SHARE in Orlando 2015

SHARE

Introduction to Mainframe (z/OS) Network Management

Monday, August 10, 1:45-2:45

Session 17736

Dean Butler (butlerde@us.ibm.com)

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides **education**, professional **networking** and industry **influence**.

CELEBRATING 60 YEARS OF SHARE

Agenda

- What is network management?
- Why manage the network on z/OS?
- z/OS networking
- Network monitoring tools and technologies
- Best practices

What is network management?

FCAPS

ISO standard model

Fault management

Goal: Keep network operational, minimize downtime.

Find and correct network problems. Identify future potential problems and prevent from occurring or reoccurring.

Configuration management

Goal: Monitor network and system configuration to track and manage changes.

Monitor and control network operation. Coordinate hardware and programming changes. Maintain an inventory.

Accounting / Admin

Goal: Distribute resources optimally and fairly among network subscribers.

Measure utilization of all important network resources. Analyze results. Regulate, bill, or charge users.

Performance management

Goal: Measure and make network performance data available to optimize performance.

Gather, analyze, and threshold performance data.

Security management

Goal: control access to assets in the network

Manage network authentication, authorization, and auditing. Maintain data security with authentication and encryption.

Network management architecture

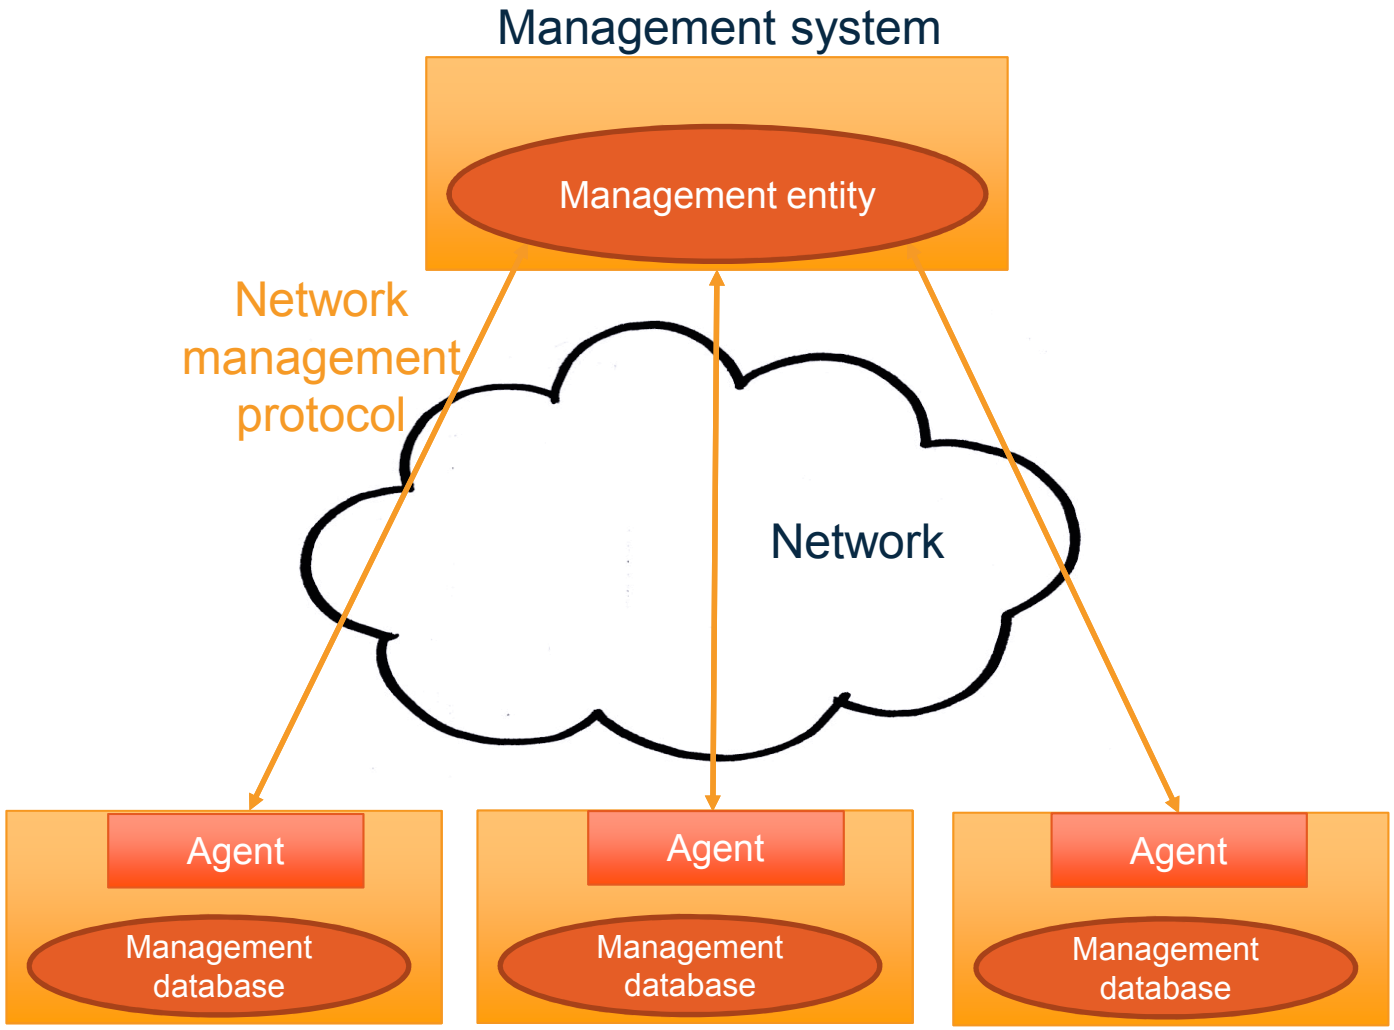

Agenda

- What is network management?
- Why manage the network on z/OS?
- z/OS networking
- Network monitoring tools and technologies
- Best practices

Why manage the network on z/OS?

Why the network?

A well-managed network will allow you to get on with running your company.

- Prevent problems
- Work efficiently
- Maintain security
- Stay up to date

Why on z/OS?

In spite of perception that z/OS networking doesn't break, SLA's are at risk:

- Application changes
- Hybrid networks
- BYOD
- Device failures
- Network congestion

Network systems programmer:
"A critical application is 'broken'. We all get on a bridge call. Everyone says 'My stuff is okay. It must be the network.' I need to be able to say it is not the network or at least not my part of the network."

Agenda

- What is network management?
- Why manage the network on z/OS?
- z/OS networking
- Network monitoring tools and technologies
- Best practices

System z network connectivity

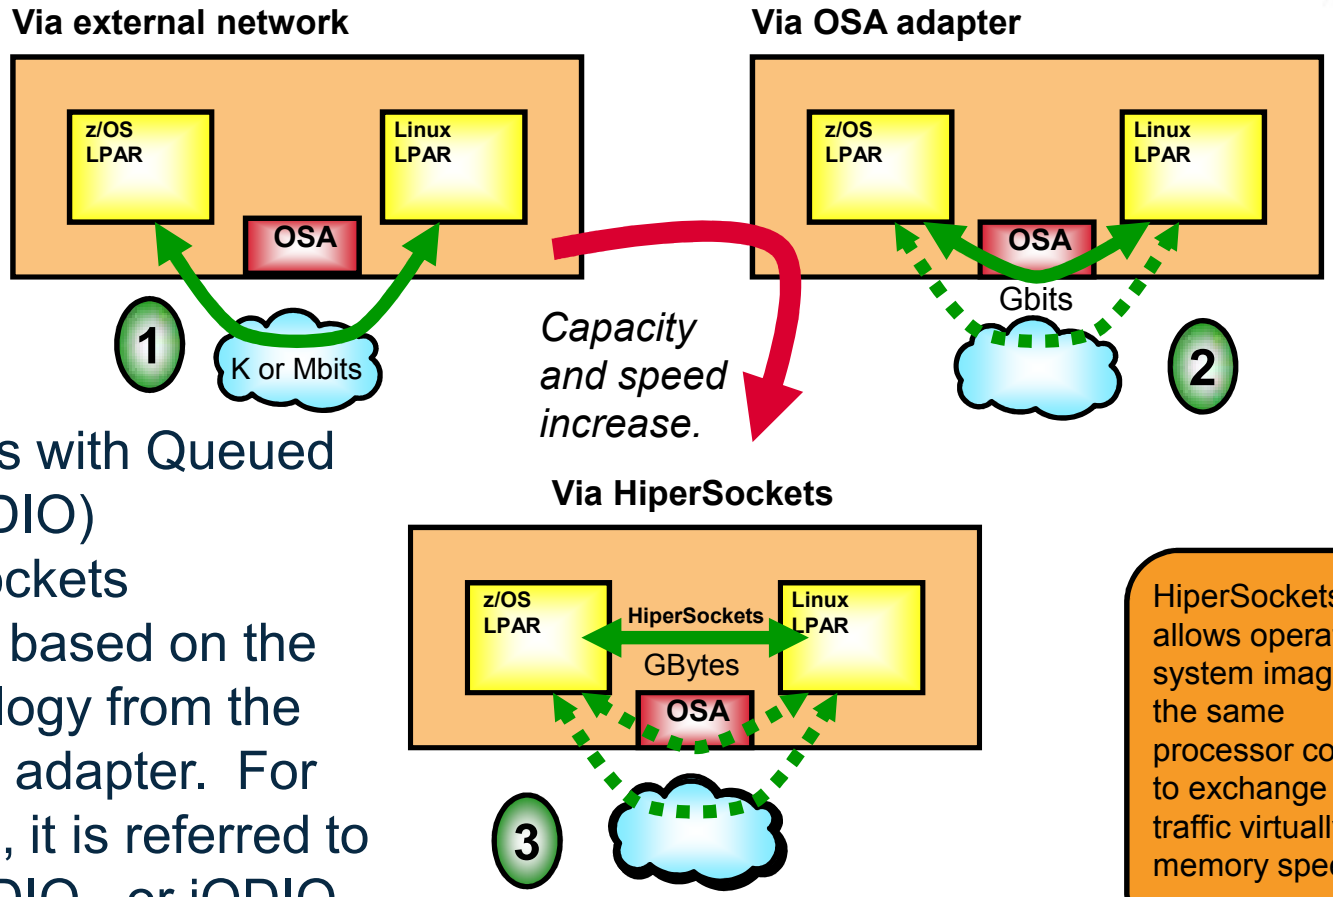

- OSA Express with Queued Direct I/O (QDIO)
- The HiperSockets technology is based on the QDIO technology from the OSA Express adapter. For HiperSockets, it is referred to as internal QDIO - or iQDIO for short.

