- OSA, HiperSockets, and RoCE are all used for data transfer between hosts
- Cross CEC zEnterprise traffic
  - OSA and RoCE can be used to send traffic between CECs.
  - HiperSockets only supports traffic between LPARs on a single CEC.
- Different zEnterprise Operating Systems
  - OSA and HiperSockets are supported by multiple Operating Systems (ie. z/OS, z/VM, Linux on System z, etc.)
  - RoCE is only supported by z/OS.
- Traffic outside of a single zEnterprise CEC (and might therefore require additional security measures)
  - OSA traffic can go over the card if shared between LPARs on a single CEC.
  - OSA traffic goes over a network if used to a non-shared-OSA partner.
  - RoCE traffic goes over a 10GbE Layer 2 LAN.
- Security exposure is limited if contained in a secure location.
  - HiperSockets traffic never goes outside a single CEC.
- Firewall with stateful packet inspection (a PCI (Payment Card Industry) requirement for some traffic)
  - OSA traffic can be sent over a LAN to a Firewall with stateful packet inspection support.
  - HiperSockets traffic cannot be sent over a Firewall with stateful packet inspection support (unless routed traffic).
  - RoCE traffic cannot be sent over a Firewall with stateful packet in support (does not support routed traffic).
- Protocol Support
  - OSA supports all TCP/IP protocols and even supports SNA protocol (natively in OSE mode, or when sent with UDP (Enterprise Extender (EE))).
  - HiperSockets only supports IP traffic, it does not support native SNA (so EE must be used to send SNA).
  - RoCE only supports TCP traffic (except IPsec), it does not support native SNA or UDP (EE).
- Required hardware feature
  - OSA feature is required.
  - HiperSockets is part of System z Firmware so it does not require any additional hardware / adapter card purchase.
  - RoCE feature is required.
    NOTE: A minimum of 2 per LPAR is recommended.
- CP Overhead
  - OSA provides many different types of offload to the adapter that reduces CP overhead (ARP, Check Sum, Segmentation, etc.)
  - HiperSockets supports zIIP offload to reduce associated cost.
  - RoCE reduces TCP/IP overhead by using RDMA protocol.
- Storage Usage
  - OSA and HiperSockets use CSM fixed storage backed by 64-bit real for data buffers and use Hardware System Area (HSA) memory for routing table.
  - RoCE uses pinned Fixed Memory (mostly 64-bit Common), not CSM managed memory.
    NOTE: The maximum amount of memory available for SMC-R with RoCE is definable. When the maximum is reached, new connections will *not* be SMC-R RoCE eligible.

Open Systems Adapter (OSA)

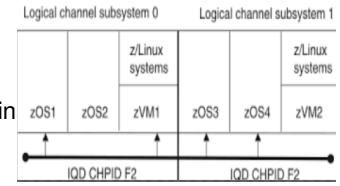

HiperSockets

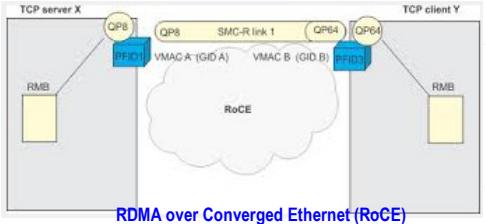

RDMA over Converged Ethernet (RoCE)

RDMA (Remote Direct Memory Access) enables a host to read or write directly from/ to a remote host's memory ***without*** involving the remote host's CPU.

The **Open Systems Adapter (OSA)** is a network controller that you can install in an IBM System z mainframe I/O cage. The adapter integrates several hardware features and supports many networking transport protocols. The OSA card is the Strategic communications device for the System z architecture. It has several key features that distinguish it from communications that are based on channel control words (CCWs).

- IP Routing
  - OSA requires IP routing.
    NOTE: OSA does provide dynamic backup with multiple OSAs to the same subnet and Multipath. When multiple OSAs are attached to different subnets or OSA and other attachements like HiperSockets are used there is no dynamic backup without a Dynamic Routing protocol (ie. OSPF).
  - HiperSockets requires IP routing so there is no dynamic backup without a Dynamic Routing protocol.
- When HiperSockets is defined as part of a DynamicXCF network routing is handled automatically.
- When HiperSockets is defined with Integration with IEDN (IQDX) routing is handled automatically.
- Backup might not be a requirement because if HiperSockets fails there is probably major CEC problems occurring.
  - Traffic is automatically / transparently switched from OSA to RoCE.
    NOTE: If RoCE connection is not available traffic is sent over OSA. RoCE has automatic / transparent backup to a different RoCE path if one exists.
- Active RoCE sessions are dropped if the RoCE connection fails and there is no alternate RoCE path, but new sessions will flow using OSA.
- IP Routed Traffic
  - OSA supports routed traffic. Routed traffic is optimized on z/OS with QDIO Accelerator.
  - HiperSockets supports routed traffic (ie. traffic may come in over OSA and then be routed over HiperSockets).
    NOTE: Routed traffic is optimized on z/OS with QDIO Accelerator.
  - RoCE does not support routed traffic. All connections are over a Layer 2 network.
- VLAN Support
  - OSA supports VLAN tagging.
  - HiperSockets supports VLAN tagging.
  - RoCE inherits VLAN IDs from all associated OSAs and supports VLAN tagging.
- Routing Overhead

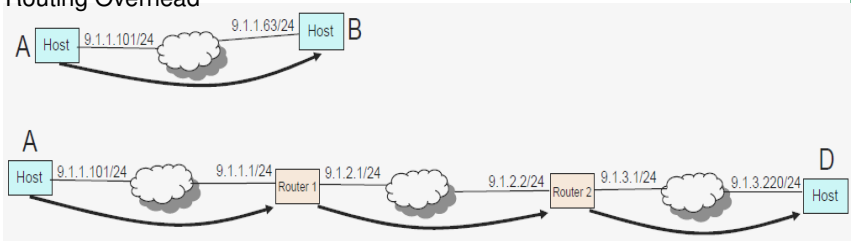

  - There is an overhead associated with routing.
    - All things being equal, performance is typically better between two hosts on the same subnet (sometimes referred to as a Layer 2 connection (or Flat network)), rather than two hosts on different subnets.
      NOTE: One of the benefits of the IBM zEnterprise zBX.

Reference: Document PRS5038