IBM Software Group
Enterprise Networking Solutions (ENS) and Transaction Processing Facility (TPF)

z/OS Communications Server V1R9 Technical Overview August 2007

Alfred B Christensen - alfredch@us.ibm.com
Sam Reynolds - samr@us.ibm.com

August 2007

IBM Systems

Trademarks and notices

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:

- Advanced Peer-to-Peer Networking®
- AIX®
- alphaWorks®
- AnyNet®
- AS/400®
- BladeCenter®
- Candle®
- CICS®
- DB2 Connect
- DB2®
- DRDA®
- e-business on demand®
- e-business (logo)
- e business(logo)®
- ESCON®
- FICON®
- GDDM®
- HiperSockets
- HPR Channel Connectivity
- HyperSwap
- i5/OS (logo)
- i5/OS®
- IBM (logo)®
- IBM®
- IMS
- IP PrintWay
- IPDS
- iSeries
- LANDP®
- Language Environment®
- MQSeries®
- MVS
- NetView®
- OMEGAMON®
- Open Power
- OpenPower
- Operating System/2®
- Operating System/400®
- OS/2®
- OS/390®
- OS/400®
- Parallel Sysplex®
- PR/SM
- pSeries®
- RACF®
- Rational Suite®
- Rational®
- Redbooks
- Redbooks (logo)
- Sysplex Timer®
- System i5
- System p5
- System x
- System z
- System z9
- Tivoli (logo)®
- Tivoli®
- VTAM®
- WebSphere®
- xSeries®
- z9
- zSeries®
- z/Architecture
- z/OS®
- z/VM®
- z/VSE

- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. 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.
- Intel, Intel Inside (logos), MMX and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Linux is a trademark of Linus Torvalds in the United States, other countries, or both.
- Red Hat is a trademark of Red Hat, Inc.
- SUSE® LINUX Professional 9.2 from Novell®
- Other company, product, or service names may be trademarks or service marks of others.
- This information is for planning purposes only. The information herein is subject to change before the products described become generally available.
- All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All performance data contained in this publication was obtained in the specific operating environment and under the conditions described and is presented as an illustration. Performance obtained in other operating environments may vary and customers should conduct their own testing.

Refer to www.ibm.com/legal/us for further legal information.

Agenda

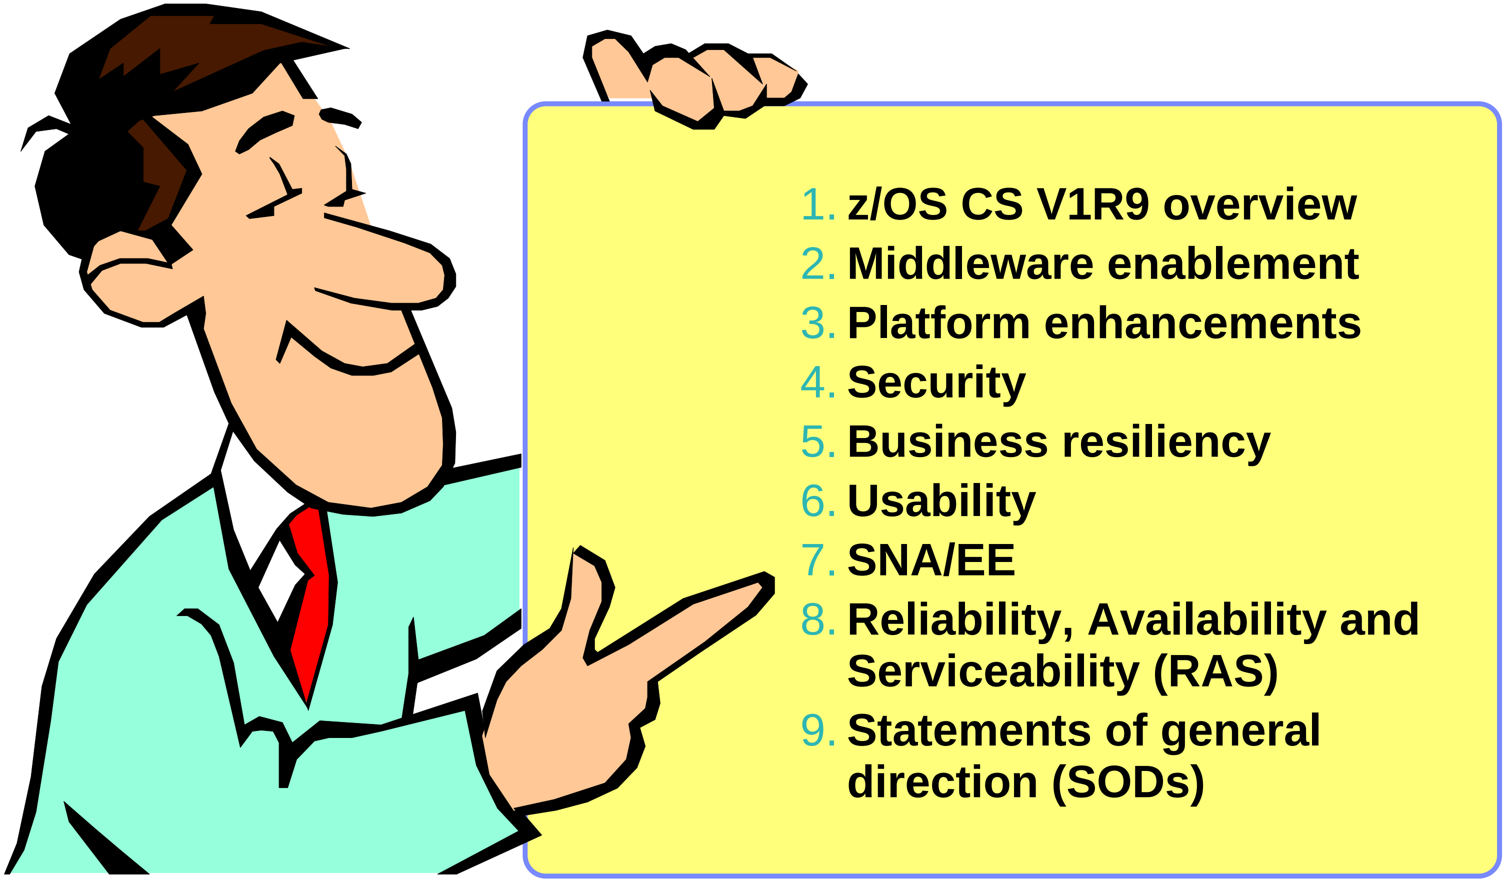

1. **z/OS CS V1R9 overview**
2. **Middleware enablement**
3. **Platform enhancements**
4. **Security**
5. **Business resiliency**
6. **Usability**
7. **SNA/EE**
8. **Reliability, Availability and Serviceability (RAS)**
9. **Statements of general direction (SODs)**

z/OS Communications Server - Disclaimer

- Plans for the z/OS Communications server are subject to change prior to general availability
- Information provided in this presentation may not reflect what is actually shipped for z/OS Communications Server
- This presentation includes an early overview of selected future z/OS Communications Server enhancements
- The focus of this presentation is the Communications Server in z/OS V1R9

Note: Plans are subject to change!

Plans may change before GA of z/OS CS V1R9!

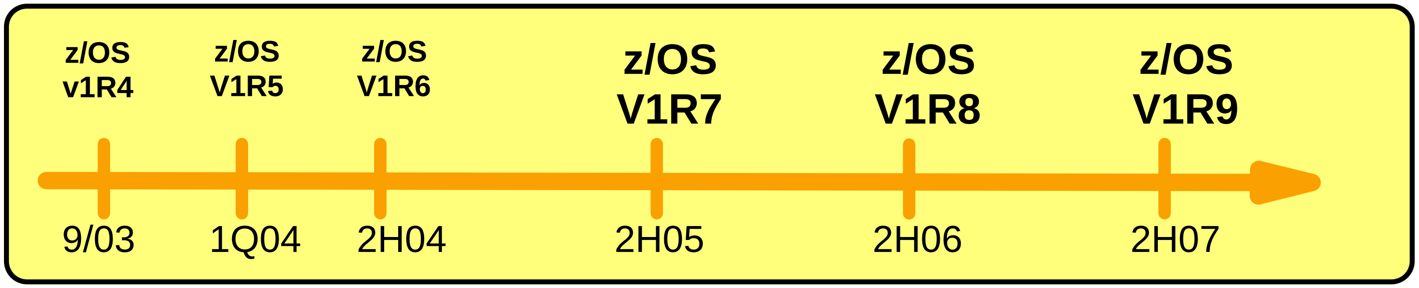

z/OS v1R4 – 9/03

z/OS V1R5 – 1Q04

z/OS V1R6 – 2H04

z/OS V1R7 – 2H05

z/OS V1R8 – 2H06

z/OS V1R9 – 2H07

z/OS CS V1R9 overview - part 1 of 3

- **Middleware enablement**
  - New SNMP manager programming interface
  - CICS sockets enhancements
  - Display TN3270 client code page
- **Platform enhancements**
  - Policy-based routing (PBR)
  - RFC currency
    - MLDv2 and IGMPv3 (source multicast filters)
    - FTP SSL/TLS support
    - IPv6 scoped address architecture API
  - FTP Unicode support
  - Policy agent retrieving centralized policies
  - Remove support for running the TN3270 server in the TCP/IP address space
- **Security**
  - IPsec network management interface (NMI) support
  - IPsec network security services (NSS)
  - AT-TLS enablement of CS-provided servers
    - AT-TLS enable the TN3270 server
    - AT-TLS enable the FTP client and server
  - FTP Kerberos single sign-on support

z/OS CS V1R9 overview - part 2 of 3

- **Business resiliency**
  - Dynamic VIPA usability enhancements
  - Support for distributed dynamic VIPA on source IP destination rules
  - TCP/IP Sysplex Dynamic XCF usability enhancements
  - Remove LDAPv2 schema support for QoS and IDS policies
  - Support for new WLM routing services
  - Support for additional distribution algorithms in Sysplex Distributor
  - Support for a port range on the VIPADISTRIBUTE statement

- **Usability**
  - Configuration Assistant
    - Policy based routing GUI configuration interface
    - Network security services GUI configuration interface
  - Allow FTP client to select source IP address
  - Allow ping to discover network MTU sizes

z/OS CS V1R9 overview - part 3 of 3

- **SNA/EE**
  - Local MTU Discovery for Enterprise Extender
  - HPR path switch enhancements
  - Generic resource enhancements
  - Enterprise Extender LDLC timers
  - Display enhancements
  - VTAM recovery of FICON CTC links
  - Request for a change in TGP capability value maximum
  - Manage the number of APPN Locate replies that may be sent
  - Remove ASUITE functions from Communications Server

- **Reliability/Availability/Serviceability**
  - Enhancements for the z/OS Health-checker support
  - Specify port number on packet trace command
  - Various RAS items

Middleware enablement

CICS Sockets - increased robustness and management

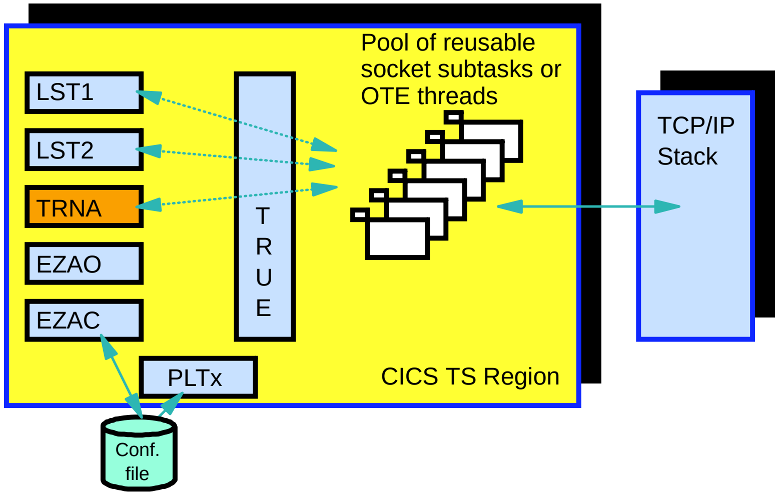

- **The CICS Sockets listeners (two flavors) are part of CICS Sockets**
  - A listener task accepts client connections and starts CICS transactions
- **If TCP/IP goes down, the listener task today terminates**
  - Operations personel or automation need to restart them when TCP/IP comes back up
- **Listeners will in z/OS V1R9 be enhanced to stay up when TCP/IP goes down and keep retrying establishing contact again until successful or an operator manually shuts the listener down.**

- **Additional changes to CICS Sockets planned for z/OS V1R9:**
  - Better error-handling by the listeners when MAXFILEPROC is reached by the CICS address space
    - Today the listener terminates on this condition
    - V1R9 will allow the listener to stay up, stop accepting new connections for a period, and then resume normal processing
  - An option for more efficient processing during CICS shutdown to allow for an immediate termination of the CICS Sockets infrastructure instead of a deferred termination
  - Ship sample ASCII/EBCDIC translation routine that can be used for single byte translation for all the languages that are covered by the hlq.SEZATCPX translation tables
    - One of the call parameters is the table name
  - Add netstat visibility into CICS Sockets information
    - CICS transaction code
    - CICS socket state (listener, passed-but-not-taken, processing in child server, etc.)

New SNMP local manager programming interface for more comprehensive network management application support

- **The osnmp command is a user interface to the SNMP protocol**
  - Uses a built-in set of functions to build, send, and receive SNMP packets
- **We are in this release implementing an open programming interface that will allow a z/OS-resident SNMP manager application to interface programmatically to the SNMP agent (OSNMPD) exchanging SNMP protocol packets over a local UDP socket**
- **Local management applications that need frequent access to high-volume management data should use the NMI interfaces:**
  - NMI provides low-overhead access to high-volume management data
  - Typically addresses performance monitoring applications
- **Local management applications that need access to SNMP management data that is not provided via NMI can use this new API:**
  - Resource monitoring
  - Availability management
  - Operations (via SNMP SET)

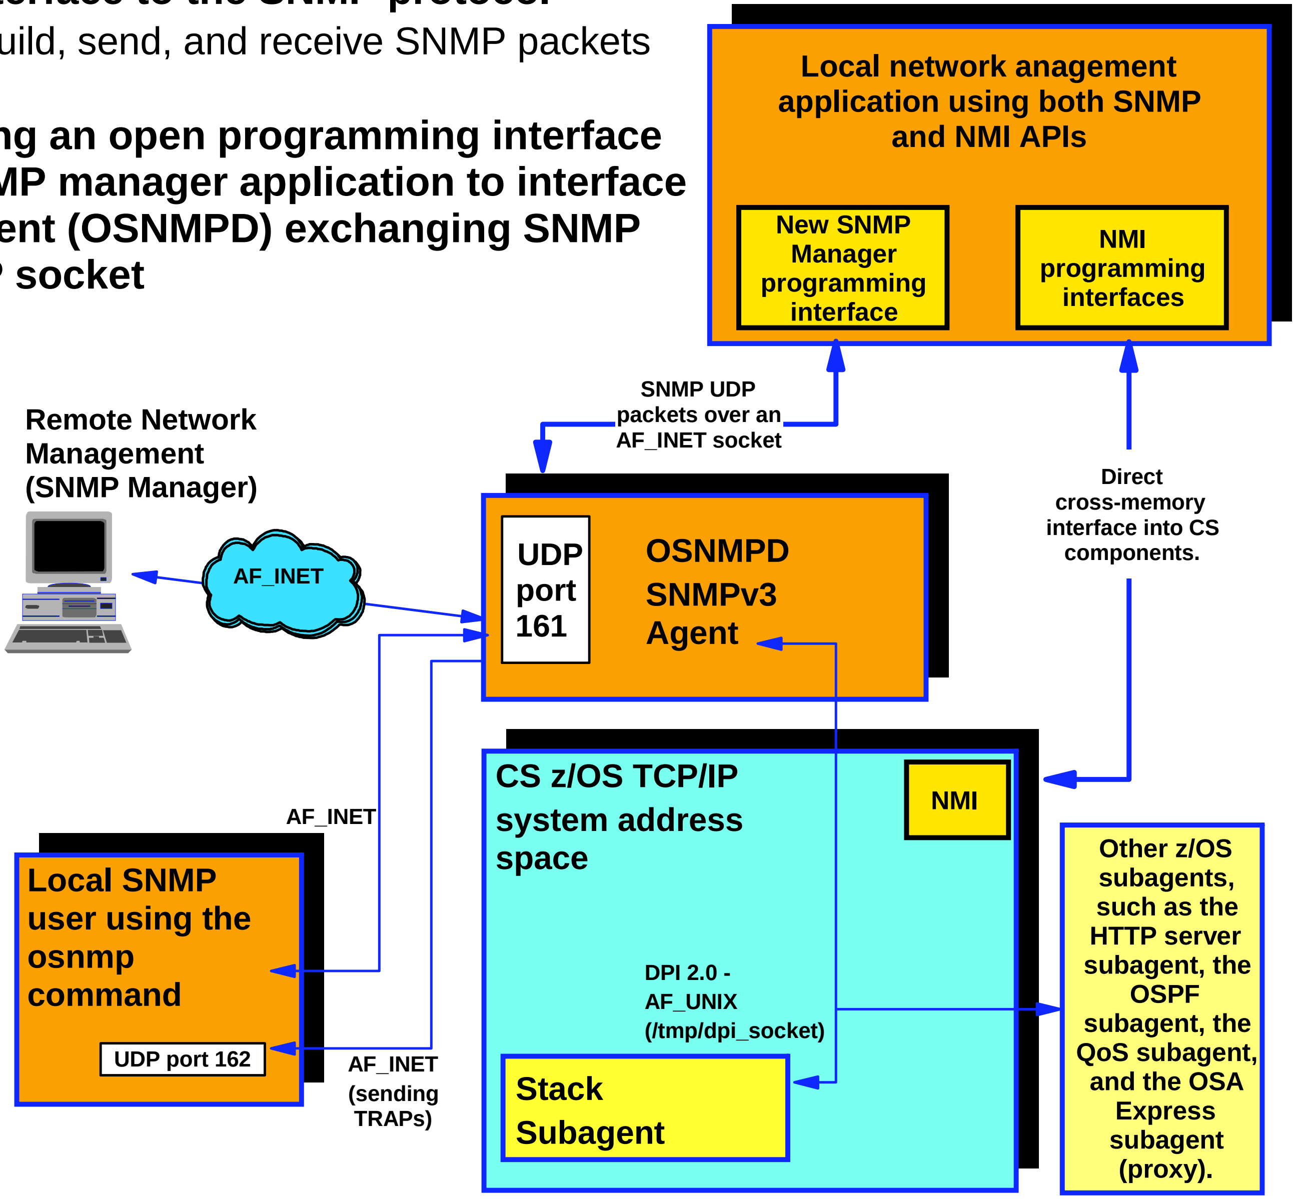

Local network agement application using both SNMP and NMI APIs

New SNMP Manager programming interface

NMI programming interfaces

SNMP UDP packets over an AF_INET socket

Direct cross-memory interface into CS components.

