The Deep View Project: Scalable, Network-centric Visualization with the Commodity Cluster

Thomas Jackman (Mgr), Peter Kirchner, Gregory Abram, James Klosowski, Christopher Morris
Visualization Systems Group
IBM Research
Hawthorne, NY

Contact: tjackman@us.ibm.com

Outline

- The Deep View Rendering Cluster
  - Anatomy of the cluster
  - The rendering nodes
  - Chromium sw for network indirection
  - Pixel readback
  - Network technology
- The Scalable Graphics Engine (SGE)
- The Scalable Software Concentrator
- Remote Network Visualization

DEEP VIEW

High Performance Cluster-Based Rendering and Media Server For Networked Visualization

… a scalable, commodity-based, cluster visualization system capable of remote delivery of image content over networks to both offices and/or collaborative displays.

Key System Components

- IBM IntelliStations w/ AGP 8X graphics for rendering/processing
- GbE networks for communication
- Chromium Open Source SW for automatic parallel rendering
- Linux

In 2000, the Visualization Systems Group of IBM Research began a project called Deep View to explore the capabilities of cluster based rendering for real-time, interactive visualization.

From the outset, our focus has been on the network for on-cluster and off-cluster communication.

This network-centric design allowed interesting variations for Deep View and insulated our software development from frequent upgrades caused by expected changes in the associated hardware.

User Benefits from the Approach

- Lower cost
- Longer lifetime for system
- Scalability
- Compatibility with legacy applications
- Versatile design supports both local and remote visualization
- Linux and Open standards/protocols

Inside the Deep View Cluster

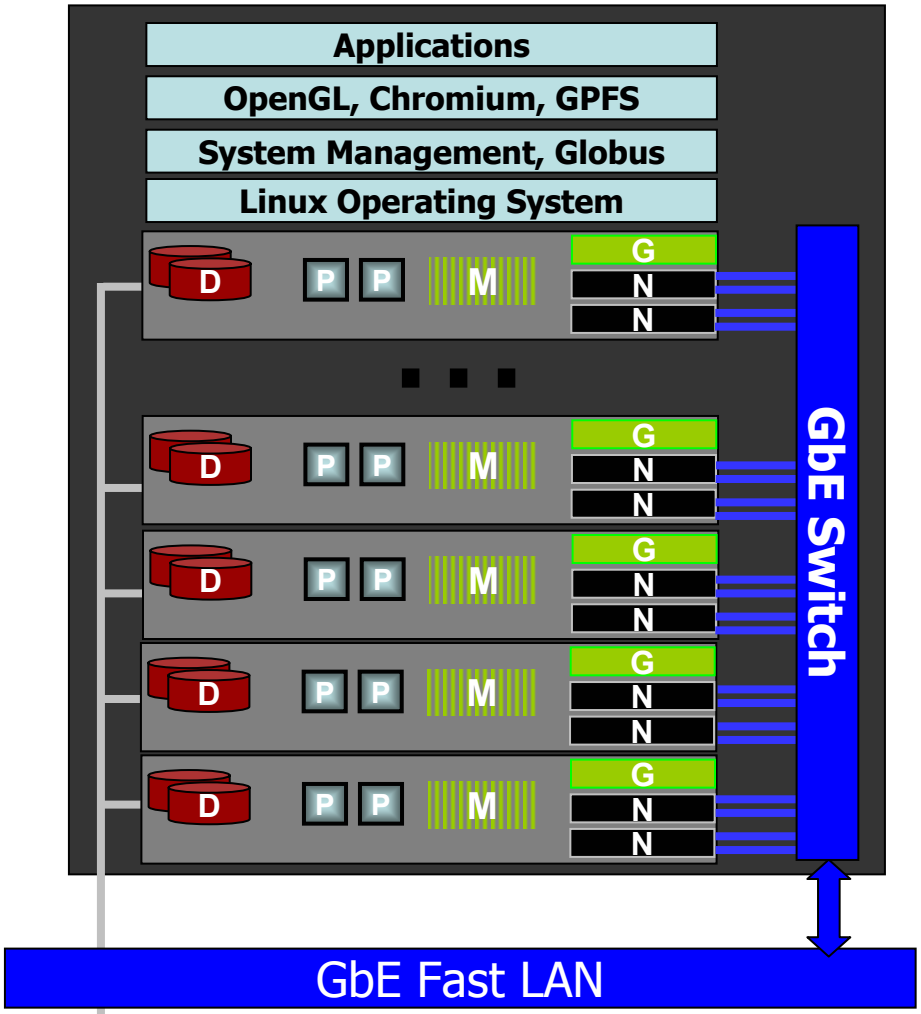

- IntelliStation Z Pro graphics workstations are rackmounted to serve as rendering nodes
- IntelliStations provide AGP graphics for hardware accelerated rendering
- Gigabit ethernet is used as cluster interconnect to transport geometry, synchronization, events, and pixels
- Visualization cluster can be integrated into compute cluster management
- Cluster can be integrated into a Grid
- Nodes are available for processing when not utilized for rendering

The Deep View Rendering Nodes
IBM IntelliStation Z Pros

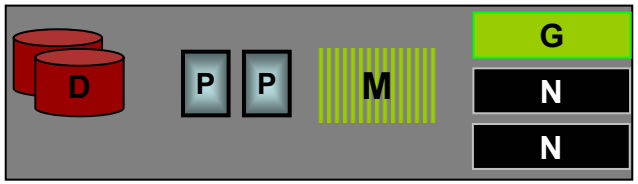

P — 2x 3.06GHz Pentium IV Xeon processors w/ 512kB L2 cache and 533MHz FSB

D — Up to 4x 145GB Ultra 320 SCSI disk drives

M — Up to 4 GB ECC DDR SDRAM

N — 2x PCI-X buses support 4x 64bit,100MHz PCI slots that can be used w/ single or dual ported Intel 1000Base-T network adapters

G — Latest AGP 8X graphics with up to 256MB of texture memory

Common SW component: Chromium

An Open Source middleware that automatically parallelizes an existing visualization application for cluster rendering.

- Any OpenGL application can be parallelized for cluster rendering simply by linking to a surrogate OpenGL library.
- Chromium is an Open Source project principally that originates from WireGL from Stanford University.
- Chromium is available on Linux, Windows and most flavors of UNIX

Chromium SW Infrastructure for Managing Parallel Rendering

Sort-first rendering (screen decomposition)

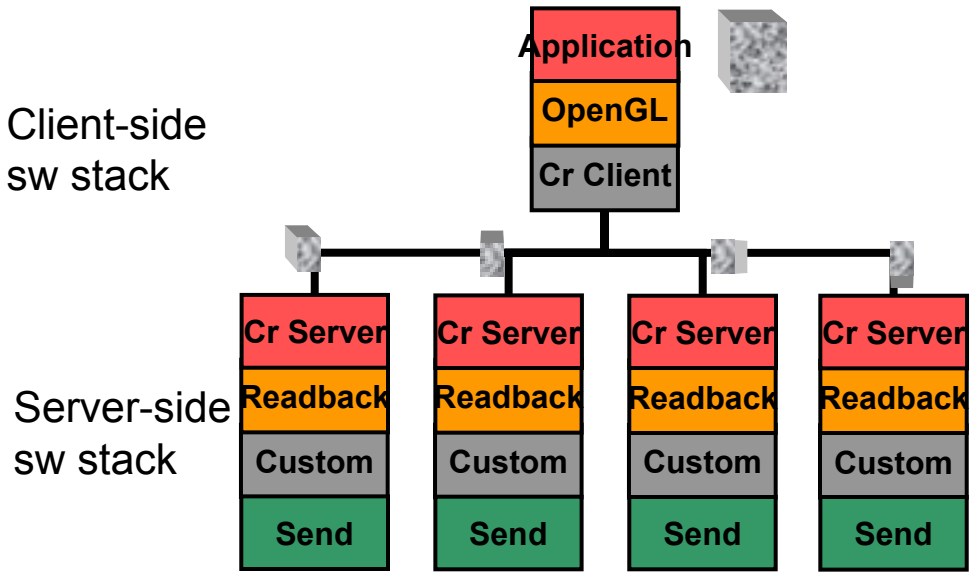

Sort-last rendering (domain decomposition)

