



RenderStream GPGPU HPC and Visualization & Data Analysis Cluster

RenderStream
Austin, TX

<http://www.renderstream.com/>
<http://blog.renderstream.com/?p=600>
info@renderstream.com



Contents

- Introduction
 - Introducing Fermi
- Node Offerings
 - 2 TFLOP
 - 4 TFLOP
 - 8 TFLOP
 - 12 TFLOP



- VDAC Workstation, HPC Clusters and Multi-tiered Storage
- Roadmap
- Conclusion



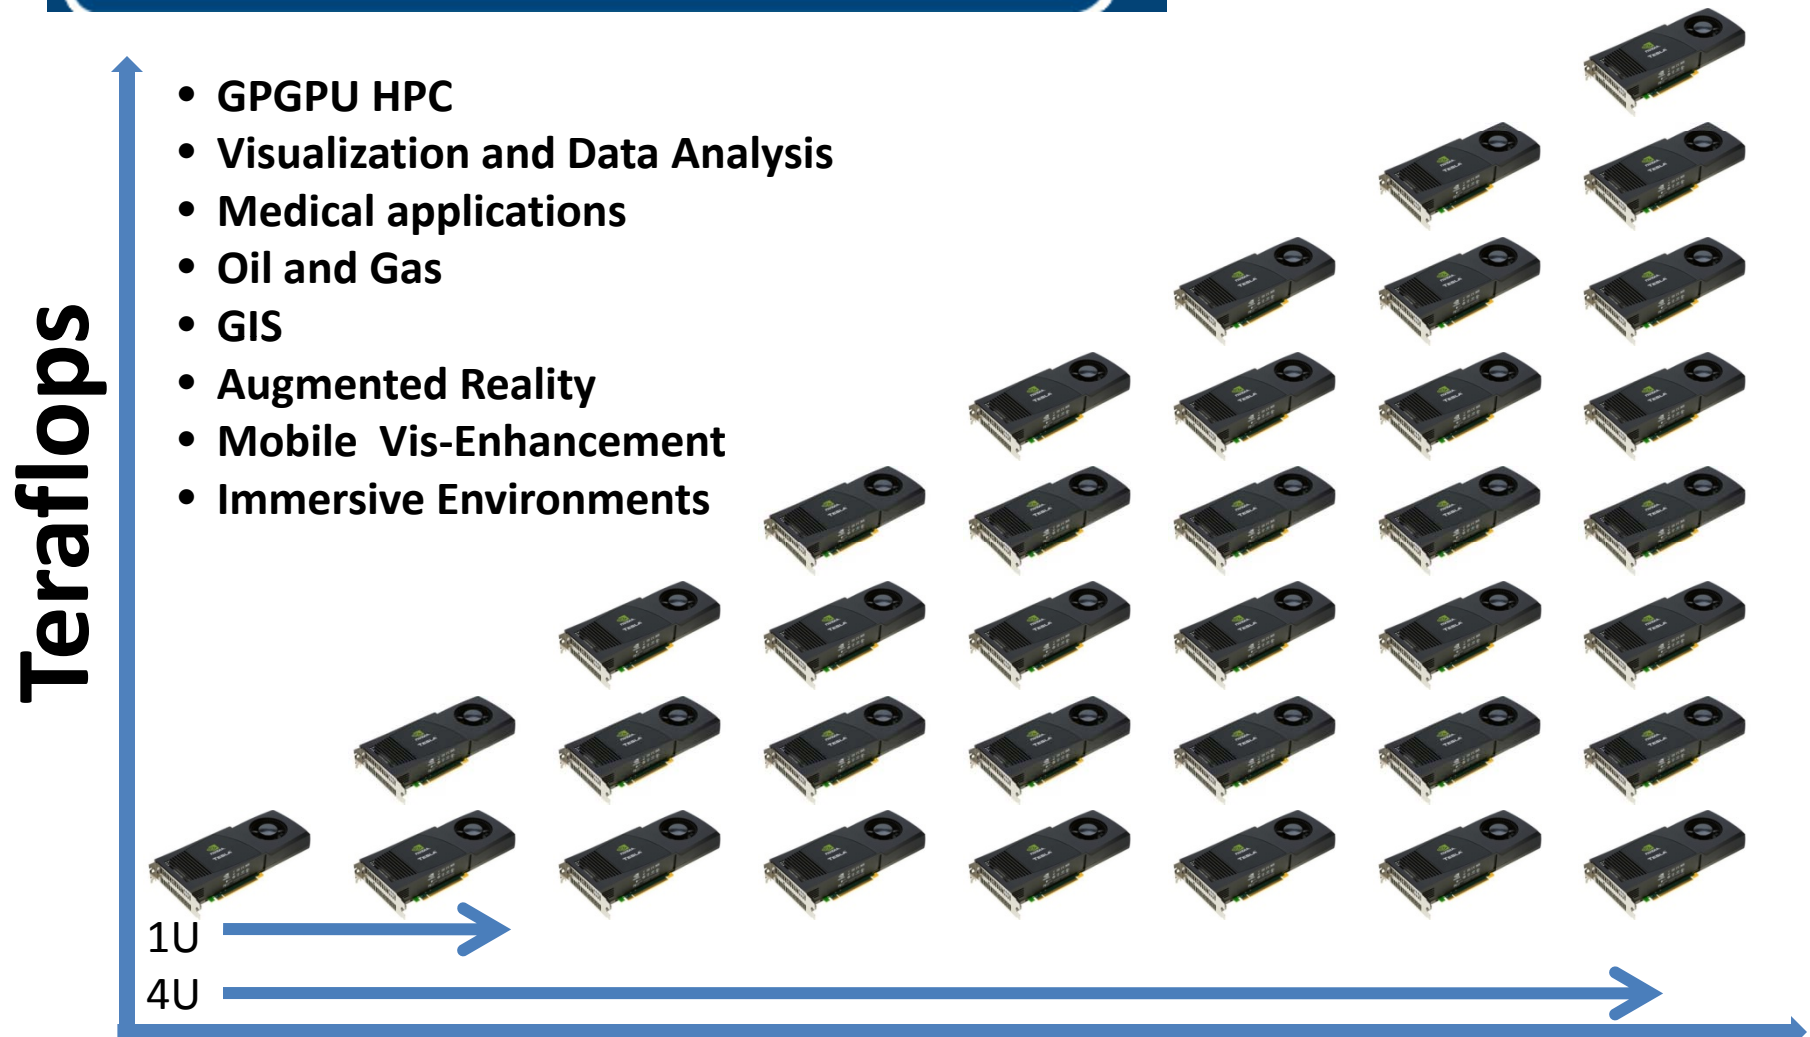
HPC Intro

High-Performance Computing Solutions For Science and Industry

- RenderStream systems use an integrated hardware-software approach tuned for peak performance.
- We research component balancing and software to determine what configuration provides the most cost effective solution for your rendering needs.
- Working with our customers our systems are fully optimized for their applications.
- We design and our OEMs build
 - All systems are assembled and burned-in in an ISO 9001 compliant manufacturing facility at 94 to 102 F for 24 hours
 - We ship worldwide (except to unfriendlyies)
 - We offer a full line of Warranties

RenderStream

Single Chassis GPGPU & VDAC




4/2/2011

Number of GPU Cards
