



LINUX
NETWORK

LS-1 LINUX SUPERSYSTEMS



The LS-1 is a family of Linux Supersystems that blends application-optimized performance with a systems approach to deliver ultimate supercomputing value for mid-range supercomputing applications. LS-1 Supersystems are delivered as production-ready systems, eliminating the risk, complexity and cost traditionally associated with first generation cluster deployment, configuration and optimization. Customers focus on delivering computational solutions to their research and design challenges – not on supercomputing cluster setup and operations.

LS-1 Supersystems are specifically designed to provide:

- **Application-optimized supercomputing** - eliminating the user resource and effort required for system design, configuration and tuning to meet diverse application needs. Each LS-1 Supersystem configuration is tuned by Linux Network and the application vendor to assure optimized Performance on Power-up.
- **Performance on Power-up** - accelerating initial deployment and powering users to focus on computational solutions, not supercomputing cluster infrastructure.
- **Ownership value** – LS-1 Supersystems are available in standard configurations featuring proven, best-of-breed hardware and software components. The LS-1 streamlines the efforts and costs associated with ongoing system scale, operations and maintenance, dramatically improving total cost of ownership.
- **Modular, flexible subsystems** – accommodating the needs of a variety of third party and custom applications.

Application Optimized Supersystems

Linux Networkx delivers LS-1 Supersystems as application-optimized systems, eliminating excessive user time and effort spent in designing, building and testing first generation clusters. LS-1 Supersystems are created from a specialized suite of subsystem components, each optimized for specific application requirements. Linux Networkx works with our customers to identify their specific application, environmental and budgetary requirements. We then select the optimum configuration of subsystem modules to deliver an application-optimized Supersystem that meets or exceeds customer expectations - on delivery.



LS-1 Supersystem

Thanks to a modular subsystem approach, LS-1 Supersystems are easily configured to deliver the application optimized-performance required by specific user application environments. The LS-1 also features a software tuning mechanism that adapts the Supersystem to optimize performance for custom applications, enabling

supercomputing users to take advantage of LS-1 systems value for their own in-house developed applications.

A range of pre-tested subsystem configurations are available – enabling rapid initial configuration and straightforward scalability and extension. A variety of subsystem modules are available in the areas of compute, interconnect, visualization, data storage and acceleration. These subsystems can be mixed and matched to deliver the optimum performance system depending on specific application requirements. LS-1 Supersystem configurations scale to meet a variety of application requirements – from personal Supersystems to large scale supercomputing centers. Available subsystems options include:

Compute – Compute modules combine processors and memory into pre-optimized configurations. A variety of configurations deliver compute capabilities to meet a range of application and customer requirements, including:

- Ultimate performance modules provide the highest flops density for compute-intensive applications such as CFD, Crash and visualization
- Energy efficient modules deliver dense supercomputing power with dramatically reduced cooling and power consumption – without the legacy requirement for special cooling or raised floor environments.
- Value compute modules deliver the optimal balance of computational capabilities and economic return

Interconnect – LS-1 Supersystems support industry standard interconnects including GigE, Infiniband and Myrinet. LS-1 interconnect configurations are tested and optimized as part of the overall system configuration to meet the specific application workload and size requirements, ahead of customer delivery.

Visualization – Visualization processing modules deliver the optimum technology for visualizing volumetric data, structure details or the ultimate in interactive visual reality. Options are available to support local or distributed visualization.