Remote Network Management (SNMP Manager)

AF_INET

UDP port 161

OSNMPD SNMPv3 Agent

CS z/OS TCP/IP system address space

NMI

DPI 2.0 - AF_UNIX (/tmp/dpi_socket)

Stack Subagent

Other z/OS subagents, such as the HTTP server subagent, the OSPF subagent, the QoS subagent, and the OSA Express subagent (proxy).

AF_INET

Local SNMP user using the osnmp command

UDP port 162

AF_INET (sending TRAPs)

Display TN3270 client-configured code page to help avoid DB2 code page inconsistencies

- **DB2 has experienced several problems with mis-configured TN3270 clients that have resulted in DB2 data corruption problems.**
  - The problems have occured in scenarios where TSO users that use TN3270 to connect to z/OS have a code page configured for their emulator session that is not compatible with the code page of the data that is stored in DB2.
  - In these scenarios, these users have attempted to edit DB2 data (using an editor or other tool) under TSO, resulting in the corruption of some of the data (that couldn't properly be displayed using the configured user code page).
  - DB2 has requested that VTAM provide a display that shows the user's configured code page so that a customer has a means of attempting to identify mis-configured clients.
    - Note that this information can only be provided for sessions that chose to use a Logmode entry that specifies Extended Data Streams (EDS) to establish the session. However, it is believed that over 95% of all TN3270 sessions use EDS.

- **Code page and character set ID information will be represented using the standard Coded Graphic Character Set Global Identifiers (CGCSGID)**
  - TSO/VTAM will allow users to query which CGCSGID a TSOUSER is using.
  - Provide the CGCSGID information in a TSO/VTAM logon exit routine
  - Provide the CGCSGID information in a new message in an existing message group when a VTAM TSOUSER display is requested.

The codepage to use is configured on the emulator workstation, not the TN3270 server node.

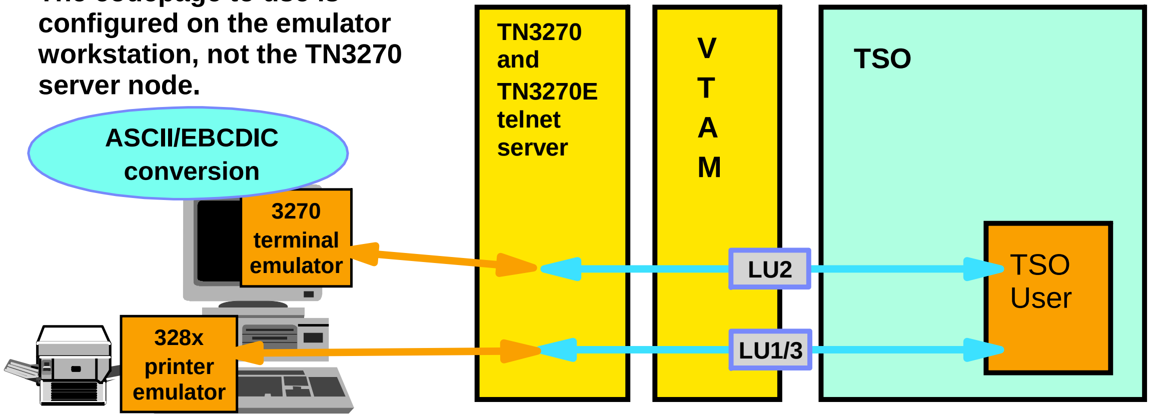

Please see *"3174 Character Set Reference"*, GA27-3831 chapter 5 for details on CGCSGID values.

Platform enhancements

System z9 - recent OSA-related enhancements

- **Selected OSA basics incl. recent new functions**
  - Layer-3 virtual MAC
    - System z9 with OSA-Express or OSA-Express2 and z/OS V1R8 with PTFs
  - Network Traffic Analyzer
    - System z9 with OSA-Express2 and z/OS V1R8 with PTFs
  - LAN Idle timer
    - System z9 with OSA-Express2 and z/OS V1R8 with PTFs
  - QDIO diagnostic synchronization
    - System z9 with OSA-Express2 and z/OS V1R8 with PTFs

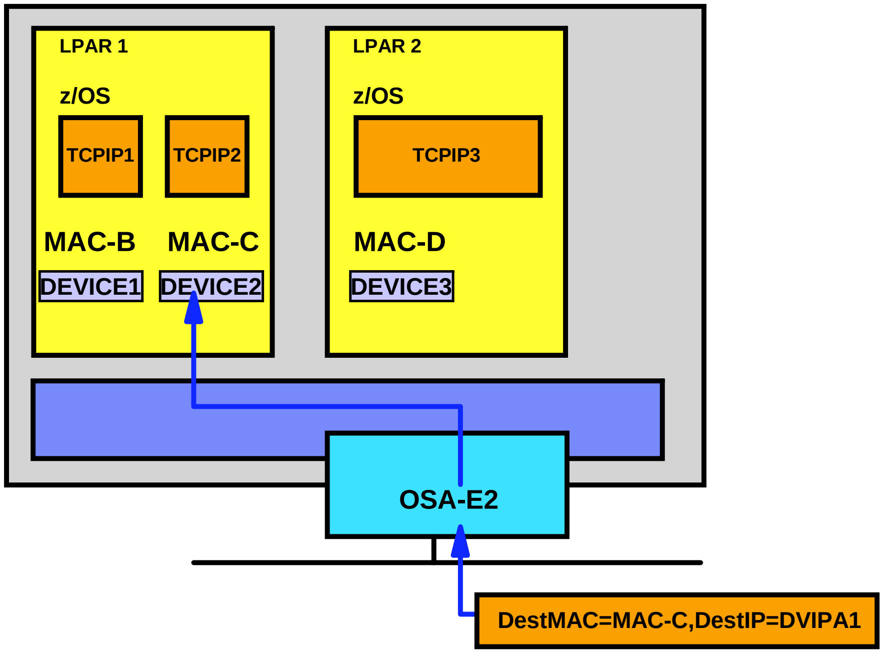

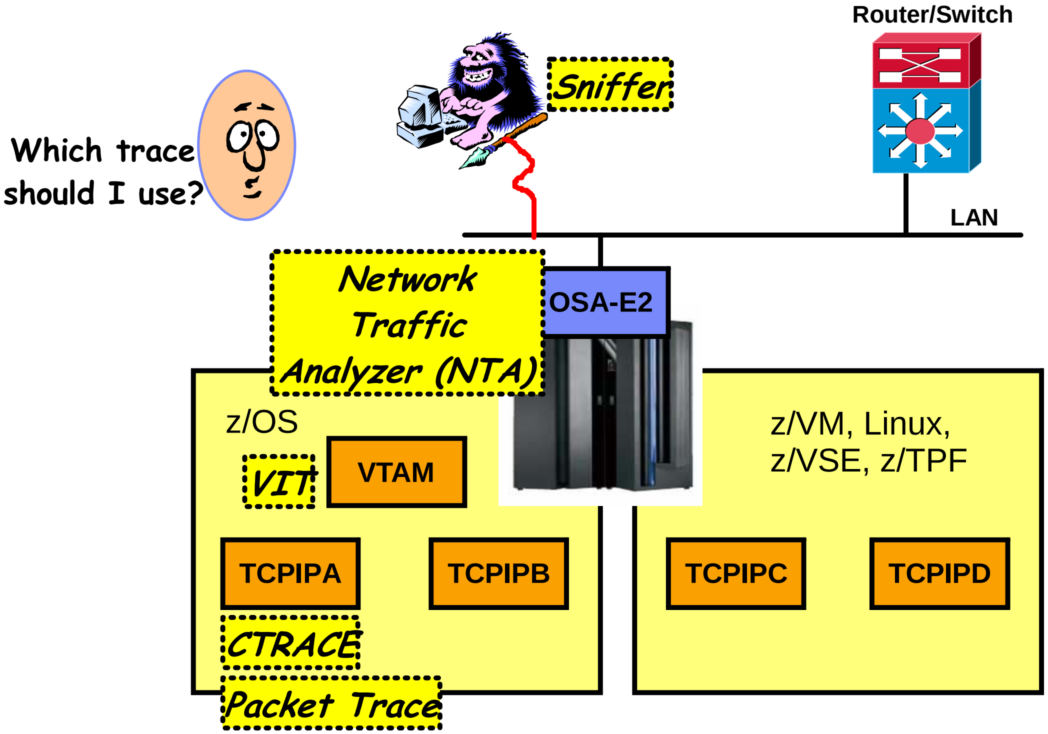

Bring FTP up to latest IETF standards level for SSL/TLS support for improved interoperability with other platforms

- **FTP was originally enabled for SSL/TLS back in z/OS V1R2**
  - Based on a draft RFC that described how the FTP protocol were to work with SSL/TLS

- **That draft RFC has since that time undergone several revisions and has now made it into official RFC status**
  - RFC 4217 "Securing FTP with TLS"

- **A few changes have been made since the draft RFC version that we used in z/OS V1R2 was written:**
  - Change USER command reply code from 232 to 230 if a password is not required
  - Allow a REIN (Re-initialize) command on a control connection that is secured through an AUTH command
    - Currently, FTP resets everything on the control connection, except the TLS status when a REIN is processed
    - RFC 4217 says that also the TLS state must be cleared when this command is processed
  - Allow an AUTH command on a control connection that is already secured through an AUTH command
    - Currently, FTP rejects such a command in this situation
    - RFC 4217 says that such an AUTH command in a sense is to be treated as a REIN command - re-initializing the control connection including the TLS state
  - Stop using out-of-band data when connections are secured
    - RFC 4217 requires all commands (also ABOR, STAT, and PROXY) to be sent over the TLS connection and not out-of-band

```
                  +--DRAFT----+
                  |           |
>--TLSRFCVERSION--+-----------+---><
                  |           |
                  +--RFC4217--+
```

Default is to use the current default (DRAFT). To use the RFC4217 level, this option must be specified in the FTP client FTP.DATA or set by a LOCSITE command for the client - and in FTP.DATA for the server (no SITE command support).

Policy-based outbound routing of traffic that originates on z/OS

- **What does Policy Based Routing (PBR) do?**
  - Choose first hop router, outbound network interface (including VLAN), and MTU
  - Choice can be based on more than the usual destination IP address/subnet
    - With PBR, the choice can be based on source/destination IP addresses, source/destination ports, TCP/UDP, etc.
- **Allows an installation to separate outbound traffic for specific applications to specific network interfaces and first-hop routers:**
  - Security related
  - Choice of network provider
  - Isolation of certain applications
    - EE traffic over one interface
    - TN3270 traffic over another interface
  - PBR policies will identify one or more routes to use
    - If none of the routes are available, options to use any available route or to discard the traffic will be provided

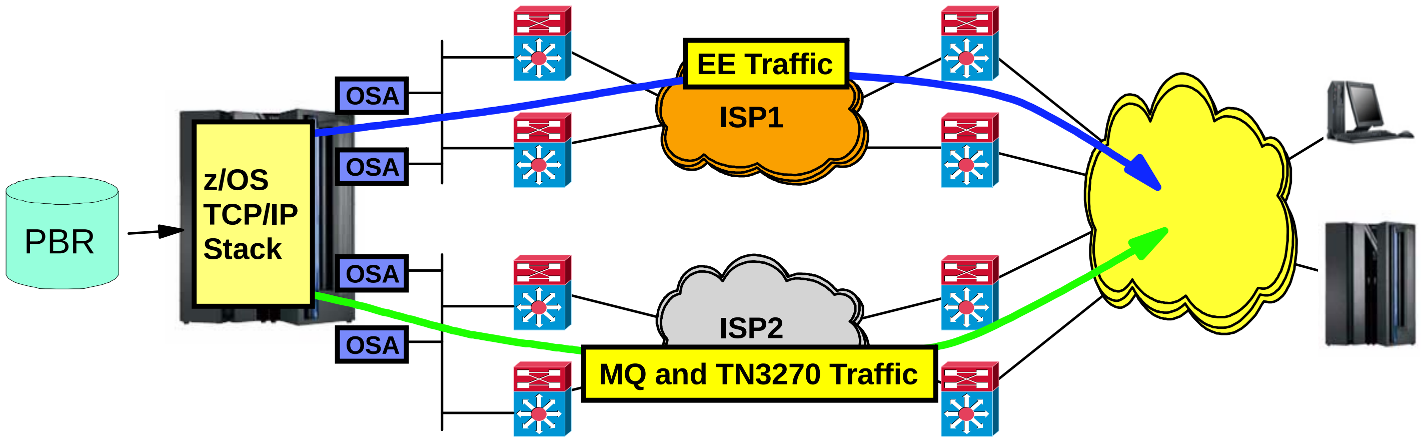

PBR technologies are a great companion to VLAN technologies for separation of traffic over different networks or network providers.

Policy-based routing - staging of the function

- **PBR Policy example**
  - EE traffic to 10.2.0.0/16 to use OSA1 and R4
  - All other traffic to 10.2.0.0/16 to use OSA2 and R5
- **What if:**
  - The route through R4 to subnet 10.2.0.0/16 isn't working?
    - How to determine if it isn't working?
      - Consult dynamic routing information?
        - Adjacencies would do for first hop router
        - Full topology (not just best route topology) will be needed to determine availability of full routing path
      - Decide based on (semi)-static information: static routing table, TCP retransmits, ICMP unreachable messages, and ICMP redirect messages?
        - What about UDP and RAW?
    - What to do if it isn't working?
      - Drop the packets?
      - Revert to normal routing table lookup (use any available route to the destination)?

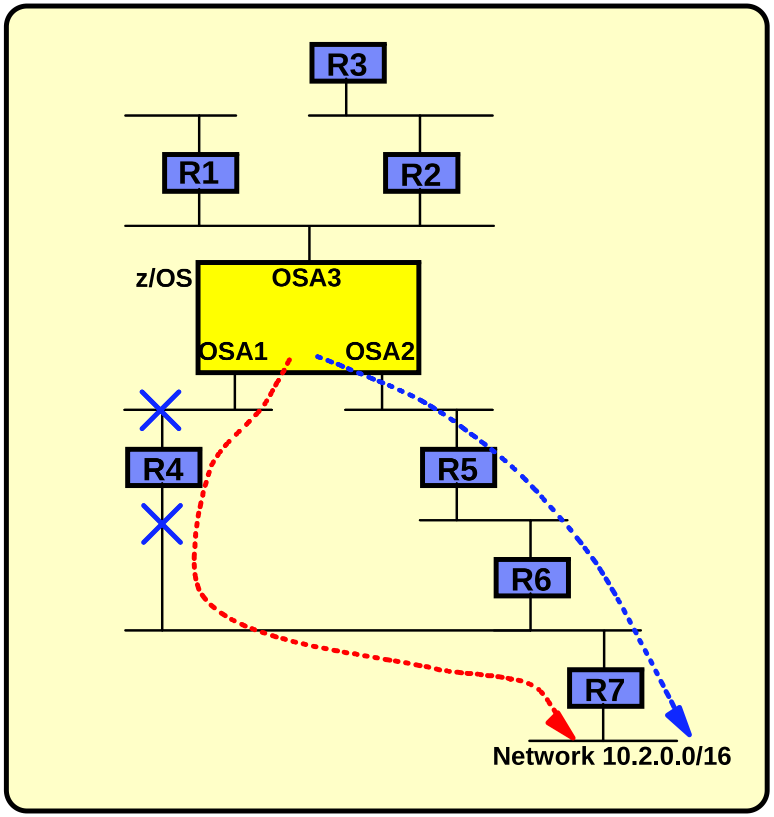

- **PBR will be implemented to work in conjunction with dynamic routing to determine if a route is available or not**
  - PBR support in z/OS CS V1R9 is limited to IPv4 destinations.
  - PBR will in a follow-on release be extended to also support IPv6 destinations.

