High-Performance Supercomputers for Scientific Research Use Fast RAID-5 SATA Disk Storage and Data Transfer

High-performance supercomputing clusters, sometimes called Beowulf Computers after a continuing parallel computing project at NASA, are the dominant design for highly computational scientific projects. Such research generates phenomenal volumes of data, which must be sent periodically to high-speed, nonvolatile disc caches. AMCC 3ware Serial ATA RAID controllers are a key enabling component of many cluster designs, making high performance multi-Terabyte caches both practical and economical.

Scientific research centers such as the University of Tokyo's Earth Simulator, the CERN advanced physics labs in Switzerland, and the biological/genealogical simulation facilities of Pacific Northwest National Labs have changed the face of supercomputing. These facilities have migrated to massively parallel computing strategies that typically use multiple high-performance Linux-based workstations linked in clusters by a distributed operating system layer.

3ware high-performance Serial ATA (SATA) controllers provide the speed, capacity and reliability needed to "capture" sequential stages of large computing projects into safe, non-volatile storage on high-capacity RAID-5 SATA disk cache systems. Such store-and-forward scratch spaces save hours of expensive recomputation when a node goes down. Off-the-shelf SATA-based caches create enormous savings compared to equivalent SCSI storage.

Market Background

Today's most powerful scientific computing systems consist of High Performance Clusters running Linux-based distributed operating systems to manage massively parallel manipulation of unimaginably huge volumes of data. Managing such clusters is sometimes called "herding penguins," in a nod to the Linux penguin logo.

These projects are not like the huge data streams of national stock markets, where real-time quotes make high throughput speed and five-nines data availability paramount. Rather, they are undertaken at scientific research centers where huge volumes of data are massaged to extend our understanding of physical processes and sharpen our approach to problems from mapping the human genome to teasing out details of how the universe works. It is work done at the very frontiers of human knowledge.

No one can say yet where the newest projects will take us, but AMCC 3ware SATA technology has become a key part of making ever more powerful computational clusters practical, feasible and economical.

Market Challenges

The distance from the frontiers of research to practical applications has shrunk incredibly since DNA was first conceptualized as a double helix more than fifty years ago. It was a major scientific accomplishment in its time, but now the human genome has been mapped and our understanding of how its structures interact expands daily. A similar explosion of knowledge is occurring everywhere in science.

Cluster Computers and MTBF

Cluster computers range from eight-node systems to massively networked environments of 1,000 or more nodes. A representative current scientific research environment will deploy up to 1024 nodes in a cluster, each with dual Xeon or Opteron processors running a Linux kernel interconnected to other nodes by cluster management software.

Failure of a single node during highly computational tasks can halt the entire process. The larger a project, the more nodes it requires — and a cluster's MTBF (mean time between failures) shortens as the number of nodes increases.

When a cluster node crashes, researchers must bring a new node online quickly and get the processing restarted in minutes, not in hours. To that end, they store data-state "snapshots" regularly in a store-and-forward cache or "scratch space." Such a cache must provide very fast disk reads and writes: the faster the snapshot, the less disruption, and the faster a halted computation comes back on line, the more effectively the project utilizes the cluster's expensive time and computational horsepower.

AMCC STORAGE

Range of Applications

3ware high-performance SATA controllers turn inexpensive SATA drives into fast, affordable caches for cluster-processed data in any size scientific project. A SATA strategy is especially cost-efficient when the body of data is huge, requiring many nodes operating in parallel. A RAID-based SATA scratch space allows dumping of data to disk at regular, scheduled stop points, typically every hour or so. Large projects become feasible without expensive, frequent reprocessing or purchase of expensive SCSI hardware.

A single 3ware SATA controller per node allows clusters to be scaled efficiently to any size project. A cluster node typically relies on a single eight-port 3ware controller to manage eight fast 250 GB ATA drives. Such a cache stores data right up to the 2 TB dataspace design limit set by the Linux 2.4 kernel.

