AccessionIndex: TCD-SCSS-T.20121208.097
Accession Date: 8-Dec-2012
Accession By: Dr.Brian Coghlan
Object name: VRengine
Vintage: c.2005
Synopsis: 9-node virtual reality engine using 600MB/s SCI 2-d toroidal interconnect.

Description:
The VRengine was created as a special-purpose engine for virtual reality. It was constructed c.2005 for Ronan Watson's PhD research in the School of Computer Science and Statistics, Trinity College Dublin.

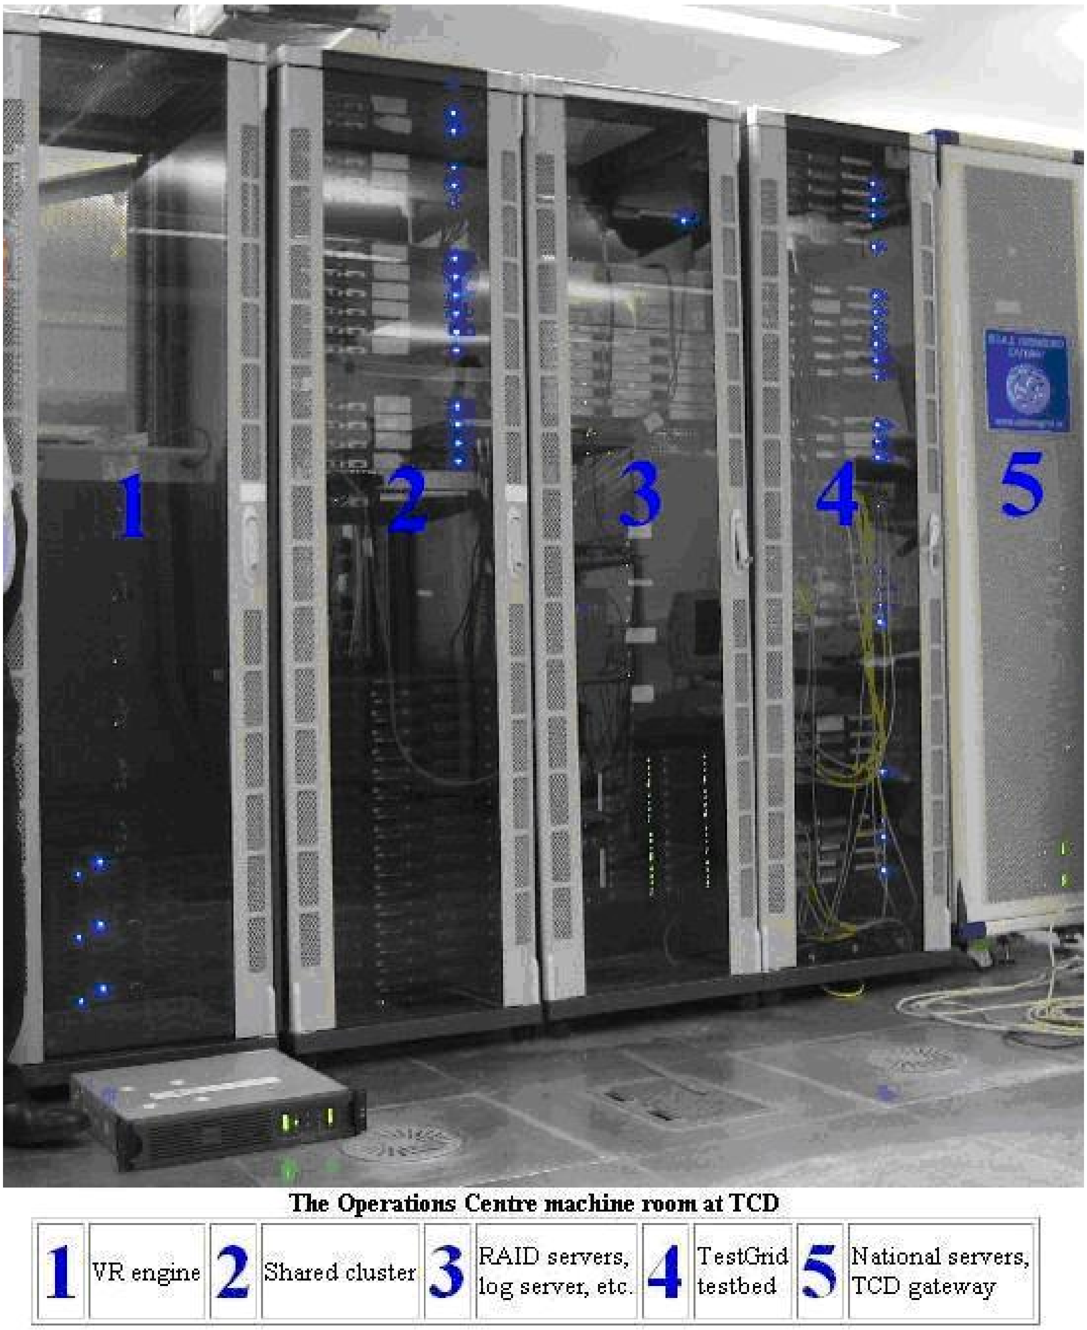

Figure 1: VRengine (left rack)

The first Irish Beowulf cluster was built in 1997 in the Dublin Institute for Advanced Studies, then the second in TCD's Department of Computer Science, the first of a number in the department, some very adventurous, as here. The intention was to explore the use of Grid for virtual reality, with multiple users using Grid for virtual reality compute tasks and sharing the VRengine amongst these users for VR rendering. Figure 2 illustrates its usage to model and render the EGEE Grid as a virtual reality world.