HiperSockets allows operating system images on the same processor complex to exchange IP traffic virtually at memory speed.

A HiperSockets network looks like an internal LAN

z/OS Communications Server

- **Integrated Services**
  - **Provide common services within CS**
    - Network attachment
    - Storage management
    - High Performance Data Transfer
  - **TCP/IP and SNA integration**
    - TN3270
    - Network access
    - Internal optimizations
    - Enterprise Extender
  - **Standard TCP/IP applications**
- **Multi-protocol Solutions**
  - **Sockets (TCP/IP) applications**
    - Unix services offers z/OS users access to a wide range of UNIX-based applications over IP networks
  - **SNA applications**
    - SNA applications are supported over SNA or IP networks

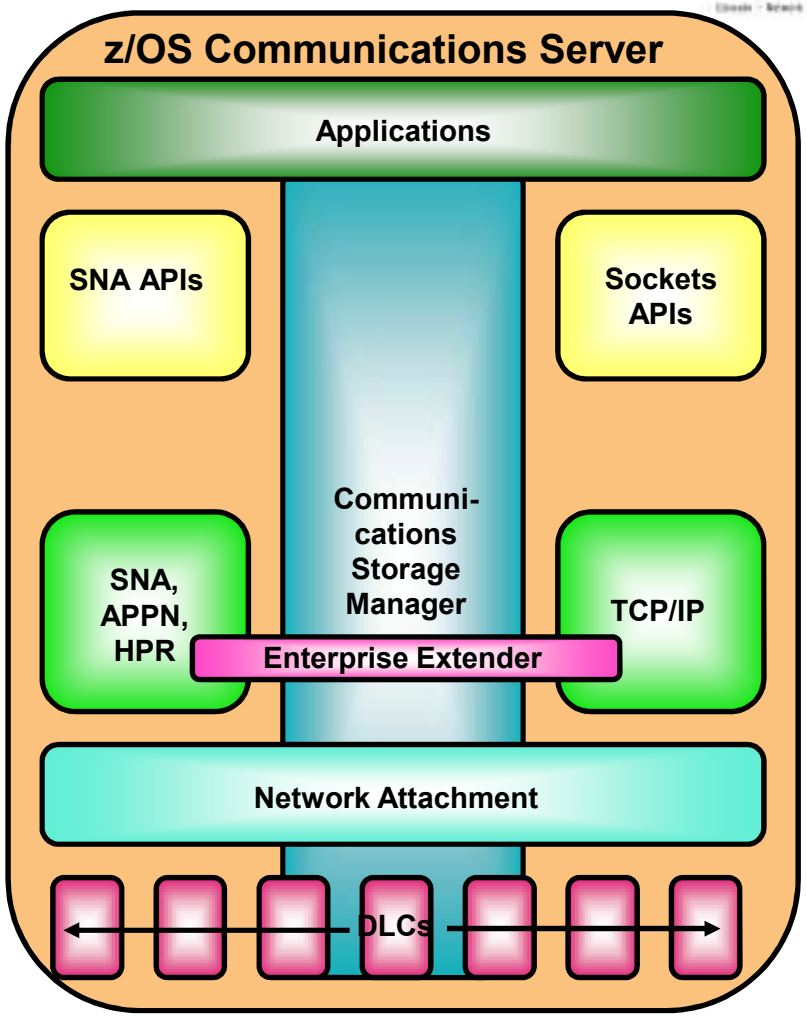

Enterprise Extender

- Allows use of IP network for SNA sessions
  - To an SNA application, the IP network looks like an HPR link
  - To the IP network, EE is just a UDP application
- SNA traffic is sent as UDP datagrams (on ports 12000-12004) over the IP network
- Supported by many platforms:
  - z/OS
  - CS/Linux, CS/AIX, CS/Windows
- PComm
  - i5/OS
  - Cisco SNASw
  - Microsoft HIS
- Typically isolates SNA footprints to the "outside" of the network.

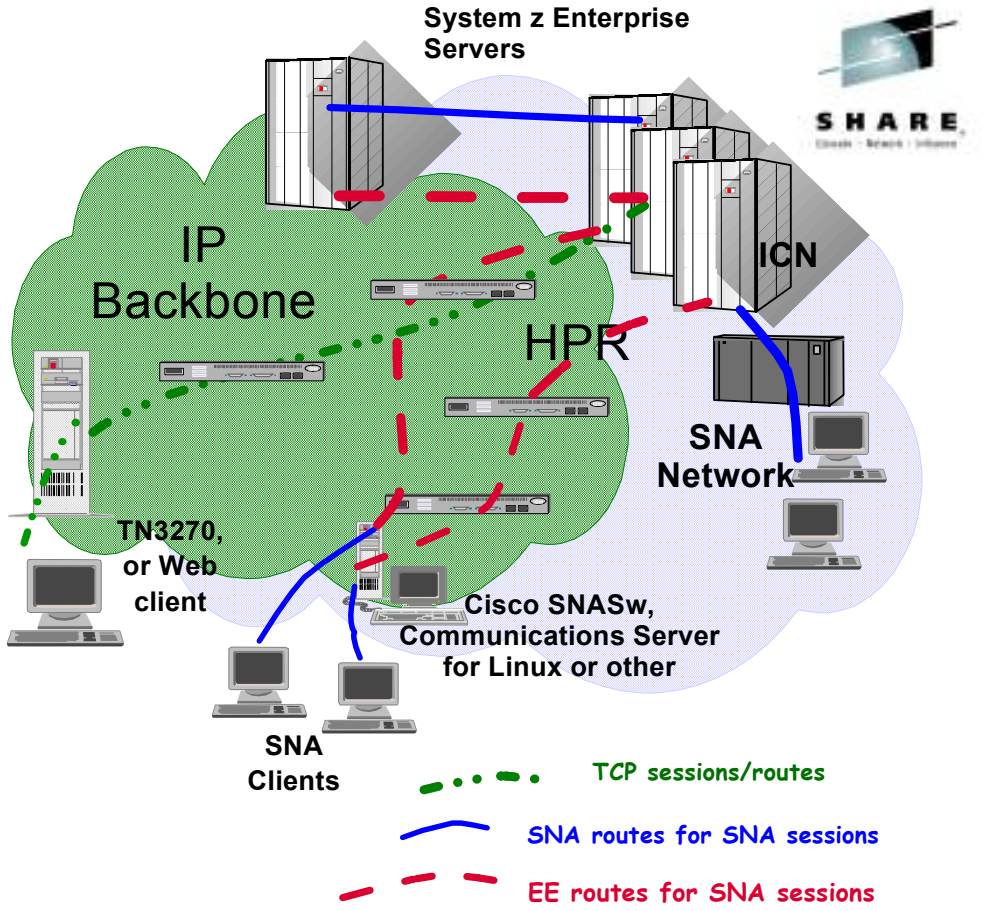

EE allows enablement of IP applications and convergence on a single network transport while preserving SNA application and endpoint investment.

How z/OS communicates with the outside world

- Enable IBM and ISV network management products to provide differentiated value on System z
- Enable network connectivity with differentiated z/OS Qualities of Service
  - Performance
  - Security
  - High Availability

