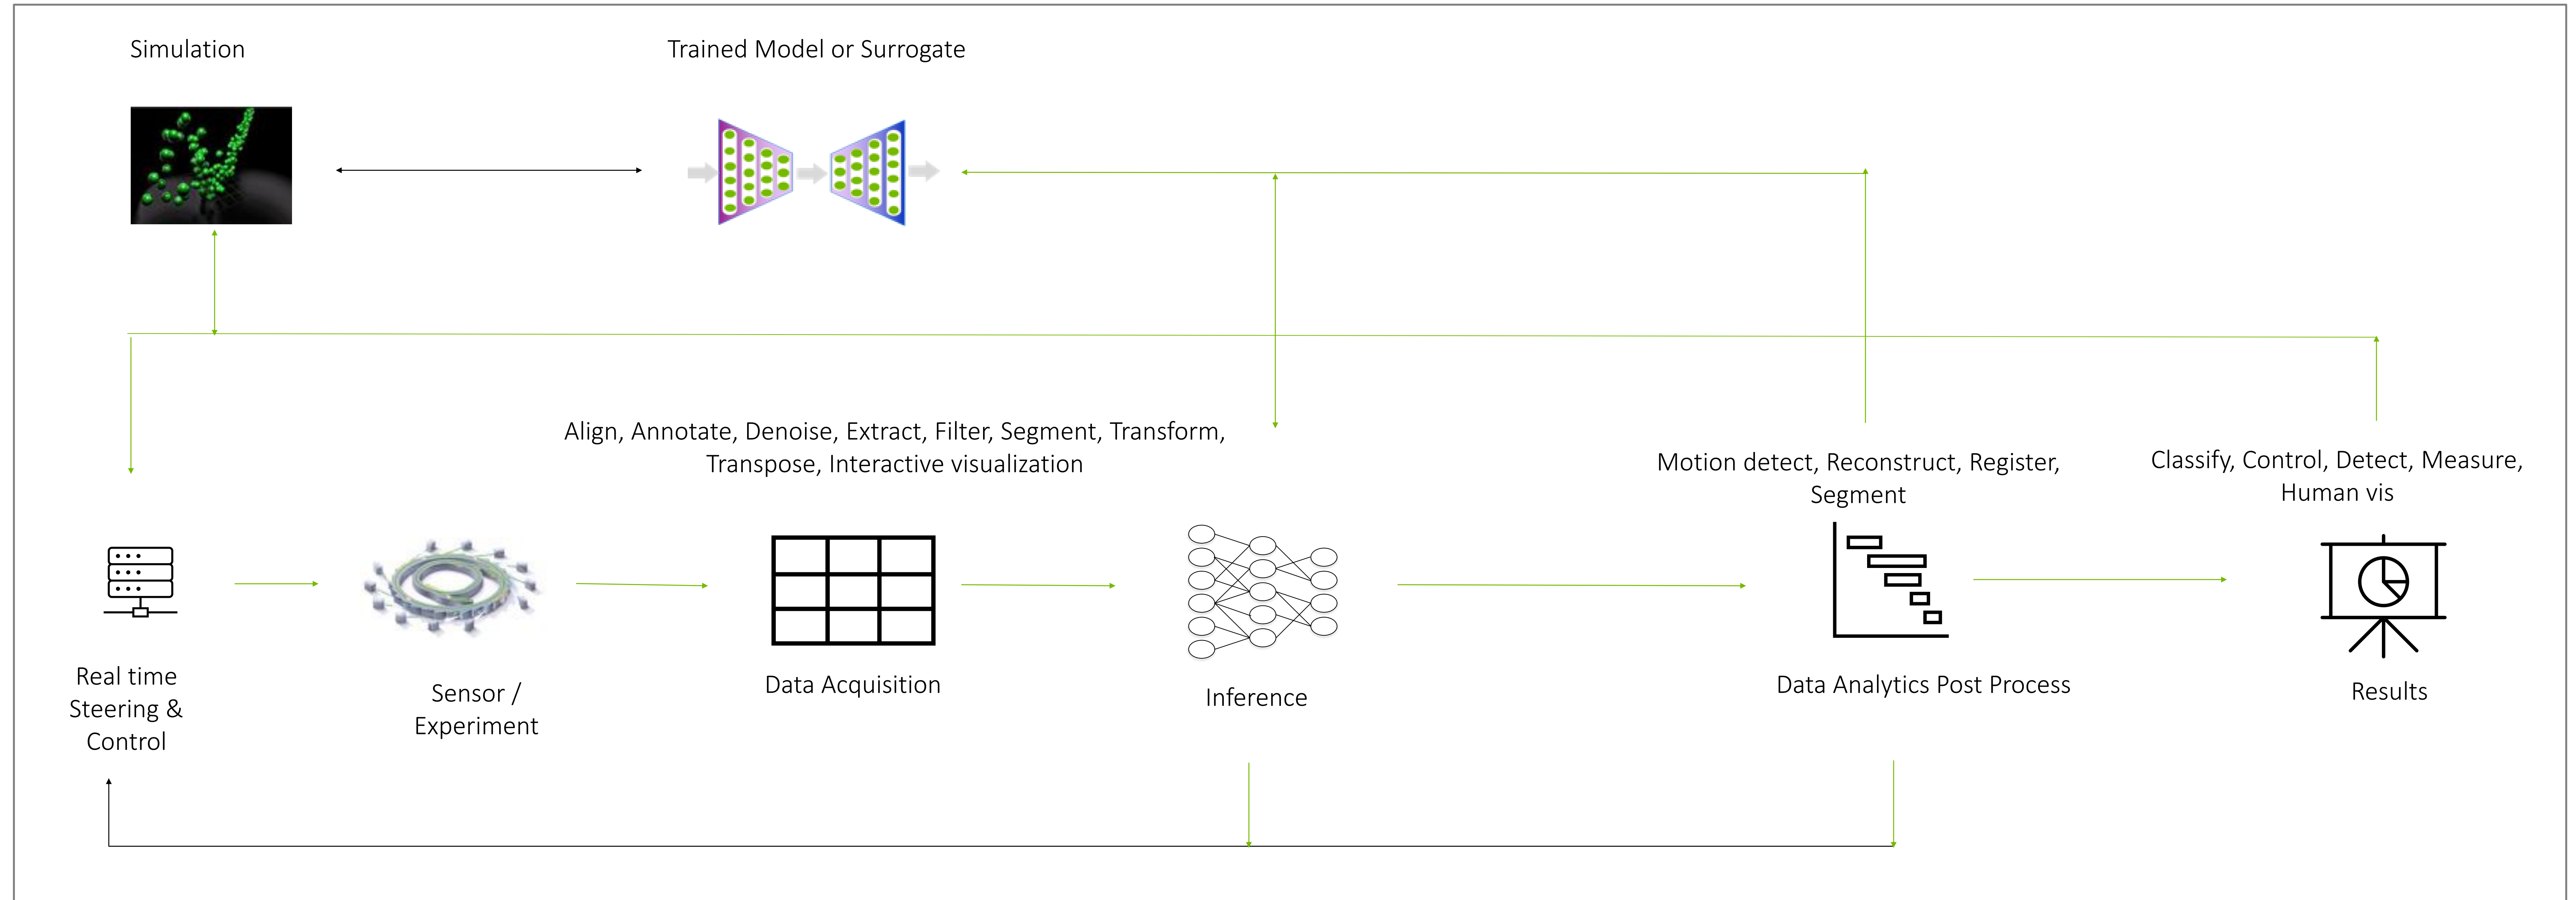
COMPOSITE WORKFLOWS EMERGE TO SOLVE GRAND CHALLENGES



NVIDIA ROADMAP EVOLVING TO MEET THE CHALLENGE

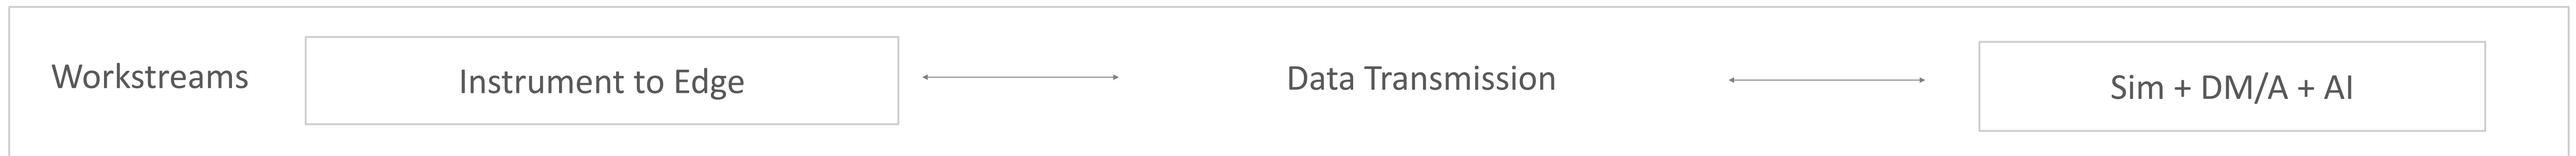
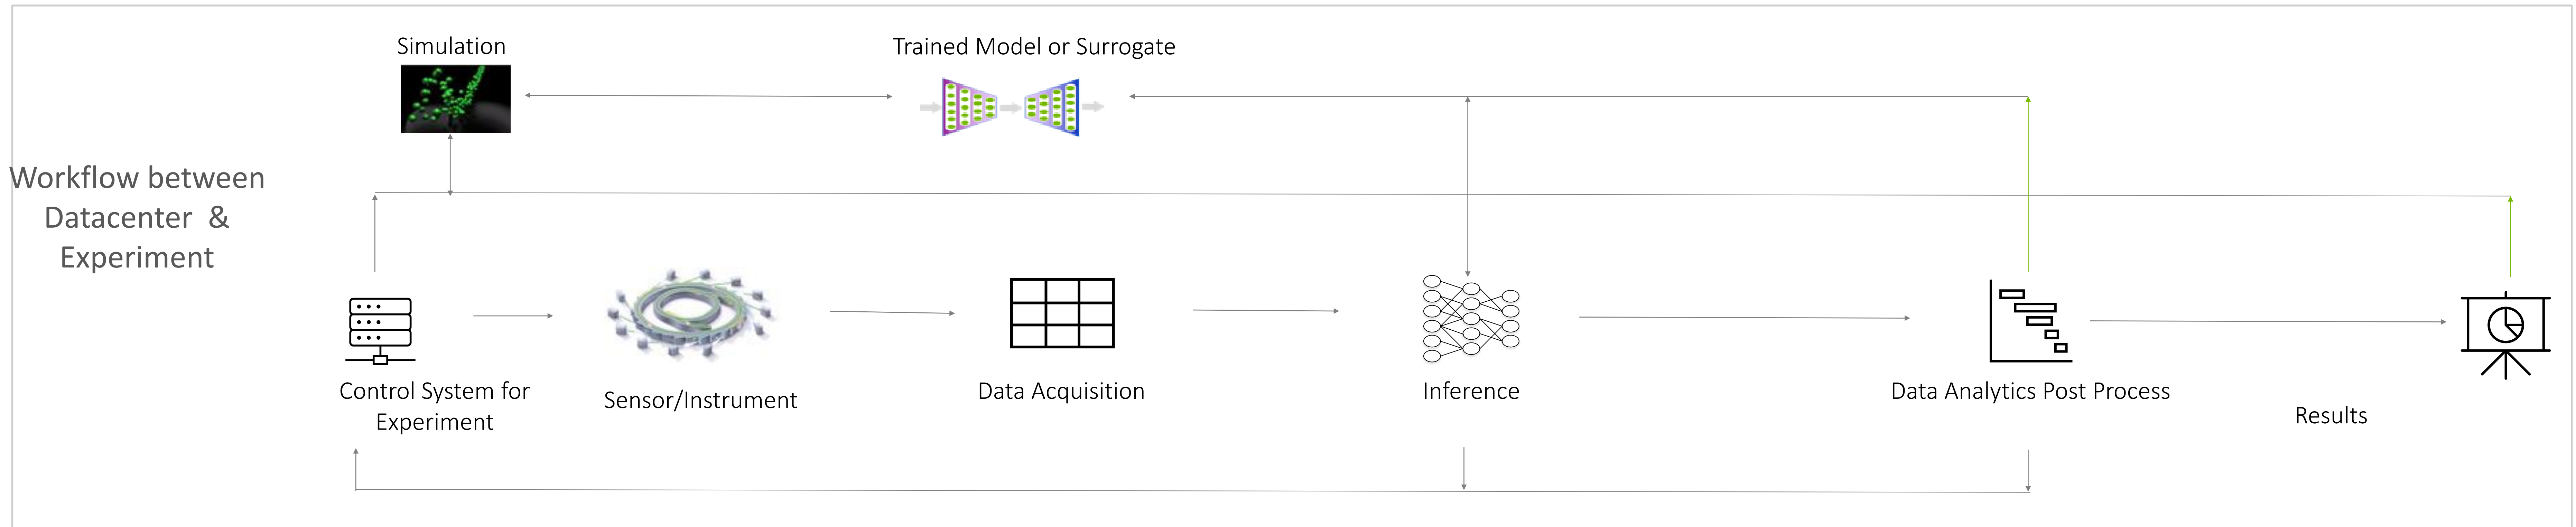


