SIGGRAPH 2013

Shaping the Future
of Visual Computing

NVIDIA IndeX™ - Enabling Interactive and Scalable Visualization for Large Data

Marc Nienhaus, NVIDIA IndeX Engineering Manager and Chief Architect

NVIDIA IndeX™ Positioning

- NVIDIA IndeX is a commercial software product initially developed to serve the Hydrocarbon market. It is a cluster-based scalable software Platform as a Service (PaaS) ready for the cloud and enables distribution of large-scale data for compute and high quality visualization of volumetric and surface data with interactive frame-rates.

http://www.nvidia-arc.com/products/nvidia-index.html

NVIDIA IndeX – Enabling Interactive and Scalable Visualization for Large Data

- NVIDIA IndeX software leverages GPU-clusters for scalable large-scale data visualization
- NVIDIA IndeX is a GPU-cluster aware solution for interactive visual computing
- NVIDIA IndeX is a commercial software solution available and already deployed by customers for data interpretation

http://www.nvidia-arc.com/products/nvidia-index.html

Example: Exploration in Hydrocarbon Industries

- Scanning earth's subsurface structure
- Cost-effective drilling for oil reservoirs
- Acquisition and preprocessing of subsurface data
  - Huge (peta bytes) subsurface dataset sizes
  - Automatic data processing

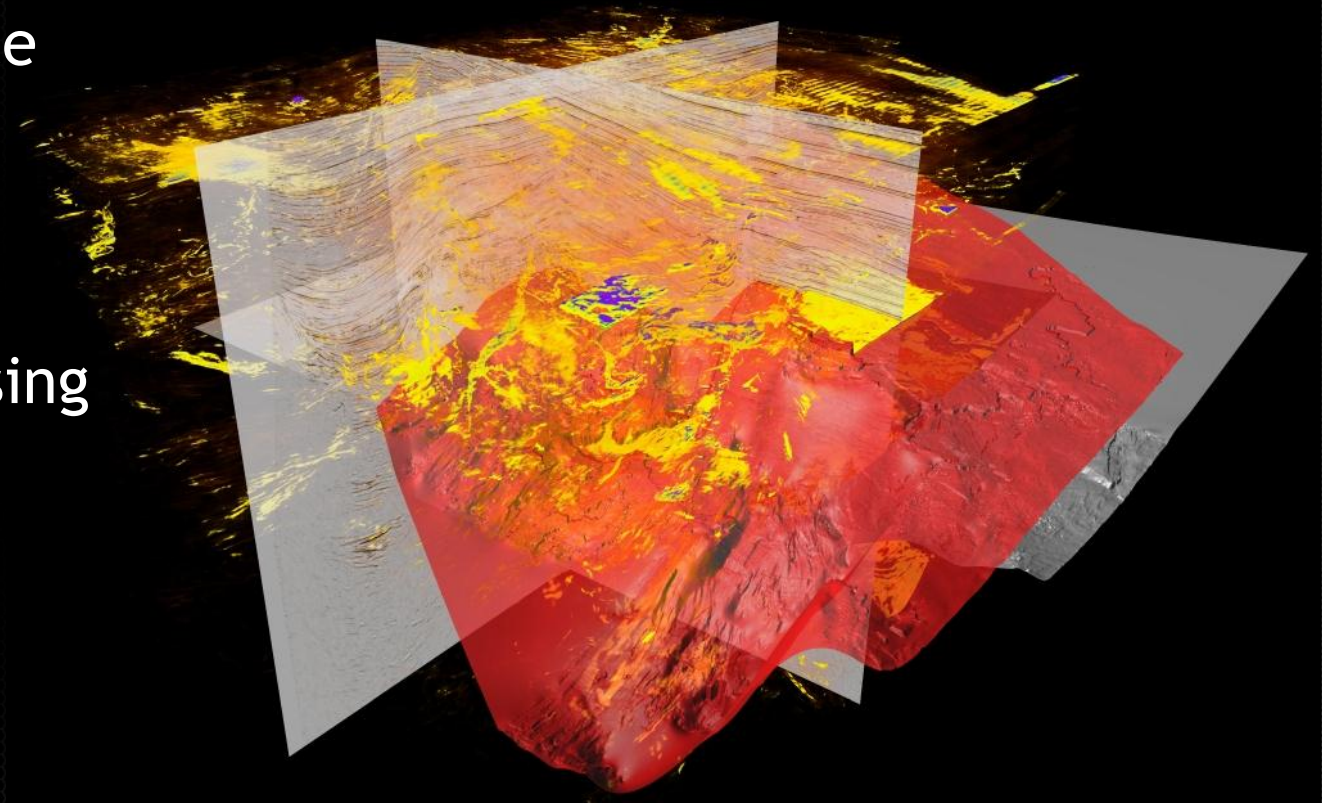

Visualization of a seismic volume with embedded height field and slices

Special thanks to Crown Minerals and the New Zealand Ministry of Economic Development
for allowing us to display this Taranaki Basin dataset.
Crown Minerals manages the New Zealand Government's oil, gas, mineral and coal resources.
More information is available at: www.crownminerals.govt.nz

Efficient Data Interpretation

- Knowledge and experience of experts in this field
- Visually assess subsurface data
- Interactive exploration
  - Real-time frame rates
  - Visual quality

- Especially in Oil & Gas domain

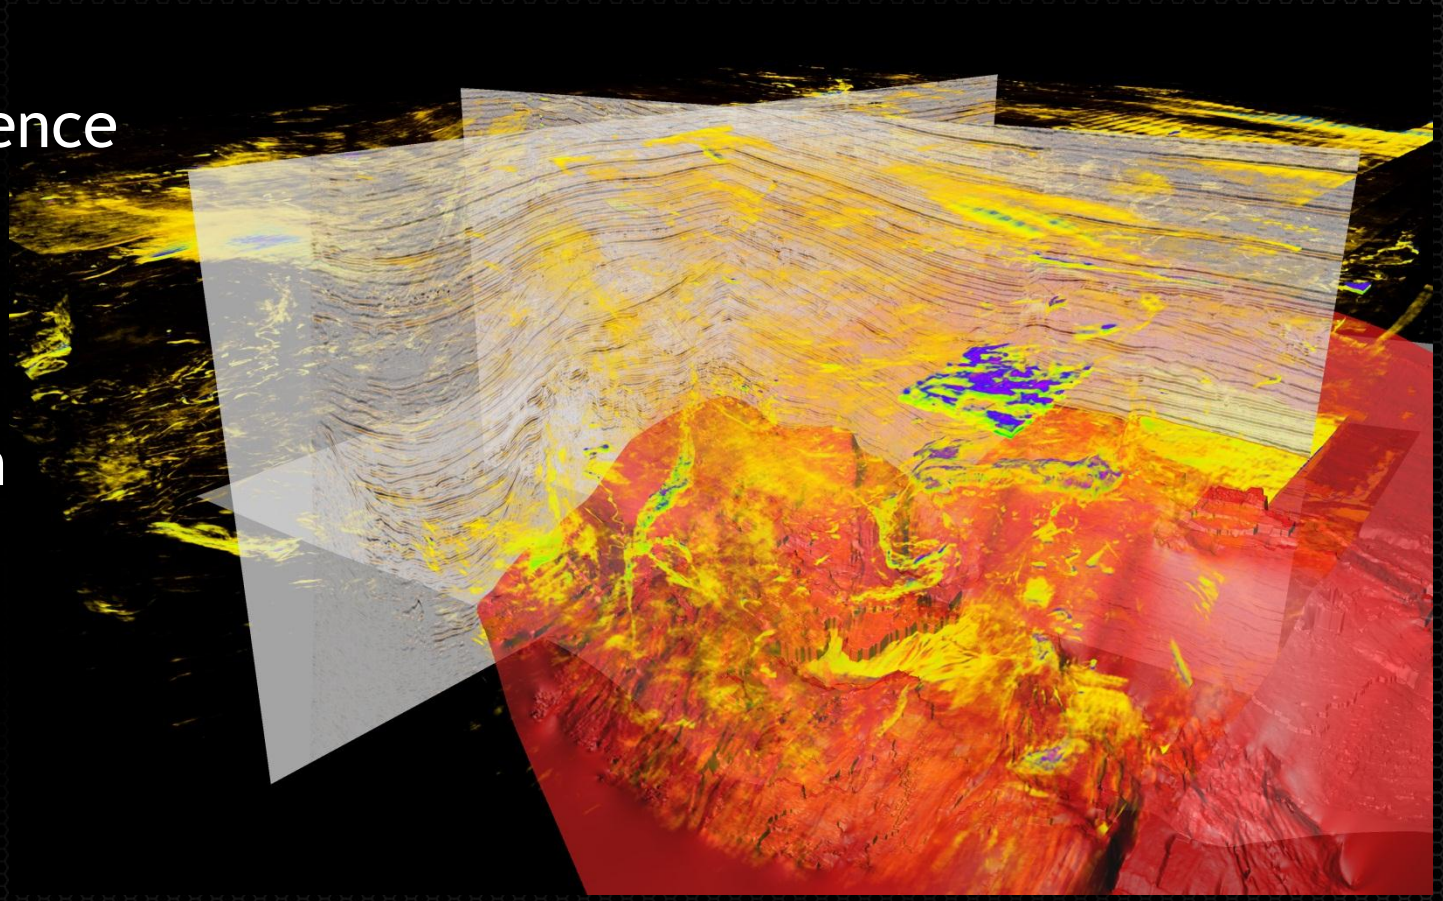

Visualization of a seismic volume with embedded height field and slices

Special thanks to Crown Minerals and the New Zealand Ministry of Economic Development
for allowing us to display this Taranaki Basin dataset.
Crown Minerals manages the New Zealand Government's oil, gas, mineral and coal resources.
More information is available at: www.crownminerals.govt.nz

NVIDIA IndeX - Scalable Interactive Large-Scale Data Visualization

- Distributed Rendering on GPU clusters
  - Supports today's and tomorrow's huge dataset sizes
- Real-Time Rendering
  - 260 GB volume
  - 14 cluster machines with 4 Tesla K10
  - 13 frames per second

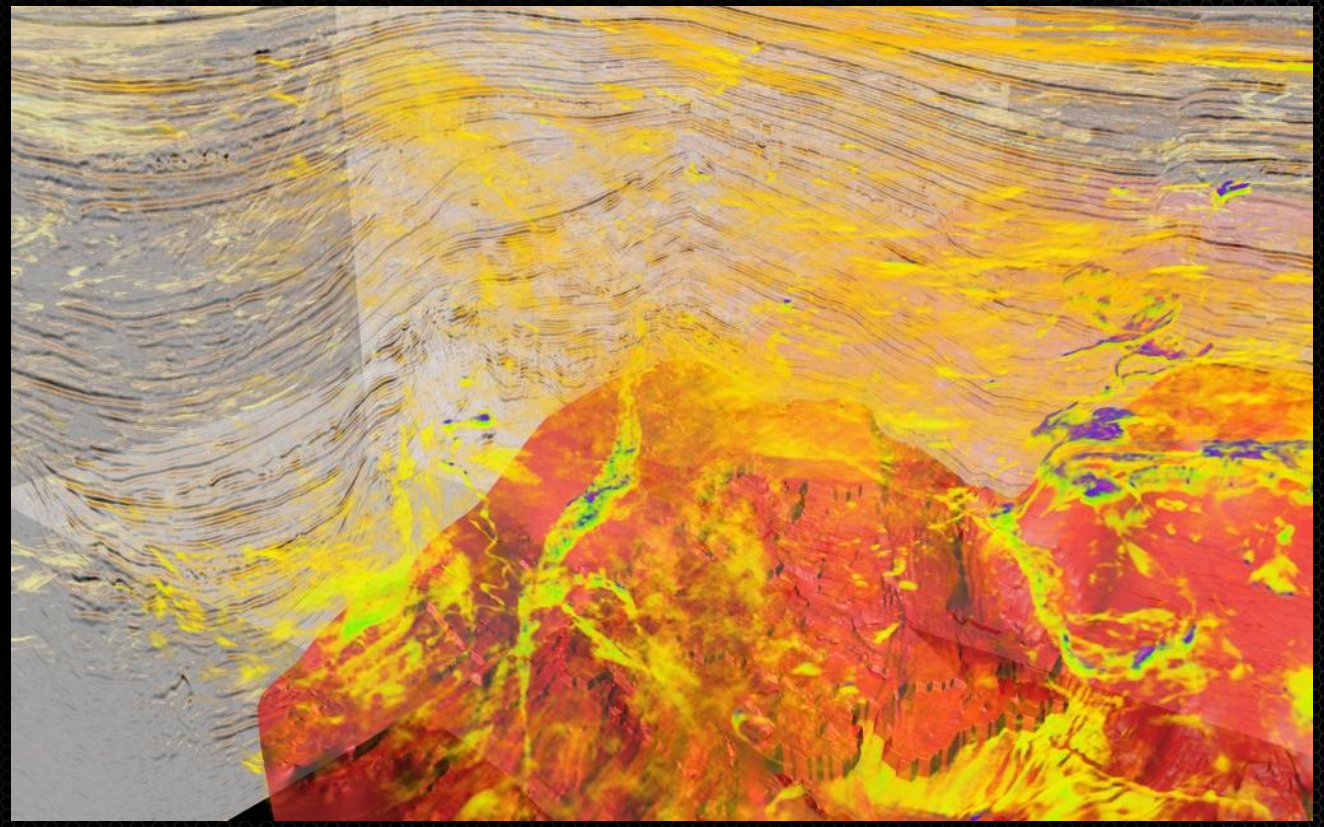

Visualization of a seismic volume with embedded height field and slices

Special thanks to Crown Minerals and the New Zealand Ministry of Economic Development
for allowing us to display this Taranaki Basin dataset.
Crown Minerals manages the New Zealand Government's oil, gas, mineral and coal resources.
More information is available at: www.crownminerals.govt.nz

Visual Quality and Accuracy

- Visualization at original data resolution
  - Avoiding distractions e.g., popping artifacts due to level-of-details
  - Highly accurate visual assessment
- Depth-correct transparency rendering
  - Example: height field embedded into volume

