z/OS Network update

Comm Server 1.9

Large Systems Update 2007

olle.zetterlund@se.ibm.com

AGENDA

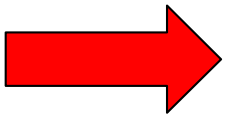

- OSA express enhancements
  - ***Virtual MAC Address***
  - ***Dynamic LAN timer***
  - ***Network traffic analyzer***
- Policy Enhancements
- Security enhancements
- Sysplex Enhancements
- Application enhancements
- Management
- EE/SNA enhancements

Background Information
Sharing of OSA-Express Features

- Allows many stacks, in different LPARs, to share bandwidth
- Even more important with high bandwidth adapters (10 gig, etc)
- Accomplished by registering IP addresses, sharing "burned in" MAC
- One stack may be PRIROUTER for unknown packets

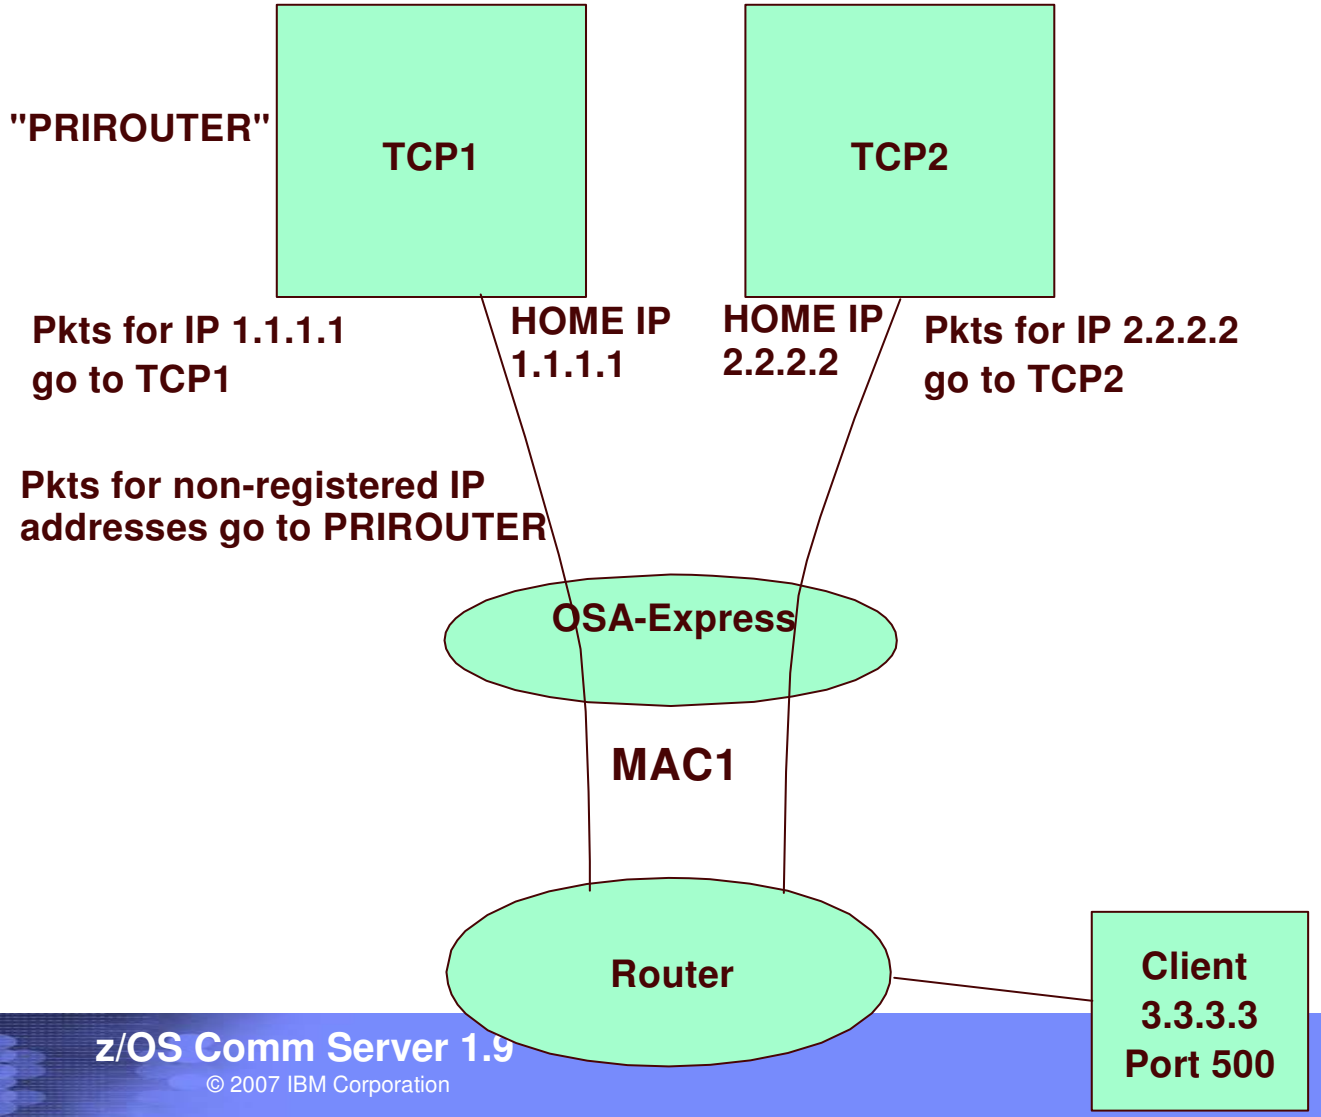

Sharing Problems

- Sharing can fail in Load Balancing solutions
  - Cisco MNLB
  - z/OS Load Balancing Advisor
  - Workaround is to use GRE or NAT
    - negative effect on performance
    - IPV6 not supported
- Only one routing stack possible

Solution: OSA-Express virtual MAC address

- Problems are solved if each stack has its own MAC
  - "virtual" MAC
  - To the network, each stack appears to have a dedicated OSA
- All IP addresses for a stack advertised with virtual MAC
  - by OSA using ARP for IPv4
  - by the stack using Neighbor Discovery (ND) for IPv6
- All external routers now forward packets to virtual MAC
  - OSA will route by virtual MAC instead of IP address
  - All stacks can be "routing" stacks instead of 1 PRIROUTER stack
- Simplifies configuration greatly
  - No PRIROUTER/SECROUTER!

OSA-Express virtual MAC rules

- Each stack may define one VMAC per protocol (IPv4 or IPv6) for each OSA
  - One VMAC for the LINK statement
  - One VMAC for the INTERFACE statement
- VMAC routing is mutually exclusive with PRIROUTER/SECROUTER routing
  - If a VMAC is defined
    - This stack will not receive any packets destined to the physical MAC
  - If VMAC is not defined
    - This stack will not receive any packets destined for a VMAC
    - Even if this stack is PRIROUTER!
- VLAN ids apply to VMACs like physical MACs

Dynamic LAN Timer Background Information

- OSA supports an inbound “blocking” function over the QDIO interface.
  - Affects how long OSA will hold packets before “presenting” those packets to the host.
  - Indirectly affects how frequent the host will be interrupted, and the payload per interrupt.
- For an OSA Express in QDIO mode device the TCP/IP profile INBPERF parameter can be specified with one of the following options:
  - **MINCPU** - a static interrupt-timing value, selected to minimize host interrupts without regard to throughput
  - **MINLATENCY** - a static interrupt-timing value, selected to minimize latency
  - **BALANCED** (default) - a static interrupt-timing value, selected to achieve reasonably high throughput and reasonably low CPU

LAN idle timer Problem/solution

- **Problem: LAN idle timer settings contributes to nw latency on zSeries**
  - **Even when the INBPERF parameter is specified with a value of MINLATENCY the permitted *inter-packet gap is set to 20 microseconds***
  - **LAN idle timer settings are *static* and can not be changed unless the connection to OSA connection is terminated and reestablished.**
- **Solution**
  - **Dynamically tune the LAN Idle timer values to reflect current workload characteristics**
  - **Allow for the minimum latency when a light interactive workload is determined**
  - **The inter-packet gap time can now be reduced as small as a microsecond**

Dynamic LAN Timer configuration

- New DYNAMIC option for the existing INBPERF parameter.
  - INBPERF parameter can be specified on the OSA-Express QDIO LINK or INTERFACE statement.
  - New option is valid for OSA-Express2 on an IBM System z9 EC or z9 BC with the corresponding Dynamic LAN Idle functional support
  - When specified for an OSA-Express device that does not support this new function then the option of BALANCED will be used for INBPERF parameter.