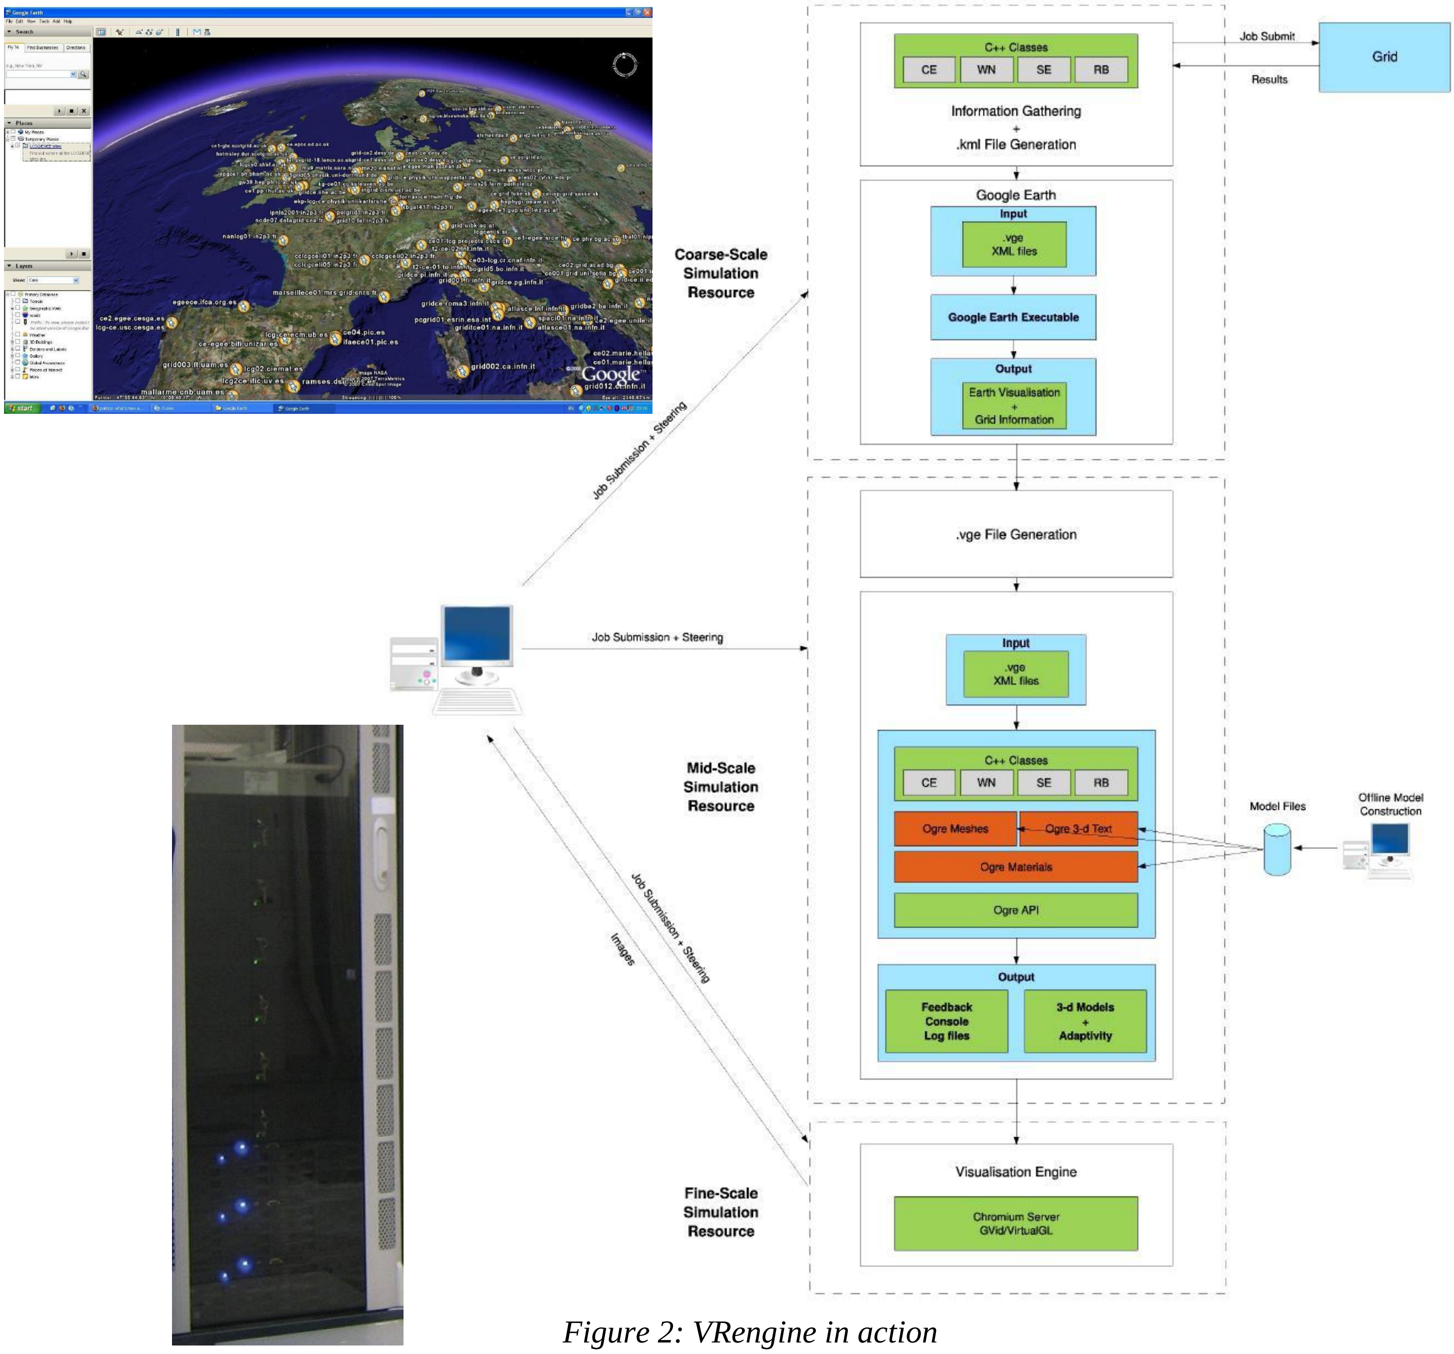

Figure 2: VRengine in action

The architecture was based on the use of two-dimensional 600MB/s 64bit/66Mhz D339 SCI (Scalable Coherent Interconnect) cluster interconnect adapters from Dolphin Interconnect Solutions in Norway. Nine compute nodes were configured as a private 3 x 3 toroidal SCI network. This private 2-d toroidal network existed in parallel with a public switched 1Gbps Ethernet network that allowed it to be viewed as a Beowulf cluster.

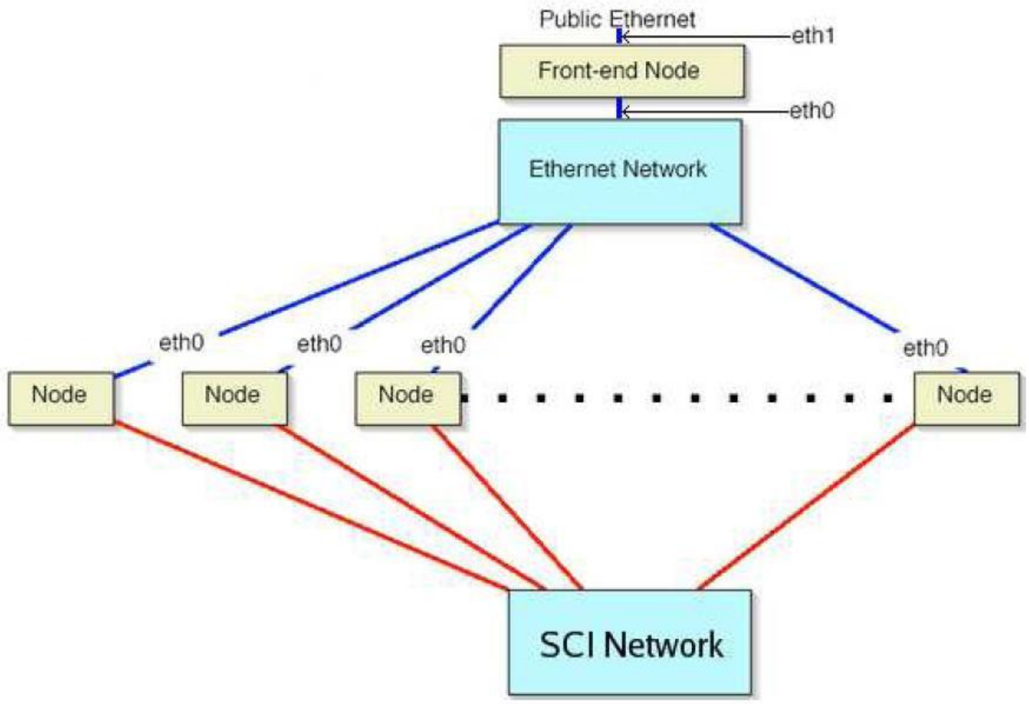

Figure 3: VRengine architecture

The VRengine was set up using the *Rocks Cluster* distribution, which included specialized software sets distributed as *Rolls*. In particular interest there was a *Vizualisation Roll* that provided a software set to enhance cluster performance in graphical computations by using *Chromium* to distribute the graphical workload throughout the cluster.

The low level user space interface for SCI was called *SISCI*, which provided basic mechanisms to share memory segments between nodes and to transfer data between them. At the time the only protocol that was supported by Chromium was TCP/IP. Therefore the SCI hardware needed drivers that would support communication over the high speed interconnect using TCP/IP. This software was provided in open source form by the manufacturers of the SCI hardware, and was called *SuperSockets*. It provided a fast and transparent way for applications that employed Berkeley sockets (TCP/UDP/IP) to use SCI as the transport medium. The major benefits were plug and play, high bandwidth and a much lower socket latency than network technologies like Gigabit Ethernet, Infiniband and Myranet.