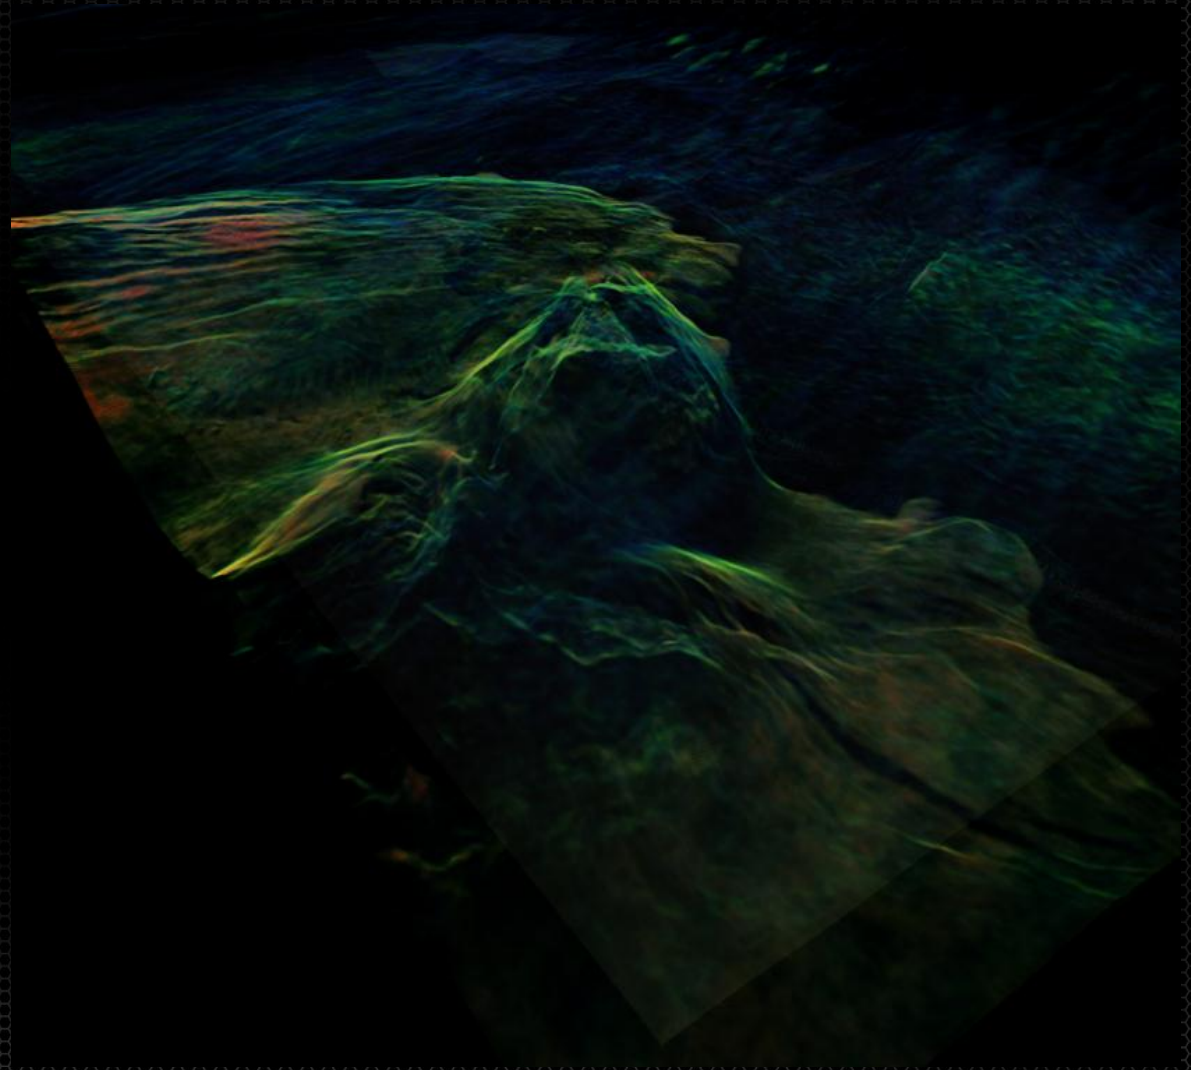

Visual Quality

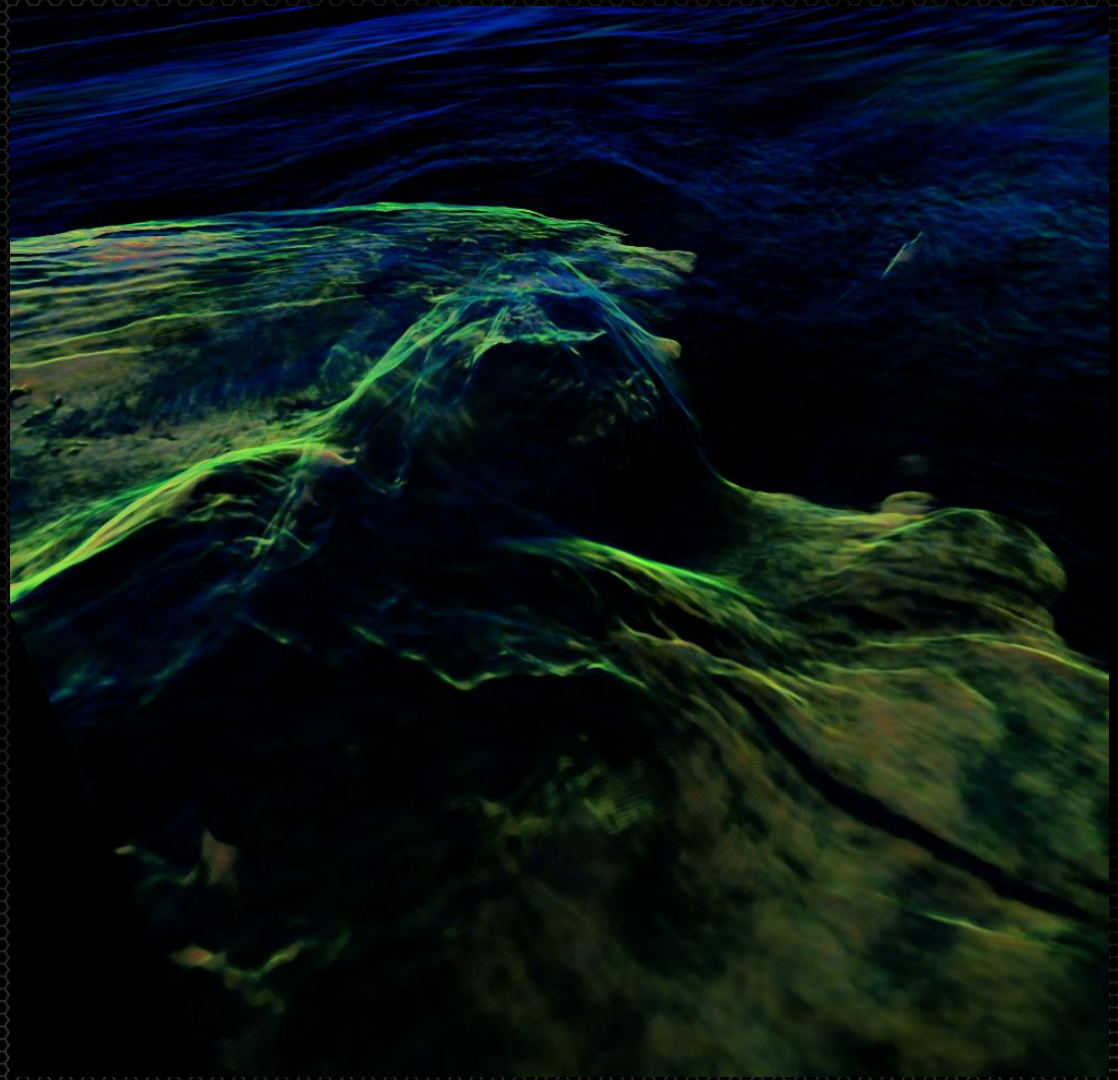

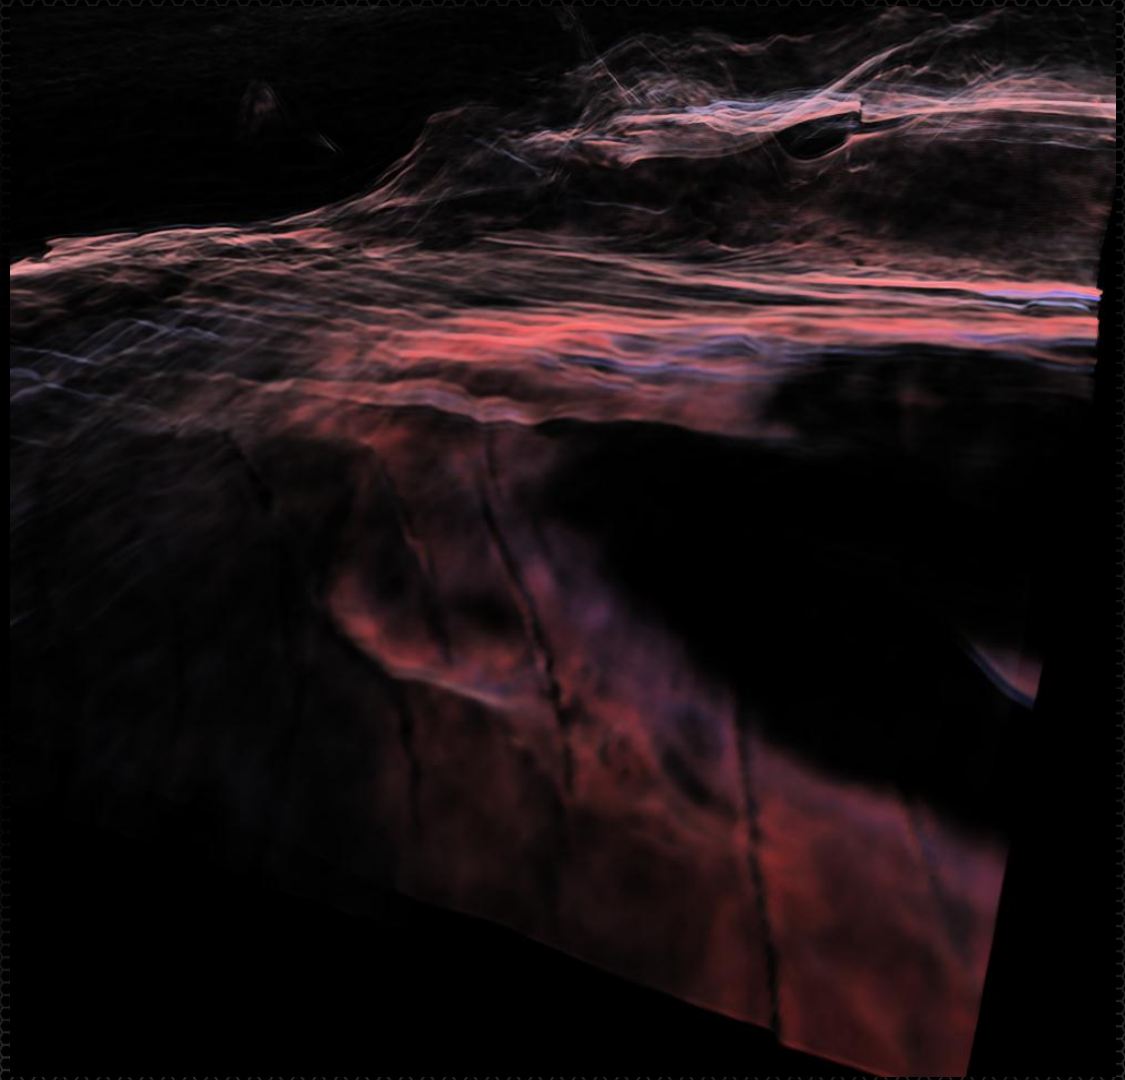

Sort-Last Approach for Distributed and Scalable Rendering

- Object-space subdivision
- Later compositing of intermediate renderings

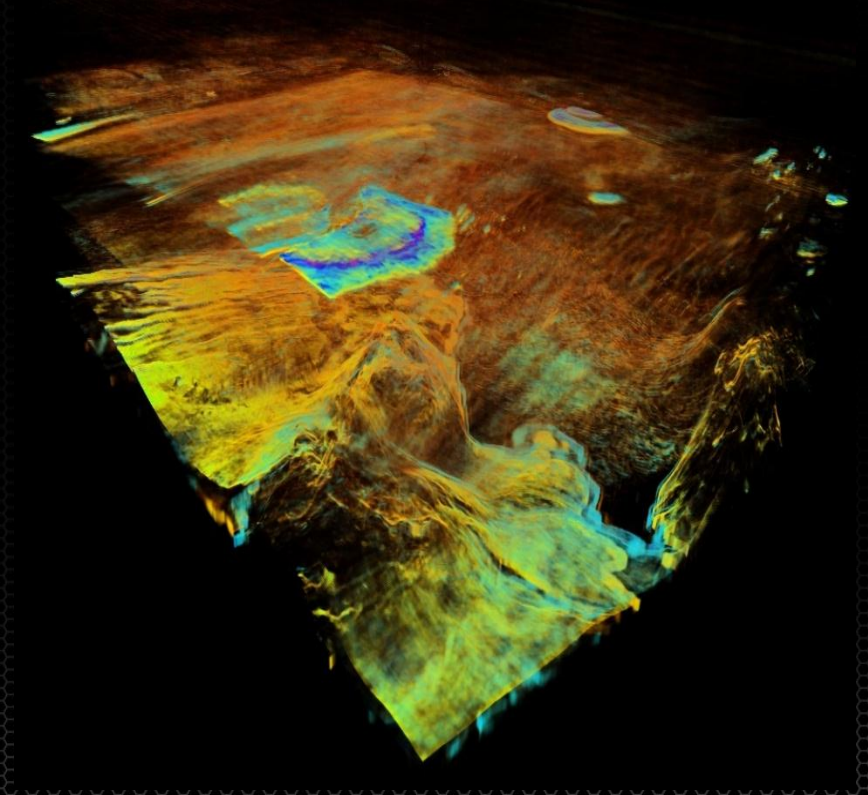

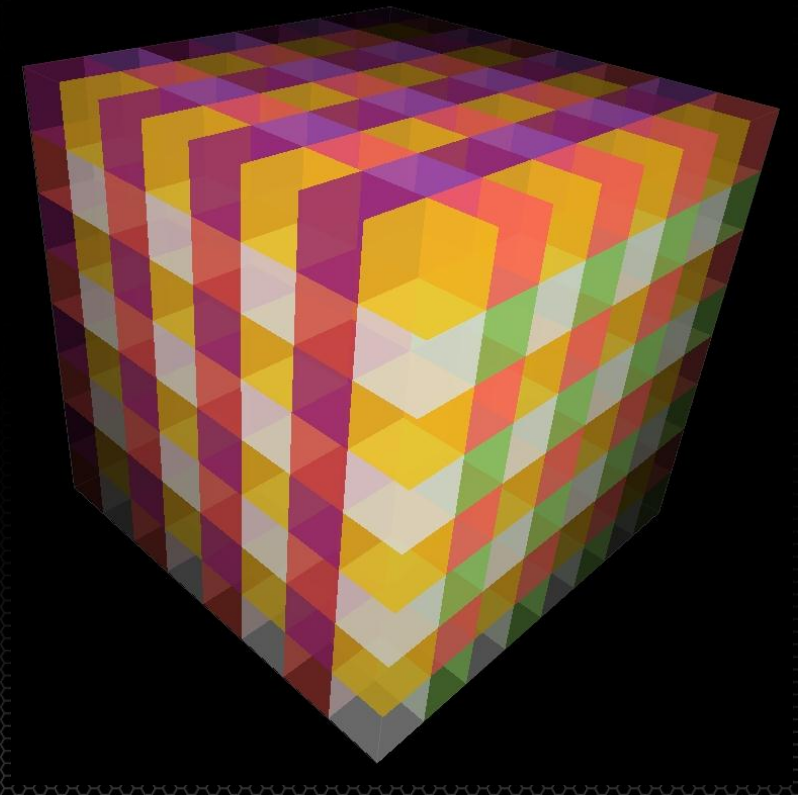

Distributed Rendering using GPU-Clusters

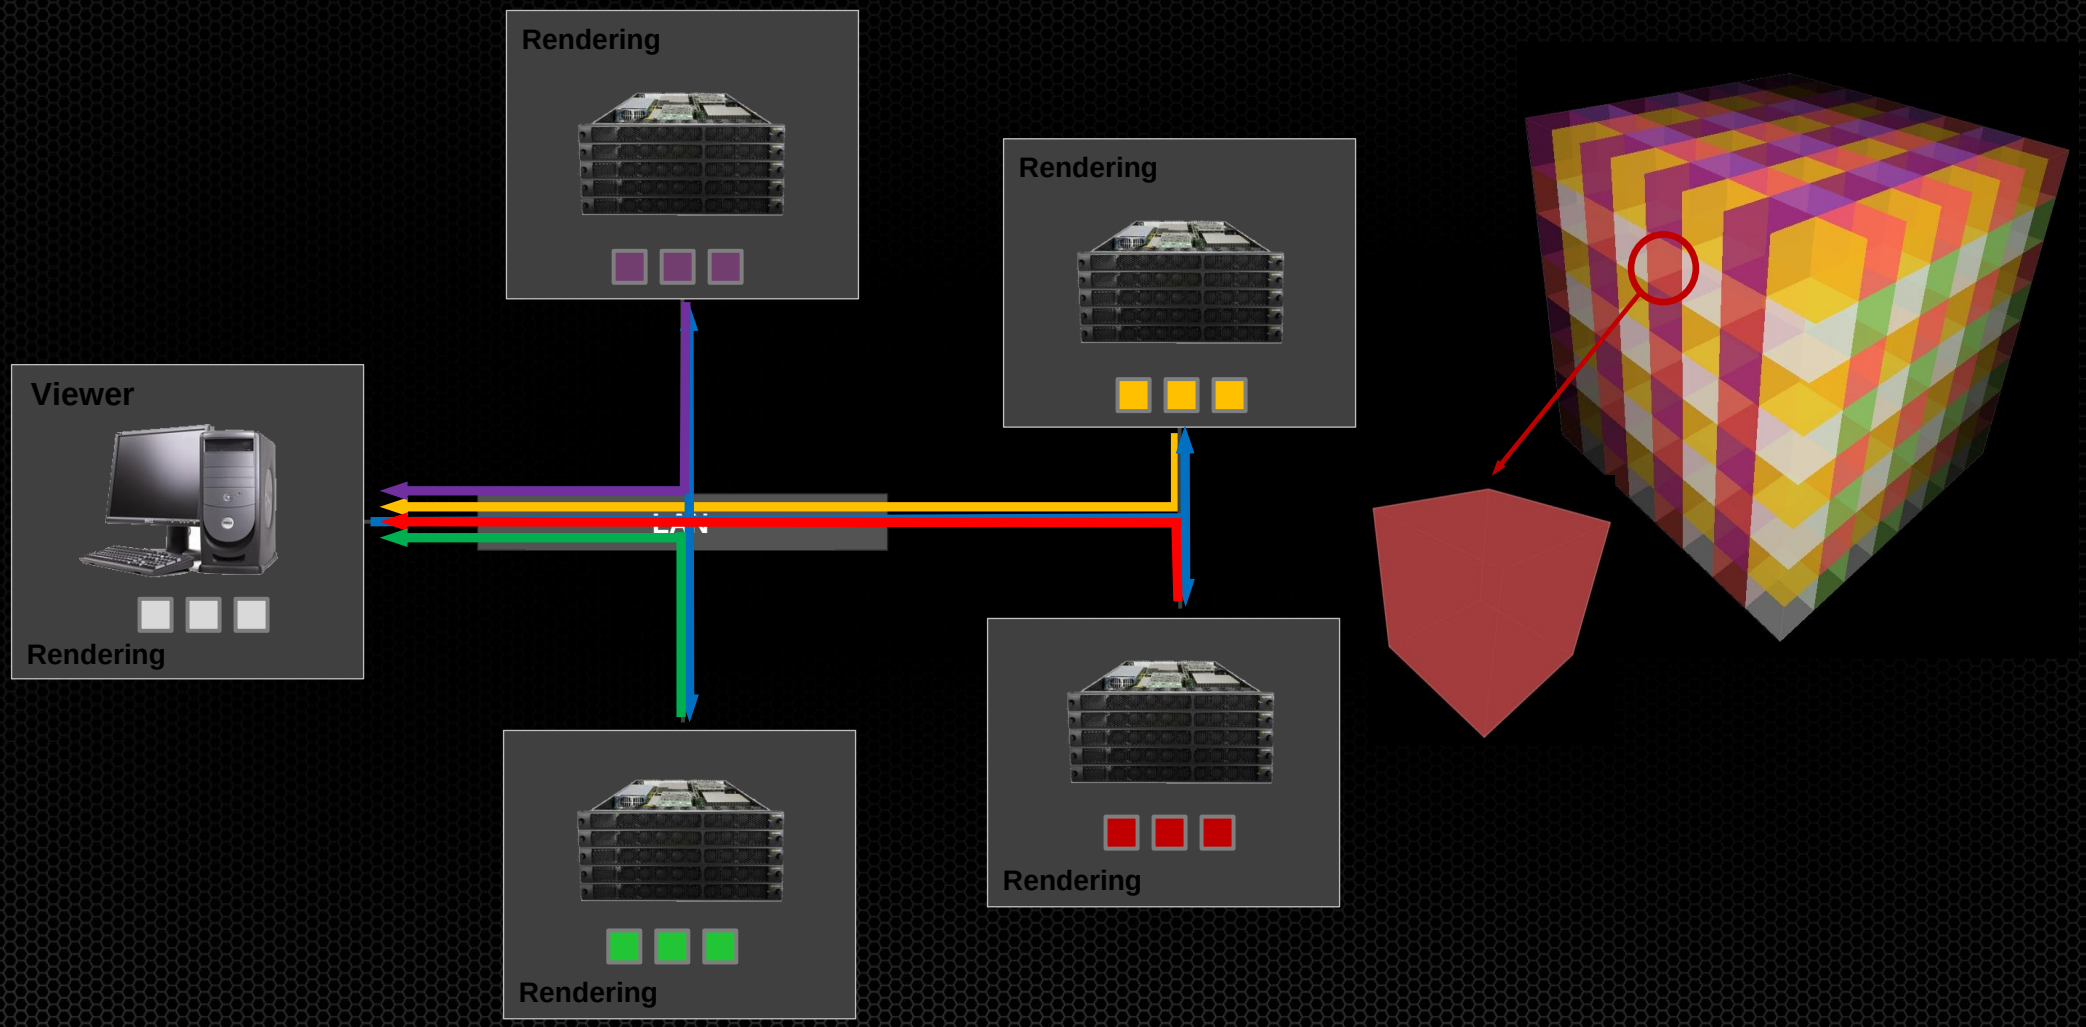

Parallel and Distributed Rendering

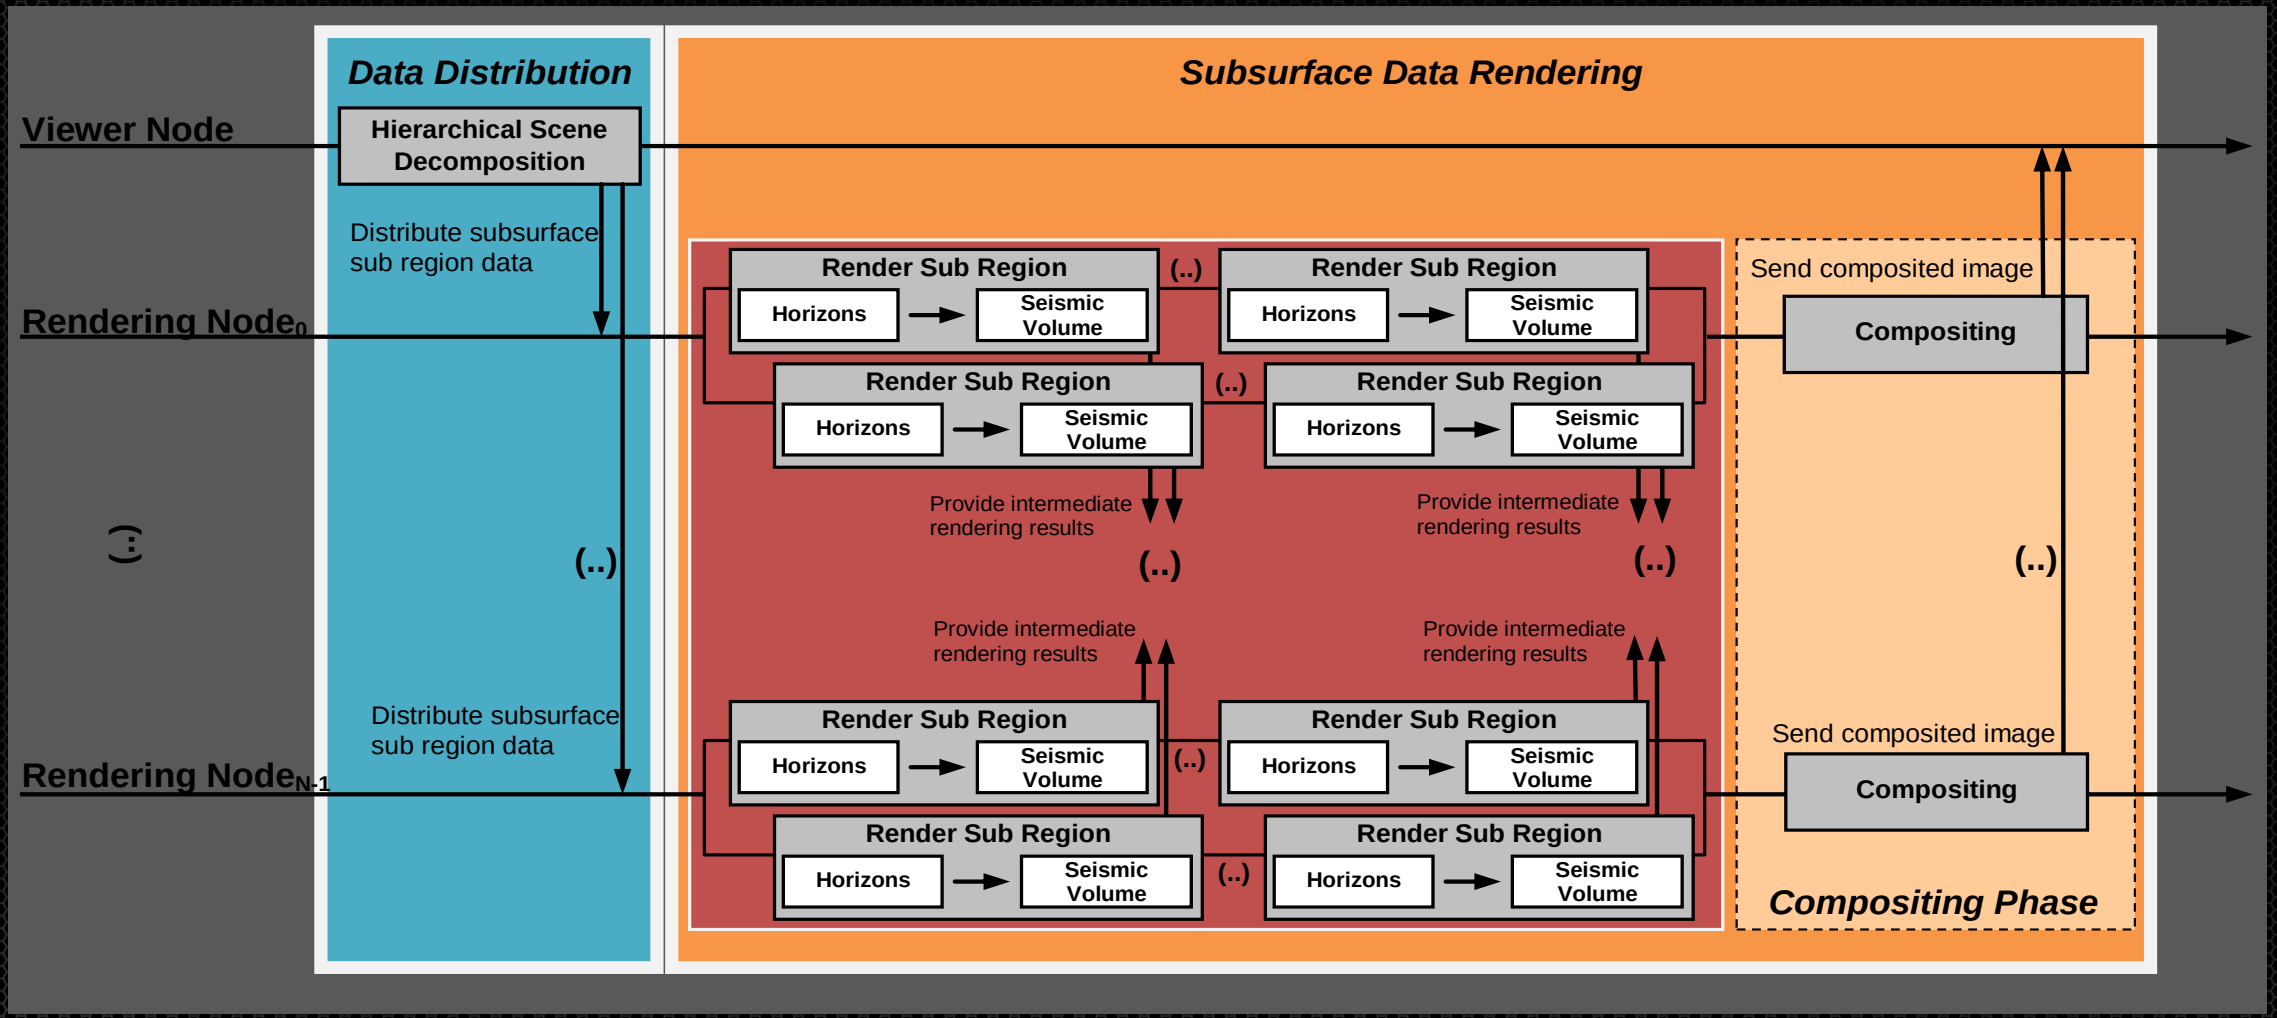

Performance and Scalability

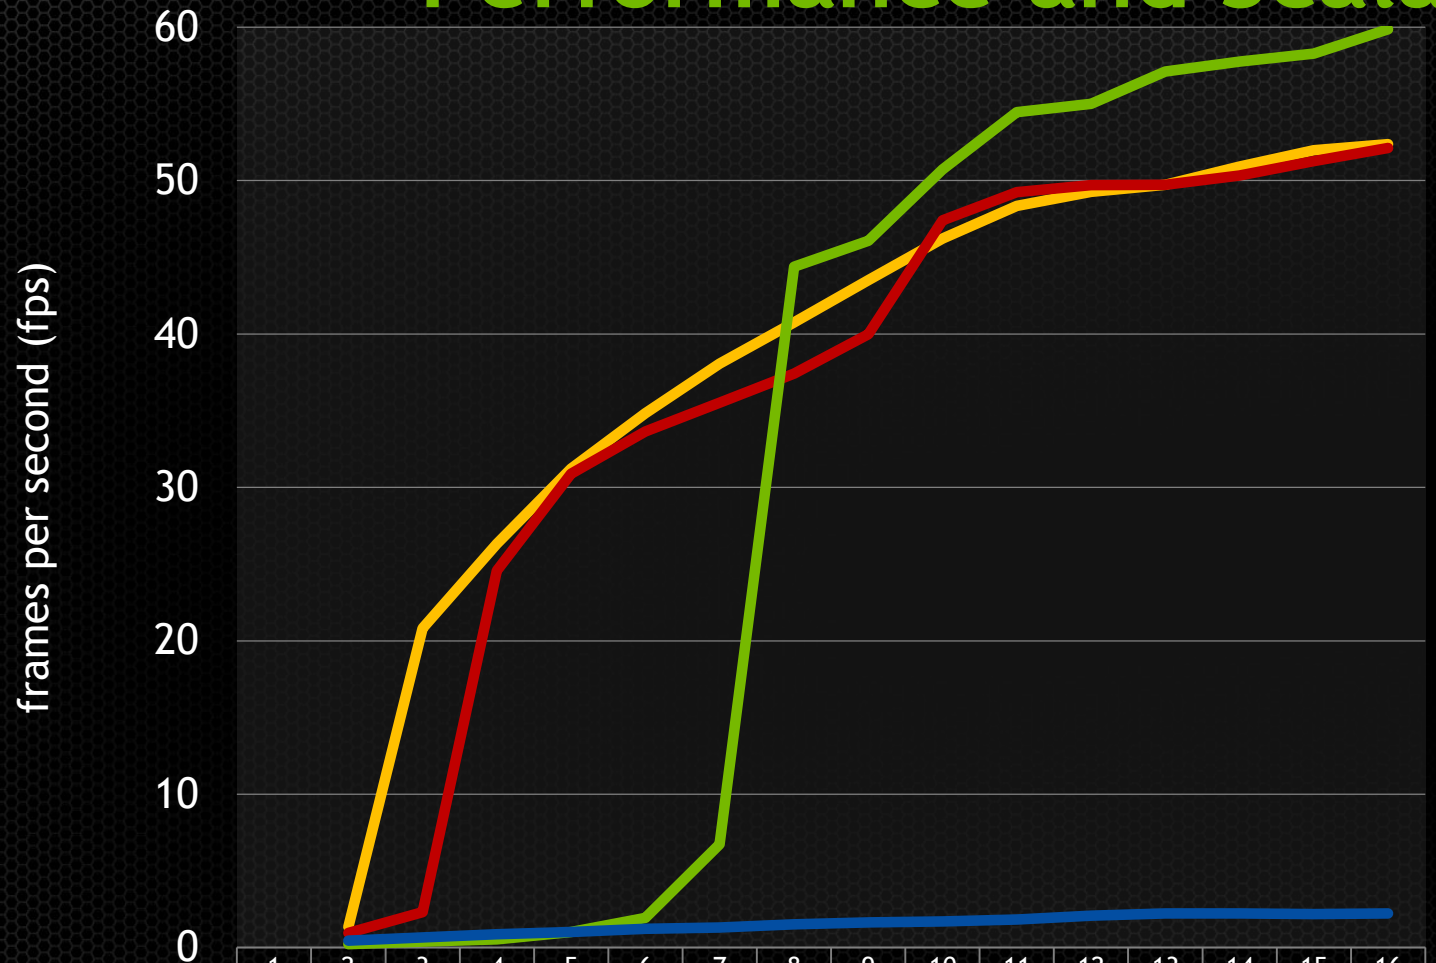

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
40 GB		1.363	20.82	26.33	31.23	34.84	38.06	40.75	43.5	46.22	48.36	49.27	49.73	50.92	51.99	52.38