QDIO problem Diagnosis
traces in different places

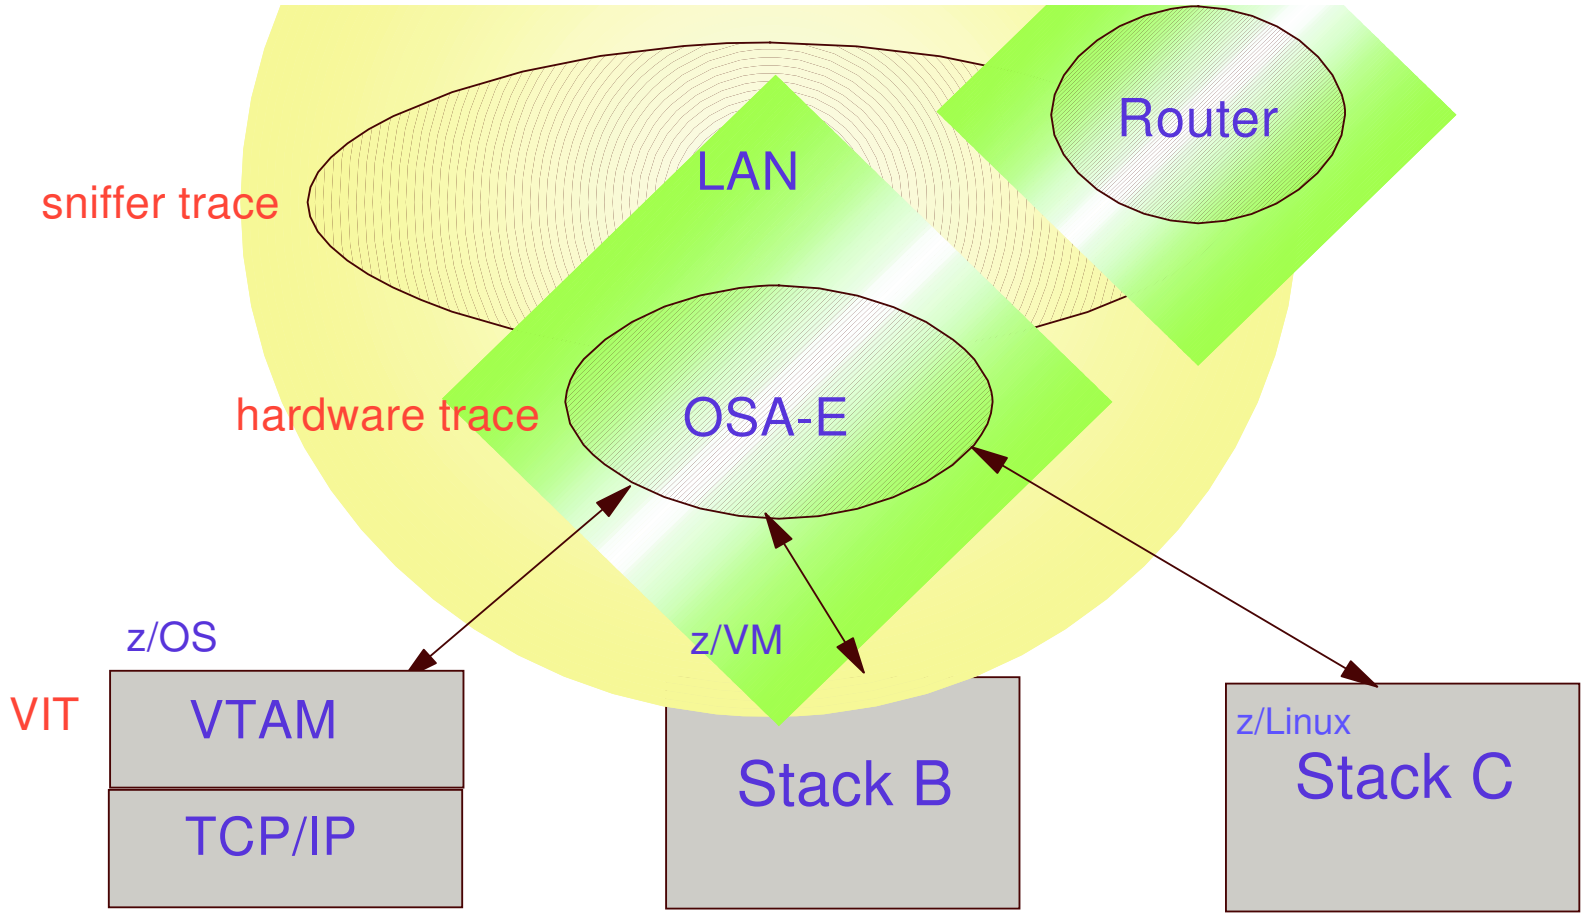

Solution: Network Traffic Analyzer

- Supported on OSA-Express2 GA3 (in QDIO mode) on z9-109.
  - Refer to the 2094DEVICE Preventive Service Planning (PSP) and the 2096DEVICE Preventive Service Planning (PSP) buckets for the level of the OSA-Express2 LIC.
- Allows z/OS Comm Server to collect Ethernet data frames from OSA
  - Not a sniffer trace (but similar in some aspects)
  - No promiscuous mode
- Minimizes the need to collect and coordinate multiple traces for diagnosis
- Minimizes the need for traces from the OSA Hardware Management Console (HMC)

Network Traffic Analyzer

- Controlled by z/OS Comm Server
  - New OSAENTA command
    - Define trace filters and parameters
    - OSA sends trace records to the z/OS stack
  - Save and format the data using existing Ctrace facilities

- Collected by OSA
  - Ability to see:
    - ARP packets
    - MAC headers (including VLAN tags)
    - Packets to/from other stacks shared by the OSA (which could be z/VM or z/Linux)
    - SNA packets

Network Traffic Analyzer

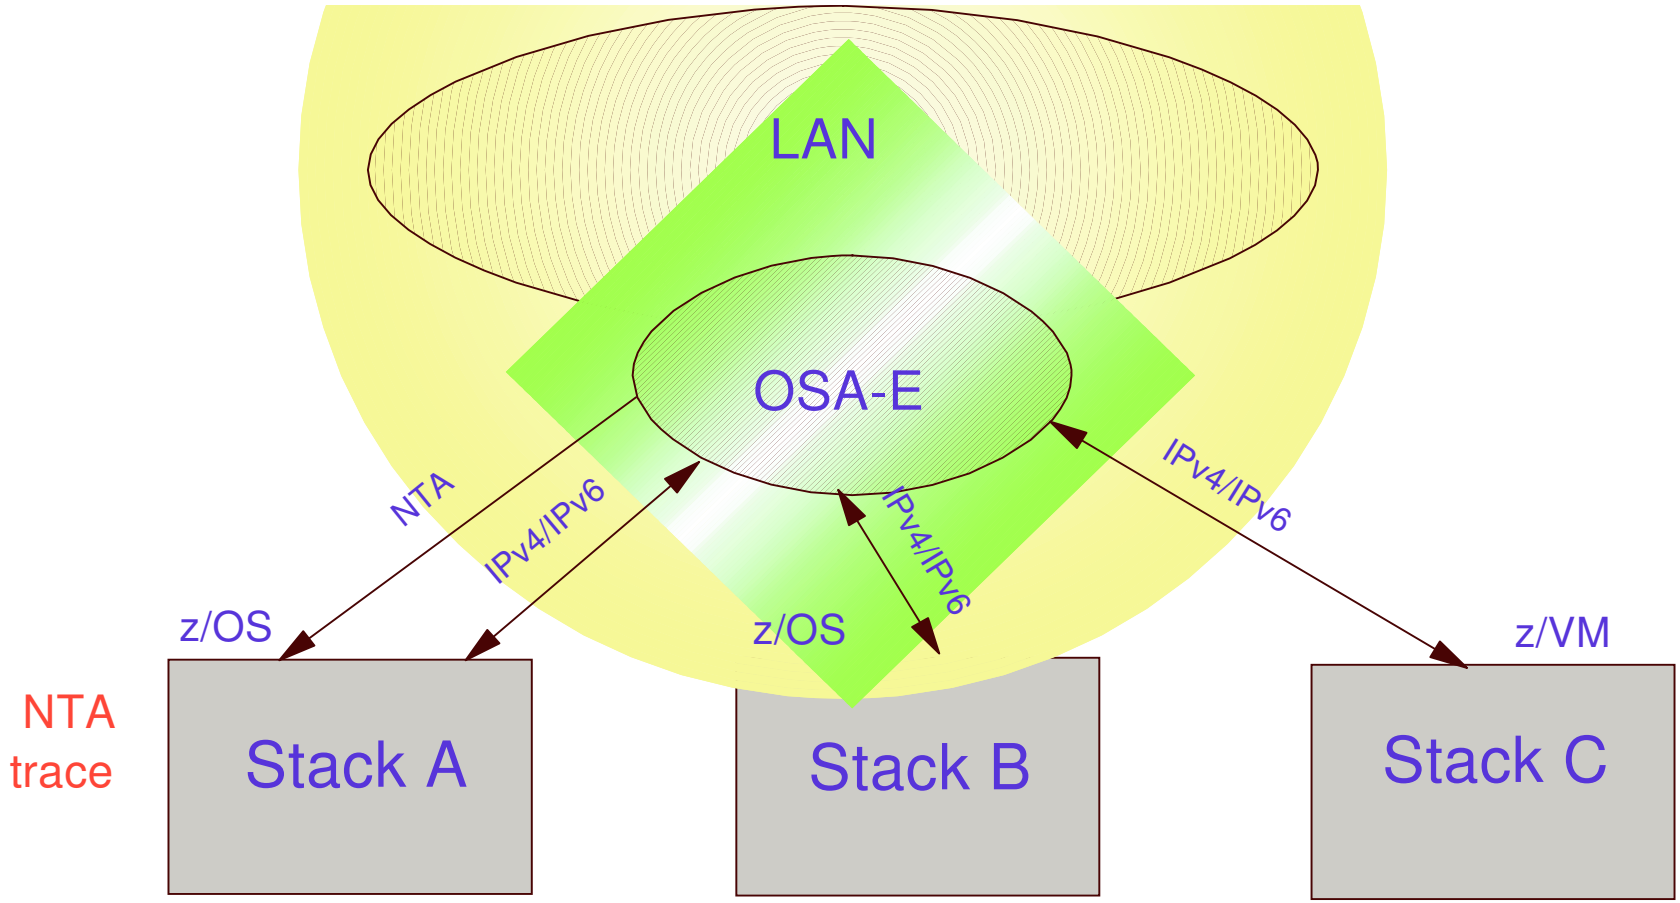

AGENDA

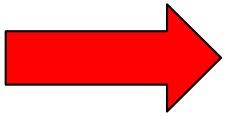

- OSA express enhancements
- Policy Enhancements
  - **Policy Based Routing**
  - **Central Policy Server**
  - **New Config Assistant**
- Security enhancements
- Sysplex Enhancements
- Application enhancements
- Management
- EE/SNA enhancements

Policy-based routing

- What is Policy Based Routing (PBR)?
  - Choose first hop router and outbound network interface
  - Choice can be based on more than the usual destination IP address/subnet

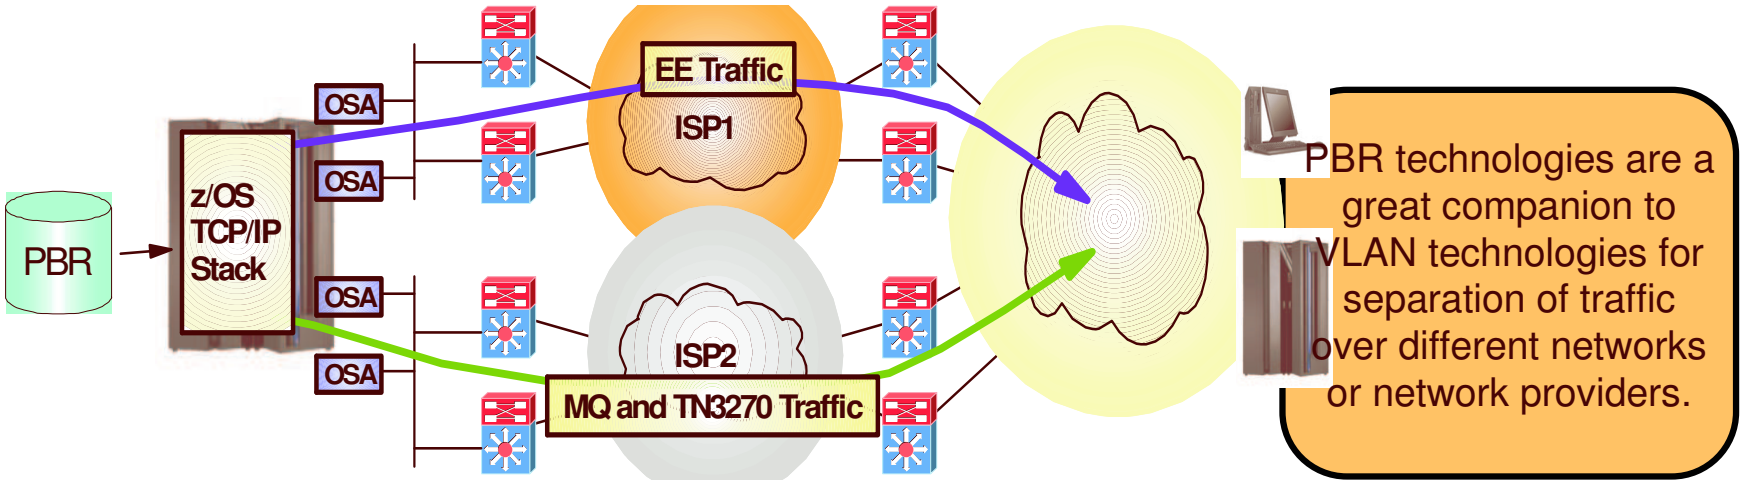

Overview
Policy-based routing

- **Problem Statement / Need Addressed:**
  - Routing decision, within the TCP/IP stack, does not take into consideration the type of application, source IP address and other criteria
  - A single metric is used to decide which route to use when sending packets
    - The shortest path (or a static route) for a destination IP address

- **Solution:**
  - Policy-based routing (PBR) which enables routing decision that takes into account other criteria in addition to destination IP address.
    - Source and destination ports
    - Protocol (TCP or UDP)
    - Source and destination IP addresses
    - Job name
    - Security zones and security labels

- **Benefit:**
  - Policy can be used to select networks with different capabilities for different applications
  - Policy can be used to ensure that secure traffic is routed to a secured network via an appropriate outbound interface.

