z/OS Communications Server V1R7
What's New

Linda Harrison
lharriso@us.ibm.com

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:

- AIX®
- AnyNet®
- AS/400®
- Candle®
- CICS®
- CICSPlex®
- CICS/ESA®
- DB2®
- DB2 Connect™
- DPI®
- DRDA®
- e business(logo)®
- ESCON®
- eServer™
- ECKD™
- FFST™

- GDDM®
- GDPS®
- HiperSockets™
- IBM®
- Infoprint®
- IMS™
- IP PrintWay™
- iSeries™
- Language Environment®
- MQSeries®
- MVS™
- MVS/ESA™
- NetView®
- OS/2®
- OS/390®
- Parallel Sysplex®

- PrintWay™
- PR/SM™
- pSeries®
- RACF®
- Redbooks™
- Redbooks (logo)™
- S/390®
- System/390®
- ThinkPad®
- Tivoli®
- Tivoli (logo)®
- VM/ESA®
- VSE/ESA™
- VTAM®
- WebSphere®
- xSeries®

- z/Architecture™
- z/OS®
- z/VM®
- zSeries®

- 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. Sysplex Enhancements

2. Application Enhancements

3. Security

4. Enterprise Extender and SNA Enhancements

5. IPv6 on z/OS Communications Server

6. Statements of Direction

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 future z/OS Communications Server enhancements

➤ The focus of this presentation is the Communications Server in V1R7

z/OS
CS

Note: Plans are subject to change!

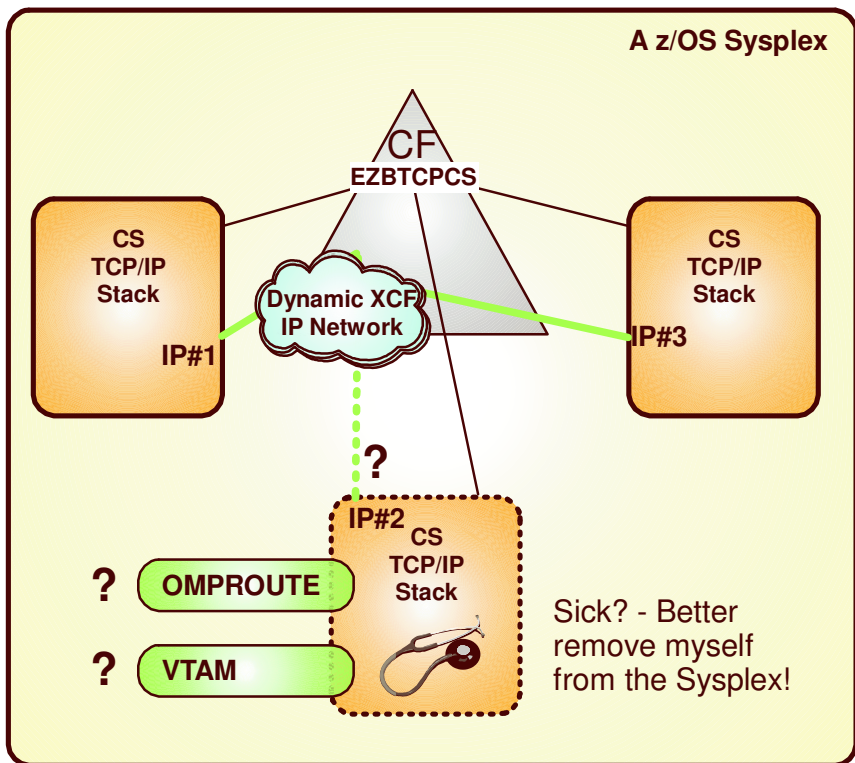| z/OS V1R2	z/OS v1R4		z/OS V1R5	z/OS V1R6	z/OS V1R7
9/01	9/02	9/03	1Q04	2H04	2H05
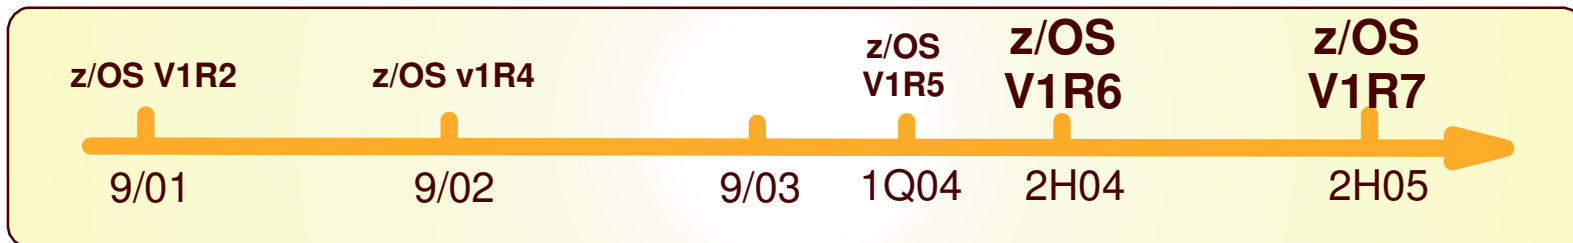
Sysplex Enhancements

➢ **Focus areas:**
  ▸ Extending the autonomic behavior in error scenarios for the IP Sysplex
  ▸ Improved quality in workload distribution decisions made by Sysplex Distributor
  ▸ Improved operator control

TCP/IP Sysplex autonomics to let TCP/IP proactively handle error conditions

A z/OS Sysplex

CF
EZBTCPCS

CS TCP/IP Stack — IP#1

Dynamic XCF IP Network

CS TCP/IP Stack — IP#3

?

IP#2
CS TCP/IP Stack

? OMPROUTE

? VTAM

Sick? - Better remove myself from the Sysplex!

The assumption is that if a TCP/IP stack determines it can no longer perform its Sysplex functions correctly, it is better for it to leave the TCP/IP XCF group and by doing so, signal the other TCP/IP stacks in the Sysplex that they are to initiate whatever recovery actions have been defined, such as moving dynamic VIPA addresses or removing application instances from distributed application groups.
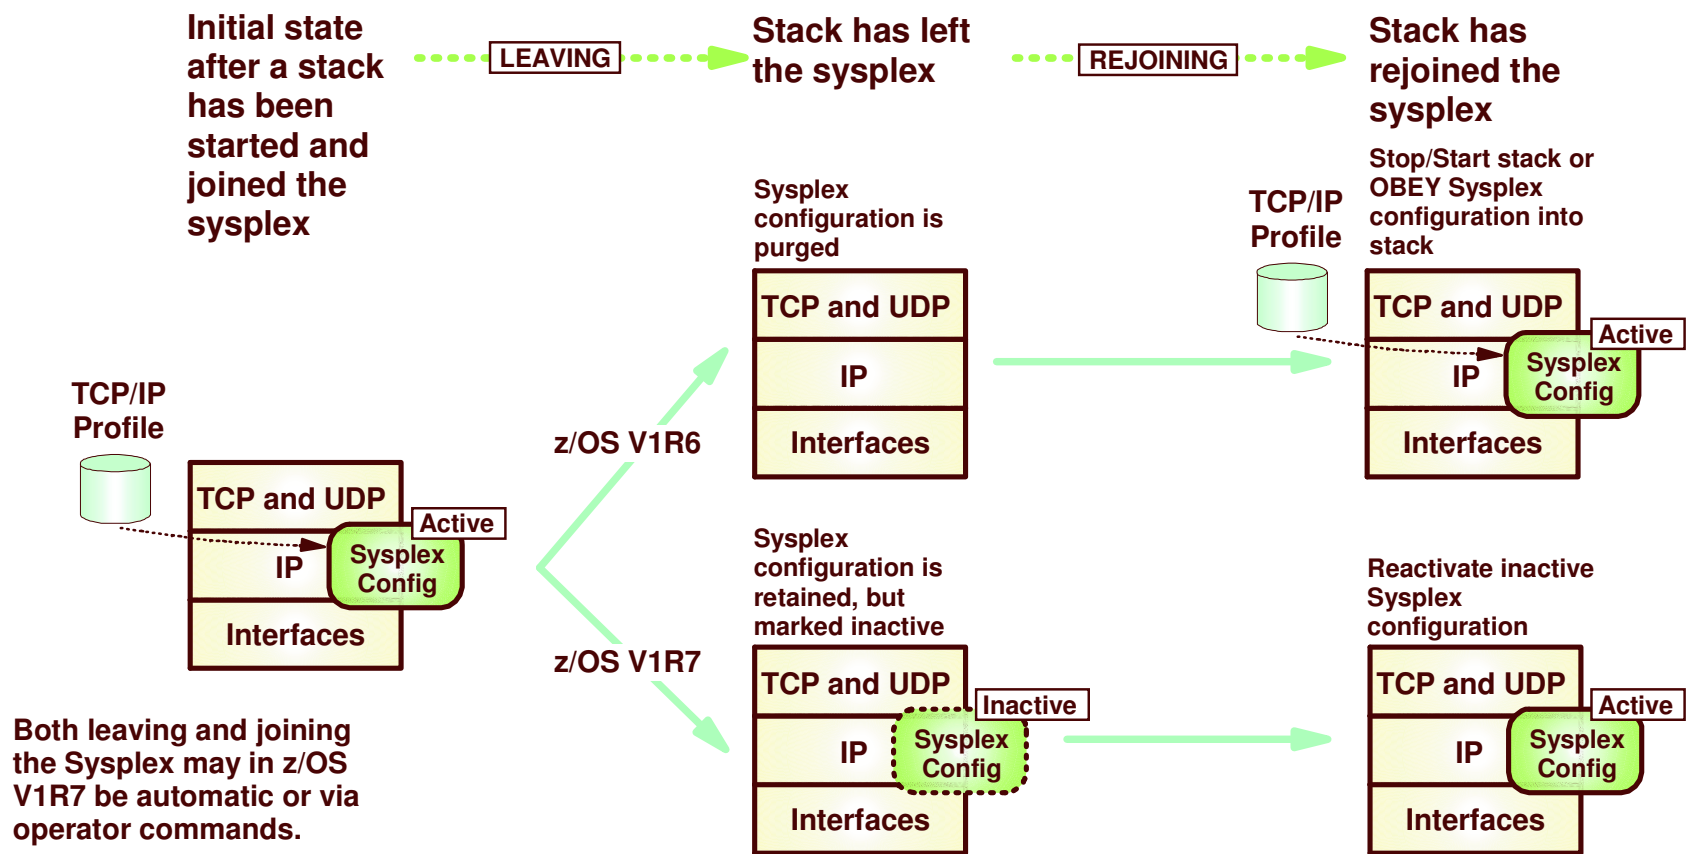
AUTOREJOIN options are being added in z/OS V1R7