SIMULATION & EXPERIMENT INTEGRATED WITH ML & AI



Integrated workflow with real-time analysis, steering and visualization for human in the loop

GLOBAL ENGAGEMENTS ON CONVERGED WORKFLOWS



ALS/LBNL

Optimizing Ptychography pipeline



CNMS/ORNL

Automating Microscopy



DIAMOND LS UK

Accelerated ptychography and tomography



DIII-D/GA

AI surrogate, Simulation, Control Digital Twin



LHCb/CERN

Design complete for using NVIDIA A40 for HLT




HED Physics/LLNL

HRR HED laser plasma experiments workflow

EXPANDED ECOSYSTEM CHALLENGES

10X - 100X MORE DATA

50+ GIANT SCALE INSTRUMENTS WW



ELT ESO



ALS @ LBNL



LIGO



APS @ ANL



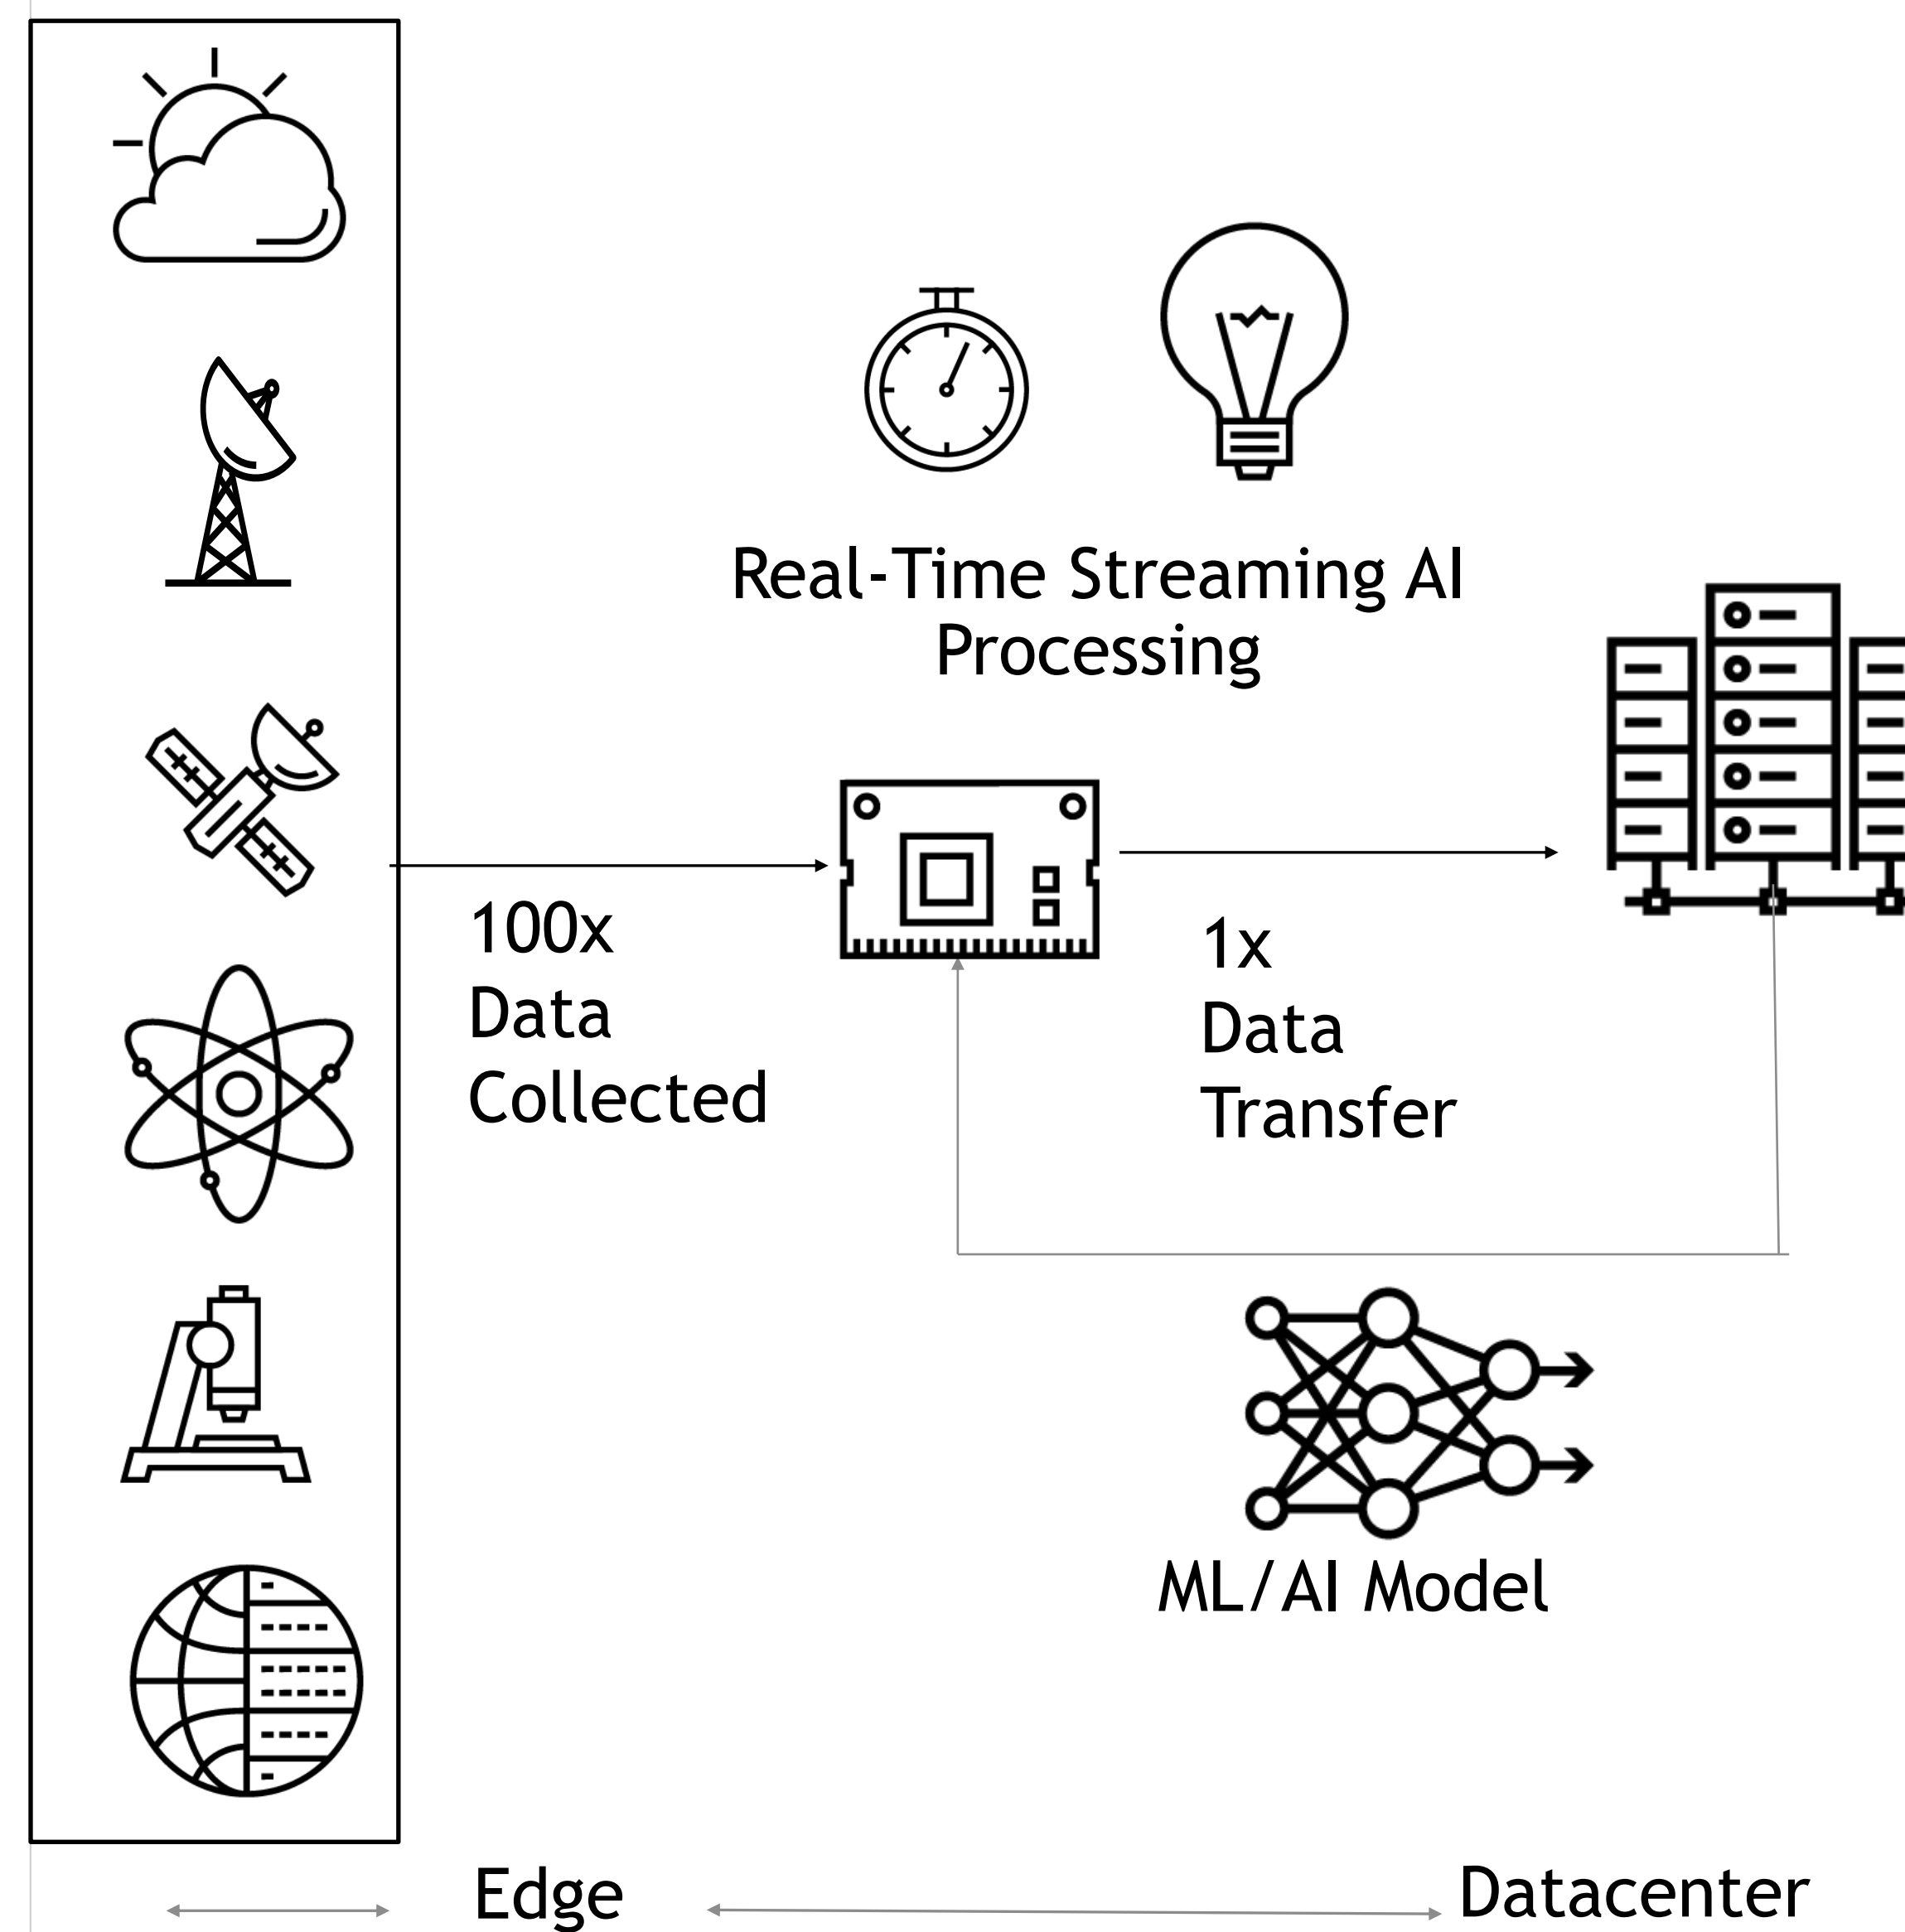
SKA



Diamond, UK

AI SUPERCOMPUTING AT THE EDGE

ENABLES REAL-TIME INSIGHTS AND CONTROL




STREAMING DATA POSES NEW CHALLENGE

FOR DATA SCIENTISTS, RESEARCHERS AND DEVOPS


Streaming Data Performance



Developer Ease-of Use




Easily Scale Implementation



Combining multiple datastreams



LHCB INTRODUCES ALLEN FOR REAL TIME TRIGGER AND TRACKING AT THE EDGE



Challenge

Apply conventional FFT and KALMAN Filter methods to perform trigger and tracking in the same process step for the LHCb Upgrade