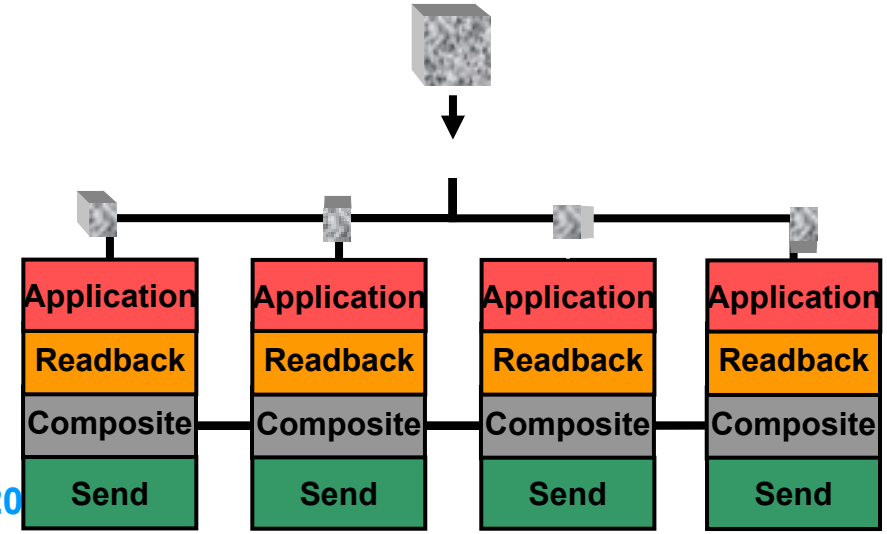

Network Indirection Through SW Virtualization

client

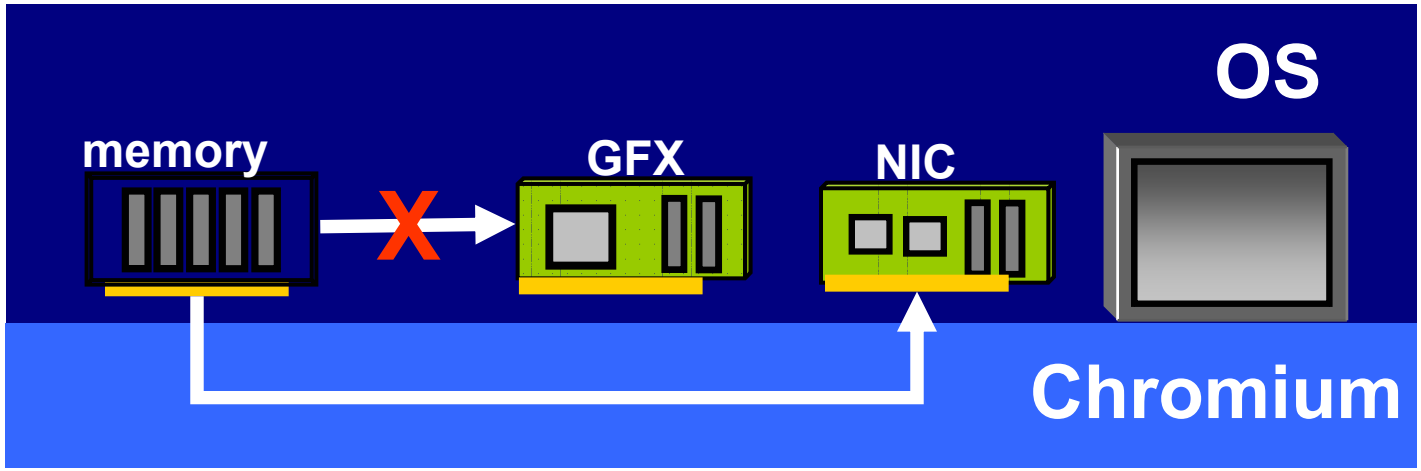

server

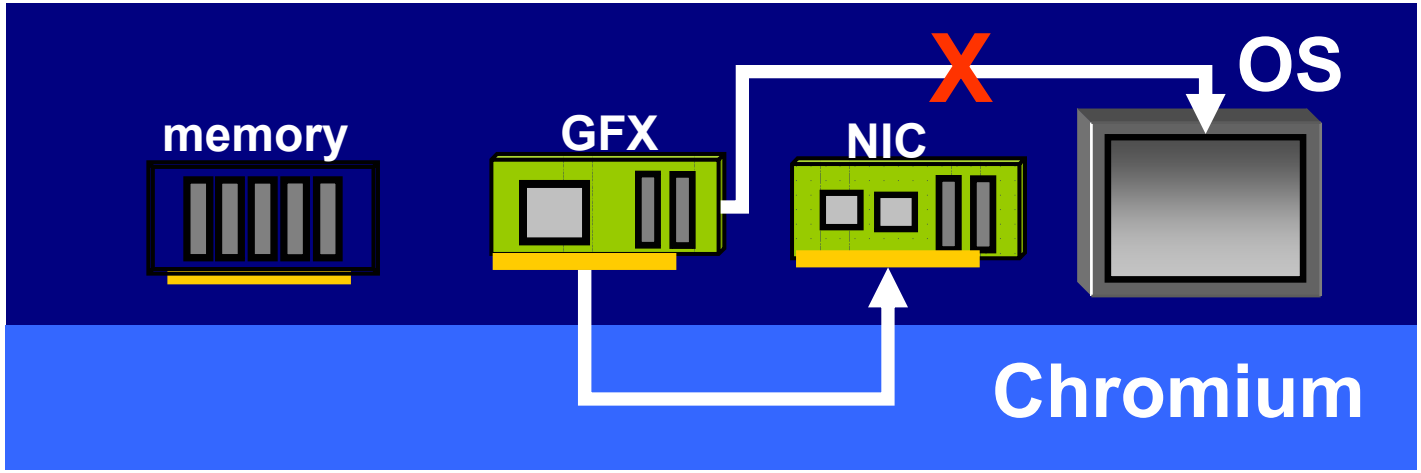

Network Transmission of Hardware Rendered Pixels

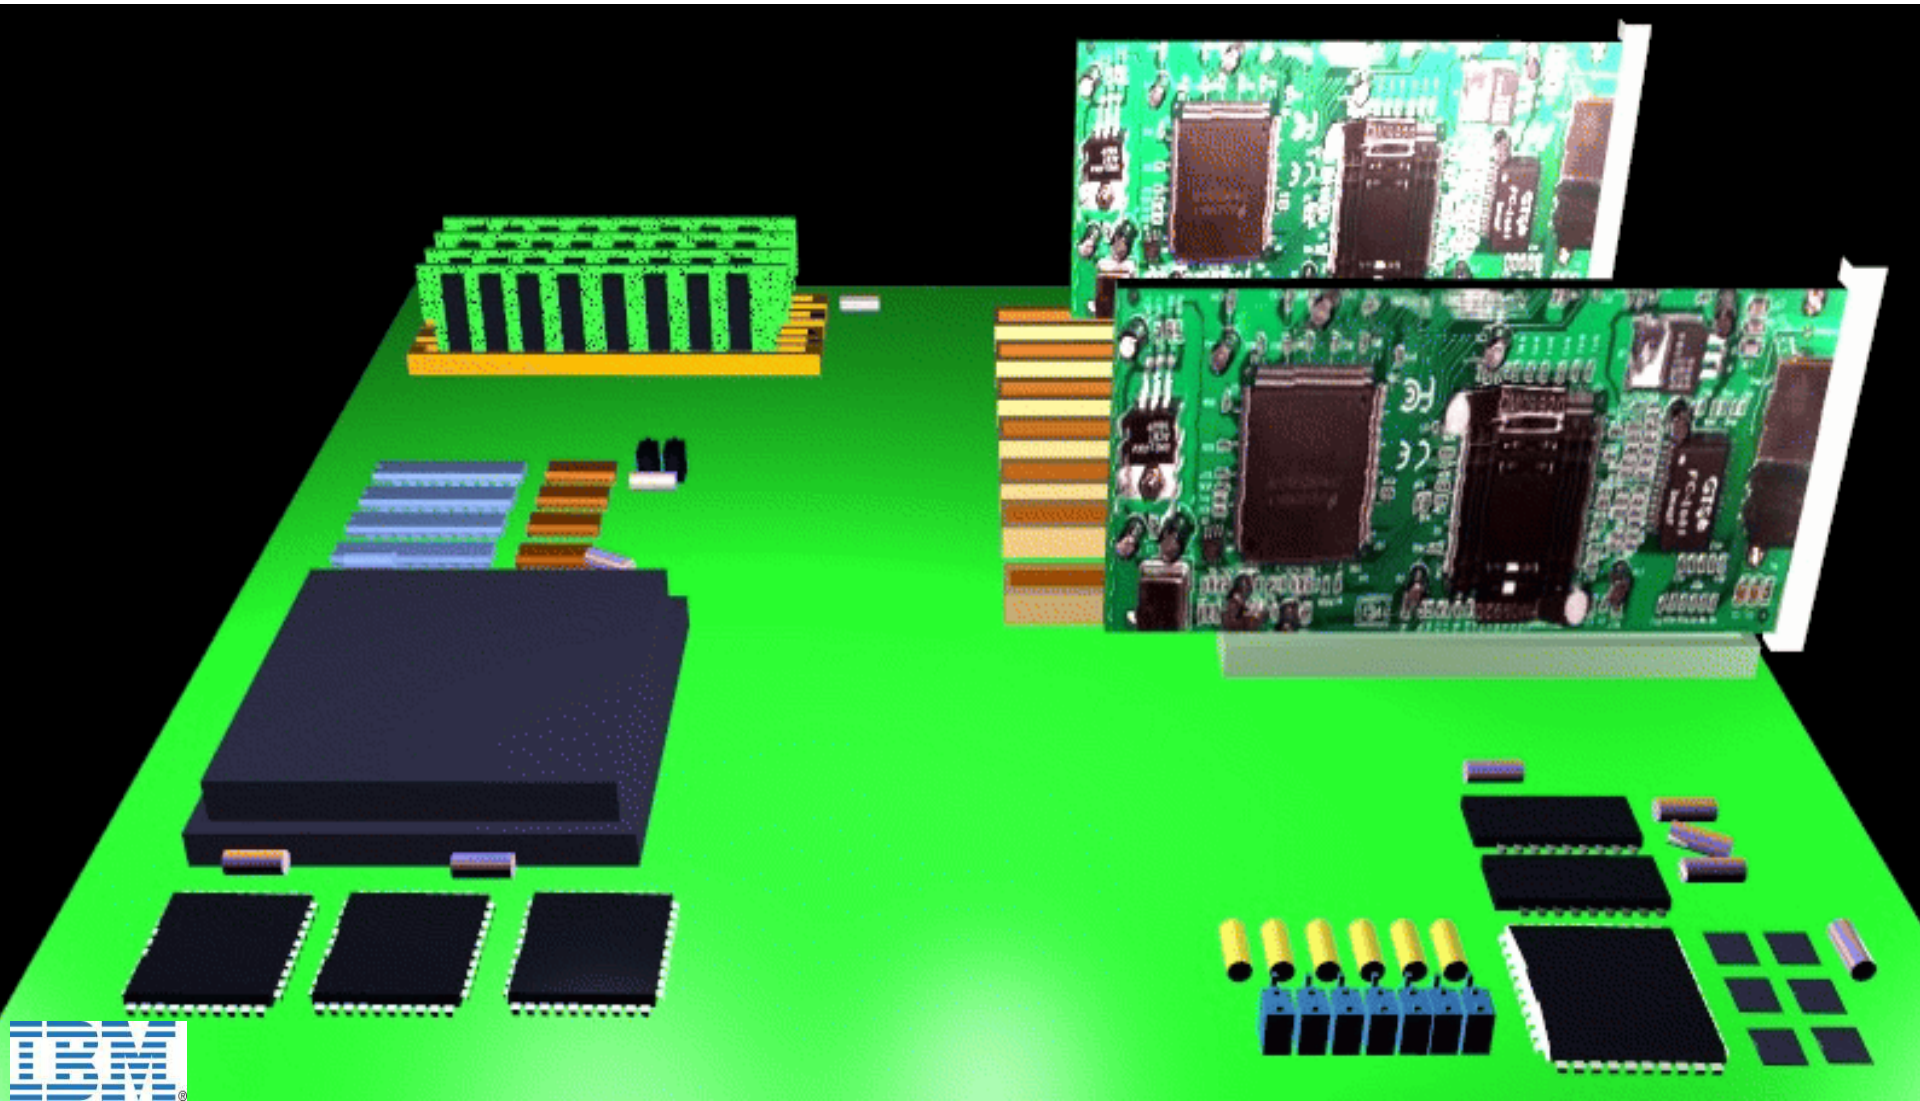

1. Transfer geometry over AGP bus
2. Render geometry on accelerator
3. glReadpixels()