Usage & Invocation
Policy-based routing

- Policy-based routing (PBR) is configured in a policy agent flat-file
  - No LDAP file support for PBR
  - Centralized policy support for PBR
- IBM Configuration Assistant for z/OS Communications Server (Configuration Assistant)
  - Can be used to generate PBR policy statements

Overview
Centralized Policy Services

- **Problem Statement / Need Addressed:**
  - The scope of Policy Agent policies continues to widen, with new policy types added over the last several releases
  - Local management of policies is therefore becoming a larger administrative burden
- **Solution:**
  - Use the Policy Agent as a centralized policy repository
- **Benefit:**
  - Centralized administration and management of policy definitions
  - The IBM Configuration Assistant only needs connectivity to the policy server, if no local policies are defined on the policy clients.

Centralized Policy Services

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

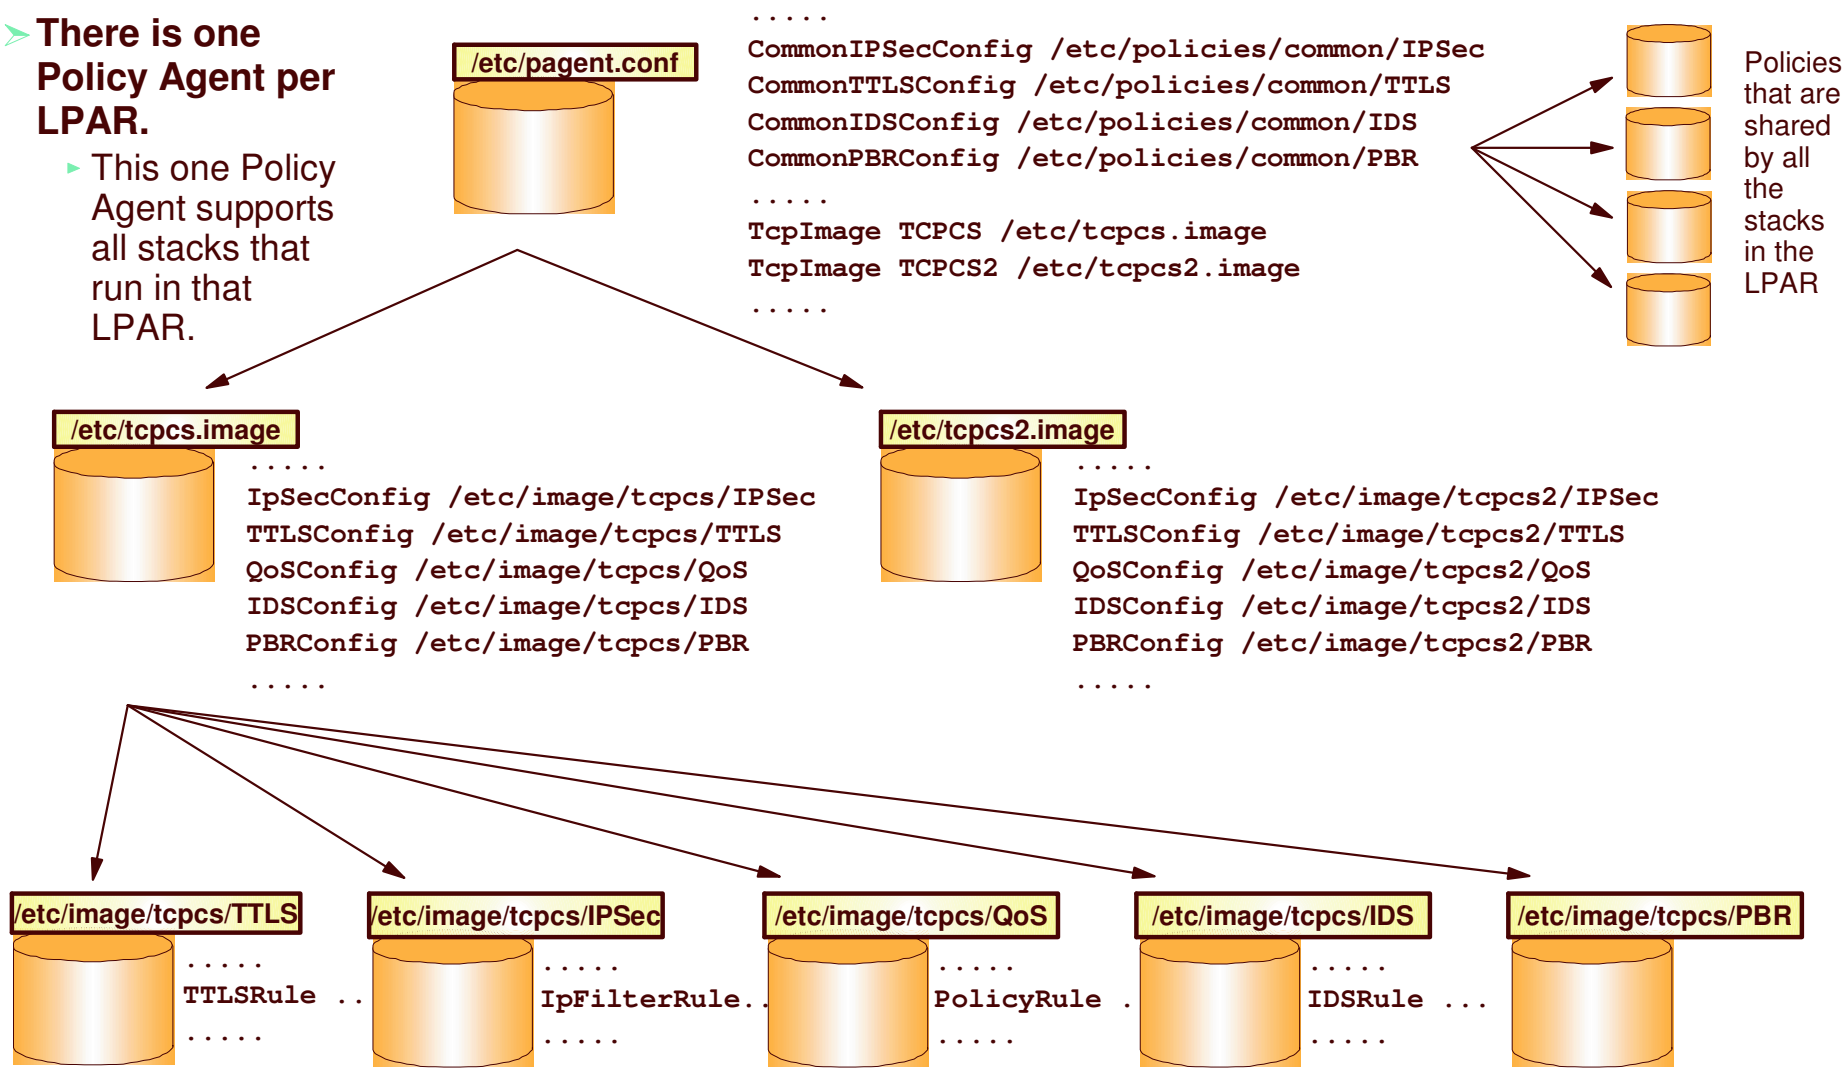

Centralized Policy Services

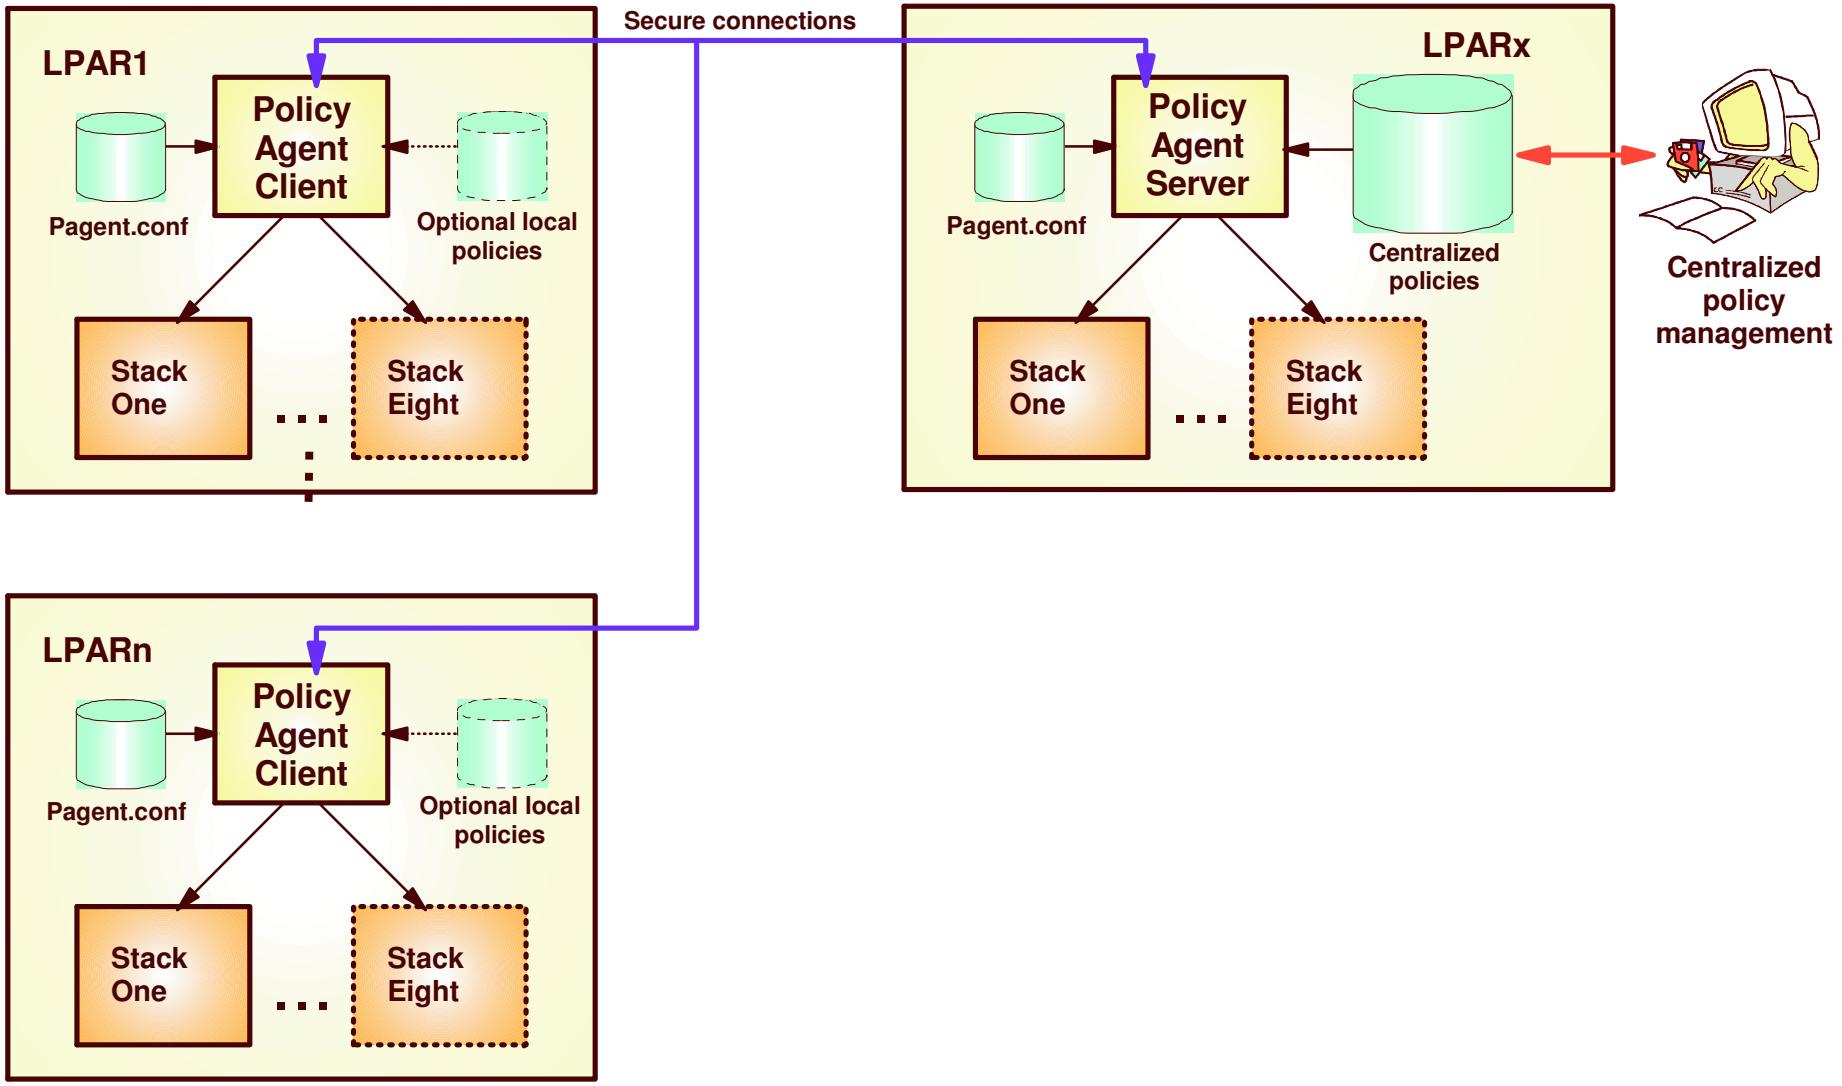

IBM Configuration Assistant for z/OS Communications Server

- The new look is centered around the images and stacks to be configured

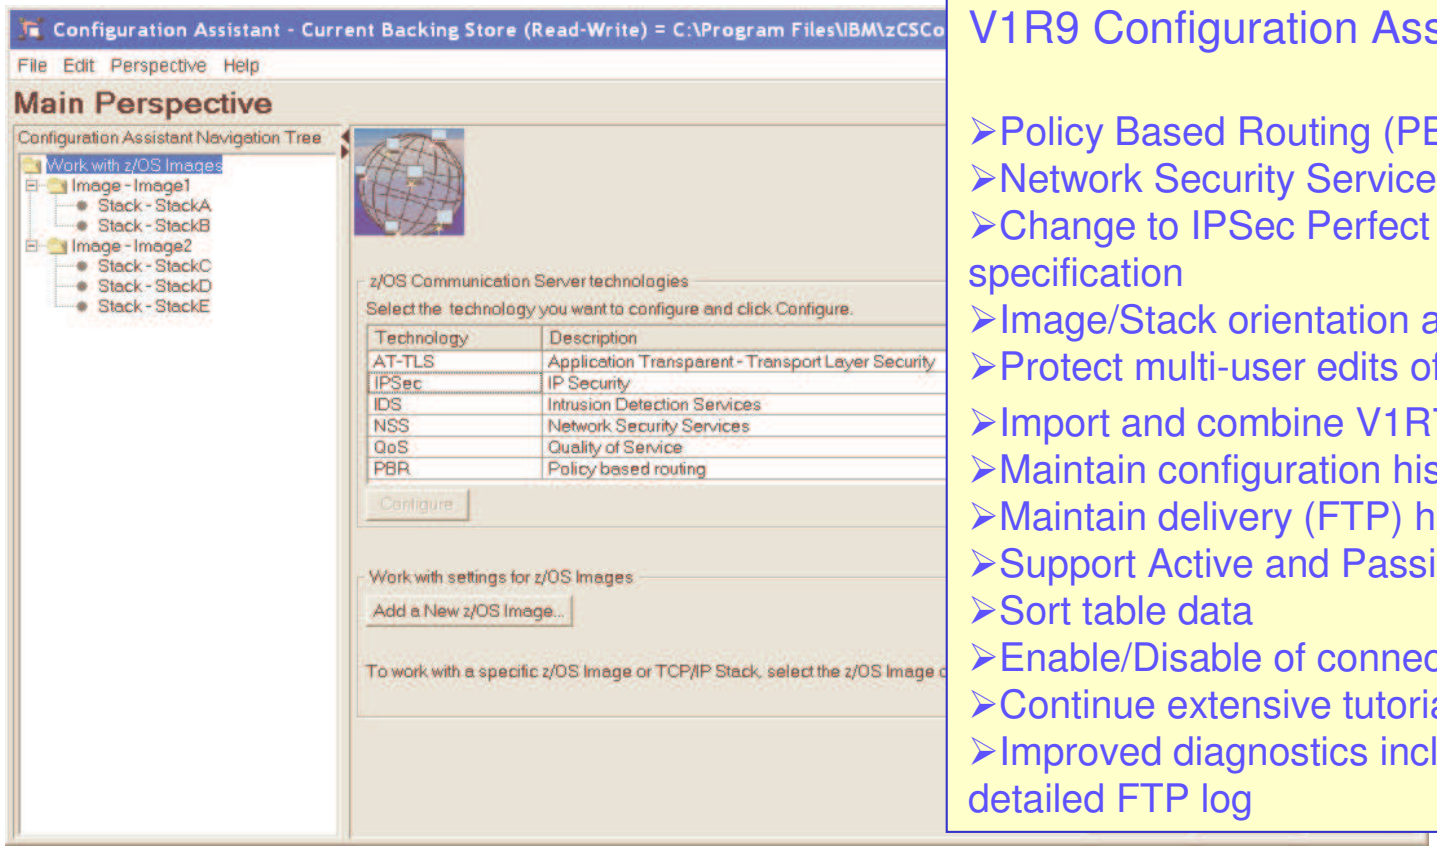

V1R9 Configuration Assistant Enhancements

- Policy Based Routing (PBR)
- Network Security Services (NSS)