Solution

An optimized suite of algorithms was developed by CERN with support from NV to

Impact


Throughput with trigger and tracking of $>60\text{kHz}$ was demonstrated

The full experiment can be supported with 500 GPUs


Run 3 commenced this quarter with 10x processing throughput



GYROKINETIC FUSION REQUIRES THE EXPANDED ECOSYSTEM AT THE EDGE

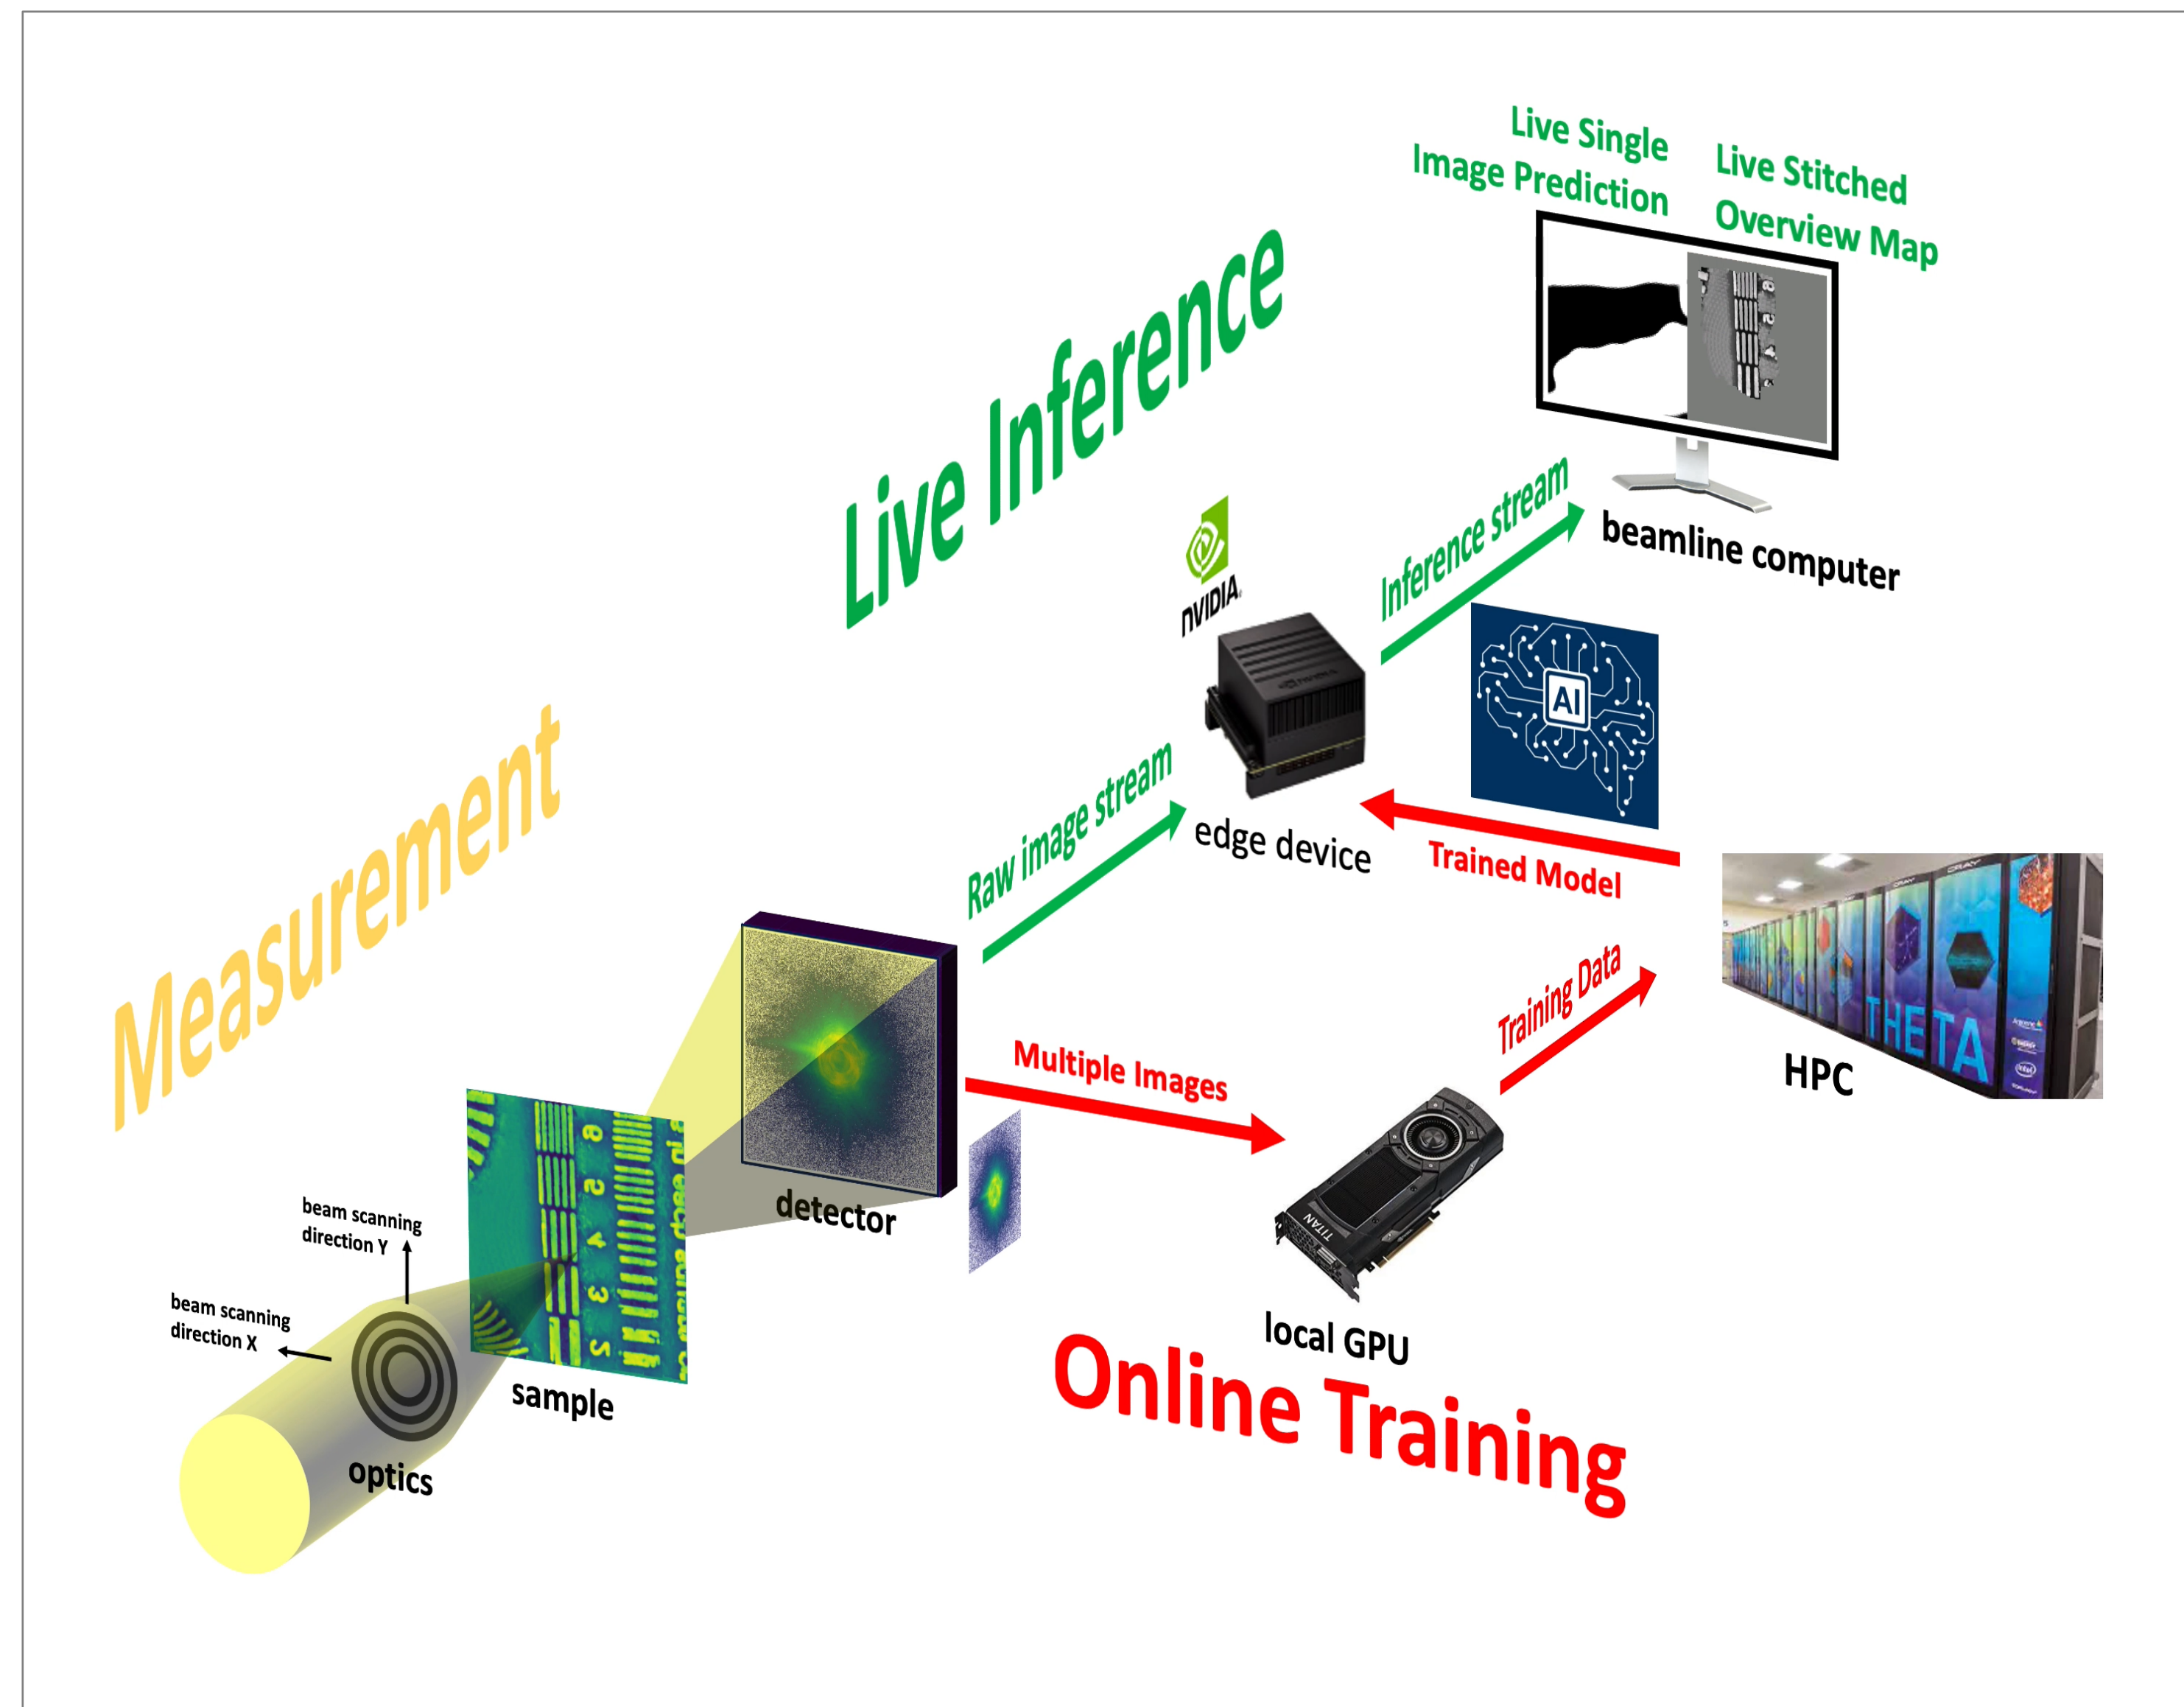


- Accelerated Simulations
 - GTC, XGC, GENE, CGYRO ... 10x + unaccelerated
- Surrogate Models
 - SGTC six orders of magnitude faster than GTC
 - QualikNN 4 orders of magnitude faster than Qualikiz
- Control System Prediction for Disruption at the Edge
 - At DIII-D
 - FRNN 86% accuracy based on diagnostics from JET experiment with live testing at DIII-D underway
 - GatedRNN and Random Forest ML in control system now
 - At TCV
 - Reinforcement Learning applied with DeepMind and DIFFER
- Digital Twin
 - Early Demonstration with MAST Experiment at UK AEA

