Shared Memory Communications – Direct Memory Access (SMC-D)

Optimize Communications with Innovative Shared Memory Communications

The digital transformation is creating a marked shift in the IT world. The proliferation of mobile and other network-connected devices is driving increases in transaction workloads, data volumes and availability requirements. Vast amounts of data are generated – needing to be managed, shared and stored. These changes are driving the need for new methods of high performance communications.

For over 15 years, the mainframe has supported HiperSockets™ for high performance communications between logical partitions within the same physical machine[1]. HiperSockets transports TCP/IP traffic between the logical partitions such as OS images, TCP/IP stacks, guest virtual servers under z/VM®, and associated middleware and applications. HiperSockets eliminates the need for any physical cabling.

Another communications option from IBM Z® is a connectivity protocol Shared Memory Communications – Direct Memory Access (SMC-D) for IBM z14™ (z14), IBM z13s™ (z13s) and IBM z13® (z13), designed to provide application transparent communications within the server. SMC-D has been shown to offer up to 61% CPU savings on the z13 for file transfers 'within-the-box' versus standard TCP/IP over HiperSockets.

Best of all there are no application changes needed and the use of SMC-D is *totally transparent* to applications.

How does SMC-D work?

The SMC-D protocol is part of z/OS® 2.2. SMC-R and SMC-D can be exploited concurrently. The protocol dynamically selects the variation of SMC based on physical location of the peer hosts.

SMC-D delivers compelling improvements in network performance for 'within-the-box' communications with reduced latency, improved throughput, reduced CPU networking costs and application transparency.

Communications within the IBM Z occurs directly using memory-to-memory with Internal Shared Memory (ISM). ISM is a firmware only implementation. No additional I/O adapters, switches or cables are required. ISM devices are defined using virtual PCI device definitions. ISM allows communications to exist across unique operating system instances in separate partitions (LPARs) within the same physical hardware footprint.

While TCP/IP is still used to establish the connection and provide the security, SMC-D eliminates TCP/IP processing in the data path. The TCP/IP connection can be provided using OSA or HiperSockets connectivity.

- Once the initial handshake is complete, communications then uses SMC-D sockets-based communications. Socket application data is exchanged while the TCP connection remains active. This model preserves critical operational and network management features of TCP/IP.
- Qualities of Service Preserved
- The qualities of service known to IBM Z are preserved with SMC-D. There are minimal operational changes, dynamic discovery of partner capabilities and dynamic connection setup. SMC-D with ISM preserves the existing network administrative and operational models and supports security configurations to help organizations reduce risk.

[1] Other types of LPAR-to-LPAR communications include OSA and SMC-R using the 10GbE RoCE Express2 or 10GbE RoCE Express.

SMC-D Can Benefit Workloads

- Use cases include multi-tiered work co-located onto a single IBM Z server without requiring extra hardware such as CICS®, DB2®, Websphere®, WebSphere MQ®.
- Initial deployment is intended for z/OS to z/OS data transfers.

SMC-D Proof points

- Up to 61% CPU savings for FTP file transfers across z/OS systems versus HiperSockets*
- Up to 800% improvement in throughput with more than a 89% decrease in CPU consumption and a 89% decrease in response time for streaming workloads versus using HiperSockets*
- Up to 2001% improvement in throughput with more than a 85% decrease in CPU consumption and a 95% decrease in response time for streaming workloads versus using OSA*
- Up to 91% improvement in throughput, up to 47% reduction in networking related CPU cost and up to 48% improvement in response time for interactive workloads versus using HiperSockets*
- Up to 1601% improvement in throughput, up to 40% reduction in networking related CPU cost and up to 94% improvement in response time for interactive workloads versus using OSA*

Hardware/Software Requirements for SMC-D

- An IBM z14, IBM z13s or IBM z13
- z/OS V2.2 or later (with PTFs)
- A z/OS V2.2 Communications Server planning tool, SMC-AT, can be used to help clients gain insight into the applicability of this technology for specific z/OS application workloads.

Value

IBM SMC-D delivers faster communications with:

- Transparent application use
- Low CPU utilization and latency
- Ability to leverage existing infrastructure
- Preserves TCP/IP security, management
- Standards-based

Summary

SMC-D represents a low latency, low CPU overhead alternative for data transfers to help organizations prepare for the huge volumes of enterprise data exchange from increase in transactional workloads and data volumes. It provides the strengths of TCP/IP for secure handshaking with the benefits of application transparency plus outstanding performance.

* All performance information was determined in a controlled environment using IBM z13. Actual results may vary. Performance information is provided "AS IS" and no warranties or guarantees are expressed or implied by IBM.

ZSB03082-USEN-01