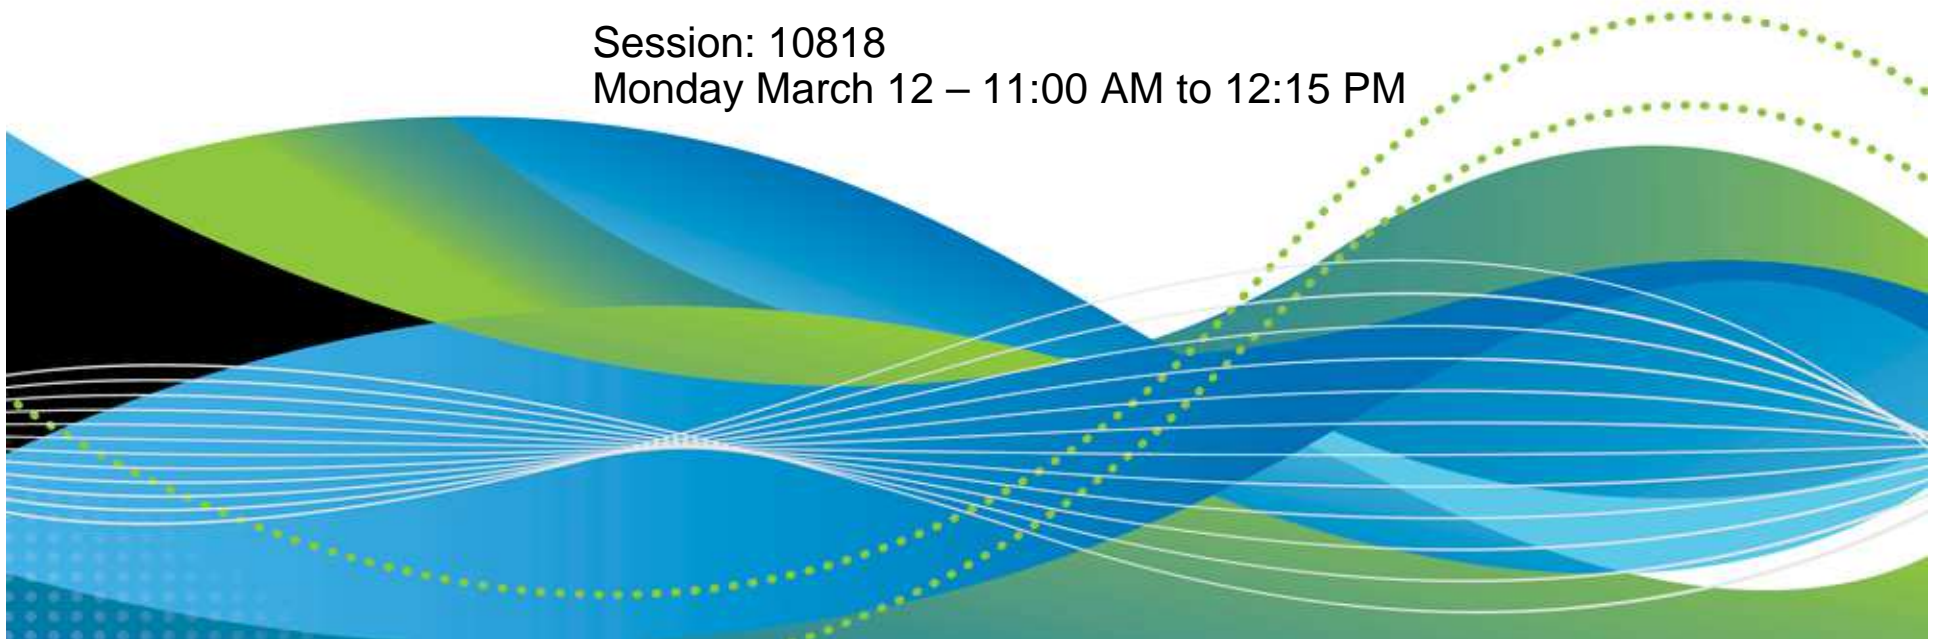# What's New in z/OS Communications Server

Gus Kassimis - kassimis@us.ibm.com
Sam Reynolds - samr@us.ibm.com
IBM Enterprise Networking Solutions
Raleigh, NC, USA

Session: 10818
Monday March 12 – 11:00 AM to 12:15 PM

What's New in z/OS Communications Server

Session number:	10818
Date and time:	Monday March 12, 2012 – 11:00 AM – 12:15 PM
Location:	Chestnut (Omni Hotel CNN Center)
Program:	Communications Infrastructure
Project:	Communications Server
Track:	Network Support and Management
Classification:	Technical
Speaker:	Gus Kassimis, IBM Sam Reynolds, IBM
Abstract:	The z/OS Communications Server combines TCP/IP and SNA support to better address the needs of today's complex networks. This session introduces new functions and capabilities for z/OS Communications Server, with a focus on the z/OS V1R13 CS release.

Trademarks, notices, and disclaimers

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®
- DataPower®
- DB2 Connect
- DB2®
- DRDA®
- e-business on demand®
- e-business (logo)
- e business(logo)®
- ESCON®
- FICON®

- GDDM®
- GDPS®
- Geographically Dispersed Parallel Sysplex
- HiperSockets
- HPR Channel Connectivity
- HyperSwap
- i5/OS (logo)
- i5/OS®
- IBM eServer
- IBM (logo)®
- IBM®
- IBM zEnterprise™ System
- IMS
- InfiniBand ®
- 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®
- POWER®
- POWER7®
- PowerVM
- PR/SM
- pSeries®
- RACF®

- Rational Suite®
- Rational®
- Redbooks
- Redbooks (logo)
- Sysplex Timer®
- System i5
- System p5
- System x®
- System z®
- System z9®
- System z10
- Tivoli (logo)®
- Tivoli®
- VTAM®
- WebSphere®
- xSeries®
- z9®
- z10 BC
- z10 EC

- zEnterprise
- zSeries®
- z/Architecture
- z/OS®
- z/VM®
- z/VSE

* All other products may be trademarks or registered trademarks of their respective companies.

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:
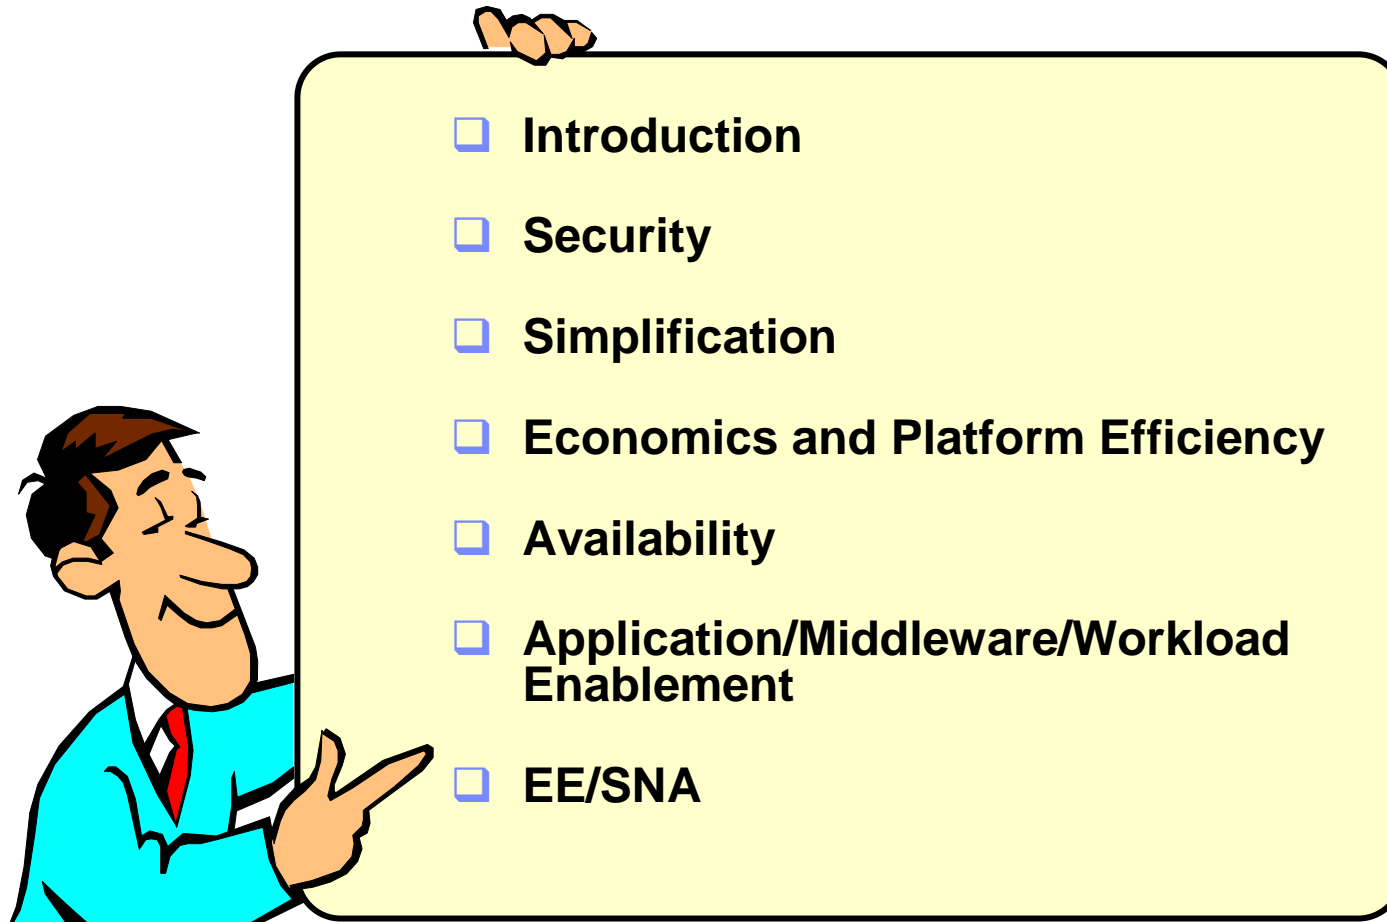- Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
- Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license there from.
- 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.
- InfiniBand is a trademark and service mark of the InfiniBand Trade Association.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
- IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

