

Intelligent data Center

Presentation will talk about InfiniBand based programmable server fabric switches which creates a new category of data center infrastructure designed to allow virtual servers to be created on demand by interconnecting industry-standard low-cost *commodity servers*, storage, and networking resources.

Today switches are divided into network or Ethernet switches, storage or fiber channel switches and server fabric or infiniband switches. The network switch, pioneered in the 80's was fundamentally designed to interconnect *clients* with shared network resources like printers and file servers into local area networks (LANs). The storage switch, pioneered in the 90's was designed to interconnect shared *storage* resources aggregated into Storage Area Network (SAN) with its servers. A server fabric switch is designed to *interconnect* industry-standard low-cost *commodity servers, storage, and networking resources into* virtual on-demand servers. Intelligent networking and storage services are extended with intelligent server fabric switching services such as server virtualization, virtual I/O, low latency RDMA, and server clustering. InfiniBand, due to the advantages in latency, IO, and the lack of TCP overhead, meets the requirements of customer deployments in high performance computing, grid and clustered applications. Classic HPC are characterized largely by labs and universities with computational-intensive applications driven to achieve the most compute power for the lowest cost. Enterprise HPC are characterized by enterprise customers with mission critical applications that are both computational and data intensive and require enterprise-class manageability and high availability. HPC applications are deployed in both lab and university research environments as well as mainstream enterprises from financial services, to industrials and manufacturing, entertainment, pharmaceuticals and Bio-IT. Large enterprises looking to dramatically reduce costs through data center consolidation and utility computing and Service Providers offering outsourced "virtual" computing capabilities to the market on an as-needed basis are looking to build enterprise grids or utility data centers with Virtualization economics and flexibility as primary value proposition. Server fabric switches delivers a compelling return on investment by promoting resource flexibility, and dramatically reducing equipment and management costs. The price/performance attributes of InfiniBand-based solution make it an ideal interconnect. Pools of stateless servers are interconnected into a flexible server "fabric" that provides connectivity amongst the servers and with their shared network and storage I/O resources. Based on policies, the server fabric switch provides the ability to "cross-connect" the right server, storage, and network resources together at the right time in order to create "virtual servers" which gives users greater flexibility in deploying their data center infrastructure and allows them to deploy and re-deploy applications to servers on-demand in a fraction of the time it takes today.

Data center networks are undergoing a three-phase evolution which consists of Consolidation, Virtualization, and Automation. Consolidation is accomplished through integration of server fabric, LAN and SAN switches with a common set of intelligent services which results in more manageable data center network. Server virtualization software makes it possible to create "virtual" servers on-demand by interconnecting commodity server, storage, and I/O resources through a server fabric switch. Automation means integration of network intelligence with 3rd party management applications in order to automate adherence to business policies.

Server fabric switch portfolio consists of several product groups including server fabric switches, I/O Chassis, I/O Gateways, Host Channel Adapters and Virtualization Software. Server fabric switches and I/O Chassis provide the backbone infrastructure for high performance grid and utility computing. The InfiniBand Switches provide pure InfiniBand connectivity and are used to build out the server fabric in the data center. The I/O Chassis provide InfiniBand switching plus Ethernet and Fiber Channel gateways to connect the server fabric with its legacy LANs and SANs or NAS. I/O gateways provide transparent connectivity between the InfiniBand-based server fabric and the LAN, SAN, or NAS devices used by the servers. I/O Gateway modules plug into the I/O chassis products. 10Gbps Host Channel Adapters provide InfiniBand connectivity to any server with a PCIx or PCIe bus. VFrame™ server virtualization software suite makes it possible to

create “virtual” servers on-demand by interconnecting server, storage, and I/O resources through a server fabric switch.

Benefits are high performance computing at a fraction of the cost leveraging commodity server components and InfiniBand, a standardized, low-latency protocol available today, better utilization of new blade server investment for server virtualization, enterprise clustered applications, and grid/utility computing, better management of data center resources by breaking down the barriers which form silos of server, storage, and networking, ability to provision new applications in seconds, ability to failover provision, and add/remove I/O or storage bandwidth on the fly

InfiniBand is an industry-standard administered by the InfiniBand Trade Association (IBTA) founded in August 1999. InfiniBand is backed by top companies in the industry, including Dell, HP, IBM, Intel, Sun, Hitachi, NEC, and Microsoft. All major server vendors have announced support for InfiniBand technology as part of their clustering and/or grid solutions for both their standard and their blade server solutions. InfiniBand's unprecedented combination of high bandwidth and low latency make it an ideal data center switching fabric for high performance server clusters, enterprise grids, and utility data centers. Compared to connection-less Ethernet, connection-oriented InfiniBand offers higher bandwidth (10Gbps today, 30Gbps coming versus 10Gbps), lower latency (6usec versus 70 usec) and remote direct memory access (RDMA) capabilities, resulting in better server-to-server performance and dramatically reduced load on the server CPU (1-3% versus 10-50%) enabling greater server efficiency.