80 GB		0.959	2.321	24.56	30.89	33.65	35.52	37.42	39.99	47.41	49.25	49.71	49.73	50.34	51.27	52.12
160 GB		0.23	0.393	0.543	1.024	1.934	6.737	44.39	46.08	50.72	54.47	55.01	57.11	57.76	58.29	59.87
40 GB (CPU)		0.46	0.66	0.88	1.02	1.22	1.32	1.52	1.65	1.7	1.83	2.08	2.23	2.23	2.18	2.23

- Cluster details
  - 2-16 cluster machines
  - 4 Tesla K10 per cluster machine
  - 8 GB per K10
  - 1k x 1k screen
  - 10 Gigabit Ethernet

Another Cluster Setup

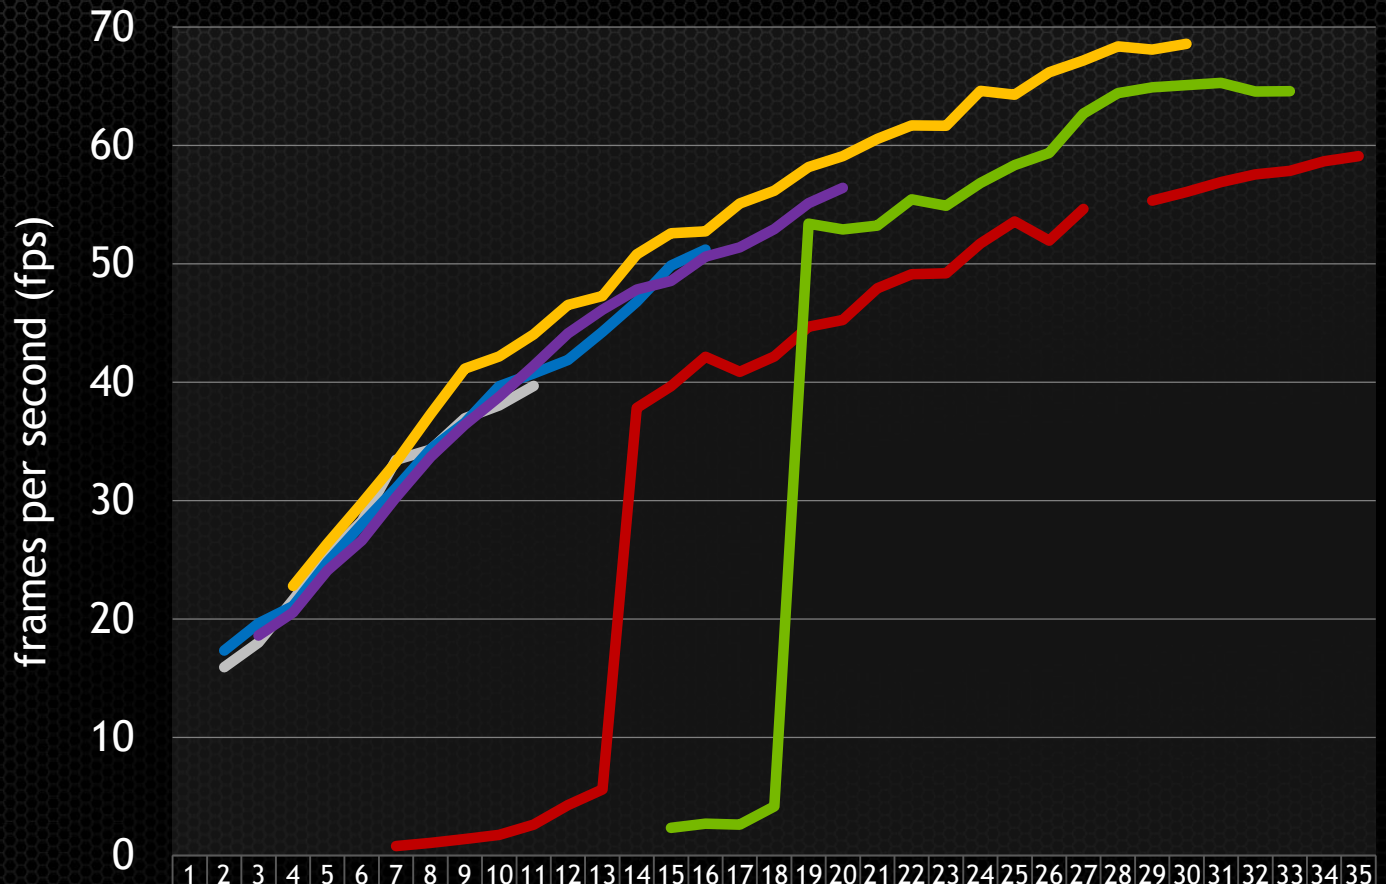

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35
5 GB		16	18	22	25	28	33	34	37	38	40																								
10 GB		17	20	21	25	28	31	34	37	40	41	42	44	47	50	51																			
