



13. Question: Where is the rendering performed in RVN?

Answer: In order to support the widest array of clients, all the rendering is done on the IntelliStation application nodes. The geometry is rendered using the Intellistation's high-end graphics adapters, fully exploiting the OpenGL hardware acceleration. Then an OpenGL readback of the image is performed, compressed with specialized algorithms, and send over the wire to the remote user. The compression is also selectable, allowing users to optimize the image quality to match the speed of their network connection.

14. Question: How does RVN work?

Answer: RVN compresses a 3D image without having to modify the application, while still utilizing the hardware acceleration of the graphics adapter. Much like SVN, we move the standard OpenGL shared object and replace it with a special RVN libGL.so. When an OpenGL call is made, DCV will intercept the call, pass it to the standard libGL.so to be rendered, then do an OpenGL readback of the data and compress it and send it to the remote user. Some RVN code is required on the client in order to decompress the data.

15. Question: What problem does RVN solve?

Answer: RVN allows enterprises to remove expensive workstations from every engineer's desk, and consolidate them into a managed research that can be shared among a larger number of workers.

© Copyright IBM Corporation 2005

IBM Systems Group
Route 100
Somers, NY 10589

Produced in the United States
09-05
All Rights Reserved.

The following terms are trademarks of International Business Machines Corporation in the United States and/or other countries: IBM®, IntelliStation®.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product and service names may be trademarks of others.

References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates.

IBM hardware products are manufactured from new parts, or new and used parts. In some cases, the hardware product may not be new and may have been previously installed. Regardless, our warranty terms apply.


All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All information in these materials is subject to change without notice. ALL INFORMATION IS PROVIDED ON AN "AS IS" BASIS, WITHOUT ANY WARRANTY OF ANY KIND.

The IBM home page on the Internet can be found at **ibm.com**

IBM Deep Computing Visualization

Frequently Asked Questions



1. Question: What is visualization?

Answer: What some believe to be the next evolution in business intelligence, visualization is the process of transforming data into insight. And as a process or communication medium, the goal of visualization is to perform reliable decision support for actionable results. In translating abstract data and rendering it into graphical information, visualization allows organizations to display that data as images and higher dimensional models in order to simplify understanding, analysis and ultimately their decision making.

2. Question: What are some examples of industries that are benefiting from visualization?

Answer:

- Upstream petroleum companies' view of a visual model of an oil reservoir, for example, can make it easier to decide on the extraction method that will yield the most usable oil.
- Automotive and aerospace design labs use a visual model of a crash simulation to draw the designer's eye to structural elements that need to be reinforced.
- Healthcare professionals' detailed view of a patient's circulatory system shows areas of improved or unchanged blood flow following a planned surgical procedure.
- Financial services firms turn to visual representations for better financial modeling, fraud detection, or high performance analytics.

3. Question: What is Deep Computing Visualization (DCV)?

Answer: DCV is a new IBM® visualization solution and strategy based on rack-mounted clusters of IntelliStations® coupled with cluster rendering software from IBM Research that enables both local and remote visual collaboration over networks.

4. Question: What is the role of visualization in Deep Computing?

Answer: The role of visualization in Deep Computing is to allow organizations to gain advanced capabilities for insight into exploding data volumes, and to enable collaboration through both large scale immersive and remote visualization.

5. Question: Why is IBM entering the visualization market?

Answer: IBM believes that now is the perfect time to leverage the power of commodity graphics adapters, fast interconnects, and Linux® clustering technology to bring forward a standards-based visualization solution. It is the intention of IBM to bring visualization capabilities to desktops around the world.

6. Question: What capabilities does DCV offer?

Answer: A scalable, flexible, and modular solution based on industry standard technologies, DCV has two capabilities – SVN or Scalable Visual Networking for immersive environments, and RVN or Remote Visual Networking for remote and collaboration applications. In the future, DCV will also be adding SPVN or Scalable Parallel Visual Networking for scalable distributed cluster visualization applications.

7. Question: Is it necessary to modify my applications in order to use DCV?

Answer: No, one of the primary value propositions of DCV is that it runs your existing applications, unmodified. Using virtualization technologies DCV runs off-the-shelf OpenGL applications and delivers both the SVN and RVN capabilities.

8. Question: How can DCV work with unmodified applications?

Answer: DCV intercepts the OpenGL commands that are used to drive the 3D graphics cards then passes them to either the SVN or RVN libraries.

9. Question: What exactly is Scalable Visual Networking (SVN)?

Answer: SVN enables large, multi-panel, high-resolution displays for immersive visualization environments using standard unmodified OpenGL applications. SVN is also extensible, allowing developers to add functionality should they desire.

10. Question: What capability does SVN provide?

Answer: SVN distributes a 3D object's geometry over a high-speed network to a cluster of IntelliStation nodes, which then render their portion of the image and display it on the screen.

11. Question: How does SVN work?

Answer: SVN scales up an image without having to modify the application. We move the standard OpenGL shared object (on Linux called libGL.so) and replace it with a special SVN libGL.so. When an OpenGL call is made, DCV will intercept the call, and send it with the geometry using MPI over InfiniBand to the appropriate IntelliStations for rendering, where it will then be displayed out the DVI port to the display device.

12. Question: What is Remote Visual Networking (RVN)?

Answer: RVN is remote desktop software specialized for 3D applications. But unlike many competing solutions, RVN exploits the OpenGL hardware acceleration of the IntelliStation's graphics adapters to provide the utmost in performance.