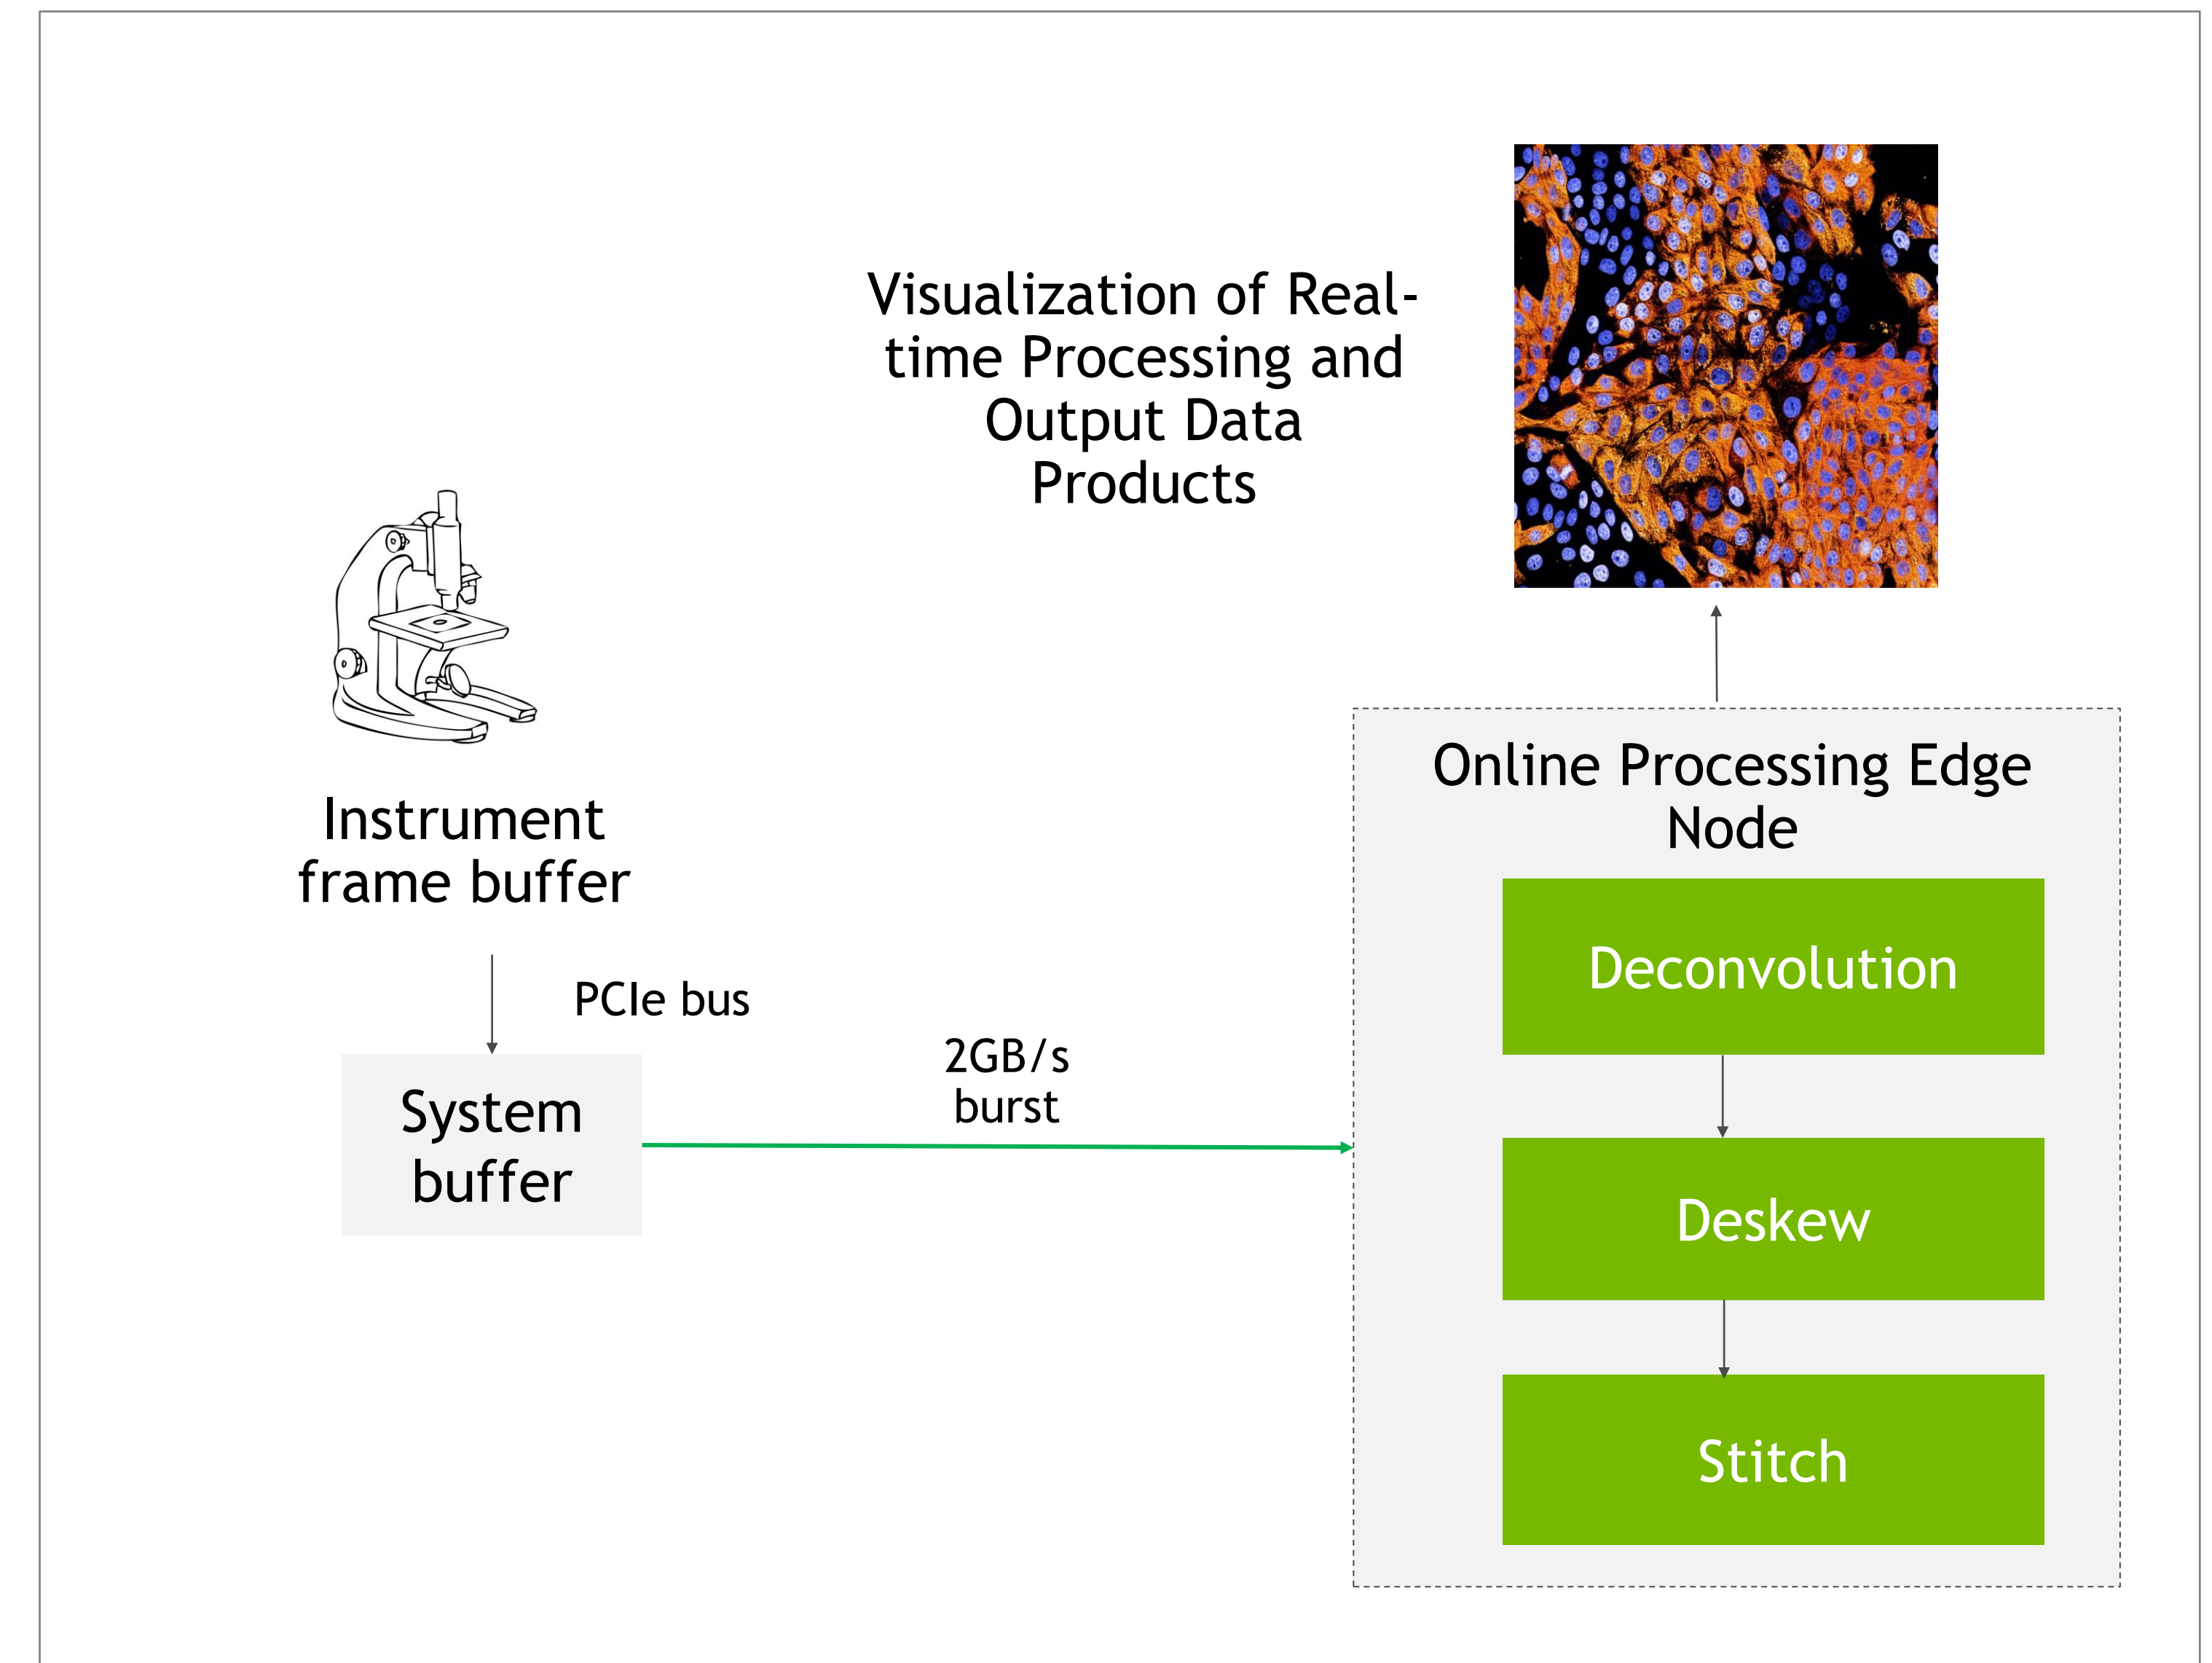


SUPERCHARGING SCIENCE EXPERIMENTS AND INSTRUMENTS

ANL/ APS ACCELERATES X-RAY PTYCHOGRAPHY 300X WITH PTYCHONN