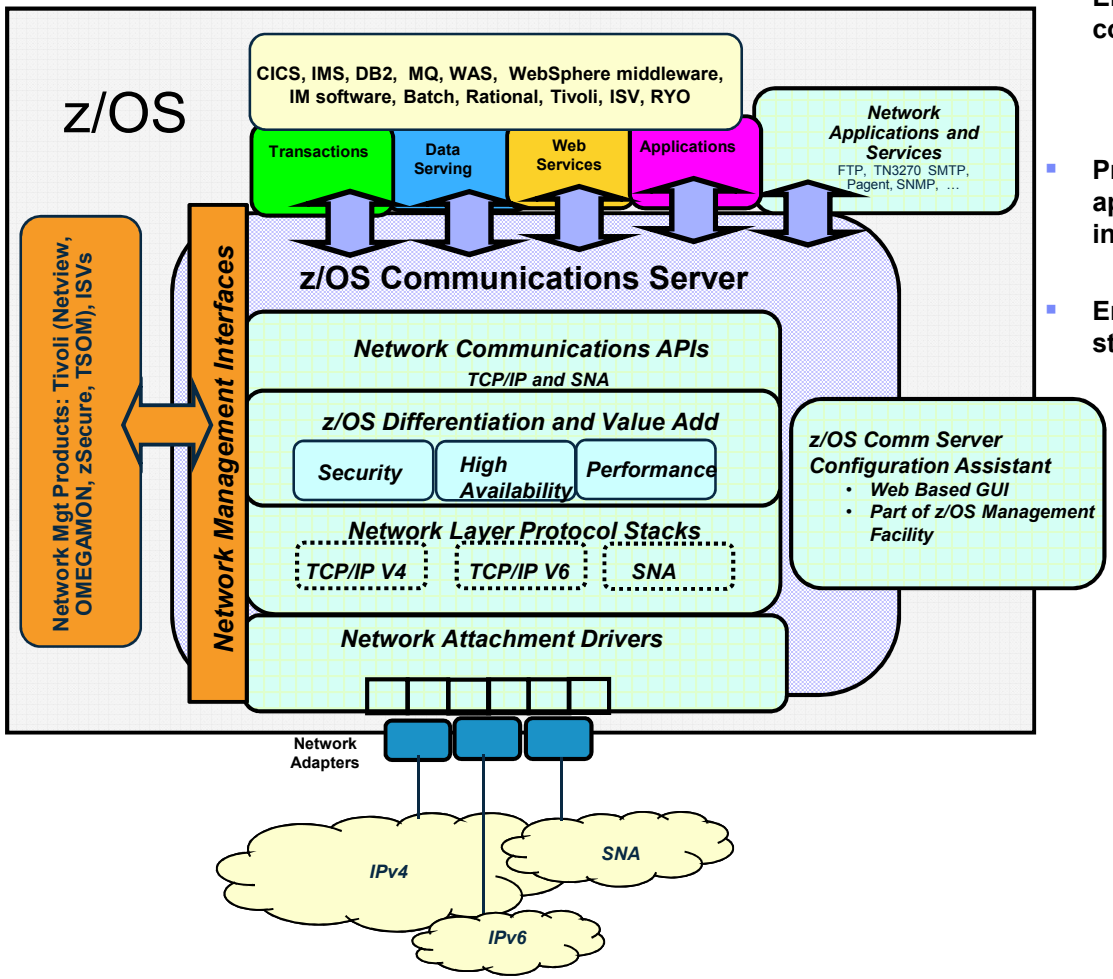

- Enables all key z/OS workload to communicate with the outside world
  - *Every request for service and response flows through Comm Server*
- Provides standard suite of TCP/IP applications *but* with close integration into z/OS
- Enables network connectivity using standards based network protocols
  - IPv4: Current internet protocol, support for latest standards
  - IPv6: Emerging next generation Internet protocol
  - SNA: Maintain existing customer application investment

Typical data center network

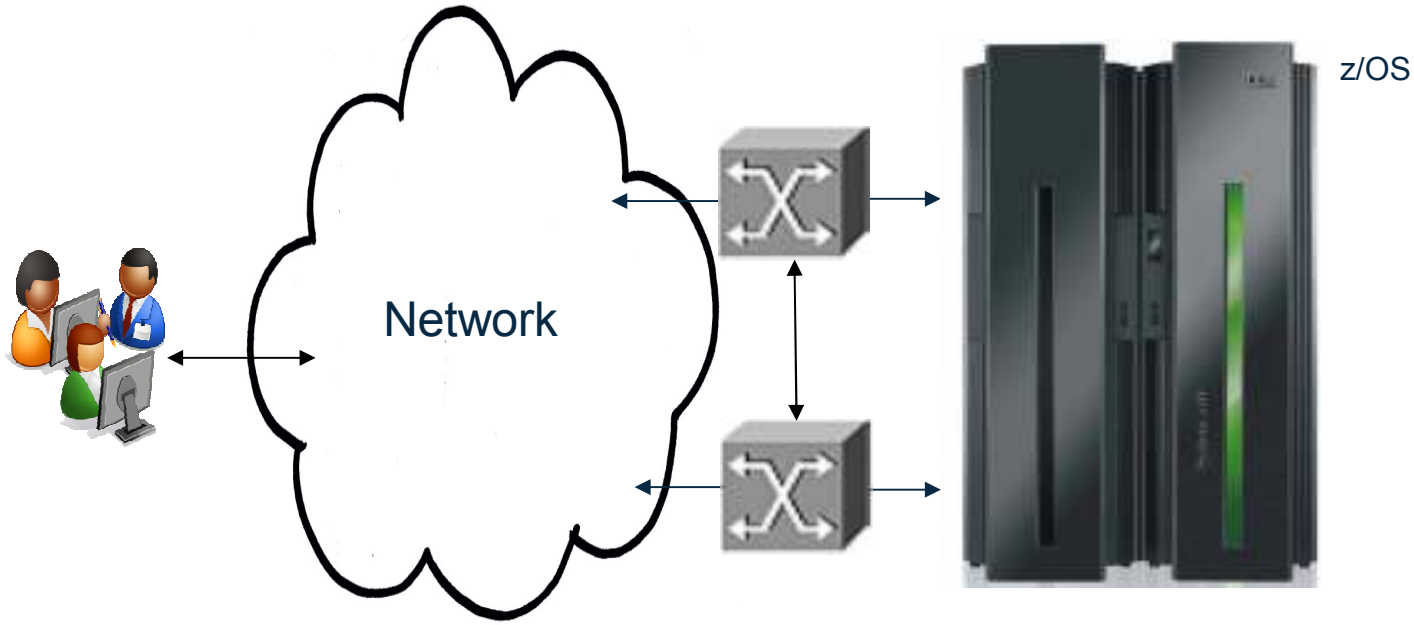

Agenda

- What is network management?
- Why manage the network on z/OS?
- z/OS networking
- Network monitoring tools and technologies
- Best practices

Network diagnostic commands

PING

NETSTAT

TRACERTE

DISPLAY

NSLOOKUP

VARY

DROP

Network diagnostic commands

- PING
  - Test connectivity to a host or device on an IP network
  - ICMP echo request / echo response
  - Measures round trip time

- TRACERTE
  - Shows most likely path to an IP device
  - Series of ICMP or UDP packets
  - Uses TTL to identify network hops in path
  - Measures round trip time to each hop

Network diagnostic commands

- NSLOOKUP
  - Query the Domain Name System (DNS)
  - Translate host name to/from IP address

- DROP
  - Terminate a TCP/IP socket endpoint
  - TCP listener, TCP connection, UDP endpoint
  - Requires CONTROL access to MVS.VARY.TCPIP.DROP

Network diagnostic commands

- NETSTAT
  - Display information about the **local host**, including TCP/IP configuration, connections, network clients, gateways, and devices
  - Drop connections (access to MVS.VARY.TCPIP.DROP)

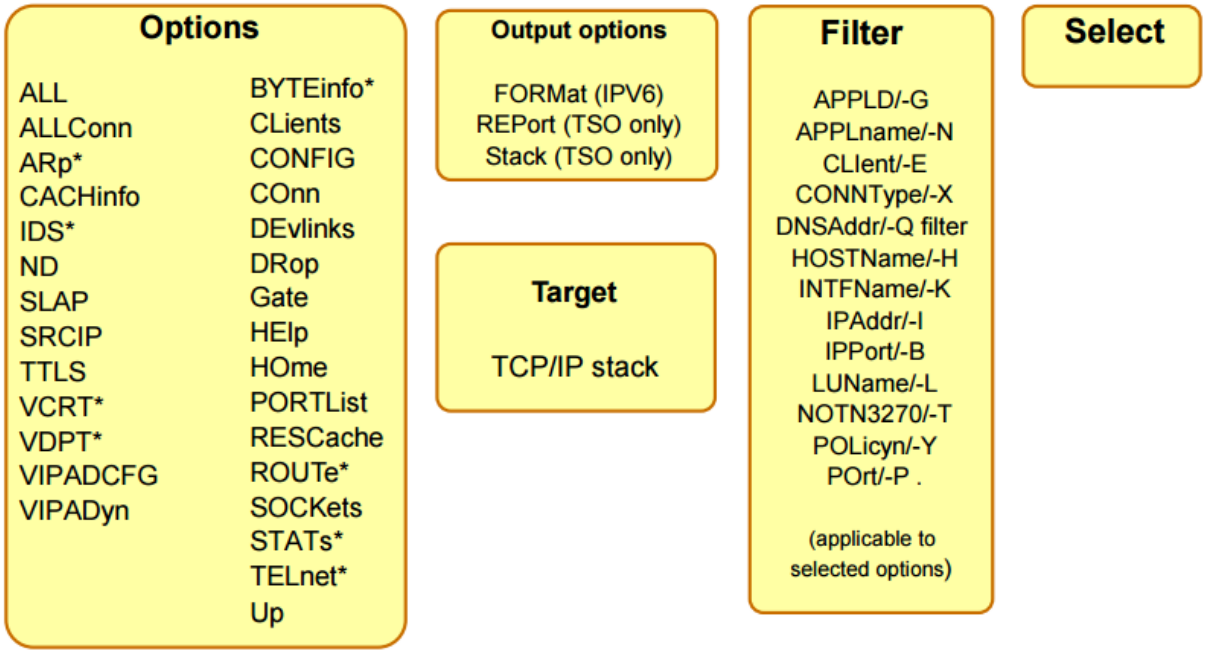

Network diagnostic commands

z/OS operator commands:

- DISPLAY TCPIP
  - Display information about the **local** TCP/IP or TELNET address spaces
- VARY TCPIP
  - Control some functions of **local** TCP/IP and TELNET address spaces

What is SNMP?

Simple Network Management Protocol

One of the most widely accepted protocols to manage and monitor network elements.