20 GB			19	21	24	27	30	34	36	39	41	44	46	48	49	51	51	53	55	56															
40 GB				23	26	30	33	37	41	42	44	47	47	51	53	53	55	56	58	59	61	62	62	65	64	66	67	68	68	69					
80 GB							1	1	1	2	3	4	6	38	40	42	41	42	45	45	48	49	49	52	54	52	55		55	56	57	58	58	59	59
160 GB															2	3	3	4	53	53	53	55	55	57	58	59	63	64	65	65	65	65	65		

- Cluster details
  - 2-35 cluster machines
  - 2 Tesla M2090 (Fermi) per cluster machine
  - 1k x 1k screen
  - 10 Gigabit Ethernet

Dataset Scalability

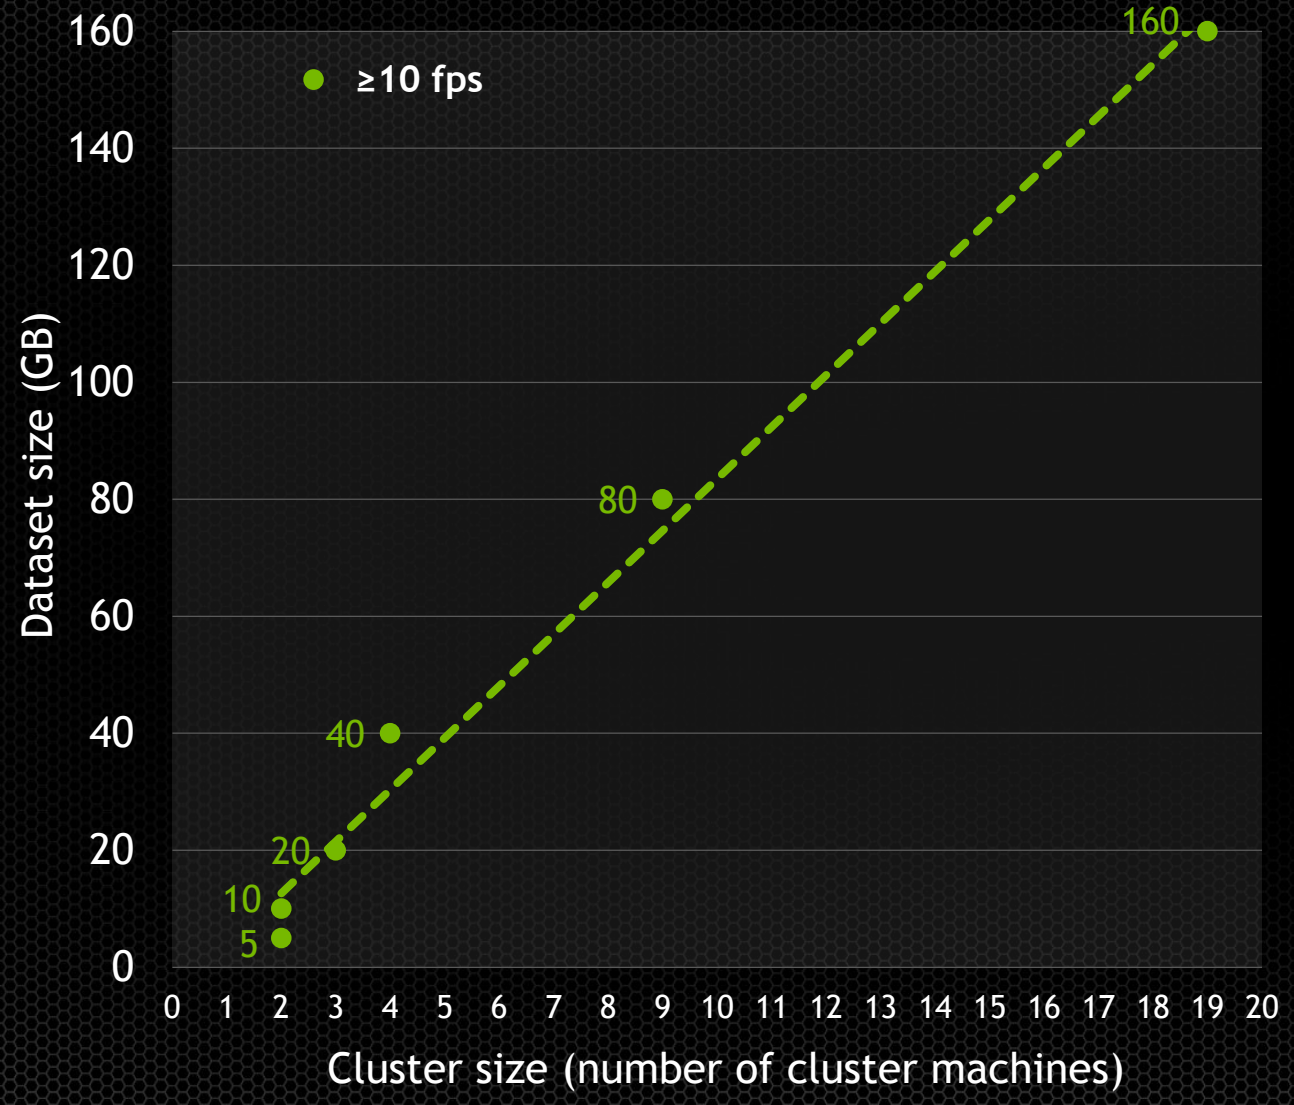

- Target performance
  - ≥10 fps @ 1024x1024
- Volume dataset sizes
- Cluster details
  - 2-35 cluster machines
  - 2 Tesla M2090 (Fermi) per cluster machine
  - 1k x 1k screen
  - 10 Gigabit Ethernet