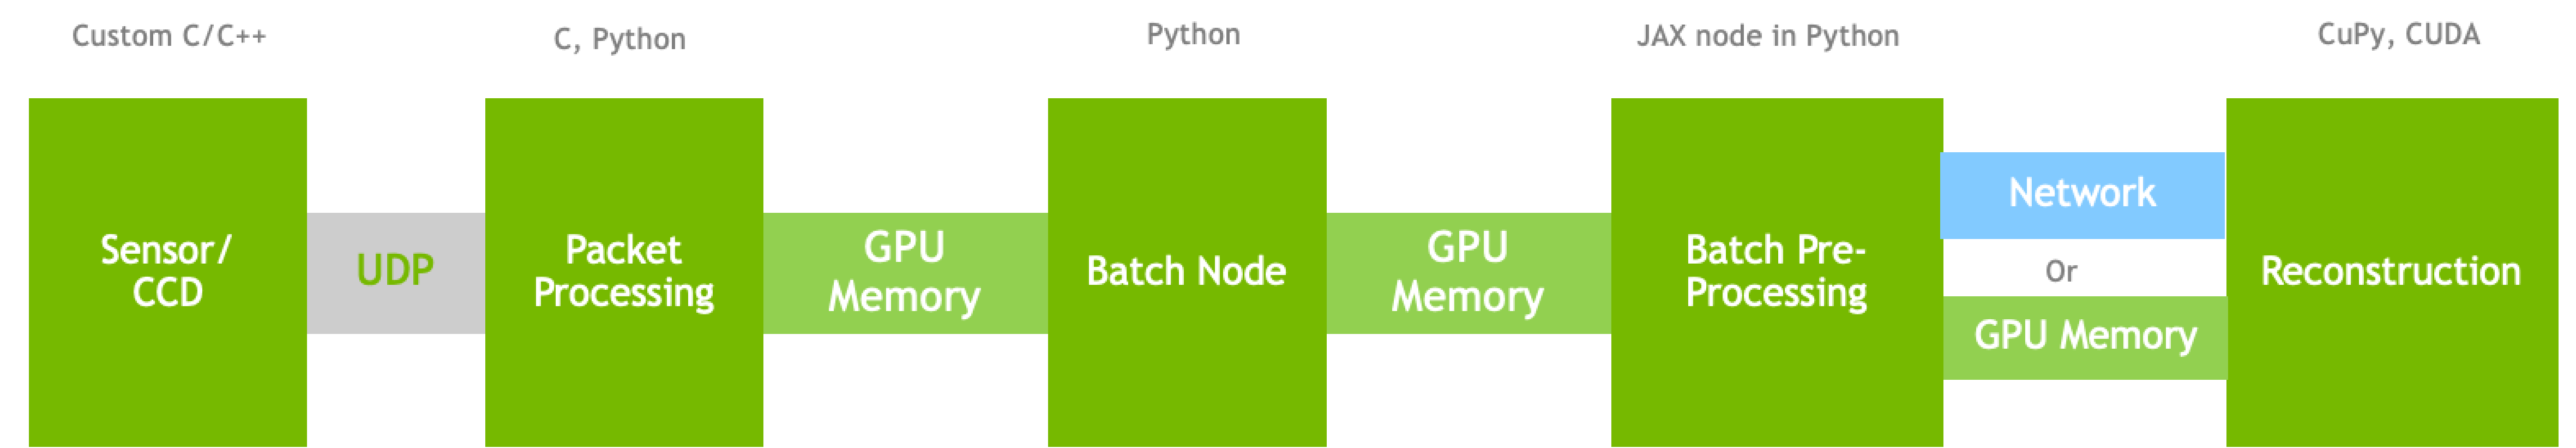
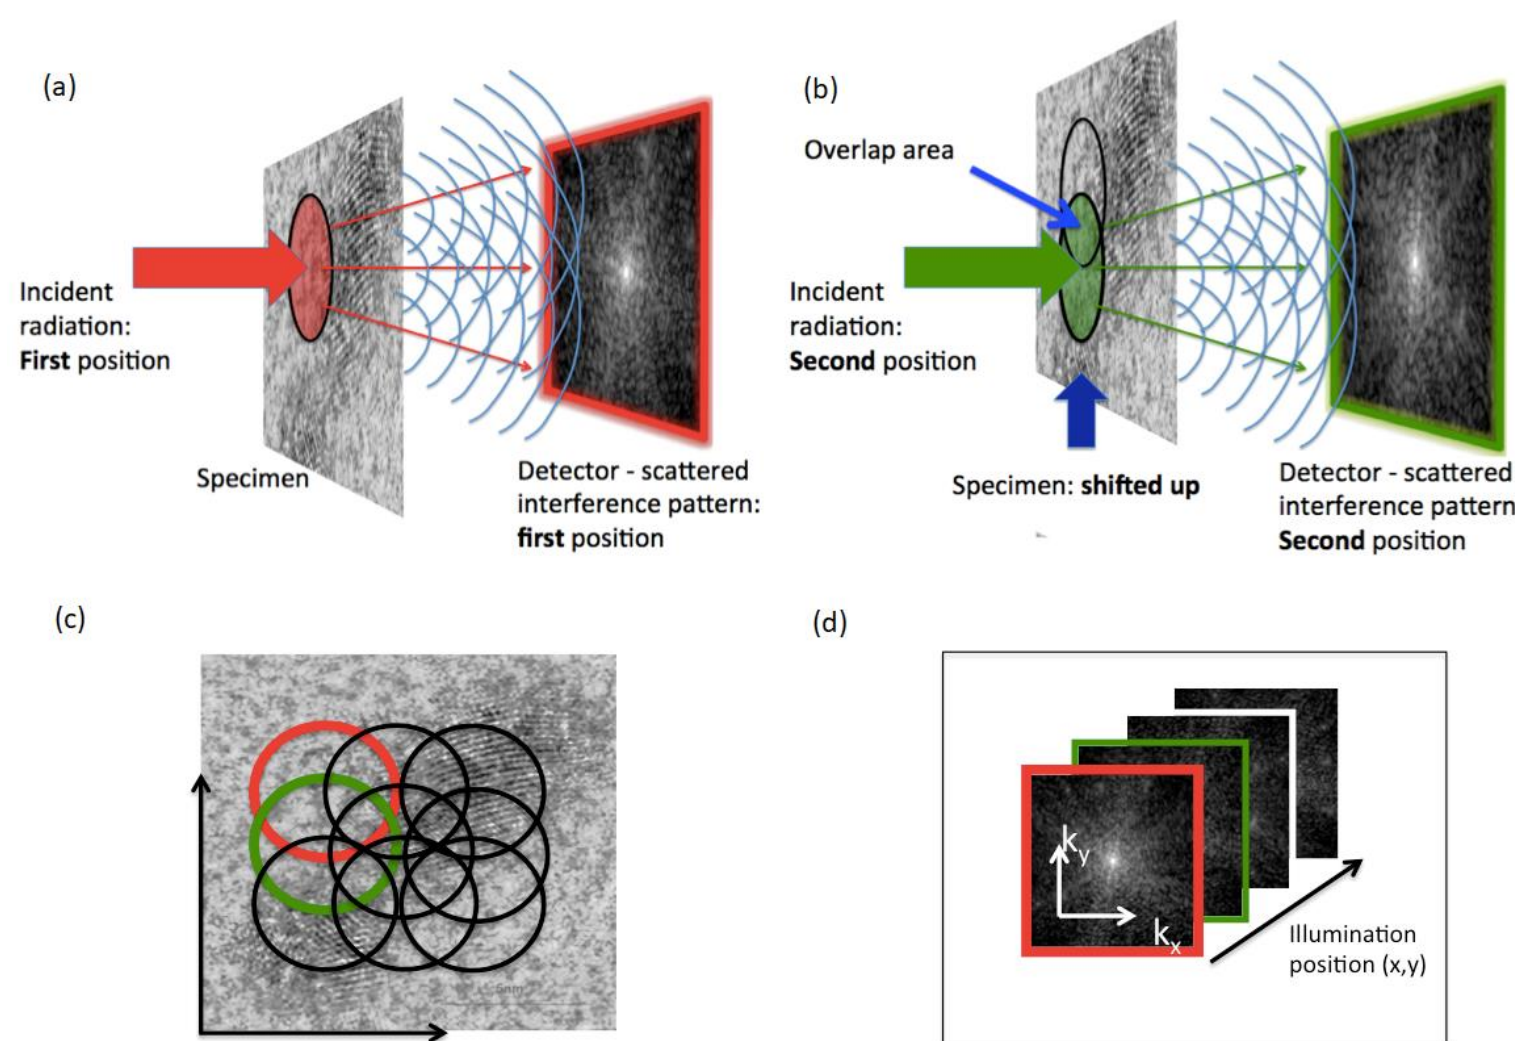
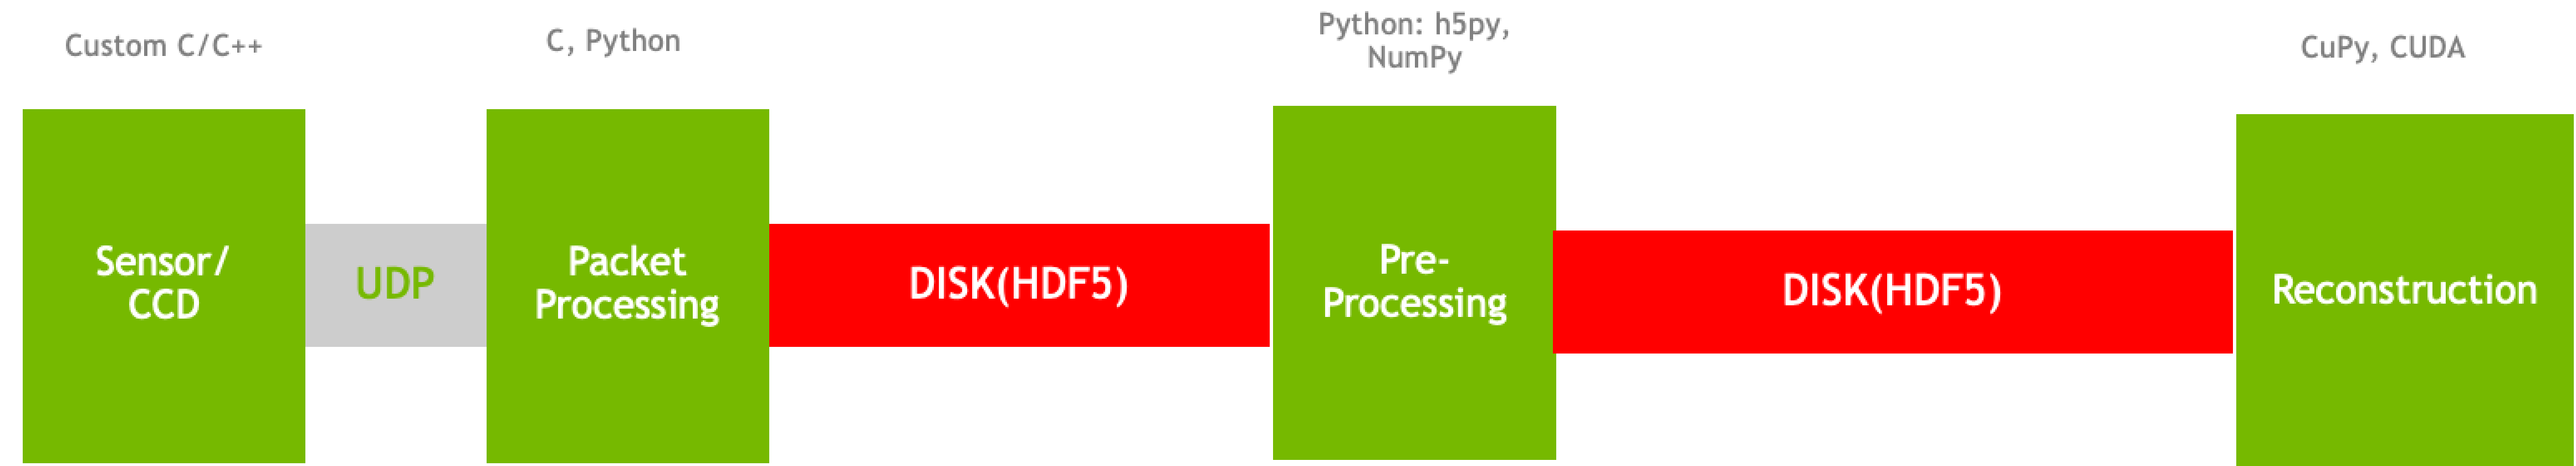