- UDP application
- Each SNMP agent maintains an information database describing the managed device parameters
- SNMP manager uses this database to request information from the SNMP agents
- SNMP agent sends some information as traps

Basic commands:

- Get, Get next, Get bulk
- Set
- Traps
- Inform
- Response

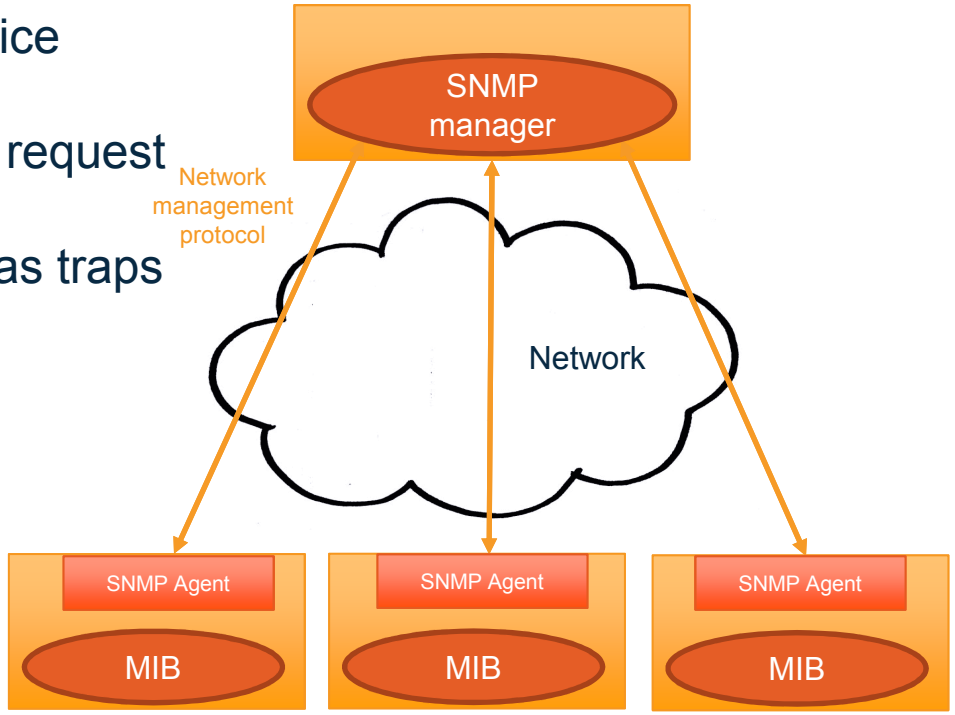

What is SNMP?

An Object Identifier (OID) identifies a managed object.

For example, the OID in RFC1213 for "sysDescr" is .1.3.6.1.2.1.1.1

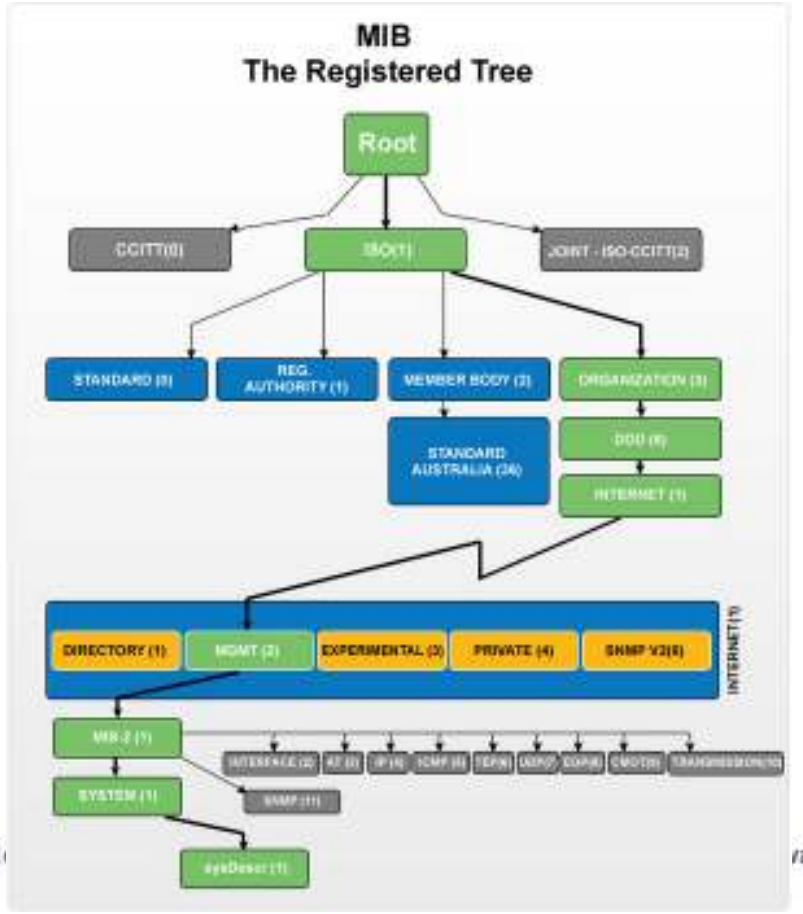

SNMP packet flows:

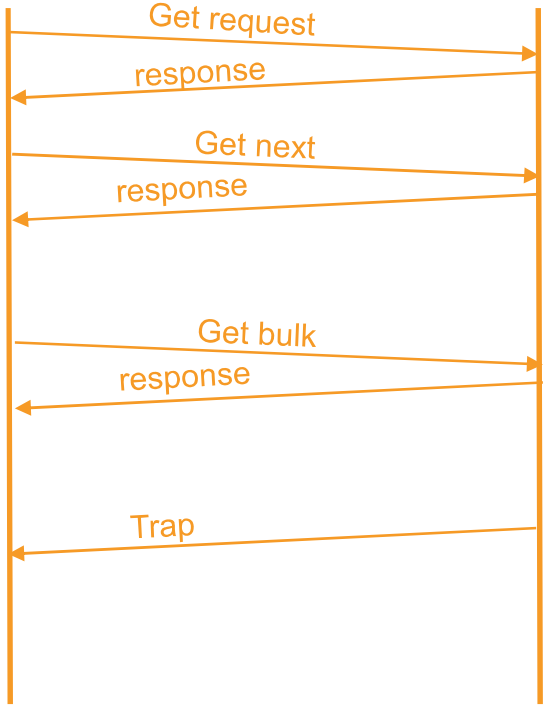

System Management Facilities (SMF)

- Standardized method for writing out records of activity to a file (or data set to use a z/OS term)
- Full "instrumentation" of all baseline activities running on an LPAR, including I/O, network activity, software usage, error conditions, processor utilization, etc.
- SMF forms the basis for many monitoring and automation utilities
- Each SMF record has a numbered type (IBM: 1-127, others: 128+)
- Installations have great control over how much or how little SMF data to collect
- TCP/IP statistics are captured in SMF 109, 118, 119

z/OS Communication Server Network Management Interfaces (NMI)

API to obtain information about TCP/IP and VTAM resources:

- Alternative for most SNMP, NETSTAT, and DISPLAY output
- Efficient!

Methods of collection:

- Callable APIs
- Event information
- Real-time packet trace

Most “modern” network management products on z/OS use the NMI to collect data.

