

IBM z/OS SMC Applicability Tool

Dave Herr (dherr@us.ibm.com)



Agenda topics

- Problem:
How to evaluate the applicability and the potential benefits of SMC-R with 10 Gbe RoCE Express in your (IBM Z customer's) environment
- SMC Applicability Tool (SMC-AT) Introduction
- SMC-AT Overview
- SMC-AT Usage and Reporting Overview
- Backup
 1. SMC-AT Configuration and Usage (starting / stopping) Details
 2. SMC-R and SMC-AT Evaluation Considerations / References
 3. SMC-R Background

Evaluating SMC-R applicability and benefits

As customers express interest in SMC-R and RoCE Express one of the initial questions asked is:

- “What benefit will SMC-R provide in my environment?”
 - Some users are well aware of significant traffic patterns that can benefit from SMC-R
 - But others are unsure of how much of their TCP traffic (in their environment) is:
 - z/OS to z/OS and
 - how much of that traffic is well suited to SMC-R
- Reviewing SMF records, using Netstat displays, Ctrace analysis and reports from various Network Management products can provide these insights...

This approach can be a time-consuming activity that requires significant expertise.

SMC Applicability Tool Introduction

A new tool called SMC Applicability Tool (SMC-AT) has been created that will help customers determine the value of SMC-R in their environment with minimal effort and minimal impact

- SMC-AT is integrated within the TCP/IP stack:
Gathers new statistics that are used to project SMC-R applicability and benefits for the current system
 - Minimal system overhead, no changes in TCP/IP network flows
 - Produces reports on potential benefits of enabling SMC-R
 - Does not require RoCE hardware or SMC-R function. No IP configuration changes are required (measures your existing TCP/IP traffic).

- Available via the service stream on existing z/OS releases:
 - V1R13: APAR PI27252 PTF UI24872
 - V2R1: APAR PI29165 PTF UI24762

SMC-AT Usage Overview

Activated by Operator command

(**Vary TCPIP,,SMCAT,dsn(smcatconfig)**) – Input dataset contains:

- Interval Duration, list of IP addresses or IP subnets of peer z/OS systems ((i.e., systems that we can use SMC-R for)
 - If subnets are used, the entire subnet must be comprised of z/OS systems that are SMC-R eligible
 - It is important that all the IP addresses used for establishing TCP connections are specified (including DVIPAs, etc.)
- At the end of the interval a summary report is generated that includes:
 1. **Percent of traffic eligible for SMC-R** (*% of TCP traffic that is eligible for SMC-R*)
 - *All traffic that matches configured IP addresses (not using IPsec or FRCA)*
 2. **Percent of traffic well suited for SMC-R** (*your eligible traffic that is also “well suited” to SMC-R, excludes workloads with very short lived TCP connections that have trivial payloads*)
 - *Includes break out of TCP traffic send sizes (i.e. how large is the payload of each send request)*
 - *Helps users quantify SMC-R benefit (reduced latency vs reduced CPU cost)*

SMC-AT Usage Overview (continued)

The Summary Report includes 2 sections based on the specified IP addresses/subnets defined in SMC-AT configuration file:

1. Total potential benefit:

All eligible¹ TCP traffic that matched the specified SMC-AT configuration (IP address/subnets). This total includes IP traffic (if present) that did not meet the direct IP route connectivity requirement.²

2. Immediate benefit:

The eligible¹ TCP traffic that can use SMC-R immediately / as is (meets SMC-R (RoCE) direct route connectivity requirements).

Detected by the tool automatically (non-routed traffic).

1. “Eligible” means TCP traffic not using IPsec or FRCA.

2. The total benefit report represents the opportunity to re-configure routing topology (single subnet) to enable SMC-R.

SMC Applicability Tool Sample Report (Direct Connections)

TCP SMC-R traffic analysis for matching direct connections

Connections meeting direct connectivity requirements

50% of connections can use SMC-R (eligible)

67% of eligible connections are well-suited for SMC-R

79% of total traffic (segments) is well-suited for SMC-R

81% of outbound traffic (segments) is well-suited for SMC-R

75% of inbound traffic (segments) is well-suited for SMC-R

Interval Details:

Total TCP Connections:	6
Total SMC-R eligible connections:	3
Total SMC-R well-suited connections:	2
Total outbound traffic (in segments)	274
SMC-R well-suited outbound traffic (in segments)	222
Total inbound traffic (in segments)	211
SMC-R well-suited inbound traffic (in segments)	159

Application send sizes used for well-suited connections:

Size	# sends	Percentage
----	-----	-----
1500 (<=1500):	1	20%
4K (>1500 and <=4k):	1	20%
8K (>4k and <= 8k):	0	0%
16K (>8k and <= 16k):	0	0%
32K (>16k and <= 32k):	0	0%
64K (>32k and <= 64k):	1	20%
256K (>64K and <= 256K):	2	40%
>256K:	0	0%

This report shows direct IP connections. Another report is produced showing all TCP connections.

How much of my TCP workload can benefit from SMC-R?

What kind of CPU savings can I expect from SMC-R (message size)?

