

Windows Server® 2012 SMB Direct Networking Acceleration: Virtualization and Database Results

Increased Performance, Scaling and Resiliency



Microsoft

The continued expansion of business-critical information and rich content within extended enterprises continues to change the storage dynamic in a wide range of industries and organizations. This market trend drives the need for higher connectivity speed and higher levels of fault tolerance. In order to address the market needs Microsoft and Mellanox worked together to provide a solution for remote storage access that allows faster and more efficient file servers for server applications such as Windows Server Hyper-V® and SQL Server®.

The solution has been implemented in Microsoft's Windows Server® 2012: The enhanced Server Message Block (SMB) protocol takes advantage of the capabilities of Mellanox's RDMA interconnect technologies in a way that provides greater performance scale and reduced CPU utilization. These RDMA-based networking technologies include Mellanox's 56Gb/s InfiniBand and 40Gb/s Ethernet with RoCE interconnect products.

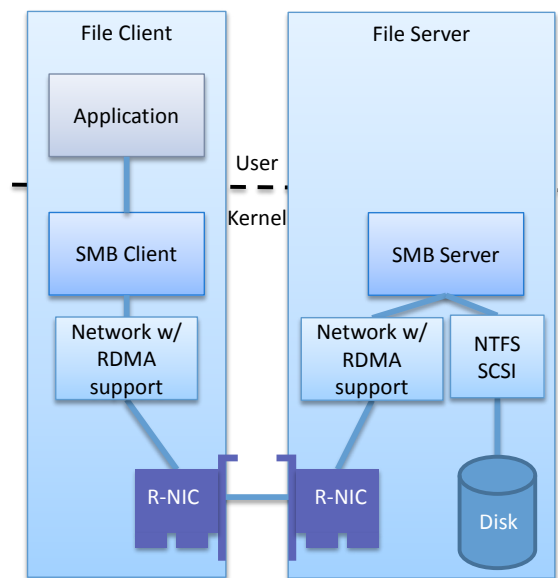


Figure 1: Basic SMB Direct (SMB over RDMA) Architecture

Microsoft Windows Server® 2012 with SMB Direct

In Windows Server® 2012, the SMB protocol for remote storage has been enhanced to allow faster and more efficient file servers for applications such as Windows Server Hyper-V and SQL Server. As part of the SMB protocol, two new features, SMB Direct and SMB Multichannel, enable customers to deploy storage for server applications on file servers that need to be cost efficient and continuously available, while delivering high performance.

SMB Direct supports the use of network adapters that have remote direct memory access (RDMA) capability. SMB Direct (SMB over RDMA) is a new storage protocol in Windows Server® 2012 that includes:

- **Increased Throughput:** Helps to optimize the throughput of high speed networks where the network adapters coordinate the transfer of large amounts of data at line speed.
- **Low Latency:** Provides extremely fast responses to network requests, and as a result, makes remote file storage respond similarly to directly attached block storage.
- **Low CPU Utilization:** Uses fewer CPU cycles when transferring data over the network, which leaves more power available to server applications.

SMB Multichannel allows file servers to use multiple network connections simultaneously and includes the following capabilities:

- **Fault Tolerance.** When using multiple network connections at the same time, the file server continues functioning despite the loss of a network connection.
- **Increased Throughput.** The file server can simultaneously transmit more data using multiple connections for high speed network adapters or multiple network adapters.

Both SMB Direct and SMB Multichannel are automatically configured by Windows Server® 2012.

Mellanox RDMA based Interconnect Solution

Mellanox's end-to-end connectivity solutions enable the highest performance and most efficient data center infrastructure and support up to 40Gb/s Ethernet with RoCE or 56Gb/s InfiniBand servers and storage interconnect. Compared to traditional hardware and software architecture that imposes a significant load on a server's CPU and memory, Mellanox products use offloads such as RDMA. This accelerates IO by allowing application software to bypass most layers of software and communicate directly with the hardware, and it enables servers to directly place information into the memory of another computer. The technology reduces latency and minimizes the CPU overhead.

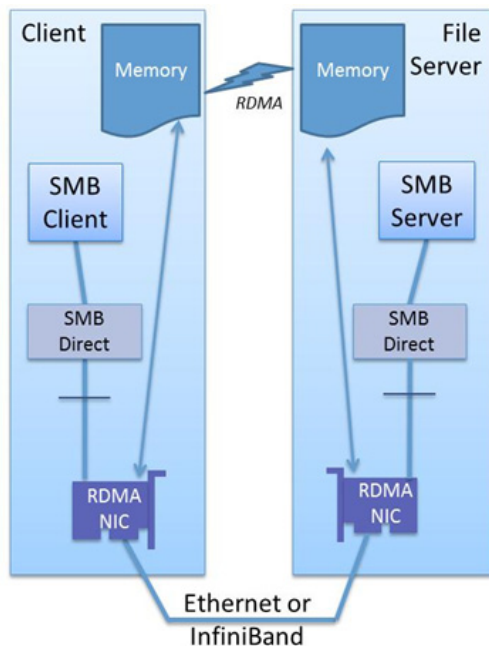


Figure 2: RDMA Operation Architecture

When an application performs an RDMA Read or Write request, no data copying is performed. The RDMA request is issued from an application running in user space to the local NIC and then carried over the network to the remote NIC without requiring any kernel involvement. Request completions might be processed either entirely in user space (by polling a user-level completion queue) or through the kernel in cases where the application needs to sleep until a completion occurs.

Solution Benefits

Using SMB Direct over Mellanox’s RDMA-enabled interconnect solutions provides a more efficient data center infrastructure, delivering higher performance at the lower cost.

With SQL Server it can become the most efficient protocol to service a variety of workload types, ranging from small, random I/O requests to large sequential I/O requests.

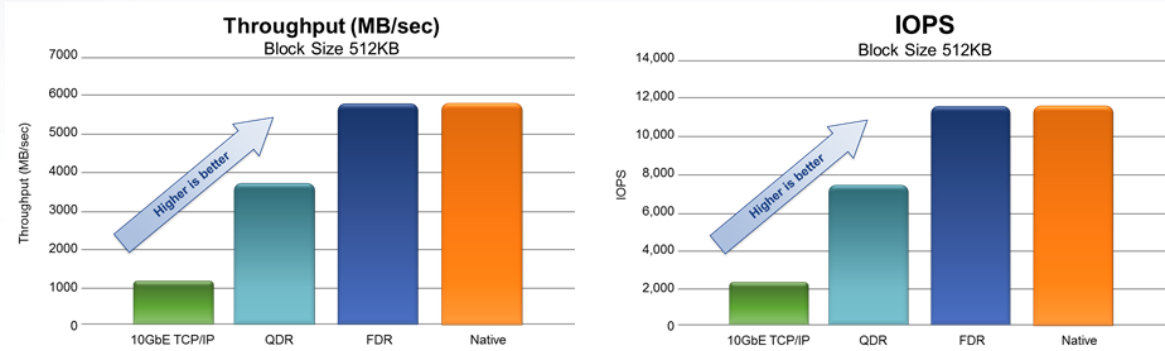


Figure 3: Benchmark results comparison between different interconnect technologies and native performance

Also within a virtualized environment using Windows Server® Hyper-V, SMB Direct can not only service a variety of workloads, but can do so simultaneously.

It also delivers the stringent requirements for storage resiliency in those environments, as well as:

- Effective scaling
- Connection Resiliency
- Higher server and network utilization

For more information on Microsoft’s and Mellanox’s joint remote storage access solution, please visit www.mellanox.com



350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
 Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com