- ➤ Autonomic functions to reduce single point of failure for distributed applications in a sysplex
    - ► Monitor CS health indicators
        - – Storage usage - CSM, TCPIP Private & ECSA
    - ► Monitor dependent networking functions
        - – OMPROUTE availability
        - – VTAM availability
        - – XCF links available
    - ► Monitor Communications Server component-specific functions

- ➤ Monitors determine if this TCPIP stack will remove itself from the sysple and allow a healthy backup to take ownership of the sysplex duties (own DVIPAs, distribute workload)

- ➤ Monitoring is always done, but configuration controls in the TCPIP Profile determine if the TCPIP stack will remove itself from the sysplex.

```
GLOBALCONFIG SYSPLEXMONITOR TIMERSECS
seconds RECOVERY|NORECOVERY
DELAYJOIN|NODELAYJOIN
AUTOREJOIN|NOAUTOREJOIN
```

- ➤ *Timersecs* - used to determine duration of the troubling condition before issuing messages or leaving the sysplex (if Recovery)
- ➤ *RECOVERY* - TCPIP removes itself from the sysplex. Recommended an is the default value.
- ➤ *NORECOVERY* - TCPIP does not remove itself from the sysplex.
- ➤ *DELAYJOIN* - Delay joining Sysplex until OMPROUTE is up
- ➤ *NODELAYJOIN* - Join Sysplex immediately
- ➤ *AUTOREJOIN* - Rejoin when condition is cleared
- ➤ *NOAUTOPREJOIN* - Let an operator decide when to rejoin

TCP/IP stacks leaving and rejoining the sysplex

Initial state after a stack has been started and joined the sysplex

- - → LEAVING - - → **Stack has left the sysplex**

- - → REJOINING - - → **Stack has rejoined the sysplex**

TCP/IP Profile

TCP and UDP

IP Sysplex Config Active

Interfaces

z/OS V1R6

Sysplex configuration is purged

TCP and UDP

IP

Interfaces

Stop/Start stack or OBEY Sysplex configuration into stack

TCP/IP Profile

TCP and UDP

IP Sysplex Config Active

Interfaces

z/OS V1R7

Sysplex configuration is retained, but marked inactive

TCP and UDP

IP Sysplex Config Inactive

Interfaces

Reactivate inactive Sysplex configuration

TCP and UDP

IP Sysplex Config Active

Interfaces

Both leaving and joining the Sysplex may in z/OS V1R7 be automatic or via operator commands.

➢ **Leaving the Sysplex in z/OS V1R6, purges sysplex configuration data from the stack's internal configuration blocks.**
  ▸ To rejoin the sysplex, the sysplex configuration data must be reapplied to the stack's active configuration through a restart or an OBEY command
➢ **In z/OS V1R7, a stack's sysplex configuration data will be retained in an inactive status when a stack leaves the sysplex**
  ▸ The inactive sysplex configuration data will be shown on the NETSTAT VIPADCFG report as inactive
  ▸ Rejoining the sysplex will then reactivate the currently inactive sysplex configuration data

How to rejoin

➤ **Rejoin can be**

▸ Automatic
  – GLOBALCONFIG SYSPLEXMONITOR AUTOREJOIN
    ● The stack will rejoin the sysplex, when the problem that caused it to automatically leave the sysplex has been relieved
    ● Is only supported in combination with the SYSPLEXMONITOR RECOVERY option (leave the sysplex automatically if a problem is detected)
    ● Automatic rejoin is triggered by the events that clear the error condition (XCF links back up, OMPROUTE restarted, etc.)
    ● Bounce prevention logic built into the storage condition logic if storage limits are set on GLOBALCONFIG

▸ Operator command initiated
  – VARY TCPIP,[stackname],SYSPLEX,JOINGROUP
    ● Matching the vary command to leave the sysplex that was introduced in z/OS V1R6
    ● Allowing full operator control over when to leave and when to rejoin the sysplex

▸ OBEYing a new sysplex configuration into a stack that currently has left the sysplex
  – VARY TCPIP,[stackname],OBEY,DSN=my.sysplex.conf
    ● Will override (replace) currently inactive sysplex configuration and rejoin the sysplex
    ● Provides compatibility for operations procedures that were established prior to z/OS V1R7

➤ **Rejoin will work under the same conditions as the initial join**

▸ If DELAYJOIN is configured, the stack will ensure OMPROUTE is up and fully functional before the rejoin will take place

SD and non-disruptive DVIPA movement forwarding of IP packets

➢ **The reasons why SD and non-disruptive DVIPA movement initially required use of DynamicXCF were:**

  ► The forwarding of packets is done without using NAT - the destination address never changes
    – This is known as MAC-level forwarding, or dispatch mode balancing
    – The destination address (the DVIPA) reside in the HOME lists of all stacks that are potential targets
  ► This mode of forwarding requires that the destination host is exactly one hop away, or in other words that all members of the z/OS Sysplex are attached to a single shared IP network
    – DynamicXCF was a convenient way to ensure that this requirement was always met with minimal customer configuration requirements

➢ **Removing the requirement for DynamicXCF means that we cannot guarantee that the target stack we're forwarding a packet to is exactly one hop away**

  ► When DynamicXCF is not used, TCP/IP will use GRE (Generic Routing Encapsulation) to forward the packet to a unique IP address on the target stack
  ► The address to forward the packet to will be configured using a new configuration option in the VIPADYNAMIC block
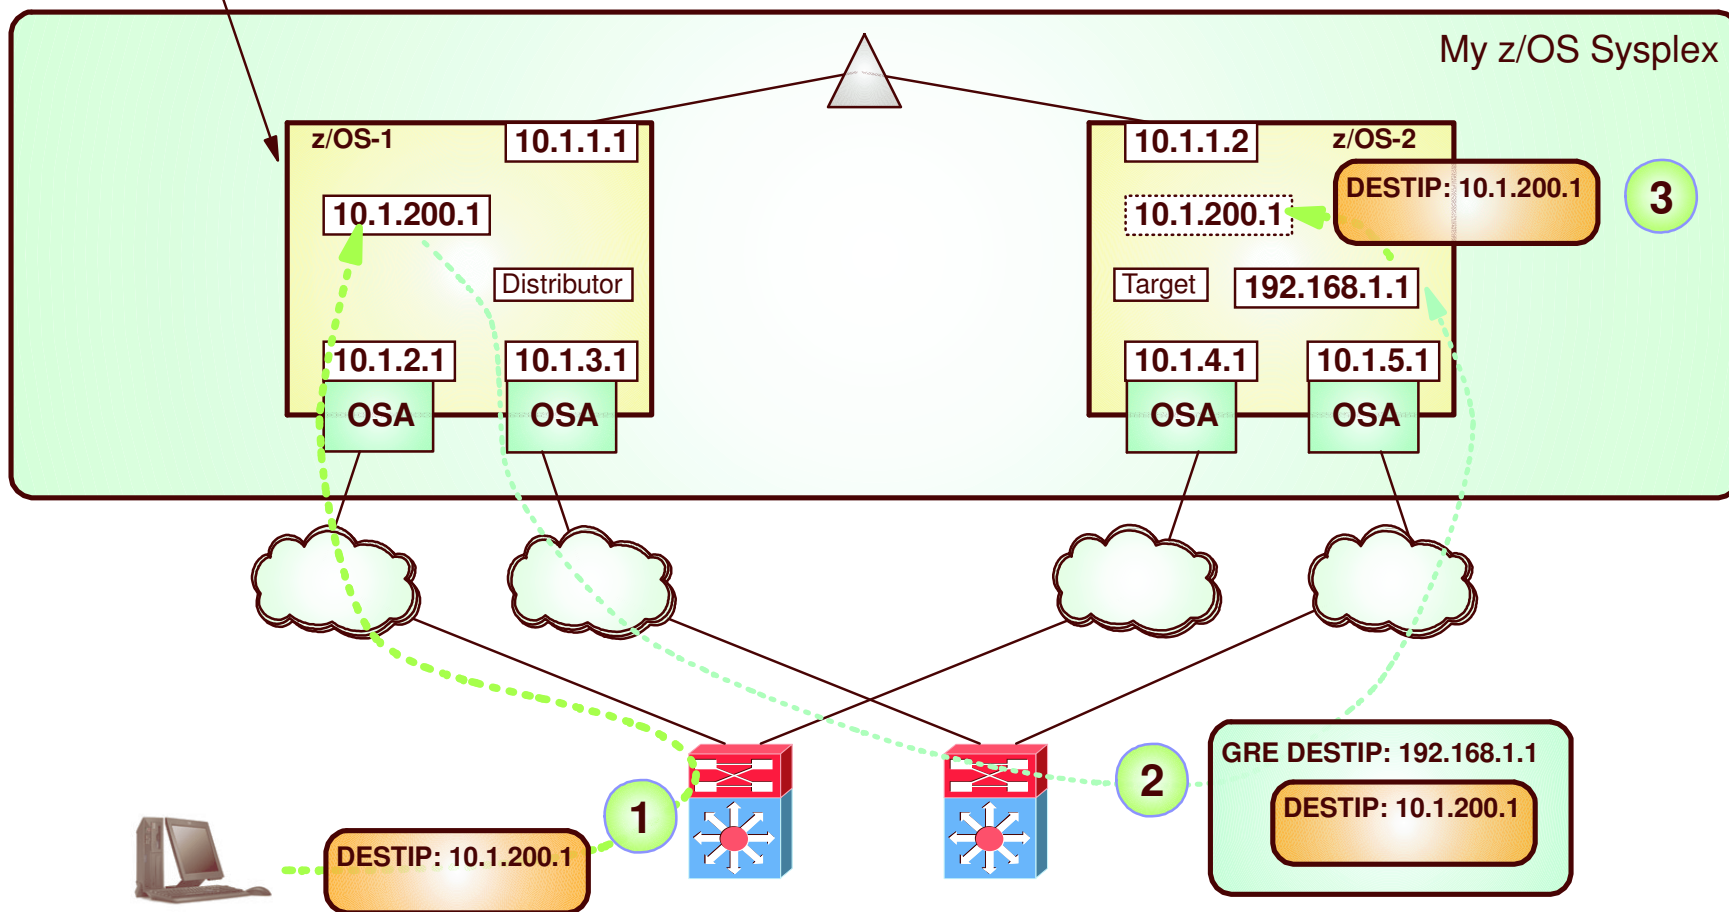
      ● VIPAROUTE DEFINE dynxcfIPaddress targetIPaddress