Multicast currency: IGMPv3 and MLDv2 support - better control over who the multicast sender is from where we will accept multicasts

- **Applications that want to receive multicasts, register with IGMP (for IPv4) or with MLD (for IPv6) using specific sockets API functions.**
  - Multicast functions are based on UDP only.
- **IGMP and MLD both recently have been enhanced to support new capabilities:**
  - Allow a receiving multicast application to specify a source IP address filter to limit which senders it is willing to receive multicasts from on the specific multicast destination IP address (group)
- **z/OS adds support for IGMPv3 (RFC 3376), MLDv2 (RFC 3810), and APIs (RFC 3678)**
  - Protocol support as a host (not a multicast router)
  - APIs to support the new functions in IGMPv3 and MLDv2

z/OS is an IPv4 and IPv6 multicast-enabled host (sending or receiving multicast datagrams).

z/OS uses IGMP or MLD to report multicast group membership to directly connected multicast-enabled routers.

The routers in the network infrastructure need to dynamcially optimize how multicast datagrams are propagated throughout the entire network in order to reach all receiving membership hosts, with the least amount of datagram duplication effort.

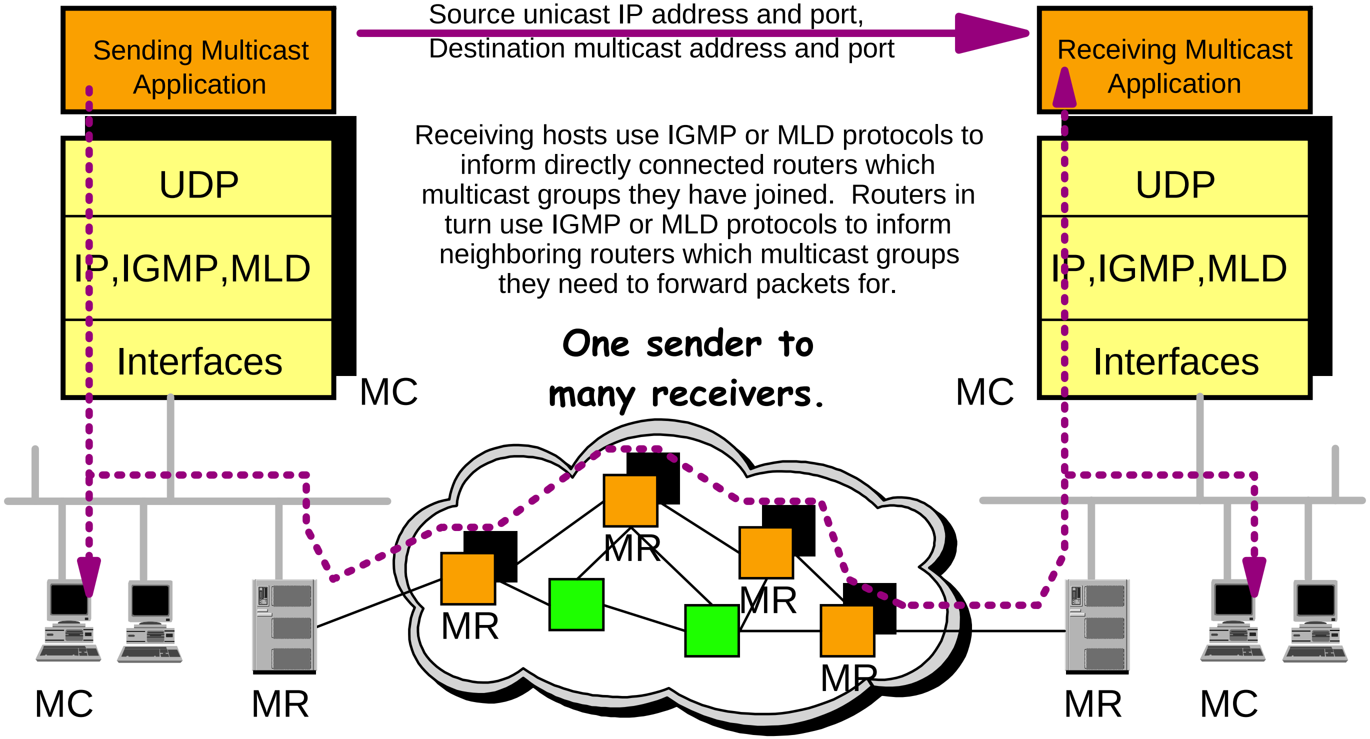

Source unicast IP address and port, Destination multicast address and port

Receiving hosts use IGMP or MLD protocols to inform directly connected routers which multicast groups they have joined. Routers in turn use IGMP or MLD protocols to inform neighboring routers which multicast groups they need to forward packets for.

One sender to many receivers.

Extending z/OS FTP support of Unicode data transfers to meet data interchange requirements from government and public agencies

- **Unicode requirements are becoming more and more common when exchanging data with**
  - partner companies
  - various public services
  - government agencies
- **FTP is an easy-to use technology that is common to all platforms making it very easy to exchange files**
- **z/OS CS V1R8 FTP:**
  - Support for transfer and store UTF-8 text data
- **z/OS CS V1R9 FTP:**
  - Support for transfer of text data in:
    - UTF-16
    - UTF-16LE
    - UTF-16BE
  - Support to store text data in:
    - UTF-16
      - always as big endian on z/OS

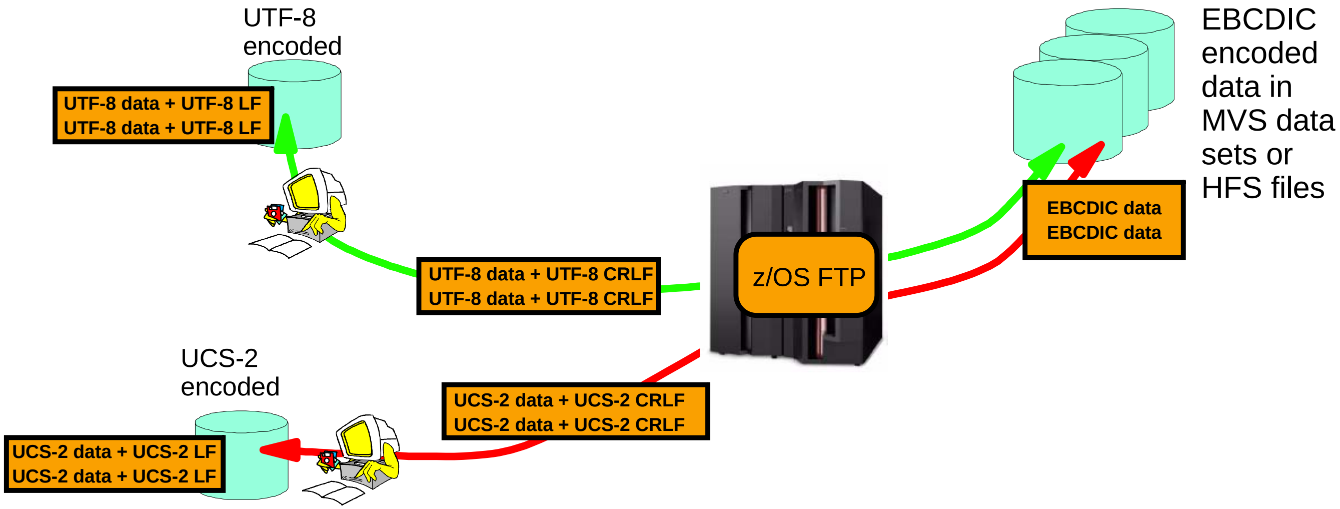

UTF-8 encoded

UTF-8 data + UTF-8 LF
UTF-8 data + UTF-8 LF

UTF-8 data + UTF-8 CRLF
UTF-8 data + UTF-8 CRLF

z/OS FTP

EBCDIC encoded data in MVS data sets or HFS files

EBCDIC data
EBCDIC data

UCS-2 data + UCS-2 CRLF
UCS-2 data + UCS-2 CRLF

UCS-2 encoded

UCS-2 data + UCS-2 LF
UCS-2 data + UCS-2 LF

A UTF-8 data stream is a Multi-Byte Character Set (MBCS) stream. Each character occupies from one to six bytes:

Single-byte:	**0xxx xxx**
Two-byte:	**110x xxxx 10xx xxxx**
Three-byte:	**1110 xxxx 10xx xxxx 10xx xxxx**
Four-byte:	**1111 0xxx 10xx xxxx 10xx xxxx 10xx xxxx**
Five-byte:	**1111 10xx 10xx xxxx 10xx xxxx 10xx xxxx 10xx xxxx**
Six-byte:	**1111 110x 10xx xxxx 10xx xxxx 10xx xxxx 10xx xxxx 10xx xxxx**

One of the attributes of UTF-8 is that it carries US-ASCII as a subset of the supported characters. Since all US-ASCII characters have the high-order bit set to zero, they are all valid single-byte UTF-8 characters. UTF-8 encoded files can be stored, and transported using the normal line-termination sequences of the operating systems (0d0a or 0a - 15 or 25).

UTF-16 uses two bytes binary to represent characters. The bytes may be ordered according to little endian (LE) or big endian (BE).

Policy configuration data structure overview in z/OS V1R9

- **There is one Policy Agent per LPAR.**
  - This one Policy Agent supports all stacks that run in that LPAR.

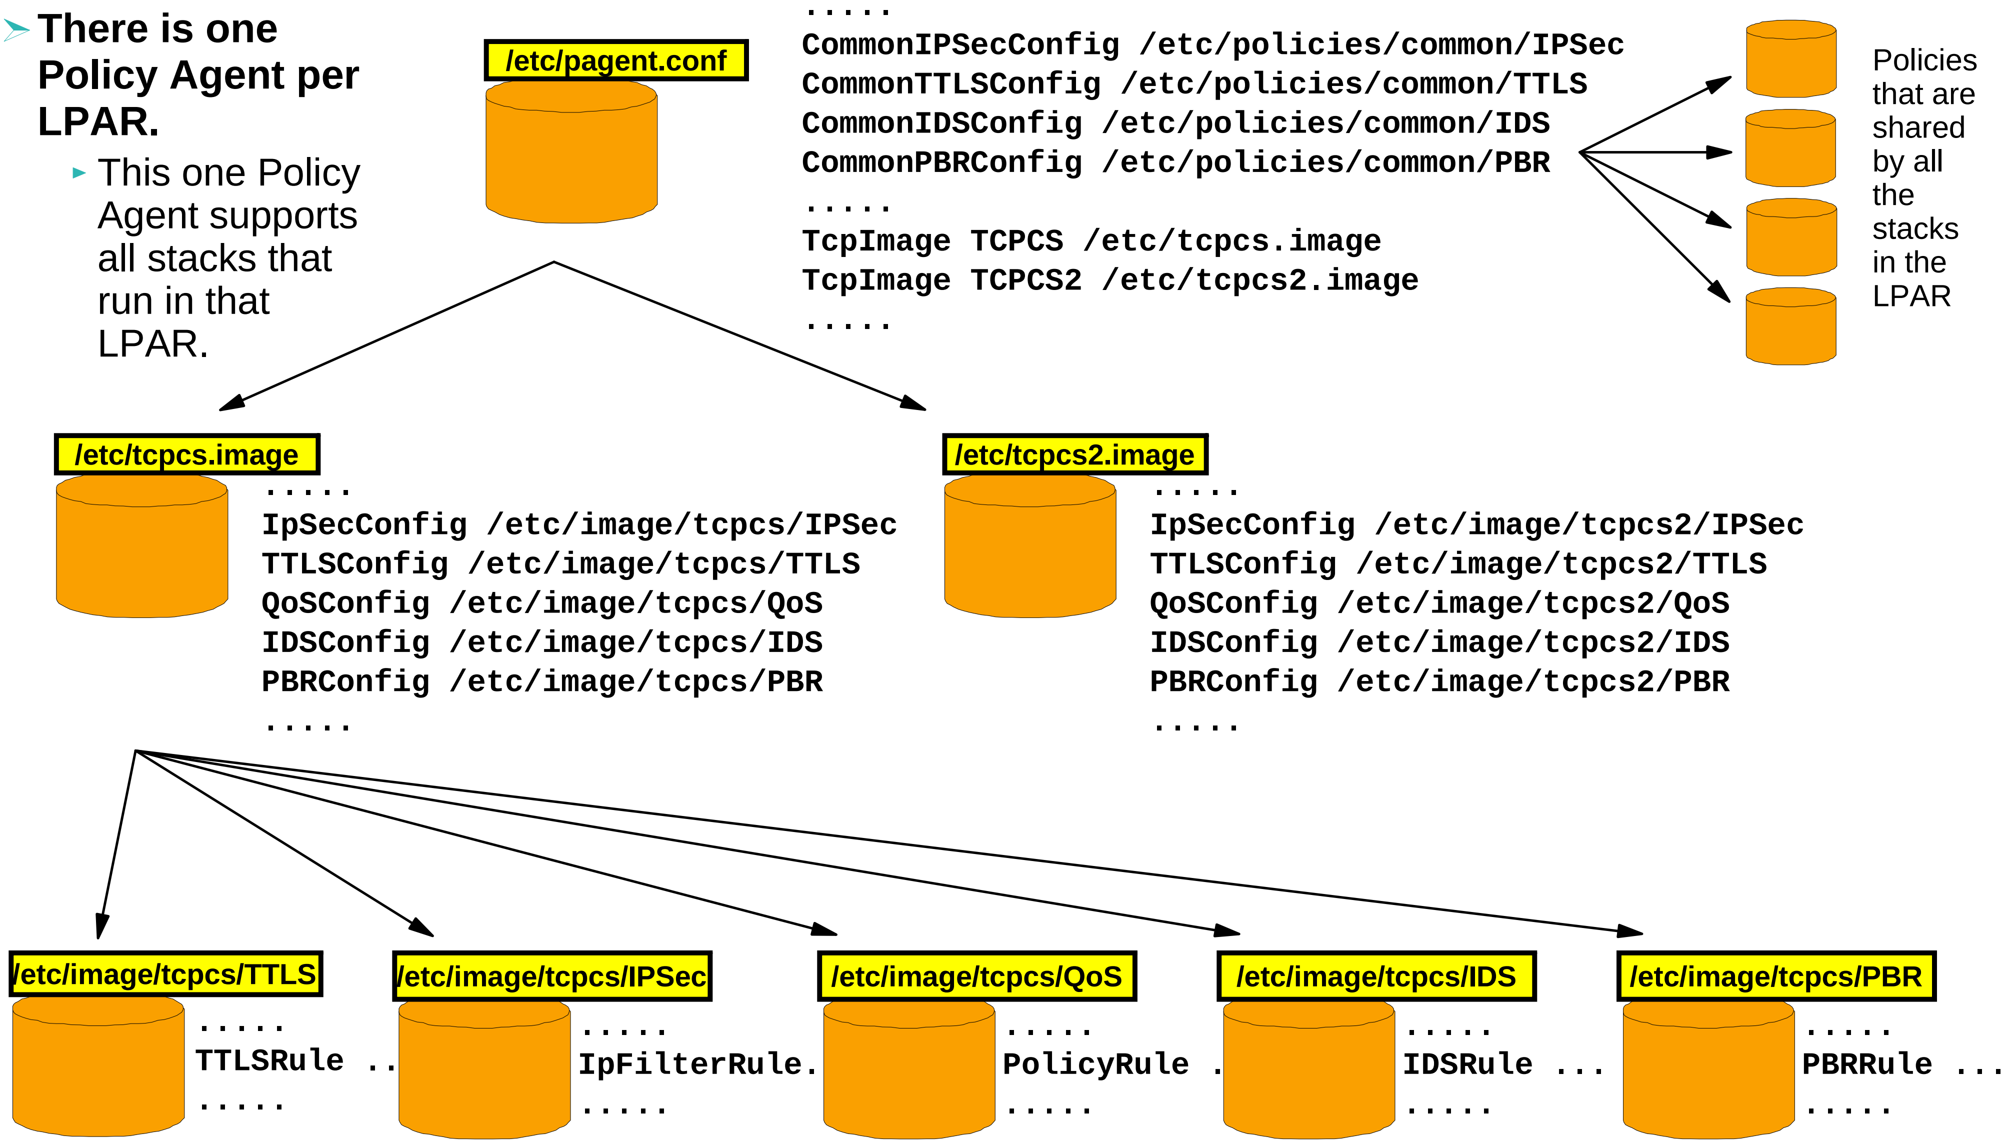

/etc/pagent.conf

```
.....
CommonIPSecConfig /etc/policies/common/IPSec
CommonTTLSConfig /etc/policies/common/TTLS
CommonIDSConfig /etc/policies/common/IDS
CommonPBRConfig /etc/policies/common/PBR
.....
TcpImage TCPCS /etc/tcpcs.image
TcpImage TCPCS2 /etc/tcpcs2.image
.....
```