End of report

Backup

Backup Topics:

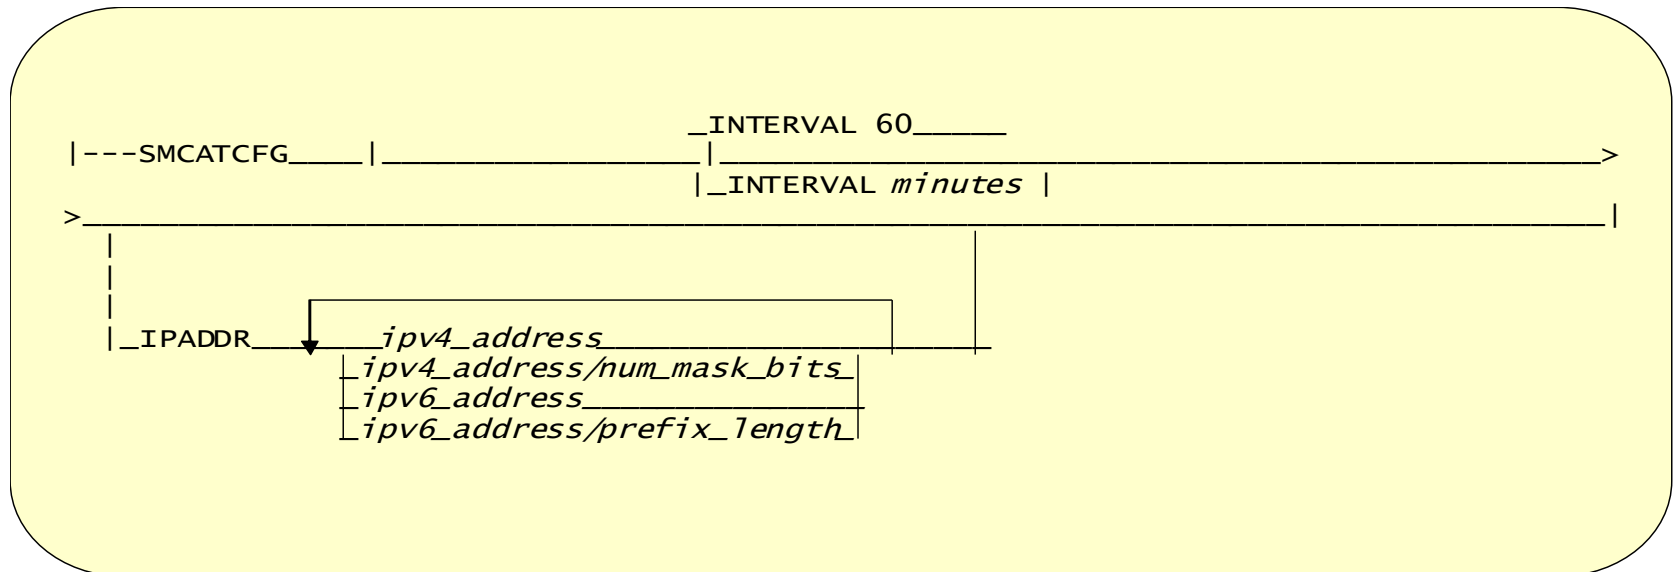
1. SMC-AT Configuration and Usage (starting / stopping) Details
2. SMC-R and SMC-AT Evaluation Considerations / References
3. SMC-R Background

Backup Topic 1 (Configuration and Usage Details)

Configuring the SMC-AT Dataset

SMC-AT data set configuration

- Interval defaults to 60 minutes
Max interval is 1440 minutes (24 hours)
- IPADDR is a list of IPv4 and Ipv6 addresses and subnets
256 max combination of addresses and subnets



SMC-AT Dataset Example

```
SMCATCFG INTERVAL 120  
IPADDR  
C5::1:2:3:4/126  
9.67.113.61
```

When SMC-AT is started using this SMC-AT configuration data set it will:

- Monitor TCP traffic for 2 hours for:
 - IPv6 prefix C5::1:2:3:4/126 and
 - IPv4 address 9.67.113.61

Starting and Stopping SMC-AT

Vary TCPIP,,SMCAT command starts and stops the monitoring tool:

- datasetname value indicates that SMC-AT is being turned on
- datasetname contains the SMCATCFG statement that specifies monitoring interval and IP addresses or subnets to be monitored
- OFF will stop SMC-AT monitoring and generate report

```
>>__Vary__TCPIP,_____,__SMCAT,____datasetname____><  
          |_procname_|          |_OFF_|
```

```
VARY TCPIP,TCPPROC,SMCAT,USER99.TCPIP.SMCAT1
```

SMC-AT Operator Messages

Key messages – Operator console:


- EZD2031I SMC APPLICABILITY TOOL HAS STARTED COLLECTING DATA
- EZD2032I SMC APPLICABILITY TOOL HAS STOPPED COLLECTING DATA

Configuration information and the SMC-AT report are sent to the system log

```
STC06578      EZD2040I TCP/IP CS V2R2 TCPIP Name: TCPIP
080           SMC Applicability Configuration Parameters - 02/04/2015, 10:09:49.08
080           Interval: 3 minutes
080           IP addresses/subnets being monitored
080
080           9.67.113.61
080           C5::1:2:3:4/126
080           End of configuration parameters
```

