

Background

As part of z/OS[®] V2R1, the updated Communications Server includes the SMC-R protocol, otherwise known as Shared Memory Communications over RDMA (Remote Direct Memory Access). This protocol enables two z/OS LPARs to communicate through 10 GbE RoCE Express features with a considerable reduction in latency.

How it works

SMC-R can be established between two LPARs on z/OS with minimal setup. These LPARs may be co-located on a single IBM System z[®] or across systems. One or more OSA cards must be configured in the LPARs as OSD. Although only one RoCE Express feature per LPAR is required, it is recommended that two or more per LPAR are installed for redundancy.

When two LPARs with enabled RoCE Express features attempt to communicate over TCP/IP, the machines will signal that the features are available, and then set up an SMC-R link. All traffic sent through new TCP/IP sockets will then be diverted through the SMC-R link for transmission over Ethernet, without passing through the internet (IP) or transport (TCP) layers of the network (see Figure 1). The data is instead remotely passed into the memory of the other LPAR, where it can be quickly processed by the application with minimal CPU overhead. This provides a significant reduction in the effective network latency experienced by an application.

Performance Benefits

Lower network latency provided by the RoCE Express feature and SMC-R offers significant improvements in overall response time. Workloads that exhibit the greatest performance benefits are those with streaming data patterns over the network, (e.g. File Transfer Protocol, etc.) due to the savings in CPU cycles consumed by packet processing. Improvements can also be seen on a variety of OLTP workloads hosted on z/OS systems, such as WebSphere[®] Application Server, CICS[®], IMS[™] and DB2[®]. This is especially noticeable on chatty, multi-tier workloads with multiple network flows per logical transaction.

The IBM Competitive Project Office (CPO) has conducted benchmarks with the Friendly Bank application running on WebSphere Application Server Liberty Profile within a z/OS LPAR. This application communicated with a DB2 database running on a different LPAR. Communications were run through separate OSA cards over TCP/IP, and then over SMC-R. Client requests were driven from an x/Windows server connected to the WAS server over standard TCP/IP, and were not subject to SMC-R and RoCE exploitation (see Figure 2). The majority of the network traffic ran between the client and the WAS server, with about 20% flowing between WAS and DB2.

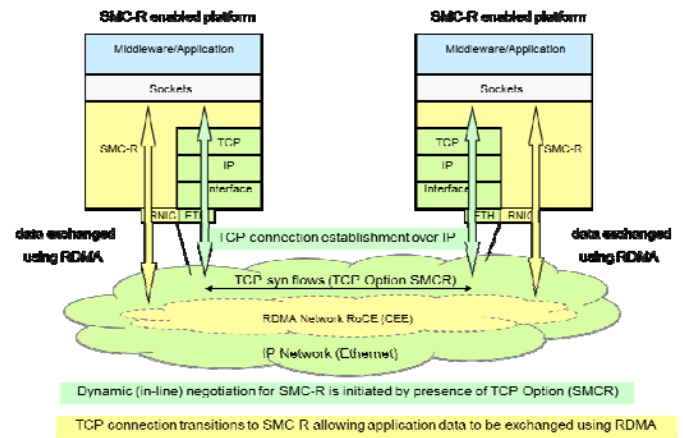


Figure 1: Overview of SMC-R establishment and communication

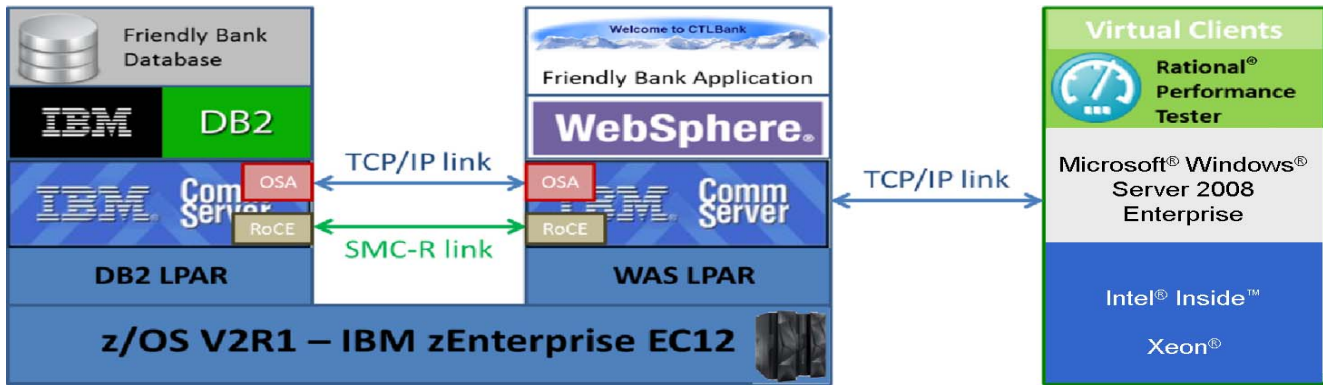


Figure 2: Setup for CPO Friendly Bank benchmark

Despite the limited use of the RoCE Express feature, benchmarks showed that SMC-R provided over a **19% reduction** in the **total response time** of a user session for a fixed workload (see Table 1). Web pages that required a higher volume of data transfer between WAS and DB2 displayed the largest reductions in response times, reaching as high as **32%**.

Table 1: Results of CPO Friendly Bank benchmark

RoCE	User Load	User Sessions per Second	Total Sessions per Second	LPAR CPU		Total Session Response Time (ms)		
				DB2	WAS	Average	Std. Dev.	Reduction
DISABLED	40	10	400	21.0%	94.0%	53.3	35.1	19.32%
ENABLED	40	10	400	22.0%	94.0%	43.0	19.4	

© Copyright IBM Corporation 2013

IBM, IBM logo, CICS, DB2, IMS, Rational, System z, WebSphere and z/OS are trademarks or registered trademarks of the International Business Machines Corporation. Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.