  ► Whenever SD or non-disruptive DVIPA is to sent a packet to a given DynamicXCF IP address and a VIPAROUTE statement is configured with that DynamicXCF IP address, a GRE envelope will be wrapped around the original packet with the destination IP address from the VIPAROUTE statement and normal IP routing logic will forward that packet (DATAGRAMFWD is *not* required)
    – Path can change based on actual network availability
    – Multipathing is supported
    – High-speed network technologies are available for SD and non-disruptive DVIPA movement forwarding

SD and non-disruptive DVIPA movement forwarding via non-DynamicXCF interfaces

```
VIPADISTRIBUTE DEFINE 10.1.200.1 PORT 23 DESTIP ALL
VIPAROUTE DEFINE 10.1.1.2 192.168.1.1
```
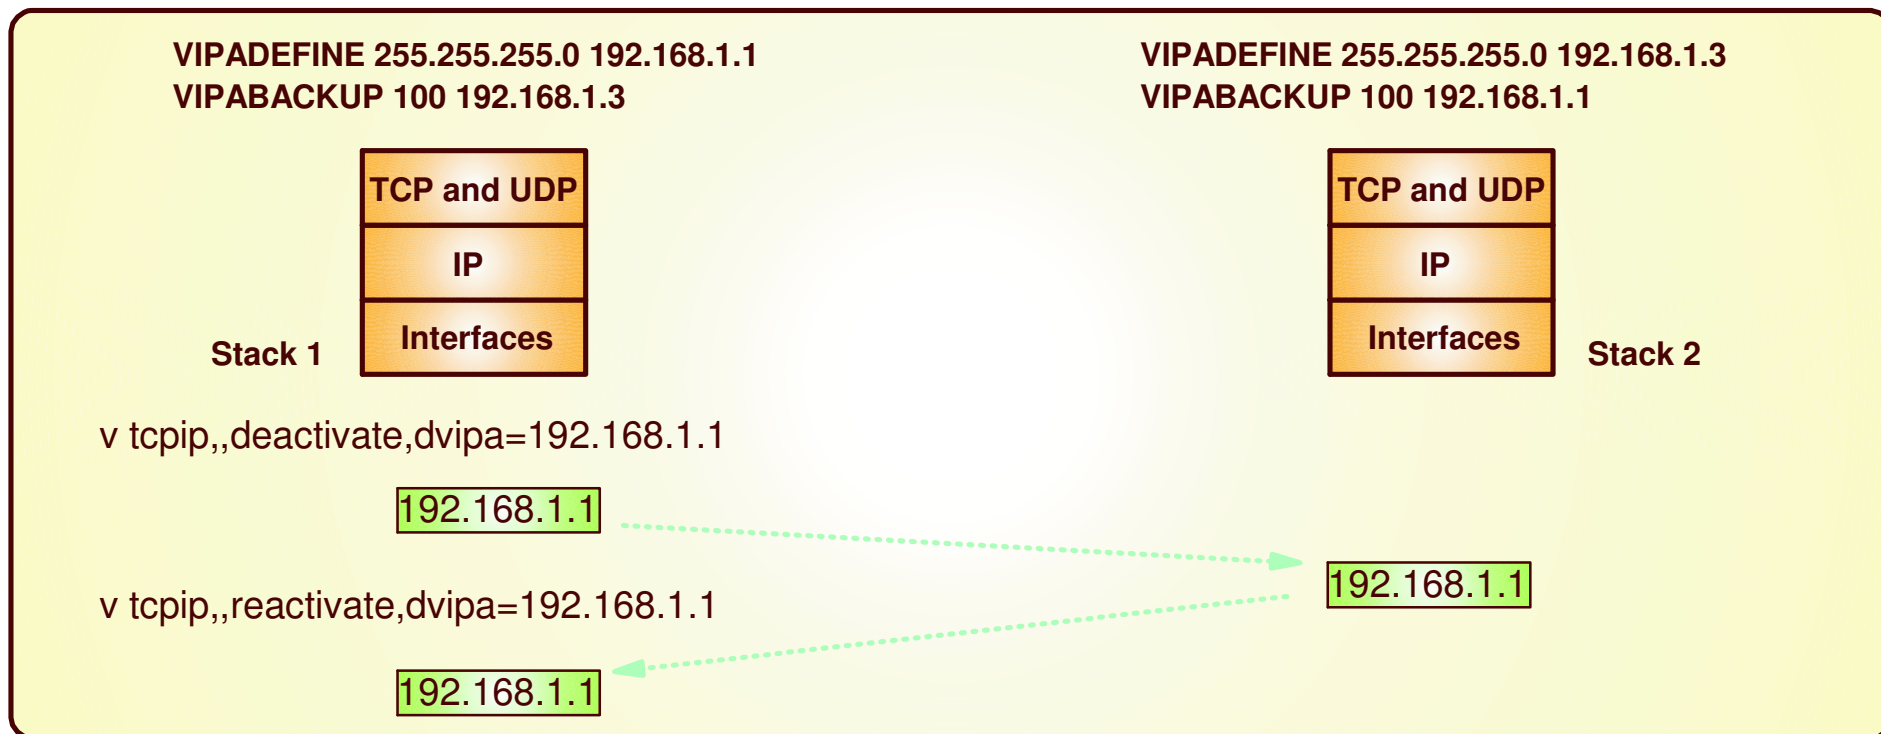
PATHMTUDISCOVERY is in general recommended
- ► On remote nodes to learn max packet size (including the GRE hop)
- ► On z/OS if the directly connected network is a Gigabit Ethernet network that uses jumbo frames

My z/OS Sysplex

z/OS-1 **10.1.1.1**

10.1.200.1

Distributor

10.1.2.1 **10.1.3.1**
OSA   OSA

10.1.1.2 z/OS-2

DESTIP: 10.1.200.1 3

10.1.200.1

Target **192.168.1.1**

10.1.4.1 **10.1.5.1**
OSA   OSA

1

DESTIP: 10.1.200.1

2

GRE DESTIP: 192.168.1.1

DESTIP: 10.1.200.1

Improved operations: operator-initiated movement of individual stack-managed dynamic VIPA addresses

VIPADEFINE 255.255.255.0 192.168.1.1
VIPABACKUP 100 192.168.1.3

VIPADEFINE 255.255.255.0 192.168.1.3
VIPABACKUP 100 192.168.1.1

```
TCP and UDP

IP

Interfaces
```
Stack 1

```
TCP and UDP

IP

Interfaces
```
Stack 2

v tcpip,,deactivate,dvipa=192.168.1.1

`192.168.1.1`

`192.168.1.1`

v tcpip,,reactivate,dvipa=192.168.1.1

`192.168.1.1`

➢ **Deactivate**
  ▸ DVIPA is deactivated and a configured backup stack will takeover the DVIPA
  ▸ Backup DVIPA can be deactivated also removing eligibility as a backup

➢ **Reactivate**
  ▸ Original owner can regain ownership
  ▸ Can also reactivate a backup DVIPA that's been deactivated
  ▸ Prior to these commands, Vary obey files were needed to cause a DVIPA takeover
  ▸ These commands can't be used on a DVIPA's created from VIPARANGE with bind, ioctl(), or the Modvipa utility

Improved operations: operator-initiated quiesce and resume of individual server applications or full target systems

- ➤ **Ability to quiesce a system or an application prior to shutdown**
  - ▸ Planned maintenance scenarios of system or application
    - – Allows existing systems or applications to drain work queue prior to shutdown
  - ▸ Relieve temporary constraints of resources on target system
  - ▸ Temporary - Does not affect Sysplex Distributor's permanent configuration
  - ▸ Issued on target system being affected
  - ▸ Can also be used to control individual server applications in a SHAREPORT group
  - ▸ Only way to achieve similar capability earlier was via temporary configuration changes based on OBEYFILE commands

- ➤ **VARY TCPIP,,SYSPLEX,QUIESCE,options**
  - ▸ TARGET - Quiesces all applications on target stack.
  - ▸ PORT=xxx - Quiesce all applications bound to the specified port on this stack
    - – JOBNAME=jobname - Allows quiesce of a single application in SHAREPORT group
    - – ASID=asid - Further qualify job being quiesced (such as when dealing with duplicate jobnames)
  - ▸ No new TCP connections sent to the quiesced target (stack or application)
    - – For all Distributed DVIPAs that the entity is a target for
  - ▸ Existing TCP connections are maintained (or in other words, the process is non-disruptive)

- ➤ **VARY TCPIP,,SYSPLEX,RESUME,options**
  - ▸ TARGET|PORT|JOBNAME|ASID
  - ▸ Allows identified target stacks and/or applications to once again be targets for distribution