Backup Topic 2 (SMC-R and SMC-AT Considerations)

Backup **SMC-R – Micro benchmark performance results**



Significant Latency reduction across all data sizes (52-88%)
Reduced CPU cost as payload increases (up to 56% CPU savings)
Impressive throughput gains across all data sizes (Up to +717%)

*Note: vs typical OSA customer configuration
 MTU (1500), Large Send disabled
 RoCE MTU: 1K*

SMC-R and SMC-AT Considerations

SMC-AT Report Terminology:

- Eligible connections:
TCP connections that match SMC-AT configured IP addresses/subnets (IPSEC and FRCA connections are excluded)
- Well-suited connections:
Eligible connections that have workloads ideally suited to exploit SMC-R (not short lived)

RoCE Distance (SMC-R) imposes limits on the distance between exploiting hosts:

- Typically deployed within a datacenter or site
- Can be deployed across sites for a distance up to 100 KM

IBM System z® Qualified Wavelength Division Multiplexer (WDM) products for Multi-site Sysplex and GDPS® solutions

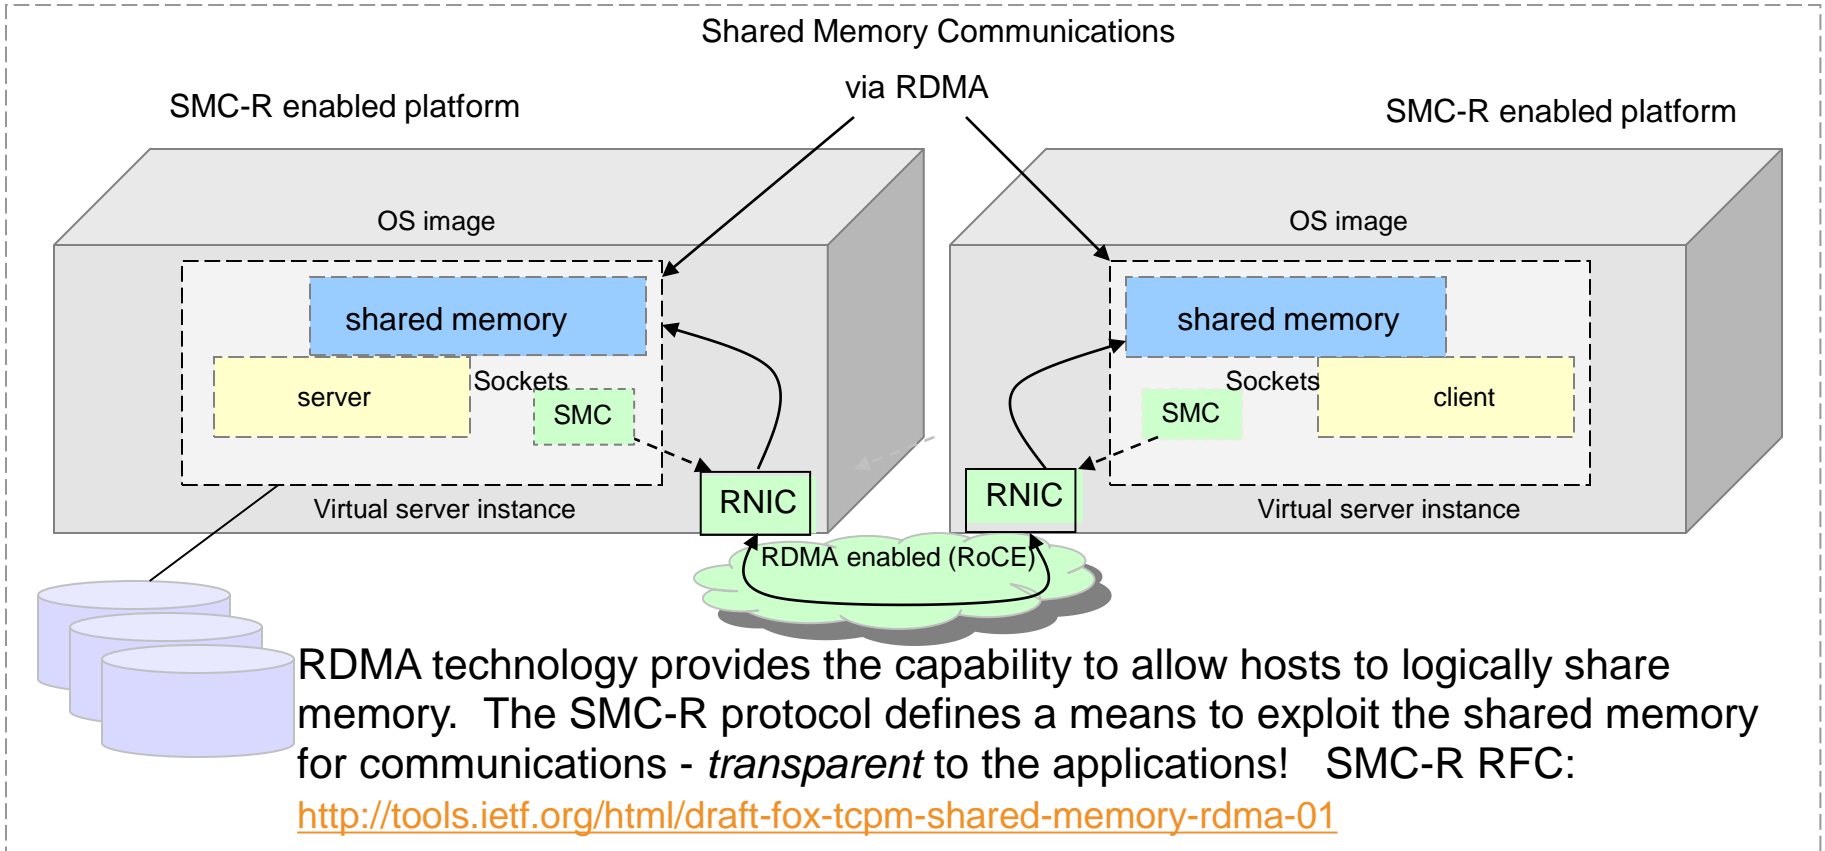