Scalability at Various Screen Resolutions

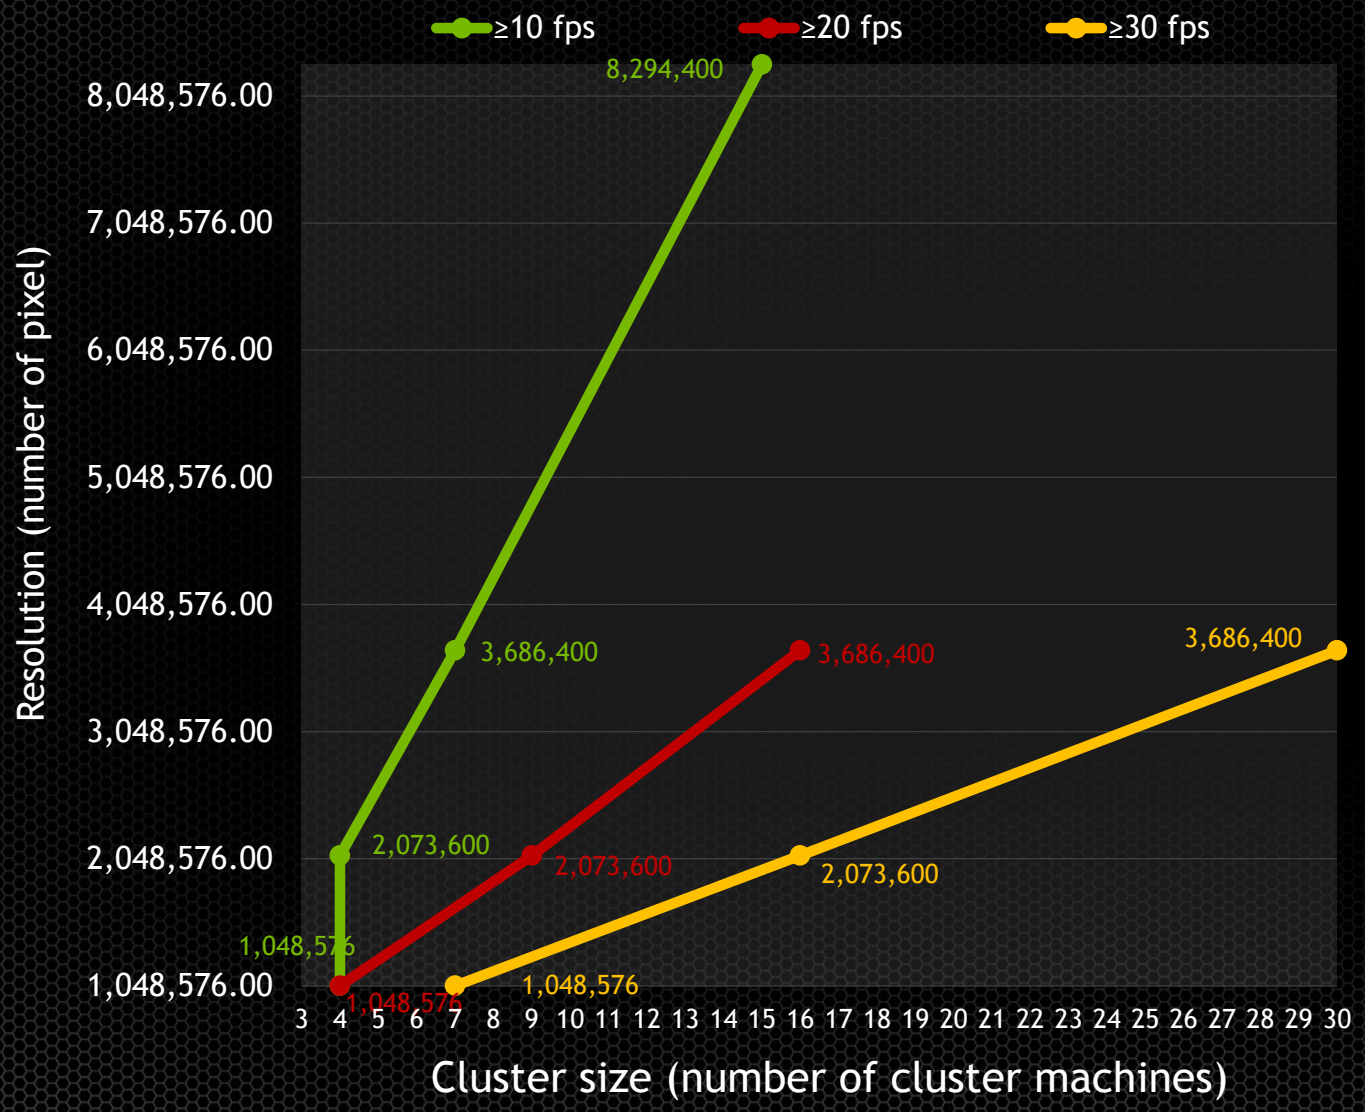

- Target performance
  - ≥10 fps, ≥20 fps, ≥30 fps
- 40 GB volume dataset
- Screen resolutions
  - 1024x1024 (Baseline)
    (1,048,576 pixels)
  - 1920x1080 (Full HD)
    (2,073,600 pixels, 1.98 x Baseline)
  - 2560x1440 (WQHD)
    (3,686,400 pixels, 3.5 x Baseline)
  - 3840x2160 (QFHD)
    (8,294,400 pixels, 7.9 x Baseline)

Interactively on Workstations, Clusters, and Clouds

Workstation

NVIDIA GRID Visual Computing Appliance (VCA)

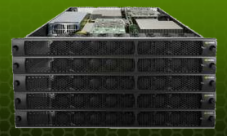

(..)

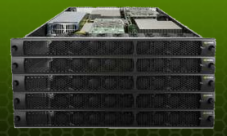

GPU Clusters

- Single or multiple GPU(s)
- Smaller dataset sizes
- Interactive rendering performance

(..)

- Unlimited number of GPUs
- Huge dataset sizes
- Increasing rendering performance

Interactive GPU-Cluster aware Visual Computing

- Interactive attribute generation for instantaneous visualization
- Applications
  - Flow simulations
  - Atmospheric dynamics visualization
  - Combustion simulation
  - Molecular dynamics simulations
  - Seismic attribute generation for survey visualization

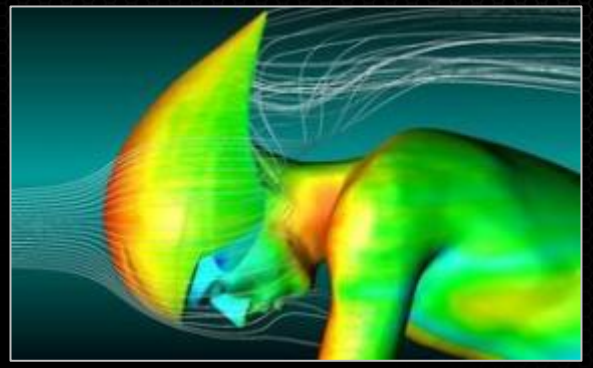

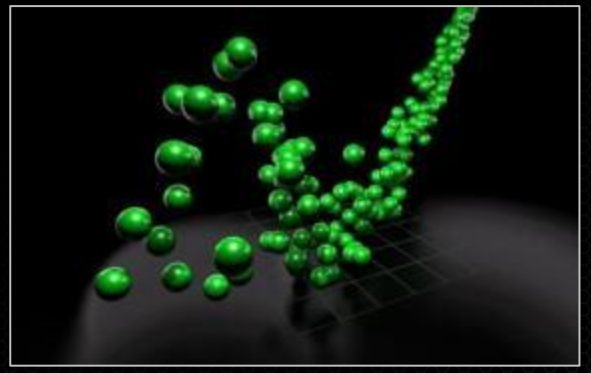

Architectural Challenges for Interactive GPU-Cluster aware Visual Computing

- Raw n-dimensional data is huge
  - Multiple times larger than generated attributes
- Process raw data using user-defined algorithms
  - Plethora of possible types of attribute
  - Manifold parallel compute algorithms
  - Diversity of algorithm-specific subdivision schemes
- Interactive attribute generation for instantaneous visualization
  - Scalability