Notes:
- Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
- IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
- All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
- This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
- All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
- Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.
- Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.

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

Agenda

- **Introduction**
- **Security**
- **Simplification**
- **Economics and Platform Efficiency**
- **Availability**
- **Application/Middleware/Workload Enablement**
- **EE/SNA**

z/OS V1R13 GA: September 2011

Disclaimer: All statements regarding IBM future direction or intent, including current product plans, are subject to change or withdrawal without notice and represent goals and objectives only. All information is provided for informational purposes only, on an "as is" basis, without warranty of any kind.

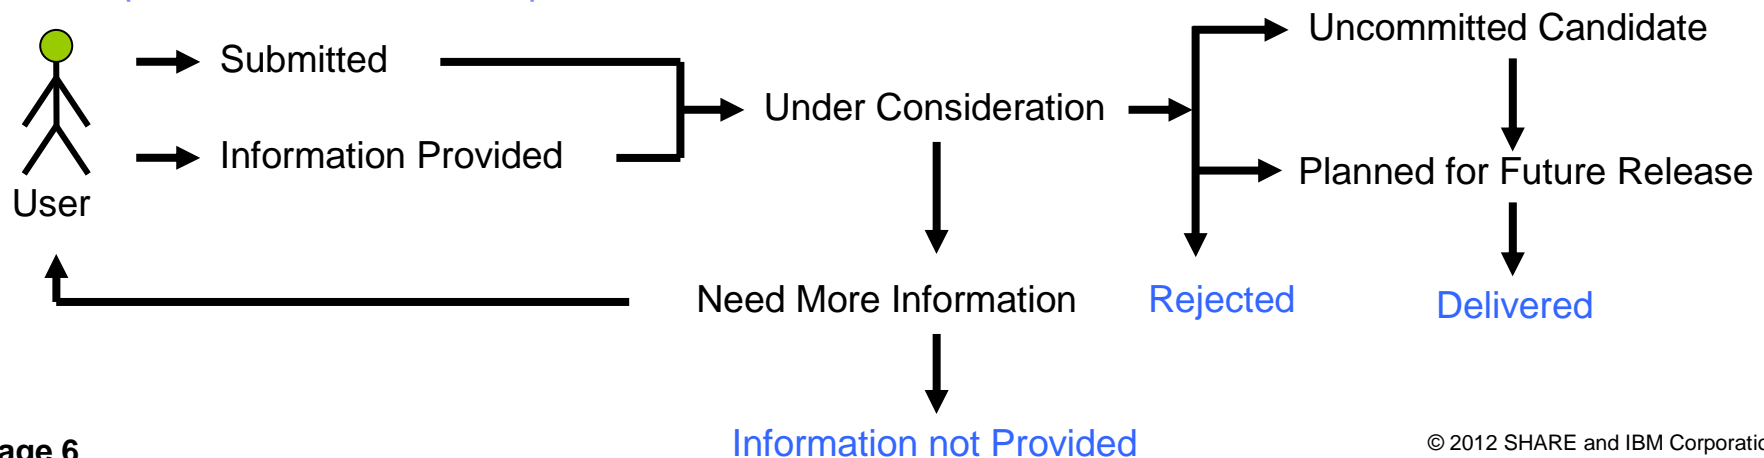# What will the z/OS community need from z/OS networking in 2012-2013?

- **System z technology is expected to continue to evolve**
  - Networking software need to support new technologies such as zEnterprise

- **Access to System z system-level skills will continue to be an issue**
  - Retiring existing people, who grew up with system z
  - New people becoming responsible for the overall system z environment – including z/OS networking
  - Note: follow the IBM Academic Initiative
    - https://www.ibm.com/developerworks/university/academicinitiative/

- **Security will continue to be a hot topic**
  - Per customer survey, over 50% of network traffic will need encryption within the next few years
  - Trade organizations and governments continue to establish security and privacy compliance requirements that must be met

- **Price/performance requirements are high priority**
  - Continued demand for reduced cost in combination with increased performance and scalability on system z

- **Demand for increased "autonomic" system integration capabilities**
  - Continued demand for improved integration with other hardware and software platforms for more complex heterogeneous solutions

- **IANA has already run out of IPV4 addresses. Regional registries are also running out (APNIC has already run out, RIPE is expected to be next)**
  - IPv6 compliance (USGv6, IPv6-Ready, TAHI test suite, etc.)

RFE: New Requirements Process for z/OS Communications Server

- **RFE (Request for Enhancement)**
  - New web-based interface for submitting requirements to some IBM products
  - Unlike FITS, RFE is available to anyone that signs up for an IBM Universal ID.
  - Can submit requirements, vote on requirements, watch and comment on requirements, and interact with IBM if more information is needed
  - RFEs submitted against z/OS CS will initially be private, but most will be converted to public status after an initial screening.
  - Even public RFEs do have certain fields that will always be private (only viewable by the submitter and IBM):
    - Customer name
    - Business justification
    - Attachments can be private (your choice at submission)
  - Can keep up with requirements you are interested in via watchlists, bookmarks, and/or RSS feeds
  - http://www.ibm.com/developerworks/rfe/?PROD_ID=498

User → Submitted

User → Information Provided

→ Under Consideration → Uncommitted Candidate

→ Planned for Future Release

Need More Information

Rejected

Delivered

Information not Provided

RFE: New Requirements Process for z/OS Communications Server …

© 2012 SHARE and IBM Corporation

RFE: New Requirements Process for z/OS Communications Server …

What's Coming in z/OS Communication Server

Security

IBM®

Review: Intrusion Detection and Prevention services on z/OS

IDS Policy administrator

IDS Policy

z/OS CS Policy infrastructure

Applications

TCP/UDP

IPv4 & IPv6

Interfaces

Attack!!!

Attack Probe

Trmdstat reporting or other auditing tools

Security Auditor

Syslogd

Detailed event messages to z/OS USS Syslogd

Intrusion Event Notification

z/OS NetView

Tivoli Security Operations Manager (TSOM)

MVS Console Operator

Selected event messages to MVS console

Automation based on MVS console messages

Dynamic packet trace of suspicious activity

CTRACE

Network Engineer – detailed P/D

- ❑ **Events detected**
  - Scans
  - Attacks against stack
  - Flooding (both TCP and UDP)
- ❑ **Defensive methods**
  - Packet discard
  - Limit connections
- ❑ **Reporting**
  - Logging
  - Event messages to local console
  - IDS packet trace
  - Notifications to Tivoli NetView and Risk Manager
- ❑ **IDS Policy**
  - Samples supplied with z/OS CS Configuration Assistant

The current IDS support is for IPv4 traffic only!