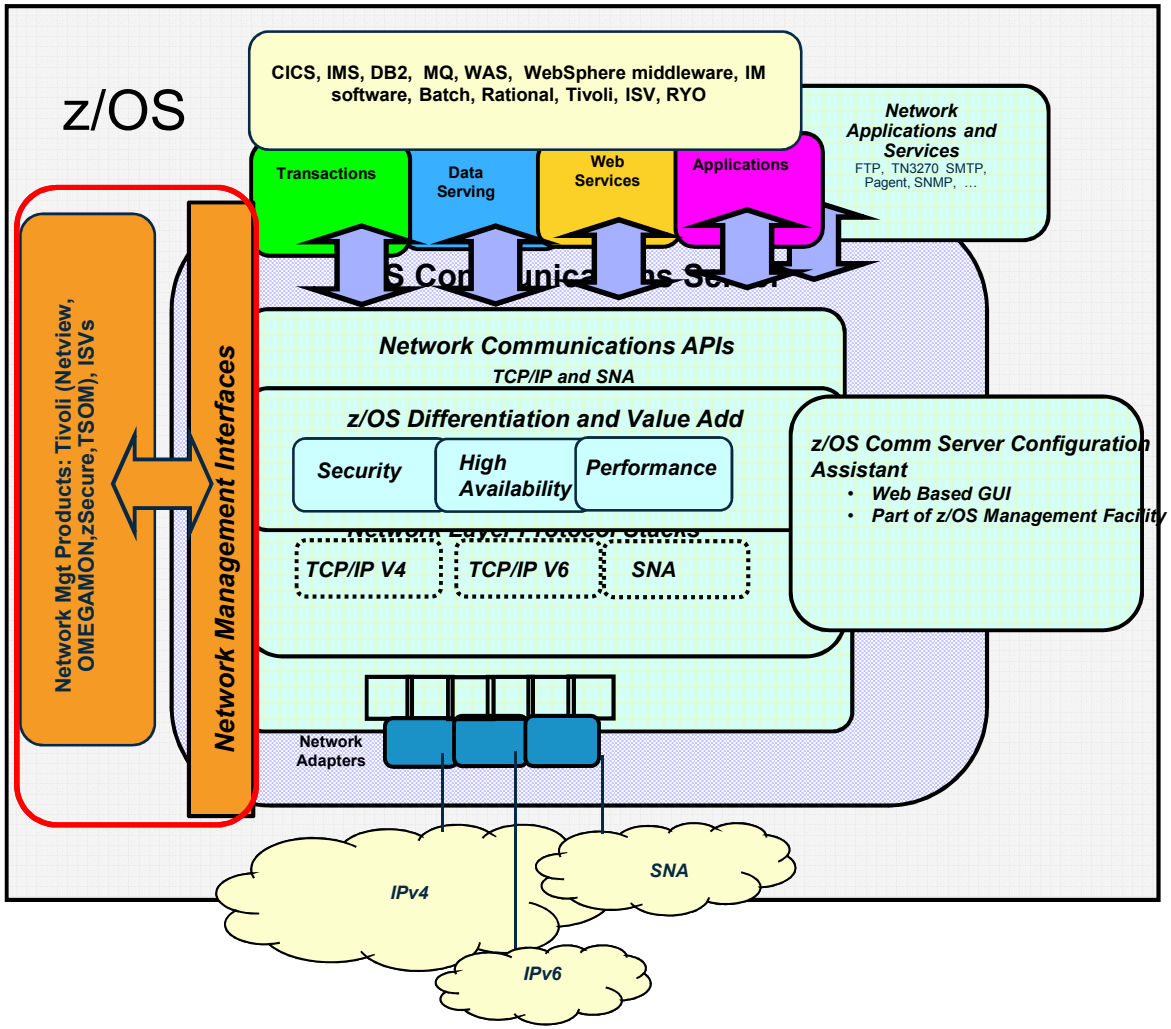

Agenda

- What is network management?
- Why manage the network on z/OS?
- z/OS networking
- Network monitoring tools and technologies
- Best practices

Network management process

Plan / Design:
- Build history
- Baseline
- Trend analysis
- Capacity planning
- Procurement
- Topology design

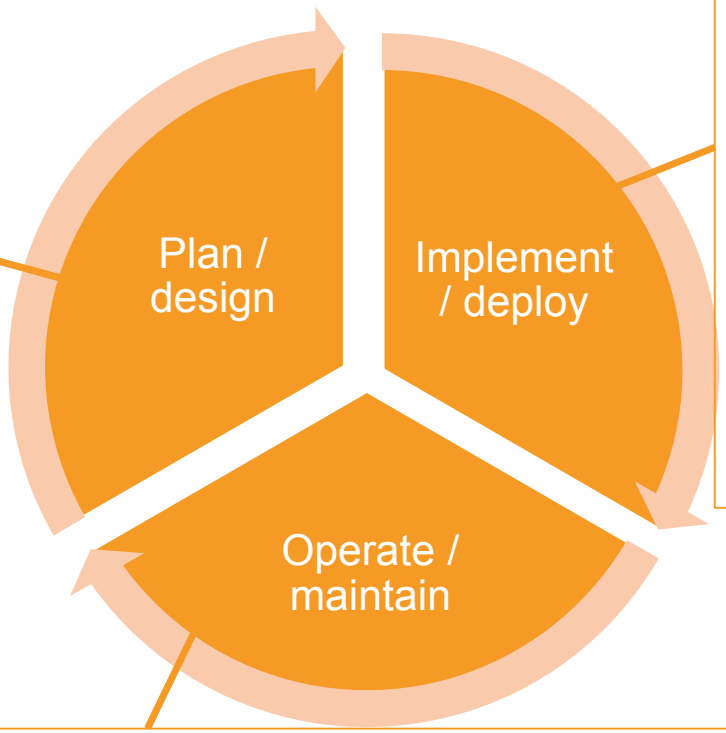

Implement / Deploy:
- Install and configure
- Address management
- Adds, moves, changes
- Security
- Accounting/billing
- Assets/inventory
- User management
- Data management

Operate / Maintain:
- Define thresholds
- Monitor exceptions
- Notify
- Correlate
- Isolate problems
- Troubleshoot
- Bypass/resolve
- Validate and report

Baseline Your Environment

- Become familiar with your network resources. Know what is:
  - Normal
  - Abnormal but uninteresting
  - Abnormal **and** interesting

- Be part of your change control system

- Review periodically!!!

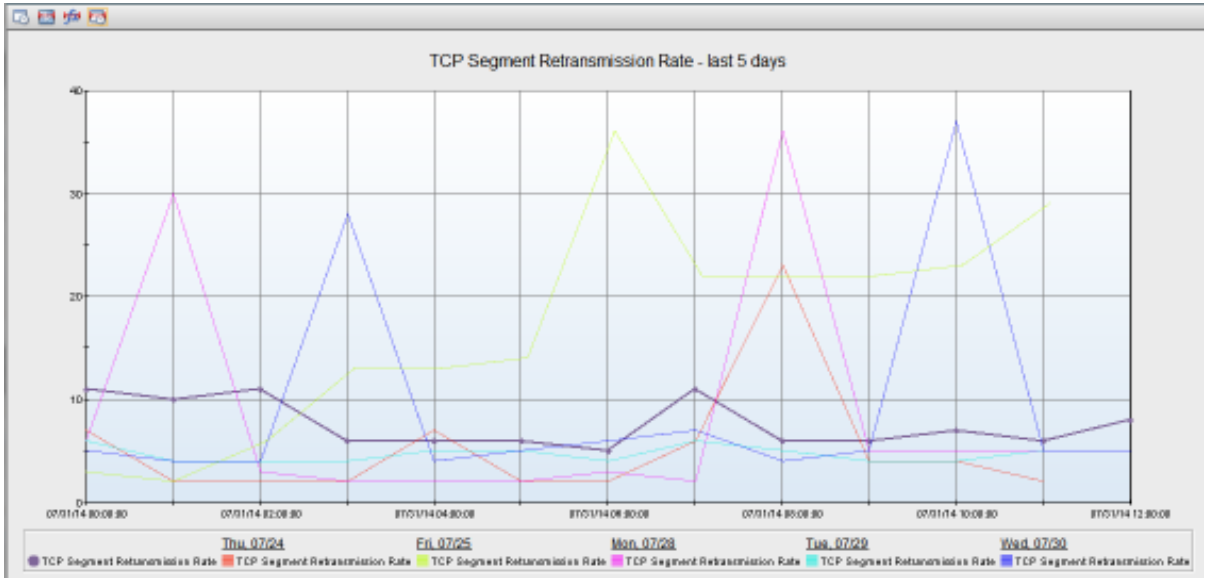

Problem isolation

z/OS Comm Server:

- TCP/IP
  - IP layer
  - TCP layer
  - UDP layer
- TCP/IP address space
- VTAM
  - Buffer pools
  - EE
  - HPR
- VTAM address space

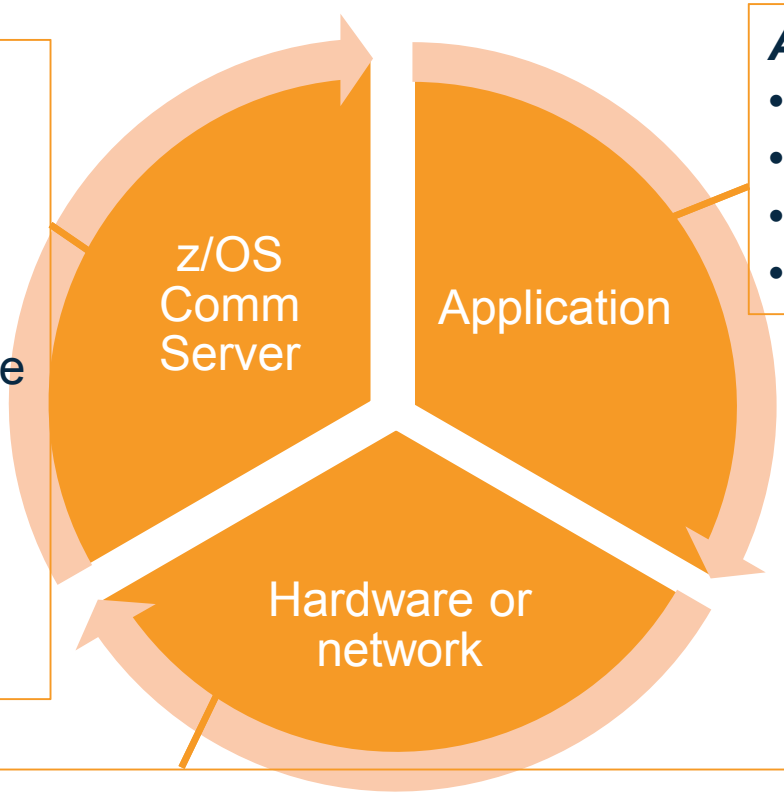

Application:

- Applications
- TCP Listeners
- TCP Connections
- UDP Endpoints

Hardware or network:

- OSA ports
- OSA channel
- HyperSockets
- Interfaces
- Data center network
- Intranet
- Internet