Policies that are shared by all the stacks in the LPAR

/etc/tcpcs.image

```
.....
IpSecConfig /etc/image/tcpcs/IPSec
TTLSConfig /etc/image/tcpcs/TTLS
QoSConfig /etc/image/tcpcs/QoS
IDSConfig /etc/image/tcpcs/IDS
PBRConfig /etc/image/tcpcs/PBR
.....
```

/etc/tcpcs2.image

```
.....
IpSecConfig /etc/image/tcpcs2/IPSec
TTLSConfig /etc/image/tcpcs2/TTLS
QoSConfig /etc/image/tcpcs2/QoS
IDSConfig /etc/image/tcpcs2/IDS
PBRConfig /etc/image/tcpcs2/PBR
.....
```

/etc/image/tcpcs/TTLS

```
.....
TTLSRule ..
.....
```

/etc/image/tcpcs/IPSec

```
.....
IpFilterRule.
.....
```

/etc/image/tcpcs/QoS

```
.....
PolicyRule .
.....
```

/etc/image/tcpcs/IDS

```
.....
IDSRule ...
.....
```

/etc/image/tcpcs/PBR

```
.....
PBRRule ...
.....
```

Centralized networking policy management

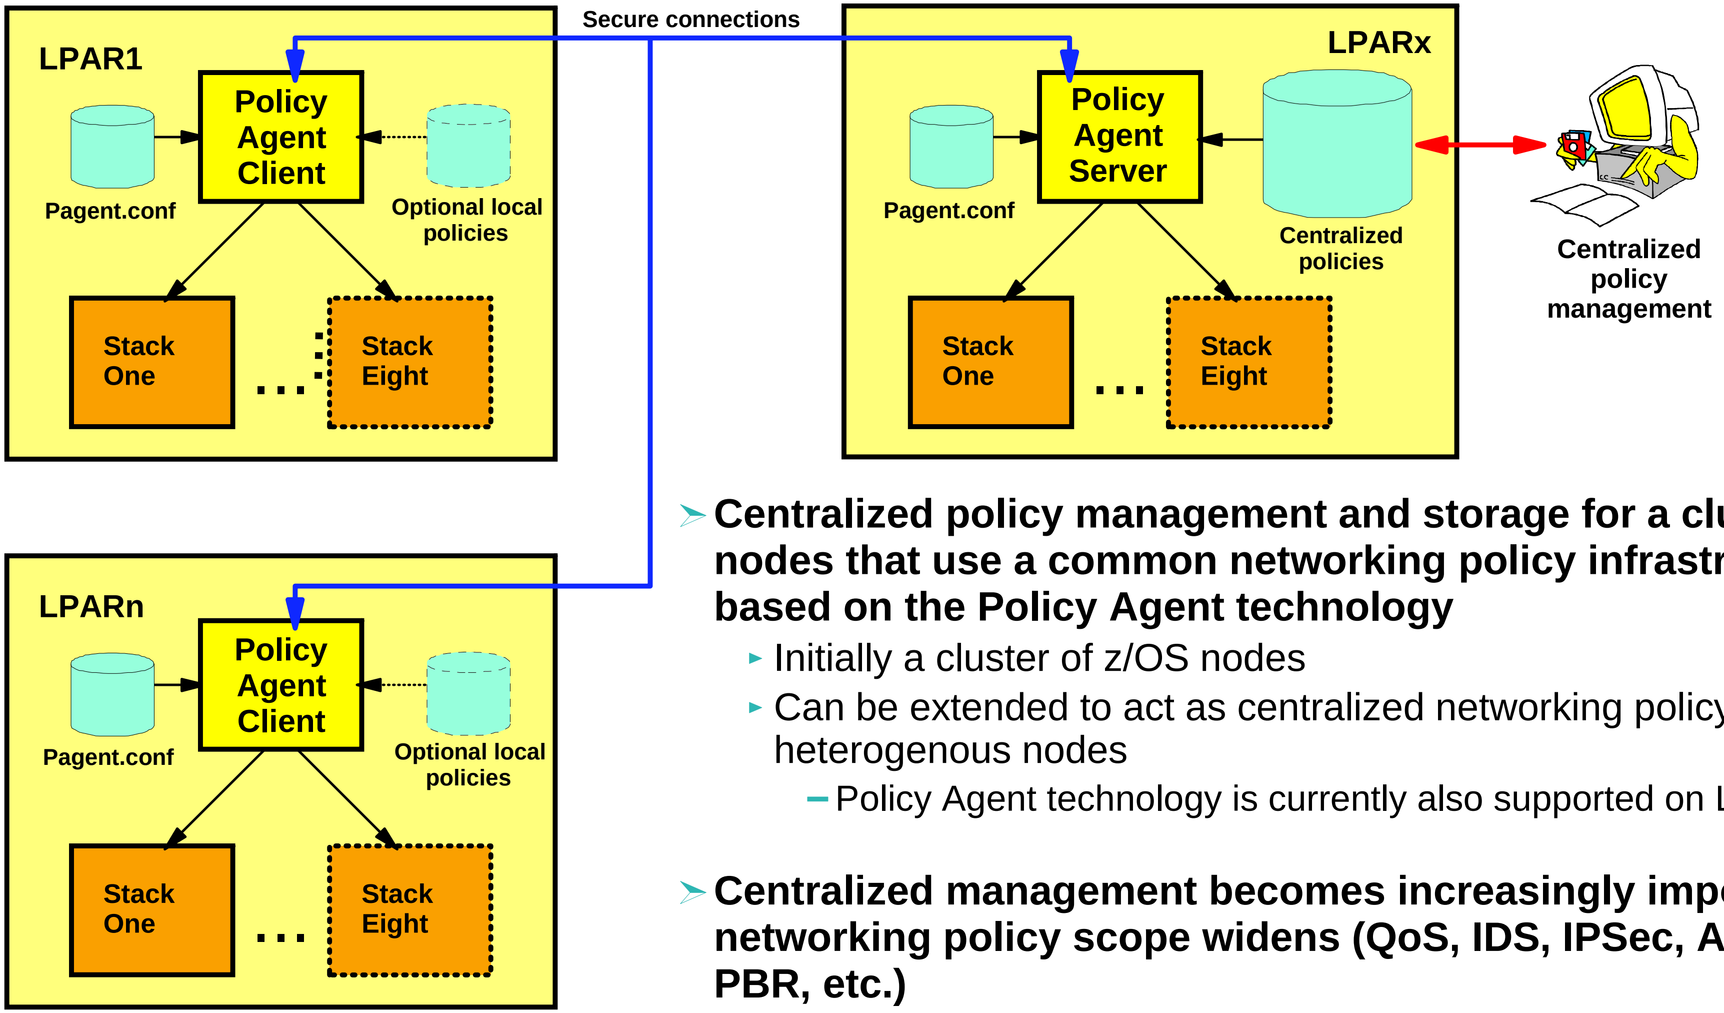

- **Centralized policy management and storage for a cluster of nodes that use a common networking policy infrastructure based on the Policy Agent technology**
  - Initially a cluster of z/OS nodes
  - Can be extended to act as centralized networking policy server for heterogenous nodes
    - Policy Agent technology is currently also supported on Linux

- **Centralized management becomes increasingly important as networking policy scope widens (QoS, IDS, IPSec, AT-TLS, PBR, etc.)**

- **Some of the policy types can benefit from additional centralized services, for which the centralized policy services will serve as a base**

Remove support for TN3270 server in TCP/IP stack

- **Before z/OS V1R6, TN3270 server ran as a subtask of the IBM TCPIP stack address space**
- **In z/OS V1R6 through V1R8, you have a choice:**
  - Run the TN3270 server as a separately started address space from TCPIP
  - Continue to run TN3270 server as a subtask of the TCPIP address space
- **From z/OS V1R9, TN3270 server runs as a separate address only**
- **Reasons why an installation may want to run the TN3270 server in a separate address space:**
  - Allows for prioritzation of TCPIP address space vs TN3270 server
  - Much less likely for TN3270 server failure to cause a total TCPIP failure
  - Allow for easier problem diagnosis for both TCPIP and TN3270
  - Easier controls for starting and stopping the server
- **Considerations**
  - Profile statements are the same and must be in a file separate from TCPIP
  - Commands are the same but must be directed to the intended TN3270 procedure name
  - Multiple TCPIP stacks supported
    - One server per stack (affinity)
    - One server associated with all stacks (Generic Server)
  - Multiple TN3270 server address spaces supported
    - Max 8 TN3270 server address spaces per LPAR
    - Only one can activate the SNMP subagent (for response time data reporting via SNMP) in a stack
      - Must have stack affinity to that stack
      - The first one started with stack affinity activates the SNMP subagent
  - Must run TN3270 server with affinity for the following functions
    - TN3270 SNMP subagent
    - WLM function
  - Requirements
    - Separate start up JCL. Sample is provided.

Remote terminal access

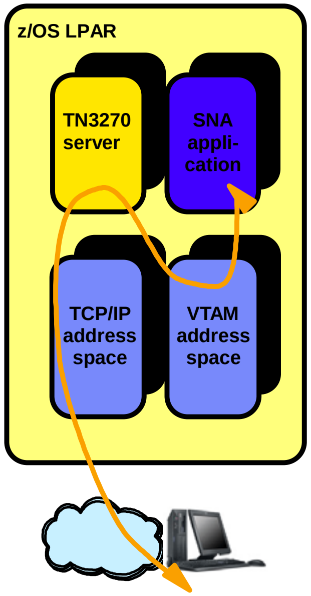

Security

IP Security network management interface (NMI) support for real-time IP security monitoring management applications

- **In z/OS V1R7, IP security network management data was made available via a UNIX shell command (the ipsec command)**
- **Adding a formalized network management programming interface to retrieve IP security manegement data, enables a network management application, such as IBM Tivoli OMEGAMON to access IP Security management data and integrate such data into the network management functions OMEGAMON XE for Mainframe Networks already provides.**

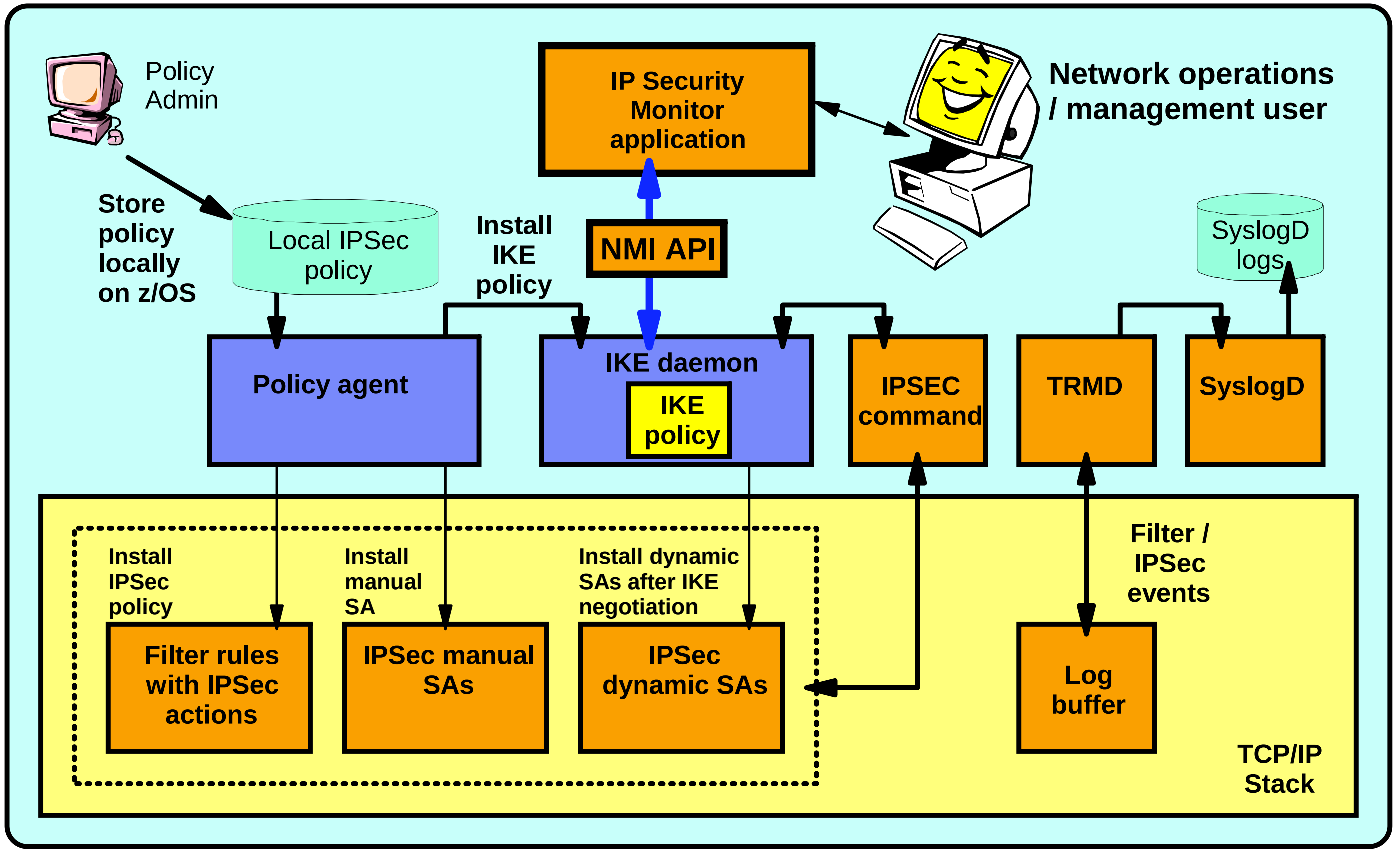

Data similar to what can be retrieved using the ipsec command will be available over the IP Security NMI interface:

1. IP filtering rules and statistics
2. IKE Phase I SA information and status
3. IKE Phase 2 SA information and status
4. Manual SA information and status
5. Port translation data

IBM Tivoli's OMEGAMON XE for Mainframe Network product is expected to utilize this interface.

Centralized Network Security Services - overview

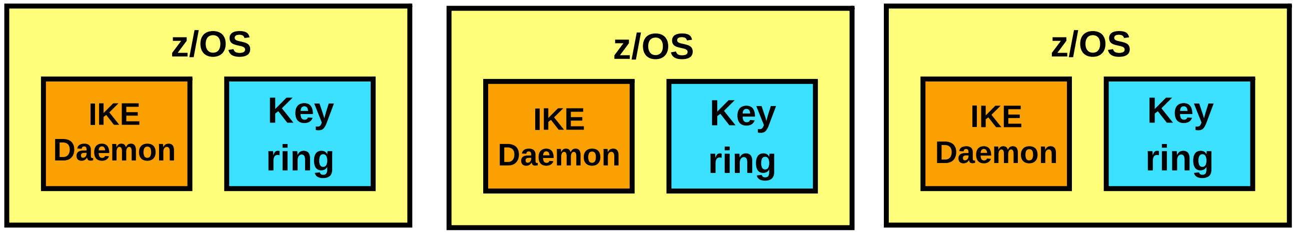

Prior to z/OS V1R9 all z/OS CS IKE Daemons have their own key ring repository.

- **Network Security Services z/OS V1R9 new feature**
  - Centralized certificate and keyring management
  - Centralized services that require access to private keys
  - Monitoring and management for IPSec security for z/OS systems within and across Sysplexes
- **IKE Daemon may be configured as a Network Security Client.**
  - Configuration is on a per-stack basis
  - Each NSS-enabled stack will appear to the Network Security Server as an independent client.
  - For TCP/IP stacks that are not configured to use Network Security Services, the IKE daemon will continue to manage certificates out of a local keyring.