Intrusion Detection Services enhanced to include IPv6 traffic

- **Attack types supported for both IPv4 and IPv6**
  - **Scan**
    - TCP and UDP scan event rules
    - ICMP scan event rule – unchanged
    - ICMPv6 scan event rule – new
    - Scan exclusion list
  - **Traffic regulation (TR)**
    - TCP TR – IPv4 and IPv6 connection requests monitored
    - UDP TR – IPv4 and IPv6 packets monitored
  - **Attack types extended to IPv6**
    - Malformed packet events – IPv6 packets dropped due to malformed headers, options, or values.
    - UDP perpetual echo
    - ICMP redirect restrictions – extended to apply to ICMPv6 redirects
  - **Flood attacks**
    - SYN flood – extended to IPv6 connection requests
    - Interface flood

- **Plus, new attack types for IPv6-specific vulnerabilities**
  - Restricted IPv6 Next Header
  - Restricted IPv6 Hop-by-hop Options
  - Restricted IPv6 Destination options

- **Note:** Defense Manager daemon already supports IPv6 and does not need upgrading

IBM

New IDS attack types implemented for both IPv4 and IPv6

- **TCP Window Attack – global stall attack**
  - Prevents an attacker from creating multiple connections with zero window sizes and keeping them open indefinitely

- **Data Hiding**
  - Prevents an attacker from hiding data in reserved fields
  - PadN options in IPv6, reserved fields in IPv4 headers

- **TCP Queue Size**
  - Helps you manage the amount of storage TCP can take up for the queues used for holding sent and received data
    - For example, out of order packets awaiting re-sequencing
  - Provides user control over storage constraint availability improvements added in z/OS V1R11
  - Helps avoid TCP causing storage constraint situations
  - IDS policy enables dropping of connections that exceed specified limits

Stalled application

Local receiving application

Full receive buffer

Local TCP/IP stack on z/OS

Remote TCP/IP stack

Send buffer

Remote sending application

Local sending application

Full send buffer

Local TCP/IP stack on z/OS

Remote TCP/IP stack

Receive buffer

Remote receiving application

Stalled application
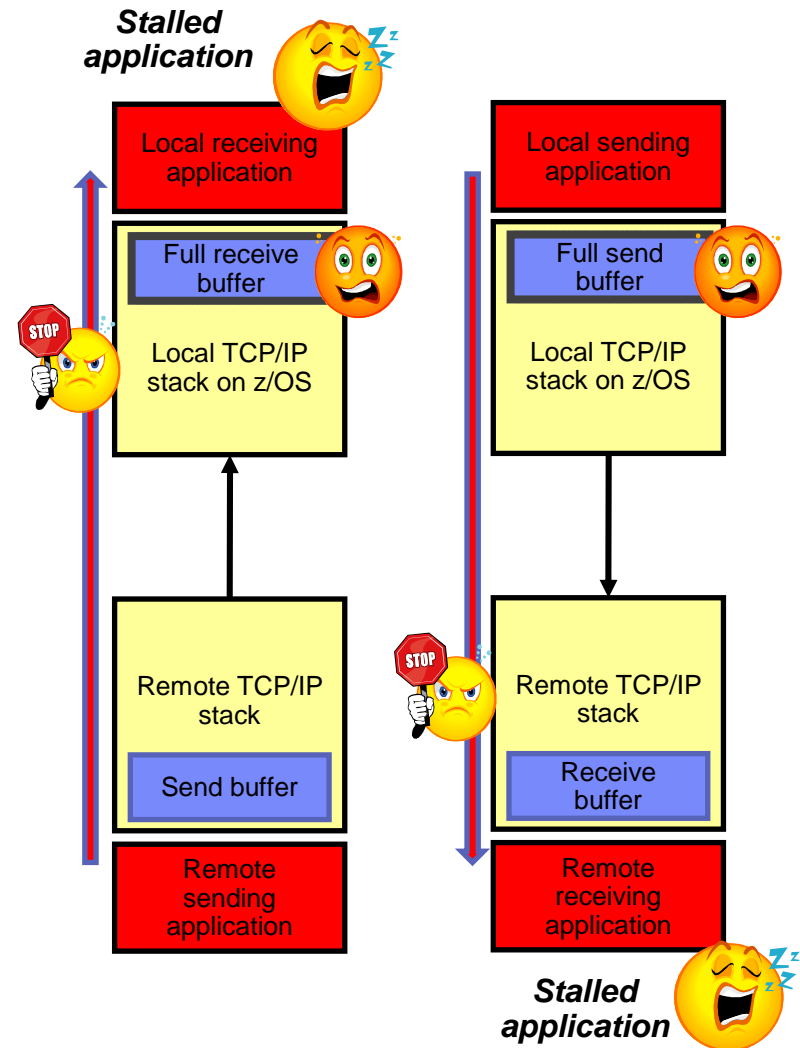
Support for Network Address Translation when using IKEv2

- Network address translation (NAT) is commonly used in enterprises to conserve IPv4 addresses

- In z/OS V1R12 we added support for IKEv2, which was required for IPv6 currency
  - IKEv2 supports both IPv4 and IPv6

- We encourage our customers to move to IKEv2, but for many IPv4 customers, NAT is a requirement
  - Currently supported on IKEv1, support being added to IKEv2 to encourage migration
    - Enables end to end network encryption in NAT configurations!

Private address space   **Public address space**   **Private address space**   z/OS

10.1.1.1

NAT   **NAT**

10.1.1.2   Intranet   Internet   Intranet

10.1.1.3

210.1.1.1   211.1.1.1
210.1.1.2
210.1.1.3

Data stream
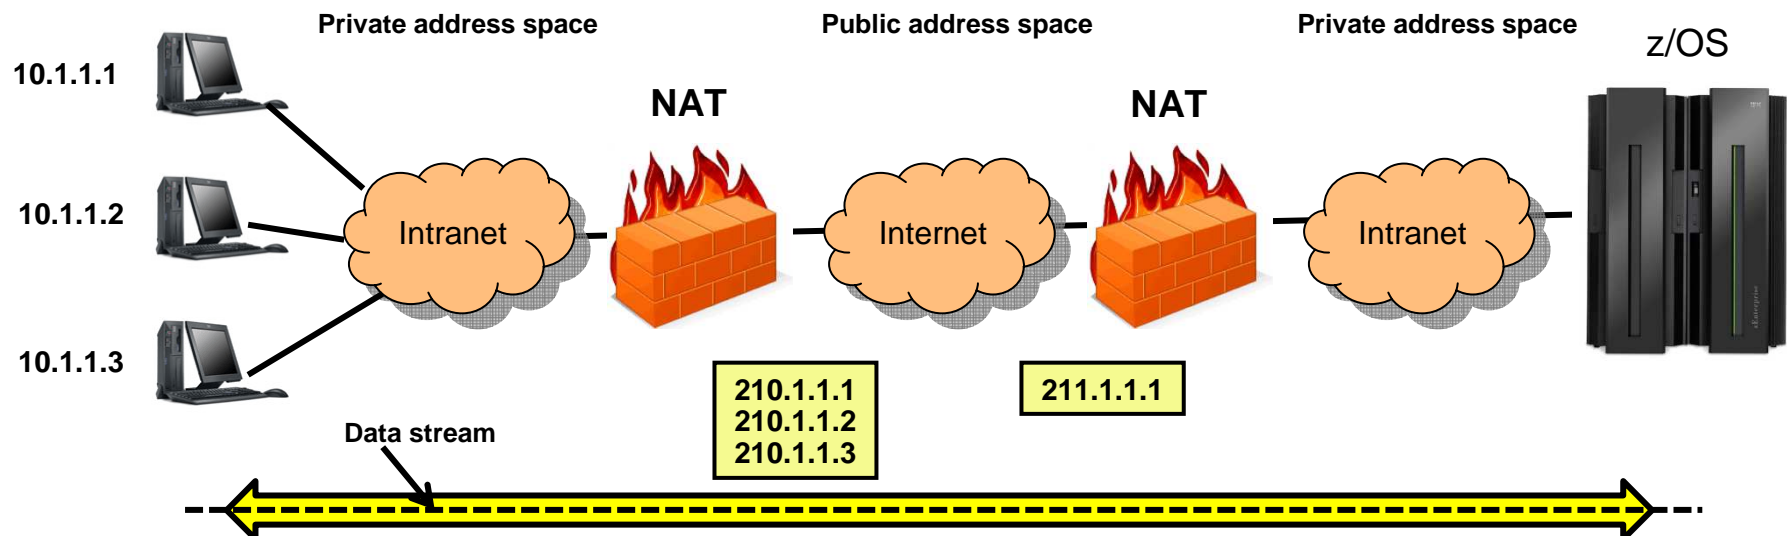
Password phrase support in selected servers

- Password
  - One to eight characters
  - Limited range of characters allowed (for example, no blanks in the password)

- Password phrase
  - Nine to one hundred characters
  - Can contain any characters allowed in the EBCDIC 1047 code page
    - Includes spaces and punctuation characters
    - But not a NULL character
    - Case sensitive

- Every user ID with a password phrase also has a password (since V1R10)
  - Applications that use passwords to validate users in RACF need to accept and use these longer passwords
  - Current z/OS Communications Server functions that verify password, restrict the password to be no longer than 8 characters
  - Support for password phrases added to FTP and TN3270 in z/OS V1R13
    - TN3270 support is for solicitor screen only.
      - Application password controls not affected
  - Enables applications and users that use these servers to exploit password phrases
    - For example, applications that call FTP through an API and want to use password phrases

LONGER PASSWORDS MAKE STRONGER PASSWORDS

Enhanced Dynamic VIPA binding security

- Application-specific Dynamic VIPAs are virtual IP addresses that are created when applications request (bind to) them and removed when they give them up.
  - Provides improved availability, for example dynamic VIPA can move around in the Sysplex, following the application when it moves, so clients are uninterrupted.