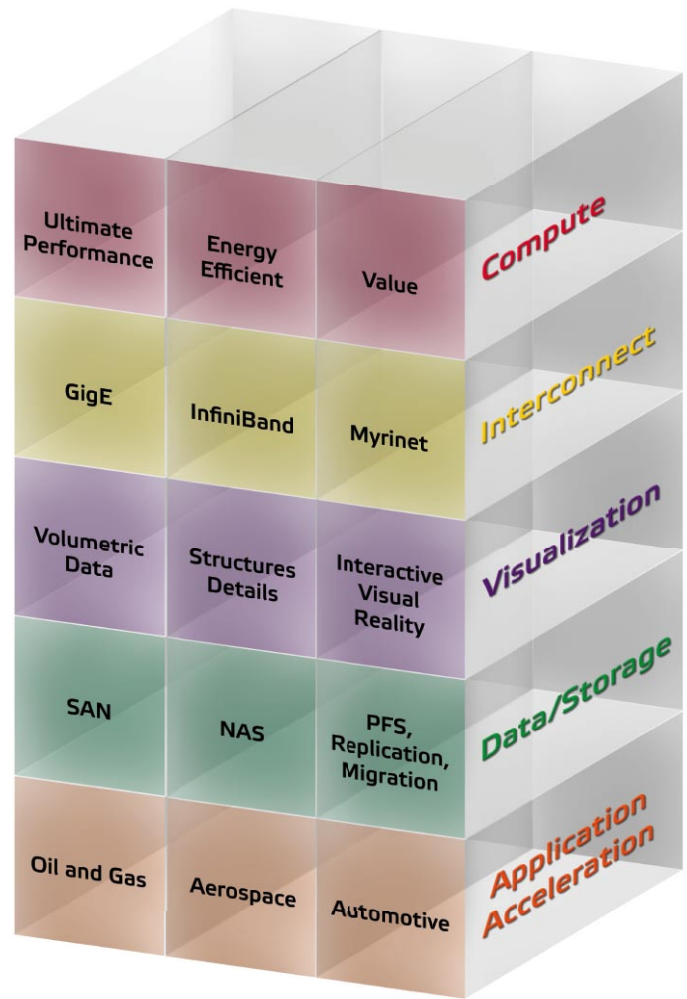
Data – Data storage subsystems include best-of-breed hardware and software components that meet the

specific needs of diverse applications and customers – from highest performance/high availability storage to mid-range storage to high capacity archival. Hardware support includes high performance SAN and NAS, with integrated software solutions for high performance parallel file systems as well as data migration and replication for resilient and effective data management.

Application Accelerators – Application Accelerators will be made available in 2006 to enhance application-specific processing for high-value, time critical computations. Accelerators are designed to enhance performance for applications codes in areas such as oil and gas exploration, aerospace and automotive design.

Systems Ownership Value

LS-1 customers enjoy a reliable, high value systems ownership experience – across the lifespan of the LS-1 Supersystem. Streamlined system delivery, featuring standardized hardware and software, takes the guesswork out of supercomputing configuration and upgrades – simplifying ongoing operations. When compared to first generation clusters, the impact is dramatic.



Available Subsystem Options

	LS-1 Supersystem	First Generation Clusters
Time to Design Build and Test	< 2 Weeks	Months
Delivery to Production Time	Same Day	2-8 weeks
Application Optimized on Delivery	Yes	User optimized - on site
Performance on Power-up	Yes	Weeks to production
Operate in lab environment	Yes	No
Installed without onsite engineering	Yes	No
Time to Scale/Upgrade	Same Day	2-8 weeks
Seamless Software Upgrade	Yes	No
Modular Hardware Scale	Yes	No
System-Wide Optimized Software Stack	Yes	User Optimized - on site
Pre-tested software upgrades & patches	Yes	User Tested - on site
Regression test ahead of software delivery	Yes	User Tested – on site
Simple, Efficient System Operations		
Integrated Rack and Node Beacon	Yes	No
Thermal Efficient Vertical Node Assembly	Yes	No
Advanced Cooling & Noise Management	Yes	No