Representative current scientific cluster applications include:

- **Genome Codes:**
  Clusters are used to simulate DNA structures with each computer acting as a piece of the mapped genome grid, to determine such unknowns as protein interactions and effects of structure modifications.
- **Pharmaceuticals:**
  Today's rapid, accurately tailored therapeutic drug designs depend on massive detailed simulations of physiological and chemical reactions.
- **Physics and Math:**
  Physicists and mathematicians now compute the answers to problems considered unsolvable only a few years ago because of the massive computing power they require.
- **Particle Interactions:**
  Researchers use cluster computers to map probabilistic quantum events such as subatomic particle interactions, at such research centers as CERN in Switzerland and the Pittsburgh Supercomputing Center in Pennsylvania.
- **Energy and Defense Labs:**
  Defense and Department of Energy facilities like Sandia and Lawrence Livermore use cluster computers to simulate the results of nuclear explosions, fuel oil storage explosions and other such events.
- **Satellite Data:**
  Satellites typically stay in a data download "window" for only about seven minutes, during which time massive amounts of data must be collected and reliably stored for later massaging and interpretation.
- **Oil and Gas Mapping:**
  Oil companies rely on cluster computers to interpret geophysical data accurately and to reduce the huge costs of physical exploratory drilling.
- **Atmospheric Mapping:**
  Clusters are used to simulate chaotic systems such as weather cells and storms, as well as to understand the likely effects of recent phenomena such as the hole in the ozone layer over the southern hemisphere.

The AMCC Storage Solution: Why 3ware Controllers Make the Difference

AMCC-supported cluster nodes permit fast cache storage of huge amounts of data, minimizing the risk of data loss when nodes drop offline in the middle of large computational runs. Fast SATA data-writing minimizes disruption when the system is paused to refill the RAID-5 scratch space, while fast SATA disk reads allow a crashed system to be brought back online in minutes, at the point where data was last cached.

Cost can be an important factor in research projects, especially those involving highly networked nodes performing parallel computations in supercomputer clusters. A cluster's overall savings per node using RAID-5 controllers can be significant, a compelling argument for VARs bidding to implement cluster designs.

3ware controllers are ideal for environments where high levels of sustained reads and writes are required, including high performance clusters. A single card can provide more than 3 TB of storage with up to 12 ports, each port dedicated to a single disk drive. The design makes massive data storage and access fast and cost-efficient.

The AMCC solution scales flawlessly, using a point-to-point topology to eliminate bus sharing and deliver 1.5 Gb/sec to each hard drive in the array. An onboard processor offloads the drain on central CPU capacity and onboard battery backup protects important data from power failures, even those suddenly and intentionally induced.

3ware controllers utilize the company's innovative StorSwitch™ architecture, which applies network packet switching technology to large-scale data storage. The non-blocking switched fabric allows data to be simultaneously routed to multiple hard disk drives with the highest efficiency, delivering the performance and reliability customers expect from enterprise solutions, with the added benefit of affordable Serial ATA storage.

The Bottom Line

AMCC is the only SATA RAID supplier with a reliable and effective approach to high-capacity, high-performance, value-priced RAID data storage solutions.

Low-cost Serial ATA hard drives and high-performance 3ware controllers in RAID-5 storage installations can range from 500 GB to 3+ TB and above, providing storage, bandwidth and access capabilities more than adequate for the most powerful Linux cluster.

3ware-enabled Serial ATA storage answers the call for an affordable storage solution while delivering the level of performance and capacity necessary for effective high performance computing tasks.

AMCC

Sales Offices for 3ware Products:

USA:	+1-877-883-9273
	+1-408-523-1000
Europe:	+00-800-3927-3000
Asia/Pacific:	+65-6826-3381
Japan:	+81-3-5532-7234

3waresales@amcc.com
www.3ware.com
www.amcc.com