- Currently there is global security around creation and destruction of dynamic VIPAs.
  - An application can be permitted to create and destroy all dynamic VIPAs
    - EZB.BINDDVIPARANGE.sysname.tcpname – all VIPARANGE defined VIPAs
    - EZB.MODDVIPA.sysname.tcpname – ability to issue MODDVIPA or SIOCSVIPA

- z/OS V1R13 adds more granularity by providing ability to control which applications can create and remove specific DVIPAs or DVIPA ranges.
  - Allow an application to create/remove its own DVIPAs but prevent it from interfering with other applications' ranges.
  - Prevent an application from inadvertently removing another application's DVIPA
  - New keyword "SAF resname" supported on the VIPARANGE statement
    - Identifies a VIPARANGE statement using this profile
    - If keyword not present, VIPARANGE statement uses existing profiles
  - EZB.BINDDVIPARANGE.sysname.tcpname.resname

```
VIPARANGE DEFINE 255.255.255.255   20.20.20.1 SAF APPL1
```

To Bind to 20.20.20.1 – user must be permitted to EZB.BINDDVIPARANGE.sysname.tcpname.APPL1
To MODDVIPA 20.20.20.1 – user must be permitted to EZB.MODDVIPA.sysname.tcpname.APPL1

What's Coming in z/OS Communication Server

Simplification

IBM

IBM

Review: IBM Configuration Assistant for z/OS Communications Server review

Policy Agent

z/OS MF

WAS

z/OS

Windows

As-Is, single user

Fully Supported, Multi User
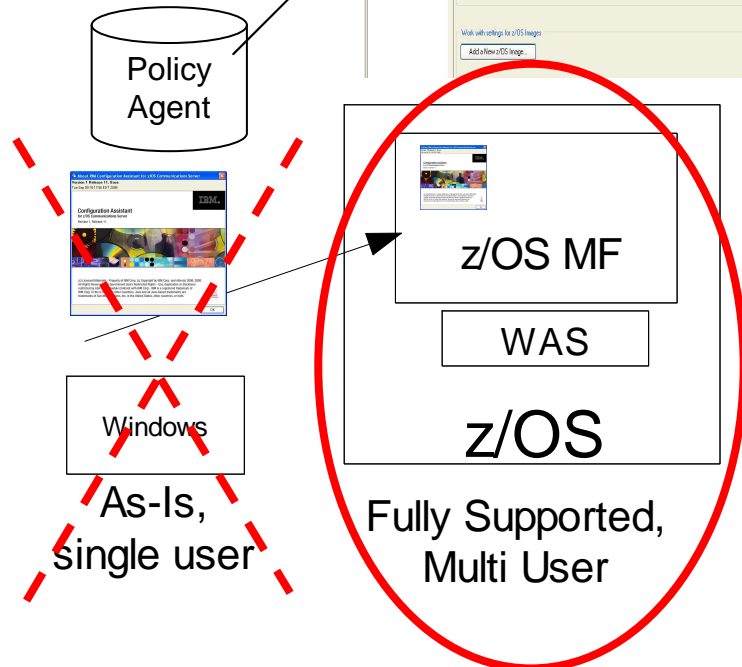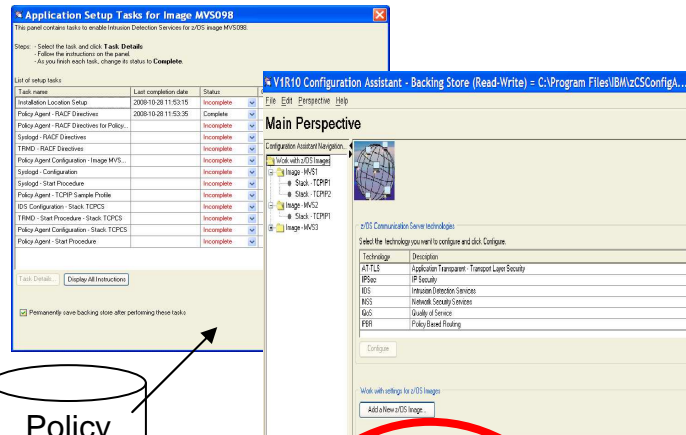
- As of z/OS V1R11, IBM Configuration Assistant for z/OS Communications Server is integrated with z/OS Management Facility (z/OSMF)
  - z/OSMF version is integrated into the product and runs on z/OS.
  - z/OSMF version is officially supported.

- The standalone Windows version is still available, but is made available as-is, without any official support:
  - http://www.ibm.com/support/docview.wss?rs=852&context=SSSN3L&dc=D400&uid=swg24013160&loc=en_US&cs=UTF-8&lang=en&rss=ct852other
    - Or
  - http://tinyurl.com/cgoqsa

Page 18

© 2012 SHARE and IBM Corporation

Statement of Direction: IBM Configuration Assistant for z/OS Communications Server

z/OS V1.13 is planned to be the final release for which the IBM Configuration Assistant for z/OS Communications Server tool that runs on Microsoft Windows will be provided by IBM. This tool is currently available as an as-is, nonwarranted web download. Customers who currently use Windows-based IBM Configuration Assistant for z/OS Communications Server tool should migrate to the z/OS Management Facility (z/OSMF) Configuration Assistant application. The IBM Configuration Assistant for z/OS Communications Server that runs within z/OSMF is part of a supported IBM product and contains all functions supported with the Windows tool.

Configuration Assistant support for multiple releases

- z/OSMF direction is for one instance in a sysplex
  - For customers who have a mixed-release environment, this means that one instance has to manage multiple releases
  - Currently, the version of Configuration Assistant that ships with a release of Communications Server can only update that version
  - Multiple release support will allow one instance of Configuration Assistant to manage hosts in a mixed-release sysplex environment.
    - Releases z/OS V1R12 and z/OS V1R13 supported in z/OS V1R13
    - z/OSMF version only – multiple release support not provided in as-is Windows version
  - When creating a new operating system image, the user can now set the release of that image
  - The release of the operating system image can be changed at any time
  - Options that only apply to the higher-level release, are ignored if the image is currently at the lower level
    - New IDS attach types are ignored if the z/OS image is z/OS V1R12, but included if the image is z/OS V1R13