SMC-R and SMC-AT Considerations

- Consider configuring a subset of IP addresses for similar workloads/similar servers and then extrapolate the SMC-AT results of your sample across your enterprise data center
- SMC-AT does not take into consideration your type of network attachment (e.g., OSA, XCF, HiperSockets, etc.). SMC-R only applies to RoCE (Ethernet). In order to exploit SMC-R only the TCP connections that use OSA are capable of using SMC-R.
- Authorize users to invoke the command by permitting user IDs for CONTROL access to the RACF profile name MVS.VARY.TCPIP.SMCAT
- The configuration data set must be cataloged and fully-qualified
 - Specified without any quotation marks
 - Can be either a sequential data set or a member in a PDS
 - Cannot be a z/OS UNIX file
- Communications Server: IP System Administrator's Commands
 - VARY TCPIP,,SMCAT section
- SMC-R references:
<http://www-01.ibm.com/software/network/commserver/SMCR/>

Backup Topic 3 (SMC-R Background)

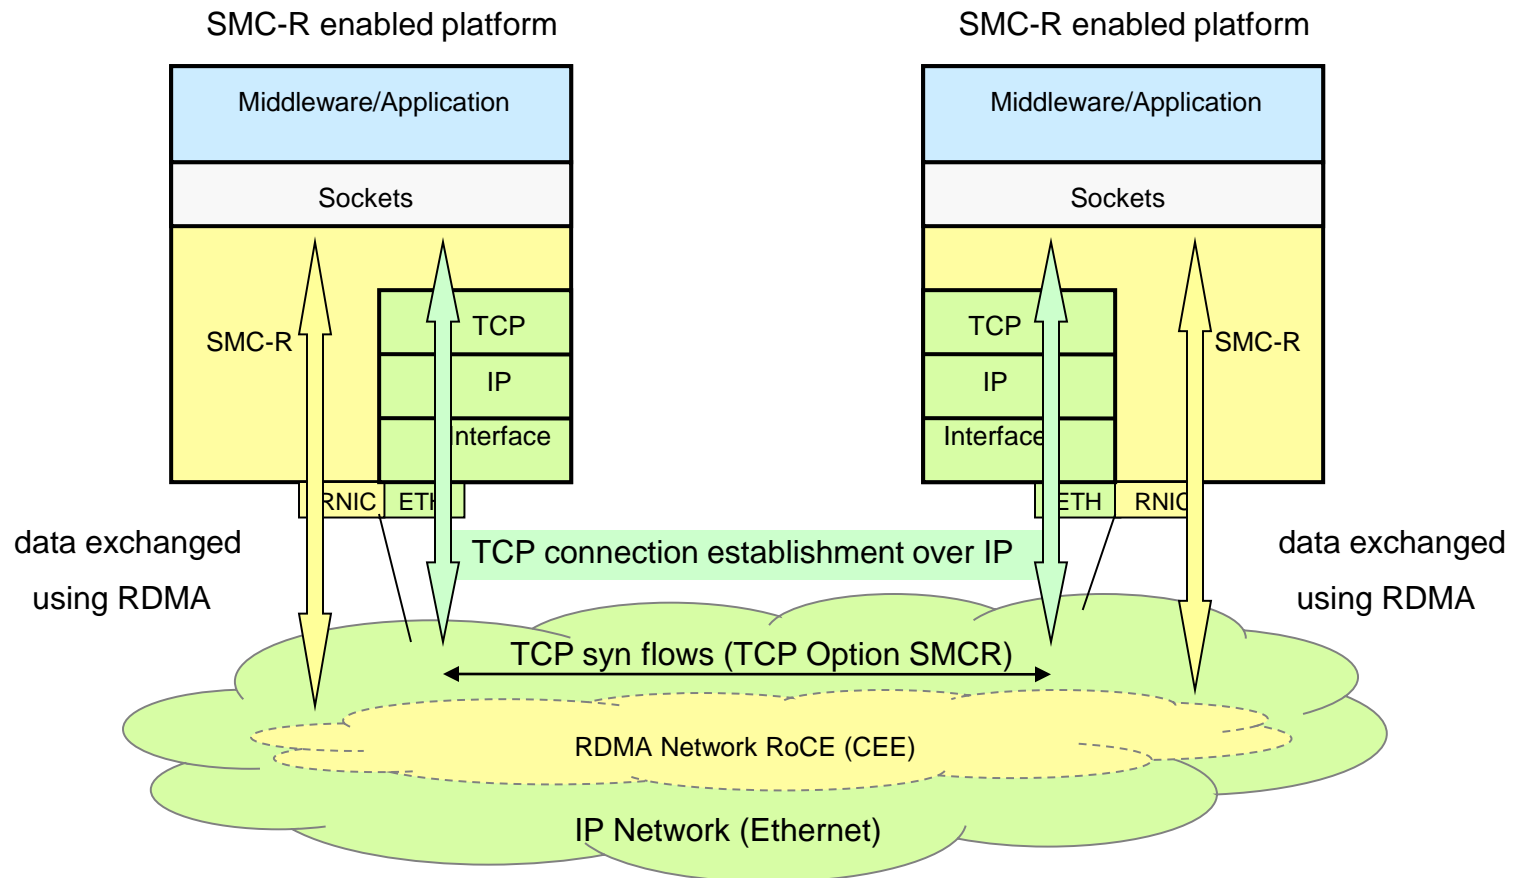
“Shared Memory Communications over RDMA” concepts