ADVANCED BIOIMAGING CENTER @UC-BERKELEY REAL TIME LIVE CELL IMAGING LIGHT SHEET MICROSCOPY



PtychoNN paper: [AI-enabled high-resolution scanning coherent diffraction imaging](#)

Link to keynote video - <https://youtu.be/rXG27G3bWzY>

ADVANCED LIGHT SOURCE : @ LBNL



From ~105 seconds to 12 seconds

- For more details watch : [Accelerating Sensor Processing Pipelines with NVIDIA Toolkits](#)
- The GTC talk may reference SRF as "Neo" which was the internal name used during initial development

COMPOSING AN HPC STREAMING DATA PIPELINE USING STREAMING REACTIVE FRAMEWORK (SRF*)




SRF* is a **reactive, network-aware, flexible, and performance-oriented** streaming data framework that standardizes building modular and reusable pipeline mixing C++, Python, JAX

- Asynchronous computation and mitigation of I/O and GPU blocking
- Distributed computation with message transfers over RMDA using UCX
- Dynamic reconfiguration to scale up and out at runtime
- Designed to mitigate backpressure with concurrent blocking queues between stages
- Hybrid HPC and Cloud Native

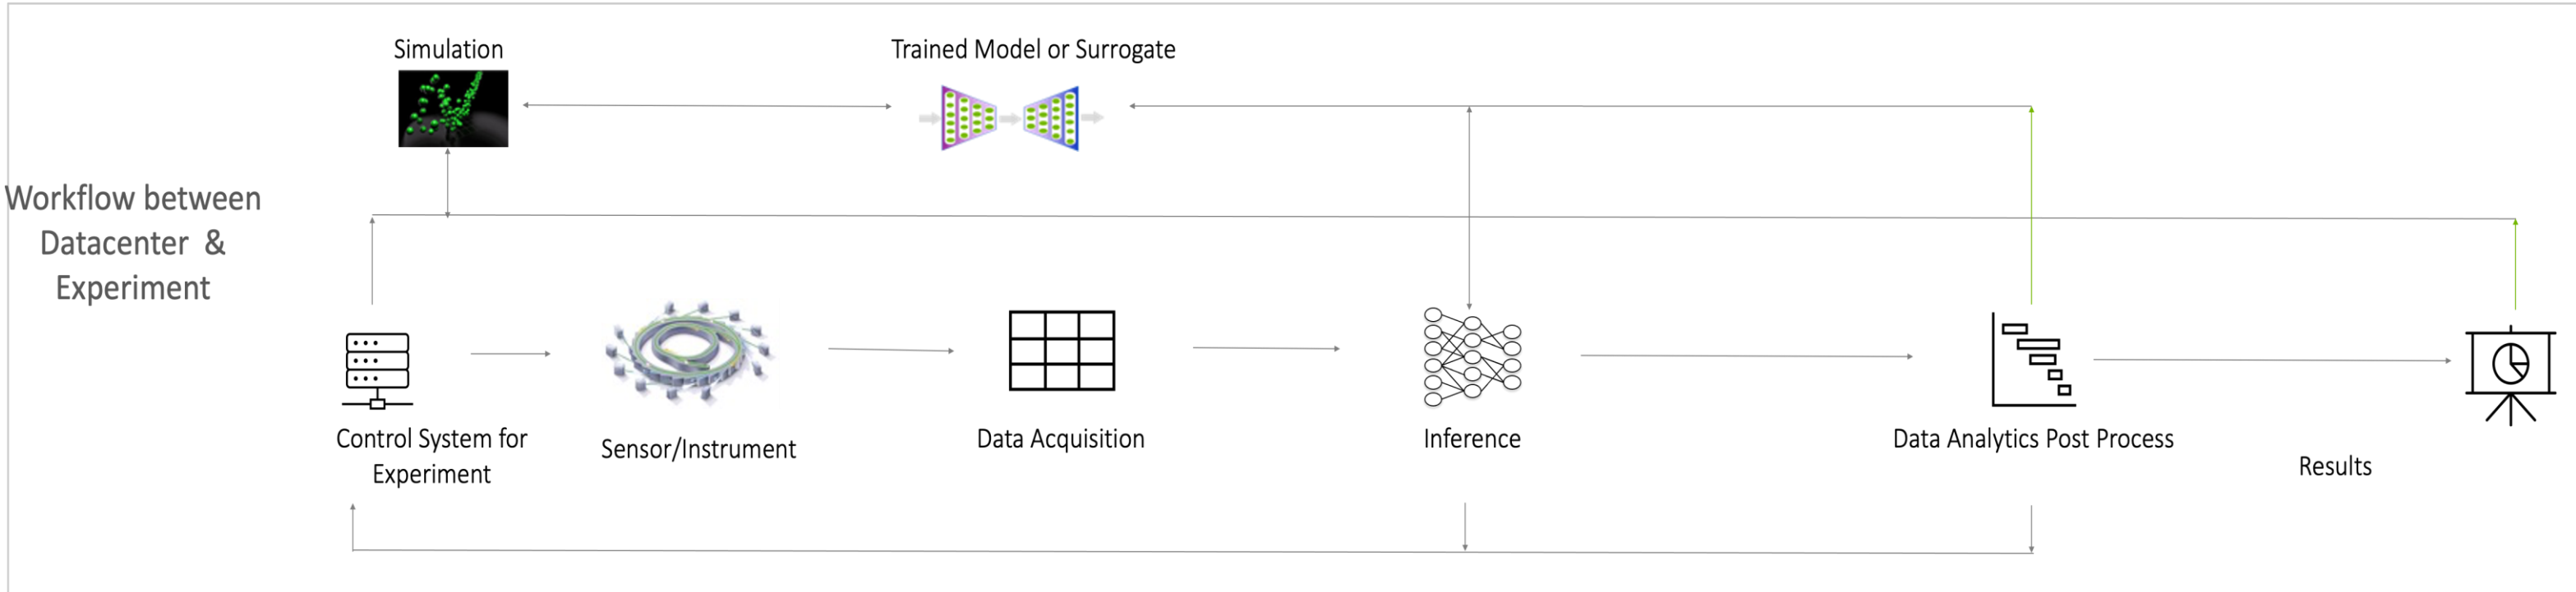
*SRF is under development. Final name subject to change