Isolate problem

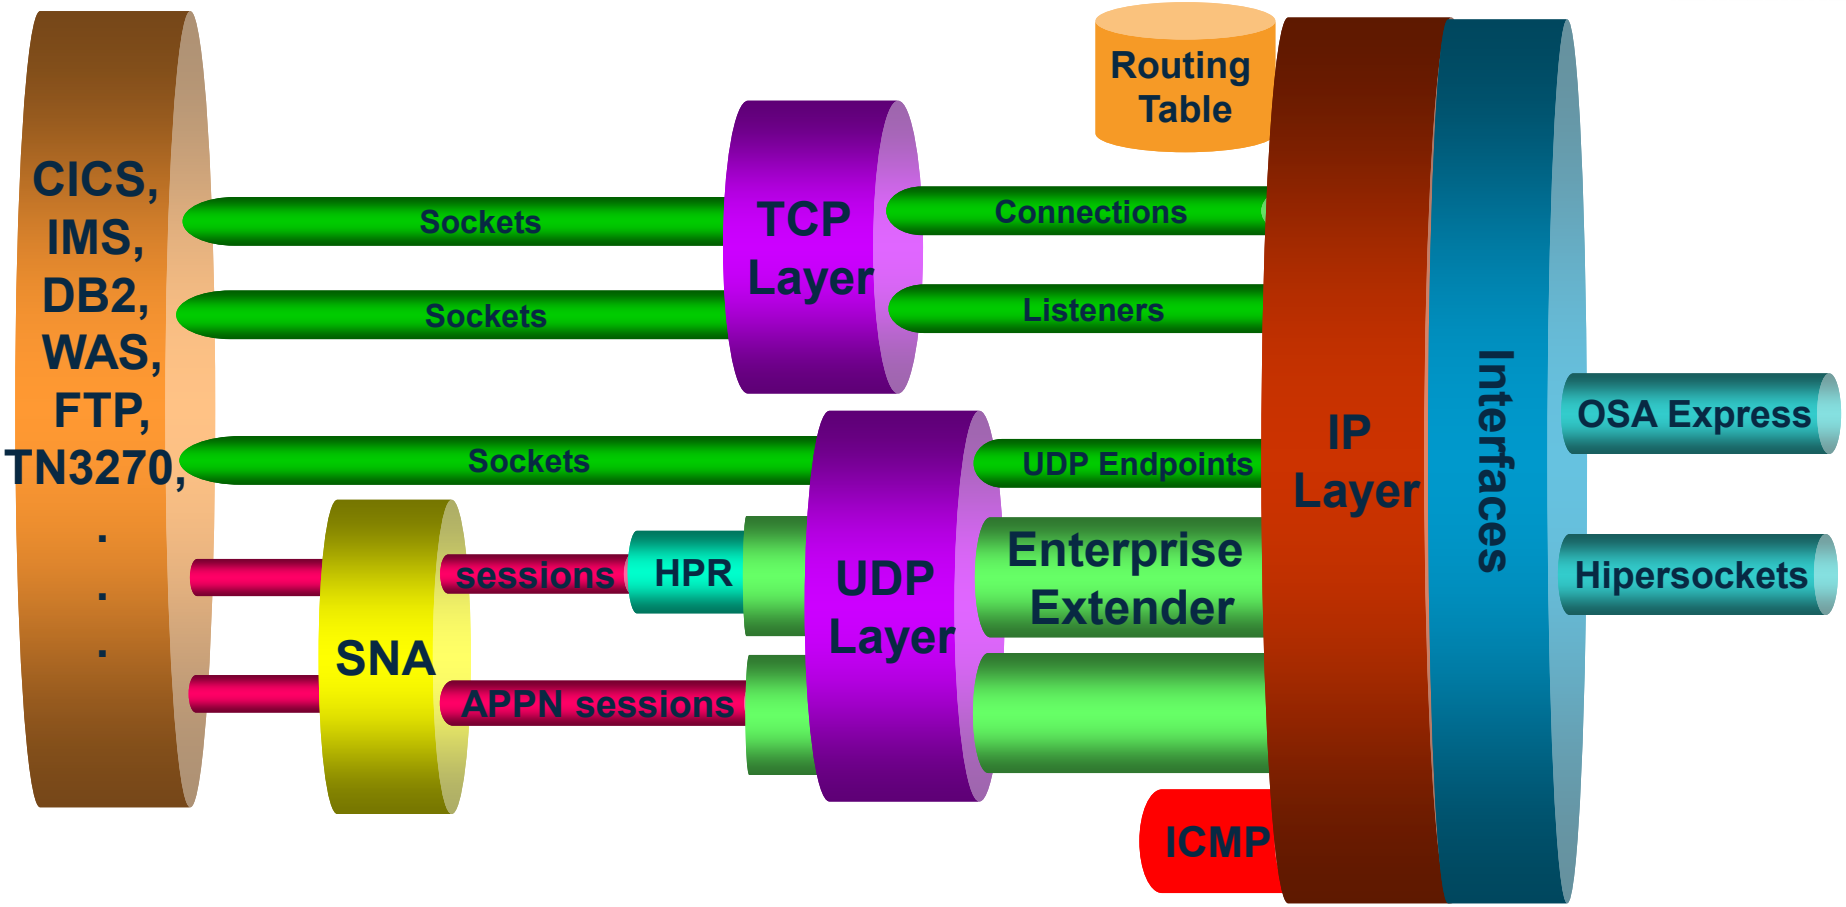

Protocol indicators

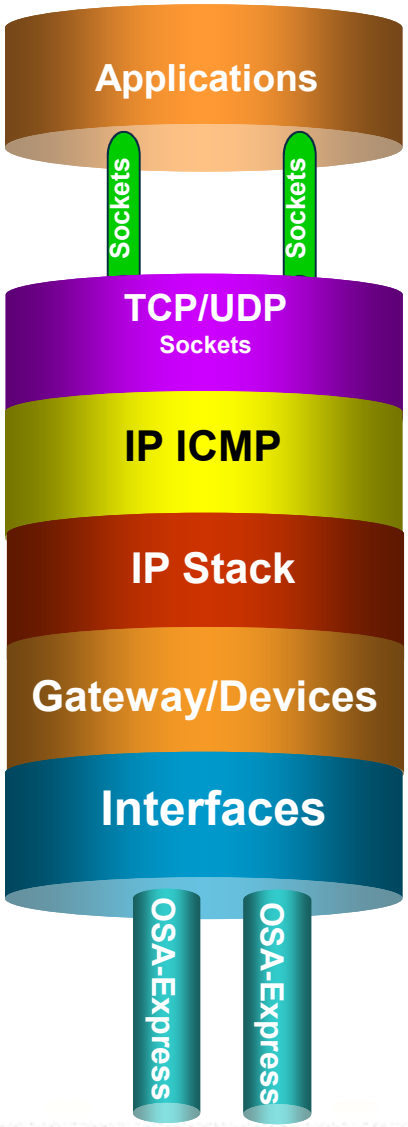

TCP/IP

- Input Discards
- Output Discards
- UDP Discard
- UDP Input Errors
- UDP No Port

OSA

- Fragments, jabber, length error, CRC, alignment
- Unknown IP Frames

Interfaces

- Inbound Packets – discarded, in error
- Outbound Packets – discarded, in error
- Utilization
- Transmission Rates
- Unknown IP Frames

Resource constraint indicators

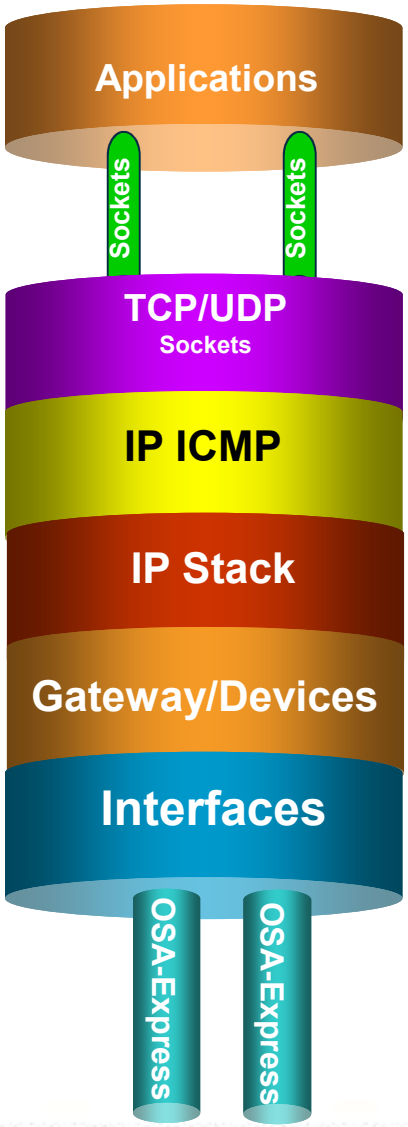

TCP/IP

- CPU Percentage
- CSA – Allocated, In Use
- Authorized Private Storage – Allocated, In Use
- ECSA storage – max, allocated, in use, pools
- Datagrams Discarded
- Backlog Connections Rejected

UDP

- Datagrams Discarded

HPR

- Throughput rate – allowed
- Unacknowledged Buffers – high water mark

Interfaces

- Receive/Transmit Bandwidth utilization

OSA

- PCI Utilization
- Processor Utilization
- Missed Packets

VTAM

- CPU Percentage
- CSA – Allocated, Allowed, In Use

Indirect indicators

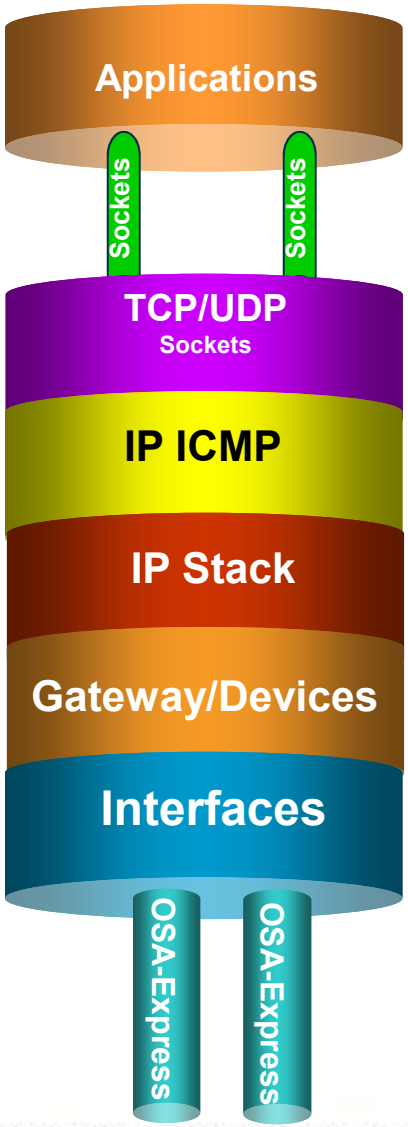

TCP

- Segments Retransmitted
- Response Time
- Fragmentation and/or reassembly
- Out of Order Segments
- Remote Window Size Frequency
- TCP Keep-Alive Drops

HPR

- Out of Sequence Buffers
- Packet Retransmission Rate
- Path Switches
- Response Time Variance
- Smoothed Round Trip Time

TN3270

- Average IP Response Time and variance
- Average SNA Response Time and variance

Scenario: User reports batch FTP failures

The setting:

A company relies on batch FTP to copy files between a mainframe at headquarters and each of its retail stores every night (local store time). Sales and inventory data is uploaded and product and pricing changes are downloaded to the stores. One morning, a systems administrator notices that some of the files have not been updated for days. He reports the problem to the IT help desk. The problem is routed to the mainframe networks systems programmer.