Configuration Assistant backing store

z/OS V1R13

z/OS V1R13

z/OSMF

z/OS V1R12 policy

z/OS V1R13 policy
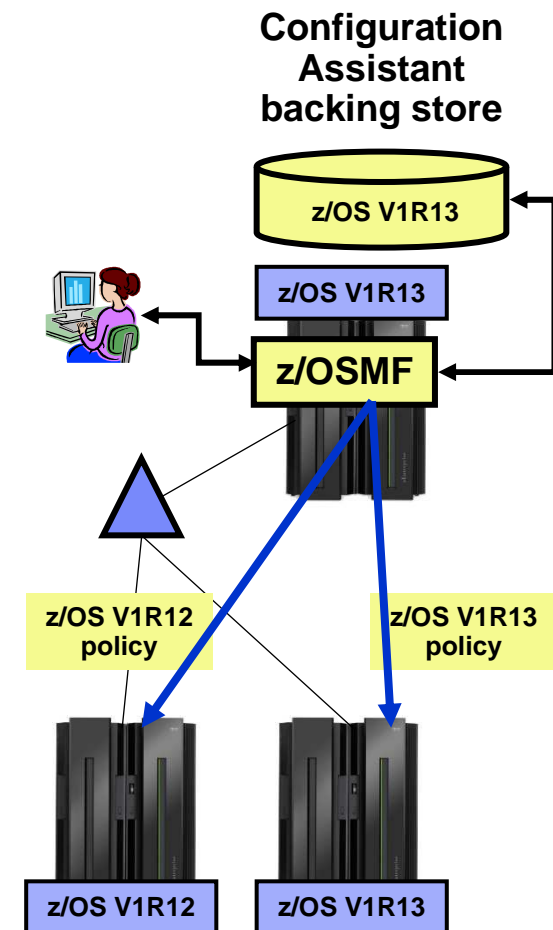
z/OS V1R12

z/OS V1R13

Configuration Assistant import of TCP/IP configuration information

TCP/IP Profile

Policy Agent

z/OSMF

WAS

z/OS
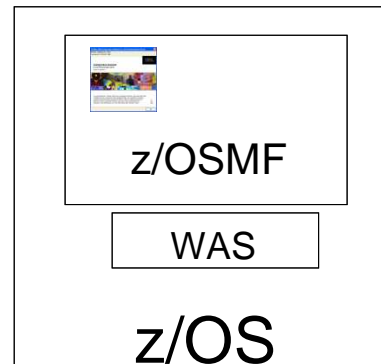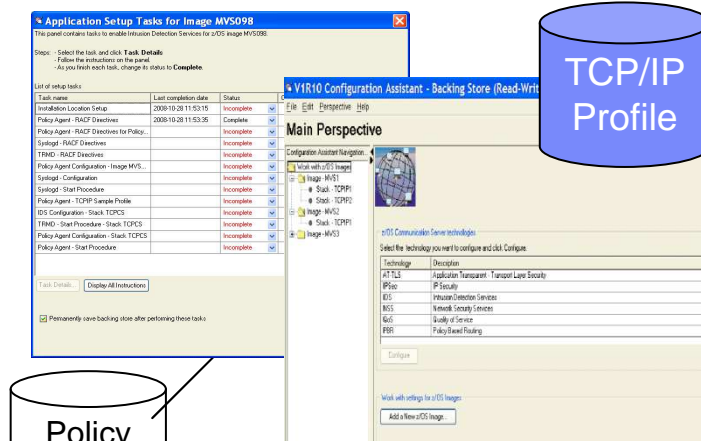
Fully Supported, Multi User

- For z/OS V1R13, Configuration Assistant will import profile information from running TCP/IP stacks
  - Will be used to help develop policy configuration
    - Examples: Learn home addresses, suggest address groups, etc.

- This function will be provided for the z/OSMF version only
  - Not supported in windows configuration assistant

- Along with this support will be support for defining a policy rule once for multiple stacks, without having to individually define every policy rule for every stack

What's Coming in z/OS Communication Server

Economics and platform efficiency

IBM®

Enhanced support for OSA VLANs

- In z/OS V1R10, Communications Server added multiple VLAN support
  - Up to 8 IPv4 and 8 IPv6 VLANs per OSA port
  - Separate INTERFACE statement and data device per VLAN
  - The value of 8 is a z/OS CS software limitation

- Raise limit from 8 to 32 VLANs per stack per OSA port
  - No impact to OSA
  - Driven by emphasis on VLANs for network isolation in the zEnterprise IEDN

```
INTERFACE1      INTERFACE2      INTERFACE3           INTERFACE32
VLAN1           VLAN2           VLAN3                VLAN32
VMAC1           VMAC2           VMAC3      . . .     VMAC32
10.1.1.0/24     10.1.2.0/24     10.1.3.0/24          10.1.32.0/24
```

**OSX
10 GbE**

VLAN1 .. VLAN32   **IEDN**
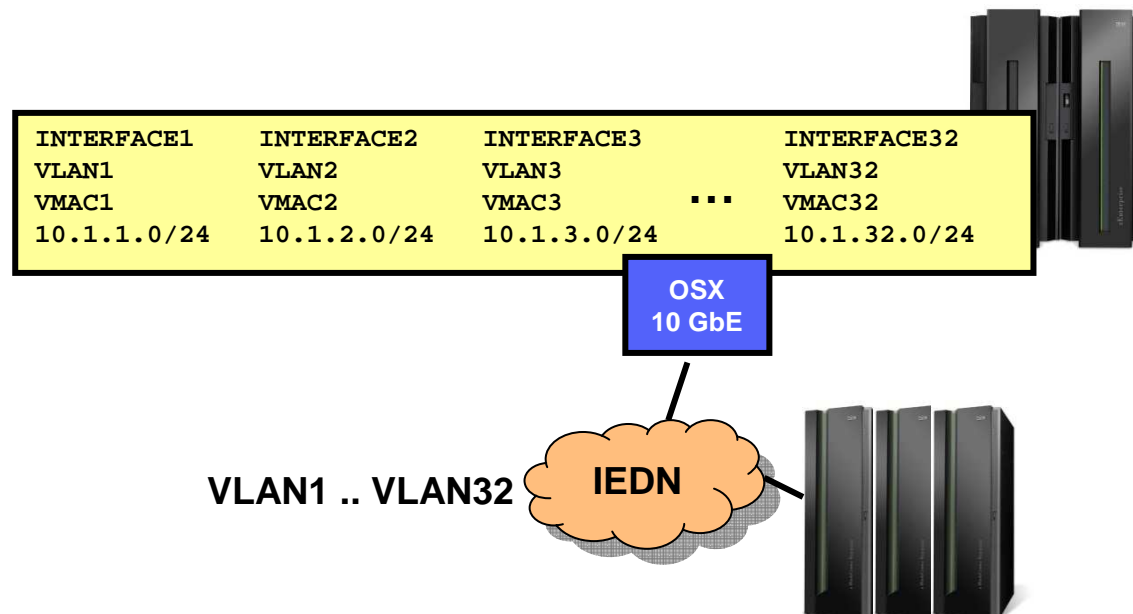
zEnterprise IEDN without Hipersockets

.... Intra Ensemble Data Network with platform managed virtualization, isolation and access controls

CEC X

z/OS LPAR1

IntfA = IP A

Host Interface
IEDN

OSA NIC

z/OS LPAR2

VLAN 5

IntfB = IP B

Host Interface
IEDN

OSA NIC

z/OS LPAR3

VLAN 4

IntfC = IP C

Host Interface
IEDN

OSA NIC

z/OS LPAR4

VLAN 5

IntfD = IP D

Host Interface
IEDN

OSA NIC

PR/SM

OSX

SE

IEDN

zBX ASB

Linux1

IntfG = IP G

Host Interface
IEDN

VNIC

Linux2

VLAN 5

IntfH = IP H

Host Interface
IEDN

VNIC

Linux3

VLAN 4

IntfI = IP I

Host Interface
IEDN

VNIC

Hypervisor

NIC

VSwitch

HMC
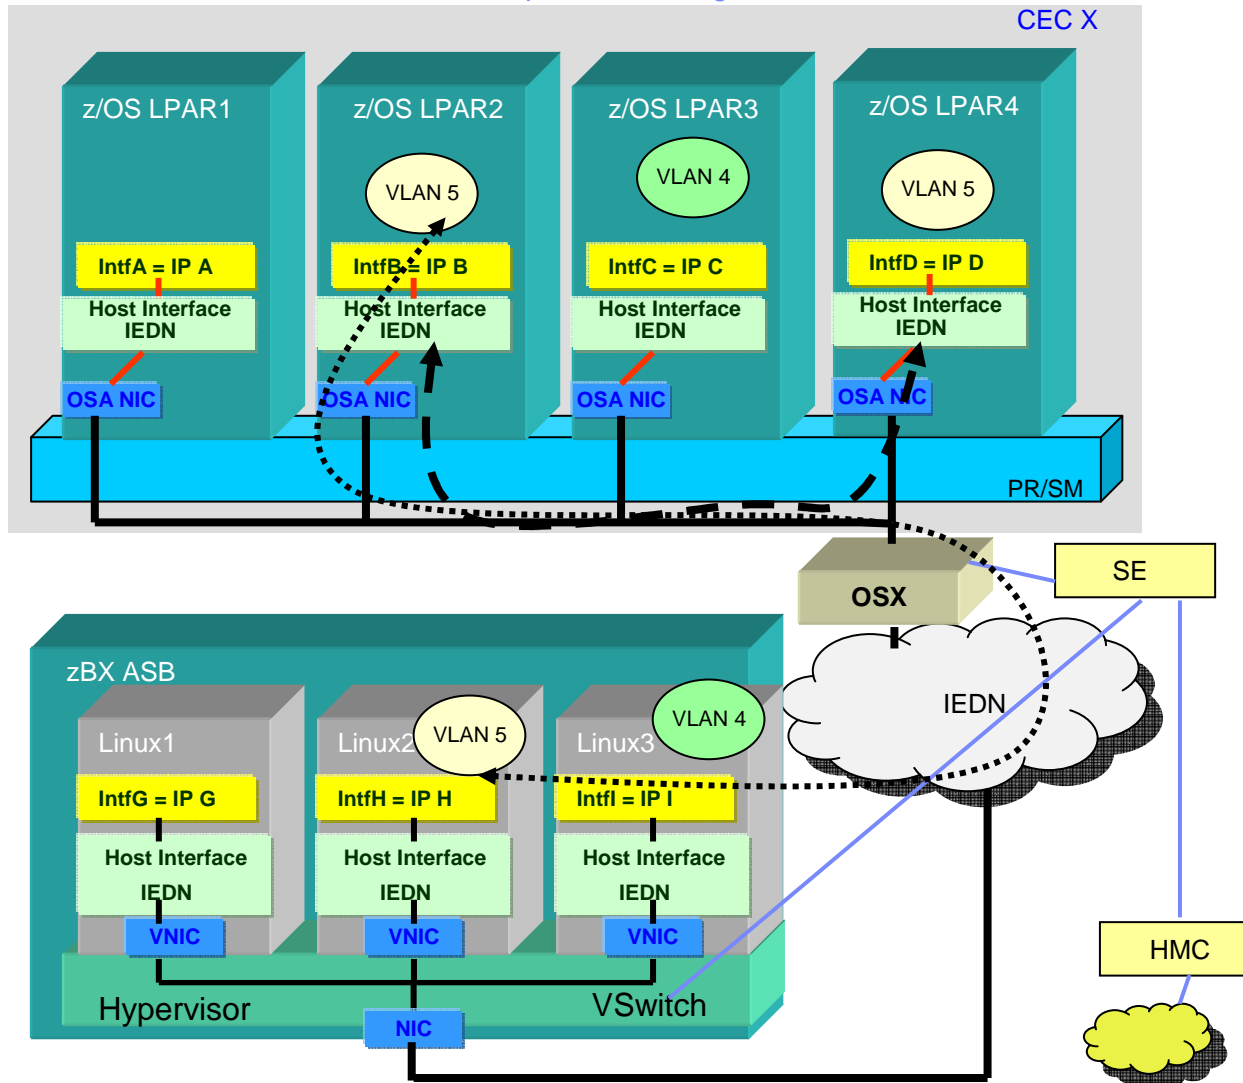
HiperSockets is another type of System z
internal network that is a
System z differentiator!