UKAE EARLY VERSION OF DIGITAL TWIN FOR TOKAMAK REACTOR

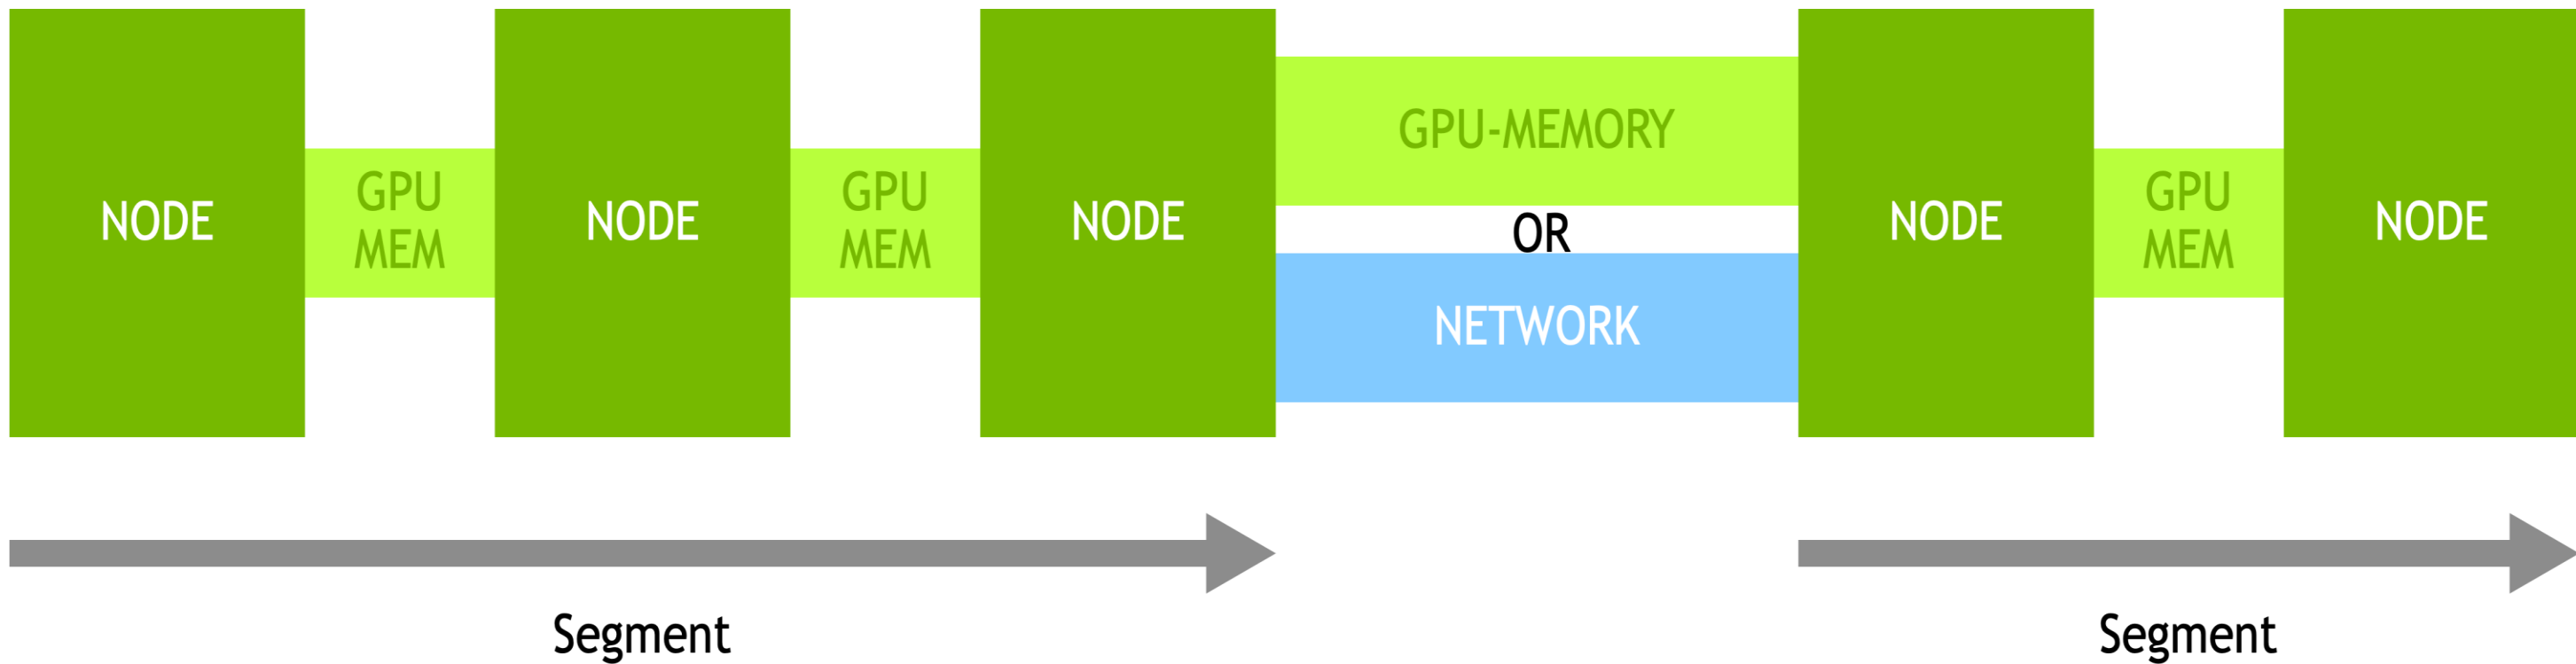


PREPARING FOR THE NEXT DECADE OF SCIENTIFIC COMPUTING


INTEGRATING THE SIMULATION +AI AND EXPERIMENT WORKFLOW




ACCELERATING THE SENSOR / EXPERIMENT DATA PROCESSING




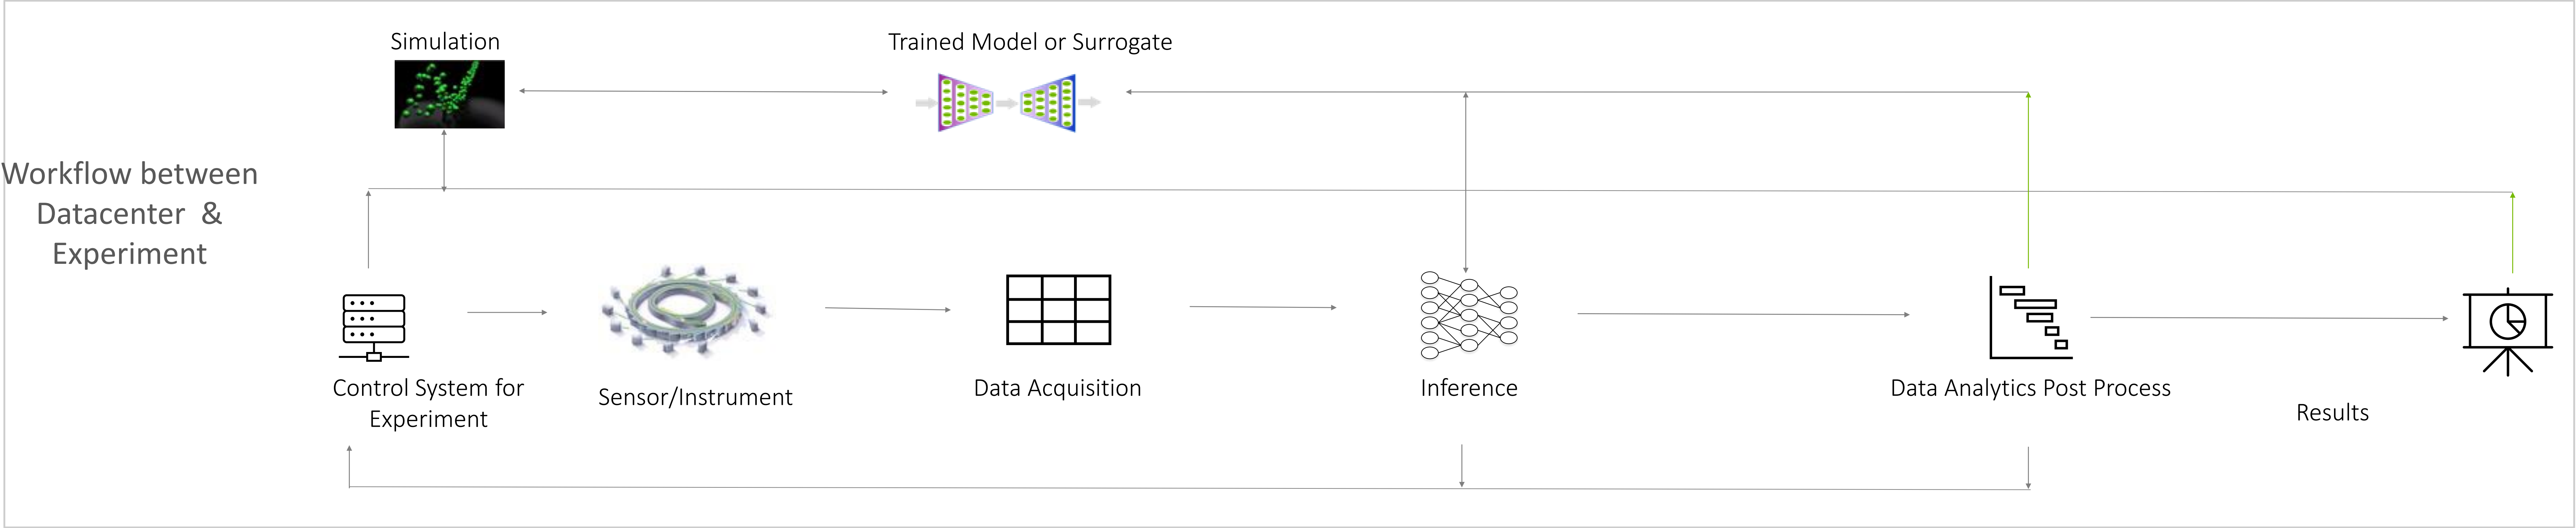
BUILDING A DIGITAL TWIN TOWARDS A SCIENCE GRAND CHALLENGE





nVIDIA®


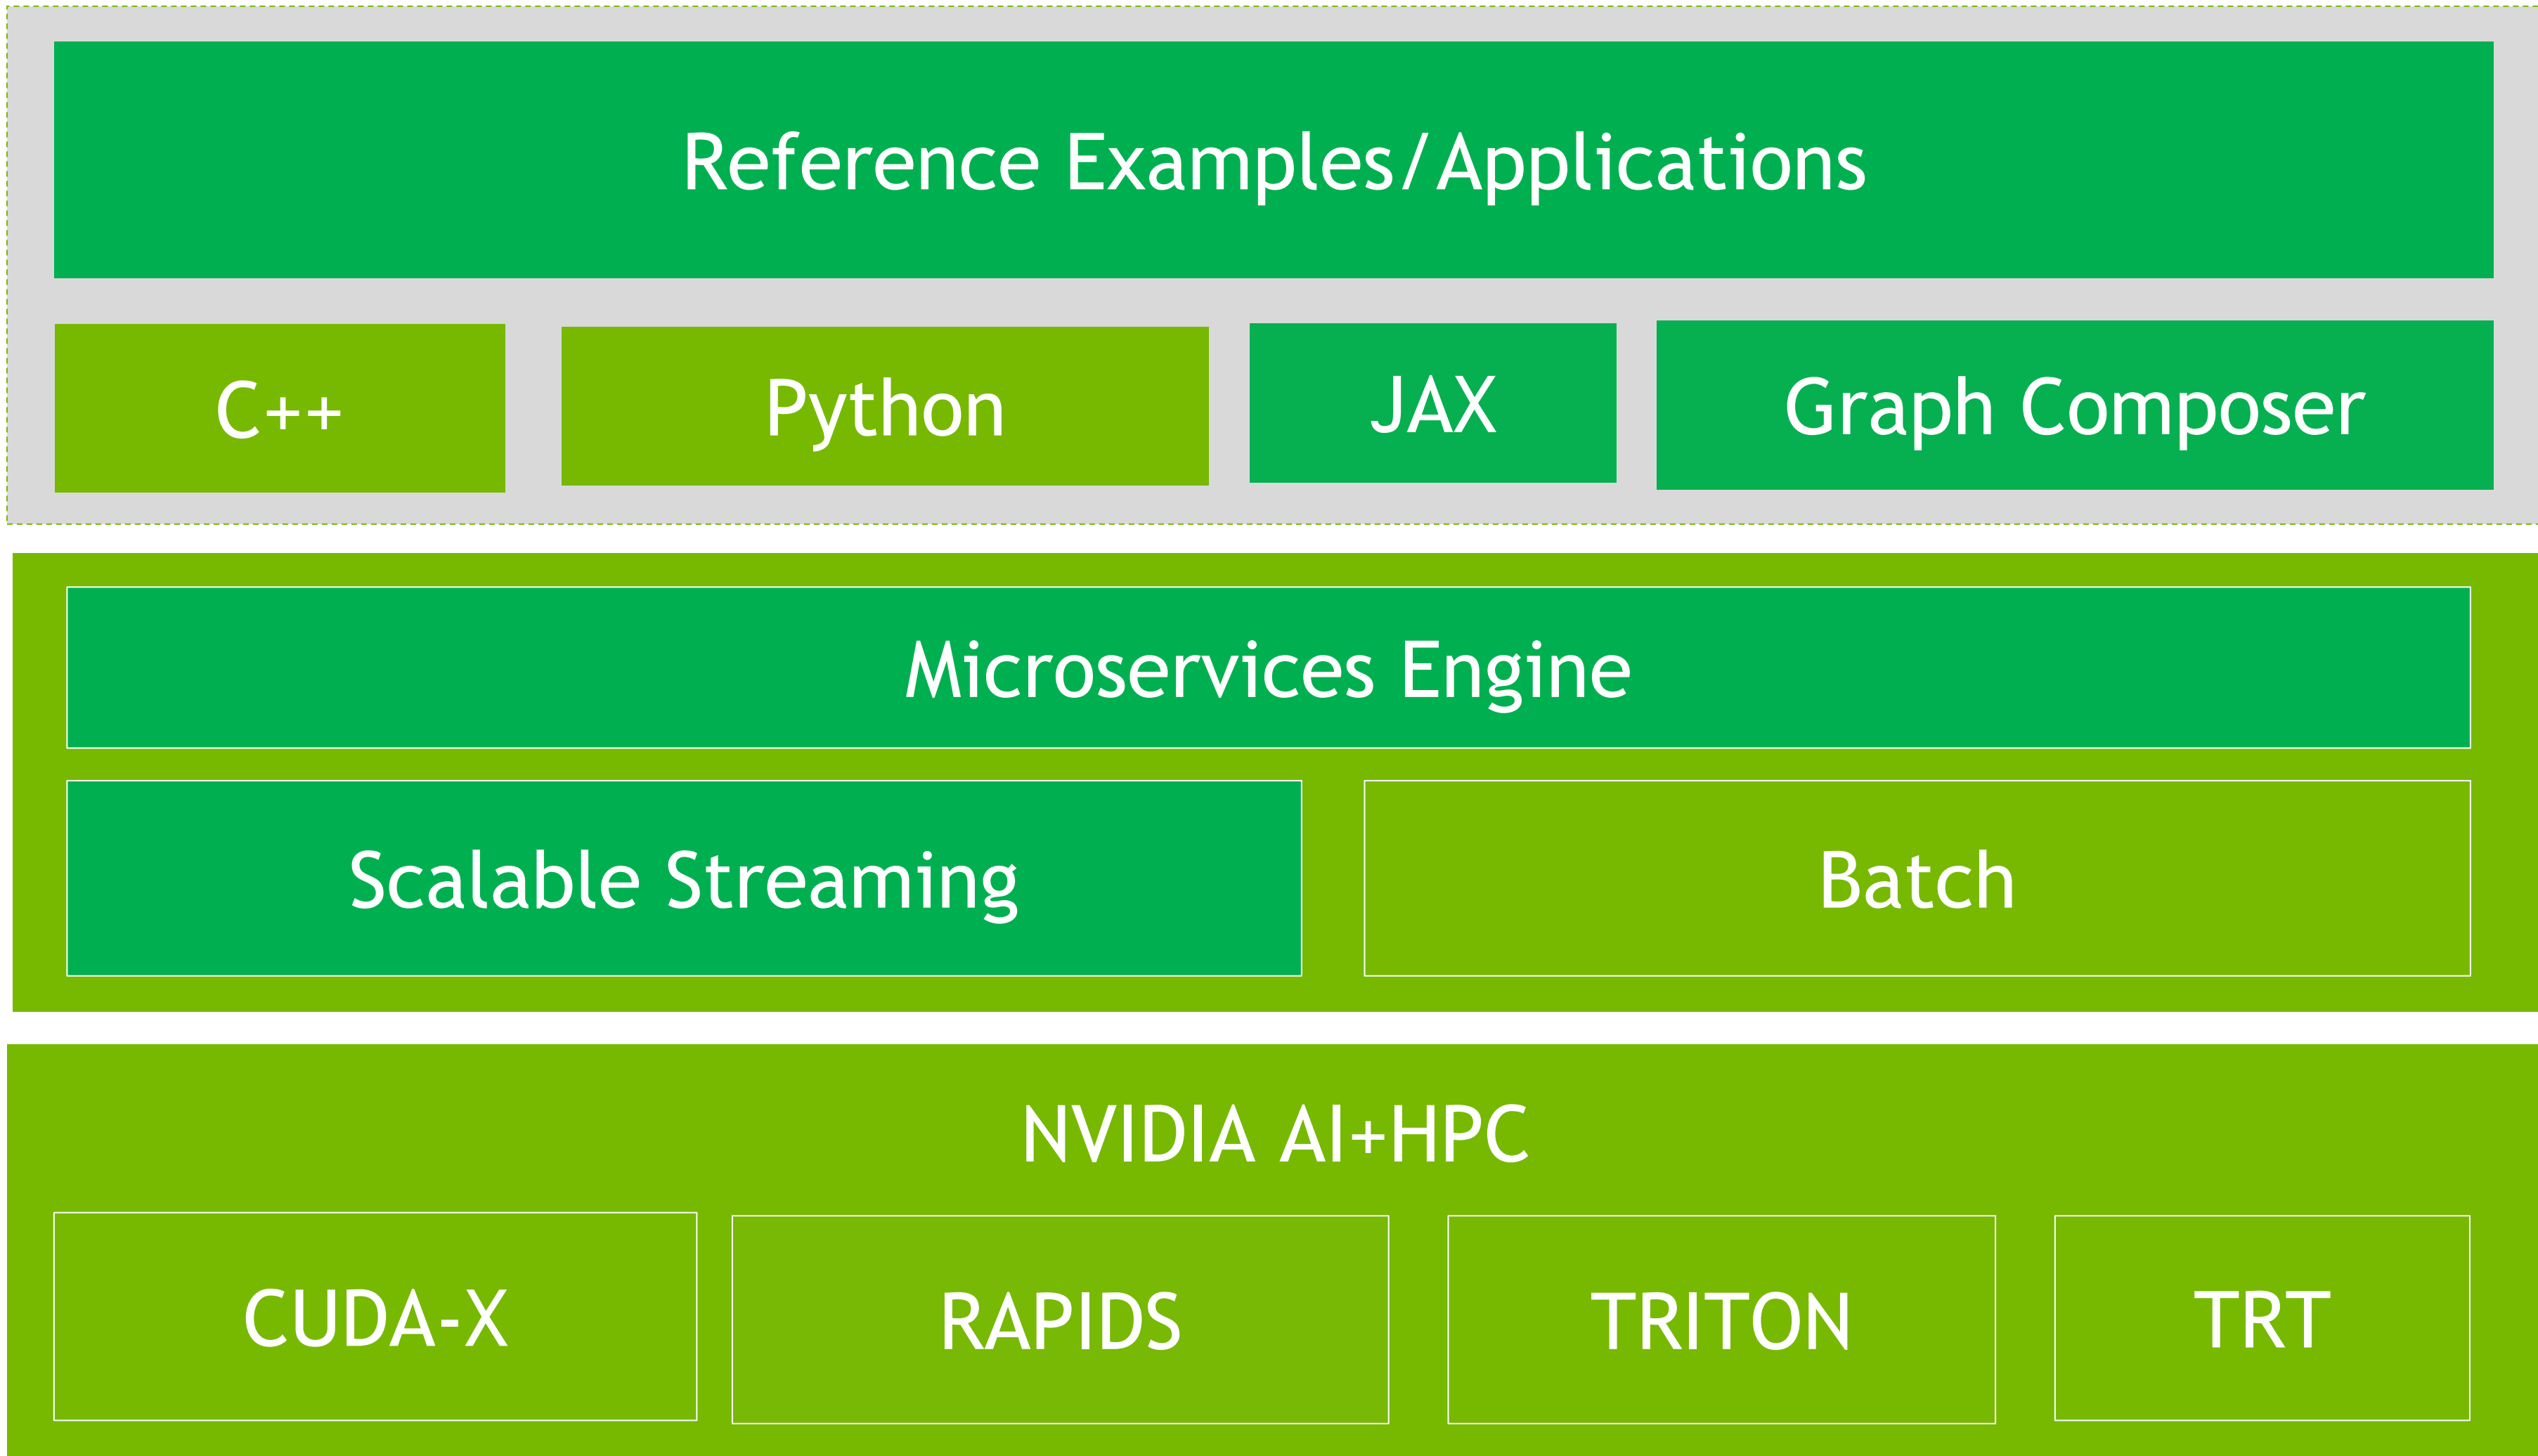

WORKFLOW TO WORKSTREAMS




Optimize and bring the best solution to that create the integrated workflow

NV SW	RAPIDS Triton/TensorRT Holoscan, Morpheus, Issac, UCF	Aerial DOCA Morpheus	HPC SDK, RAPIDS, DL FW, Omniverse, Modulus, UC
NV HW	Jetson, AGX,EGX DGX Station/Server	DPU MetroX	DGX Server/SuperPOD HGX

HOLOSCAN SDK : INTEGRATING DATA STREAMING FROM THE EDGE TO THE DATACENTER



AGX Orin



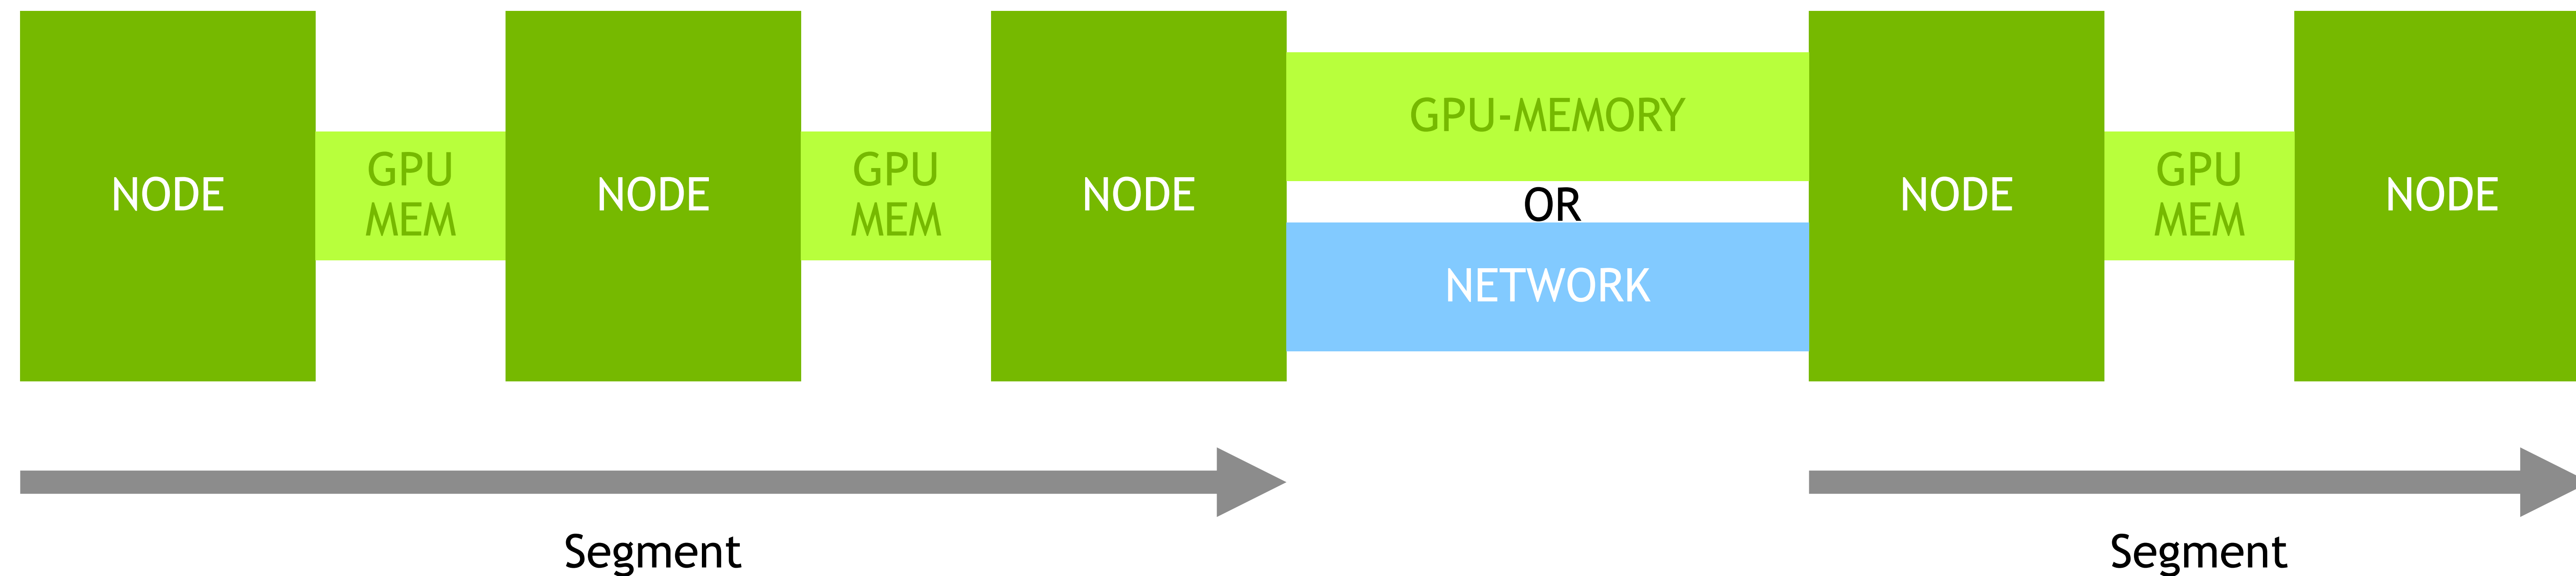
DGX Workstation



DGX/OEM Server

- Build their applications using a mix of C++, Python, JAX
- Develop AI microservices combining low-latency data streaming while passing more complex tasks to data center resources
- Scale from embedded to datacenter

ANATOMY OF A SRF PIPELINE



Definitions

- A SRF pipeline is composed of **Segments**
- **Segments** are composed of **Sources**, **Sinks**, and **Nodes** (Source + Sink)