4. Packetize and transmit pixels

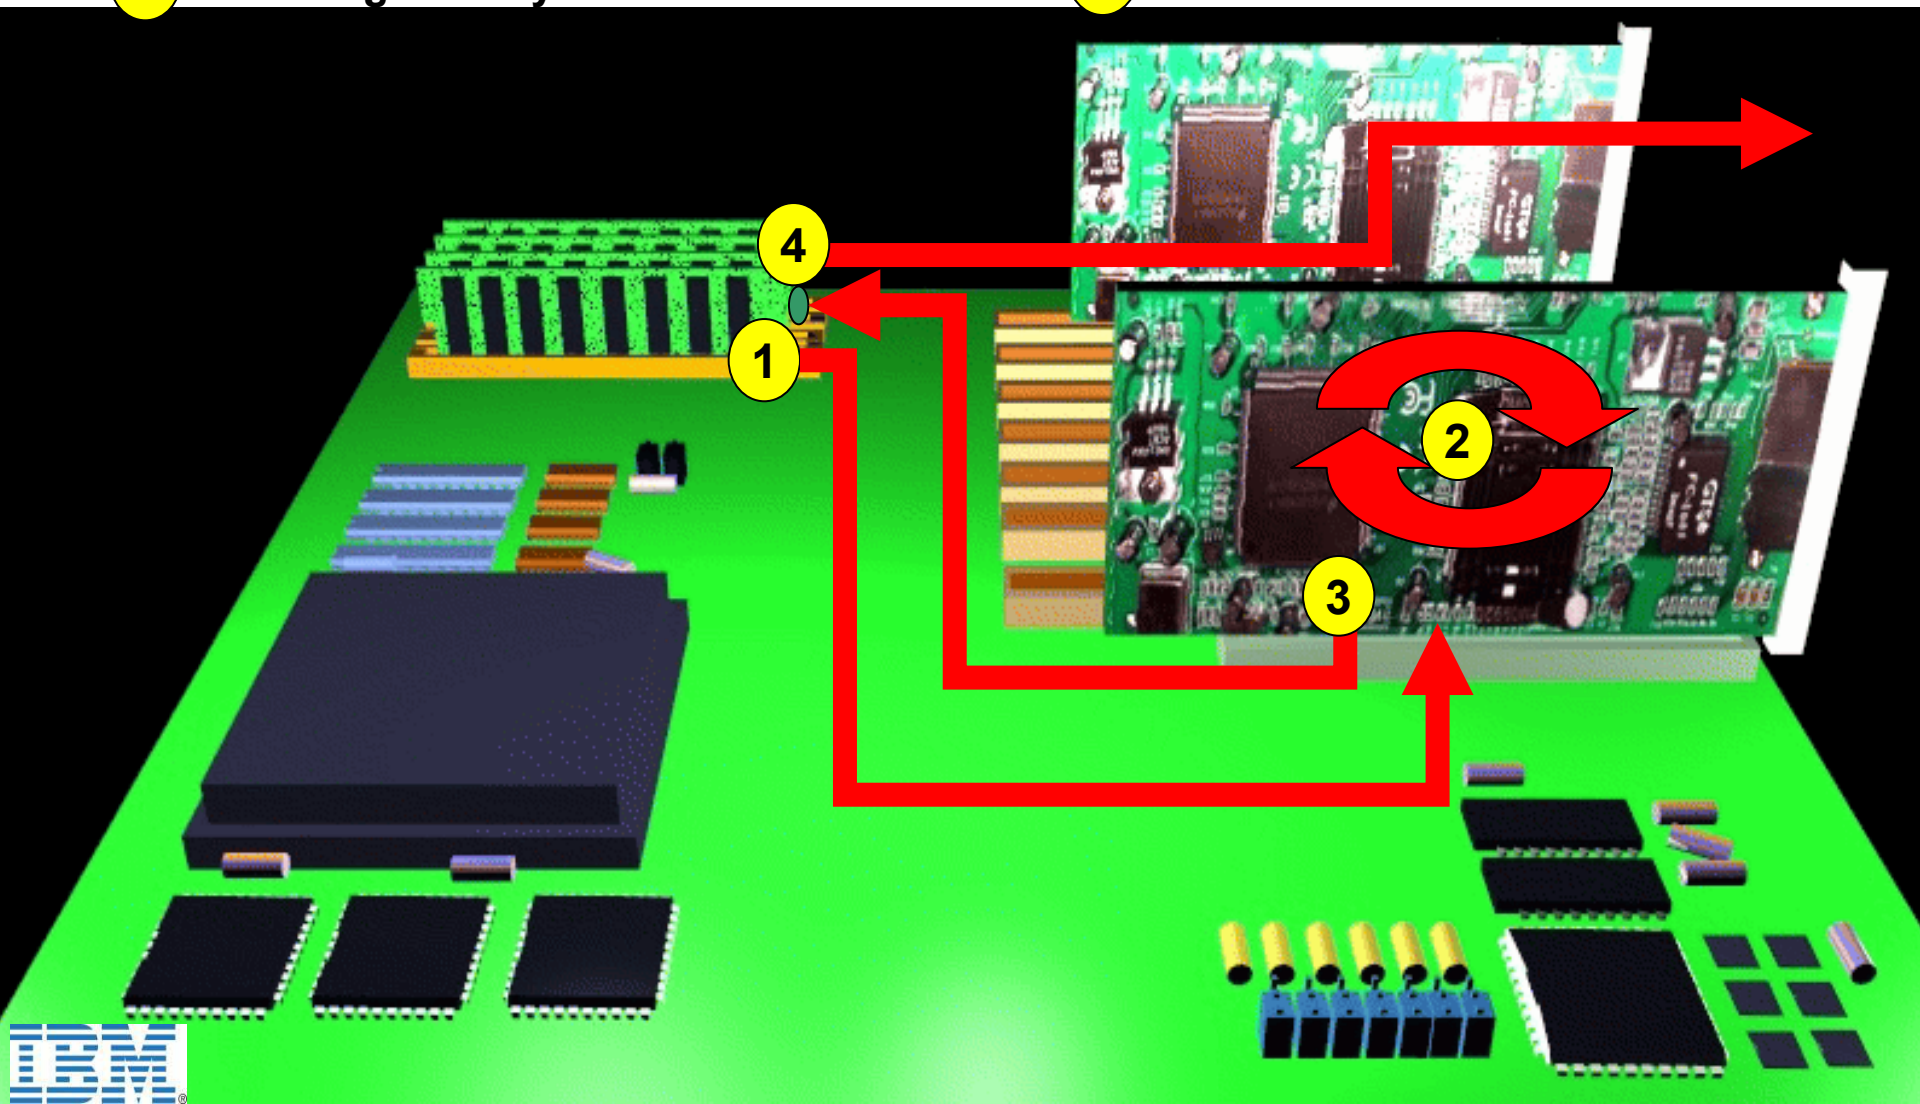

Performance Enablers

- Message Passing Interface (MPI)
- User Datagram Protocol (UDP)
- Link Aggregation
- Jumbo Frames
- Staged Pixel Readback

Three Visualization Environments for Deep View

1. The Scalable Graphics Engine (SGE)

Specialized network attached video frame buffer hw along w/ specialized sw for enabling cluster visualization and high resolution displays.

2. The Scalable Software Concentrator

Specialized software for enabling cluster visualization and SGE-like capability on IBM IntelliStations.

3. Remote Visual Networking

Stack of software technologies for enabling interactive, remote visualization to thin clients attached to low bandwidth connections.

