

The SeaMicro SM10000: Redundancy and Reliability Overview

Anil Rao, June 2010
anil@seamicro.com, www.seamicro.com

Overview

The SeaMicro SM10000 brings together in a single standards-based x86 system, compute, storage, switching, server management, and load balancing. SeaMicro has reconceived the volume server as a single box cluster computer built around 512, 1.6 GHz Intel® Atom™ processors, with an integrated load balancer front end. The result is a system that uses 1/4 the power and 1/4 the space of the best in class volume server. The SM10000 is 10 rack units tall and draws less than two kilowatts of power. Four SM10000s fit in a standard 19 inch, 8 KW rack, enabling 2,048 physical servers to be deployed in a single rack. It is standards based, x86, plug and play, and runs existing operating systems, applications, and management tools without the need for custom drivers or any software modifications or recompilation.

In this paper, we discuss hardware, software, and solution level reliability features of the SM10000, and compare them to those offered by today's volume servers.

Hardware Reliability

The SeaMicro SM10000 is a 10 RU standard 19" rack mount server comprised of several major hardware blocks.

- 64 hot-swappable compute cards, each of which holds eight 1.6 GHz Intel Atom processors for a total of 512 CPUs
- Up to 8 hot-swappable uplink I/O cards, each of which holds either eight 1 Gigabit Ethernet uplinks or two 10 Gigabit Ethernet uplinks, for a maximum of 64 x 1 GbE ports or 16 x 10 GbE ports
- Up to 8 hot-swappable storage cards, each of which holds 8 hot swappable SATAHDD or SSDs. The system can be configured to run with no disk, or with up to 64 disks

These blocks are tied together with a super computer-style fabric created by interconnecting 512 SeaMicro ASICs. The fabric is a kary n-cube, more commonly known as a torus, with $k=8$ (nodes per dimension) and $n=3$ (dimensions) for a total of 512 nodes. (Each CPU is paired with a SeaMicro custom ASIC and DRAM to create a "compute node"). The fabric is low latency, high bandwidth and massively fault tolerant, and looks like a 3-dimensional cubeshaped donut. The 512 Intel Atom CPUs, Ethernet I/O, and storage cards all connect to this resilient fabric.

Packets are routed over the fabric using a strict deterministic dimensional order routing (traverse completely in one-dimension before moving to another dimension) in order to prevent fabric dead-locks. The SeaMicro fabric can be programmed for traffic to traverse a secondary path in case of primary path failure. Secondary routes can use any of the six available links from each node, providing sufficient redundancy within the fabric. The SeaMicro fabric also implements packet classes and virtual channels in order to better utilize the link bandwidth and eliminate packet loops. Implementation of virtual channels and secondary routes enables a

higher level of resiliency and also eliminates the need for error-prone link-level loop prevention and redundancy protocols such as the spanning tree.

I/O and Storage Cards

Connecting to the fabric at multiple locations are the I/O and storage cards. The I/O and storage cards have a similar architecture. Both the I/O and the storage card contain two FPGAs and each FPGA connects to four different "nodes" in the fabric. In the case of a FPGA link failure, the SeaMicro fabric routing automatically forwards traffic via one of the remaining three unaffected paths.

The SM10000 can be configured with up-to 8 hot-swappable I/O cards. To protect application traffic, the Ethernet interfaces can be bundled together in a Link Aggregation Group (LAG). A LAG can be created by using interfaces on different Ethernet cards, enabling automatic link level resiliency in the case of I/O card failure. Multiple LAGs can be configured to be in a primary/secondary mode to provide for resiliency in case of a downstream switch/router failure. SeaMicro software automatically switches traffic from the primary to the secondary LAG upon detecting any failure in the primary LAG.

The SM10000 can also be configured with up to 8 hot swappable storage cards. Each storage card holds up to 8 physical disks and each physical disk can be sliced into multiple virtual disks. Individual disks or the entire storage card can be hot-swapped. The two FPGAs that reside on the storage card create the hardware path for data access from the virtual disks and implement transparent hardware RAID to protect against disk failure. In the case of a disk failure in a configured RAID group, the SeaMicro system software provides the ability to rebuild data on a newly inserted disk. In addition to RAID, multiple virtual disks can be assigned to each of the 512 CPUs providing the ability to store data on distinct disks on different storage modules to protect against failure.