…yet HS is missing from the IEDN … in order to exploit HS it requires explicit and separate network config (IP address,

IP route, OS config etc.)

IEDN enabled HiperSockets

.... Intra Ensemble Data Network with platform managed virtualization, isolation and access controls
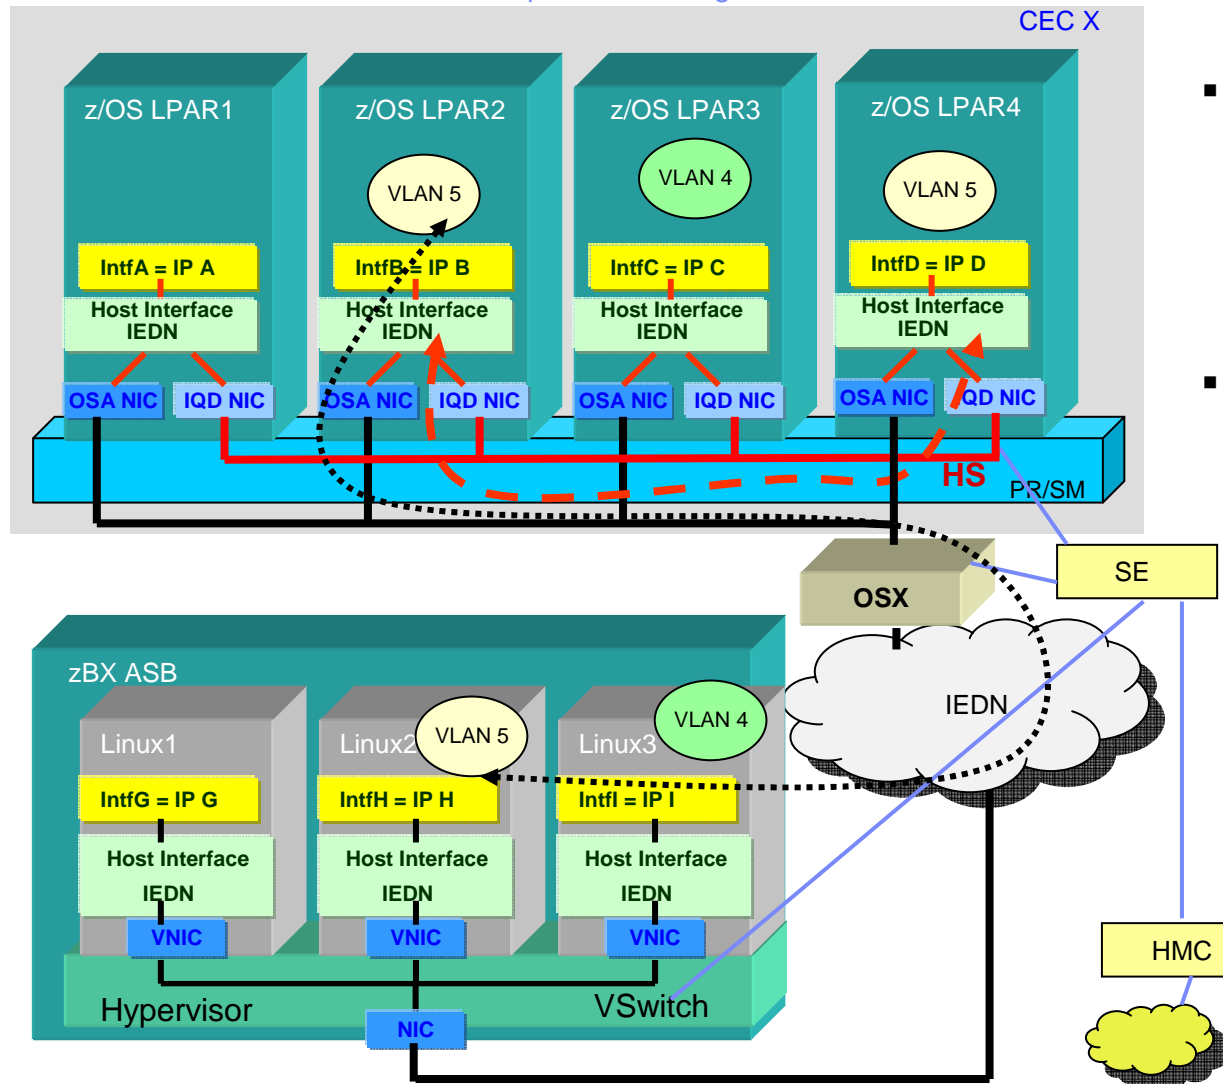
- **HiperSockets becomes part of the IEDN**
  - z/OS support in V1R13
  - Announcement of support on zEnterprise (now) and zVM 6.2 support (April 2012)

- **In a transparent manner**
  - The virtual servers present a single IP address (their IEDN address) for both internal (HiperSockets) and external (IEDN) access
    - No IP topology changes or routing changes required
    - The optimal path is selected automatically without requiring unique routing configuration
  - Also enables relocation of System z virtual servers across z CECs without reconfiguration
    - Same IP address used
    - Current HiperSockets IP topology is CEC specific
      - Moving to another CEC requires IP address and routing changes.

IEDN enabled HiperSockets - z/OS "Converged IQDX Link" Concepts

z/OS

All IP traffic to / from this IP subnet

Single IP Address (single Network Interface)

Converged

Interface

The z/OS administrator only configures / manages the OSX interface

The zEnterprise Ensemble network administrator only configures VLANs for OSX

Communications Server transparently splits and converges network traffic to this interface

OSX IQDX

With **AUTOIQDX…**
the IQDX interface is dynamically / transparently configured / added ("tucked" under OSX)

OSX

IQD (IQDX)

Only OSX connectivity must be configured!

IQDX connectivity is transparent!

Internal IEDN

IEDN

The OSX link provides access to the "External IEDN fabric"

… and provides access to all servers on both the Internal and External Network

The IQDX link provides access to the "Internal CPC IEDN fabric" (optimal path)

… and provides access to just the subset of servers on the "Internal (IQDX portion of the) IEDN"

HiperSockets *IEDN Access IQDX* Configuration
zVM HiperSockets VSwitch Bridge

CEC X

z/OS LPAR
Server A (native)

Linux LPAR
Server B (native)

z/VM LPAR

real = dedicated

z/OS Guest
Server C (real)

Linux Guest
Server D (real)

Linux Guest
Server E (simulated)

Linux Guest
Server F (simulated)

z/OS Converged Interface

If 1

OSX NIC   IQD NIC   OSX NIC   OSX NIC   IQD NIC   IQD NIC

OSX NIC
Guest Port

OSX NIC
Guest Port

HS Bridge Port

OSA Uplink Port

IEDN VSwitch

HiperSockets (IQDX)

PR/SM

z/OS is a non bridge capable guest

bridge capable guest (QEBSM)

OSX

Primary
OSX

Backup
OSX

IEDN

For more information on recent announcement:
http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS112-026/ENUS112-026.PDF

© 2012 SHARE and IBM Corporation

Pre V1R12 OSA inbound/outbound processing overview

- Queued Direct IO (QDIO) uses multiple write queues for outbound traffic separation
  - Outbound traffic is separated by priority (policy or WLM)
  - Multiple CPs can be used to manage the write queues