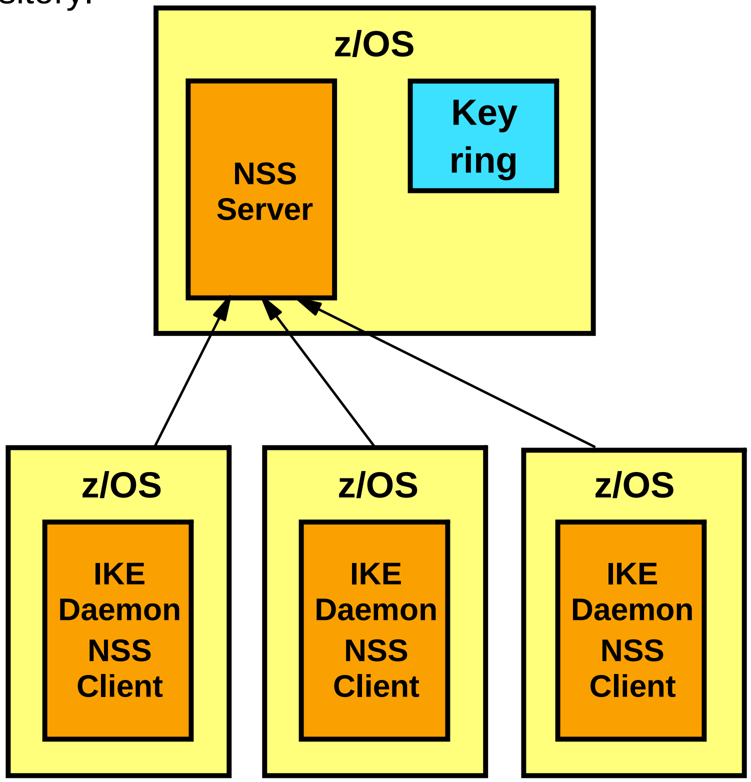

z/OS CS V1R9 IKE Daemons may have their IPSec certificates stored on a single system.

SSL/TLS enabling applications on z/OS - technology choices after AT-TLS was introduced in z/OS V1R7

- **SSL/TLS support can since z/OS V1R7 be implemented using one of two methods on z/OS:**
  - Change applications that need SSL/TLS security to:
    - Support appropriate application-specific configuration options to specify SSL/TLS options, such as keyring, cipher suites, application-specific security options, etc.
    - Replace selected socket calls with calls to system SSL (C/C++/Java only)
  - Let the Application-Transparent TLS layer inside the TCP/IP stack handle all SSL/TLS processing:
    - Common SSL/TLS configuration for all applications through an AT-TLS policy (managed by the Policy agent)
    - Use optimized SSL/TLS code within the TCP/IP stack that interfaces to system SSL to implement the SSL/TLS functions
    - In most cases SSL/TLS support can be added without application changes

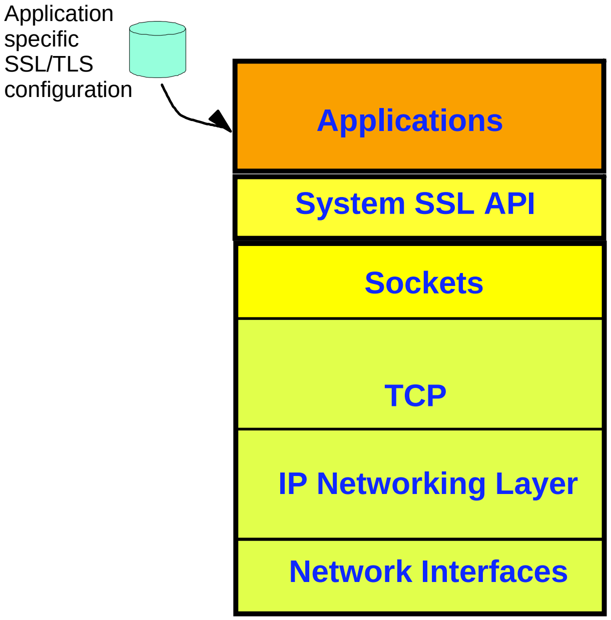

- **Using AT-TLS has several advantages:**
  - AT-TLS provides SSL/TLS features above and beyond what most SSL/TLS applications choose to support - such as, support for Certificate Revocation Lists (CRLs), multiple keyrings per server, optional use of system SSL cache, etc.
  - AT-TLS uses an optimized SSL/TLS infrastructure that in most cases performs better than when SSL/TLS is implemented directly in the applications
  - Support of new SSL/TLS functions, such as new ciphersuites, can be added without application changes
  - Allows SSL/TLS-enabling non-C sockets applications on z/OS, such as CICS Sockets, Assembler- and Callable sockets, etc.

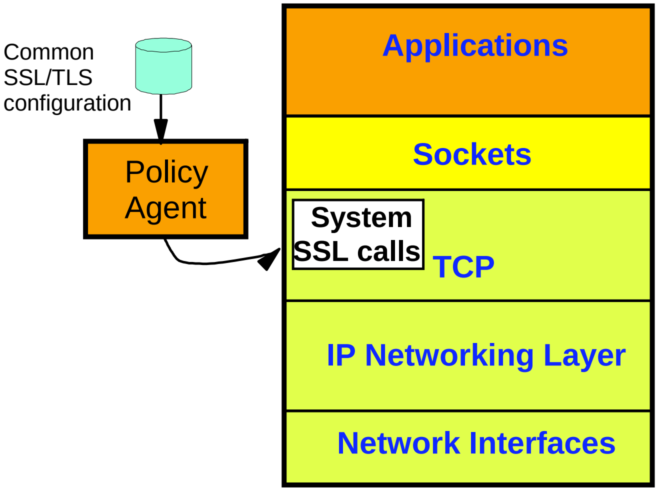

AT-TLS enabling the TN3270 server and the FTP client and server

- **Both the FTP server and client, and the TN3270 server on z/OS have in the past implemented SSL/TLS support**
  - With the advantages of AT-TLS, it is desirable to migrate that SSL/TLS support to AT-TLS
- **"Move" the SSL/TLS-specific configuration into the common AT-TLS policy format**
  - Remove application-specific configuration syntax and processing logic
  - One common policy format where new options can be added without changes to all applications

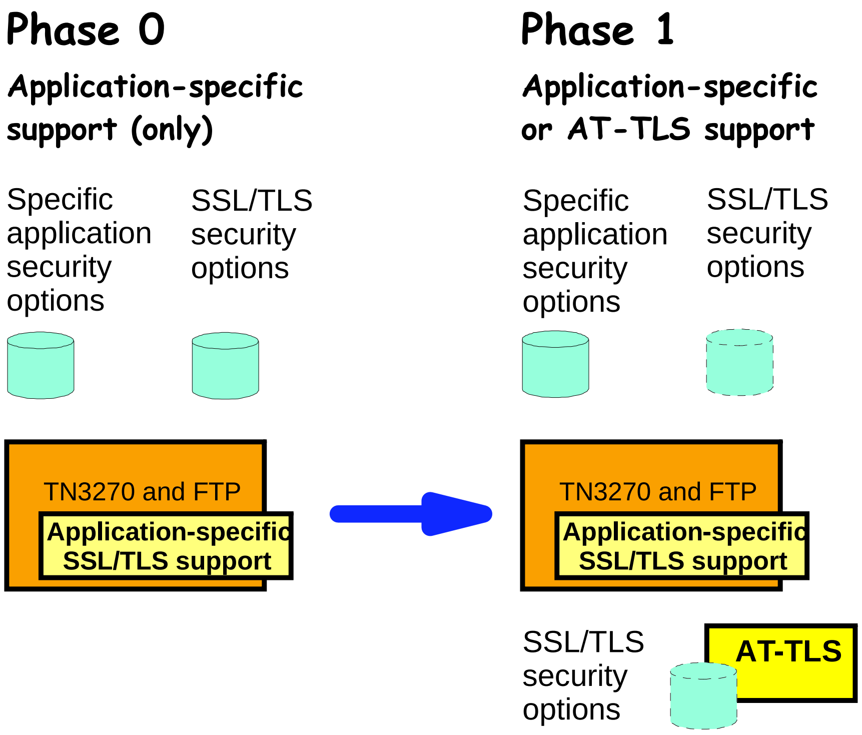
Phase 0
Application-specific support (only)
Specific application security options
SSL/TLS security options
TN3270 and FTP
Application-specific SSL/TLS support
Phase 1
Application-specific or AT-TLS support
Specific application security options
SSL/TLS security options
TN3270 and FTP
Application-specific SSL/TLS support
SSL/TLS security options
AT-TLS

AT-TLS enabling FTP

- **A new FTP.DATA option to instruct the FTP server or client to use AT-TLS instead of FTP's own system SSL calls is being implemented for Phase 1:**
  - TLSMECHANISM (Client and Server)

```
                    +--FTP---+
                    |        |
-->TLSMECHANISM--+--------+--->< 
                    |        +
                    +--ATTLS-+
```

- **When ATTLS is specified as TLS mechanism:**
  - FTP becomes an AT-TLS controlling and AT-TLS aware application
  - All the FTP-specific security options will continue to impact how FTP operates
  - The SSL/TLS security options in FTP.DATA will be ignored.
    - Matching AT-TLS policies need to be defined before enabling AT-TLS support in FTP
    - Enable "secondary mapping" on the FTP control port number only

- **FTP-specific security options:**
  - EXTENSIONS AUTH_TLS (Server)
  - SECURE_CTRLCONN (Client and Server)
  - SECURE_DATACONN (Client and Server)
  - SECURE_FTP (Client and Server)
  - SECURE_HOSTNAME (Client)
  - SECURE_LOGIN (Server)
  - SECURE_MECHANISM (Client)
  - SECURE_PASSWORD (Server)
  - SECUREIMPLICITZOS (Client)
  - TLSPORT (Client and Server)

- **FTP SSL/TLS security options**
  - CIPHERSUITE (Client and Server)
  - KEYRING (Client and Server)
  - TLSTIMEOUT (Client and Server)

AT-TLS enabling TN3270

- **A new TN3270 server option to indicate use of AT-TLS instead of the TN3270 server's own system SSL calls is being implemented for Phase 1:**
  - TTLSPORT
    - CONNTYPE retains its current meaning for a TTLSPORT
- **When TTLSPORT is used for a TN3270 server port:**
  - The TN3270 server becomes an AT-TLS controlling and AT-TLS aware application
  - All the TN3270-specific security options will continue to impact how TN3270 operates
  - Any TN3270 server SSL/TLS security options will be ignored.
    - Matching AT-TLS policies need to be defined before enabling AT-TLS support for the TN3270 server

- **TN3270-specific security options:**
  - SECUREPORT (use of this option will indicate to TN3270 that it is to use its existing application-specific SSL/TLS support, and not AT-TLS for the specified port number)
  - CONNTYPE
    - SECURE
    - NEGTSECURE
    - ANY
    - BASIC
  - EXPRESSLOGON
  - RESTRICTAPPL CERTAUTH
- **TN3270 SSL/TLS security options**
  - KEYRING
  - CRLLDAPSERVER
  - CLIENTAUTH
    - SSLCERT
    - SAFCERT
  - ENCRYPTION
  - SSLTIMEOUT
  - SSLV2/SSLNOV2

z/OS FTP server and Kerberos - single sign-on (no prompt for password)

- **One of the main benefits, and often the main reason why people use Kerberos, is the single sign-on capability:**
  - Users sign on to the Kerberos Authentication Server
  - Users are then granted access to other servers through a "ticket" approach
  - When connecting to a Kerberos-enabled server and presenting the user's "ticket", the user may be signed on implicitly
- **FTP on z/OS was Kerberos-enabled in z/OS V1R2, but continued to always require both a user ID and password.**
- **FTP protocol prevents us from bypassing the request for a user ID.**
- **If the entered FTP user ID matches the user ID in the Kerberos ticket, the prompt for an FTP password may be bypassed**
  - New FTP server configuration option to control this behavior:

```
SECURE_KERBEROS_PASSWORD {REQUIRED/OPTIONAL}
```

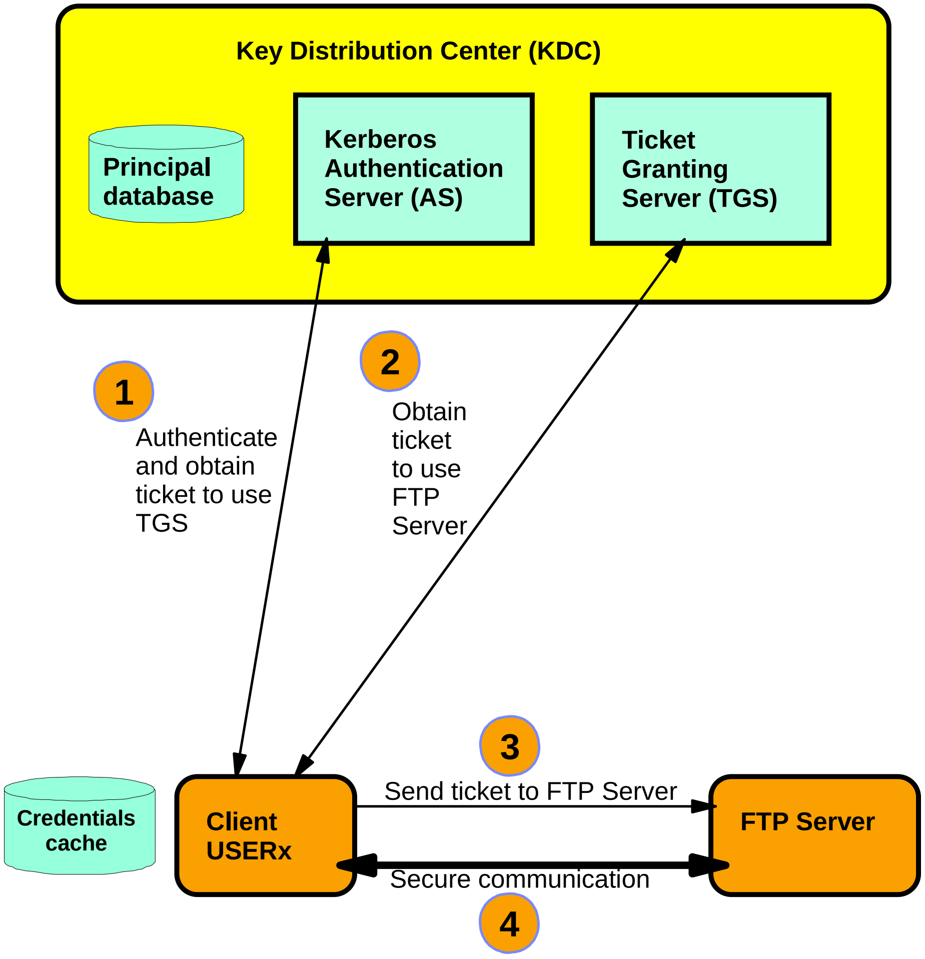
Key Distribution Center (KDC)
Principal database
Kerberos Authentication Server (AS)
Ticket Granting Server (TGS)
1
Authenticate and obtain ticket to use TGS
2
Obtain ticket to use FTP Server
3
Send ticket to FTP Server
4
Secure communication
Credentials cache
Client USERx
FTP Server

Enables easier use of z/OS FTP Server in a Kerberos-based single sign-on environment.

Business resiliency

Source IP address selection (a never-ending saga!) - distributed DVIPA as source IP address for DESTIP-based SRCIP rules

```
SRCIP
     Jobname CUSTAJOB   9.85.112.1
     Jobname CUSTBJOB   9.85.113.1
     Jobname User1*     888:555::222
     DESTIP             10.1.1.0/24 9.85.114.1
     DESTIP             10.1.2.0/24 9.85.115.1
ENDSRCIP
```

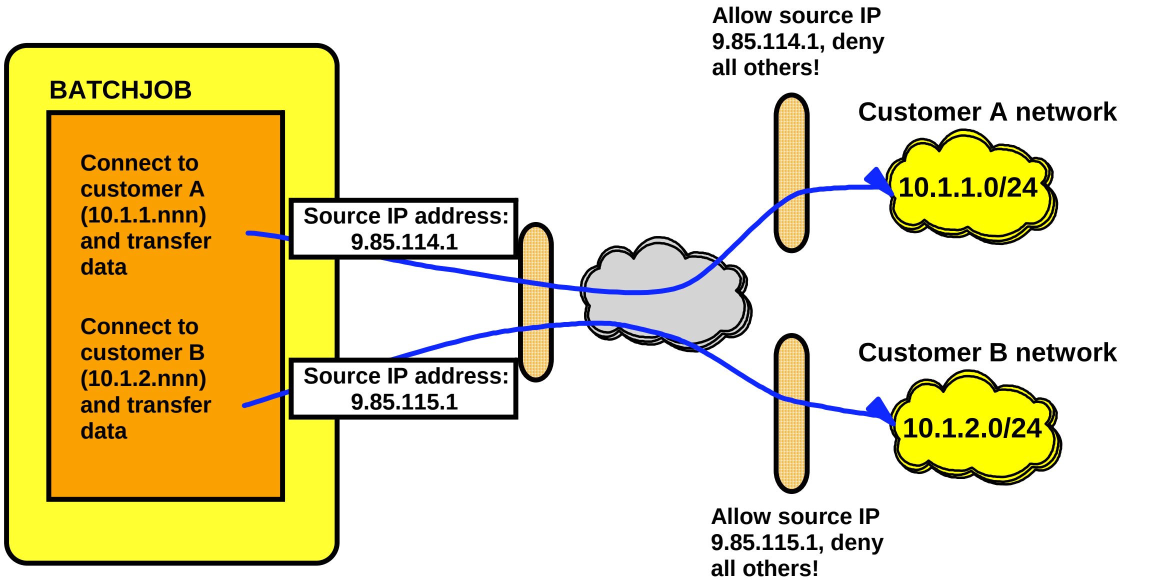
Allow source IP
9.85.114.1, deny
all others!
BATCHJOB
Connect to customer A (10.1.1.nnn) and transfer data
Connect to customer B (10.1.2.nnn) and transfer data
Source IP address: 9.85.114.1
Source IP address: 9.85.115.1
Customer A network
10.1.1.0/24
Customer B network
10.1.2.0/24
Allow source IP
9.85.115.1, deny
all others!