- Change to IPSec Perfect Forward Secrecy specification
- Image/Stack orientation across multiple technologies
- Protect multi-user edits of the same backing store
- Import and combine V1R7/R8/R9 backing store data
- Maintain configuration history for audit / tracking
- Maintain delivery (FTP) history for audit / tracking
- Support Active and Passive mode FTP
- Sort table data
- Enable/Disable of connectivity rules
- Continue extensive tutorials
- Improved diagnostics including log levels and a detailed FTP log

AGENDA

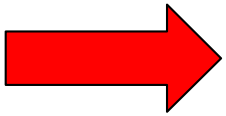

- OSA express enhancements
- Policy Enhancements
- Security enhancements
  - ***AT-TLS enhancements***
  - ***enable TN3270/FTP to AT-TLS***
  - ***IPSec NMI***
  - ***IPSec NSS***
- Sysplex Enhancements
- Application enhancements
- Management
- EE/SNA enhancements

AT-TLS Application Enhancements

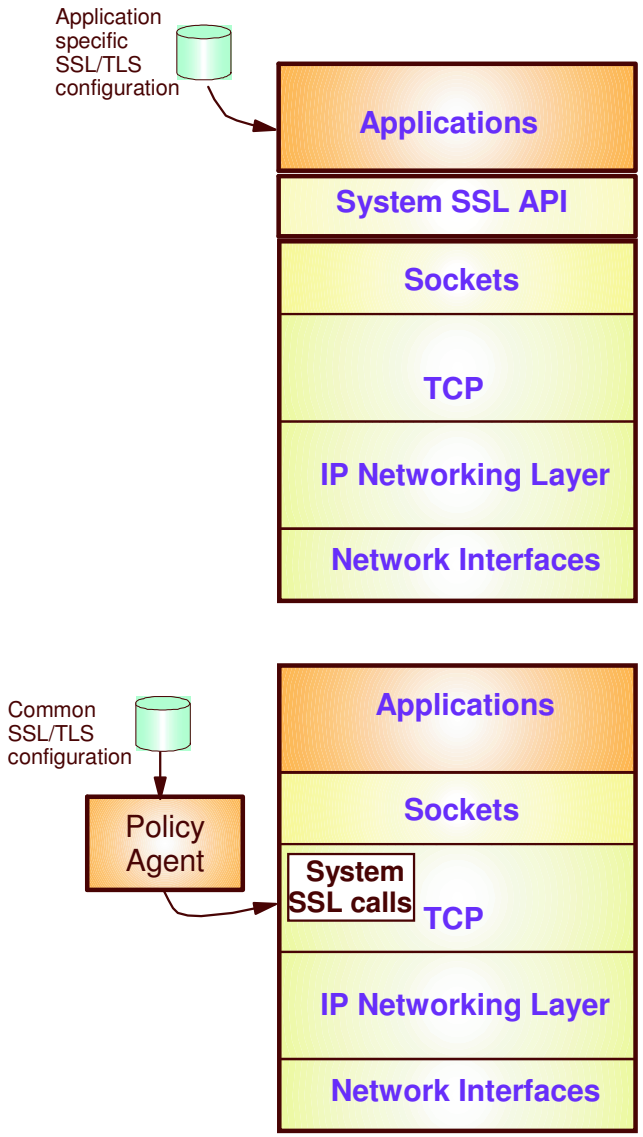

Enable AT-TLS for FTP and TN3270

- **Problem Statement / Need Addressed:**
  - FTP's and TN3270´s SSL implementation do not exploit all the functions of System SSL
- **Solution:**
  - The FTP client and server and TN3270 server can now be configured to use AT-TLS to support SSL/TLS connections.
    - ✓FTP: **TLSMECHANISM ATTLS**
    - ✓TN3270: **TTLSPORT nnnnn**
- **Benefit:**
  - More System SSL functions available to the FTP and TN3270