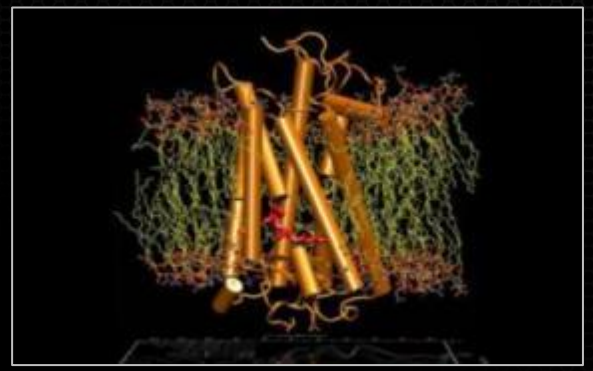

Mapping Attribute onto Scene Geometries

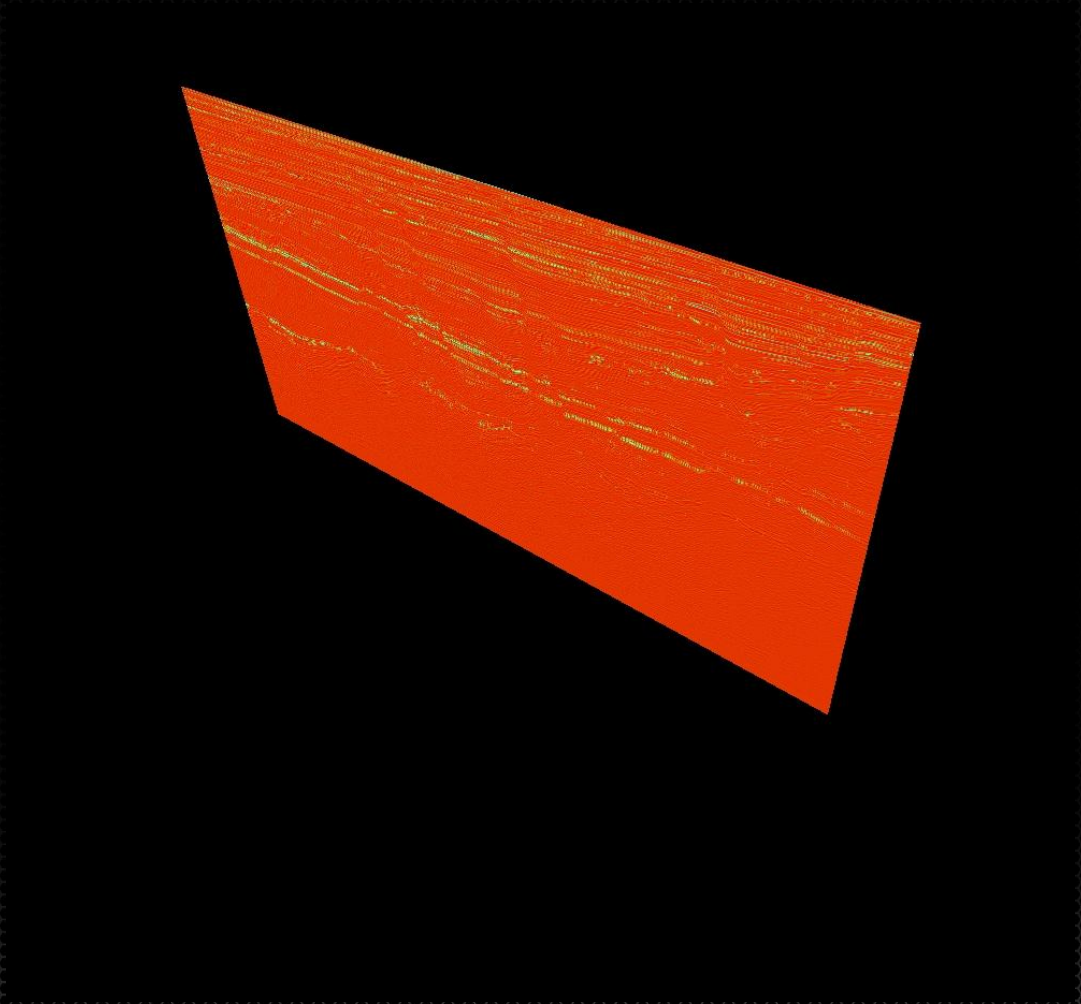

Mapping Attribute onto Scene Geometries

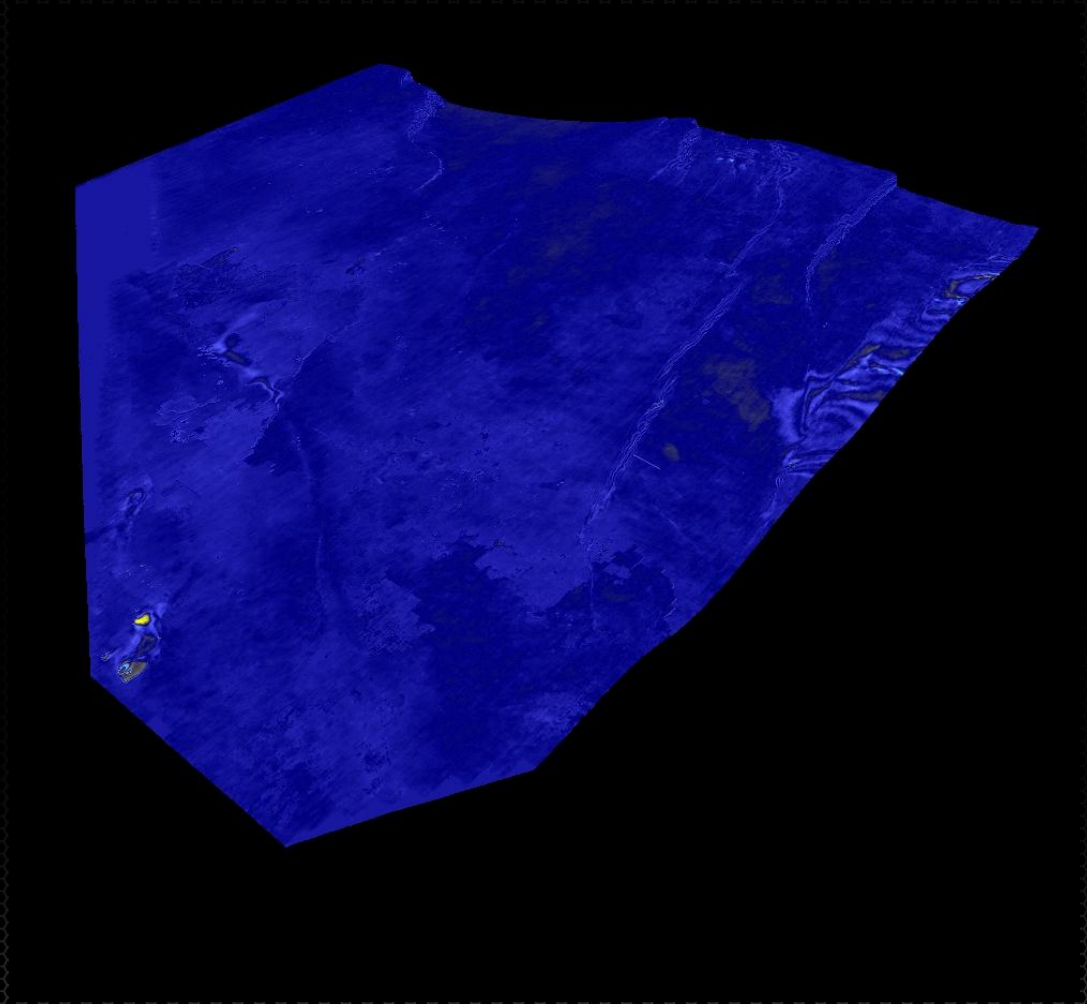

Proxy Shapes for Attribute Visualization

- Proxy shapes
  - Slices
  - Height fields
  - Triangle meshes
  - Volumes
- Part of the scene description
- Canvas for attribute visualization

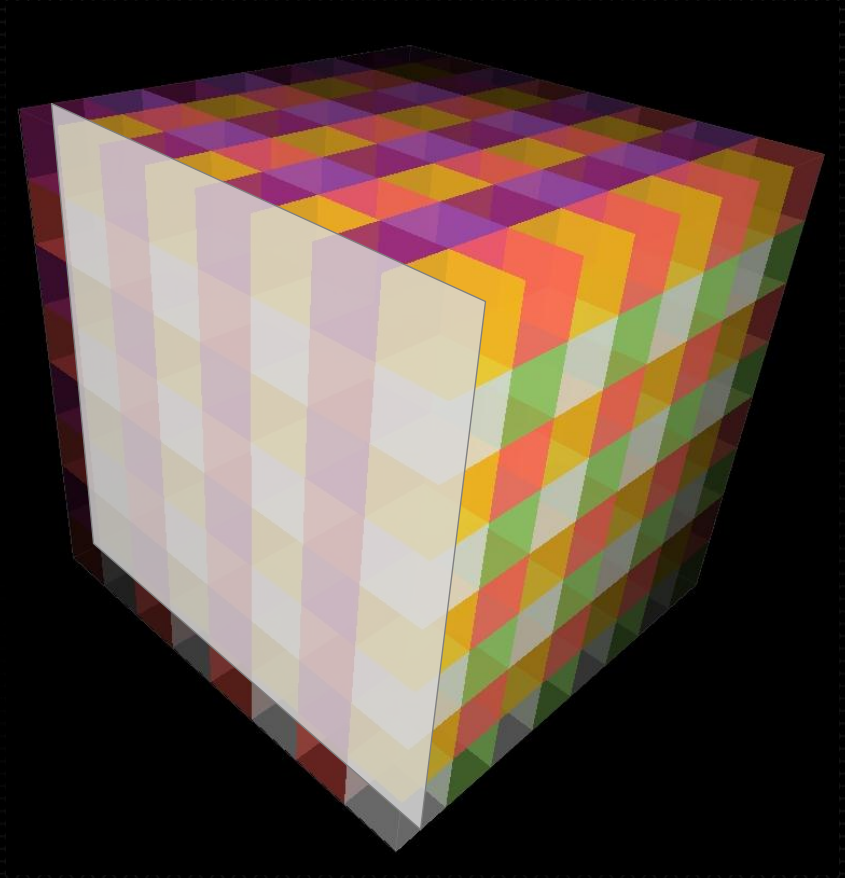

Distributed Attribute Visualization Process

- User-defined attribute computation
  - Compute jobs launched per portion
  - Proxy shape intersection
  - Algorithm-specific subdivision schemes
- Attribute mapping
  - Rendering proxy shapes
  - Analogy: procedural texturing

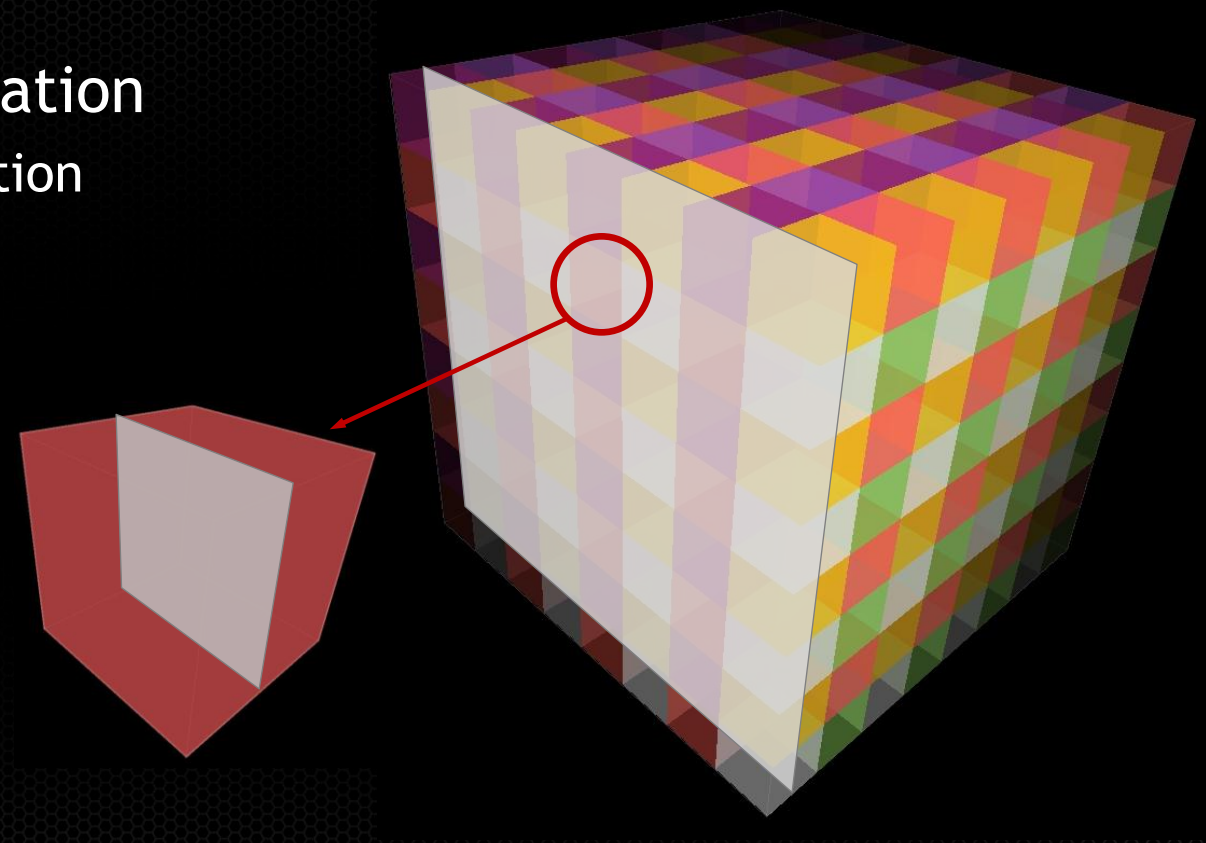

Attribute Generation and Visualization Process

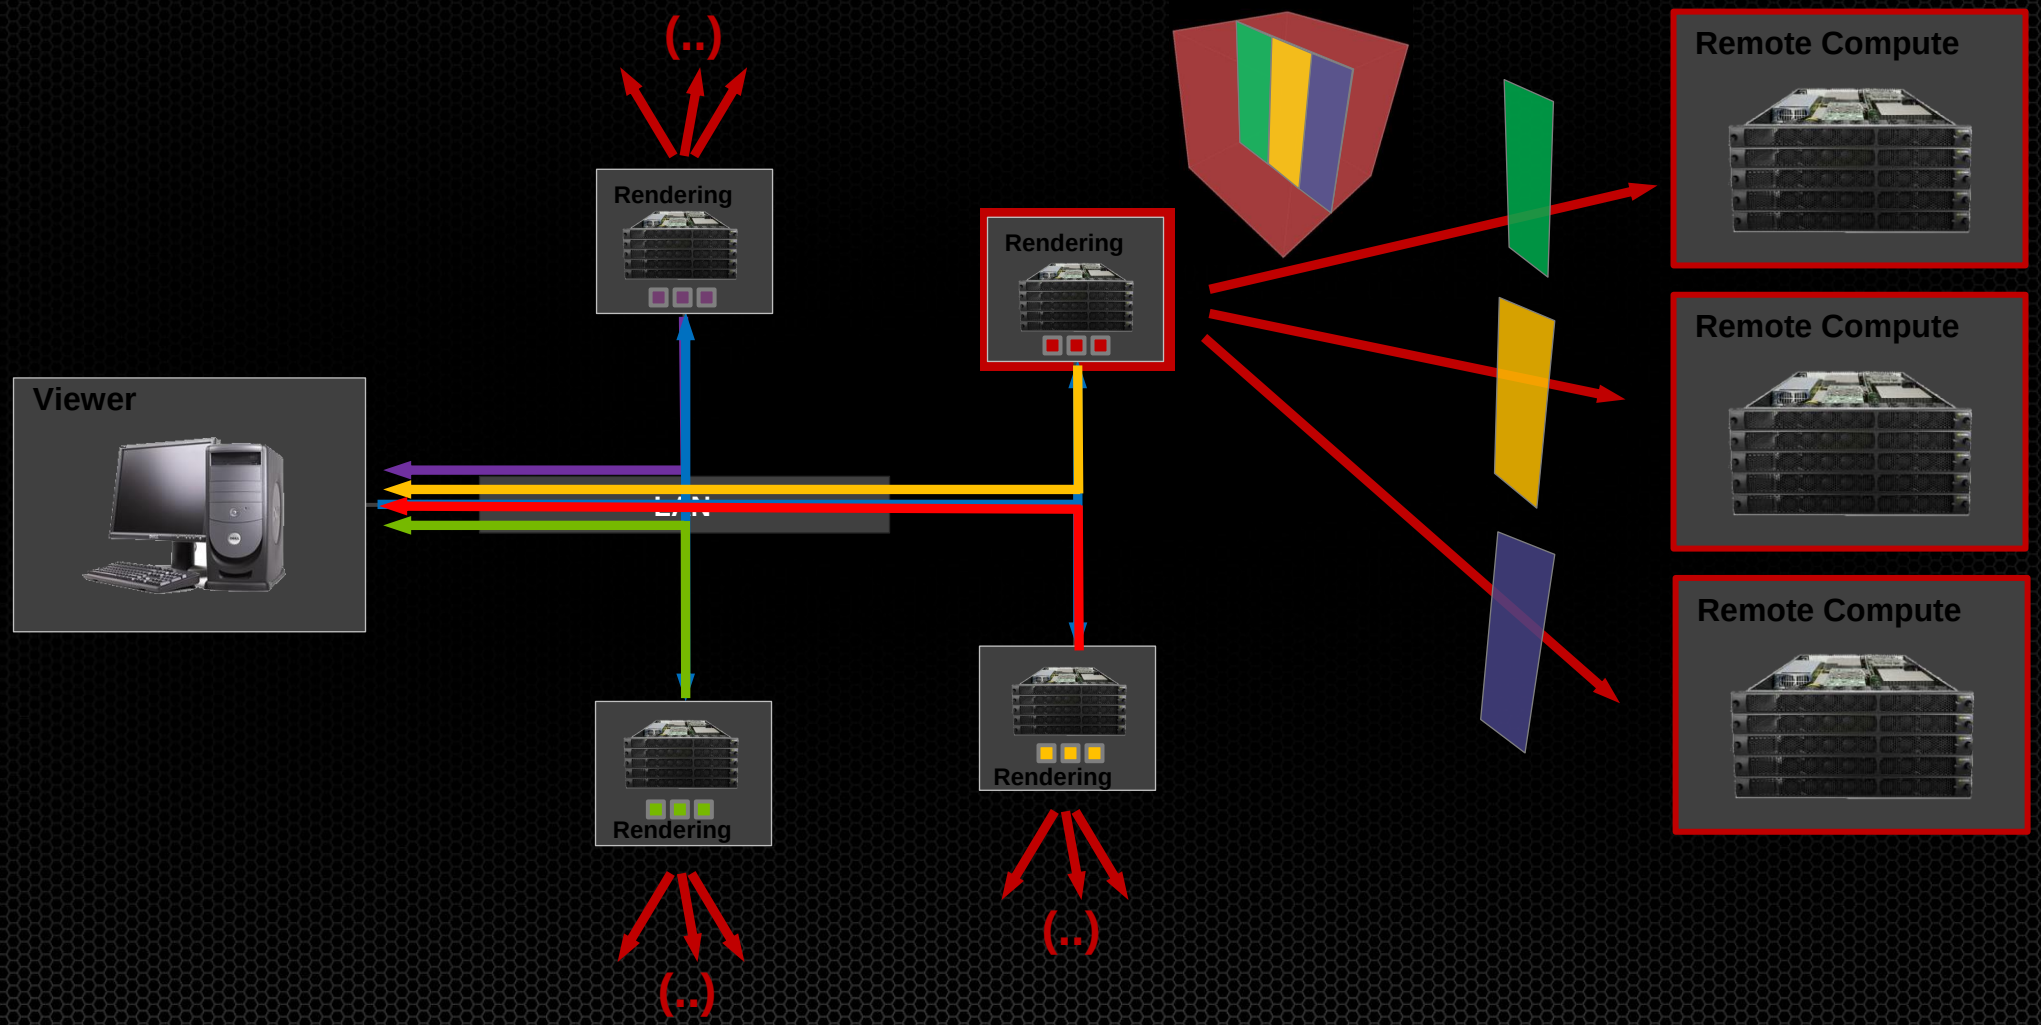

GPU Cluster Setup for Scalable Visual Computing

- Asynchronous compute maximizes performance
  - Rendering and compute process run in parallel
  - Compute integration into rendering