CINET: Supported if only one stack configured - or multiple stacks configured but all applications have stack affinity

- **In z/OS V1R8, we introduced an option to select source IP address based on the destination IP address a connection was directed towards.**
  - But we specifically excluded support for that source IP address to be a Sysplex-wide source IP address (a distributed DVIPA)

- **If installations need to be able to submit multiple jobs, that all need to connect to business partners and the jobs may execute in parallel on multiple LPARs in the Sysplex - we need a distributed DVIPA as source IP address!**

- **z/OS CS V1R9 will extend the destination-based source IP address selection to include a distributed DVIPA:**
  - Participating stacks will reserve a coordinated range of port numbers for this use - new option on GLOBALCONFIG
  - If an application issues an explicit bind to INADDR_ANY or INADDR6_ANY and port 0, the stack has SYSPLEXPORTS enabled, and the stack has SRCIP rules - a port from this new range will be requested

WLM and SD may not always get it the way you want it to be!

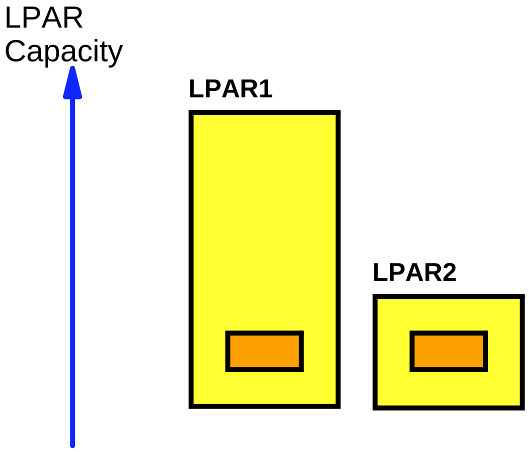

- **Target systems vary significantly in terms of capacity (small systems along with larger systems).**
  - WLM recommendations may favor the larger systems significantly.
  - However, the target application may not scale well to larger systems, being unable to take full advantage of the additional capacity on the larger systems.
  - The result can be that these types of servers when running on larger systems get inflated WLM recommendations and as a result they get overloaded with work.

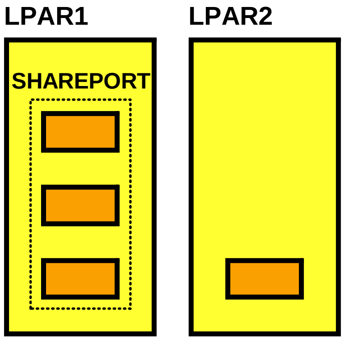

- **SHAREPORT is deployed, yet not all systems have the same number of SHAREPORT server instances (one has three the other has only one).**
  - The current Round Robin (RR) or WLM recommendations do not change distribution based on the number of server instances on each target.
  - RR distributes 1 connection per target stack regardless of the number of shareport server instances on that stack.
  - WLM Server-specific weights from a target stack with multiple server instances reflect the average weight.

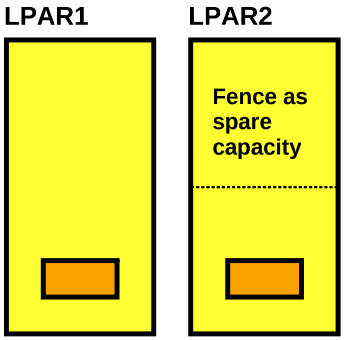

- **Customers would like to reserve some capacity on certain systems for batch type of workloads that get injected into the system during specific time periods and which have specific time window completion requirements.**
  - If that system is also a target for long running distributed connections, WLM recommendations will allow that available capacity to be consumed and thereby potentially impact the completion of the batch jobs or vice versa (the connections on that system may suffer from a performance perspective when those jobs are running).

Sysplex Distributor to add support for a weighted distribution method based on currently active connections

- **Add new weighted distribution method**
  - WEIGHTEDACTIVE method
  - Weights to be configured on the VIPADISTRIBUTE statement per destination IP address
  - Weights to balance active connections, not incoming connections
    - Objective is to keep the number of active connections distributed according to the configured weights
    - More optimal than traditional round robin or weighted round robin algorithms

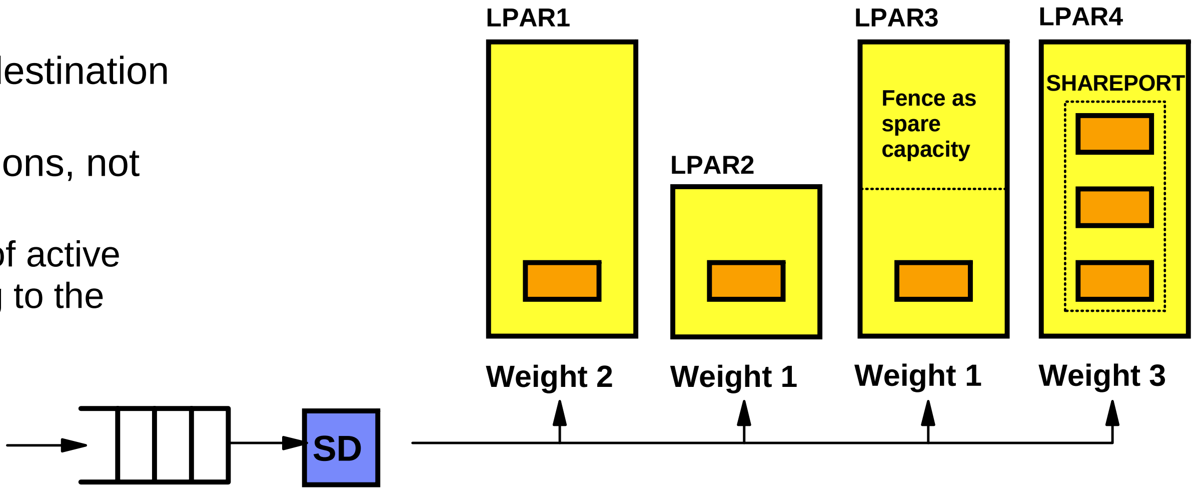

Case 1	Configu-red weights	Current number of active connec-tions	Norma-lized	Status
LPAR1	2	15	1.5	below
LPAR2	1	10	1	on target
LPAR3	1	10	1	on target
LPAR4	3	30	3	on target

Case 2	Configu-red weights	Current number of active connec-tions	Norma-lized	Status
LPAR1	2	30	3	above
LPAR2	1	10	1	on target
LPAR3	1	10	1	on target
LPAR4	3	20	2	below

Include System z specialty processor capacity in workload management decisions made by SD and LBA (if relevant!)

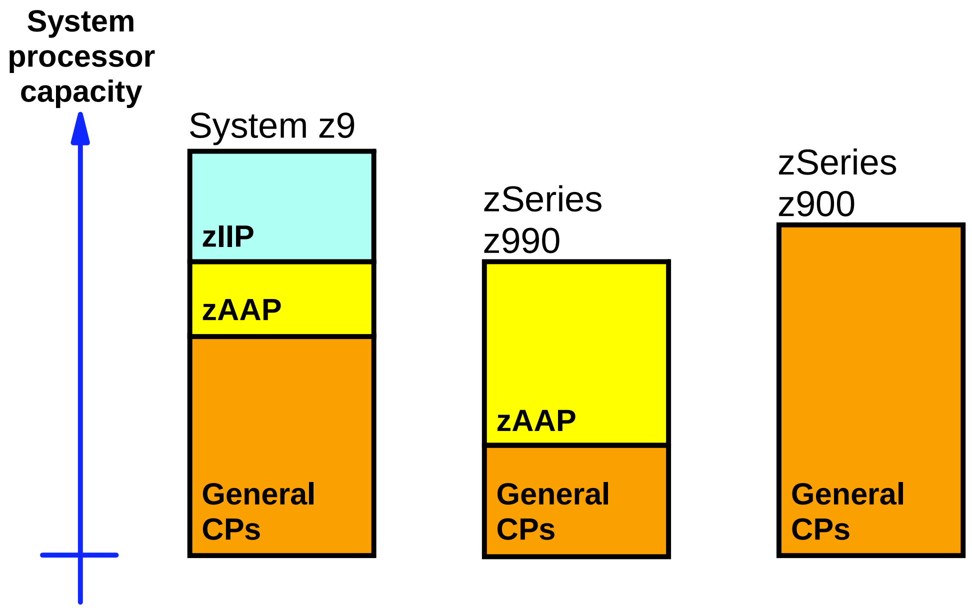

- **What is the available capacity of each System z server node (CPC) and how is that capacity reflected in the WLM weights that are used by:**
  - Sysplex Distributor
  - z/OS Load Balancing Advisor
- **So far, the capacity of specialty processors, such as zAAP and zIIP, has not been factored in by WLM when calculating the weights that are used by the above CS functions.**
- **WLM implements new functions and interfaces in z/OS V1R9 that allow CS access to improved weight information.**

- **BASEWLM - system weights (Sysplex Distributor and Load Balancing Advisor)**
  - Optionally use combined WLM weights that factor in capacity of specialty processors based on proportion by which they are used on each server node as determined by WLM
  - Use separate WLM weights from all processor types and combine them based on CS configuration of how each processor type should impact the combined weights
    - WLM and CS do not know the mix of processor type that is needed by the application workload is distributed to, so an option to configure has been added to both Sysplex Distributor and LBA

```
VIPADISTRIBUTE ... BASEWLM PROCTYPE(CP=3,zAAP=1,zIIP=0)
```

- **SERVERWLM - server-specific weights (Sysplex Distributor and Load Balancing Advisor)**
  - Use combined WLM weights that factor in capacity of specialty processors based on proportion by which the specific server address space uses these specialty processors as determined by WLM
    - No configuration changes needed for SERVERWLM

Getting the start-up sequence of the TCP/IP stack right!

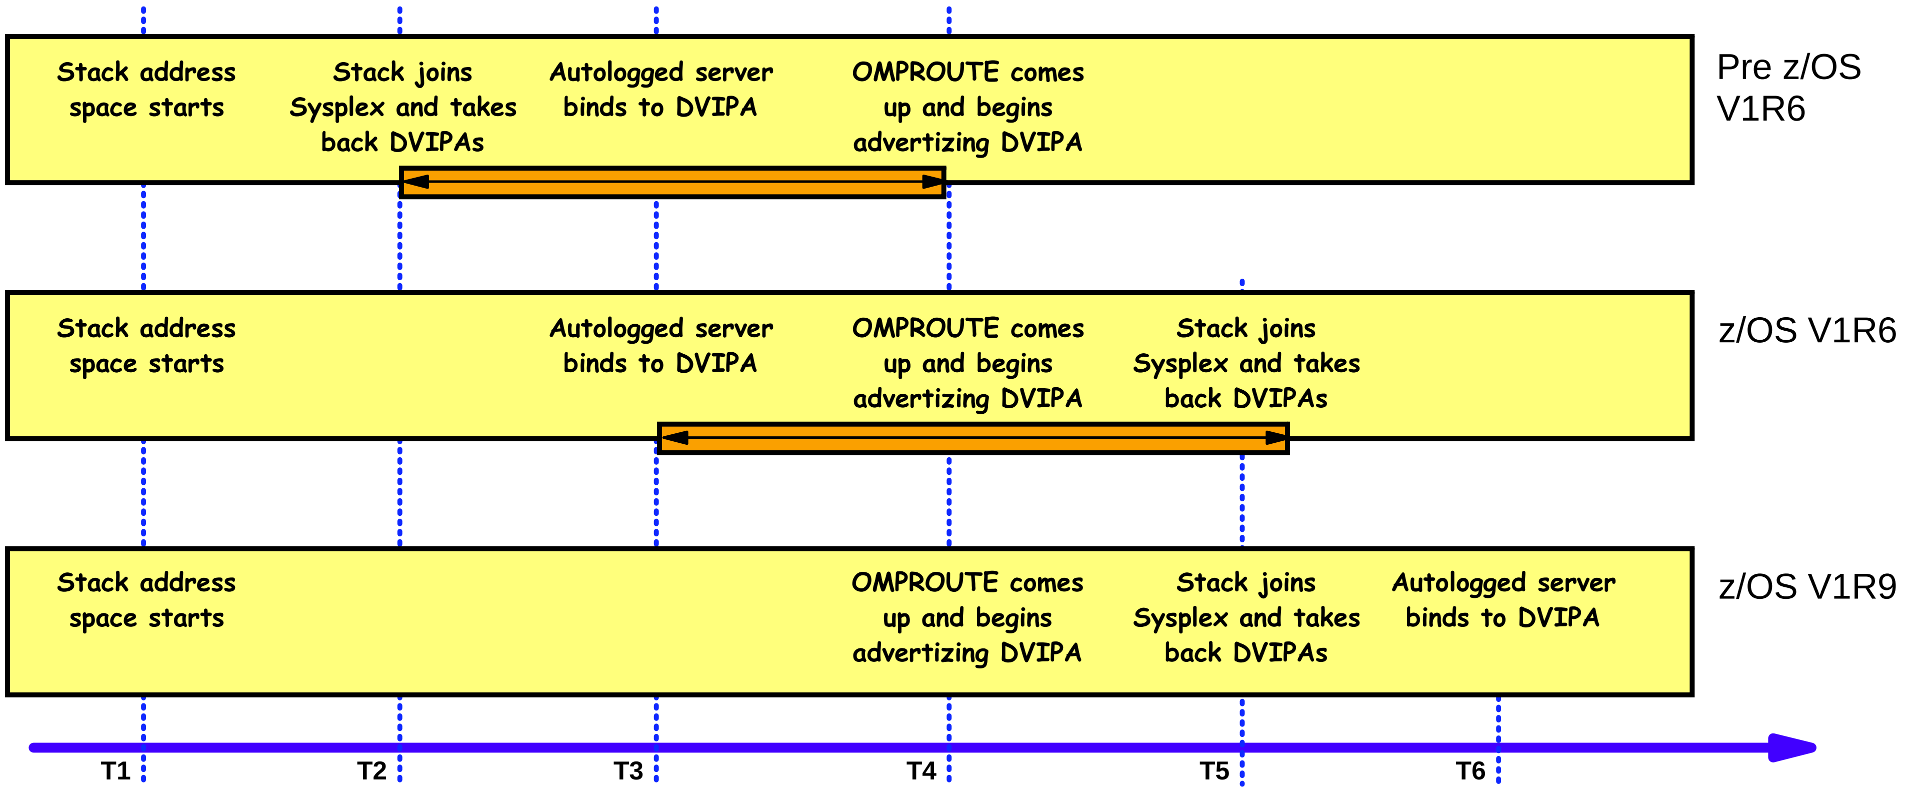
Stack address space starts
Stack joins Sysplex and takes back DVIPAs
Autologged server binds to DVIPA
OMPROUTE comes up and begins advertizing DVIPA
Pre z/OS V1R6
Stack address space starts
Autologged server binds to DVIPA
OMPROUTE comes up and begins advertizing DVIPA
Stack joins Sysplex and takes back DVIPAs
z/OS V1R6
Stack address space starts
OMPROUTE comes up and begins advertizing DVIPA
Stack joins Sysplex and takes back DVIPAs
Autologged server binds to DVIPA
z/OS V1R9
T1
T2
T3
T4
T5
T6

- **Pre z/OS V1R6:**
  - DVIPAs unavailable between T2 and T4
  - Availability issue especially if the DVIPAs were taken back from backup stacks
- **z/OS V1R6:**
  - GLOBALCONFIG DELAYJOIN delays joining the Sysplex until after OMPROUTE has initialized
  - Solves the issue with taking back DVIPAs too early
  - Autologged servers now fail when they try to bind to their DVIPAs at T3, and need to be restarted after T5
- **z/OS V1R9:**
  - AUTOLOG entry option added to specify DELAYSTART to delay Autolog until after the stack has joined the Sysplex

Configurable WLM polling interval

- **The TCP/IP stack(s) poll WLM every 60 seconds for weights**
  - BASEWLM - the distributor polls WLM for system weights
  - SERVERWLM - each target polls WLM for server-specific weights

- **WLM calculates new weights**
  - Based on a comparison of the last 10 seconds of CPU utilization on registered Sysplex systems or servers
  - It keeps a 3 minute rolling average of these calculations
  - The 3 minute rolling average is used to smooth the results
  - This average is returned by WLM when it is polled

- **The TCP/IP polling interval assumes weights will not change significantly from minute to minute.**
  - This proved to be a wrong assumption in certain scenarios
    - Many short-lived connections
    - System close to 100% utilized