(All combinations can be clustered w/ 1 or 2 GPU per card)



Introducing Fermi

Based on the next generation CUDA™ architecture, code named “Fermi”, the new Tesla 20-series GPU Computing Solutions offer the most advanced GPU computing architecture ever built. It features innovative technology including up to 8X double precision performance, more parallel processing capabilities, ECC implementation, and the NVIDIA GigaThread™ engine — delivering results for the most demanding HPC applications faster and more accurately.





Introducing Summary Table

How Fermi matches up to other Nvidia products

GPU Model	GPU	Single Precision (GFLOP)	Double Precision (GFLOP)	GPU Memory Bandwidth (GB/s)	Hardware ECC Capable	GPU #	GPU Cores per GPU
GTX 285	GT200b	1062	89	159	No	1	240
GTX 285	GT200b	1062	89	159	No	1	240
GTX 295	GT200b	1789	149	224	No	2	240
GTX 470	GF100	1088	136	133.9	No	1	448
GTX 480	GF100	1345	168	177.4	No	1	480
GTX 580	GF110	1581	197.6	192.4	No	1	512
C1060	G200	933	78	102.0	No	1	240
C2050/M2050	T20	1040	520	160.0	Yes	1	448
C2070/M2070	T20	1260	630	179.0	Yes	1	448

Fermi's double precision is nearly 8X faster than Tesla!