Improved workload distribution quality focus in z/OS V1R7

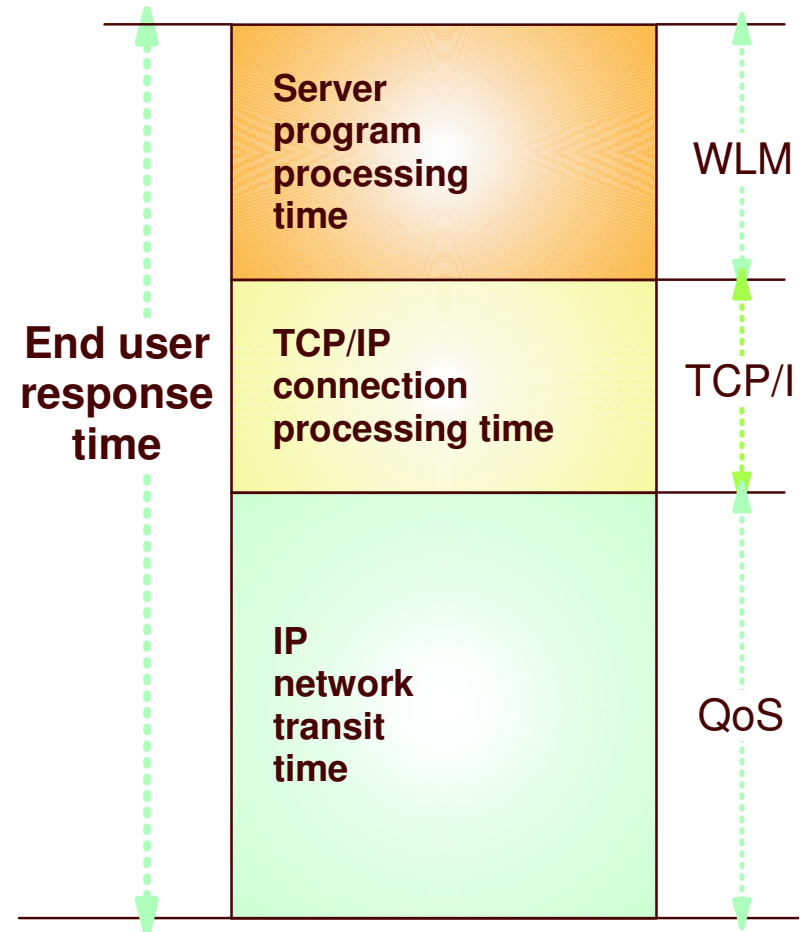➢ **Sysplex Distributor uses Server-specific WLM Interfaces to determine if target server is meeting its goal**
  ▸ Extracts WLM recommendations for each distributed server to determine which server(s) get new connections
  ▸ More precise than existing WLM method which uses recommendations base upon displaceable capacity of the system

➢ **Sysplex Distributor will detect target server unresponsiveness**
  ▸ Target stacks push key TCP/IP "health" statistics for target application(s) to distributor, such as number of connections dropped due to backlog.
  ▸ When load balancing, the distributor uses these indicators along with values for WLM and QoS to determine which stack gets the connection
  ▸ Strengthens overall evaluation of a server's health

➢ **Allow Sysplex Distributor to route over any available route**
  ▸ Removes the need for sysplex distributor to route only over coupling facility links
  ▸ Allows for use of high-speed links such as ethernet and OSA Express QDIO interfaces
  ▸ Can use any interface on the target except can not specify a dynamic VIPA

End user response time

Server program processing time	WLM
TCP/IP connection processing time	TCP/I
IP network transit time	QoS

Adresses some storm-drain scenarios, but not all.

Sysplex Distributor use of WLM and QoS feedback in z/OS V1R7

➢ **Workload Manager feedback has so far been a reflection of how much displaceable capacity the target LPARs have available at any point in time**

  ▸ It has not been a reflection of how well the individual server address space meets its WLM performance goals

➢ **In z/OS V1R7, WLM will provide new interfaces that will allow Sysplex Distributor to query performance information for individual address spaces**

  ▸ The information from WLM will reflect how well the address space meets its WLM performance goals
    – Base weight is still LPAR displaceable capacity.
    – If server address space meets its WLM performance goal, WLM will report the LPAR displaceable capacity based weight
    – If server address space does not meet its WLM performance goals, WLM will augment the LPAR displaceable capacity based weight with a fraction that represents how much below the goal this address space currently performs
  ▸ Sysplex Distributor will in z/OS V1R7 make use of these enhanced WLM interfaces to obtain server-specific WLM recommendations

➢ **Sysplex Distributor will continue to support modification of the WLM recommendations based on feedback from the Policy Agent about QoS:**

  ▸ Loss ratio
  ▸ Time-out
  ▸ Connection limit threshold

Sysplex Distributor to factor in TCP/IP connection processing performance in z/OS V1R7

➢ **Sysplex Distributor will in z/OS V1R7 factor in new weight fractions that reflect how well TCP/IP connection processing is performing:**

  ▸ Lost forwarded connections to the target stack (distributing stack forwards connection request, but doesn't receive a notification that target stack received the connection request)
  ▸ Target stack unable to actually establish a connection with the client (can't complete 3-way TCP handshake)
  ▸ Connections dropped due to server backlog queue full condition
  ▸ A server instance building up a backlog queue while appearing to be "hanging", but not yet dropping connections due to backlog queue full condition

➢ **SHAREPORT logic will also in z/OS V1R7 be enhanced to factor in how well the individual server instances process new connections**
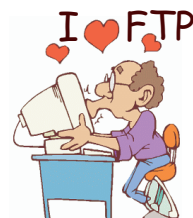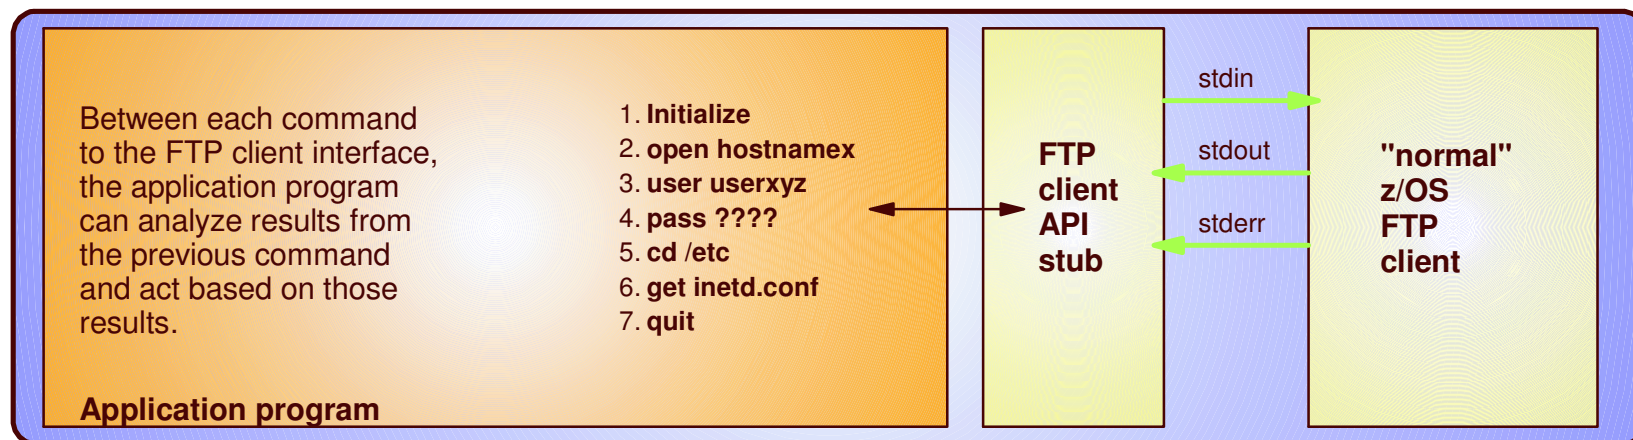
Application Enhancements

➤ **Focus areas:**
  ‣ FTP
  ‣ CICS Sockets

z/OS FTP Client programming interface for improved automation and integration of z/OS file transfers

➤ Provides an interface that allows an application to programatically invoke the FTP client on z/OS from common environments (unix shell, TSO, or MVS batch job)

➤ Characteristics of the interface:
  ‣ z/OS V1R6 provides a callable interface to be used from Assembler, Cobol, PL/I (or any z/OS supported programming language that supports a call interface) -
    – *z/OS V1R7 adds a C API version of this programming interface*
  ‣ Interface is reentrant and does support multiple parallel FTP client sessions by tasks within an address space
  ‣ For communication between the program and the interface, a simple set of commands and data areas are used. (Mappings for common programming languages are provided)
  ‣ Both blocking (wait for a response), and non-blocking (polling-mode) calls are supported
  ‣ In non-blocking mode, progress notifications can be returned to the calling application as the transfer progresses
  ‣ The simple commands tell the interface what to do, for example: initialize, terminate, execute an FTP client command, process output from the FTP client command that was executed, poll for command completion.
  ‣ Results are returned as structured fields in communication area control blocks (return codes from interface and server replies or possibly local command) along with free-format replies from the FTP client code
  ‣ Debugging options are provided

Between each command to the FTP client interface, the application program can analyze results from the previous command and act based on those results.