Clustered Systems



This solution is referred to as *SMC-R* (Shared Memory Communications over RDMA). SMC-R represents a sockets over RDMA protocol that provides a foundation for a complete solution meeting all of the described objectives. SMC-R is an RDMA model exploiting RDMA-writes (only) for all data movement.

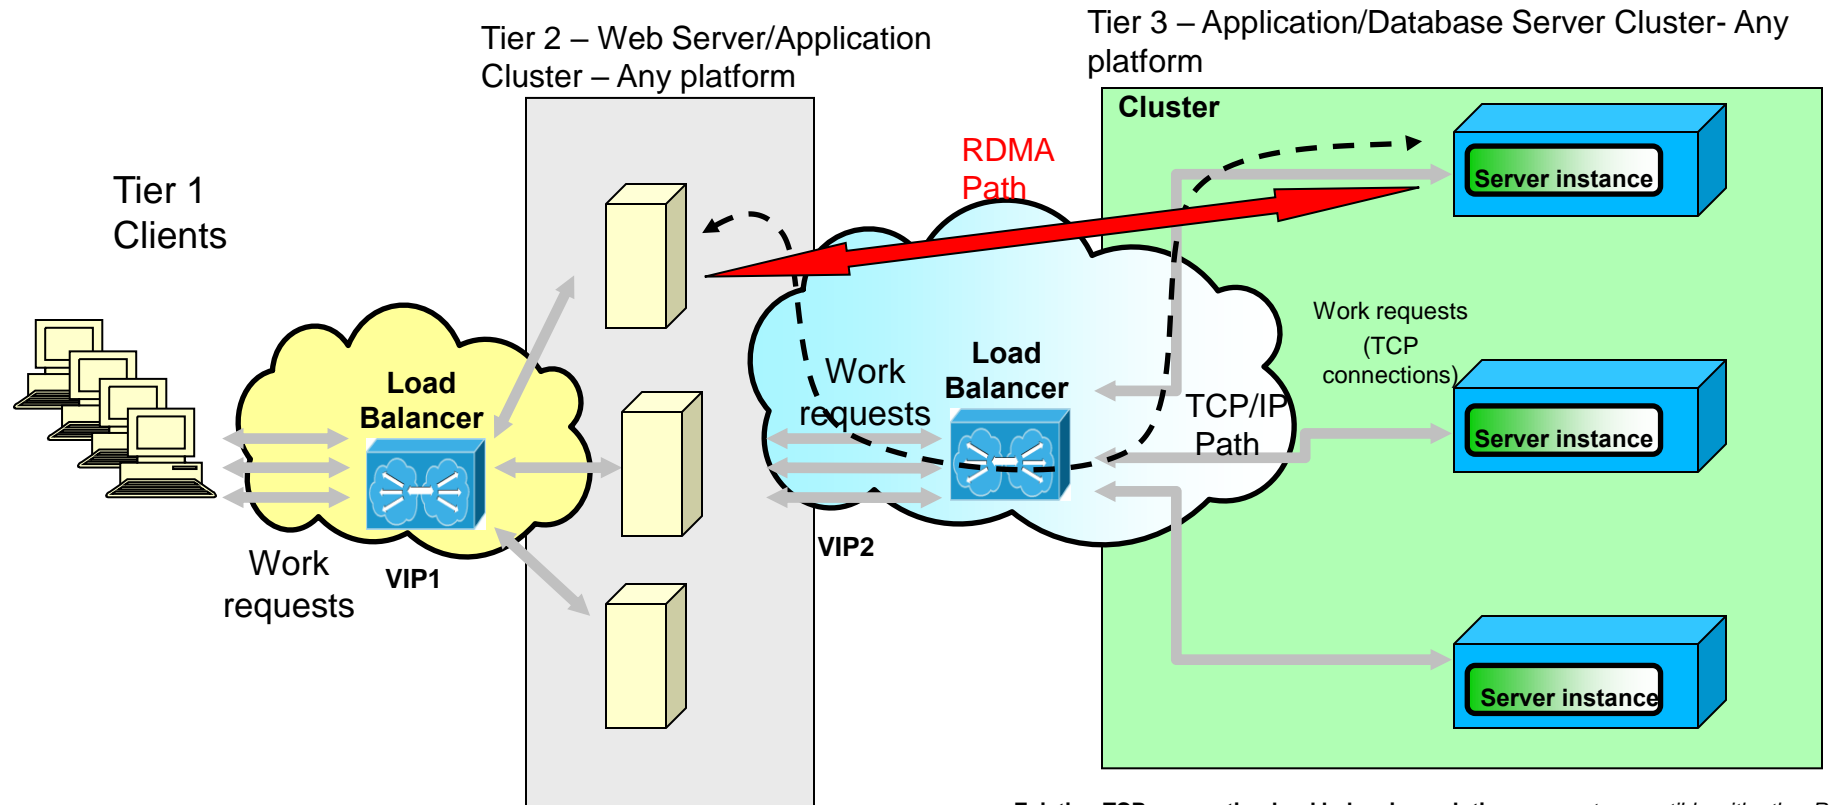
Dynamic Transition from TCP to SMC-R



Dynamic (in-line) negotiation for SMC-R is initiated by presence of TCP Option (SMCR)

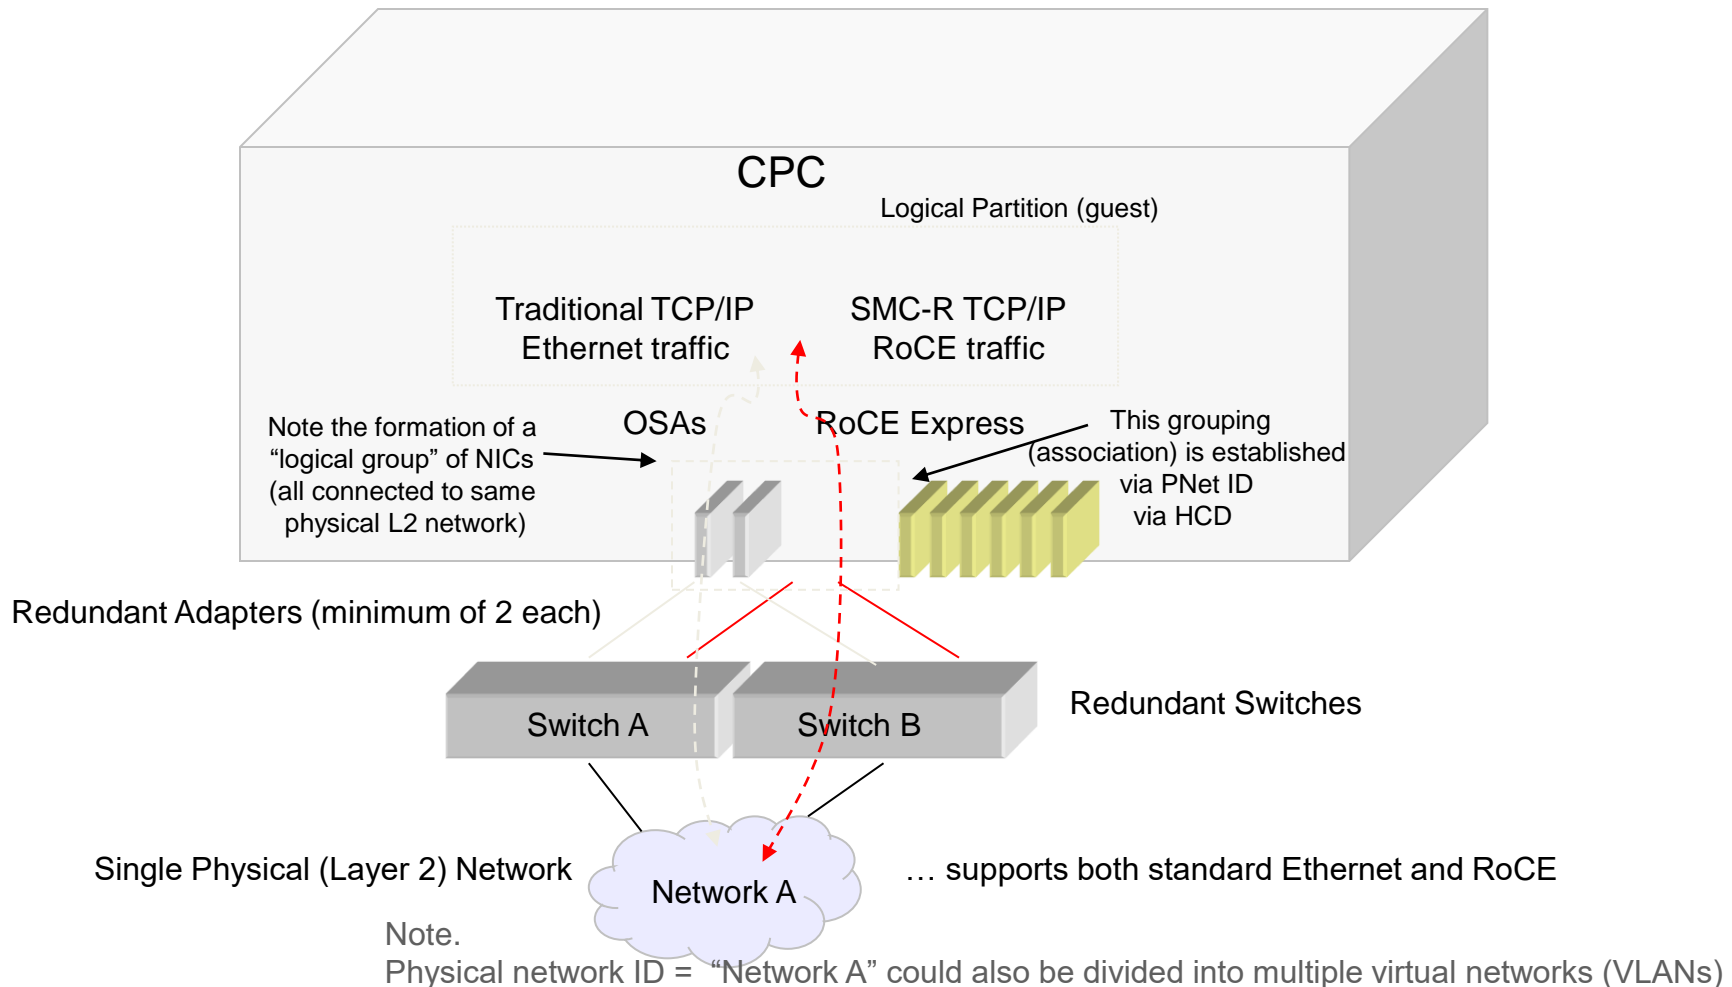
TCP connection transitions to SMC-R allowing application data to be exchanged using RDMA