- QDIO uses only one read queue
  - All inbound traffic is received on the single read queue
  - Multiple CPs are used only when data is accumulating on the queue
    - During bursts of inbound data
  - Single process for initial interrupt and read buffer packaging
  - TCP/IP stack performs inbound data separation
    - Sysplex distributor traffic
    - Bulk inbound, such as FTP
    - IPv4/IPv6
    - EE traffic
    - Etc.
  - z/OS Communications Server is becoming the bottleneck as OSA nears 10GbE line speed
    - Inject latency
    - Increase processor utilization
    - Impede scalability
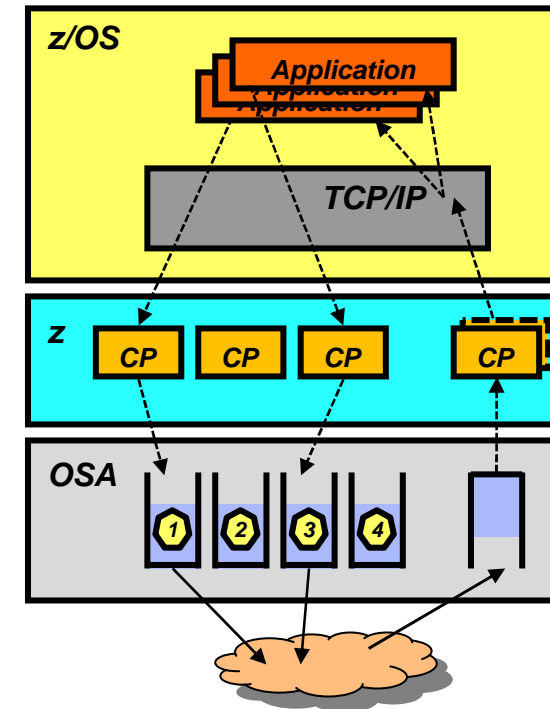
Performance problems observed for bulk inbound traffic:

- *Multiple processes run when data is accumulating on the read queue*
- *Inbound data for a single TCP connection can arrive at the TCP layer out of order*
- *TCP transmits a duplicate ACK every time it sees out of order data*
- *Sending side enters fast retransmit recovery*

OSA Inbound Workload Queueing (IWQ): Improved performance for mixed traffic patterns

- Allow inbound QDIO traffic separation by supporting multiple read queues
  - "Register" with OSA which traffic goes to which queue
 - OSA-Express Data Router function routes to the correct queue

- Each input queue can be serviced by a separate process
  - Primary input queue for general traffic
  - One or more ancillary input queues (AIQs) for specific traffic types
  - Dynamic LAN idle timer updated per queue

- Supported traffic types (z/OS V1R12)
  - Bulk data traffic queue
    - Serviced from a single process - eliminates the out of order delivery issue
  - Sysplex distributor traffic queue
    - SD traffic efficiently accelerated or presented to target application
  - All other traffic not backed up behind bulk data or SD traffic

- ***New for z/OS V1R13 – Unique inbound queue for Enterprise Extender traffic***
  - Improved performance for EE traffic
  - Supported on OSA-Express3 and new OSA-Express4S (CHPID type OSD or OSX)

- Significant performance improvement for mixed workloads/traffic patterns – for more details see:
  - http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS111-136/ENUS111-136.PDF
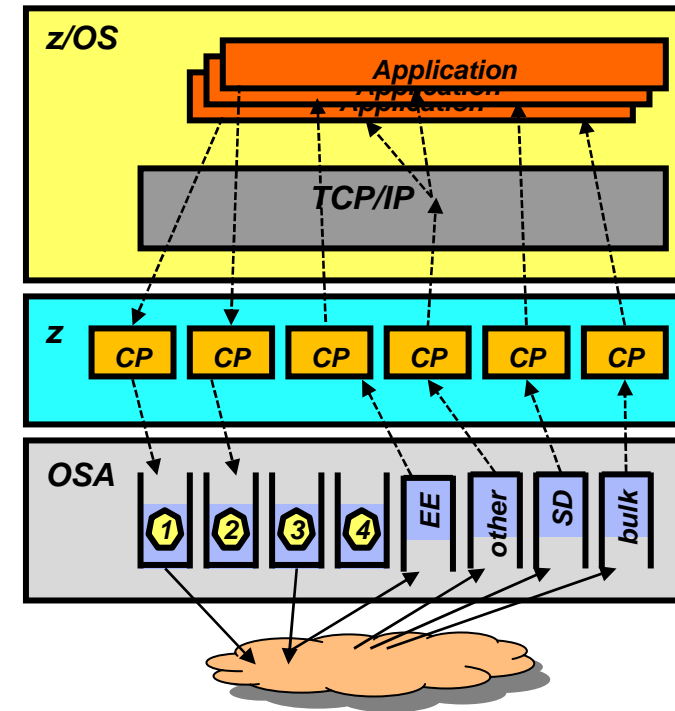
TCP/IP defines and assigns traffic to queues dynamically based on local IP address and port

Bulk traffic
- *Application sets send or receive buffer to at least 180K*
- *Registered per connection (5-tuple)*

SD traffic
- *Based on active VIPADISTRIBUTE definitions*
- *Registered on DVIPA address*

OSA-Express4S – Support of new features

z/OS 1   z/OS 2

IP-1   IP-2

OSA-A

NIC

LAN

NIC

OSA-B   IP-3
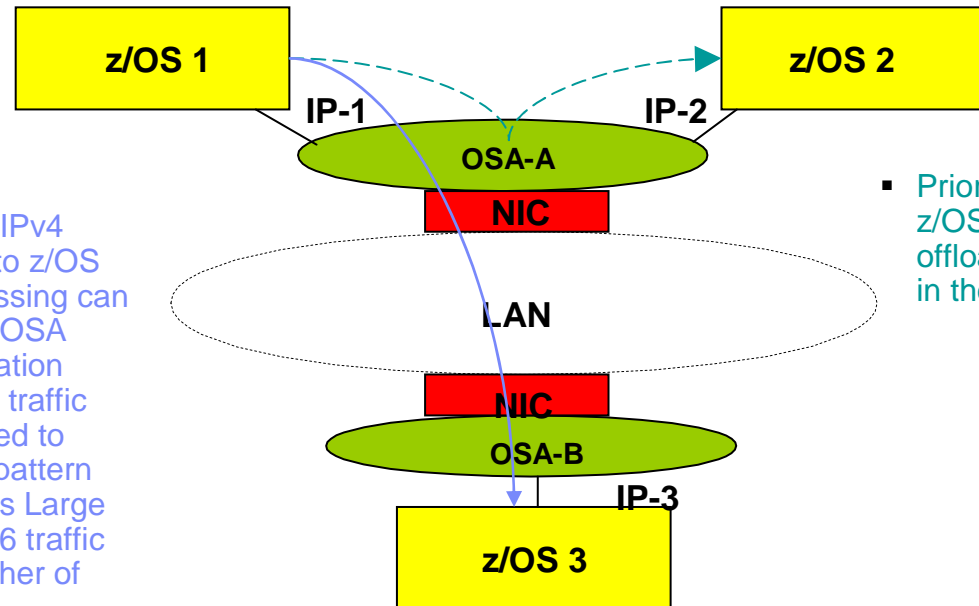
z/OS 3

- Prior to V1R13, for IPv4 traffic from z/OS 1 to z/OS 3, checksum processing can be offloaded to the OSA NIC. TCP segmentation processing for IPv4 traffic can also be offloaded to OSA on this traffic pattern (using OSA-Express Large Send support). IPv6 traffic could not exploit either of these features.

- Prior to V1R13, traffic from z/OS 1 to z/OS 2 could not exploit checksum offload. This processing would occur in the TCP/IP stack layer.