Application program

1. **Initialize**
2. **open hostnamex**
3. **user userxyz**
4. **pass ????**
5. **cd /etc**
6. **get inetd.conf**
7. **quit**

FTP client API stub

stdin
stdout
stderr

"normal" z/OS FTP client

I ♥ FTP

CICS Sockets focus in z/OS V1R7 for improved performance and security capabilities
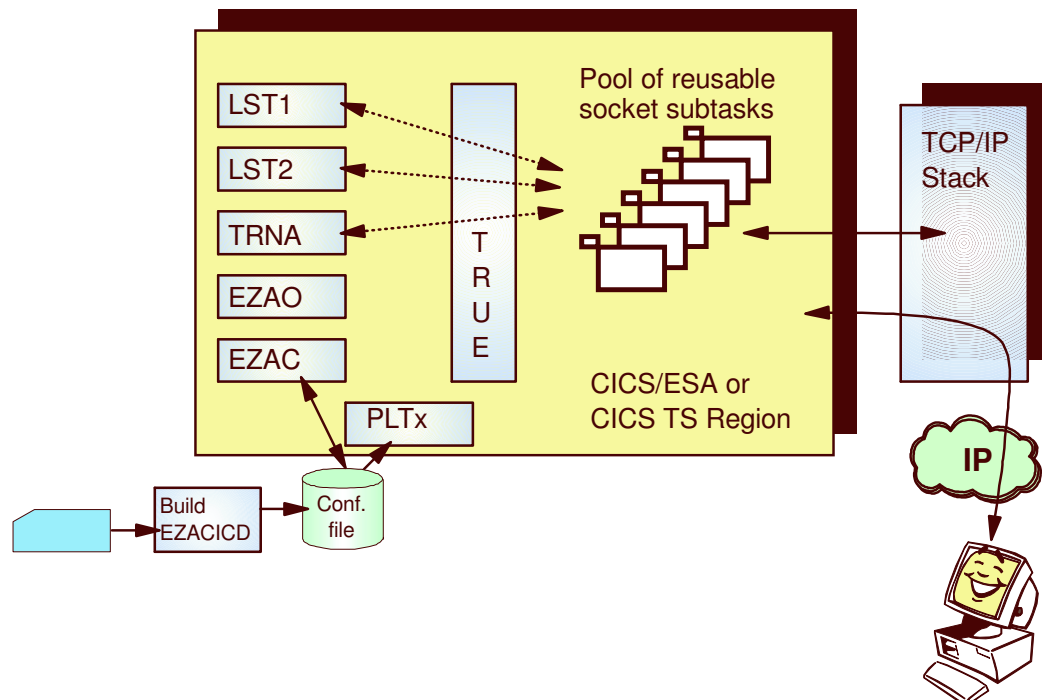
➤ **Several performance improvements:**
  ▸ New configuration option allows user to disable CICS sockets tracing to reduce overhead
  ▸ CICS sockets Transaction Related User Exit (TRUE) and IBM supplied listener event monitoring logic have been optimized
  ▸ Virtual storage relief by allowing CICS sockets TRUE to be loaded above the 16MB line
  ▸ Support new CICS Open Transaction Environment (OTE)
  ▸ CICS OTE defines new L8 task mode that will reduce task switches that occur in Quasi-reentrant (QR) mode
    – New configuration option indicates which mode CICS sockets TRUE will run in
    – Users of OTE's L8 API, must examine application resource usage for serialization issues, since applications must be thread-safe

➤ **CICS Listener supports AT-TLS awareness**
  ▸ New configuration option supported to allow the listener to provide additional controls for client authentication based on client digital certificate exchange

➤ **Enhancement to the IBM supplied CICS Listener to allow it to be associated with a specific user ID.**
  ▸ Allows unique and predictable user IDs to be associated with each CICS Listener instance in a CICS region regardless of the method used to start the listener (ex: using PLPTI or via EZAO transaction)

Pool of reusable socket subtasks

LST1
LST2
TRNA
EZAO
EZAC

T R U E

TCP/IP Stack

CICS/ESA or CICS TS Region

PLTx

Build EZACICD

Conf. file

IP

Various FTP and TN3270 enhancements in z/OS V1R7

➤ **AES encryption for FTP and TN3270 TLS/SSL connections**
  - ▸ Advanced Encryption Standard - National Institute of Standards and Technology touts this as the replacement for DES
  - ▸ The cryptography scheme is a symmetric block cipher that encrypts and decrypts 128-bit blocks of data. Lengths of 128, 192, and 256 bits are standard key lengths used by AES. Only 128 is supported

➤ **FTP Enable/Disable extended directory search**
  - ▸ LISTSUBDir in FTP.DATA determines if the following commands apply to the current directory (FALSE) or spans one subdirectory (TRUE):
    - – For the client if mput *
    - – For the server if ls *, mget *, or mdelete *
  - ▸ New SITE/LOCSITE commands are also supported for LISTSUBDIR/NOLISTSUBDir
  - ▸ Provides flexibility in how HFS subdirectories are processed.

➤ **FTP server exploits the RACF-Delegated feature for cryptographic resources in the CFSERV and CFSKEYS classes for secure connections using TLS/SSL.**
  - ▸ Delegating profiles in these classes removes the need to permit each FTP client to these profiles.
  - ▸ If you are securing connections using TLS/SSL for your z/OS FTP Server, you must define a RACF delegated profile for the cryptographic resources.

➤ **Improve the confidence level of a file transfer**
  - ▸ Extra checks are optionally done to determine if a successful file transfer completion is in doubt
  - ▸ Status reported through the FTPOSTPR exit routine

➤ **FTP support for alternate end of line (EOL) sequence for ASCII transfers <CRLF>**
  - ▸ Configured via FTP.DATA configuration file or set via SITE or LOCSITE commands
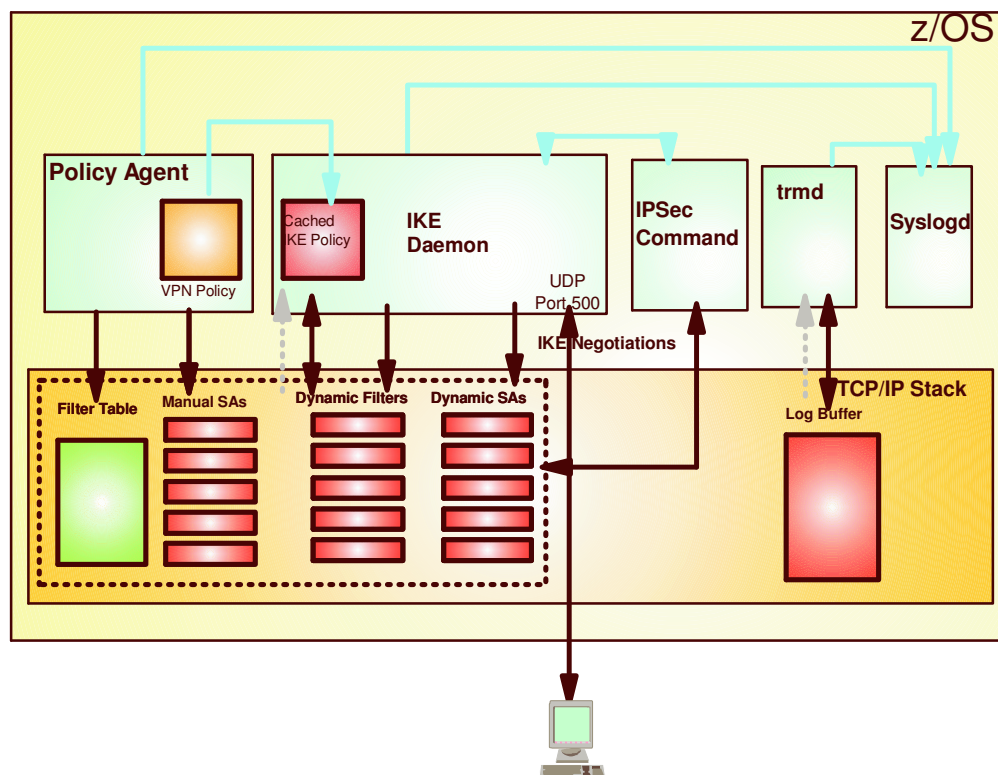
Security

> **Focus areas:**
>> ► Application-transparent IP security technologies
>>> – IPSec
>>> – Application-Transparent TLS (AT-TLS)

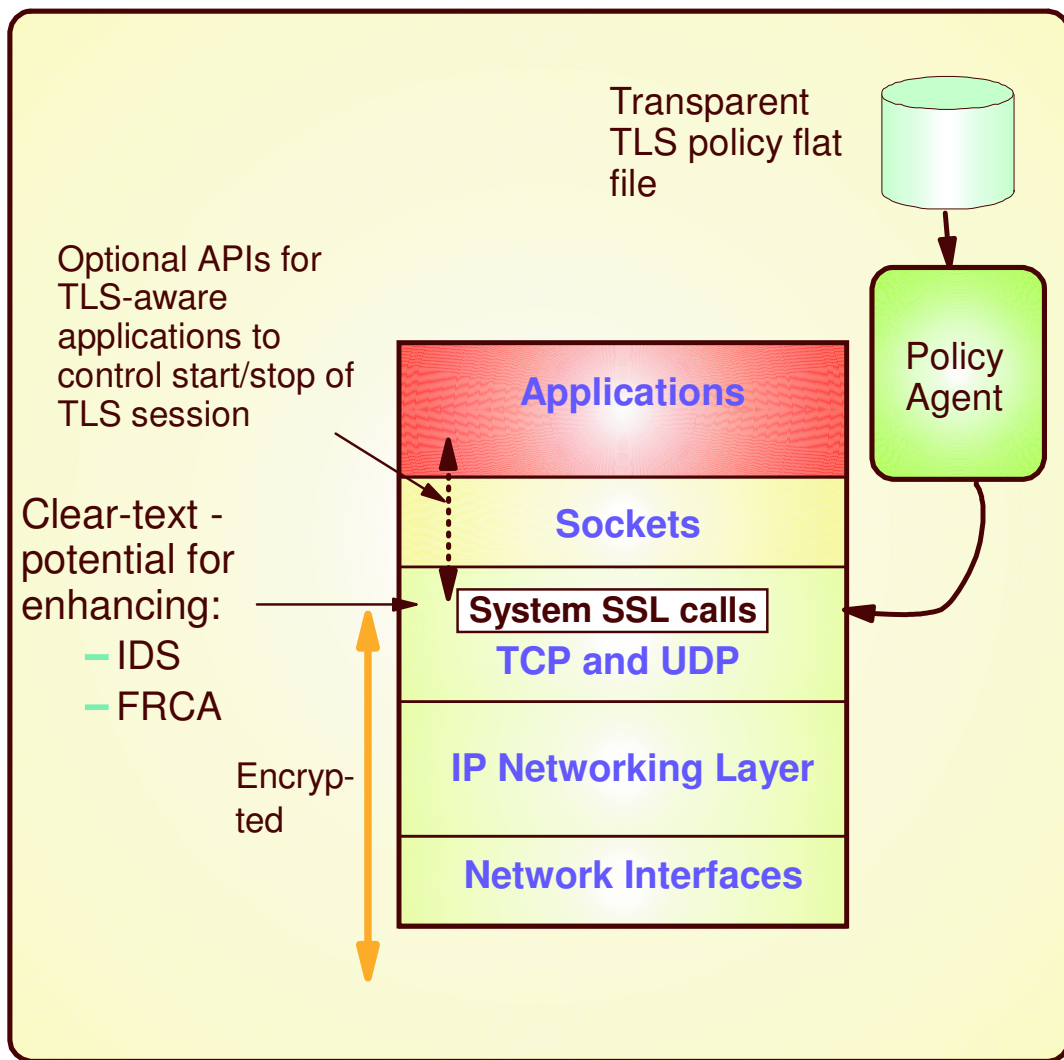# Integrated IPSec/VPN support including NAT traversal - ease of use and performance