VDAC Intro

Visualization and Data Analysis Cluster Nodes

- Is a heterogeneous MPU/GPU system designed for remote, interactive visualization and data analysis.
- Supports production, compute-intensive calculations on both the MPUs and GPUs.
- The large, per-node memory supports serial and parallel visualization and analysis applications that take advantage of:
 - multiple MPU
 - multiple graphics processors.
 - large memories with RamDisk option
 - fast RAID0 storage including PCI-e storage

We offer a full line of MPU and MPU/GPU clusters today we emphasize VDAC.



VDAC Intro

Visualization & Data Analysis

Offerings

- VDAC2
 - 2 + TFLOP
 - WS, Luggable & rackmount
 - Blades
- VDAC4
 - 4+ TFLOP
 - WS & rackmount
- VDACTr8
 - 8 to 12 TFLOP
 - WS, Luggable & rackmount
 - w/PCI-e adapters and GPU expanders extends VDAC

Applications


- GPGPU HPC
- Visualization & Data Analysis
- Bioinformatics
- Medical applications
- Oil and Gas
- Geographic Information Systems
- Augmented Reality
- Mobile Vis-Enhancement
- Immersive Environments

The remaining slides show VDAC systems which tend to need more memory than other applications so keep in mind our systems are not limited to just these examples.



VDAC2
w/2 to 4 GPU
2+ TFLOP

- 2 x Intel® Xeon® X5650 2.66 GHz 6-core up to X5670 2.93 GHz 6-core
- 2 x Nvidia Quadro™ 5000
 - GTX 470, 480, and 580
 - GTX 460 Win2 and 590 dual GPU cards
 - Quadro 6000, 5000 and 4000
 - FX 5800, 4800 and 3800
 - Tesla™ C1060, C2050 and C2070
- 48GB DDR3 1333MHz ECC Registered memory (12 x 4GB)
- 150GB Velociraptor 10K RPM Hard drive Slim SATA DVD or SSD
- 40Gb/s QDR Single Port Infiniband adapter



RenderStream VDAC2 with soon to be released Nvidia Quadro FX 6000



VDACBlade2 w/2 to 4 GPU 2+ TFLOP

- 2 x Intel® Xeon® X5650 2.66 GHz 6-core up to X5670 2.93 GHz 6-core
- 2 x Nvidia Quadro™ 5000 or
 - GTX 470, 480, and 580
 - Quadro 6000, 5000 and 4000
 - FX 5800, 4800 and 3800
 - Tesla™ C1060, C2050 and C2070
- Up to 20 Tesla GPU + 20 CPU per 7U!
- Dual IOH per blade
- Up to 96GB DDR3 1333/1066 ECC DIMM
- 1 internal SATA DOM & 1 USB flash drive
- Onboard BMC for IPMI 2.0 support
 - KVM over IP, remote Virtual Media, etc.
- Dual 40Gb InfiniBand, 10Gb Ethernet, or 8Gb FCoE supported via optional mezzanine card
- Dual-port Gigabit Ethernet NIC
- Redundant GBX connectors



RenderStream VDACBlade2

7U Enclosure not shown



VDAC4
w/4-8 GPU
4+ TFLOP

- 2 x Intel Xeon X5650 2.66GHz 6-core to X5680 3.33 GHz 6-core
- 144GB DDR3 1333MHz ECC Registered memory (18 x 8GB)
- 150GB Velociraptor 10K RPM Hard drive or SSD
- SATA DVD
- 4 x Nvidia
 - GTX 470, 480 and 580
 - GTX 460 Win2 and 590 dual GPU
 - Quadro 6000, 5000 and 4000
 - FX 5800, 4800 and 3800
 - Tesla™ C1060, C2050 and C2070
- 40Gb/s QDR Single Port Infiniband adapter
- Tower or 4U rackmount w/rails

RenderStream VDAC4



Fermi Ready


Westmere Ready



GPGPU w/8 GPU
7.2 TFLOP

- Intel i7-975 3.33 GHz
- 6GB DDR3 1333MHz ECC Registered memory (3 x 2GB)
- 4X 250GB 7,200 RPM Hard drive, RAID0
- SATA DVD
- 4 x Nvidia GTX 295
- 40Gb/s QDR Single Port Infiniband adapter
- Tower or 4U rackmount w/rails (VDAC4x3 uses same chassis).