User reports batch FTP failures

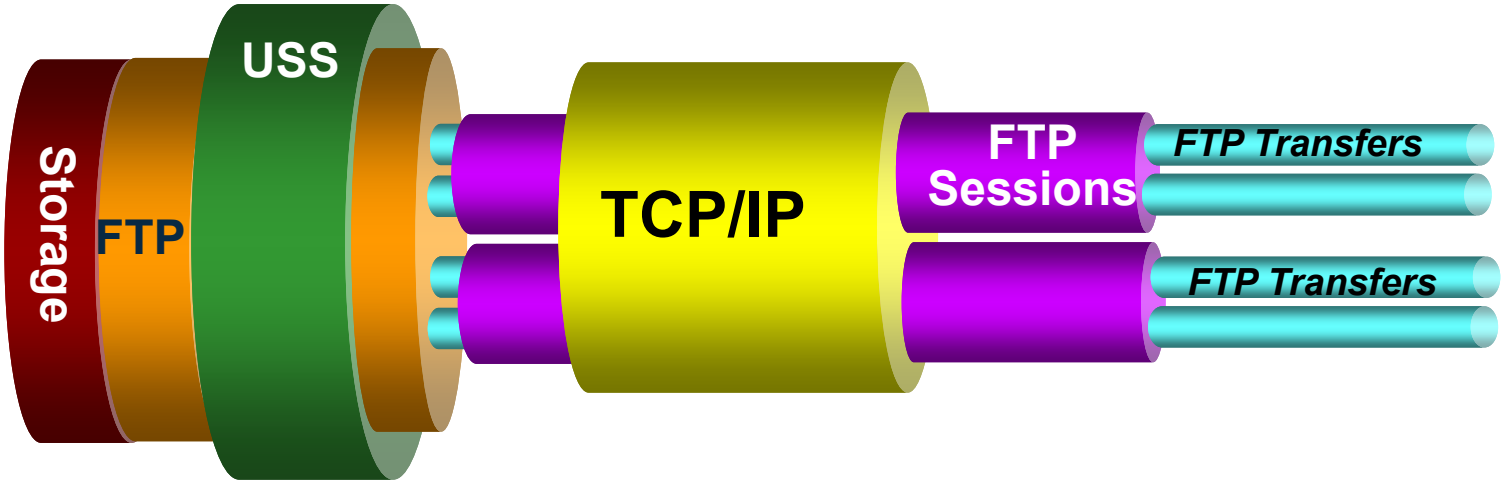

- Start with checking current activity: FTP transfers & FTP sessions

FTP Session Summary Table

Collection Time	Application Name	FTP Type	Remote IP Address	Remote Port	Local IP Address	Local Port	User ID on Server	Client User ID	Session Start	Session End	Session Duration
08/11/08 23:52:22	FTPD1	Server	9.65.126.164	3000	9.42.45.179	21	USER2		08/11/08 23:52:18		0

FTP Transfer Summary Table

Collection Time	Remote IP Address	Remote IP Port	Local IP Address	Local IP Port	User ID on Server	Client User ID	Role	Transmission Start	Transmission End	Transmission Duration	Bytes Transmitted (in GB)	Bytes Transmitted
08/11/08 23:53:03	9.65.126.164	3006	9.42.45.179	20	USER2		Server	08/11/08 23:52:49	08/11/08 23:53:01	11420	0	1440054
08/11/08 23:53:39	9.65.126.164	3010	9.42.45.179	20	USER2		Server	08/11/08 23:53:39		0	0	0

Check TCP/IP Stack

- Retransmits – network congestion
- Out of Order – routing issue or network congestion
- Fragmentation – MTU size
- Discards – resource constraints
- Timeouts – connectivity
- UDP input errors – attack
- UDP discards – no application running

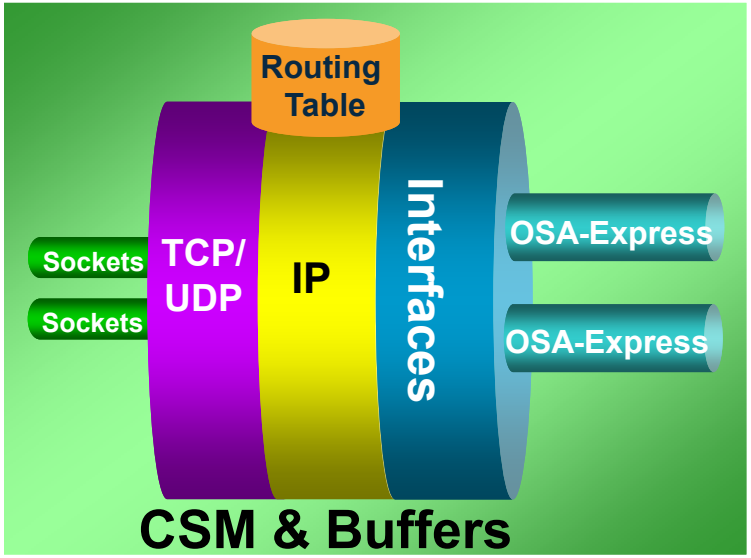

- High storage utilization - could indicate network congestion
  - This can result in requests backing up in storage

TCPIP Stack Layers Summary Table

utput Routes	Reassembly Count	Reassembly Percentage	Reassembly Failure Count	Reassembly Failure Percentage	Fragmentation Count	Fragments To Be Reassembled	Fragmenta Percentag
0	6453	0	256	2	0	0	

Check OSA and Interfaces

OSA

- Online Status
- Configuration
- Microcode Level
- Utilization
- Transmission Rates
- Unknown IP Frames
- By LPARS
- By Ports

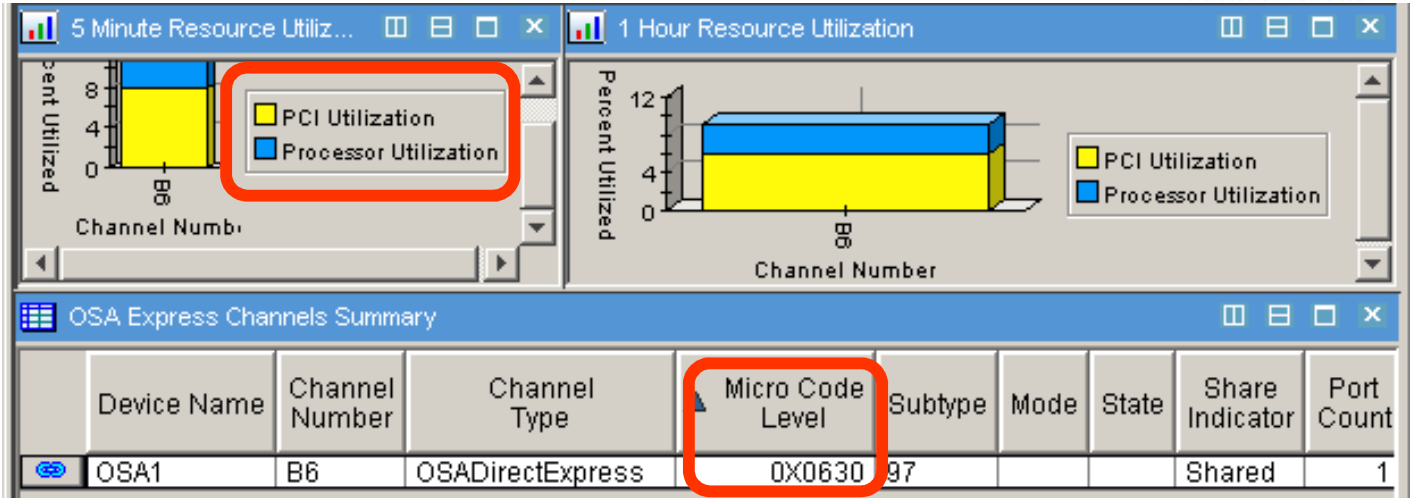

Device Name	Channel Number	Channel Type	Micro Code Level	Subtype	Mode	State	Share Indicator	Port Count
OSA1	B6	OSADirectExpress	0X0630	97			Shared	1