➢ Features

- Configuration support
  - ▸ Optimized for z/OS host-to-host and z/OS host-to-gateway (z/OS gateway still supported)
  - ▸ NAT traversal support

- Simplified infrastructure
  - ▸ Eliminates need for FW technologies daemons

- Simplified configuration
  - ▸ New configuration GUI for both new and expert users
  - ▸ Direct file edit into local configuration file
  - ▸ Reduced definition, more "wildcarding"

- Improved serviceability
  - ▸ Improved messages and traces

- Default filters part of TCP profile
  - ▸ More granular control before policy is loaded

- Administrative controls
  - ▸ pasearch, new IPSec command

➢ Complete IPSec, filtering, and IKE solution part of z/OS Communications Server
  - Alternative to firewall technologies
    - ▸ New IKE daemon and configuration
➢ Makes use of existing Communications Server Infrastructure
  - TCP/IP stack - IPSec and IP filtering
  - Policy agent - reads and manages IPSec and IKE policy
  - trmd - monitors TCP/IP stacks for log messages

Transparent application security: policy-controlled transparent SSL/TLS support - SSL/TLS for all z/OS sockets applications

Transparent TLS policy flat file

Policy Agent

Optional APIs for TLS-aware applications to control start/stop of TLS session

Applications

Clear-text - potential for enhancing:
- IDS
- FRCA

Sockets

System SSL calls

TCP and UDP

Encryp-ted

IP Networking Layer

Network Interfaces

➤ **Basic TCP/IP stack-based TLS**
  ▸ TLS process performed at TCP layer without requiring any application change (transparent)
  ▸ All connections to specified port are designated as TLS required
    ● Can be further qualified by source/destination IP addresses
  ▸ Transparent TLS policies managed via Policy Agent

➤ **Transparent TLS can be requested by aplication**
 ▸ Application issues transparent TLS API calls to indicate that connection should start/stop using TLS

➤ **TCP/IP stack-based TLS with client identification services for application**
  ▸ Application issues TLS API calls to receive user identity information based on X.509 client certificate

➤ **Available to any TCP application**
  ▸ CICS Sockets and JES/NJE are primary focus of this support
  ▸ All programming languages supported

Policy-controlled application transparent network security
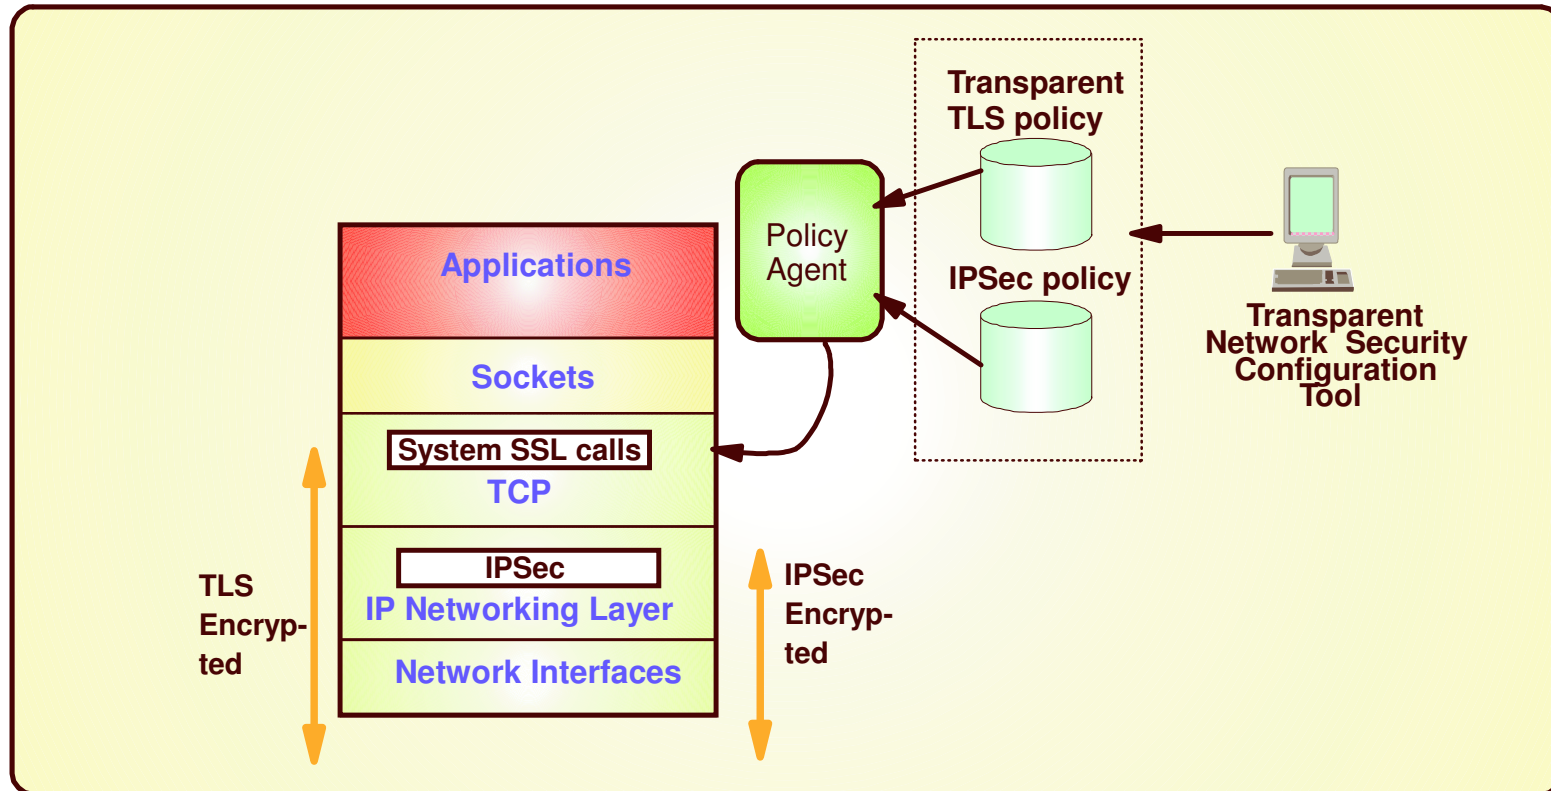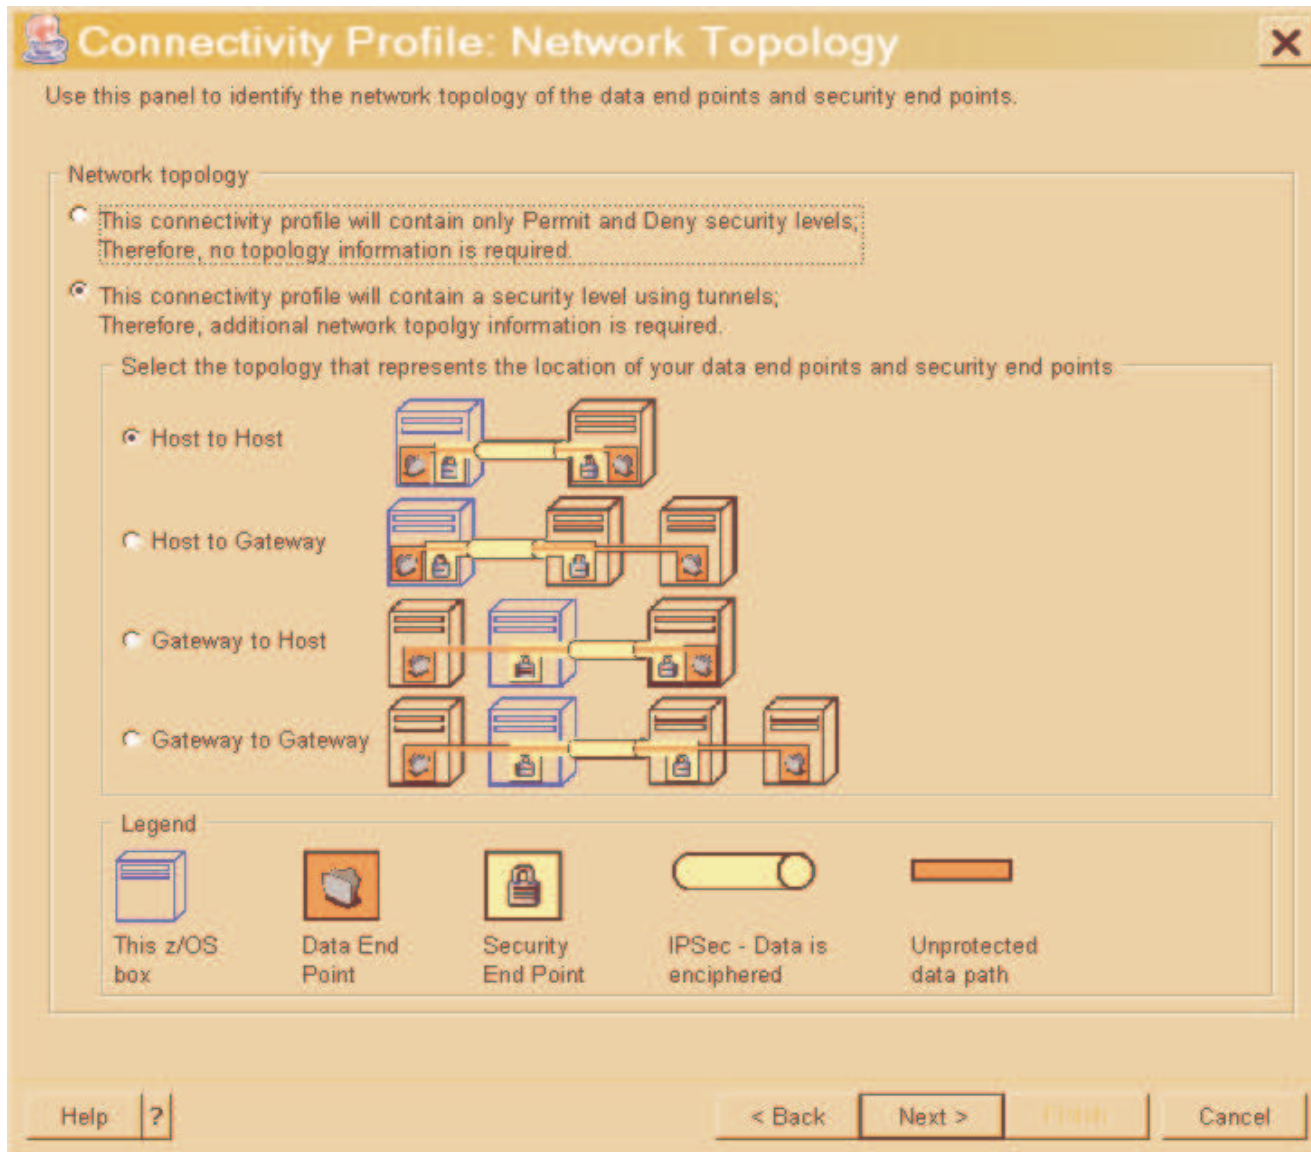