RenderStream 4XGTX295EE GPGPU




Two units mounted in a rack



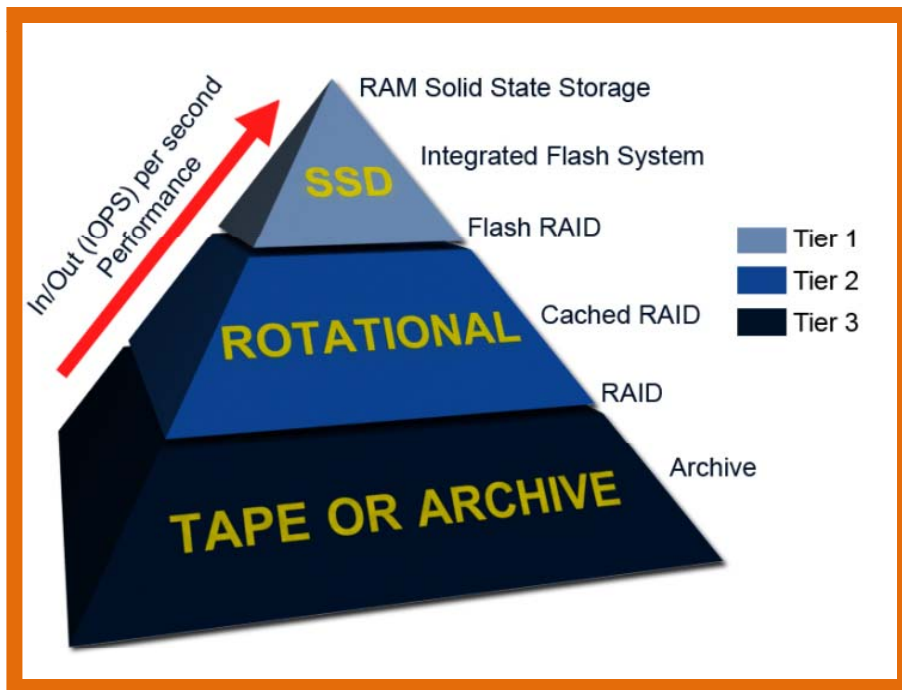
VDACTr8
w/8 GPU
8 to 12 TFLOP
RenderStream VDACTr8

- RenderStream Tr8GPU VDAC & HPC
- 2 x Intel® Westmere Xeon® X5680 3.33GHz
- 8 x Nvidia Video Cards w/All-PCI-e X16
 - GTX 470, 480 and 580
 - Quadro 6000, 5000 and 4000
 - Tesla™ C1060, C2050 and C2070
- 2X-4X Intel 510 250GB SSD
- 144GB DDR3 1333MHz ECC Registered memory (18 x 8GB)
 - 48GB Main Memory
 - 96GB RamDisk
- 3.0 kW PSU X 3





- **VDAC workstations**
- **HPC Clusters**
- **Multi-tiered storage**





VDAC Workstation

- I/O Performance balanced for MPU, graphics and storage
- 2 x Intel Westmere 6-core Xeon X5600-series
- 2 -4x Nvidia Quadro 5000 or 6000 in SLI or External Quadruplex
- 32GB+ DDR3 1333MHz ECC Registered memory for main memory and optional RAMDisk
- Boot Drive: Intel 510 250GB SSD
- Media Cache: 6x rotational HDD
 - Adaptec 32GB SSD Cache + RAID Controller

Or

- 2x PCI-e Storage in 160 GB RAID0 (replaces two GPU cards)





VDAC w/6 GPU Workstation 6 to 8 TFLOP

RenderStream VDACTr8-WS

- RenderStream VDAC6 Workstation
- 2 X Intel® Westmere Xeon® X5680 3.33GHz
- Nvidia Video -Cards w/All-PCI-e X16
 - 1X Quadro or GTX Video card (no SLI available)
 - Or external Quadruplex via PCI-e adapter
- 5 X Nvidia Cards w/All-PCI-e X16
 - 5x Tesla™ C1060, C2050, C2070
- Boot drive: 1X Intel 510 250GB SSD
- Media Cache: 2X PCI-e Storage 160 GB RAID0
- 144GB DDR3 1333MHz ECC Registered memory (18 x 8GB)
 - 48GB Main Memory
 - 96GB RamDisk (Optional)
- 3.0 kW PSU (1.0 kW X 3)