Three Visualization Environments for Deep View

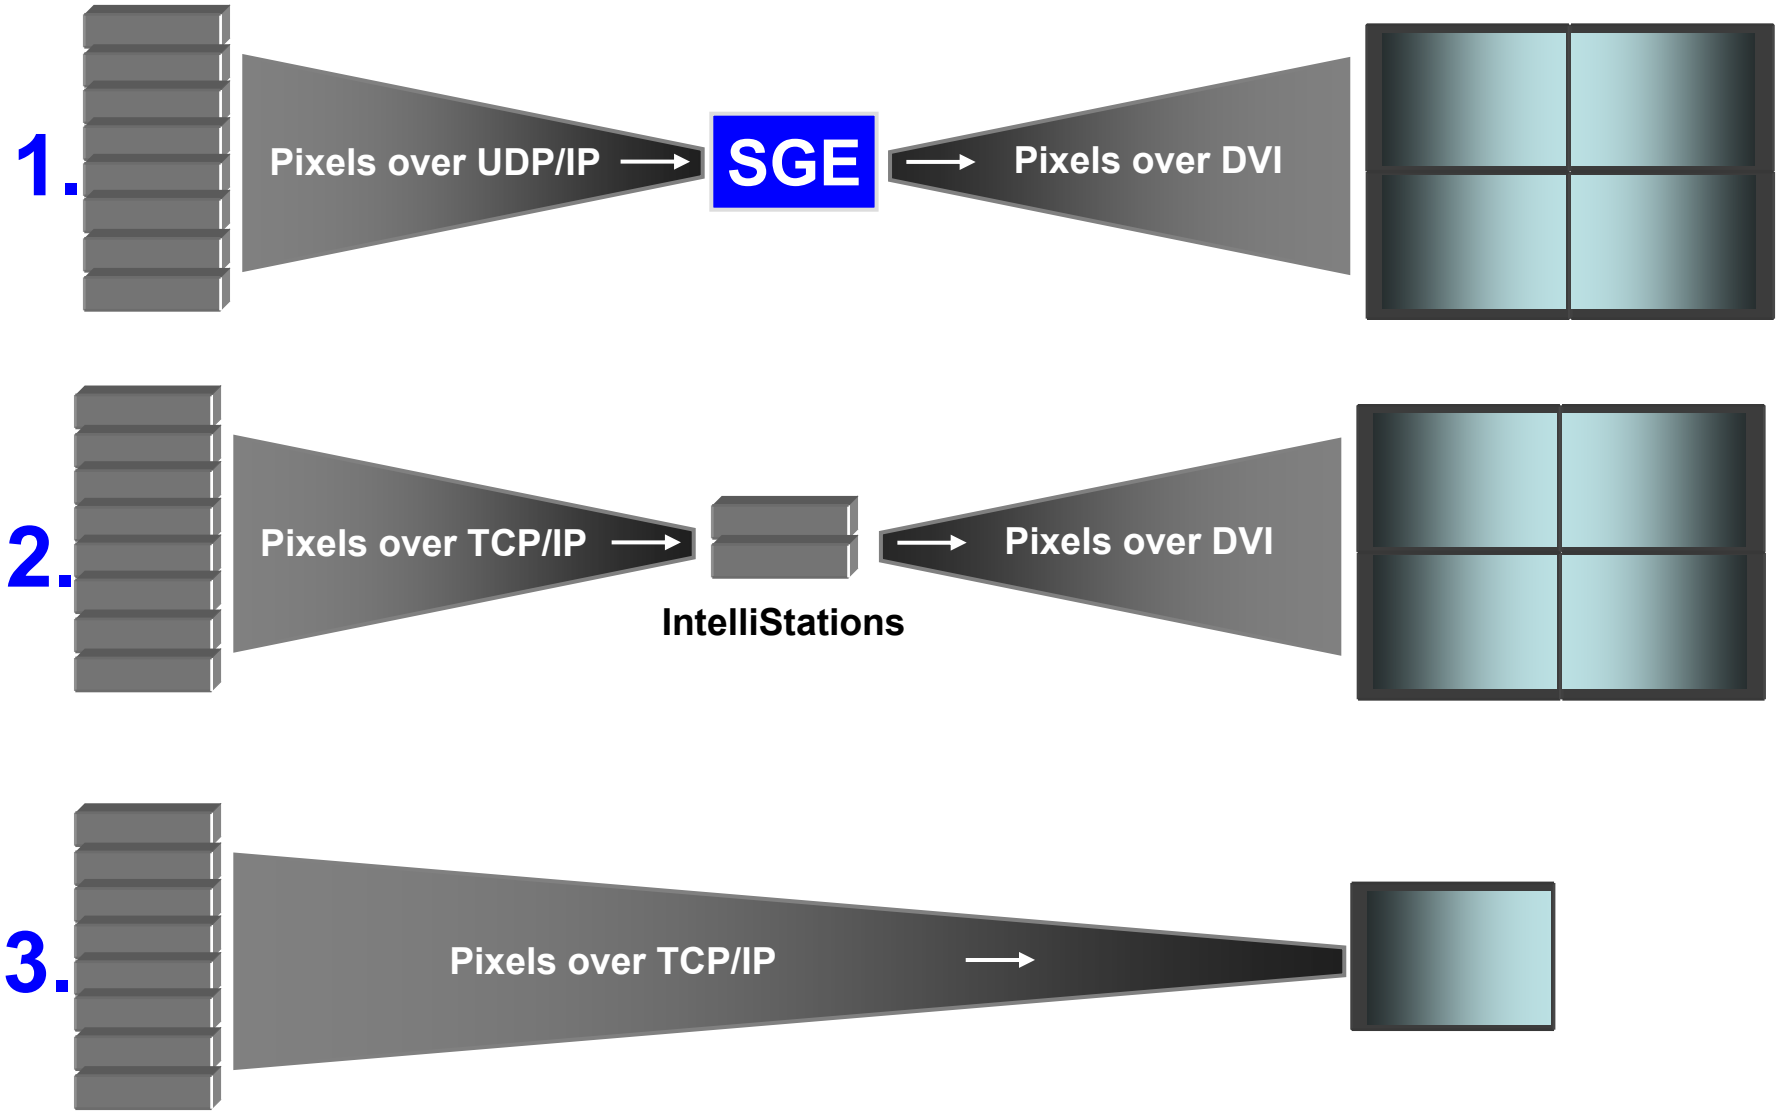

1. The SGE

The SGE3 desktop framebuffer w/ 9.2Mpel IBM T221 display

- Supports up to 16x gigabit ethernet inputs for a total peak throughput of 2 GB/sec.
- 4x Dual-link DVI outputs

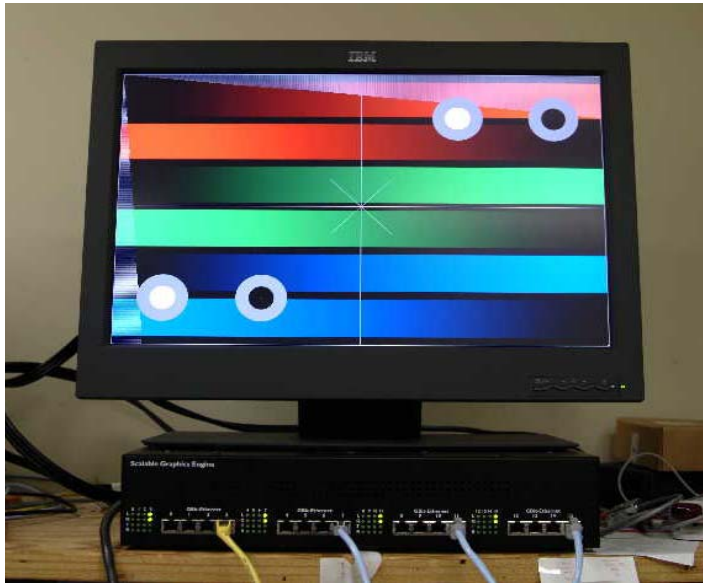

SGE can be used to enable:

- High end cluster visualization
- Remote thin client desktops

SGE3 Packaging: Desktop or Rackmountable

Front

Back

SGE3 Internal Design

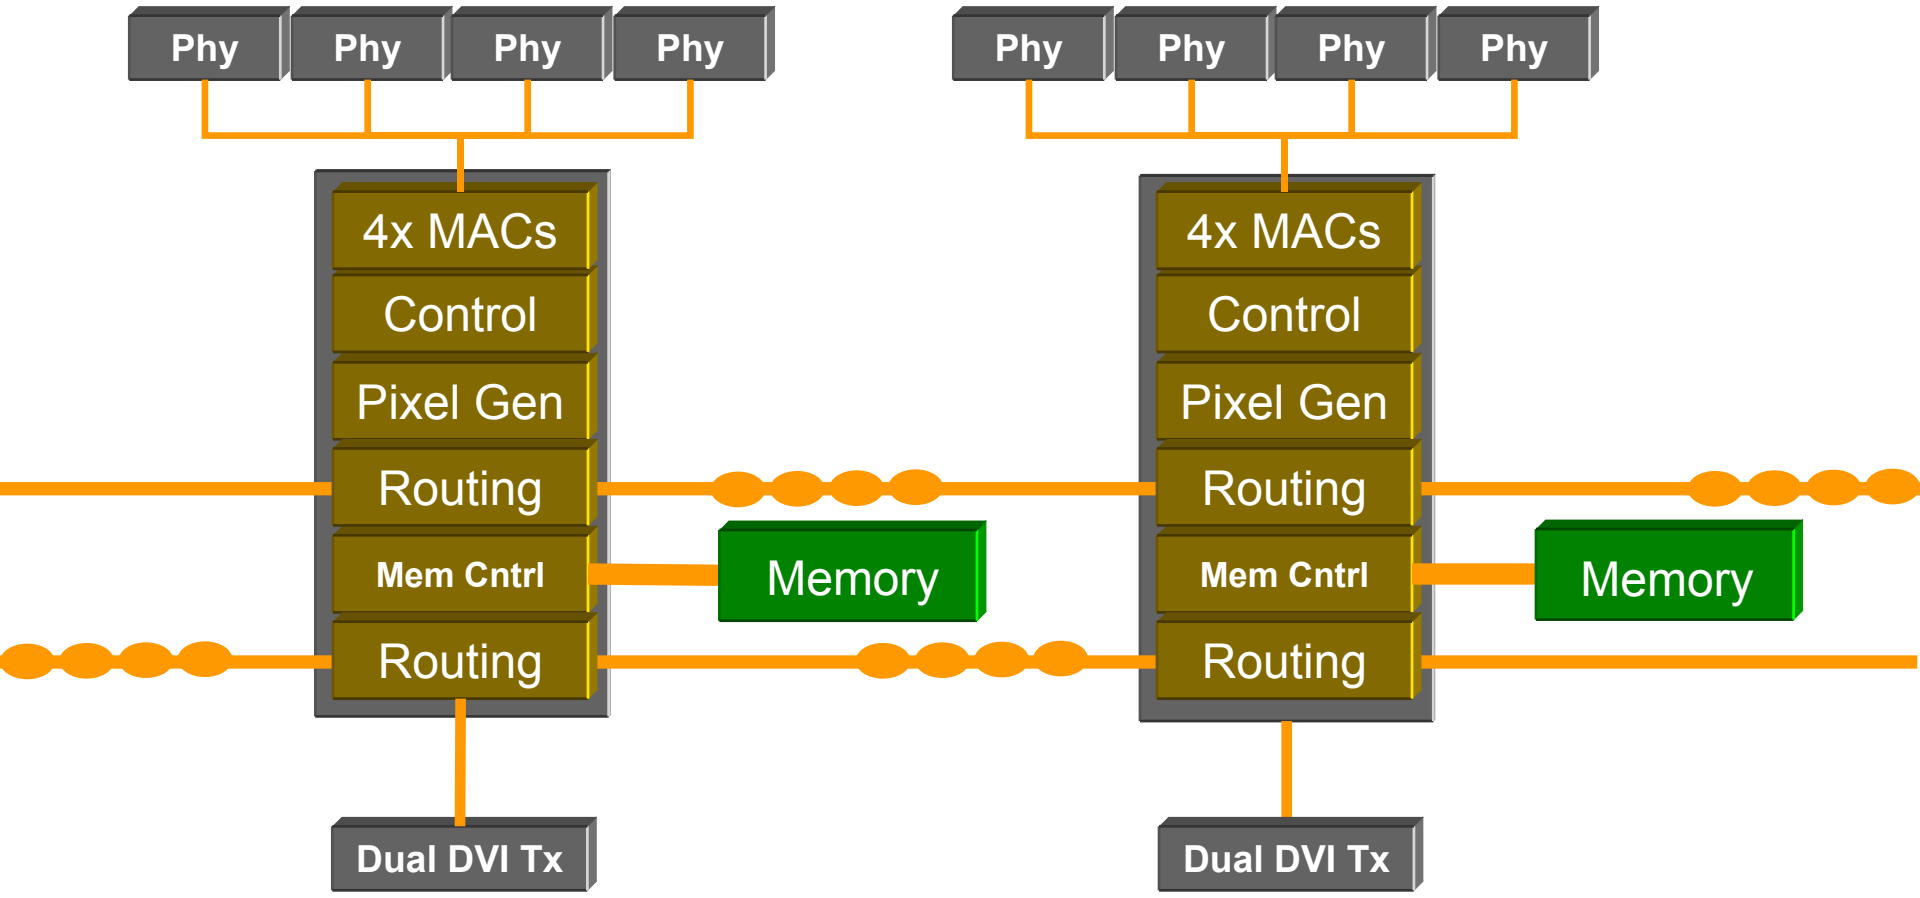

The SGE Software

- SgeControl software for configuring the environment
- A modified X11 server to allow pixel tunneling
- A client library for communicating with SGE
- sgeSPU
- SGE simulator for development
- Sample applications
- MPI (required)