For the VREngine the Vizualization (Viz) Roll included the software that was felt necessary for high performance virtual reality environments. The Viz Roll added a new Rocks appliance type called a *Tile*. Just as a Rocks *Compute* appliance could be used to build compute clusters, this new Tile appliance could be used to build tiled-display visualization clusters. The core software to drive the tiled display was called DMX, and this created a single unified X11 desktop from multiple graphics cards, monitors, and computers. Visualization clusters could be built to create larger displays up to thousands-by-thousands of pixels, especially for OpenGL applications.

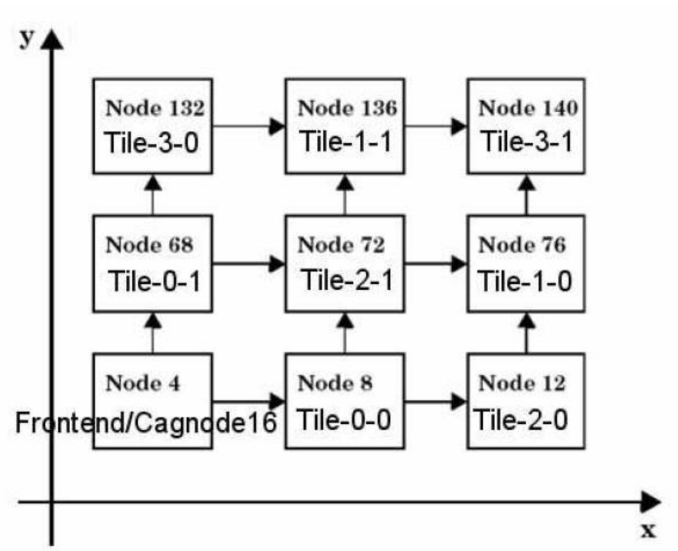

Figure 4: VRengine tiled display topology and hostname numbering

Chromium provided graphics scalability on three fundamental axes: quantity of data, rendering capability, and display capability. Chromium's data scalability meant it could handle increasingly larger datasets on increasingly larger clusters, rendering more triangles per second or more pixels per second. Its rendering capability could be increased by aggregating together multiple commodity hardware to draw an image. Its display capability could be increased by scaling its rendering capabilities to generate larger output images, for example via tiled displays made up of multiple individual displays or projectors.