> **Network security without requiring application changes**
>   ▸ IPSec
>   ▸ Transparent TLS
> **Configuration single administrative task**
>   ▸ Higher level of abstraction
>   – Focus on what traffic to protect and how to protect
>   – Less focus on low level details (though available on expert panels)

# IPSec configuration dialog example

Enterprise Extender and SNA Enhancements

➤ **Focus areas:**
- ▸ Enhanced operations
- ▸ Enhanced configuration control

Display EEDIAG Command

➢ **V1R7 provides a new operator command to provide additional Enterprise Extender problem determination and management capability**

➢ **Two basic forms:**
  ► Display EE connections which have a packet retransmission rate that exceeds a specified threshold
  ► Display EE connections which have exceeded a specified number of LDLC "TEST" frame retries (SRQRETRY)

EE

EE Tool Box
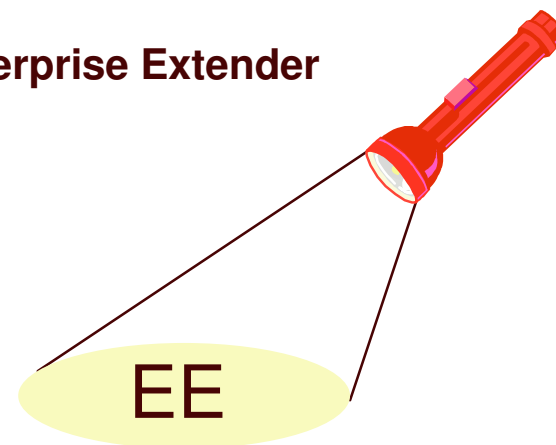
Example: Display EEDIAG

➢ **Find all Enterprise Extender connections associated with a particular VIPA with retransmission rates that meet or exceed 5% and display the output in detail format.**

```
D NET,EEDIAG,REXMIT=5,IP=9::67:1:1,LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2065I ENTERPRISE EXTENDER CONNECTION REXMIT INFORMATION
IST2067I EEDIAG DISPLAY ISSUED ON 08/27/04 AT 13:31:05
IST924I ------------------------------------------------------------
IST1680I LOCAL IP ADDRESS 9::67:1:1
IST1910I LOCAL HOSTNAME IP.SSCP1AV6
IST1680I REMOTE IP ADDRESS 9::67:1:6
IST1909I REMOTE HOSTNAME IP.SSCP2AV8
 .
 .
IST924I ------------------------------------------------------------
IST2032I PORT PRIORITY = HIGH
IST2036I   NLPS SENT            =                  134  ( 000K )
IST2038I   NLPS RETRANSMITTED   =                   67  ( 000K )
IST2068I   NLP RETRANSMIT RATE  =                   50%
IST924I ------------------------------------------------------------
 .
 .
IST924I ------------------------------------------------------------
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I   NLPS SENT            =                 1948  ( 001K )
IST2038I   NLPS RETRANSMITTED   =                   67  ( 000K )
IST2068I   NLP RETRANSMIT RATE  =                    3%
IST2069I REXMIT COUNTERS LAST CLEARED ON 08/27/04 AT 13:20:42
IST2042I 1 OF 1 EE CONNECTIONS DISPLAYED
IST314I END
```

Model Definition of VTAM CDRSCs

➤ **VTAM will now allow model CDRSCs (cross-domain resources) to be defined in a CDRSC major node.**
  - ▸ Model CDRSC definitions can be created similar to the way model APPL definitions are created, based on the knowledge of the naming conventions of the network to which connections are desired.

➤ **Using the appropriate model CDRSC definitions, VTAM will dynamically create clone CDRSCs as needed.**
  - ▸ Unaffected by CDRDYN start option and CDRSC keyword.

➤ **Benefits:**
  - ▸ Amount of definition required when connecting to a new network is reduced without totally giving up control.
  - ▸ More parameters can be specified for model CDRSCs than dynamic CDRSCs.
    - – e.g., DLOGMOD can be different for various model CDRSCs.
  - ▸ Provides control of when clone CDRSCs are deleted.

➤ **'?' and '*' may be specified within the name field of the CDRSC definition statement:**

```
CDRSCMOD  VBUILD  TYPE=CDRSC
?APPL     CDRSC
APPL*     CDRSC
APPL?     CDRSC
          NETWORK NETID=NETA
APPL1*    CDRSC DELETE=NO
APPL1?    CDRSC
ABCD*     CDRSC DELETE=YES
EF?G*     CDRSC DELETE=NO
```
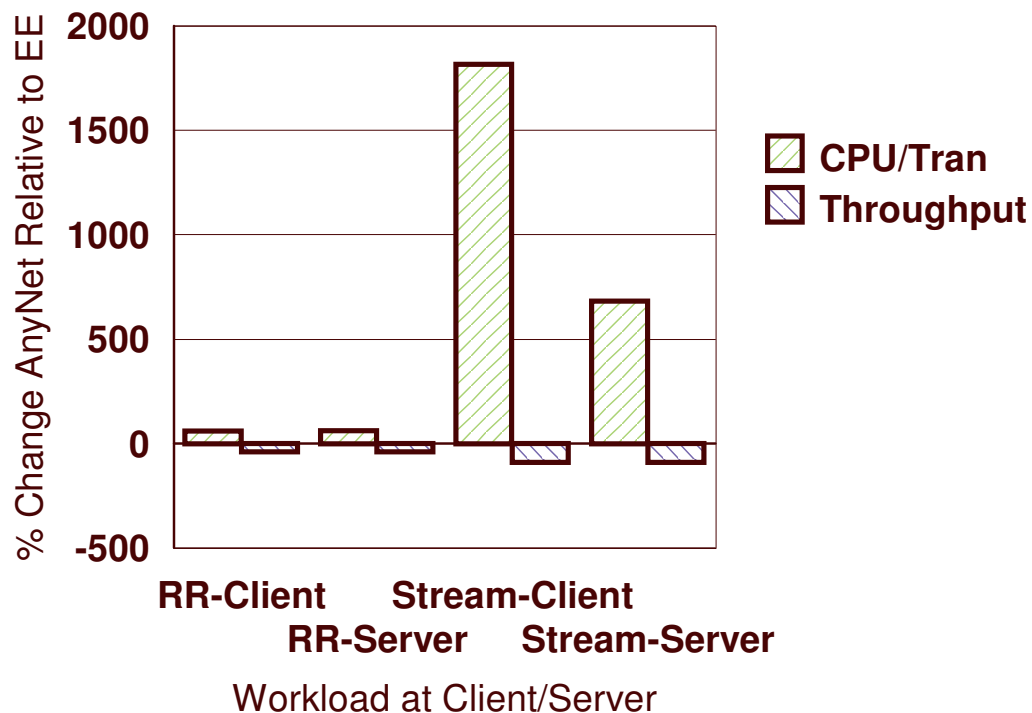
Other EE/SNA Enhancements

➢ **Display the APPN route used when a session fails to establish**
  ▸ Controlled by a new RSIRMSG start option

➢ **Provide a command to delete a generic resource (GR)**
  ▸ May be useful if transitioning from a Generic Resource back to a uservar
  ▸ Cleanup GR's in the Coupling Facility

```
IST663I CINIT REQUEST FAILED, SENSE=08010000
IST664I REAL  OLU=NETB.APPLBA1        REAL
DLU=NETA.APPLAA1
IST889I SID = F6AF18D8B39E92EF
IST2103I RSCV TOWARDS SLU
IST1460I TGN  CPNAME              TG TYPE      HPR
IST1461I  21  NETA.SSCP1A         APPN         RTP
IST1461I  21  NETB.SSCP2A         ISL          RTP
IST1461I  21  NETB.SSCPBA         APPN         RTP
IST314I END
```

➢ **Enhancements to AUTOLOGON processing**
  ▸ Provide command to restart normally-terminated autologon sessions

➢ **Allow subarea VTAMs to support TCP/IP XCF interfaces**
  ▸ VTAM provides the device layer support for TCP/IP's XCF interfaces (i.e. Coupling Facility links) and has restricted the definitions of these types of links to only APPN-enabled VTAMs.
  ▸ With this release, z/OS Communications Server will allow:
    – TCP/IP connectivity through XCF on APPN nodes without having to first establish APPN connections
    – TCP/IP connectivity through XCF from pure subarea nodes

➢ **DISPLAY SESSIONS,SID= displays the RU sizes (PLU and SLU direction) used for the session**
  ▸ Improves serviceability

➢ **DISPLAY ID=rtpname,HPRDIAG=YES displays RTP work queue depth**
  ▸ Current depth and high-water mark