Pixel Tunneling w/ SGE

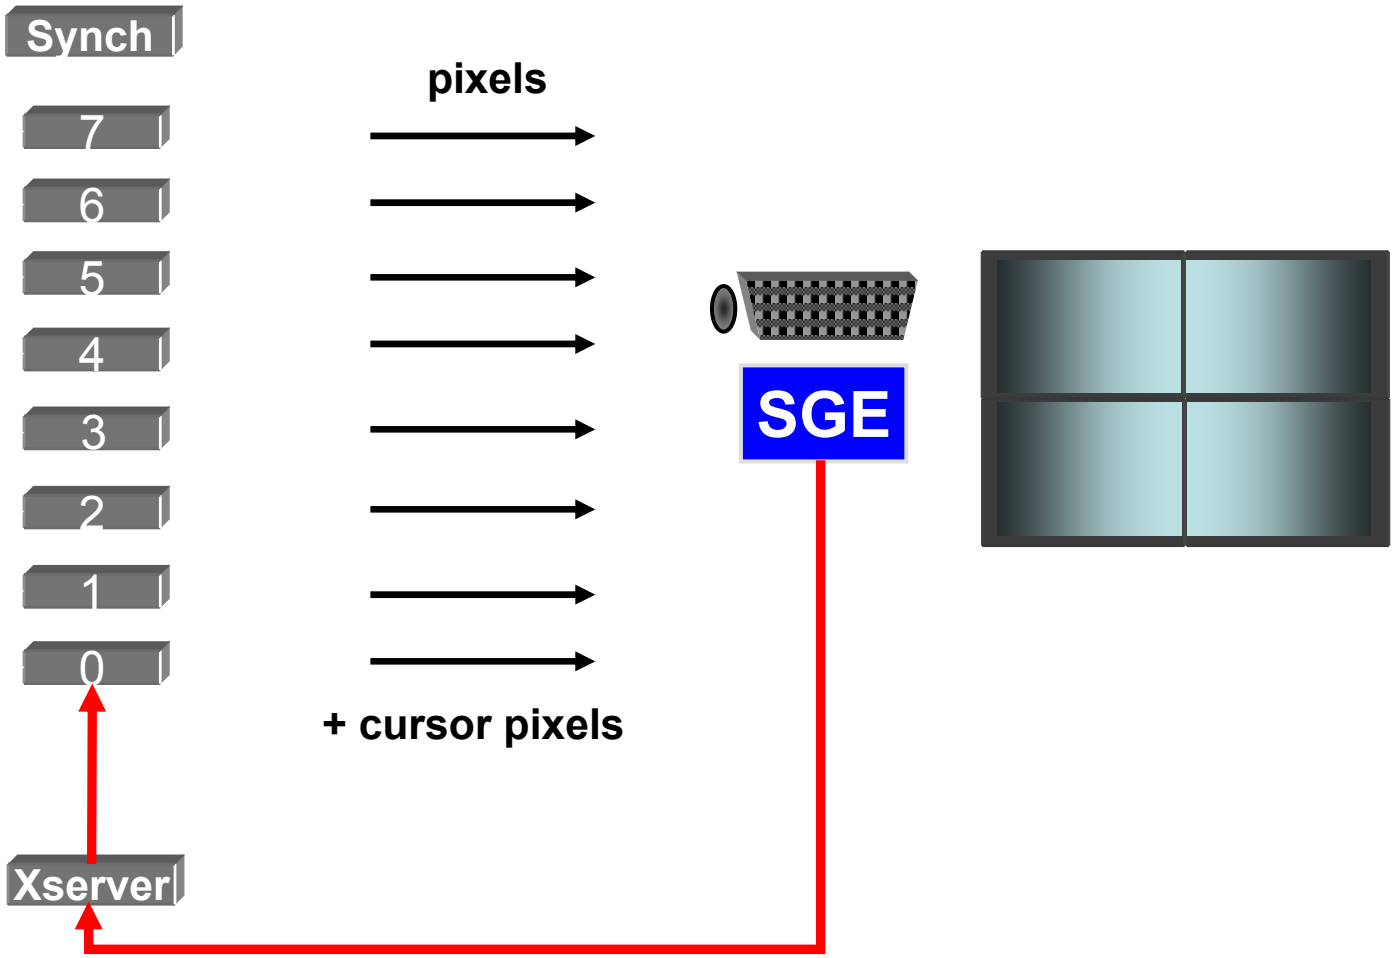

2. The Scalable SW Concentrator

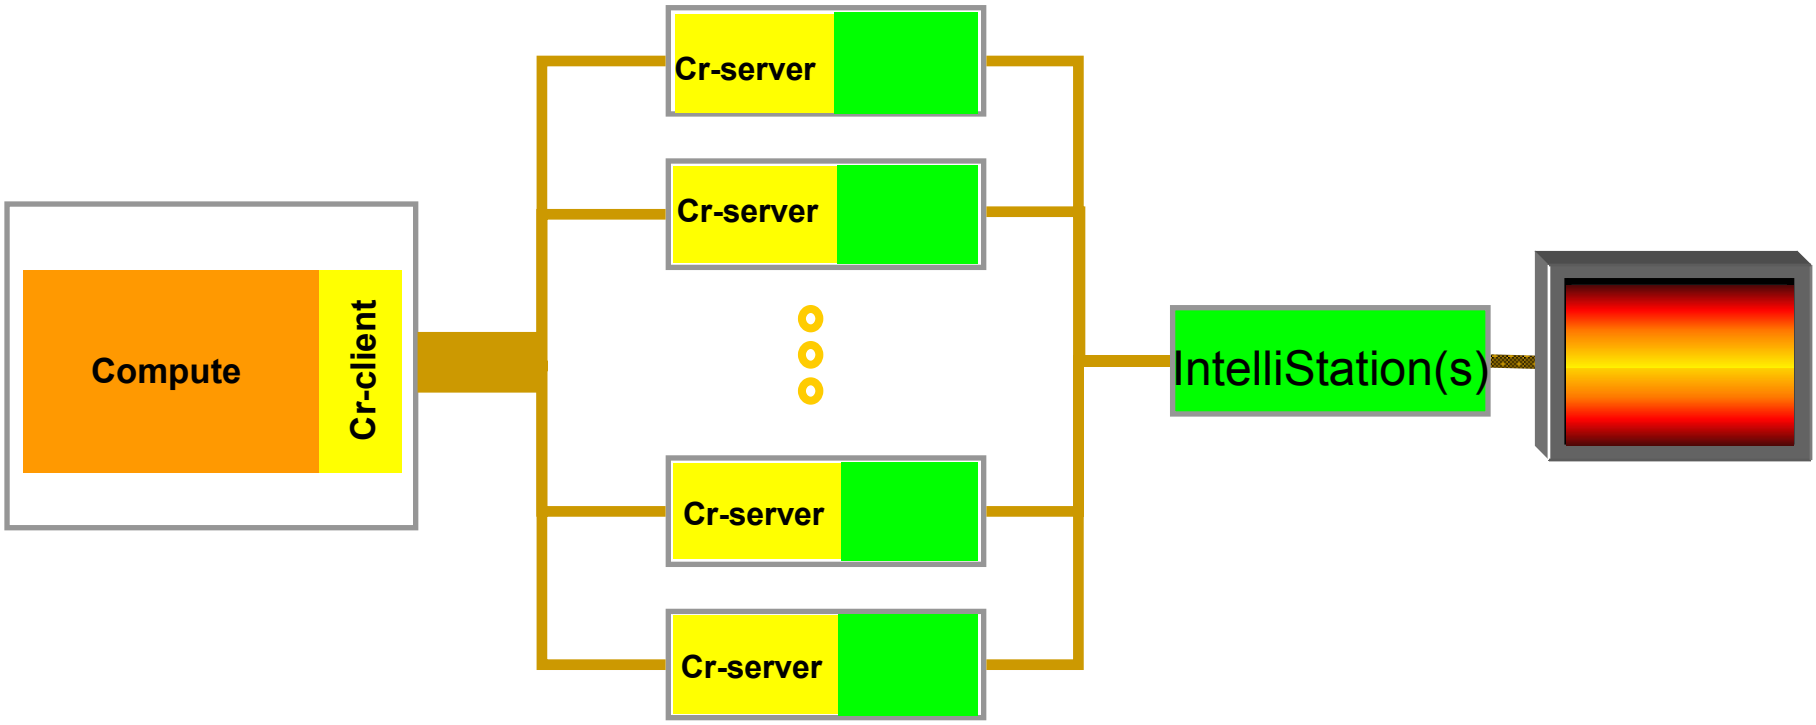

Deep View Cluster

3. Remote Visual Networking

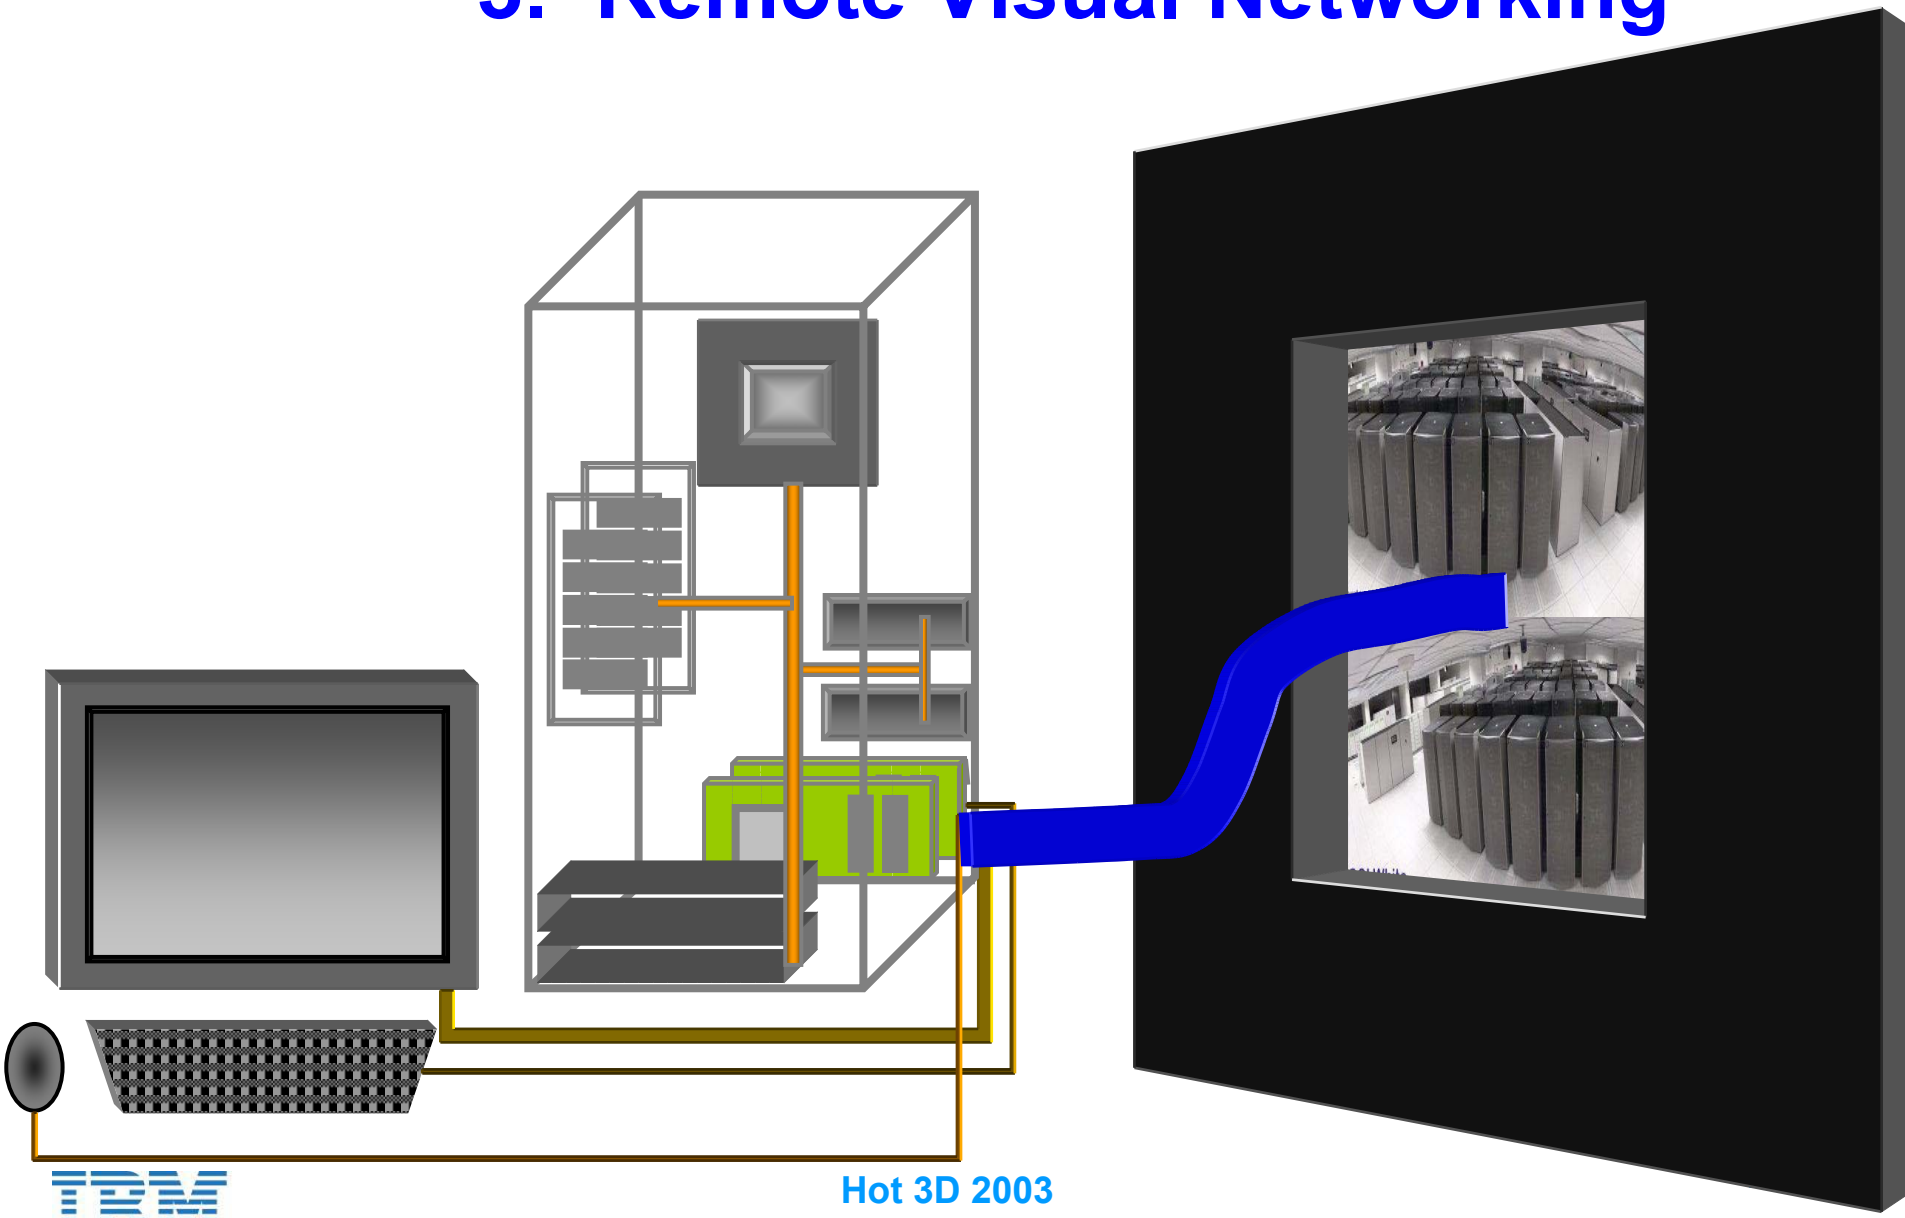