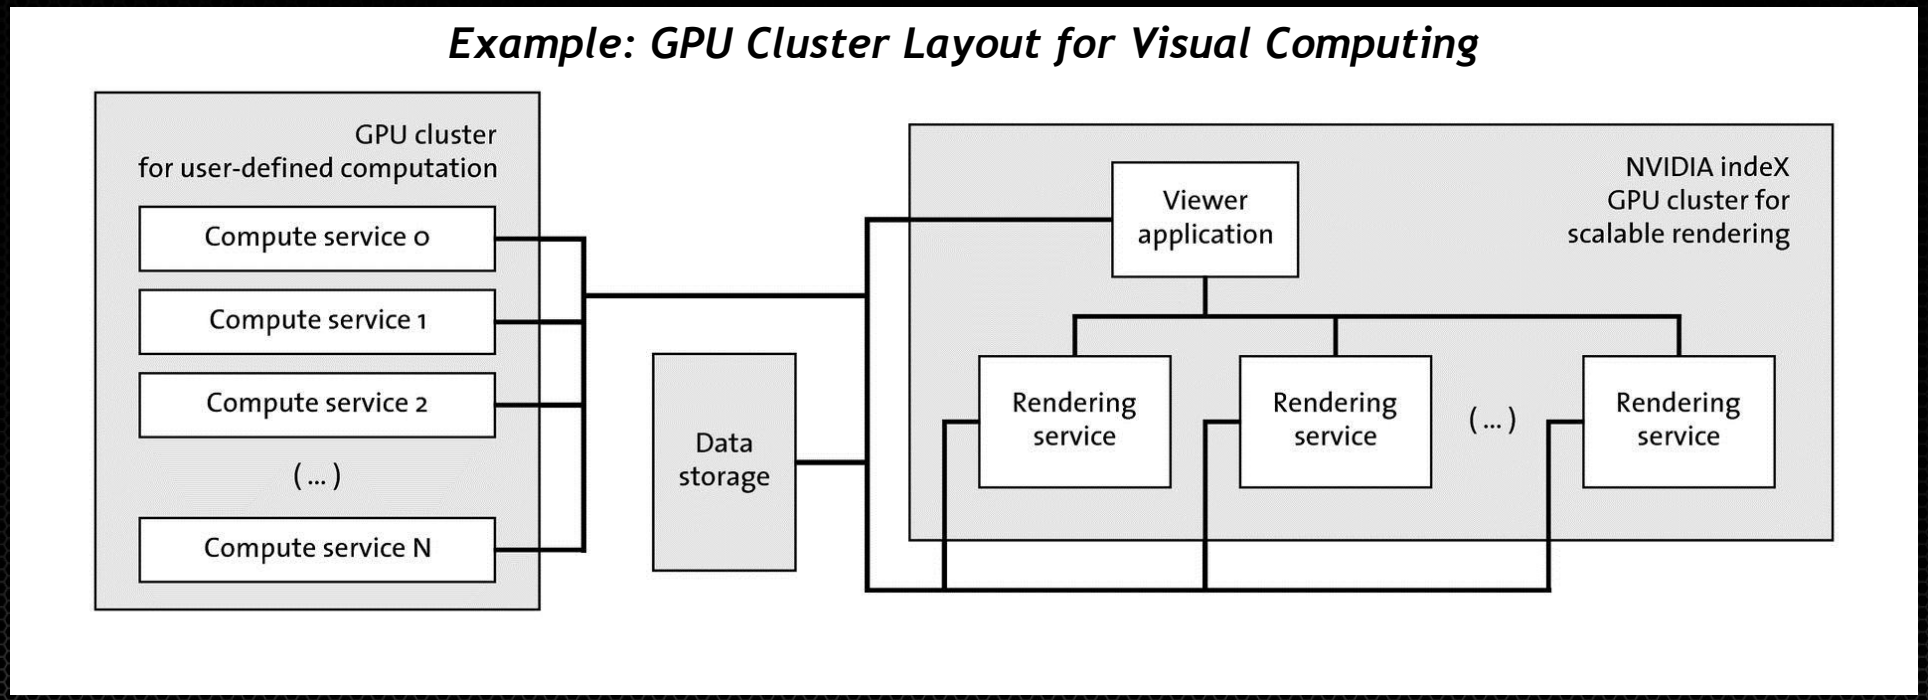

Extensible Software Architecture

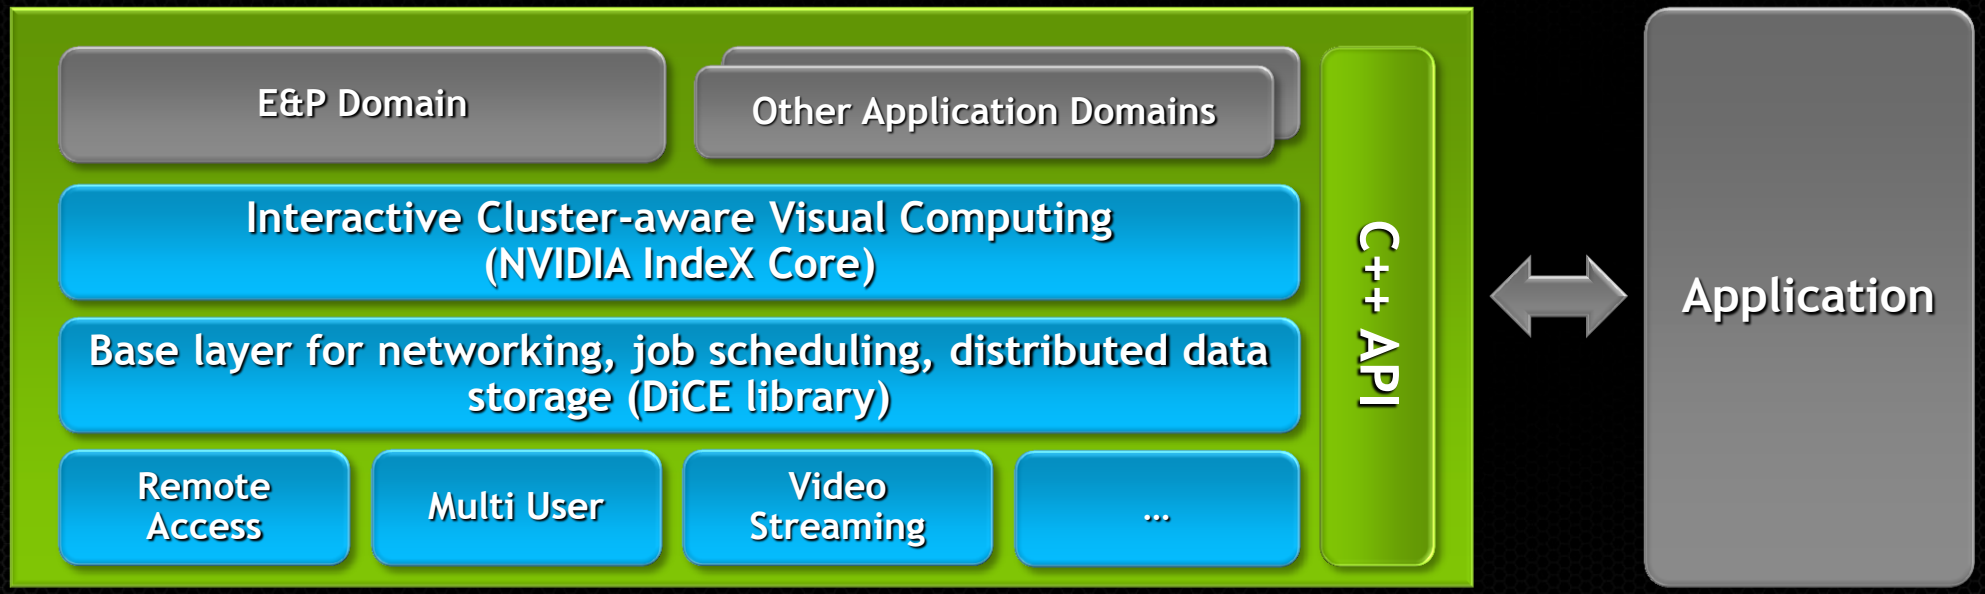

NVIDIA IndeX's Building Blocks for Managing Distributed Data

- Data locality information
  - *Spatial query* tells which cluster machine stores which portions of data
  - Depends on dataset type
- Accessing large-scale data
  - Assemble from cluster machines
  - Can be restricted to portions
- Editing large-scale data
  - Direct editing/compute to the distributed data
  - Can be restricted to portions
  - Simple example: user-defined filter

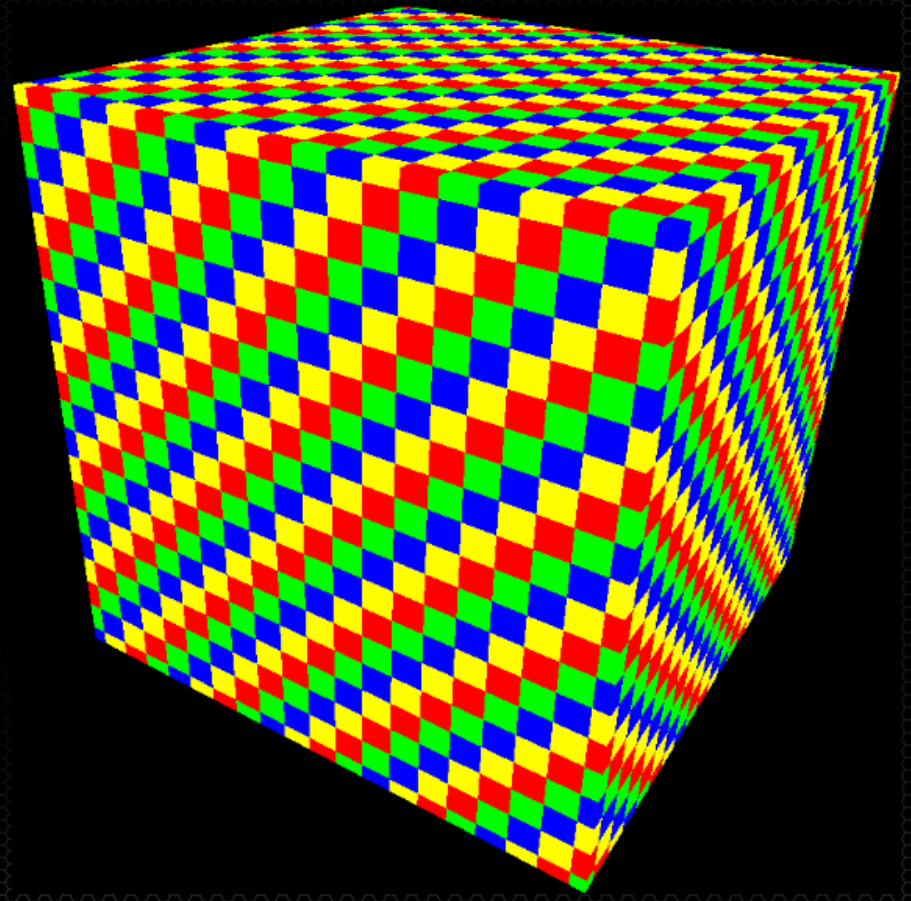

OpenGL Integration

1. Rasterize opaque geometry
   - Capture depth and color buffers
2. Volume ray casting
   - OpenGL depth buffer
3. Alpha-based color blending
   - Result of NVIDIA IndeX's large-scale data rendering with alpha
   - OpenGL color buffer

Pseudo code

```
// OpenGL rendering to color- and
// depth-buffer on local machine
(gl_buffer_RGBA, gl_buffer_z) ← renderGL();

// Distributed, scalable rendering
// with depth-buffer input
result_RGBA ← nvidia_index->render(buffer_z);

// Blending rendering results
blend(gl_buffer_RGBA, result_RGBA);
```

Proof-of-Concept Example

Opaque OpenGL geometry integrated in NVIDIA IndeX

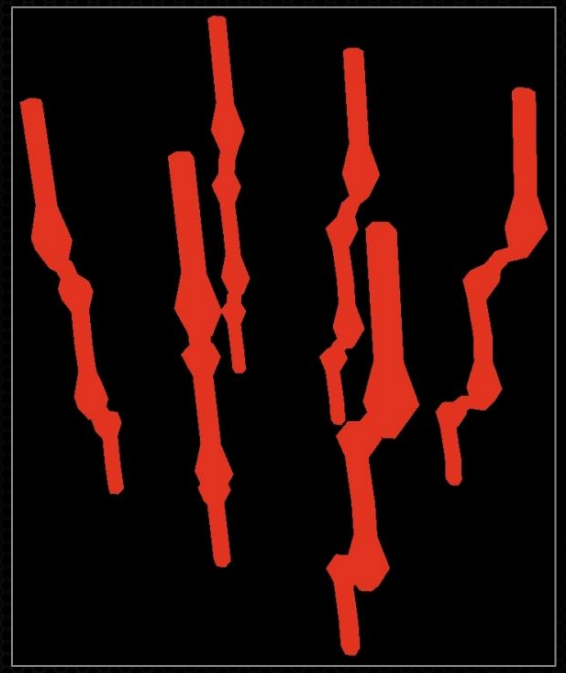

OpenGL color buffer

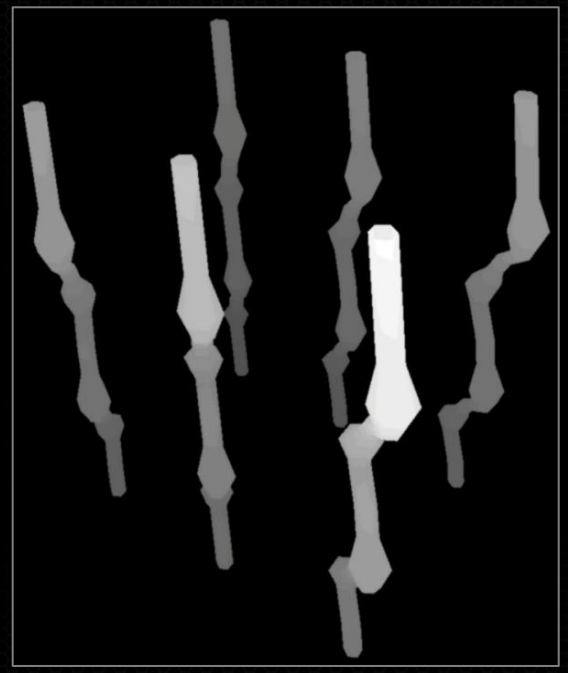

OpenGL depth buffer (inverted)

