APPRO
Supercomputer Solutions

Appro HyperGreen™ Clusters

- Power-Efficient Green Clusters
- Based on Appro GreenBlade™ System
- For Medium to Large-Scale HPC Installations
- Fully Integrated and Pre-Packaged Solution

Outstanding Performance, Flexibility, and Choice for Medium to Large-scale HPC Installations

The Appro HyperGreen™ Cluster solution is based on the Appro GreenBlade™ system solution building block offering a modular, flexible and energy-efficient cluster architecture that addresses today's natural business growth from mid to large-scale high performance and high-density computing applications.

The Appro HyperGreen Cluster provides flexibility with a choice of servers, networking, interconnects with an open commercially supported cluster management solution that can be easily configured and pre-integrated as a part of a complete package to include HPC professional services and support.

Ideal Environment

Best for high performance computing applications such as electronic design automation, aerospace and automotive engineering simulations, petroleum exploration and production, scientific visualization for oil discovery and recovery, research in seismic, health and sciences, research projects, financial modeling and risk management.

Appro Support and Services

- Closed-loop solution management, up-front consulting, pre-integration and testing with pre-sales and post-sales support.
- Support for Linux operating environments, with configurations that provide exceptional performance across a wide range of HPC applications.

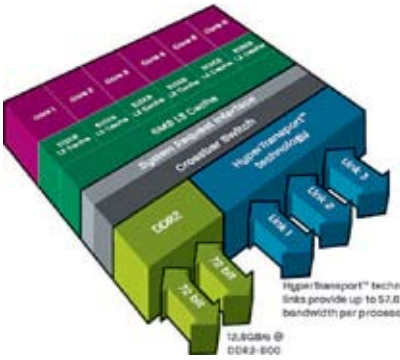

Six-Core
AMD Opteron™
Processor Design

Appro HyperGreen based on AMD Opteron™ processor delivers intelligent performance:

- **AMD Smart Fetch Technology** - Can reduce power consumption by allowing idle cores to enter a "halt" state.
- **AMD PowerCap Manager** - Allows IT datacenter managers to set a fixed limit on a server's processor power consumption.
- **AMD PowerNow!™ Technology with Independent Dynamic Core Technology** - Allows processors and cores to dynamically operate at lower power and frequencies, depending on usage and workload to help reduce TCO and to lower power consumption in the datacenter.
- **Enhanced Performance-per-watt** - 50% more compute cores vs. quad-core within the same power envelope.
- **AMD CoolCore™ Technology** - Can reduce processor energy consumption by dynamically turning off sections of the processor when inactive.
- **Dual Dynamic Power Management™** - Enables more granular power management capabilities to reduce processor energy consumption. Separate power planes for cores and memory controller.

Simplifying Cluster Integration for Faster Deployment

The Appro HyperGreen Clusters based on the GreenBlade System is simply a smarter way to build cluster solutions providing the capacity computing that HPC applications require while offering better value compared to traditional rack mounted servers.

Each GreenBlade system represents the evolution of traditional turnkey rack mounted server clusters. It consolidates server, storage, network, power and simplified management capabilities into one building block solution making it easy to test, configure and integrate as a pre-packaged turnkey solution.

Mid-to Large Data Centers Scaling up to 1000 nodes

Appro HyperGreen™ Cluster

Appro GreenBlade™ System

Refreshing ROI with a Green and Modular Architecture

- Consolidates Server, storage, network – all in a shared environment utilizing 90%+ efficient power supplies.
- Increase server count up to 60% while reducing electricity costs up to 30%
- It boasts up to 20% power consumption reduction per node compared to equally configured 1U servers, significantly reducing energy bills in the datacenter.
- Delivers hot-swappable and redundant core components such as cooling fans, power supplies and blade nodes offering superior reliability, availability and serviceability.
- Takes full advantage of the latest AMD Technologies minimizing power consumption, increasing memory capacity while improving system I/O bandwidth for high performance computing applications.