HPC Offerings

Ridiculous -Speed 2

Ludicrous-Speed 20



➤ 2 computers in 1u

➤ 20 computers in 7u

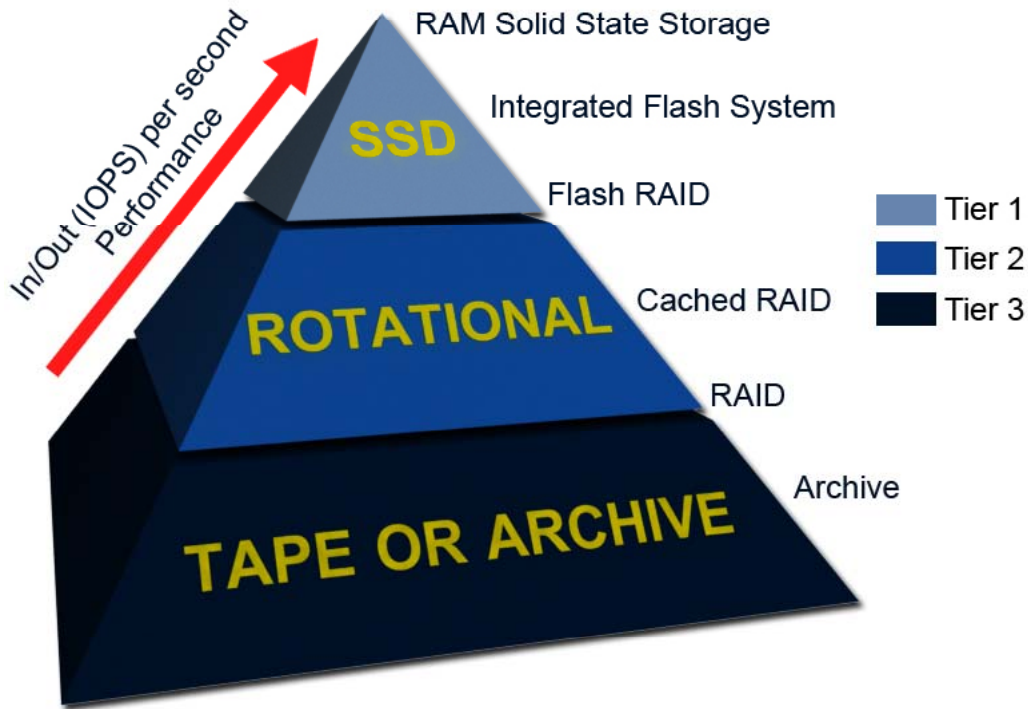
- Dual Nehalem and Westmere including 6-core
- 48 GB DDR3 ECC registered memory maximum (recommend more than 12 GB)
- RS-2 up to 4 2.5" or 2 3.5" HDD; LS-20 up to 2 2.5" HDD
- Dual NICS per computer
- Windows 7 (with downgrades available) or Linux
- On board video card
- ISO 9001 Compliant

4/2/2011

Product names are homage to Mel Brooks and Spaceballs (1987).



Multi-Tiered Storage



TMS RamSan-630



Tier 1

2TB to 10TB/ **500,000 IOPS**

RS iSCSI 3000 w/SSD Cache



Tier 2

150TB / **15,000 IOPS** and More

IBM TS3200 Tape Library Model L4U

Archive Example



Tier 3

144TB to 1,000 TB possible



HPC: VDAC Conclusion

- RenderStream offers 2 to 12 TFLOP in a single chassis and all systems are Fermi compatible.
 - Development Systems using GTX-cards
 - Production Systems using Tesla and Quadro-cards
 - Cluster nodes using Tesla and Quadro-cards
 - VDACBlade2 is now available
 - Appliances w/toll manufacturing
- Serves many communities and markets
- Offers full Warranties and On-Site service worldwide