

GPU Workbench

GPU Workbench™ is a complete platform for developing and deploying real-time applications that use NVIDIA CUDA technology. Based on the latest available GPU and CPU products, GPU Workbench systems are powered by Concurrent's RedHawk Linux operating system specially optimized for real-time CUDA performance. GPU Workbench systems can be configured with up to 96 Teraflops of double-precision compute power.

GPU technology is widely used by scientists and engineers in applications that require massive parallel processing. Compute intensive processes in molecular biology, cosmology, particle physics, radar and sonar data analysis, economics, medical imaging and many other disciplines can now execute in a small fraction of the time needed to run them on a CPU. In addition to providing a CUDA platform, GPU Workbench offers key features for time-critical CUDA applications that require guaranteed hard real-time response to external events.

RedHawk Linux Performance and Determinism

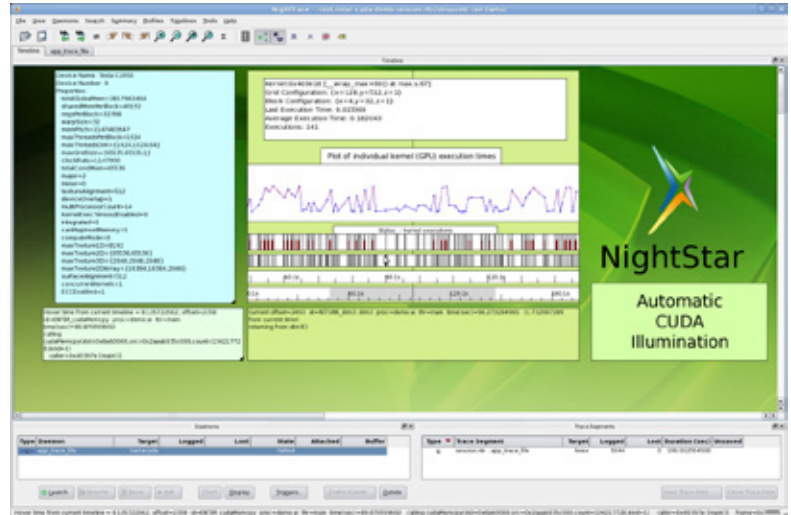
At the heart of each GPU Workbench solution is Concurrent's RedHawk Linux real-time operating system. Compatible with the popular Red Hat® Linux, CentOS and Ubuntu distributions, RedHawk provides high I/O throughput, fast response to external events, optimized interprocess communication and NUMA memory management. RedHawk is the ideal Linux environment for complex real-time CUDA applications. Proprietary GPU drivers supplied by NVIDIA frequently make demands upon kernel resources that can dramatically impact real-time performance. RedHawk addresses these special challenges and guarantees the performance of real-time processes when CUDA processes are concurrently running on a system.

RedHawk Linux, which includes the latest version of the NVIDIA CUDA SDK, reduces the process dispatch latency of real-time processes in CUDA applications from hundreds of microseconds to under 10 microseconds. RedHawk minimizes jitter and provides better overall performance.

NightStar Debugging and Analysis Tools

GPU Workbench supports Concurrent's powerful NightStar development tool kit. Users can debug, analyze, monitor, schedule and tune their real-time applications on GPU Workbench systems or remotely. Each tool runs on the target system nonintrusively, thus preserving the execution behavior of the real-time application.

NightStar's NightTrace tool includes an application illumination feature allows programmers to automatically trace CUDA API function calls and examine the values of



parameters passed and returned as well as get detailed timing analysis of CUDA kernel executions. NightTrace also permits a user to graphically view the interaction between the Linux kernel and multiple application threads in real-time. In addition, NightTrace allows users to add trace points into the CUDA kernels that are executed by the GPU.

Customized GPU Platforms

Concurrent offers fully customized real-time GPU Workbench platforms that can contain up to eight of the latest NVIDIA Tesla and graphics cards and up to four CPU sockets. Systems come in standard tower, rackmount or desktop enclosures with up to 11 integral PCIe slots and optional expansion chassis. Configurations can contain up to 24 disk drives with optional RAID. Multiple GPU Workbench systems can be interconnected via 10 Gbit Ethernet or high-speed fabrics.

Real-Time Clock & Interrupt Module

The Real-Time Clock & Interrupt Module (RCIM) is a multifunction PCIe card designed for time-critical applications that require rapid response to external events. The RCIM includes a synchronized clock readable by multiple GPU Workbench systems, eight programmable timers, and twelve input and twelve output external interrupt lines.