Flexibility, Performance and Efficiency to meet Customers Needs

- Clusters based on the Appro GreenBlade™ System (total of 10 server blades per 5U System) deliver scalability, flexibility and power efficiency in a dense cluster architecture
- Supports Dual/Quad-Core AMD Opteron™ Processors series offering up to 80 GreenBlades, 960 cores in a standard 42U rack configuration, doubling the density per rack compared to 1U servers
- Supports up to 1.0TB SATA with up to 64GB of DDR2 memory per compute blade while offering improved RAS with hot-swappable and redundant fans, power supplies and blade nodes.
- Supports a variety of configuration and interconnect options and choice of management options such as Server Mgmt: Standards-based IPMI or Appro BladeDome remote server management software and/or Cluster Mgmt options based on open source software solutions such as Rocks+ and MOAB from ClusterCorp.
- Fully integrated, tested and pre-packaged solution to include HPC professional services and support

Appro HyperGreen™ Clusters

(Number of Scalable Compute Nodes per Standard 42U Rack Cabinet)

	1 Rack	2 Racks	4 Racks	8 Racks	12 Racks
Number of Cores	960	1,920	3,840	7,680	11,520
Number of Processors	160	320	640	1,280	1,920
Number of Nodes	80	160	320	640	960
Max Memory Capacity	5TB	10TB	20TB	40TB	60TB

Computational Density and Flexible Interconnect Options

The Appro HyperGreen Cluster provides an extremely flexible and manageable way to create high-density computing. It is designed to scale, with little effort or re-configuration. It is based on the Appro GreenBlade System building block solution supporting up to 10 dual socket blades per subrack allowing 8 subracks in one standard rack cabinet with 2 rack unit spaces left for switches. Working through the cores-per-rack arithmetic, (8x2x10x6) yields 80 blades with 960 processing cores per 42U rack cabinet, twice the density that you can get with 1U servers.

The Appro HyperGreen cluster offers a variety of interconnect options to include only standard 1U Ethernet or 1U Infiniband (IB) switches. This is accomplished by pre-allocating space in each rack cabinet for two standard 24-port Infiniband or Ethernet switches. This offers flexibility for customers who may choose to deploy either an Ethernet or IB-based HPC cluster. Appro also offers variations in the Infiniband Interconnect with ConnectX DDR single IB, DDR dual IB or QDR single IB configurations. This provides excellent network bandwidth and low latency to many HPC workloads that require higher performance for their applications.

Better Virtualization Support in I/O Devices

As businesses deploy more and more applications in virtualized environments, and as they take advantage of live migration to save power or boost availability, the demands on virtualized I/O increase significantly. AMD Virtualization™ (AMD-V™) is a set of unique on-chip features that help AMD processor-based servers and clients run multiple operating systems and applications on a single machine by improving the efficiency of virtualization software. AMD-V™ technology allows you to better utilize your resources, which makes your client systems, servers, and datacenters more effective.

AMD-V™ technology leverages the AMD Opteron™ processor with Direct Connect Architecture to provide fast and efficient memory handling, which is a must-have for memory intensive applications like virtualization. The combined effects of Direct Connect Architecture, an integrated memory controller, HyperTransport™ technology, and AMD-V technology with Rapid Virtualization Indexing help to streamline the datacenter and maximize IT investment - reduce power consumption, support more users, more transactions, and more resource-intensive applications and achieve higher levels of efficiency and utilization in your virtual environment.

Appro offers server management options such as standards-based IPMI or Appro BladeDome remote server management and/or cluster management options based on open source software solutions such as Rocks+ and MOAB from ClusterCorp.

Open Standards and Commercially Supported Cluster Management Options

Rocks+ Flexibility and Scalability

The design specification of a Rocks+ Cluster can vary from a small 16-node system built with a local Gigabit Ethernet network to a large-scale supercomputing cluster with thousands of cores and high-end interconnect such as InfiniBand. Rocks+ can manage each type of system with equal efficiency. Rocks+ includes important additional software (i.e. Rolls – modular plug-ins to Rocks), which add a number of items to Rocks, to include: PGI Roll (Portland Group Compilers) Moab Roll (Cluster Resources), LSF Roll (Platform) TotalView Roll (Debugger), CUDA Roll NVIDIA/Tesla) Absoft Roll (Compilers) etc.

Rocks+ and MOAB Cluster Suite® = Efficiency

Rocks+ is the only licensed open source commercial solution based on Rocks+ that is designed specifically for HPC clusters. Moab Cluster Suite® is a professional cluster management solution that integrates scheduling, managing, monitoring and reporting of cluster workloads. Moab Cluster Suite simplifies and unifies management across hardware, operating system, storage, network, license and resource manager environments to increase the ROI of cluster investments. Its task oriented graphical management and flexible policy capabilities provide an intelligent management layer that guarantees service levels and speeds job processing.