- **The WLM polling interval is now configurable**
  - PTFed back to z/OS V1R6 via APAR PK24752

```
                  +-SYSPLEXWLMPoll 60 ----+
>>-GLOBALCONFig---+-----------------------+------><
                  +-SYSPLEXWLMPoll value -'
```

Usability

IBM Configuration Assistant for z/OS Communications Server

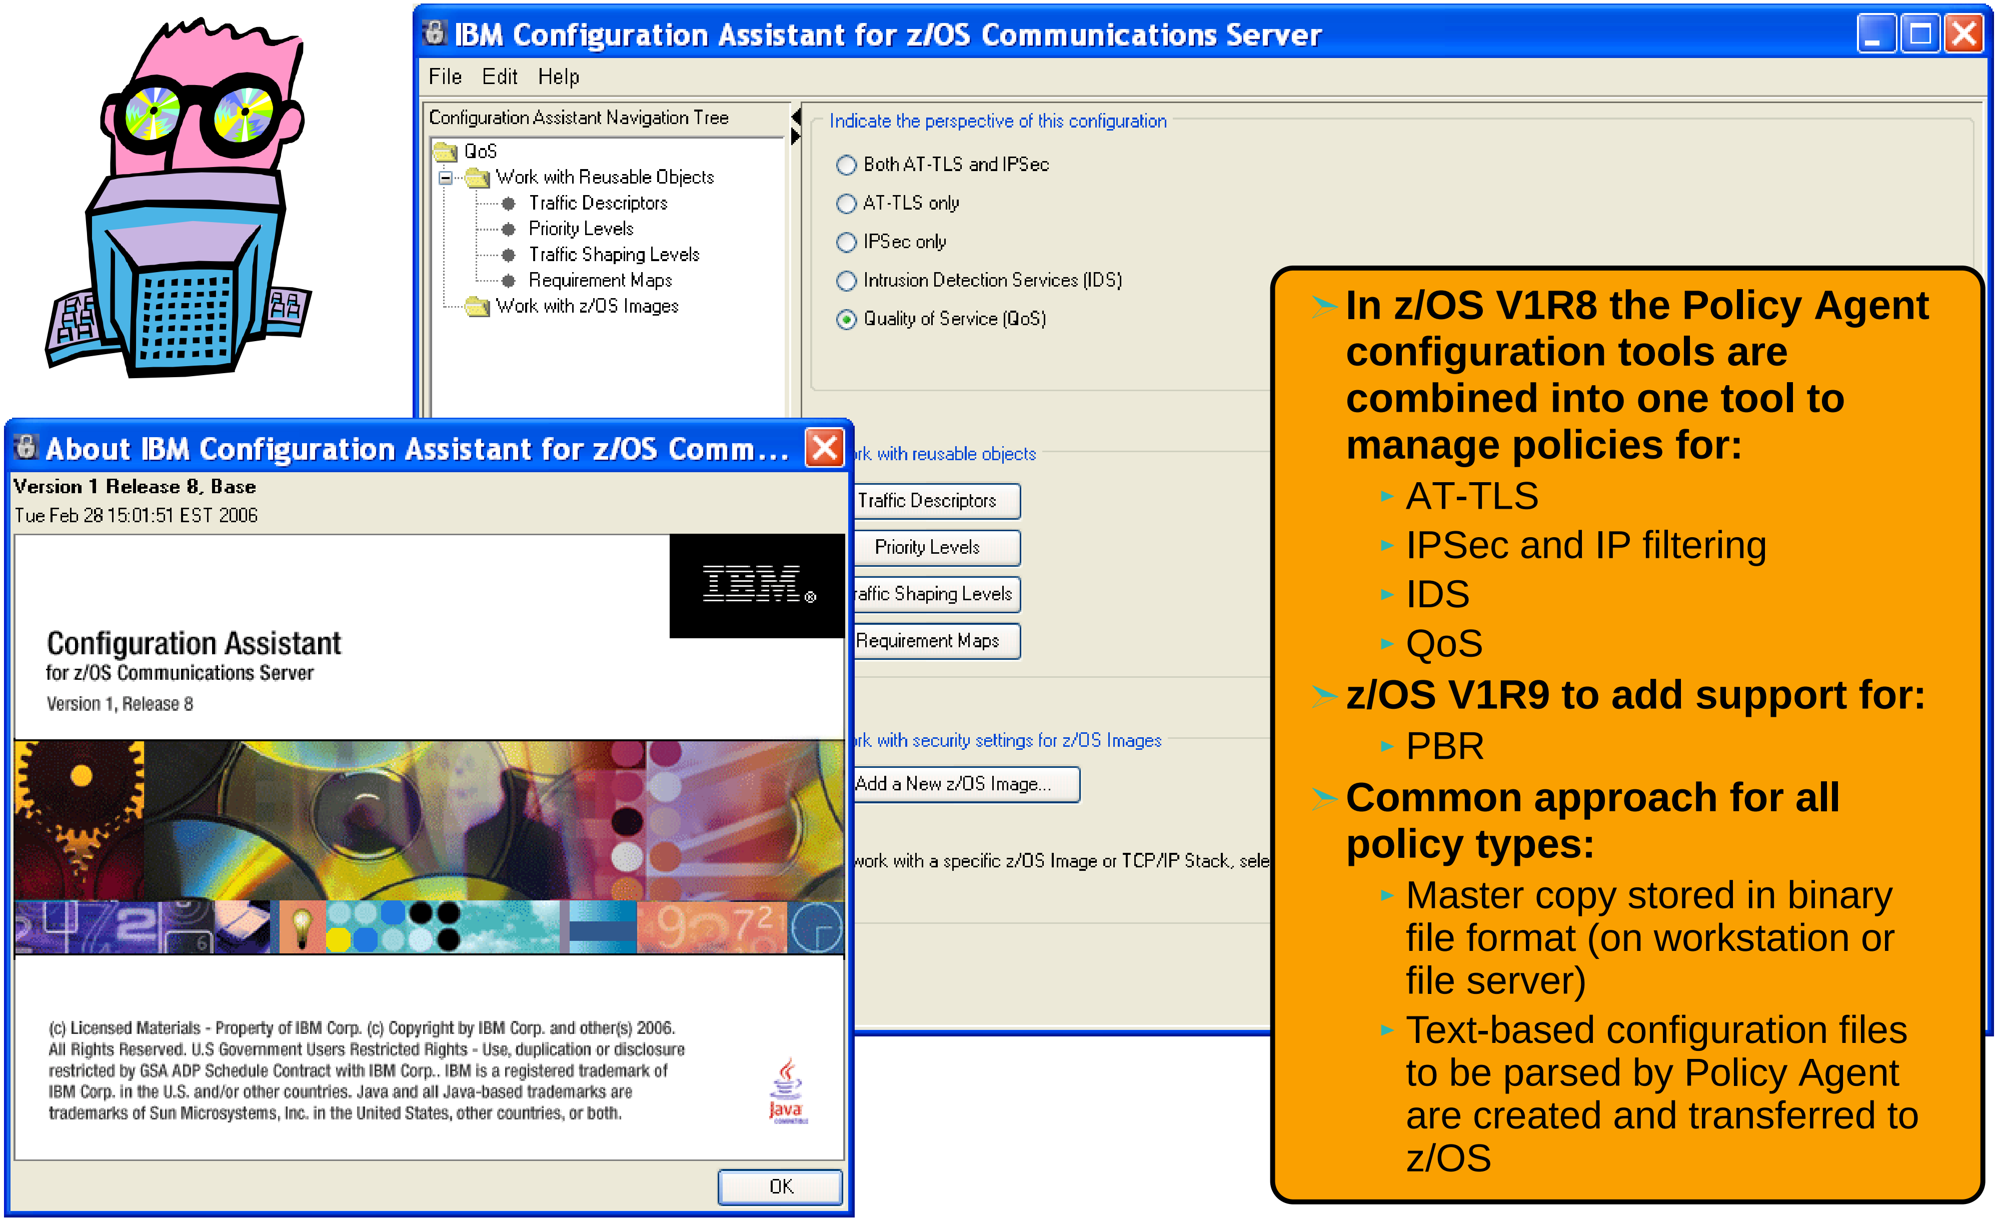

IBM Configuration Assistant for z/OS Communications Server

Configuration Assistant Navigation Tree

- QoS
  - Work with Reusable Objects
    - Traffic Descriptors
    - Priority Levels
    - Traffic Shaping Levels
    - Requirement Maps
  - Work with z/OS Images

Indicate the perspective of this configuration

○ Both AT-TLS and IPSec
○ AT-TLS only
○ IPSec only
○ Intrusion Detection Services (IDS)
● Quality of Service (QoS)

rk with reusable objects

rk with security settings for z/OS Images

work with a specific z/OS Image or TCP/IP Stack, sele

About IBM Configuration Assistant for z/OS Comm...

Version 1 Release 8, Base
Tue Feb 28 15:01:51 EST 2006

Configuration Assistant
for z/OS Communications Server
Version 1, Release 8

(c) Licensed Materials - Property of IBM Corp. (c) Copyright by IBM Corp. and other(s) 2006. All Rights Reserved. U.S Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.. IBM is a registered trademark of IBM Corp. in the U.S. and/or other countries. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

- **In z/OS V1R8 the Policy Agent configuration tools are combined into one tool to manage policies for:**
  - AT-TLS
  - IPSec and IP filtering
  - IDS
  - QoS
- **z/OS V1R9 to add support for:**
  - PBR
- **Common approach for all policy types:**
  - Master copy stored in binary file format (on workstation or file server)
  - Text-based configuration files to be parsed by Policy Agent are created and transferred to z/OS

FTP client to allow user to specify source IP address

- **Despite significant flexibility in how to direct TCP/IP to choose source IP address for outbound connections, there remains a need to be be able to specify which source IP address a given FTP client invocation should use when connecting to an FTP server.**
  - If batch FTP client job preparation is done by a group of people who do not have access to update (or maybe even view) the source IP address rules (SRCIP) in the TCP/IP profile, a need for them to specify a specific source IP address when preparing the batch jobs still exists.
- **The z/OS FTP client will implement a new command line option where a user can specify which local IP address the connection to the FTP server should come from.**
  - It is the user's responsibility to verify that the chosen address is a valid local IPv4 address that is reachable from the FTP server node
- **When a source IP address is specified in the command line invocation of the FTP client, that address will override all other source IP address selection rules.**

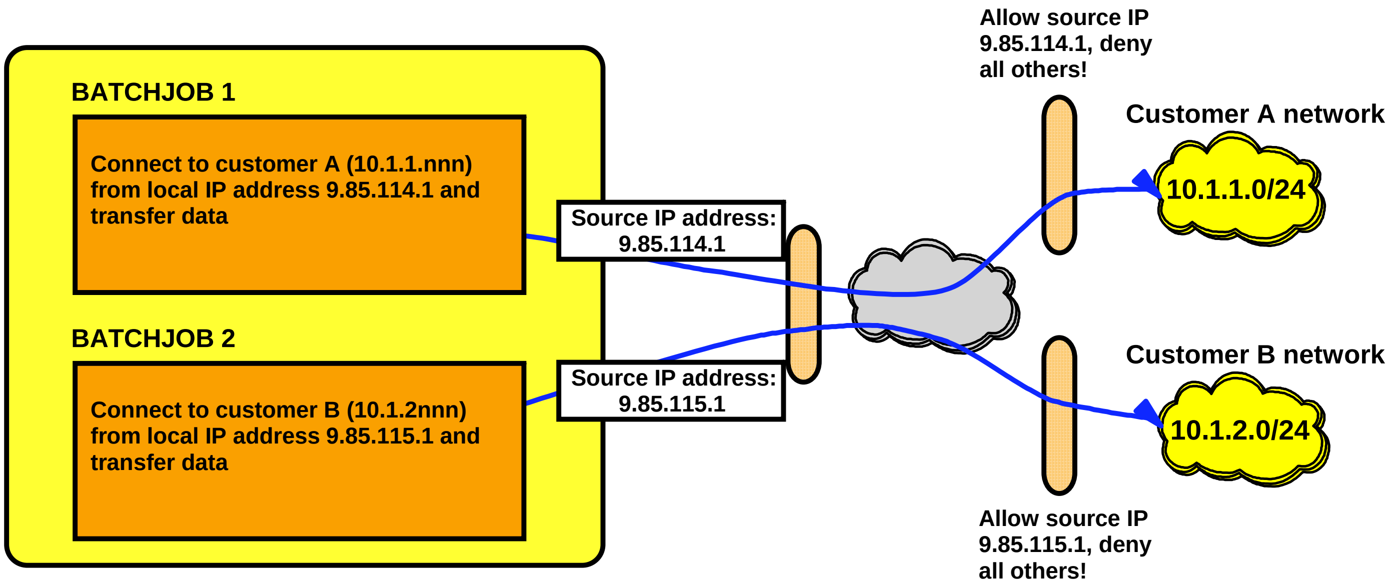
BATCHJOB 1
Connect to customer A (10.1.1.nnn) from local IP address 9.85.114.1 and transfer data
BATCHJOB 2
Connect to customer B (10.1.2nnn) from local IP address 9.85.115.1 and transfer data
Source IP address: 9.85.114.1
Source IP address: 9.85.115.1
Allow source IP 9.85.114.1, deny all others!
Customer A network
10.1.1.0/24
Customer B network
10.1.2.0/24
Allow source IP 9.85.115.1, deny all others!

Allow ping to discover network MTU size

- Ping is a useful and frequently used tool to do simple connectivity problem determination
- Today ping on z/OS can let you know if you can reach your target or not, but doesn't help you diagnose MTU problems on the path between you and the target
- z/OS V1R9 will add a "don't fragment" option to ping and have ping correlate any incoming ICMP Messages about fragmentation problems to provide improved diagnostics of where on the path an MTU size problem might exist

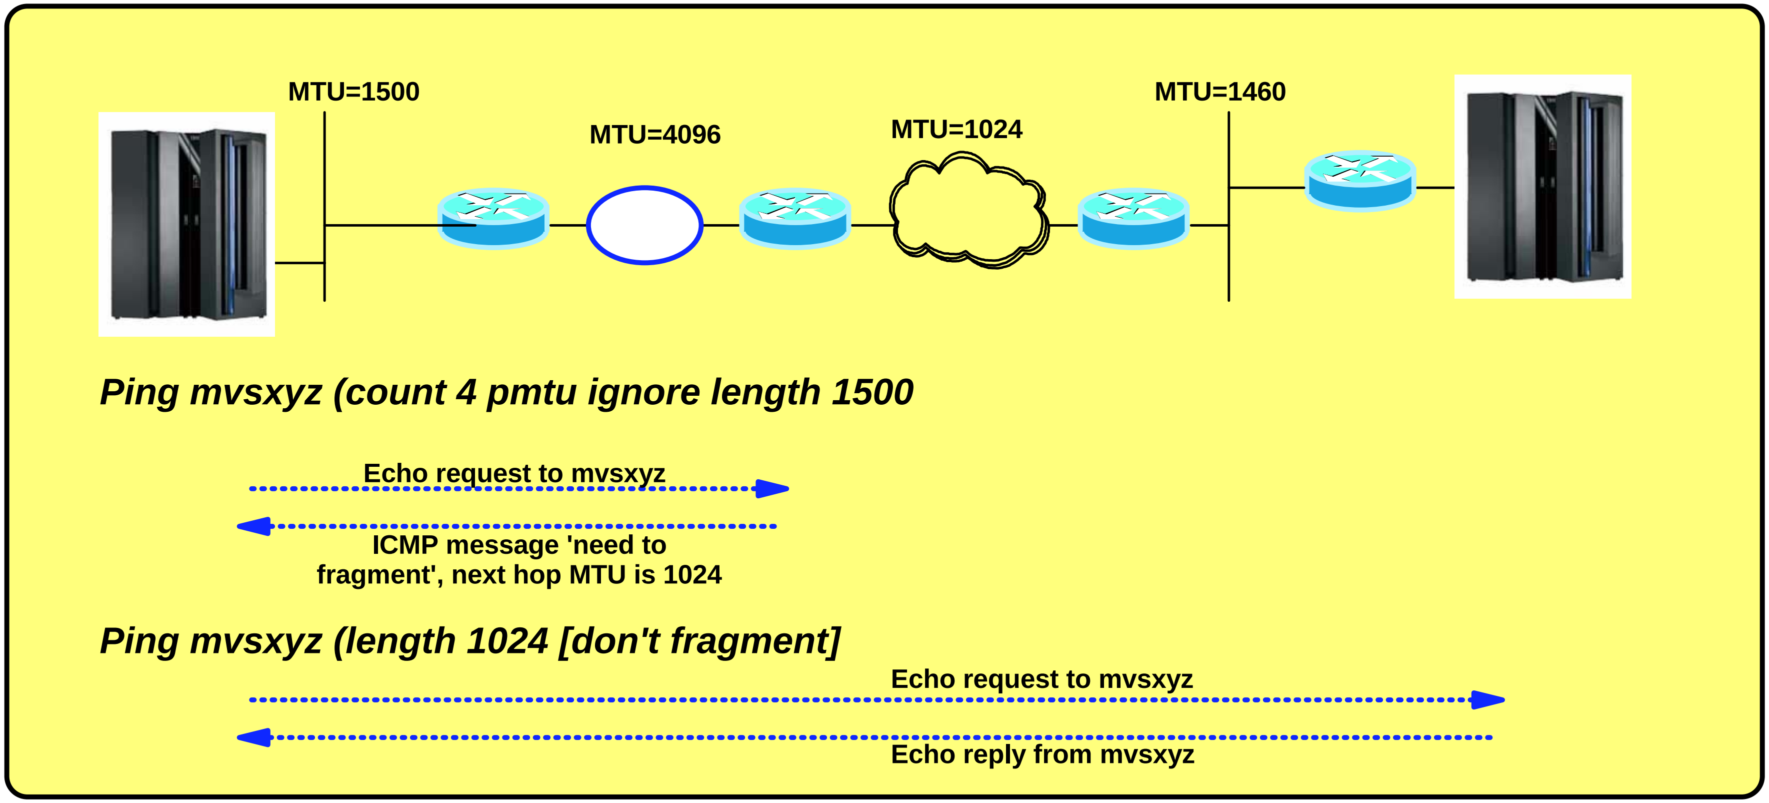

Ping mvsxyz (count 4 pmtu ignore length 1500

Echo request to mvsxyz

ICMP message 'need to fragment', next hop MTU is 1024

Ping mvsxyz (length 1024 [don't fragment]