An optional, on-board GPS module is available to align the RCIM's synchronized clock to GPS standard time. One GPS-equipped RCIM can synchronize multiple systems in a cabled RCIM chain, or multiple systems equipped with the GPS module can operate from a common time base without any cable connections between the systems. POSIX timers based on absolute GPS time can be used to simultaneously start the execution of programs on systems which are not physically connected.

Custom Engineering from Concurrent

Concurrent is available to design and deliver customized GPU Workbench configurations for customers who require complete competitive solutions for demanding real-time applications. Concurrent engineers can provide special packaging including peripherals and enclosures, integrate third-party I/O cards, develop and integrate RedHawk Linux drivers, and perform application rehosting. Hardware and software is designed and developed to exact customer specifications.

Hardware

- 1 to 8 double or single-wide GPU cards
 - NVIDIA® Tesla™
 - NVIDIA Quadro and GeForce
- Up to 96 Teraflops of double-precision performance
- 1 to 4 x86 CPUs
- Rackmount or tower chassis
- 1 Gb and 10 Gb Ethernet
- SATA, SAS and SSD drives
- Real-Time Clock & Interrupt Module
- A wide range of industry I/O cards

Software

- RedHawk™ Linux® operating system optimized for real-time CUDA
- Latest NVIDIA CUDA SDK
- NightStar™ real-time debugging and analysis tools
 - NightView™ source-level debugger
 - NightTrace™ analyzer
 - NightSim™ periodic scheduler
 - NightProbe™ data monitor
 - NightTune™ system and application tuner

Specifications

Multicore SMP Processors

- 1 to 4 Intel® Xeon® CPUs
- Intel Core i5/i7

Memory

- Up to 1TB

I/O Busses

- Up to 12 PCIe slots
- 7/13-slot PCI and PCIe expansion chassis
- VME64 (via PCI and PCIe-to-VME bridge)

Integral I/O

- 10/100/1000BaseT Ethernet

- RS-232 serial ports
- USB ports
- SAS and SATA controllers with RAID options
- On-board video port

Real-Time Clock & Interrupt Module

- One 64-bit synchronization clock
- Eight 32-bit real-time clocks
- Twelve external input interrupt lines
- Twelve external output interrupt lines
- GPS option

Standard Peripherals

- Up to 8 TB SATA disks
- Up to 8 TB SAS disks
- Up to 3.8 TB SSD (PCIe or SATA)
- DVD-ROM/DVD-RW/Blu-Ray
- Keyboard/Mouse

Optional PCIe I/O Controllers

- Analog input and output boards
- Digital input and output boards
- AFDX® / ARINC 664
- ARINC 429
- CANbus
- FlexRay
- EtherCat
- MIL-STD-1553
- RVDT/LVTD
- Resolver
- IRIG-B
- Counter Timer
- Reflective memory
- Resistor simulator
- Serial
- Pulse generator
- Pulse width modulation
- Network I/O
- Change-of-State

Enclosures

- Mini-tower and full-tower chassis
- Rackmountable chassis (1U, 2U, 3U,4U)
- Up to 24 peripheral bays
- Rackmount cabinets (14U, 25U, 34U, 38U and 43U)

Environmental

- Operating Temperature: 10° to 35°C (50° to 95°F)
- Non-operating Temperature: -40°

to 70° C (-40° to 158°F)

- Operating Relative Humidity: 8% to 90% (non-condensing)
- Non-operating Relative Humidity: 5 to 95% (non-condensing)

Regulatory

- USA - UL listed, FCC
- Canada - CUL listed
- China - CCC Mark
- Europe/CE Mark
- EN 60950/IEC 60950-Compliant
- Germany - TUV Certified
- RoHS Compliance

Service and Support

- On-site or Return-to-Factory (RTF) maintenance
- Extended warranty
- Software support
 - Telephone advisory support
 - Product improvements
 - New releases
 - Patches to reported problems
- Other support options
 - Field installation
 - Per-call maintenance service
 - Consulting services
 - Migration assistance
 - Training at a Concurrent facility or on-site
- Custom engineering
 - Hardware/software integration
 - Device drivers
 - Customized packaging

©2017 Concurrent Real-Time, Inc. Information subject to change without notice. Concurrent Real-Time and its logo are registered trademarks of Concurrent. All Concurrent product names are trademarks or registered trademarks of Concurrent, while all other product names are trademarks or registered trademarks of their respective owners. The registered trademark of Linux is used pursuant to a sublicense from the Linux Mark Institute, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. All rights reserved.