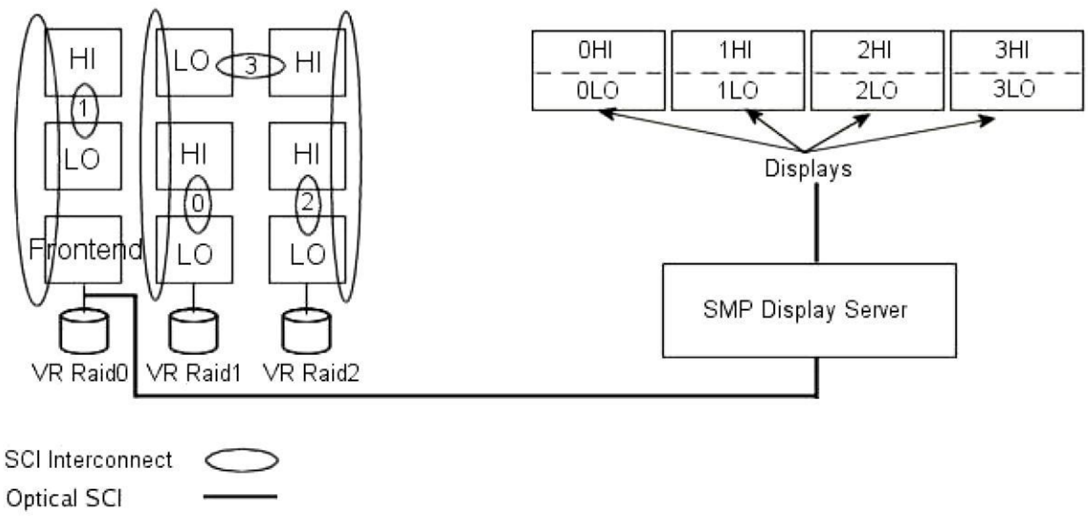

Figure 5: VRengine, showing intended tiled display topology

Chromium combined a number of key attributes into a single package:
(a) Novel method for synchronizing parallel graphics commands;
(b) Streaming graphics pipeline based on the industry standard OpenGL API;
(c) Support for multiple physical display devices clustered together;
(d) Support for aggregation of the output of multiple graphics cards;.
(e) Application-transparent plugin mechanisms supporting custom graphics pipelines.

These attributes were exploited using the VRengine. The related PhD research work proposed a generic framework for grid-enabled visualisation and computational steering that would be amenable to characterisation using mathematical models based on benchmarked resources. The models would enable designers, providers or users to select the most suited interactive computational steering and rendering resources available to them on the Grid. The framework for visualisation aimed to give greater performance than a desktop computer could provide, at a cost much cheaper than a CAVE, while allowing simultaneous usage by multiple users.