AT-TLS enabling the TN3270E Telnet server and the FTP client and server

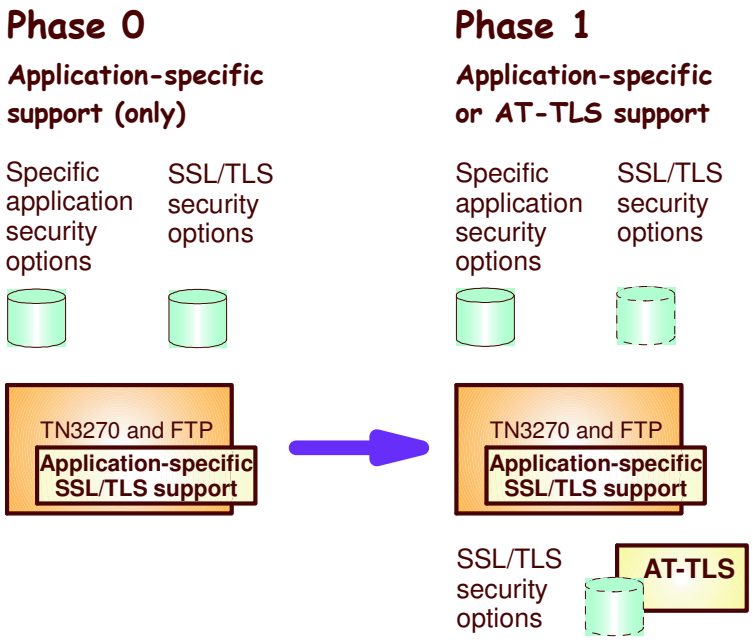

IPSec Network Management Interface support

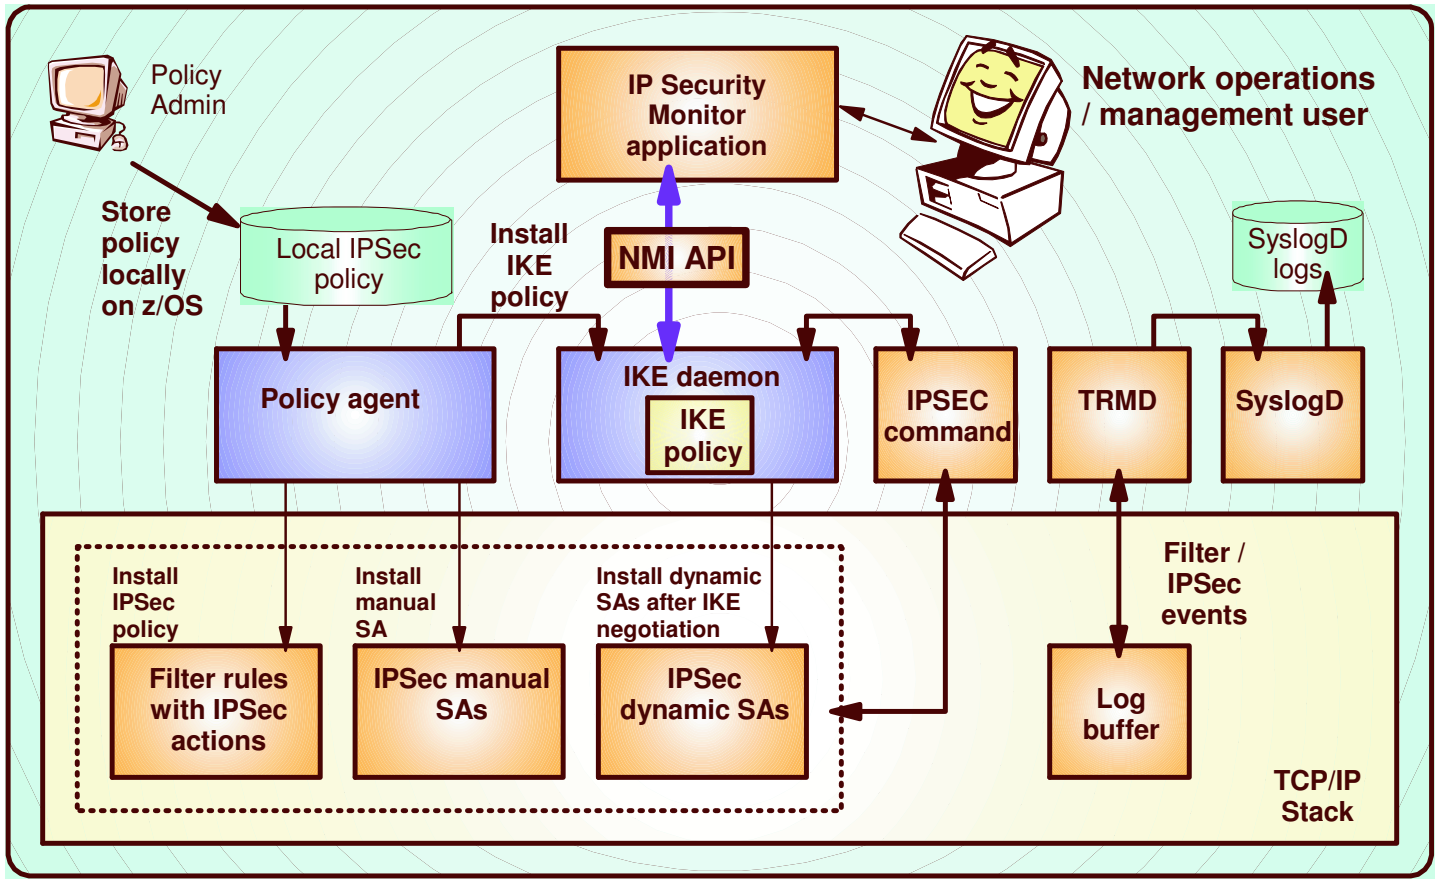

Overview
IPSec Network Security Services

- **Problem Statement / Need Addressed:**
  - Storing sensitive data like private keys in less trusted zones can create security vulnerabilities
  - Administration of certificates across many security endpoints can be cumbersome and error-prone
- **Solution:**
  - Provide centralized certificate services, monitoring and management for IPSec
- **Benefit:**
  - Centralize and reduce configuration and deployment complexity
  - Enables monitoring and management of remote IPSec endpoints

IPSec Network Security Services (NSS)

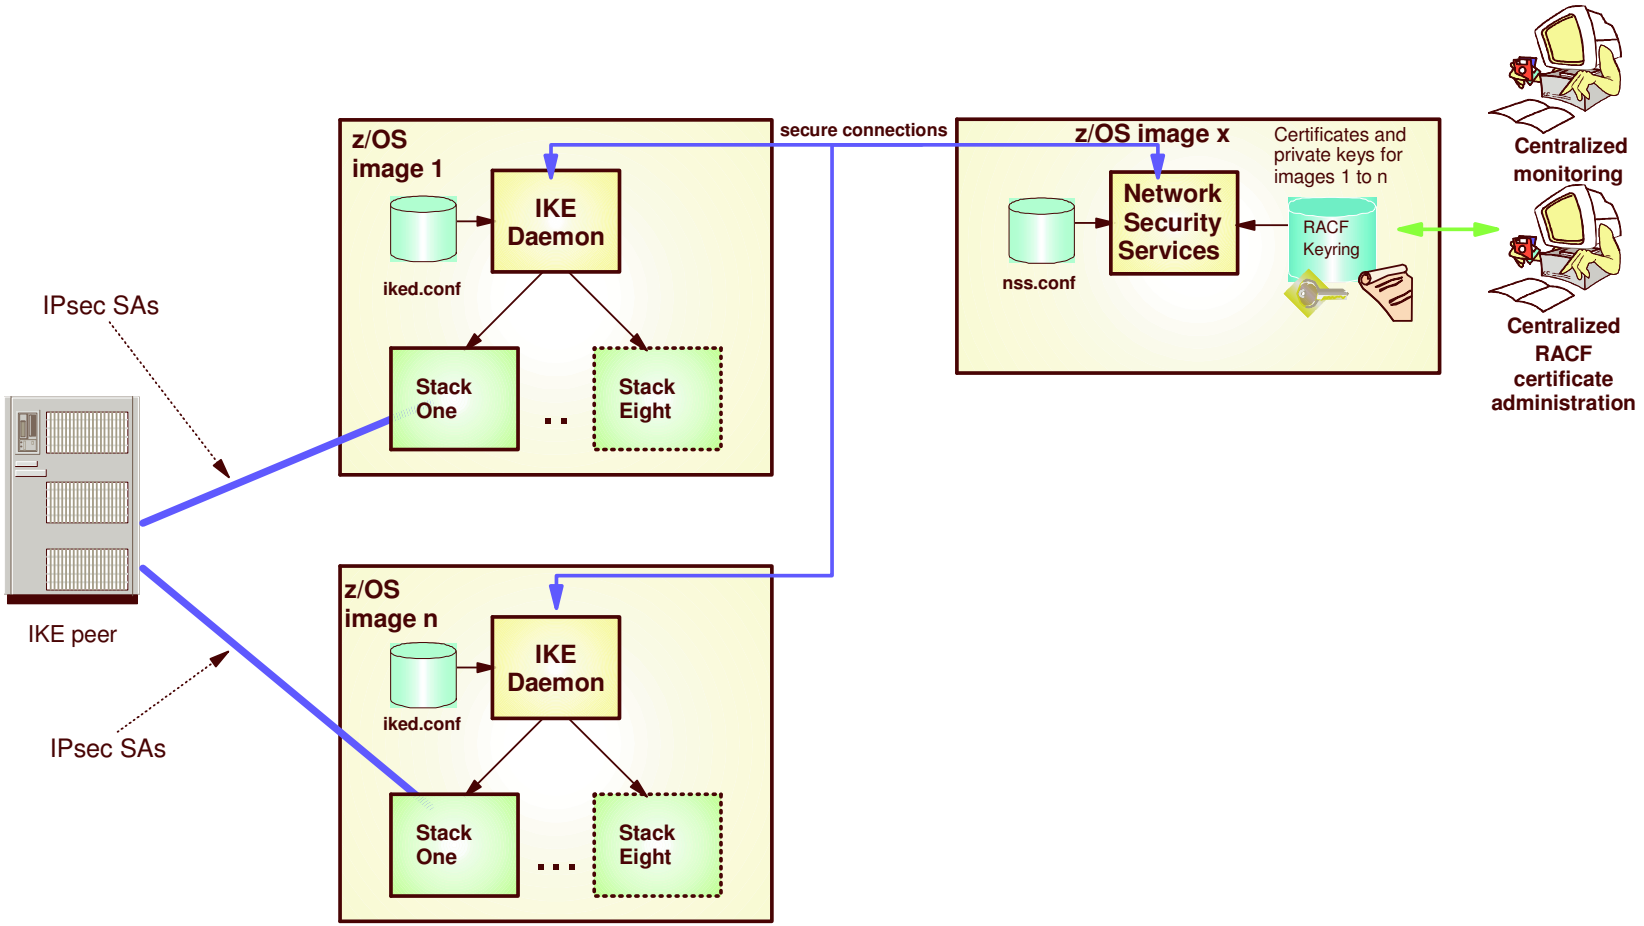

AGENDA

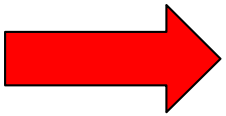

- OSA express enhancements
- Policy Enhancements
- Security enhancements
- Sysplex Enhancements
  - ***Source IP enhancements***
  - ***DVIPA usability enhancements***
  - ***V TCPIP,,SYSPLEX***
- Application enhancements
- Management
- EE/SNA enhancements

Source IP (SRCIP) Enhancements

```
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
```

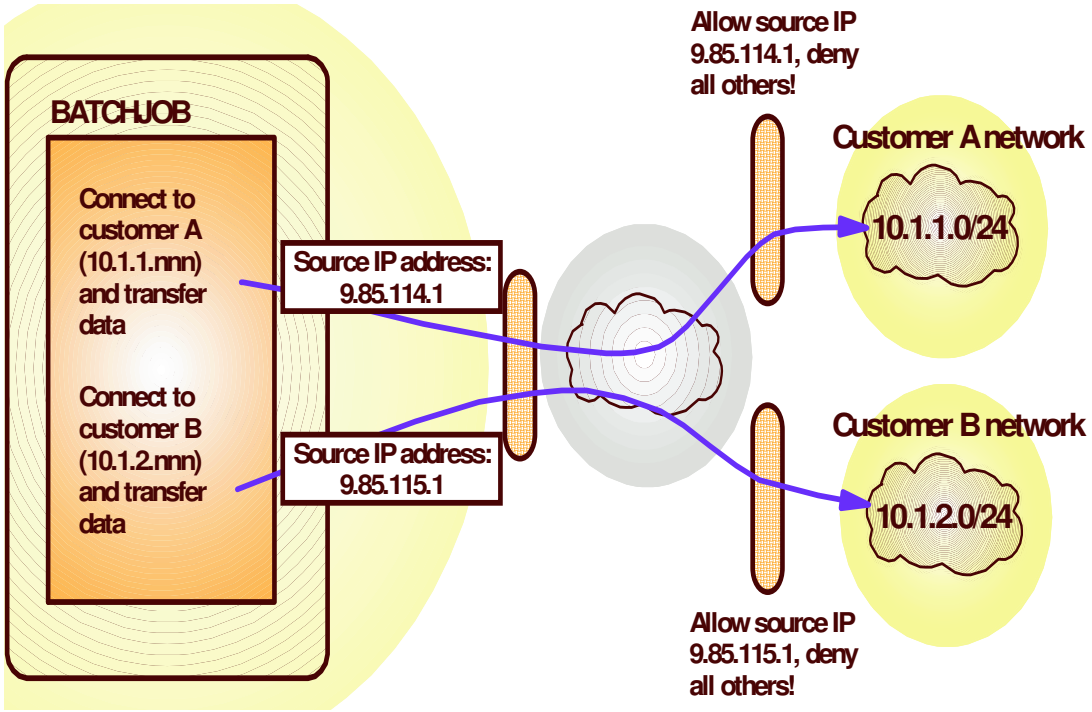

Source IP (SRCIP) Enhancements

- **Problem Statement / Need Addressed:**
  - Source IP address on DESTINATION rule in SRCIP block cannot be a distributed DVIPA
    - ✓ Source ports allocated for connections are not guaranteed to be unique across the sysplex
- **Solution:**
  - Allow a distributed DVIPA as the source IP address
    - ✓ Establish a pool of sysplex-wide unique ports
- **Benefit:**
  - SRCIP support based on destination IP address is more functional