Removal of AnyNet

➢ **Enterprise Extender, TN3270, and distributed Communications Server Remote API functions are the strategic protocols for SNA/IP integration**

  ▸ AnyNet has not been enhanced in years

➢ **EE is functionally superior, but also significantly outperforms AnyNet by all measures:**

  ▸ AnyNet exhibits lower throughput and higher CPU utilization relative to EE:
    – Interactive workloads
      • Throughput down 39%
      • CPU utilization up 63%
    – Stream workloads
      • Throughput down 89%
      • CPU utilization up 682-1817%

V1R4 AnyNet vs. V1R4 EE (EE as Base

Legend:
- □ CPU/Tran
- ▨ Throughput

y-axis: % Change AnyNet Relative to EE (2000, 1500, 1000, 500, 0, -500)

x-axis: RR-Client, RR-Server, Stream-Client, Stream-Server

Workload at Client/Server

➢ **z/OS V1R7 is the last release of z/OS to include AnyNet as a component of Communications Server**

IPv6 on z/OS Communications Server

> **Focus areas:**
>   ► IPv6 network management

Visible IPv4 hosts on the Internet through the last 12 years

Hosts on the Internet
Advertized by DNS servers
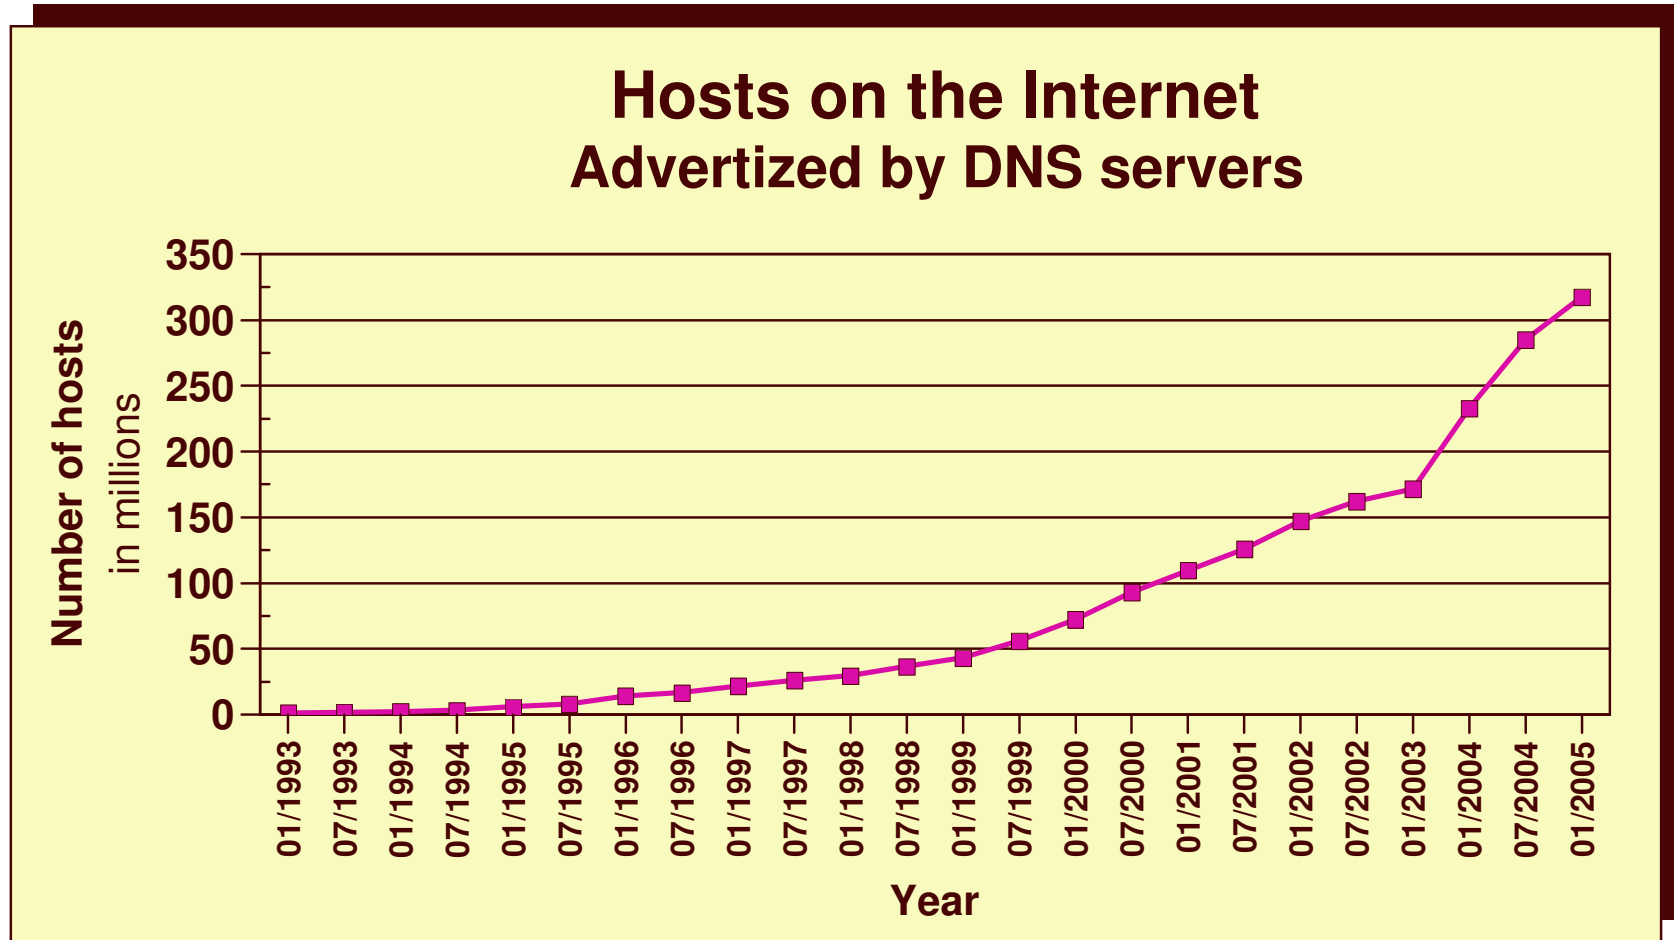
What is the upper practical limit (the ultimate pain threshold) for number of assigned IPv4 addresses? Some predictions say 250,000,000 (250 million), others go up to 1,000,000,000 (one billion or one milliard).

Both z/OS V1R5 and V1R6 have been certified with the IPv6 Ready logo

The Journey to IPv6 for z/OS Communications Server

The first phase (z/OS V1R4)
- Stack support for IPv6 base functions - (APIs, Protocol layers)
- Resolver
- High speed attach (OSA Express QDIO))
- Service tools (Trace, Dump, etc.)
- Configuration and netstat, ping, traceroute, SMF
- Static Routing
- FTP, otelnetd,unix rexec, unix rshd/rexecd

The second phase (z/OS V1R5)
- Network Management
  - Applications and DPI
  - Version-neutral Tcp/Ip Standard MIBs
  - Additional SMF records
- Applications/Clients/APIs
  - Tn3270 server,CICS sockets, sendmail,ntp,dcas, rxserve,rsh client
- Enterprise Extender
- Point to Point - type DLCS
- Dynamic Routing Protocol w/ OMPROUTE (only RIPng)

The third phase (z/OS V1R6)
- Sysplex Exploitation (Dynamic VIPA, Sysplex Distributor functions)
- Dynamic Routing Protocol w/ OMPROUTE (OSPFv3)
- Additonal Network Management MIBs

The fourth phase (z/OS V1R7)
- SNMP UDP standard MIB (RFC2013) and IBM MVS TCP/IP Enterprise-specific MIB for UDP
- Advanced Socket API support - RFC3542
- IPv6 Two Default Routers - required for IPv6 compliance

After z/OS V1R7
- HiperSockets DLC
- Integrated IPSec
- Complete Advanced Socket APIs
- Extended Stats MIB, OSPFv3 MIB
- Intrusion Detection Services
- IPv6 mobility support

The Internet - a worldwide digital utility.

Objective is to have IPv6 production ready on the platform when you need it!

Connectivity for *anyone* from *anywhere* (car, home, office) to *anything*!

Statements of Direction

Statements of Direction

➢ z/OS V1R7 is planned to be the last release in which z/OS Communications Server will support AnyNet. After z/OS V1R7, the function will be removed from the product. Customers may implement other IBM solutions such as Enterprise Extender (EE) as a replacement for AnyNet.

➢ z/OS V1R7 is planned to be the last release in which Communications Server will support the definition of parallel Enterprise Extender TGs by specifying multiple SAP addresses. After z/OS V1R7 this capability will be removed from the product. Beginning in V1R5, parallel EE TGs may be defined by using different EE VIPAs on one (or both) of the endpoints.

➢ z/OS V1R7 is planned to be the last release in which z/OS Communications Server will support the SNMP Service Level Agreement (SLA) Version 1 MIB and the PAGTSNMP subagent. After z/OS V1R7, these functions will be removed from the product. In z/OS V1R5, Communications Server provided a new SNMP SLA Version 2 MIB and NSLAPM2 subagent. It is recommended that users migrate to the new version 2 MIB.

➢ z/OS V1R7 is planned to be the last release in which z/OS Communications Server will support TCP/IP configuration profile block definition statements: ASSORTEDPARMS, ENDASSORTEDPARMS, KEEPALIVEOPTIONS and ENDKEEPALIVEOPTIONS. After z/OS V1R7 this capability will be removed from the product. Equivalent capability is provided by the GLOBALCONFIG, IPCONFIG, TCPCONFIG, and UDPCONFIG statements instead of the ASSORTEDPARMS statements. Equivalent capability is provided by INTERVAL and SENDGARBAGE on the TCPCONFIG statement instead of the KEEPALIVEOPTIONS statements. Beginning in z/OS V1R2, message, EZZ07171, has been issued for these statements as notification that this support would be removed in a future release. Refer to the z/OS Communications Server IP Configuration Reference for information on TCP profile and configuration statements.

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)