Server clustering and TCP connection load balancing



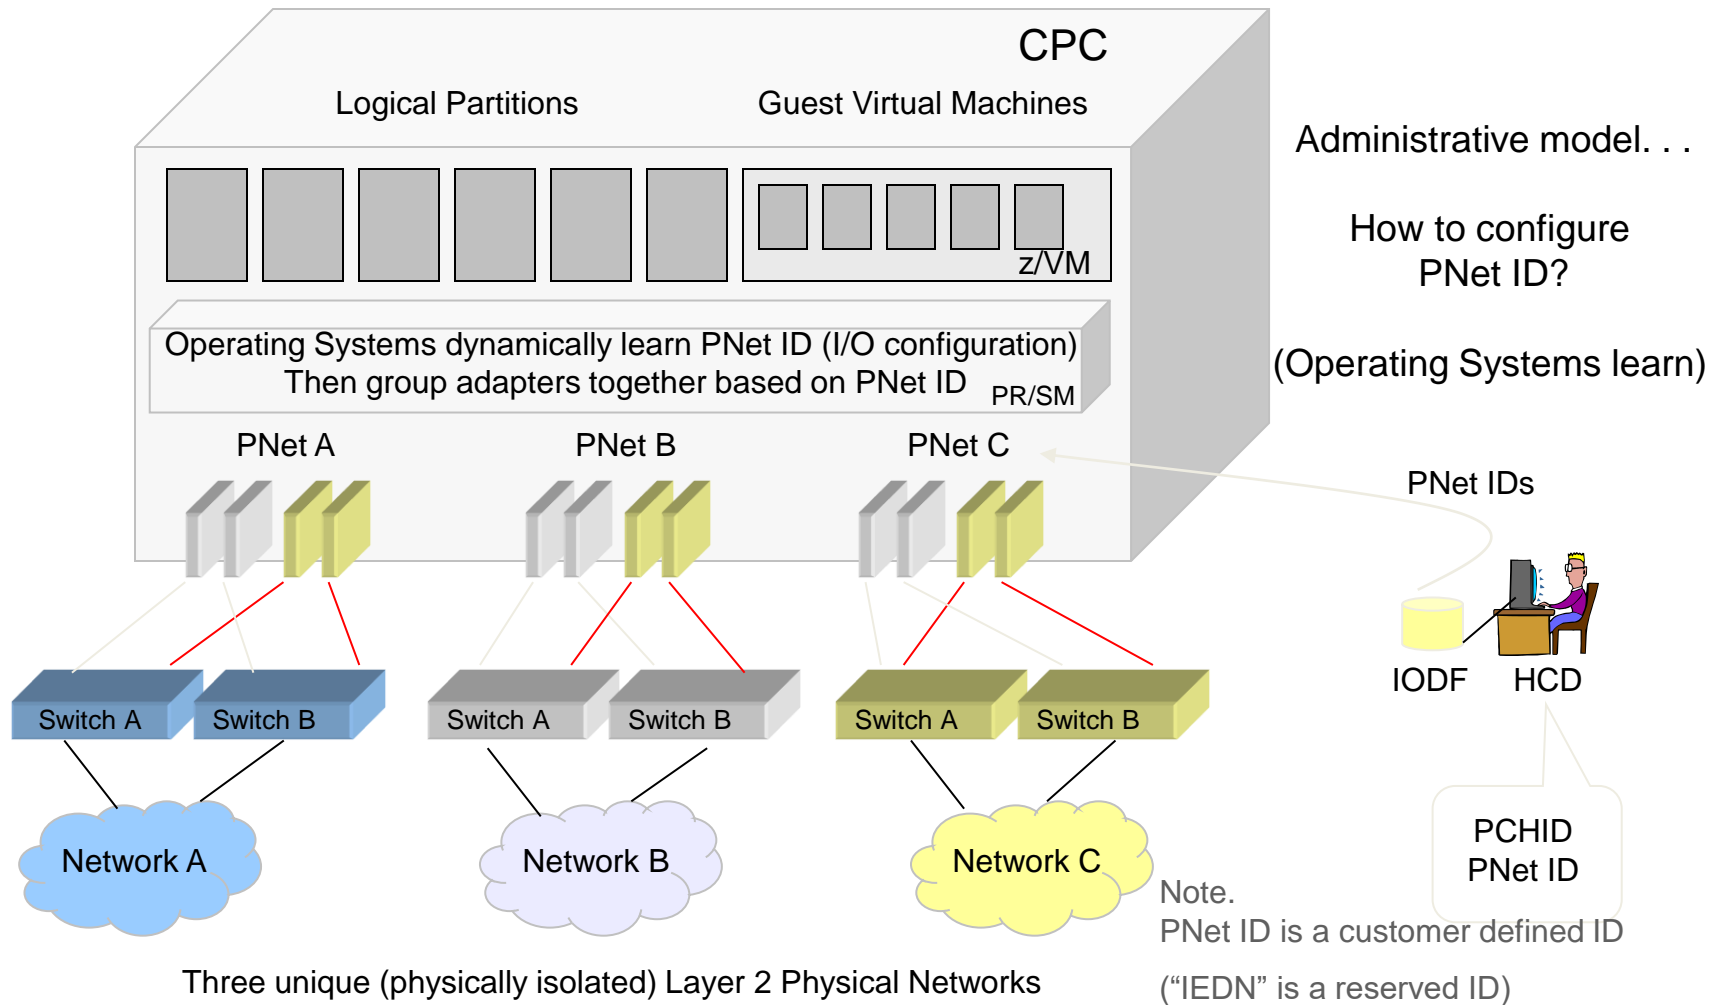
- **Server clustering is a prevalent deployment pattern for Enterprise class servers**
 - Provide High Available, eliminate single points of failure, ability to grow/shrink capacity dynamically, ability to perform non-disruptive planned maintenance, etc.
- **TCP connection load balancing is a key solution for load balancing within a cluster environment**
 - External or Internal load balancers provide this capability
- **Existing TCP connection load balancing solutions are not compatible with other RDMA solutions**
- They are not aware of the RDMA protocol **AND** RDMA flows **can not** flow through intermediate nodes
- **The SMC-R protocol allows existing TCP load balancing solutions to be deployed with no changes**
 - TCP Connection load balancing for SMC-R connections is actually more efficient than normal TCP/IP connections
 - Load balancer selects optimal back end server, data flows can then bypass the load balancer

Network Physical Connectivity for the IBM 10 Gbe RoCE Express feature System z network adapters (Requires OSAs + RoCE Express)



Multiple Physical RoCE Networks (Physical Network IDs)

Associating 10 GbE RoCE Express features with their physical networks



Value Summary: IBM 10 GbE RoCE Express with SMC-R

Summary

- z/OS application workloads transparently exploit 10 GbE RoCE Express feature using z/OS V2R1 SMC-R. When SMC-R is enabled:
 - transactional workloads (WAS, CICS, IMS, MQ, etc.) can potentially see an increase in their overall transaction rate (i.e., transactions per second) with a slight savings in CPU.
 - Streaming workloads (e.g., FTP) will see a CPU savings and a throughput improvement.
 - SD workloads (where the client and server both are within the SYSPLEX (i.e., WAS to DB2)) will benefit by eliminating SD in the network path (reduces cost in the SD host and reduces latency by avoiding the trip through the SD host)

Value

- Reduced network latency resulting in improved transaction rates
- Reduce CPU cost for bulk or streaming workloads or transaction workloads with larger messages (e.g., web services protocols)
- Time to value.... optimized network performance without:
 - requiring application changes
 - requiring network IP topology or security changes
 - sacrificing existing TCP/IP qualities of services (e.g., network resiliency)
 - significant operational (day to day administrative) changes

Thank You