Overview
Dynamic VIPA usability enhancements

- **Problem Statement / Need Addressed:**
  - Servers started from AUTOLOG, binding to a DVIPA fail initialization when DELAYJOIN is coded
  - Port range not allowed for VIPADISTRIBUTE ports
    - ✓ Ports have to be listed one at a time

- **Solution:**
  - Delay the starting of procedures that bind to a dynamic VIPA
    - ✓ Started after TCP/IP has joined the sysplex and created DVIPAs
  - Allow a range of ports to be configured on VIPADISTRIBUTE statement

- **Benefit:**
  - Customers can configure GLOBALCONFIG DELAYJOIN and specify servers that are to be automatically started in the AUTOLOG block
  - Customers can configure a range of port numbers on the VIPADISTRIBUTE PORT statement

Usage & Invocation
Dynamic VIPA usability enhancements

- The AUTOLOG support is enabled by:
  - Coding a new optional parameter, DELAYSTART on the AUTOLOG statement.
- The VIPADISTRIBUTE port range support is enabled by coding a range of ports on the statement
  - VIPADISTRIBUTE 203.1.1.94 **PORT 3006 3008-3010** DESTIP ALL

Overview
VARY TCP/IP,,SYSPLEX enhancements

- **Problem Statement / Need Addressed:**
  - Users can not quiesce or resume more than one unique application instance, listening on different ports, with one command
    - ✓ A command must be issued for each port bound by that application
  - Only Quiesce or Resume TARGET can be used for an application with multiple listeners on the same port
    - ✓ The Quiesce or Resume Port command is rejected if more than one instance of a listening application matches the values provided
- **Solution:**
  - Quiesce or resume all matching listeners regardless of port
    - ✓ New Quiesce/Resume JOBNAME command
  - Enhance Quiesce/Resume Port to quiesce or resume all matching listeners
  - All matching listeners **must** have the same jobname and asid
- **Benefit:**
  - More flexibility is available with the VARY TCP/IP,,SYSPLEX command

AGENDA

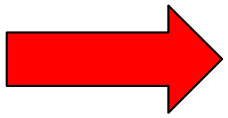

- OSA express enhancements
- Policy Enhancements
- Security enhancements
- Sysplex Enhancements
- Application enhancements
  - ***FTP RFC compliance***
  - ***FTP Kerberos SSO***
  - ***FTP Unicode support***
  - ***FTP select source addr***
  - ***TN3270 only in separate addr space***
- Management
- EE/SNA enhancements

Overview
FTP TLS/SSL Compliance

- **Problem Statement / Need Addressed:**
  - FTP's TLS/SSL support is not at the RFC 4217 level
  - Non-compliance may cause interoperability problems with other platforms
- **Solution:**
  - Implement RFC 4217
    - ✓ New TLSRFCLEVEL statement supported
- **Benefit:**
  - Full RFC 4217 functionality is available to z/OS FTP users
  - Customers don't have to be concerned with interoperability problems

FTP TLS/SSL Compliance

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

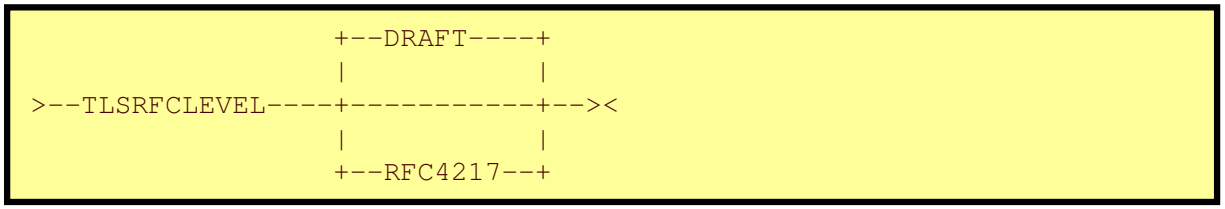

Overview
FTP Kerberos single sign-on support

- **Problem Statement / Need Addressed:**
  - The z/OS FTP server is not enabled for single sign-on when using Kerberos
- **Solution:**
  - Allow users to login to the z/OS FTP server without having to re-enter the password
    - ✓ Userid is still required
    - ✓ New FTP server statement SECURE_PASSWORD_KERBEROS
- **Benefit:**
  - Enables easier use of z/OS FTP server in Kerberos-based single sign-on environment.

FTP Kerberos single sign-on support

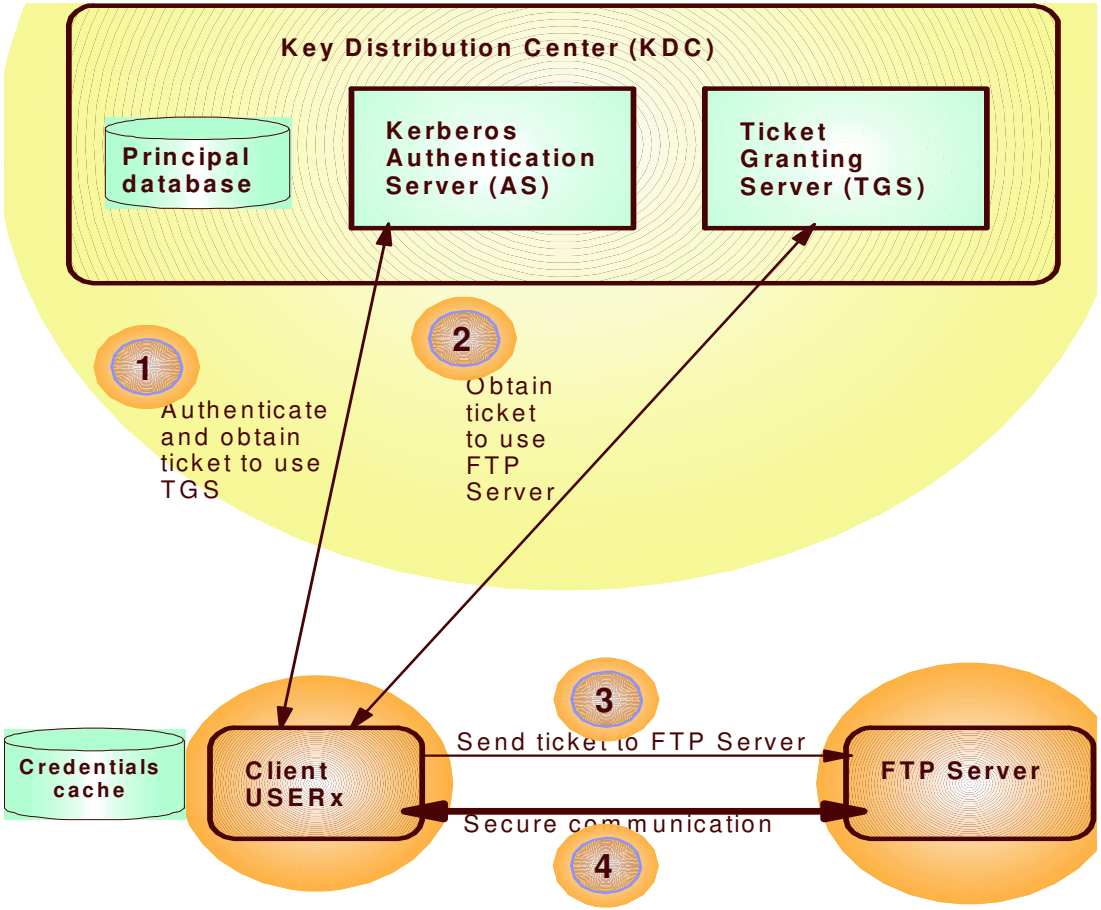

SECURE_PASSWORD_KERBEROS {REQUIRED | OPTIONAL}

Usage & Invocation
FTP Kerberos single sign-on support

- The support is enabled by a new FTP server statement
  - SECURE_PASSWORD_KERBEROS REQUIRED | OPTIONAL
    - If OPTIONAL is specified then the user ID used to logon must be the same as the user ID used to authenticate to the Kerberos KDC
      - If not, z/OS FTP server prompts for the password

FTP Unicode Support

- **Problem Statement / Need Addressed:**
  - FTP supports Unicode file transfer and storage for UTF-8 only
- **Solution:**
  - Support UTF-16 for file transfer and storage
- **Benefit:**
  - z/OS FTP can be used when exchanging data of different encodings between government and public agencies