Interfaces

- Packet Errors
- Bandwidth Utilization
- MTU Size

Interface Name	Description	Interface Type	Current State	MTU Size	Transmit Packet Rate	Receive Packet Rate
TCPIPLINK	IP Assist QDIO Ethernet	ethernetCsmacd	Up	1492	4312	74909
LOOPBACK	Loopback	softwareLoopback	Up	65535	890	890
LOOPBACK	Loopback Device	propVirtual	Up	0	890	890
OSA1	Multipath Channel IP Assist Device	propVirtual	Up	0	4312	74909
EZAXCFSA	Multipath Channel Point-to-Point	mpc	Down	55296	0	0

Check Applications and Connections

Applications:
- Accepting connections?
  - Rate, Backlog, Rejections
- Last activity time
- Response time
- Retransmissions
- Transmit / Receive Rates
- Out of order segments

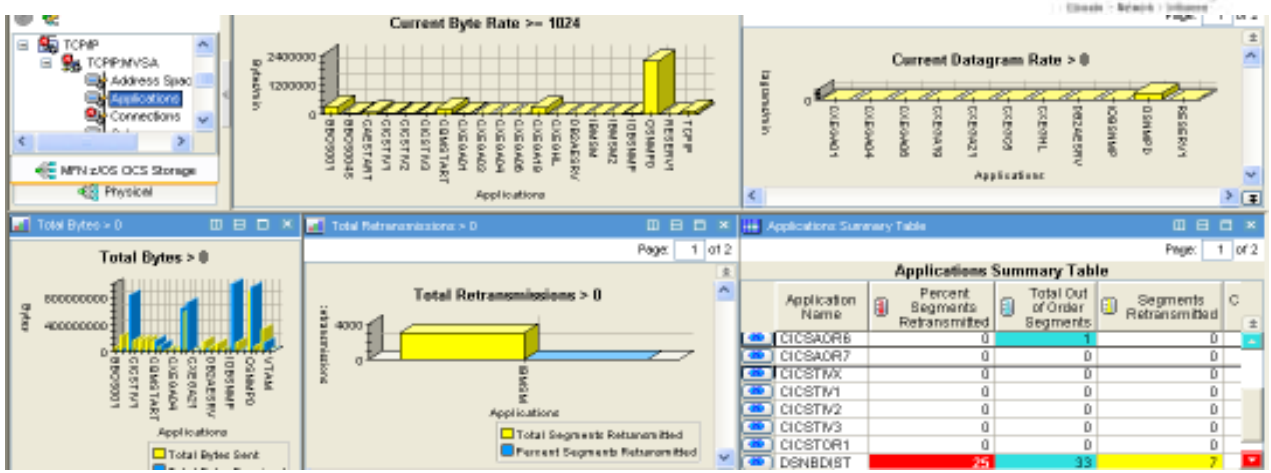

Connections:
- Start time/duration
- Response Time
- Retransmissions
- Transmit / Receive Rates
- Out of order segments

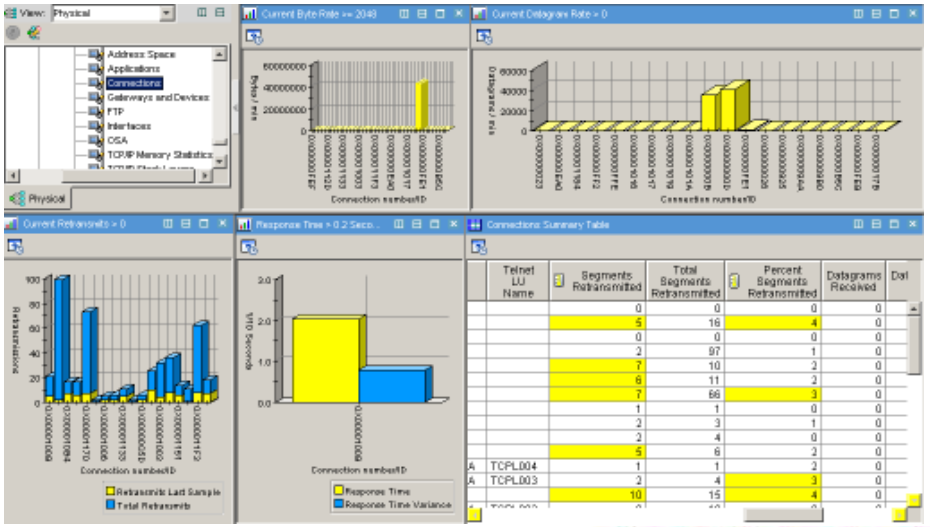

Network is OK, then where is the problem?

Check job logs, SYSLOG, and syslogd for clues.

Dataset access issue? FTP session login failure?

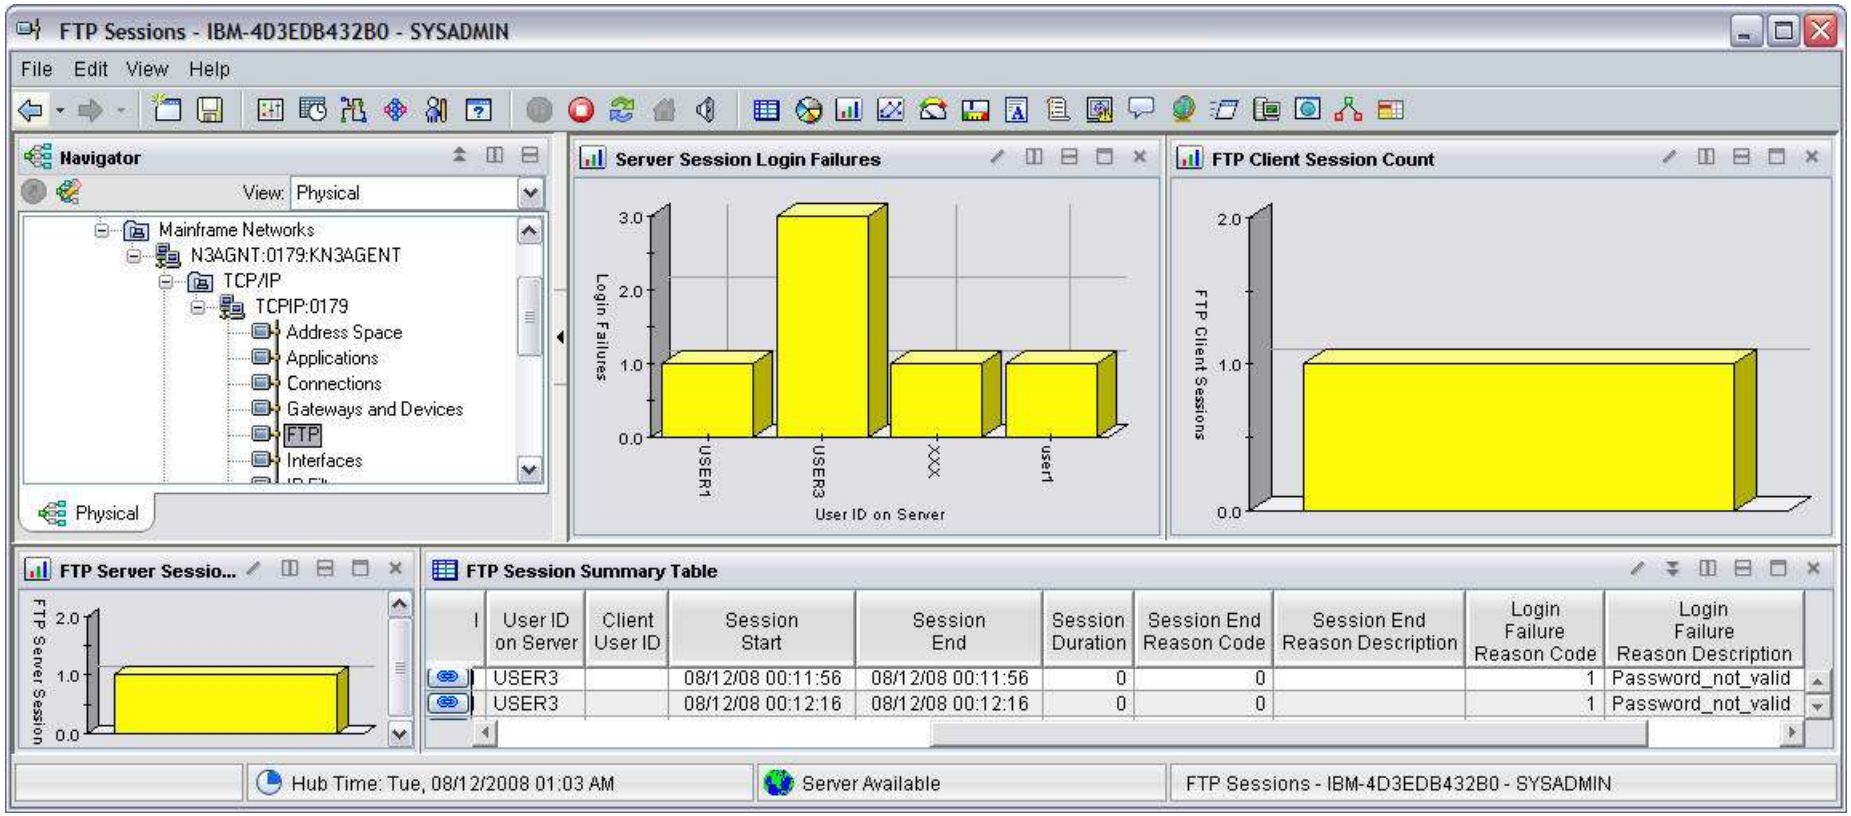

Introduction to Mainframe (z/OS) Network Management

Session 17736

Dean Butler (butlerde@us.ibm.com)