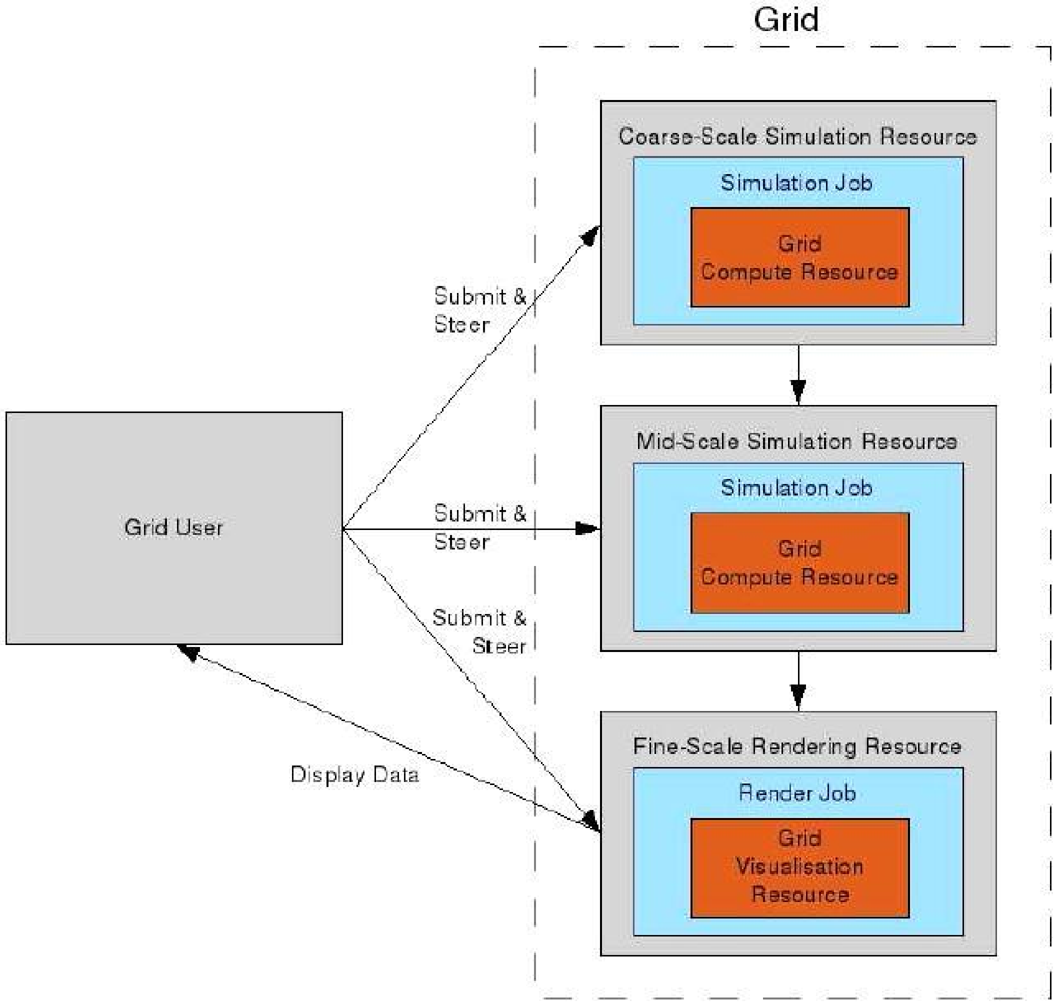

Figure 6: Proposed virtual reality execution framework

Associated grid job submissions closely matched the three scales of the framework as shown in Fig.7. A virtual reality task submitted from the user's workstation was first partitioned and then its subtasks sent to the appropriate scale of computing resource by the grid resource broker (RB). Each resource's grid front-end computing element (CE) then scheduled the subtasks on its grid worker nodes (WN), which engaged any special-purpose facility (such as the VRengine) as needed. The result was conveyed back to the user's workstation. The human-in-the-loop was supported by techniques perfected in the Int.EU.Grid project in which TCD was a participant.

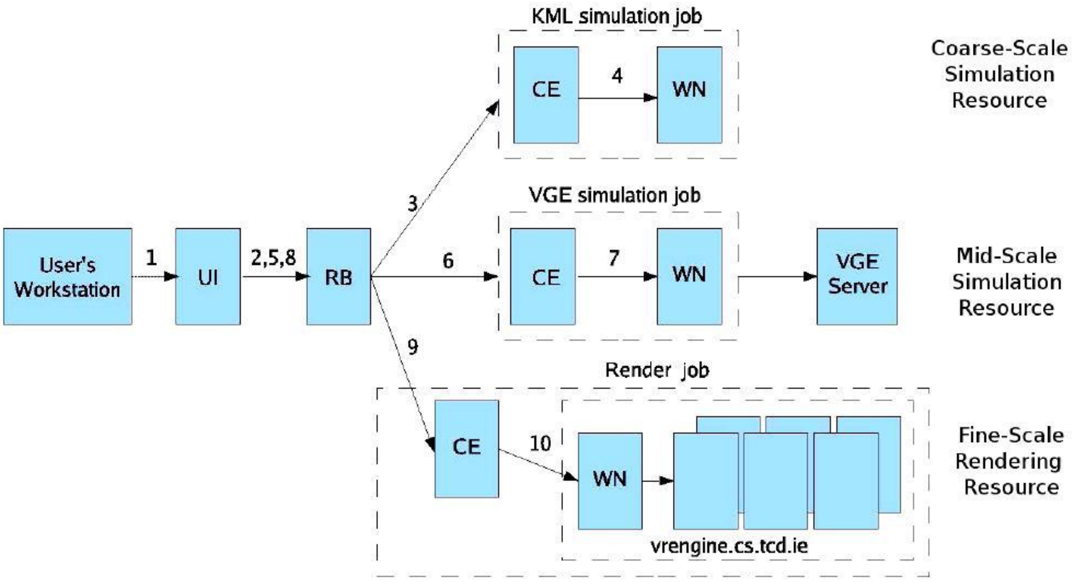

Figure 7: Grid job submissions for proposed virtual reality framework

Related experimental work included initial benchmarking trials of grid resources in the UK (UK NGS), Spain (IFCA), Poland (PSNC) and the Netherlands (NIKHEF), followed by a request for permission to benchmark all resources on the EGEE Grid (270 sites in over 50 countries), which was granted, allowing the construction of models using a larger sample set, based on a very large set of resource benchmarks, and allowing examination of the influence of sample set size. This was a very unusual level of generosity extended to a PhD student by the EGEE Grid project management as well as all the EGEE Regional Operations Centre (ROC) managers who granted permission to use their site resources.