- **Segments** also guaranteed compute within a single node, can connect nodes via network (Edge, Cloud, or Datacenter), and contain MPI support
- **Nodes** process an input stream, create an output stream, and can be implemented with Python or C++
- Components are linked by **Edges** which are implemented as **Channels**
- **Channels** move data from sources to sinks and provide a backpressure policy

TO LEARN MORE ABOUT SENSOR DATA PROCESSING

- Blog [Facing the Edge Data Challenge with HPC + AI](#)
- GTC Spring 2022 [High Performance Geospatial Image Processing at the Edge*](#)
 - Geospatial image analysis using DPUs in an edge device designed to meet the Size-Weight-and-Power requirements for aircraft deployment.
- PtychoNN paper: [AI-enabled high-resolution scanning coherent diffraction imaging](#)
 - The Advanced Photon Source at Argonne National Laboratory runs PtychoNN on an Orin AGX at the x-ray detector. It is available for use at other light sources around the world.
- GTC Spring 2022 [Accelerating Sensor Processing Pipelines with NVIDIA Toolkits*](#)
 - Faster imaging pipelines by using JAX and SRF to processing streaming data with applications in Ptychography and Microscopy
- See the SRF description above and the [GitHub](#) page

*The GTC talk may reference SRF as "Neo" which was the internal name used during initial development