

# Deploying and Operating Scientific Visualization Tools on HPC Systems

Silvio Rizzi

Computer Scientist

Argonne Leadership Computing Facility

September 26, 2025

# Argonne National Laboratory

## A U.S. Department of Energy National Laboratory

- ◉ The first science and engineering research national laboratory in the U.S.
- ◉ Argonne integrates world-class science, engineering, and user facilities to deliver innovative research and technologies.
- ◉ Argonne creates new knowledge that addresses the scientific and societal needs of our nation.





# Aurora

## Leadership Computing Facility Exascale Supercomputer

Peak Performance  
 **$\geq 2$  Exaflops DP**

Intel GPU  
**Intel® Data Center  
GPU Max Series**

Intel Xeon Processor  
**4<sup>th</sup> Gen Intel XEON  
Max Series CPU**  
with High Bandwidth Memory

Platform  
**HPE Cray-Ex**

### Compute Node

Two 4<sup>th</sup> Gen Intel XEON Max Series CPUs  
Six Intel® Data Center GPU Max Series  
Node Unified Memory Architecture  
Eight fabric endpoints

### GPU Architecture

Intel® Data Center GPU Max Series  
architecture  
High Bandwidth Memory Stacks

### Node Performance

>130 TF

### System Size

>10,000 nodes

### Aggregate System Memory

>10 PB aggregate System Memory

### System Interconnect

HPE Slingshot 11  
Dragonfly topology with adaptive routing

### Network Switch

25.6 Tb/s per switch (64 200 Gb/s ports)  
Links with 25 GB/s per direction

### High-Performance Storage

220 PB  
 $\geq 25$  TB/s DAOS bandwidth

### Software Environment

- C/C++
- Fortran
- SYCL/DPC++
- OpenMP offload
- Kokkos
- RAJA
- Intel Performance Tools

# Scientific Visualization on HPC Resources

- Visit and ParaView as open source tools
- Community efforts evolving for 20+ years
- Built on VTK (Visualization Toolkit)
- Viskores for acceleration



# VTK Basics

- Provides common data models: polygonal data, points, image stacks, unstructured grids
- Parallel readers: read native simulation data and convert to VTK model
- Renderers: project data geometry into 2D images

# DevOps Perspective: Visit and ParaView

- Complex applications with many dependencies
- Most building blocks are open source
- Domain decomposition: MPI ranks manage dataset chunks
- Final images assembled via compositing (collective MPI operations)

# Rendering on a Supercomputer

- Headless rendering (no displays on compute nodes)
- OpenGL if available (requires X server context)
- OSMesa for CPU-based rendering (no GPU/driver)
- OSPRay ray tracing as an alternative (CPU-based)

# Building and Deploying: Overview

- Complex build systems
- *build\_visit* (bash script) for Visit
- ParaView *superbuild* (CMake-based) for ParaView
- Automate retrieval and build of dependencies from source

# Build Components

- CMake, Python, NumPy, SciPy, many Python libraries
- LLVM required for OSMesa
- VTK and Viskores for visualization algorithms
- Numerous parallel data readers
- Example: National labs (e.g., LANL) dedicate staff to maintain ParaView

# Ingesting Data into Visit and ParaView

- Multiple data readers available
- Simulation data models may not map directly
- Example: Nek5000 spectral element data → needs conversion to unstructured grid → large data expansion

# Client/Server Mode: Workflow

- Most common interactive mode
- Client launches connection → login node → queued job on HPC
- Once running, parallel app reverse connects to client
- Typical chain: Client ↔ Login node ↔ Head compute node
- Tools like socat often needed for connection management

# QUESTIONS?

Silvio Rizzi  
[srizzi@anl.gov](mailto:srizzi@anl.gov)