Three Visualization Environments for Deep View

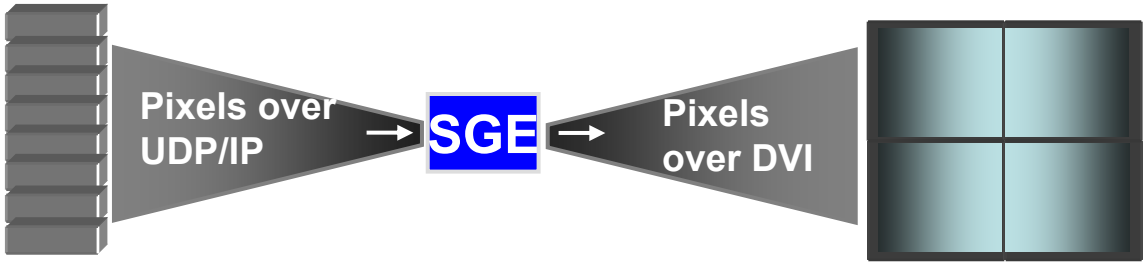

1. Deep View + SGE

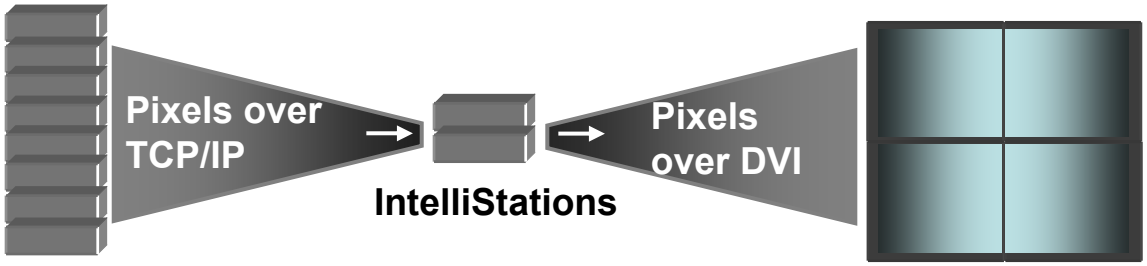

2. Deep View + Scalable SW Concentrator

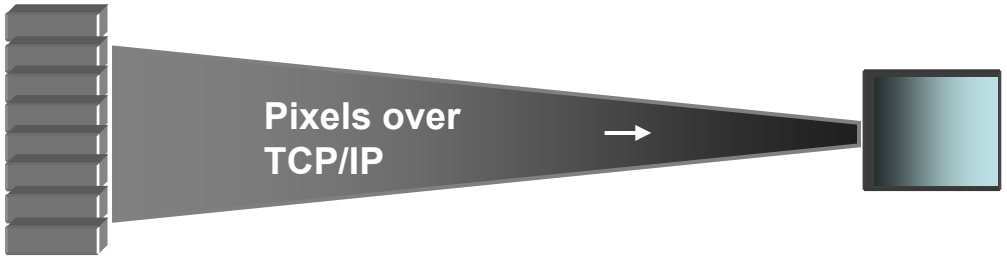

3. Remote Visual Networking