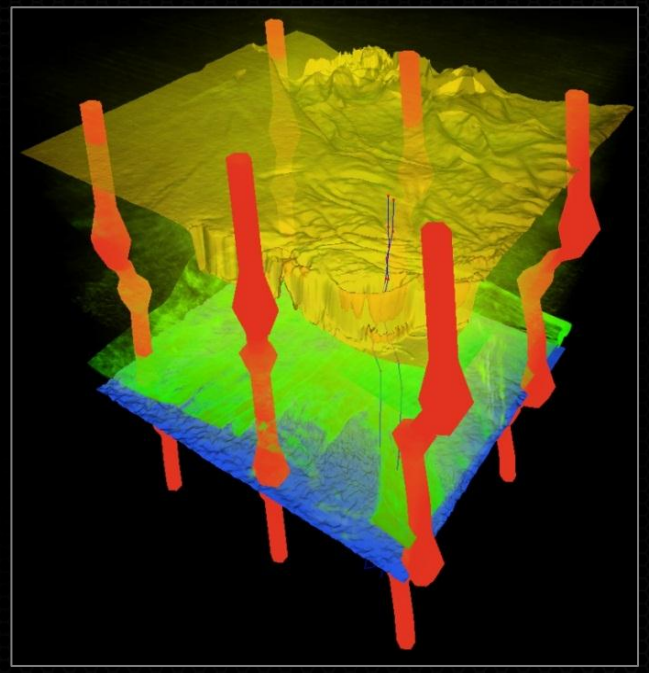

Final combined rendering

Depth-correct OpenGL Integration

- Z-buffer compressions
- Multicasting z-buffer data

- Possible extensions
  - Transparent OpenGL
  - InfiniBand and RDMA
  - GPUdirect support

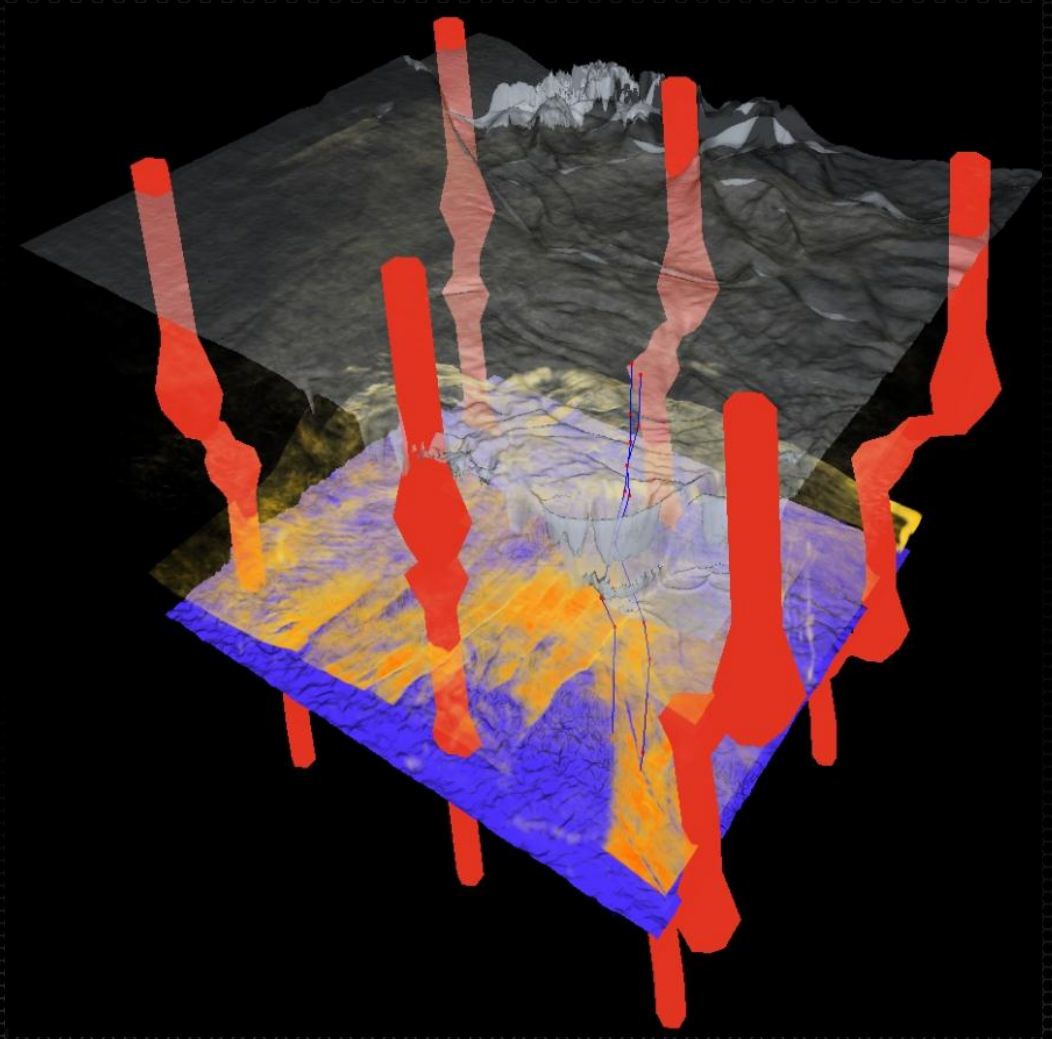

Remote Visualization

- Video streaming
  - H.264 video encoding
  - Hardware-accelerated on Kepler
- Private and public clouds
  - Web-based applications
  - Thin clients (tablets)
- Multi-user support for world-wide collaboration

Live Demo

- GPU cluster located in Berlin, Germany
  - 8 cluster nodes
  - 2 Tesla M2090 each
  - 82 GB volume
  - Height field with 250 million triangles
  - H.264-based video streaming

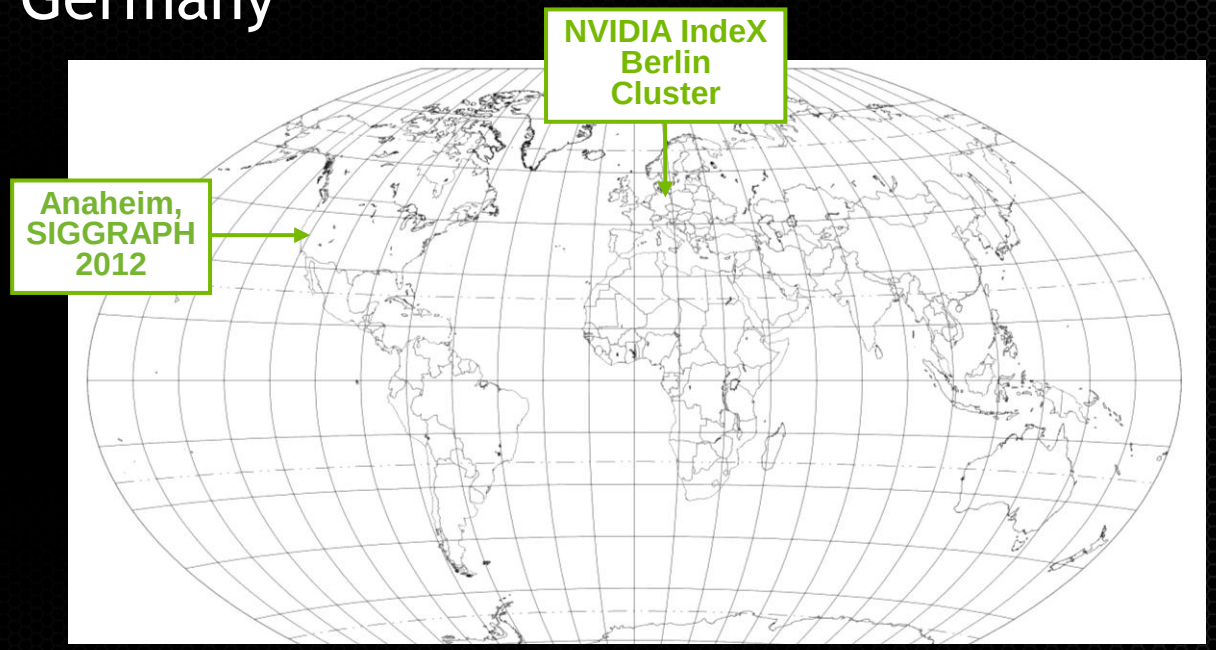

http://www.nvidia-arc.com/products/nvidia-index.html

Thank you …

Marc Nienhaus
NVIDIA IndeX Engineering Manager and Chief Architect

Christopher Lux
Sr. Graphics Software Engineer, NVIDIA IndeX

Jörg Mensmann
Sr. Graphics Software Engineer, NVIDIA IndeX

Eduardo Olivares
Sr. Graphics Software Engineer, NVIDIA IndeX

Mahendra Gopaludu Roopa
Graphics Software Engineer

Gunter Sprenger
Sr. Graphics Software Engineer

Hitoshi Yamauchi
Sr. Graphics Software Engineer, NVIDIA IndeX

Runa Löber
Director Software Engineering

Tom-Michael Thamm
Director Software Product Management
Product Manager NVIDIA IndeX

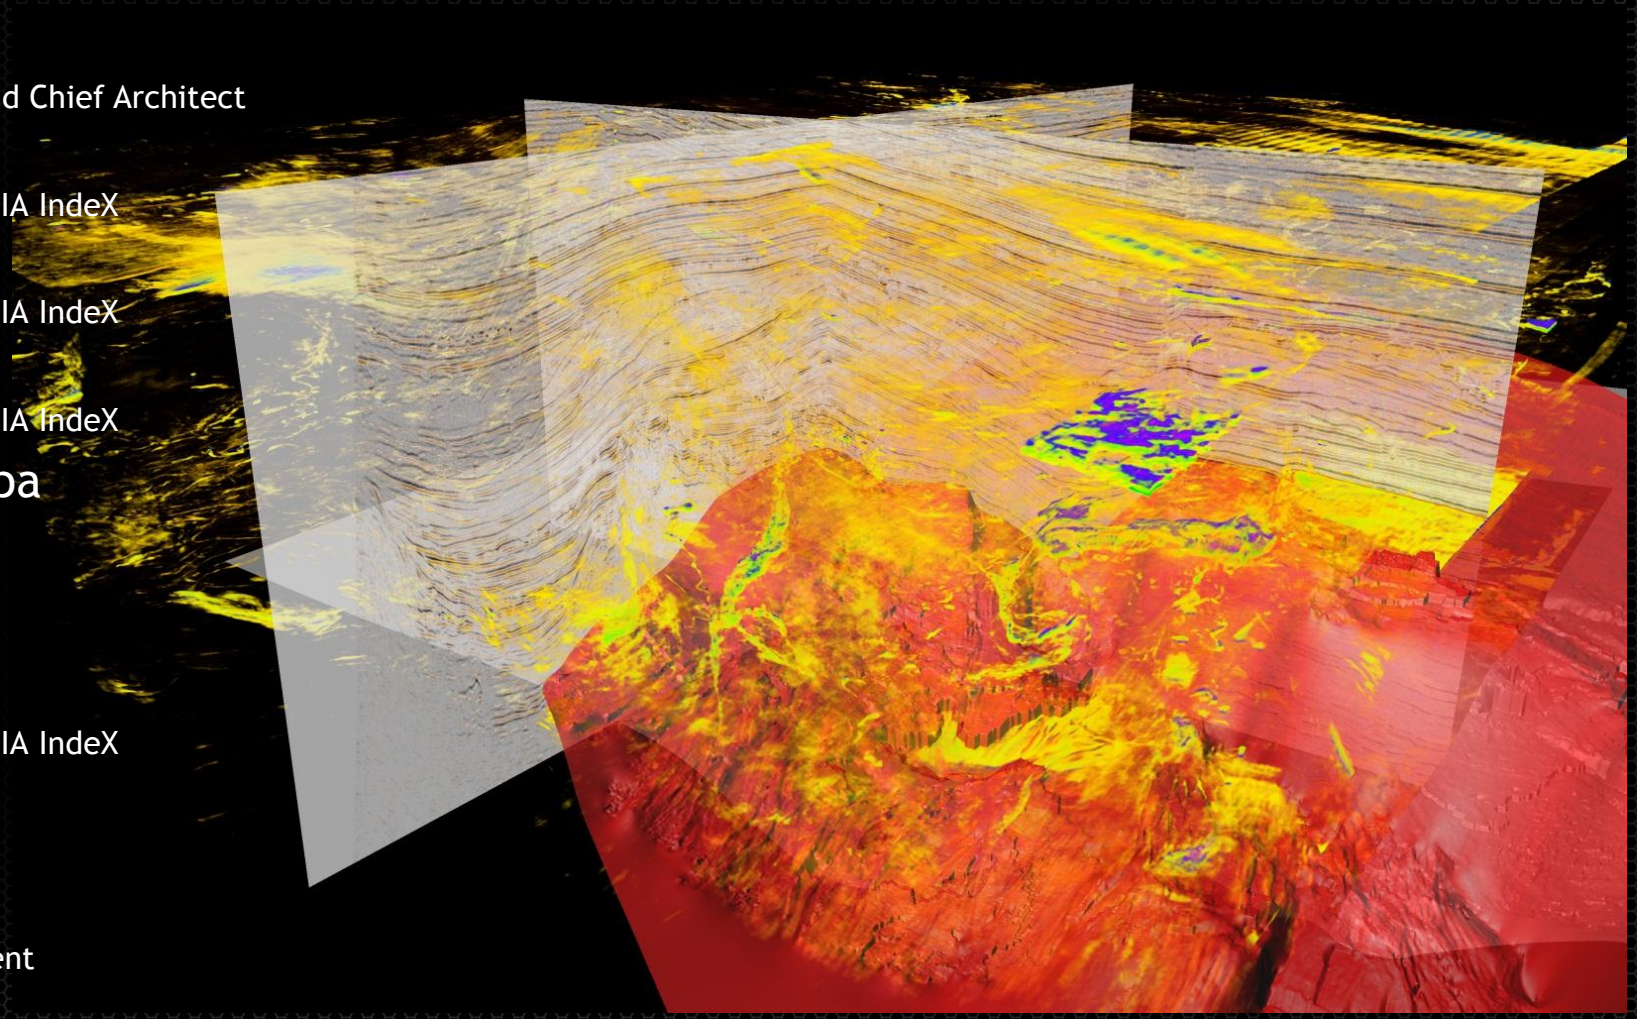

http://www.nvidia-arc.com/products/nvidia-index.html