nCore

# BROWNDWARF


## BrownDwarf Y-Class System

**ARM+DSP Supercomputer**

### Features

- Advanced parallel computing architecture
- High Performance, Low Latency Interconnect
- Unified Programming Model using C/C++
- Superior Power Telemetry Capability
- Precise Power and Cooling
- Reliable, Configurable and Hot-swap Capable

The nCore Y-Class supercomputer combines unprecedented low power computational performance with a low latency interconnect and a unified programming model.

This combination meets users' needs for eco-friendly energy consumption, powerful computational ability and mitigates technical risk by preserving existing investment in developed applications.

Designed to fulfill the demands of high performance applications in wide ranging fields, the Y-class system allows scientists and engineers to gain insight into complex problems while delivering real-world results.

## Brown Dwarf Y-Class

**The nCore BrownDwarf Y-Class system unifies COTS technologies, high performance SoCs, advanced low latency interconnects, and optimized software to create a supercomputer delivering exceptional performance, reliability, power telemetry, reconfigurability, and programmability at significantly reduced power levels.**

**The heart of the system is an AdvancedTCA compliant AMC card, which is a self contained high performance compute node running the Linux operating system.**

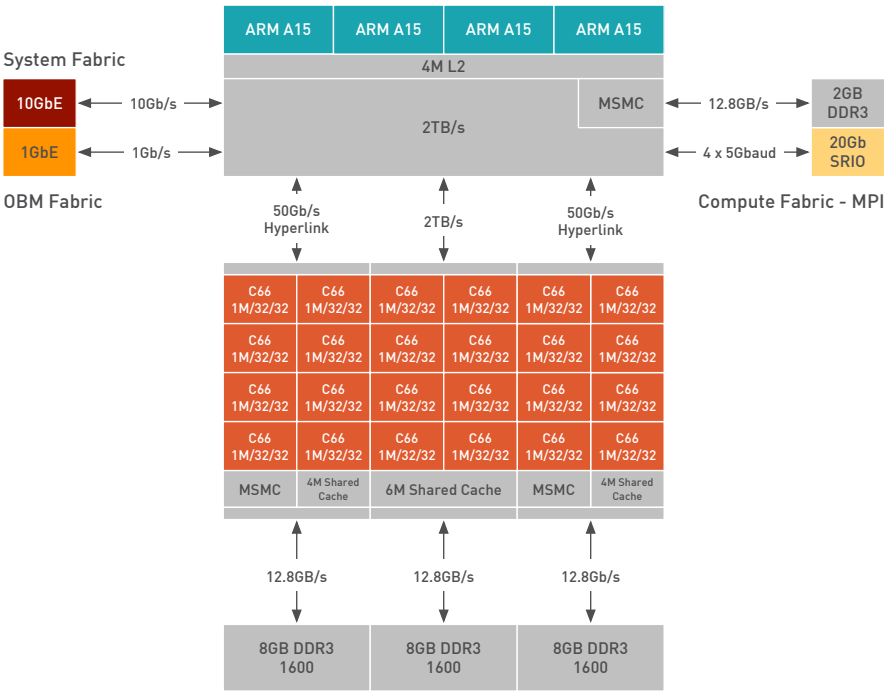


## Modular Compute Performance

Highly modular and reconfigurable, the base unit is a single Y-class chassis. Up to 3 chassis can be installed in a rack. Multirack configurations scale to tens of thousands of compute cores.

| | Chassis | Rack |
|---|---|---|
| Compute Processors (ARM/DSP) | 192/1152 | 576/3456 |
| Accelerator Performance (SP/DP) | 23.1/6.1TFLOPS | 69.5/18.2DTFLOPS |
| Aggregate Switching Capacity | 2.2Tb/s | 7.0Tb/s |
| Aggregate Memory Bandwidth | 2.4TB/s | 7.3TB/s |
| Maximum ECC Memory | 1.2TB | 3.7TB |
| Power Telemetry Measurement Points | 1632 | 4896 |
| Storage | Configurable to 45TB /Rack | |

* peak theoretical performance with 1.2Ghz A15 cores and 1.4Ghz TI C66x cores

## Compute Node

4 ARM A15 cores and 24 C66x DSP cores are organized into 48 nodes per chassis running Linux on the ARM cores. Each chassis delivers 23.1TFLOPS (SP) of ARM+DSP processing power. The ARM and DSP cores communicate via shared memory over a packet based switching fabric at 2 terabytes (TB) per second. Each node has 26GB of ECC DRAM accessed via 51.2 gigabytes (GB) per second of memory bandwidth to contain the largest problems while maximizing computational performance.

## RapidIO Interconnect

The extreme low latency RapidIO interconnect provides 20 Gigabits per second (Gbps) of non-blocking, point-to-point, bi-directional RDMA between system nodes. Each four node blade has 320Gbps of non-blocking switch bandwidth. Each chassis switch blade moves 560Gbps for a total of 2.2Tbps of total switching capability while externally presenting 14 RapidIO Quad-lanes via QSFP+ connectors for an aggregate of 280Gbps per second between chassis.

## Programming Model

The Y-class system programming model is a unified compute off-load model where the DSP cores accelerate critical algorithm components via industry standard programming APIs. This preserves the users' investment in existing applications and significantly reduces porting and tuning requirements.

## Precise Power and Cooling

Brown Dwarf leverages industry standard, redundant chassis management systems to monitor and control the power, cooling, interconnect fabric, compute nodes and system events. A separate management network provides active out-of-band communication to control all system nodes. Users will enjoy unparalleled power telemetry and sample trigger clocking capability.

