

RapidIO Accelerating Software Defined Virtualized Data Centers

Ravi Thummarukudy, CEO of Mobiveil, Inc.

Anyone attending the **Linley Data Center Conference 2015** Held on February 25 and 26 at the Hyatt Regency Hotel in Santa Clara, California would come away with the impression that data centers are increasingly becoming virtualized software-defined entities. This is increasingly being made possible by the widespread availability of relatively low-cost standard hardware components: powerful interchangeable compute engines, low cost solid state and hard drive storage and ubiquitous networking. Datacenter virtualization is achieved through a layer of software above virtualized hardware: computing infrastructure, storage farms, networks, and associated security and management components.

Yuval Bachar, Hardware Network Architect, at Facebook, during the afternoon **Session on February 26th Future Directions in Cloud Computing**, made the point that his company constructed its data centers using standard hardware components and has virtualized hardware infrastructure to allow rapid deployment and maintenance of additional capacity anywhere in the world. Some data centers are being constructed in regions of the world where year round low temperature saves the power that would otherwise be needed to lower the high costs of cooling these huge hardware installations.

Companies such as Facebook, Twitter, YouTube and many others are continuously generating enormous amounts of data. [EMC Corp.](#) Declares that by 2020 there will be nearly as many digital bits as there are stars in the universe. "It is doubling in size every two years, and by 2020 the digital universe – the data we create and copy annually – will reach 44 zettabytes, or 44 trillion gigabytes." According to [HP](#), "to handle the explosive amount of data being generated, an estimated 8 to 10 million new servers in the equivalent of 200 football-field-sized data centers, are needed for the cloud over the next three years. "

And this data is constantly being accessed, processed, and written and rewritten. As the speed of data movement increases, bottlenecks become glaringly apparent and just as one bottleneck is broken another one crops up to take its place. The best example of this phenomenon is the rapid deployment of solid-state storage elements to cache frequently accessed data from disk farms. While the flash storage eliminated the rotational latency of the hard drives, the interfaces between server and flash subsystem quickly became another bottleneck, just recently broken using the newly developed NVMe interface. [NVMe](#) reduces latency overhead by more than 200% percent over an interface based on SATA/SAS that flash SSDs were using previously.

Yet another bottleneck cropping up in data centers is produced by the need to perform real time analysis of data. In his informative presentation at the Linley Conference, **The Growing Diversity in Today's Data Center**, Tom Bradicich, Vice President of Server Engineering at Hewlett-Packard cited several real world examples that needed different workload optimizations. Some of the examples cited were real time data Analytics by Paypal, Extreme file transfer at 20th Century Fox studios and molecular dynamics to hydro dynamics and data analysis at Sandya National Laboratories.

Bradicich cited the Moonshot server platform as the data center architecture HP is offering to provide real-time data analysis as well as the ability to accommodate the enormous growth of information content flooding into data centers at Facebook and other social media sites. Moonshot is an example of the software-defined server built with low-power processors and SoCs from a variety of suppliers. A single Moonshot chassis contains 1,440 DSP cores, 760 ARM cores and up to 11.5TB of storage connected via a unified 5Gbs per lane [RapidIO fabric](#). Up to 1,800 HP Moonshot servers can fit in a single server rack. The RapidIO interface was mentioned as a key contributor that enabled such large scale connectivity within the moonshot platform.

As a member of the RapidIO community for many years, I was extremely pleased to hear how RapidIO enabled such large scale connectivity and thus accelerating exciting applications in the next generation data centers. **RapidIO**, the unified fabric for performance critical computing, addresses needs in data center and high performance computing, communications infrastructure, industrial automation and military and aerospace markets by offering high reliability and low latency unified fabric. RapidIO provides chip-to-chip, board-to-board and shelf to shelf peer to peer connectivity at performance levels scaling to 100s of Gigabits per second and beyond.