The VRengine was decommissioned in 2012, and the 2-d SCI cards repurposed for another computer architecture project.

The homepage for this catalog is at: https://www.scss.tcd.ie/SCSSTreasuresCatalog/
Click *'Accession Index'* (1st column listed) for related folder, or *'About'* for further guidance. Some of the items below may be more properly part of other categories of this catalog, but are listed here for convenience.

Accession Index	Object and Identification
TCD-SCSS-T.20121208.097	VRengine, 9-node virtual reality engine using 600MB/s SCI 2-d toroidal interconnect, c.2005.
TCD-SCSS-T.20121208.094	Experimental SCSI Cluster, 4-node prototype cluster using SCSI as interconnect, the first cluster constructed in the Department of Computer Science, Trinity College Dublin, and second cluster constructed in the Republic of Ireland, 1997.
TCD-SCSS-T.20121208.095	csTCDie Beowulf Cluster, Departmental cluster using 100Mbps Ethernet as interconnect, the second cluster constructed in the Department of Computer Science, Trinity College Dublin, 1998.

TCD-SCSS-T.20141120.003	csTCDie Grid-Ireland SCI Cluster, 16-node cluster using 400MB/s SCI switched interconnect, the third cluster constructed in the Department of Computer Science, Trinity College Dublin, c.1999.
TCD-SCSS-T.20121208.098	csTCDie Grid Site Beowulf Clusters and Datastore, Complex of clusters & storage (1500 cores/600 TB) using 1Gbps Ethernet interconnect and 10Gbps backbone, participant in DataGrid, EGEE, EGI, and CERN LHC computing. From 2013 repurposed as SCSS Cloud, c.2009.
TCD-SCSS-T.20121208.106	csTCDie PS3 Cluster, Ten nodes from a 16-node Sony Playstation PS3 cluster plus build machine, using 1Gbps Ethernet interconnect and running Yellow Dog Linux, c.2009.
TCD-SCSS-T.20121208.099	csTCDie GPU Cluster, 64-core/32-GPU/16-node cluster using 1Gbps Ethernet interconnect, c.2011.

References:

1. Wikipedia, *Beowulf cluster*, see: https://en.wikipedia.org/wiki/Beowulf_cluster Last accessed 30-May-2023.

2. Watson, R., Maad, S., and Coghlan, B., *Integrating a Common Visualization Service into a Metagrid*, Workshop on State-of-the-Art in Scientific and Parallel Computing (PARA06), Umea, Sweden , June 18-21, 2006.

3. Watson, R. and Maad, S., and Coghlan, B., *Multiscale multimodal visualization on a grid*, Proceedings Crakow Grid Workshop 2006, Crackow, Poland, 15-18 October, 2006, pp535 - 542.

4. Watson, Ronan and Maad, Soha, and Coghlan, Brian, *Virtual grid: adaptive visualization of grids*, Proceedings Crakow Grid Workshop 2007, Crackow, Poland, 15-17 October, 2007, edited by Bubak, Marian and Turala, Michal and Wiatr, Kazimierz, pp204 – 211.

5. Ronan Watson, *A generic framework for grid-enabled visualisation and computational steering, and its characterisation*, PhD Thesis, School of Computer Science and Statistics, Trinity College Dublin, 2010, see: http://www.tara.tcd.ie/handle/2262/77673 Also see: http://www.cs.tcd.ie/coghlan/PhDtheses/RonanWatson-thesis-FINAL-as-sent-to-binders-20110316-1119.pdf Last accessed 30-May-2023.

6. EU ESPRIT Project P23174, *Software Infrastructure for SCI (SISCI)*, 1998.

7. Hugo Kohmann Friedrich Seifert, *SCI SOCKET - A Fast Socket Implementation over SCI*, Dolphin Interconnect Solutions, Norway.