# Brown Dwarf Y-Class Supercomputer Technical Data

| | |
|---|---|
| CPU | 192 ARM A15 cores + 1152 TI C66x DSP cores / chassis |
| Cache | ARM (32KB L1 I/D Cache + 4MB L2 Cache) DSP (32KB L1D/32KB L1I/1M L2 per core)<br>6MB SRAM shared memory between ARM and DSP |
| FLOPS | 23.1 SP/6.1 DP Teraflops per chassis (DSP+ARM Combined Peak) |
| SMP | ARM 4-way SMP per node |
| Main Memory | 1.24TB per chassis - 26GB ECC DDR3 memory per node (ARM 2GB/DSP 24GB) |
| Memory Bandwidth | 51.2GB/s per node |
| Interconnect | 1 x 20Gbps bi-directional, non-blocking RapidIO compute fabric connection per node |
| | 1 x 1GbE OBM connection per node |
| | 2.2Tbps fully non-blocking, point-to-point RDMA RapidIO chassis switch fabric |
| | 280Gbps external interchassis links (14 x QSFP+ interconnect cables) |
| Node External I/O | 4 x 5GBaud (total 16Gbps) lanes of Bi-Directional RapidIO |
| | 10GbE System Fabric (1GbE System Fabric on Gen 1 YC-NODE) |
| | 1GbE OBM Network |
| Application Acceleration | ARM execution with compute off-load to DSP cores via OpenMP Accelerator Model |
| System Administration | Graphical and command line tools for fault management, configuration, software and firmware updates, telemetry and provisioning. |
| | Compute fabric (RapidIO) topology is configurable via software |
| | System can be partitioned into several logical computers |
| | All system components are completely hot-swappable |
| | Workload management via 3rd party tools |
| | Automatic RapidIO endpoint discovery and enumeration |
| | Parallel cluster booting capability |
| Reliability and Telemetry | 1632 power (VDC/ADC) sampling measurement points on each chassis, triggered by an advanced clocking network |
| | Precise power and cooling control through active chassis management. |
| | Detection of out-of-specification FRUs (CPUs, fans, power supplies, interconnects) |
| | Remote firmware and software update of all systems nodes |
| | Shelf Managers, Intelligent Platform Management Controllers<br>Independent, redundant chassis management bus (IPMB) |
| Operating System | nCore optimized Linux |
| File System | EXT2/3, NFS3/4, iSCSI |
| Parallel Processing | OpenMPI |
| Shared Memory Access | OpenMP 3.0 with Accelerator Model |
| Compilers | ARM gcc and C66x C/C++ optimizing compilers |
| Power | Redundant front accessible, self cooled, hot-swappable, integral 1600W power supplies. |
| Dimensions | 12.75" (13U) H x 17.637" (19" rack mount) W x 385.8mm (15.18") Depth- not including handles & cable holders. |

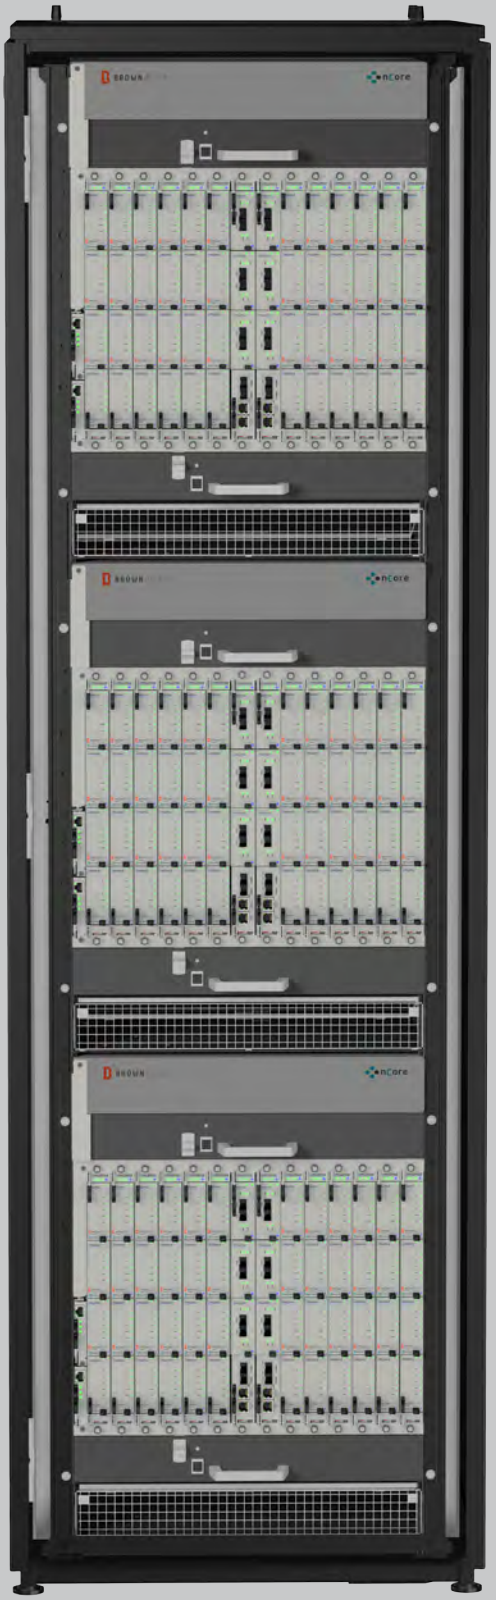

nCore

**nCore HPC**
**info@ncorehpc.com**

1270 Petaluma Blvd North
Suite C
Petaluma, CA 94952
T +1.877.666.2236