Allow FTP client to select source IP address

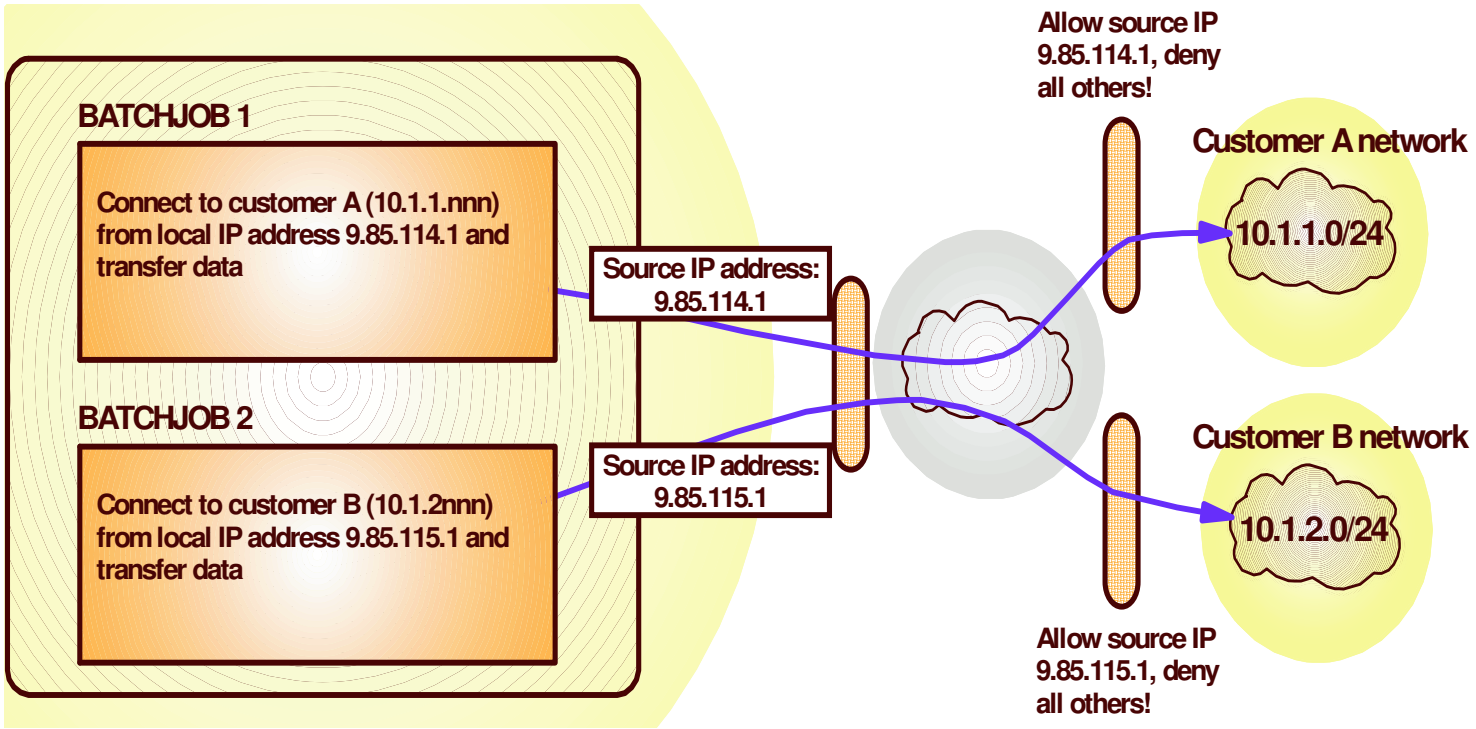

Usage & Invocation
Allow FTP client to select source IP address

- The support is invoked by the new command line parameter
  - ftp –s *srcip*
    - ✓ *srcip* must be a unicast IPv4 or IPv6 address.
    - ✓ Multicast, INADDR_ANY, IN6ADDR_ANY and IPv4-mapped IPv6 addresses are not supported.

Allow the TN3270E Telnet server only in a separate address space

- Prior to z/OS V1R6, the TN3270 server runs as a subtask of the IBM TCP/IP stack address space
- In z/OS V1R6, provide customers with a choice:
  - Run the TN3270 server as a separately started address space from TCP/IP (TSASO)
  - Continue to run TN3270 server as a subtask of the TCP/IP address space

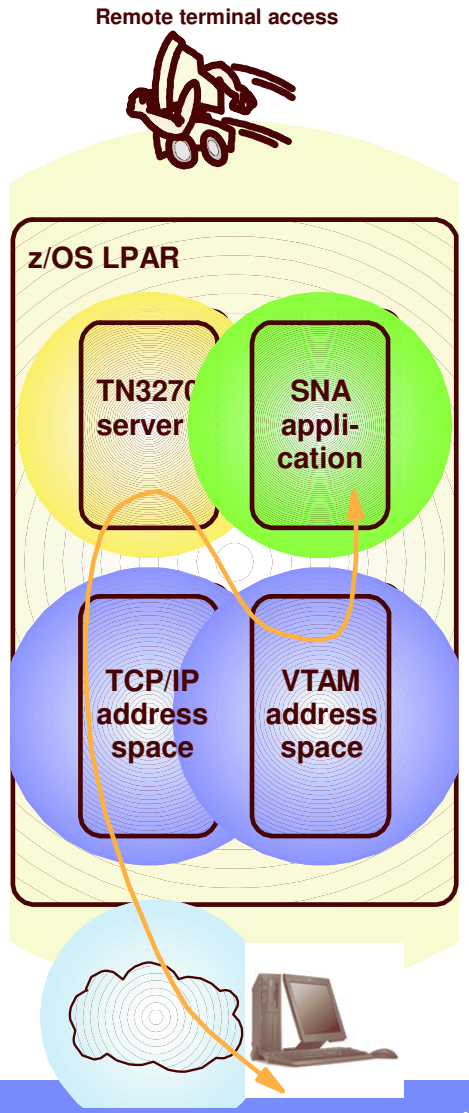

Allow the TN3270E Telnet server only in a separate address space

- **Problem Statement / Need Addressed:**
  - Dual support creates confusion
  - Duplicate development, test, and support effort
- **Solution:**
  - Stop supporting Telnet in the TCP/IP address space
- **Benefit:**
  - Allows for prioritization of TCP/IP address space vs. TN3270 server
  - Much less likely for TN3270 server failure to cause a total TCP/IP failure
  - Allow for easier problem diagnosis for both TCP/IP and TN3270
  - Easier controls for starting and stopping the server

AGENDA

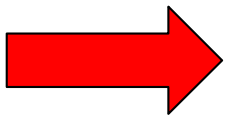

- OSA express enhancements
- Policy Enhancements
- Security enhancements
- Sysplex Enhancements
- Application enhancements
- Management
  - ***SNMP manager API***
  - ***Ping MTU discovery***
- EE/SNA enhancements

Provide a programming interface for SNMP manager

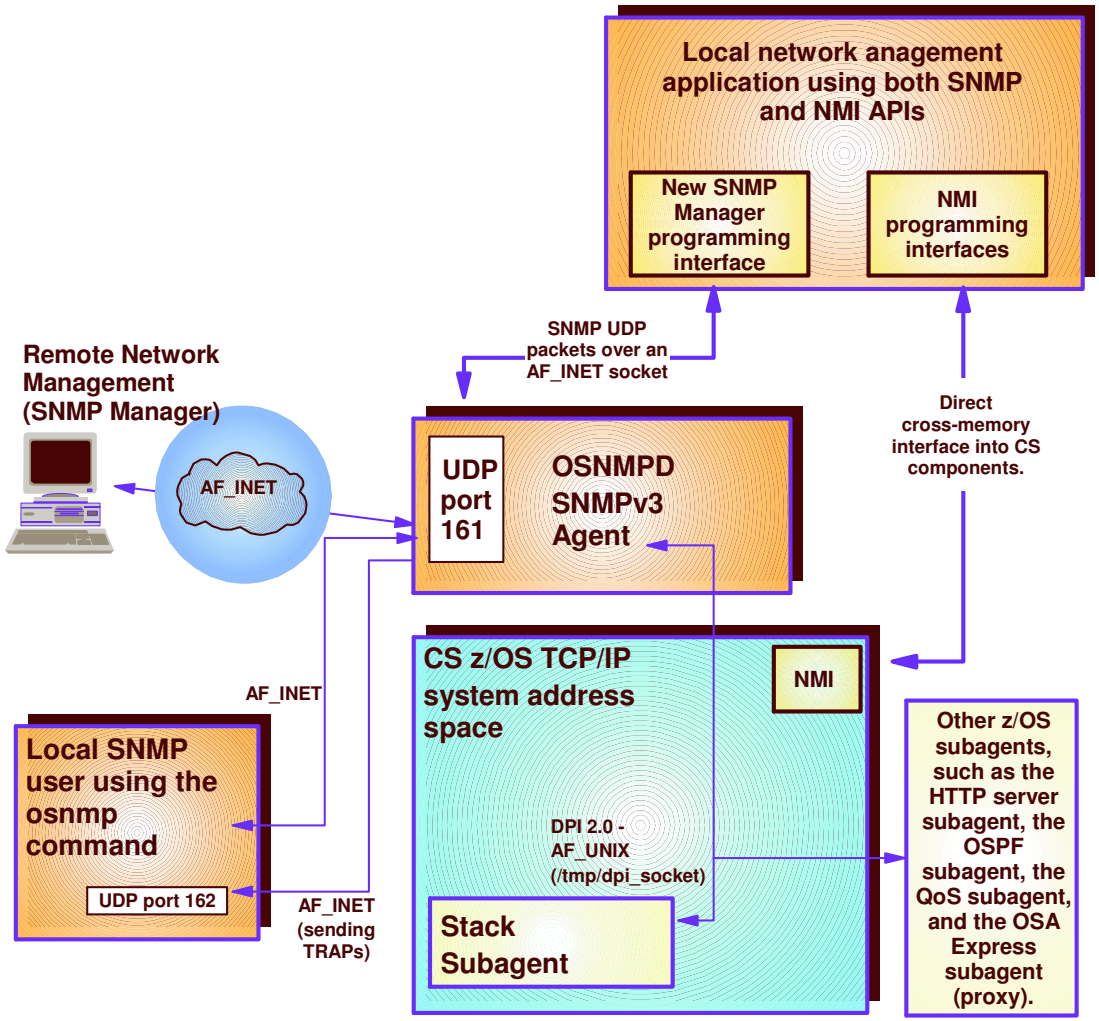

Usage & Invocation
Provide a programming interface for SNMP manager

- The function is enabled via the new API
  - Header file **snmpmgr.h** located in /usr/include/
  - Header file **snmpntfy.h** located in /usr/include/
  - A SNMP Manager sample, **snmpSMgr.c**, is provided
    - Located in /usr/lpp/TCP/IP/samples

Ping command detection of network MTU

- IPv4 Ping from Host A to Host V, with a length of 1400, will be fragmented by Host U
- IPv6 Ping from Host A to Host V will time out