Echo request to mvsxyz

Echo reply from mvsxyz

SNA/EE

Local MTU discovery for Enterprise Extender

- **Currently an Enterprise Extender (EE) connection will obtain the maximum MTU size permitted for packets being transmitted to a remote EE endpoint at the time that the connection is initially established.**

- **This MTU size is obtained from TCP/IP when VTAM obtains route information for a new EE connection and never changes for the life of the connection.**

- **During the life of this EE connection the TCP/IP stack could change the actual route that is associated with the EE connection, but EE never learns about that and never modifies the initially chosen MTU size.**

- **If EE initializes early during start of a TCP/IP stack, the available route from which EE chooses the MTU size, may not be the optimal route**
  - OMPROUTE may not have started yet and the only routes available might be some default route with an MTU size that doesn't match the MTU size of the actual network the EE traffic later will be sent over
  - Currently, the only way for EE to learn about a better route and MTU size is to stop and restart the EE connection

- **EE will be changed to dynamically learn when local routing information, that pertains to EE connections, is changed - and dynamically modify the MTU size that is used for that EE connection**

- **Note: RTP pipe(s) routed over an EE connection will only learn of changes in the MTU size when the RTP endpoint and the EE connection reside in the same node. EE is not aware of MTU sizes within the network.**

HPR path switch enhancements

- **While HPR path switch works fine in simple environments with a small number of pipes, there has been concern voiced over the impact of a large number of pipes entering path switch state at approximately the same time, such as what occurs during major network failure scenarios.**

- **The major focus areas to be addressed and controlled are:**
  - Optimize the handling of HPR's path switch timers
  - Reduce amount of storage allocation and deallocation occurring during path switch scenarios
  - Implement Path Switch Message Summarization to reduce the number of path switch messages flooding the console
    - HPR Path Switch Summarization reduces the number of path switch message groups VTAM issues across 60-second interval
    - At the end of the 60-second interval, a summarization report is issued on total path switch activity during the interval
    - Avoids WTO buffer shortages

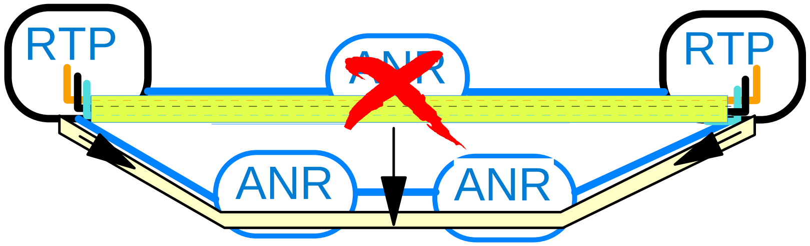

```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/05/06 AT 09:45:14
IST924I -------------------------------------------------------------
IST2192I STARTED   =    12
IST2193I   TGINOP  =    12    SRQTIMER =     0    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   NETWORK =     3  HIGH =     3  MEDIUM =     3  LOW =     3
IST924I -------------------------------------------------------------
IST2196I COMPLETED =     8
IST2195I   NETWORK =     2  HIGH =     2  MEDIUM =     2  LOW =     2
IST924I -------------------------------------------------------------
IST2197I FAILED    =     4
IST2195I   NETWORK =     1  HIGH =     1  MEDIUM =     1  LOW =     1
IST924I -------------------------------------------------------------
IST2198I NETID          STARTED          COMPLETED          FAILED
IST2199I   CPNAME    NET  HI MED LOW   NET  HI MED LOW   NET  HI MED LOW
IST2205I ---------  ---------------  ---------------  ---------------
IST2200I NETA         2   2   2   2     1   1   1   1     1   1   1   1
IST2201I   SSCP2A     1   1   1   1     1   1   1   1     0   0   0   0
IST2201I   SSCP7A     1   1   1   1     0   0   0   0     1   1   1   1
IST2205I ---------  ---------------  ---------------  ---------------
IST2200I NETB         1   1   1   1     1   1   1   1     0   0   0   0
IST2201I   SSCP99     1   1   1   1     1   1   1   1     0   0   0   0
IST924I -------------------------------------------------------------
IST2206I 24 PATH SWITCH EVENTS FOR 3 CPS IN 2 NETIDS
IST314I END
```

Generic resource enhancements

- **Recently we have received a number of requirements for greater control over how generic resources work.**
  - Since the only real control available today is through the Generic Resources exit routine (ISTEXCGR) , those requirements invariably request some additional capability for the exit.
  - However, many of our customers have provided feedback that they do NOT like to write exit code, even simple exit code that just sets/clears a bit to enable/disable a function.
  - Therefore, it seems likely that any functional capability added to the exit will be used by very few customers.
- **This line item will provide these enhancements, along with some of the existing exit capability, via traditional VTAM definitions :**

```
           VBUILD TYPE=GRPREFS
gname1     GNAME  PREFS=LOCAL,OLU,NO_WLM
gname2     GNAME  PREFS=LOCAL,NO_OLU,WLM
gname3     GNAME  PREFS= ...
```

- **In particular, the focus here is to provide control over some or all of the following functions:**
  - Preference for local instance or application (current exit function)
  - Use session counts only, instead of WLM (current exit function)
  - Do not drive exit (current exit function)
  - Favor instance on OLU host (CLSDST PASS scenario)
  - Choose whether to use a local instance of an application without affecting other generic resource name resolutions for applications.
- **With the implementation of this line item, direction is that NO further enhancements will be made to ISTEXCGR function (but current exit will continue to be supported)**

Improved MPC activation processing

- **Activation of an MPC group fails when a minimum number of subchannels are not available**
  - At least one online read and one online write subchannel must be available
- **So far there has been no automated mechanism to recover after subchannels become available**
  - Manual reactivation by operator or automations scripts
- **Situation typically arises whenever a FICON-connected host is down**
- **z/OS CS V1R9 implements enhancements to the MPC activation process:**
  - Activation of an MPC group that fails to meet the one read/one write requirement is put on hold, provided any needed read and/or write subchannel is an offline CTC or one that has no valid path available to the connecting host.
  - The suspension continues until the required minimum number of subchannels becomes available or the group is deactivated.
  - New messages signal when the hold begins and when activation resumes.
- **The display of an MPC group indicates when its activation is on hold.**
  - Other existing output in that display identifies the offline subchannels, so appropriate action can be taken to bring enough of them online to cause activation of the MPC group to complete.
- **Subarea MPC groups must support MPC Dynamics**
  - Use MPCDYN=YES specification
- **This function is enabled by default.**
  - A new start option, MPCACT=WAIT/NOWAIT can be used to disable the function

MPC activation sequence with missing subchannels

```
V NET,ACT,ID=MPCLN1,E
 IST097I VARY ACCEPTED
 IST1631I MPCLN1 SUBCHANNEL 0F1B OFFLINE
 IST1631I MPCLN1 SUBCHANNEL 0F1C OFFLINE
 IST1631I MPCLN1 SUBCHANNEL 0F1D OFFLINE
 IST2219I MPCLN1 ACTIVATION WAITING FOR MINIMUM NUMBER OF DEVICES


D NET,ID=MPCLN1,E
 IST097I DISPLAY ACCEPTED
 IST075I NAME = MPCLN1, TYPE = LINE
 IST486I STATUS= PALNK, DESIRED STATE= ACTIV
         :
 IST2219I MPCLN1 ACTIVATION WAITING FOR MINIMUM NUMBER OF DEVICES
 IST1221I WRITE DEV = 0F1A STATUS = RESET      STATE = ONLINE
 IST1221I WRITE DEV = 0F1B STATUS = RESET      STATE = OFFLINE
 IST1221I READ  DEV = 0F1C STATUS = RESET      STATE = OFFLINE
 IST1221I READ  DEV = 0F1D STATUS = RESET      STATE = OFFLINE
         :
 IST396I LNKSTA   STATUS     CTG GTG  ADJNODE ADJSA    NETID    ADJLS
 IST397I MPCPU1   NEVAC        1   1
 IST314I END


V 0F1B,ONLINE
 IEE302I 0F1B        ONLINE
V 0F1C,ONLINE
 IEE302I 0F1C        ONLINE
 IST2220I MPCLN1 ACTIVATION RESUMING - ONLINE DEVICES DETECTED
 IST093I MPCLN1 ACTIVE
```

- **Initial activation fails because not enough subchannels are online to meet the minimum requirements.**

- **A display indicates the status of the activation request and which subchannels are online and which are not available.**

- **When 0F1C is varied online, the minimum requirements are met and the activation completes.**

Enhancements to better control cross-subnet searches

- **The order of all cross-subnet searching has so far been the same for all NETIDs**
  - Controlled by the BNORD start option
- **A method to selectively restrict searches to individual nodes during cross-sub network searching has not been available**
  - Only way to impact was to code BNDYN=NONE and NEXTCP statements
- **z/OS CS V1R9 will support configuration options to specify the order of cross-subnet searching for each NETID**
  - BNORD supported on NETWORK statement
    - PRIORITY - preference given to nodes with recent successful searches
    - DEFINED - searches are done in the order specified in the list
- **z/OS CS V1R9 will also provide a method to restrict searches to a specific node**
  - New OMITCP operand added to the NEXTCP statement
    - Yes - node specified on the CPNAME operand not included in Subnetwork Routing List (SRL)
    - No - node specified on the CPNAME operand included in SRL
- **The enhancement allows you to:**
  - Control of the order of cross-subnet searching by NETID
  - Selectively restrict searches to a specific node

Additional EE/SNA enhancements

- **HPR Message Enhancements - additional information about RTP pipes on activation and deactivation**
  - New start option HPRITMSG = BASE/ENHANCED

```
IST1488I ACTIVATION OF RTP CNR00001 AS ACTIVE TO NETA.SSCP2A
IST1963I APPNCOS = #INTER - PRIORITY = HIGH
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I 21   NETA.SSCPAA        INTERM       RTP
IST1461I 21   NETA.SSCP2A        INTERM       RTP
IST314I END
```

- **EE Enhanced Packet Loss Tolerance**
  - Implements several enhancements to reduce the effect of packet loss on EE:
    - More aggressive reporting of gaps - similar to TCP Fast Retransmit
    - Support "burst" timers as small as 1 mSec
      - Paces the data across the connection better, reducing the height of the bursts
    - Additional internal code optimizations

APAR OA20923 for z/OS CS V1R8

- **V1R9 provides the ability to specify a different set of Logical Data Link Control parameters for each local EE VIPA:**
  - LIVTIME, SRQRETRY, and SRQTIME
  - LDLC parameters are specified on the PORT and/or on the GROUP statement where the EE VIPA is specified

```
XCAEE     VBUILD TYPE=XCA
PORTEE    PORT   MEDIUM=HPRIP,LIVTIME=(10,20),SRQTIME=15,SRQRETRY=3
GRPEE1    GROUP  DIAL=YES,AUTOGEN=(10,E,X),IPADDR=10.1.1.1,            *
                 CALL=INOUT,ISTATUS=ACTIVE
GRPEE2    GROUP  DIAL=YES,AUTOGEN=(10,G,V),IPADDR=10.1.1.2,            *
                 CALL=INOUT,ISTATUS=ACTIVE,LIVTIME=(10,60),SRQRETRY=5
```

Reliability, Availability and Serviceability (RAS)

z/OS CS V1R9 health-checker extensions

- **z/OS CS V1R8 implemented initial support for and use of the z/OS Health-checker infrastructure:**
  - The default size of the TCP receive buffer size (the TCPMAXRCVBUFRSIZE option in the TCP/IP Profile)
  - The default set of options for CTRACE
  - Maximum amount of fixed CSM storage (the MAXFIX option in IVTPRMxx)
  - Maximum amount of ECSA CSM storage (the MAXECSA option in IVTPRMxx)

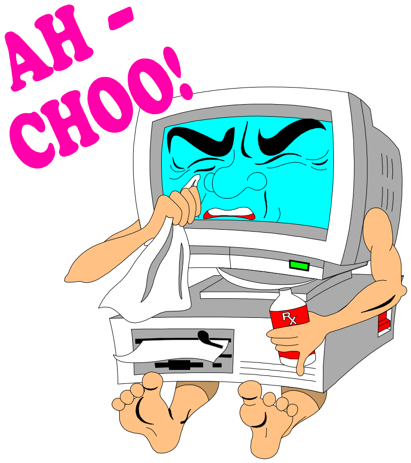

- **z/OS CS V1R9 extends support for and use of the z/OS Health-checker infrastructure:**
  - If a TCP/IP stack has DYNAMICXCF defined in combination with GLOBALCONFIG SYSPLEXMONITOR without the RECOVERY option, a warning will be issued since this is not a recommend best practice combination. In this case, the automated RECOVERY option should be used.
  - Various VTAM checks are being added:
    - Check that the VIT SIZE is not lower than the default value of 999
    - Check that the VIT options PSS (VTAM's Process Scheduling Services) and SMS (VTAM's Storage Management Services) are turned on - they are always needed to service a problem
    - Check that DSPZISE is at least 5 (50 Megabytes)
    - Check to see if someone has OPT=ALL specified (requesting all VIT trace options be turned on) - this may not be optimal, unless requested by VTAM services personnel
    - Check that the Enterprise Extender QDIO/iQDIO buffer pools (T1Buf and T2Buf pools) are of reasonable sizes

Statements of Direction

Statements of general direction of interest to z/OS CS community

Traffic Regulation Policies

- z/OS V1.9 is planned to be the last release of z/OS Communications Server which will support the configuration of Traffic Regulation (TR) policy as part of the Quality of Service discipline. The TR configuration function remains supported, but IBM recommends that you implement it as part of the Intrusion Detection Services (IDS) policy configuration made available in z/OS V1.8. This change is only for the TR policy configuration. The TR policy functions themselves remain unaffected. For more information, please refer to z/OS V1.8 Communications Server's IP Configuration Guide, chapter 16, "Intrusion Detection Services", and IP Configuration Reference, chapter 23, "Intrusion Detection Services policy".

JES3 NJE over TCP/IP

- IBM intends to improve the availability date for the support of JES3 NJE communications using TCP/IP which is planned to be made available in z/OS V1.8. as originally planned and announced in the Software Announcement 206-190 (RFA43269), dated August 8, 2006, from 1H2007 to March 30, 2007. This function includes support for IPv6, secure sockets (SSL/TLS), and all the NJE constructs (ENDNODE, SUBNET, Store-and-Forward) supported by the owning JES. This new support is in addition to the SNA and BSC protocols currently supported by JES3. In order to use the NJE/TCP support, both sides of the connection are required to support NJE/TCP. In order to enable JES3 NJE over TCP/IP in z/OS V1.8, you must install the PTF for APAR OA16527 and its prerequisites.

All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

For more information....

URL	Content
http://www.ibm.com/servers/eserver/zseries	IBM eServer zSeries Mainframe Servers
http://www.ibm.com/servers/eserver/zseries/networking	Networking: IBM zSeries Servers
http://www.ibm.com/servers/eserver/zseries/networking/technology.html	IBM Enterprise Servers: Networking Technologies
http://www.ibm.com/software/network/commserver	Communications Server product overview
http://www.ibm.com/software/network/commserver/zos/	z/OS Communications Server
http://www.ibm.com/software/network/commserver/z_lin/	Communications Server for Linux on zSeries
http://www.ibm.com/software/network/ccl	Communication Controller for Linux on zSeries
http://www.ibm.com/software/network/commserver/library	Communications Server products - white papers, product documentation, etc.
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/support	Communications Server technical Support
http://www.ibm.com/support/techdocs/	Technical support documentation (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)