Power Supplies and Fans

The SM10000 follows the data center standard of front-to-rear air flow using two fan trays. Each fan tray hosts three dual-stacked fans to pull air from the front and exhaust through the rear. Fan trays are hot swappable and can be replaced on failure. Fan redundancy ensures that if a fan fails, the system continues to operate without impact to the system.

On the back side of the SM10000 are four independent slots for the hot-swappable AC-DC power supplies. A fully-configured system requires no more than three AC-DC power supplies providing a 3+1 redundancy. All the power supplies connect to a common passive power bus that distributes power within the system. The failure of a single power supply has no impact to system availability and operation.

Software Resiliency

In addition to hardware redundancy ensuring no single point of failure, SeaMicro provides an array of redundant and resilient software features. The SeaMicro management software runs on a dedicated CPU on the I/O card (not one of the 512 Atom processors). Adding a second I/O card automatically provides control plane redundancy – active/standby.

In the SM10000, individual CPUs can be placed into pools of compute by the built-in hardware based load balancer. The stateful load balancer, in combination with the management software, keeps track of the health of the system, the CPUs, and the applications. The load balancer distributes work to multiple processors via user-configured algorithms such as round-robin, least connection, and max-connection. If one of the CPUs fails, the application healthcheck automatically isolates the failed processor and notifies the load balancer to redirect user flows to other healthy CPUs in the system. In fact, SeaMicro customers can proactively isolate a CPU from user traffic in order to perform software upgrades.

The SeaMicro management software also provides users with the ability to identify processors that are running critical applications, and provision a group of “spare” CPUs. If the CPU running a criti-

cal application experiences a failure, the management software in the SM10000 automatically assigns the virtual disks and the environmental parameters to a spare processor and restores application access to the end user.

The SM10000 also protects against the possibility that the management software or the hardware that runs the management software fails. The management software is built with two levels of resiliency. First, because of its modular nature, if a software module fails it is automatically restarted without any disruption. Second, the SM10000 fully supports redundant management cards. The management software continually monitors the health of the hardware and software on the standby management card. If the hardware or software on the primary card fails, the secondary card takes over and a new secondary card is elected with no manual intervention. In the case of either failure, SeaMicro software sends out failure notification through the management user interfaces. The system also maintains sufficient error logs for debugging and troubleshooting purposes.

Below is a comparison between a SeaMicro SM10000 and 40 dual-socket, quad-core, one disk, one RU systems, a terminal server, a load balancer, and two rack switches. For the purposes of comparison, a SeaMicro server is configured with 40 disks.

Failure Description	Traditional 1U Rack	SeaMicro Solution	Reliability Advantage
Disk (with no RAID)	The server connected to the disk fails. 1/40th compute capacity is lost.	No server failure if multiple virtual disks are created and assigned to a server.	SeaMicro
Disk (with RAID)	A traditional server needs to be configured with 80 disks to accommodate for RAID, increasing the cost and power.	Transparent RAID allows 40 disks to be configured without increasing cost or power.	SeaMicro
Fans (or Fan Trays)	The server with the failed fan is lost. 1/40th compute capacity of a rack is lost.	Redundant fan takes over. Fan tray can be replaced without bringing servers down.	SeaMicro
Power Entry Modules	The server connected to the power entry module fails. 1/40th the compute capacity is lost.	Redundant power entry module kicks in. No impact to servers.	SeaMicro
CPU	Since each 1U server is built using dual socket, quad core CPUs, 1/80th compute capacity of the rack is lost.	Only the failed CPU is isolated. 1/512th the capacity of the rack is lost.	SeaMicro
CPU Motherboard	The server using the CPU motherboard fails. 1/40th the compute capacity of a rack is lost.	The failed CPU board brings down 8 CPUs. 1/64th of the compute capacity is lost.	SeaMicro
1 Ethernet Switch	Loss of a switch results in massive loss of bandwidth (50%) between servers. 8 uplinks are lost.	Only 8 uplinks are lost. There is no impact to bandwidth between servers.	SeaMicro
2 Ethernet Switches	Complete loss of a rack of servers. It is difficult and very expensive to configure more than 2 Ethernet switches.	Can be configured with up to 8 uplink modules. Loss of 2 uplink modules has no impact to bandwidth between servers and the remaining 6 uplink modules.	SeaMicro