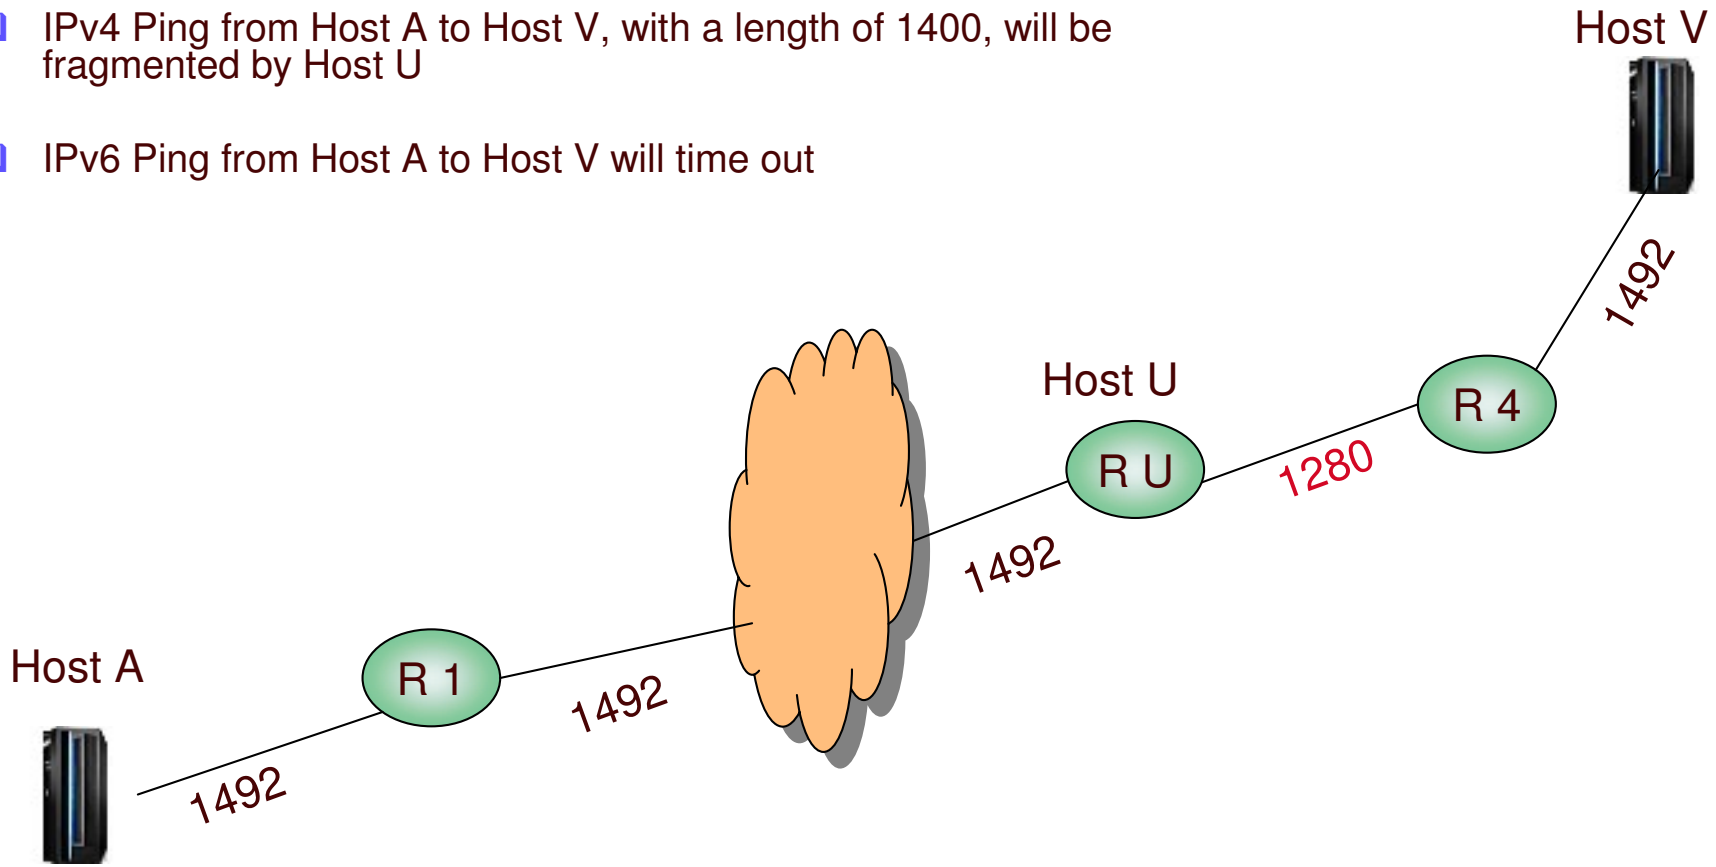

Ping command detection of network MTU

- **Problem Statement / Need Addressed:**
  - Difficult to determine where MTU problems exist in a network
- **Solution:**
  - Enhance the Ping command to detect MTU and fragmentation problems in a network
    - ✓ The host name, IP address of the host, and next-hop MTU value displayed
- **Benefit:**
  - MTU problems in the network are now easily detected

Usage & Invocation
Ping command detection of network MTU

- This support is enabled by specifying PMTU on TSO PING command
  - **PMTU YES | IGNORE**
- This support is also enabled by specifying –P on the z/OS Unix ping command
  - **-P yes | ignore**
- The current path MTU discovery value is used if **yes** is specified
- The current path MTU discovery value is ignored if **ignore** is specified
  - Allows a determination of where the MTU problem exist in the network

AGENDA

- OSA express enhancements
- Policy Enhancements
- Security enhancements
- Sysplex Enhancements
- Application enhancements
- Management
- EE/SNA enhancements
  - ***EE MTU discovery***
  - ***GRPREFS***
  - ***.....***

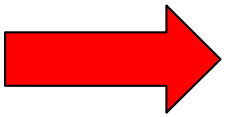

Local MTU Discovery for Enterprise Extender

- **Problem Statement / Need Addressed:**
  - The MTU size being utilized for an Enterprise Extender (EE) connection may not represent the current value.
  - An EE connection may not be utilizing an optimal route between the two endpoints.
- **Solution:**
  - EE dynamically learns when routing information, that pertains to EE connections, is changed and dynamically modify the MTU size that is used for that EE connection
    - ✓ The MTU size for the 1st hop is discovered
  - Utilize more optimal routes for an EE connection when they become available
- **Benefit:**
  - Avoid packet fragmentation
  - Avoid locking into a small MTU size for the life of the EE connection

Usage & Invocation
MTU Discovery for Enterprise Extender

- This function is automatically enabled when Enterprise Extender is being used.

Generic resources sample
Duplicate Load Balancing DDVIPAA and GRCICS

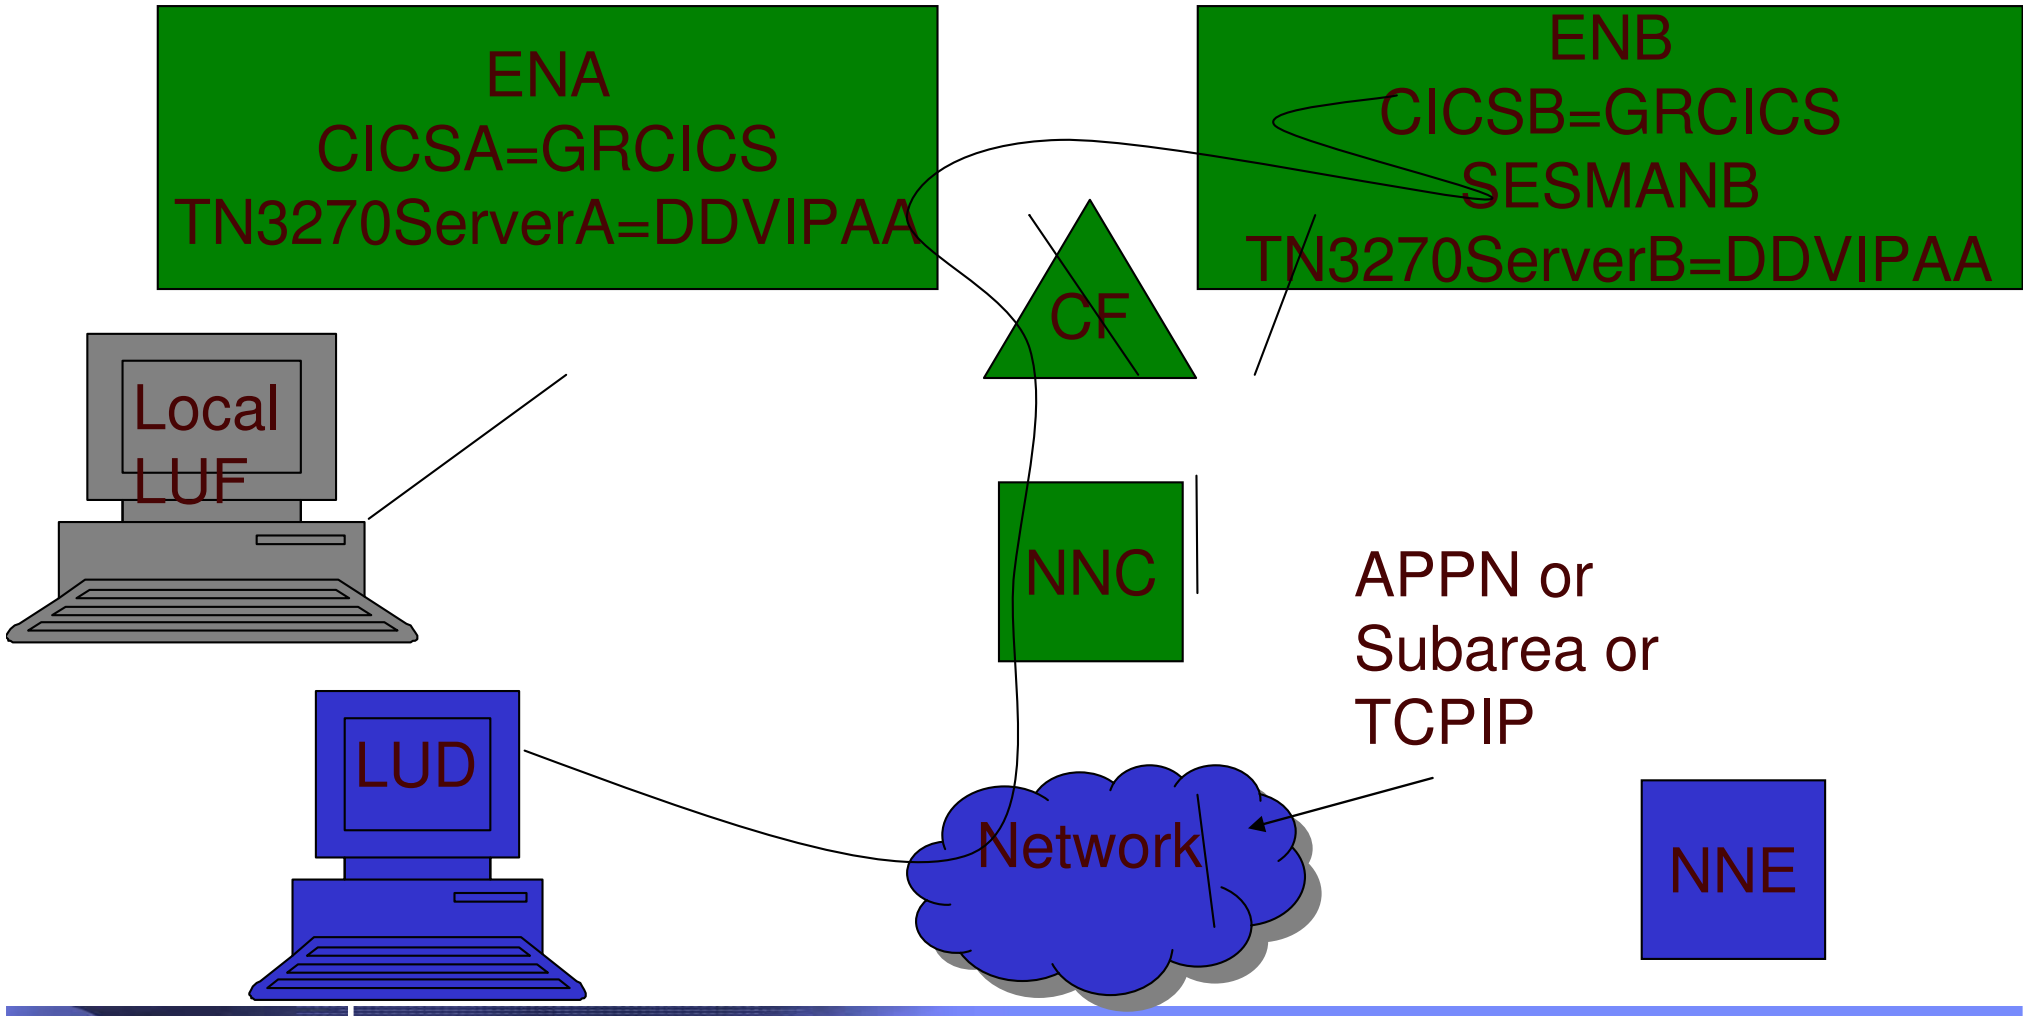

GRPREFS (1.9) Definition example

- Example of defining a GRPREFS Table named GRHOST01 with default GR preferences and GR preferences for generic resources GRCICS and GRTSO.

```
GRHOST01   VBUILD    TYPE=GRPREFS

           GRPREF     GREXIT=NO,WLM=YES,LOCAPPL=YES,LOCLU=YES,PASSOLU=NO
GRCICS     GRPREF     GREXIT=NO,WLM=NO,LOCAPPL=YES,LOCLU=YES,PASSOLU=YES
GRTSO      GRPREF     GREXIT=YES,WLM=YES,LOCAPPL=YES,LOCLU=YES,PASSOLU=NO
```

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)