Simplified Server Management Options

The Appro BladeDome Server Management software is designed exclusively for the Appro GreenBlade™ System to provide system administrators with in-depth understanding and notification of the entire system while delivering a simplified and cost-effective management solution.

Features

- Offers platform management and monitoring with status summary, and remote management capabilities
- Remotely manages up to thousands of compute nodes through support for management and sub-management servers
- Security via user ID/Password using SSH/SHA
- Utilizes system and node management monitoring through an iSCB (intelligent Shelf Controller Board) that is part of the GreenBlade system
- Web interface option

Appro HyperGreen™ Cluster, Blade System, and Node Configuration

Cluster Configuration

Processors	160 AMD Opteron™ processors
Cores	960 processing cores per rack cabinet
Memory Capacity	7.6TB max. per system
Storage Capacity	Up to 160 internal 2.5" HDDs, equal to 80TB of local storage
Networking	Gigabit or InfiniBand high speed interconnects, networking, storage, operating system, cluster management software
Rack Configuration	Standard 42U/19" rack, 2U rack space available for switches
RAS Features	Improved Reliability, Availability, and Serviceability with: • Hot-swappable blade nodes • Hot-swappable, redundant PS • Hot-swappable, reduandant cooling system
Rack Level Power	16kW to 32kW, depending on system configuration

Subrack System

Device Bays	Up to 10 blade nodes
Form Factor	5U
Power Supply	Up to four 1625W high-efficiency PSUs in N+1 configuration, Maximum power delivery - 4.707kW
Cooling	Up to 3 cooling fan units (CFU) Each CFU has redundant cooling fans
Ethernet I/O	On-board 2 port GbE LAN (RJ45)
Fibre I/O	• Optional add-on low-profile FC HCA • PCIe x8 Gen1
InfinBand I/O	• Optional one Mellanox MT25408A0-FCC-QI Dual-port DDR • Optional two Mellanox MT25408A0-FCC-QI Dual-port QDR • Optional Mellanox or QLogic add-on low-profile IB HCA • PCIe x8 Gen1
Dimensions	8.75"H x 19"W x 26"D
Weight	173.6 lbs (78.4kg) max.
Management	Support for high-speed interconnects, Appro Cluster Management, Windows or Linux OS

Appro gB122H - Blade Server

Processor	AMD Opteron™ processors 2000 series
Processor Capacity	Two; Choice of two, four or six cores per processor
Chipset	NVIDIA nForce pro3600
System Bus	2.0GT/s HyperTransport Link
Memory Type	Registered ECC DDR2 1066MHz
Memory Capacity	Up to 64GB in 16 DIMMs sockets
Disk Controller	NVIDIA nForce Professional 3600 SATA controller
Drive Bays	Two fixed 2.5" SATA HDDs
Storage Capacity	1.0TB SATA
Graphics	16MB, XGI Volari Z9 graphics controller
Network Interface	Two Intel 82571EB dual port Gigabit LAN
Input/ Output	Two USB 2.0 compliant ports, four RJ-45 LAN ports, one DB-15 video port, serial port *Optional: Single port InfiniBand (DDR) with CX4 connector*
Expansion Slots	One PCIe Gen1.0 x8 slot on riser
Weight	10.8 lbs (4.9kg) per node
Dimensions	5"Hx 1.75"W x 25"D (127 x 44.5 x 635 mm)
Temperature	Operating: 10 - 35°C, Storage: 70°C
Remote Server Mgmt	IPMI (optional)

APPRO
Supercomputer Solutions

Appro International, Inc.
446 S. Abbott Ave. | Milpitas, CA 95035, USA
1.800.927.5464 (US only) | 1.408.941.8100 Main
1.408.941.8111 Fax | **www.appro.com**

AMD Opteron™ 64

Appro International, Inc. | 446 South Abbott Ave. | Milpitas, CA 95035, USA
1.800.927.5464 (US only) | 1.408.941.8100 Main | 1.408.941.8111 Fax
info@appro.com | **www.appro.com**