Performance on Power-up

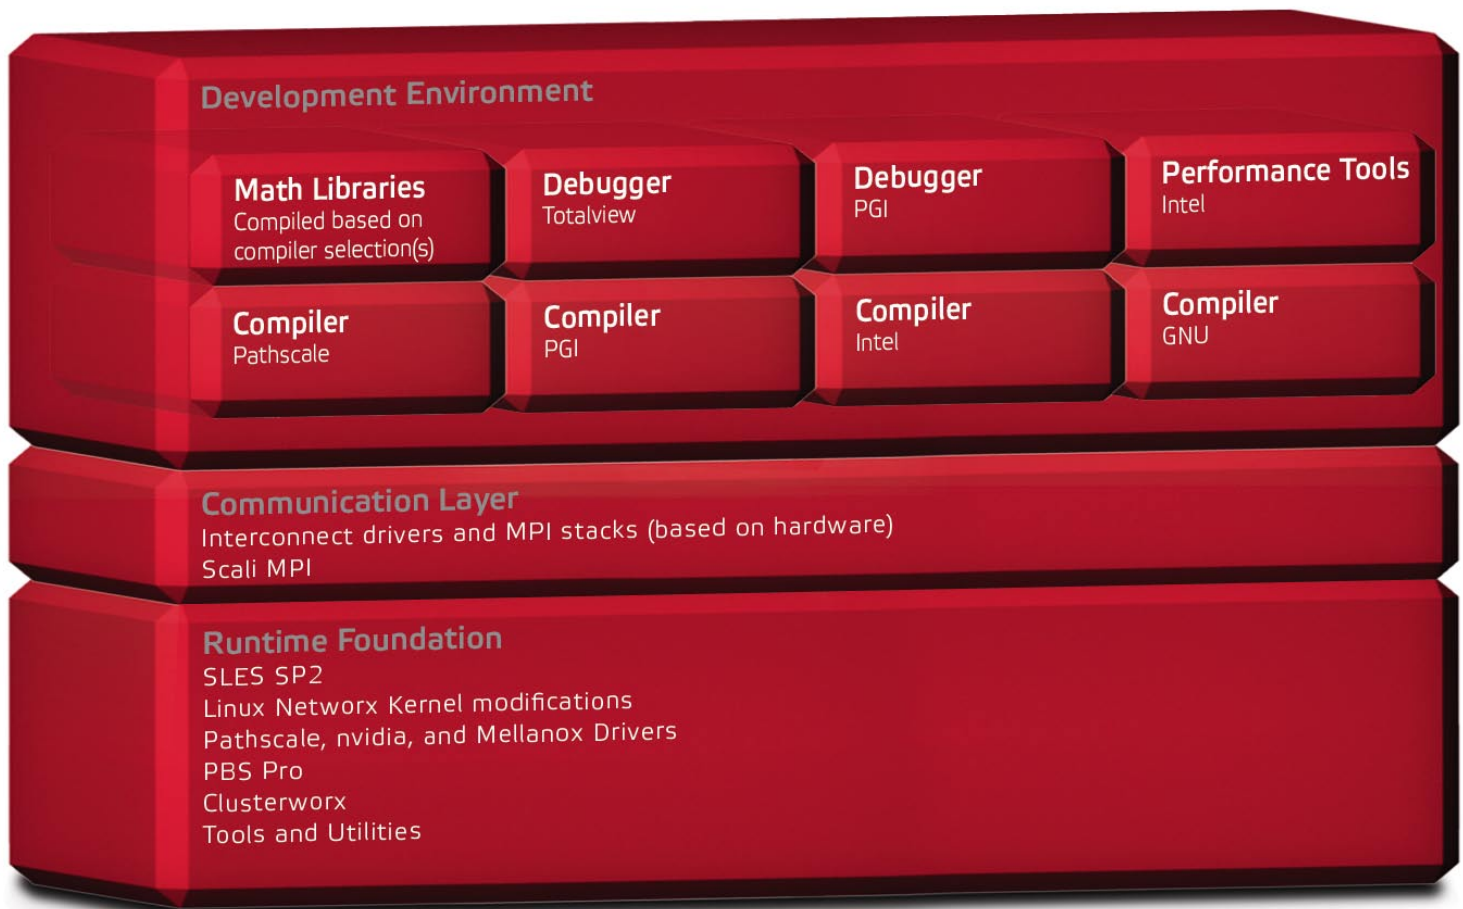
LS-1 Supersystems are optimized for maximum performance before delivery. There is no need for excessive system tuning and configuration by users. Linux Networkx tests and validates LS-1 Supersystems based on specific user and application requirements. Users enjoy Performance at Power-up with assured optimal performance. No special facilities or cluster experience is required – simply plug in the LS-1 Supersystem, power it up and begin loading datasets.

Modular Approach Delivers Flexible Scale

Initial Supersystem configurations are easily scaled by adding or substituting new subsystem modules. Simply select, install and configure the appropriate modules, and the LS-1 scales to meet expanding application requirements.

Standardized Software Simplifies Ongoing Operations

A standard production software environment comprised of proven, best-of-breed software components enhances flexibility while simplifying operations. All LS-1 configurations share this proven software foundation - including a tuned operational environment specific to applications. A robust set of development environment software options is also available for user-developed applications. The foundation and optional software components are fully validated, tested, and tuned at a system level on each LS-1 Supersystem prior to delivery. Ongoing validation and testing of new software releases ensures continuous support and seamless software upgrades within the production environment.



LS-1 Proven Best-of-Breed Software Environment

LS-1 Proven Best-of-Breed Software Environment

Each software component and option represents best-in-class technology selected based on performance, reliability, support, and interoperability. Standard software includes three component layers including:

Proven Software Foundation:

- SUSE Linux Enterprise Server distribution – High performance, reliable, secure and supported Linux distribution stack
- Industry standard Portable Batch System - Providing robust job scheduling and resource management
- Clusterworx management utilities - Enabling system-wide rapid provisioning, monitoring and administration
- Cluster tools - Selected to ensure ease of administration and high productivity

Communication Layer:

- Interconnect Drivers and MPI stack - Selected and optimized for specific interconnects
- Scali MPI Connect – An optional high performance, interconnect independent MPI layer that provides both a single MPI implementation across interconnects but also increased performance over most other MPI implementations

Development Environment:

- HPC Math Libraries – Provides developers the latest and greatest tuned math libraries to increase the performance of the most computationally intensive applications
- Compiler Options:
 - GCC – The GNU compiler suite providing a solid, simple basis for application development
 - Pathscale Compilers – Optimized for the Opteron processor, the Pathscale C, C++, and Fortran 77/9x compilers provide highest levels of performance while maintaining compatibility with the GNU environment
 - PGI Compilers – Long standing standard compiler suite used throughout the HPC industry offering outstanding reliability, support, and compatibility with multiple architectures

- Debuggers:
 - Etnus TotalView – A graphical debugger designed specifically for parallel programming that provides development engineers insight into their MPI and Open MP codes
 - PGI – Able to support both shared memory and distributed memory heterogeneous machines, the Portland Group debugger allows the ability to separately debug and control OpenMP threads, pthreads, and MPI processes
- Performance Tools:
 - Intel's Trace Collector and Trace Analyzer – Formally Vampir, Intel's performance tools provide developers a means to graphically analyze their distributed parallel codes during run-time giving the insight into how to modify codes for optimal performance

LS-1 LINUX SUPERSYSTEMS SPECIFICATIONS

CPU	16-256 64-bit AMD Opteron 200 series single or dual core processors. Scalable to 128 nodes/ 512 cores
Cache	64K L1 instruction cache, 64K L1 data cache, 1 MB L2 cache per core
FLOPS (Per Node)	Up to 19.2 GFLOPS theoretical peak performance [Opteron 280 Dual Core]
SMP	2 socket single and dual core 4 socket single and dual core
Main Memory	2-8 GB PC3200 (DDR400) Registered ECC SDRAM per chassis or 32 GB PC2700 (DDR333) Registered ECC SDRAM per chassis (1- 8 GB per socket)
Memory Bandwidth	Up to 12.8 GB/s
Interconnect	Gigabit Ethernet: 800Mb/s bandwidth, <40µs latency Myrinet: 256 MB/s + 256 MB/s bandwidth, <6µs latency Infiniband: 1 GB/s full duplex throughput, 4µs latency
External I/O Options	Vertical 1.6U - 4 PCI Low Profile Slots Vertical .8U - 1 PCI Low Profile Slot
PCI-X	<ul style="list-style-type: none"> • 2 PCI-X 64bit/133/100MHz slots on one bus • 2 PCI-X 64bit/66/33MHz slots on one bus • 2 PCI 32bit/33MHz slots on one bus
PCI-E	<ul style="list-style-type: none"> • 2 PCI-Express (x16) slots • 2 PCI-Express (x4 can support x8 cards) slots • 3 PCI (32bit/33MHz 5V) slots on one bus
Local Disk	.8U 3.5 inch Serial ATA drives (400GB 7200 RPM or 250 GB 7200 RPM)
Reliability	Independent 100 Mb/s management fabric within and between chassis
	Temperature monitoring and fan regulation maintains thermal control
	Proactive detection of impending failures triggers automatic isolation of affected components

System Administration	Clusterworx 3.3 for system installation, monitoring and management. <ul style="list-style-type: none"> • Version-controlled image management, high-speed multicast provisioning and total system monitoring • Partitioning of system into multiple logical computers • Command and control administers entire partition as a single system • Transaction processing assures configuration consistency • Workload management • Automatic interconnect topology verification • Automatic configuration of L2 and IP networking • Automatic response to component failures: isolation of hard failures, re-initialization on soft failures, redundant component switching., automatic job re-start
Operating System	SLES9 Service Pack 2 foundation with HPC-optimized Linux additions
File Systems	Support for NFS , Lustre, GFS, GPFS
External Storage	Fibre Channel disk controllers and drive enclosures
Parallel Processing Shared Memory Access	LAM, Scali MPI, OMPI 1.2 MPI-IO, MPICH, MPICH-GigE, MPICH2, PVM, XPVM, MPICH-GM, MPICH-MX, Mellanox IB Gold, MVA PICH, OpenMP,
Compilers	GCC, PGI 6.0, PathScale 2.2.1, Intel 9.0
Power	500 Watt peak per node
Dimensions	LS-1 vertical 1U node: 1.75" W x 13.25" H x 26" D. LS-1 vertical 2U node: 2.5" W x 13.25" H x 26" D. LS-1 Sub Rack: 19" H x 14" W x 26" D.