- New OSA-Express4S features exploited by z/OS V1R13 Communications Server
  - Additional checksum offload support
    - For IPv6 traffic
    - For LPAR-to-LPAR traffic (IPv4 and IPv6)
  - Large Send support for IPv6 traffic (aka "TCP segmentation offload")
    - Note: LPAR-to-LPAR traffic (IPv4 or IPv6) cannot exploit Large Send support.
- OSA-Express4S
  - New OSA-Express, smaller form factor, exploits new I/O drawer enabled for PCIe Gen2 (increased capacity, granularity and bandwidth)
- For more details, refer to IBM US Hardware Announcement 111-136, dated July 12, 2011

  http://www-01.ibm.com/common/ssi/rep_ca/6/897/ENUS111-136/ENUS111-136.PDF

z/OS virtual storage map

Private	**High User Region**	←	16EB
Shared	Default Shared Memory Addressing		512TB
Low User Private	**Low User Region**	←	2TB
			4GB
	Reserved		2GB

"The Bar" ······→

Extended Private	Extended LSAQ/SWA Sub-pools 229,230		
	Extended User Region	←	
Extended Common	Extended CSA (ECSA)		
	Extended PLPA/FLPA/MLPA		
	Extended SQA (ESQA)		
	Extended Nucleus		

"The Line" ······→ 16MB

Common	Nucleus		
	SQA		
	PLPA/FLPA/MLPA		
	CSA		
Private	LSAQ/SWA Sub-pools 229,230		
	User Region	←	
	System Region		24KB
Common	PSA		8KB
			0

z – 64-bit addressing

XA – 31-bit addressing

S/370 – 24-bit addressing

The user region may consist of up to four virtual storage areas:

1. *User region*
2. *Extended user region*
3. *Low user region*
4. *High user region*

Additional 64-bit exploitation

- Multiple trace buffers relocated to take advantage of 64 common storage
  - VTAM internal trace (VIT) is moved from ECSA to 64 bit common storage
    - Transparent to you if you use external VIT to obtain trace records
  - Multiple CTRACE components are moved from data-spaces to 64 bit common storage. The table below summarizes the changes.
    - These moves are transparent to you as long as you use the NMI interface to obtain trace data

CTRACE Component	Current location	z/OS V1R13 change	User
SYSTCPIP	TCPIPDS1 Dataspace	64 bit common	Stack
SYSTCPDA	TCPIPDS1 Dataspace	64 bit common	Stack (NMI)
SYSTCPIS	TCPIPDS1 Dataspace	64 bit common	Stack
SYSTCPCN	TCPIPDS1 Dataspace	64 bit common	Stack (NMI only)
SYSTCPSM	TCPIPDS1 Dataspace	64 bit common	Stack (NMI only)
SYSTCPRE	Private SP229	No Change	RESOLVER
SYSTCPRT	OMPROUTE Private storage	No Change	OMPROUTE
SYSTCPIK	IKE daemon Private storage	No Change	IKESMP
SYSTCPOT	TCPIPDS1 Dataspace	64 bit common	OSAENTA
SYSTCPNS	NSS daemon's private storage	No Change	Security Server

Miscellaneous TCP/IP performance improvements

- Communications Server Development strives to improve performance and throughput in every release by focusing on software pathlength

- Goal is 5% performance improvement in TCP/IP per release

- Examples of improvements planned for z/OS V1R13:
  - More use of compiler optimization
  - Modifying the layout of internal data structures for better caching
  - Separate IPv4 and IPv6 modules (reduces "IF IPv6" checking)
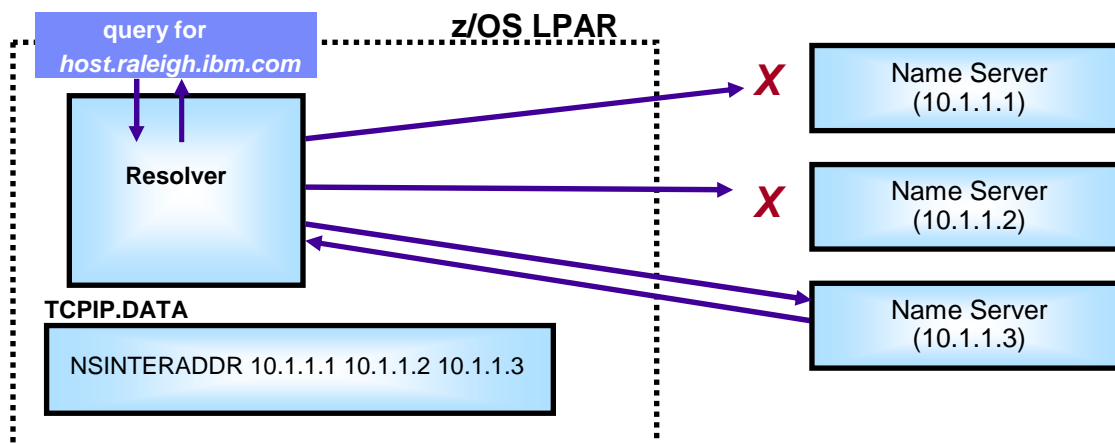 - Improved TN3270 performance for long data streams (> 90 bytes)

Throughput

Cost
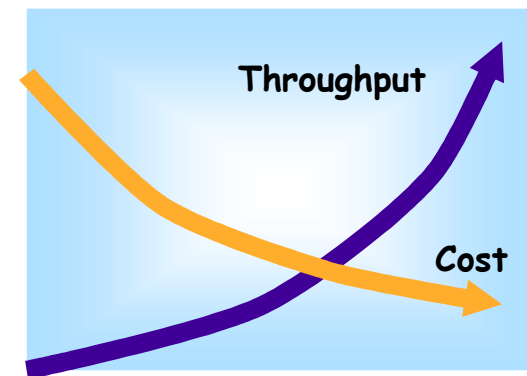
What's Coming in z/OS Communication Server

Availability

IBM®

Resolver autonomics for unresponsive name servers

- In z/OS V1R12, the resolver monitors name servers for responsiveness to queries
  - Network operator notification when a name server becomes unresponsive
  - Responsiveness is calculated on a sliding 5-minute window of statistics
  - Although the resolver detected the unresponsive name server, new queries were still sent to that name server

- In z/OS V1R13, the resolver may be configured to stop sending queries to unresponsive name servers
  - The resolver polls the unresponsive name server to detect when it becomes responsive again
  - Operator notified of condition using messages similar to those used in V1R12

query for
host.raleigh.ibm.com

z/OS LPAR

Resolver

TCPIP.DATA

NSINTERADDR 10.1.1.1 10.1.1.2 10.1.1.3

X Name Server (10.1.1.1)

X Name Server (10.1.1.2)

Name Server (10.1.1.3)

V1R12: Operator notification
V1R13: Autonomic quiescing

The autonomic quiescing function must be explicitly enabled in the resolver setup file.

- You specify what "unresponsive" means by coding a threshold failure rate in the resolver setup file
- A global TCPIP.DATA file is required.

Resolver autonomics for unresponsive name servers – main differences between z/OS V1R12 and z/OS V1R13 support

	Network Operator Notification (V1R12)	Autonomic Quiescing (V1R13)
Frequency of statistical calculation?	60 seconds	30 seconds
Number of intervals used to determine responsiveness?	5	1 or 2
Minimum sample size?	1 to become unresponsive, 0 to become responsive	10
Default setting?	Function active using 25% threshold value	Function not active (network operator notification remains the default)
Application queries sent to unresponsive name server?	Yes	No, with one exception
Vehicle for reporting statistics?	New EZZ9310I message issued every five minutes that a name server remains unresponsive	MODIFY RESOLVER command displays most recent failure rate for each name server
GLOBALTCPIPDATA required?	No	Yes

	Failure Rate (FR) < 1%	1% <= FR < Threshold	FR >= Threshold
Resolver polling name server?	No	Yes, to ensure proper sample set is available at next checkpoint	Yes, to determine whether name server is now responsive
Application queries forwarded to name server?	Yes	Yes	No, unless all name servers are considered to be unresponsive

Health Checker for the autonomic quiescing function

- Three new checks were added to Health Checker for the autonomic quiescing function:

  - **CSRES_AUTOQ_GLOBALTCPIPDATA**
    - Checks that you have coded the GLOBALTCPIPDATA setup statement if AUTOQUIESCE is coded on the UNRESPONSIVETHRESHOLD setup statement

  - **CSRES_AUTOQ_TIMEOUT**
    - Checks, by default, if you have specified a value greater than five (seconds) for RESOLVERTIMEOUT when autonomic quiescing is enabled
    - You can change the check to have a different value than five seconds if your installation uses a larger timeout value

  - **CSRES_AUTOQ_RESOLVEVIA**
    - Checks if you have specified RESOLVEVIA TCP when autonomic quiescing is enabled

- These checks are performed when the resolver is started and when a MODIFY RESOLVER,REFRESH command is issued
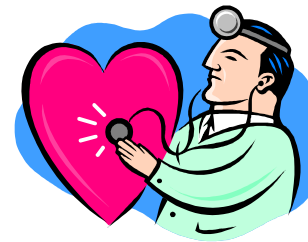
Migration note: Resolver requires an OMVS segment in z/OS V1R13

- Starting in z/OS V1R13, the system resolver uses z/OS Unix System Services within the Resolver address space for monitoring unresponsive name servers and providing NMI information. The use of z/OS Unix System Services cause the resolver to be an OMVS process, which additionally requires that the resolver have a RACF user identity to provide access to z/OS Unix, either explicitly or through the default userid. The resolver uses these z/OS Unix System Services even if you do **not use** the monitoring or NMI functions.

  You must take action if you do not have a user ID defined for resolver, otherwise resolver initialization will fail.
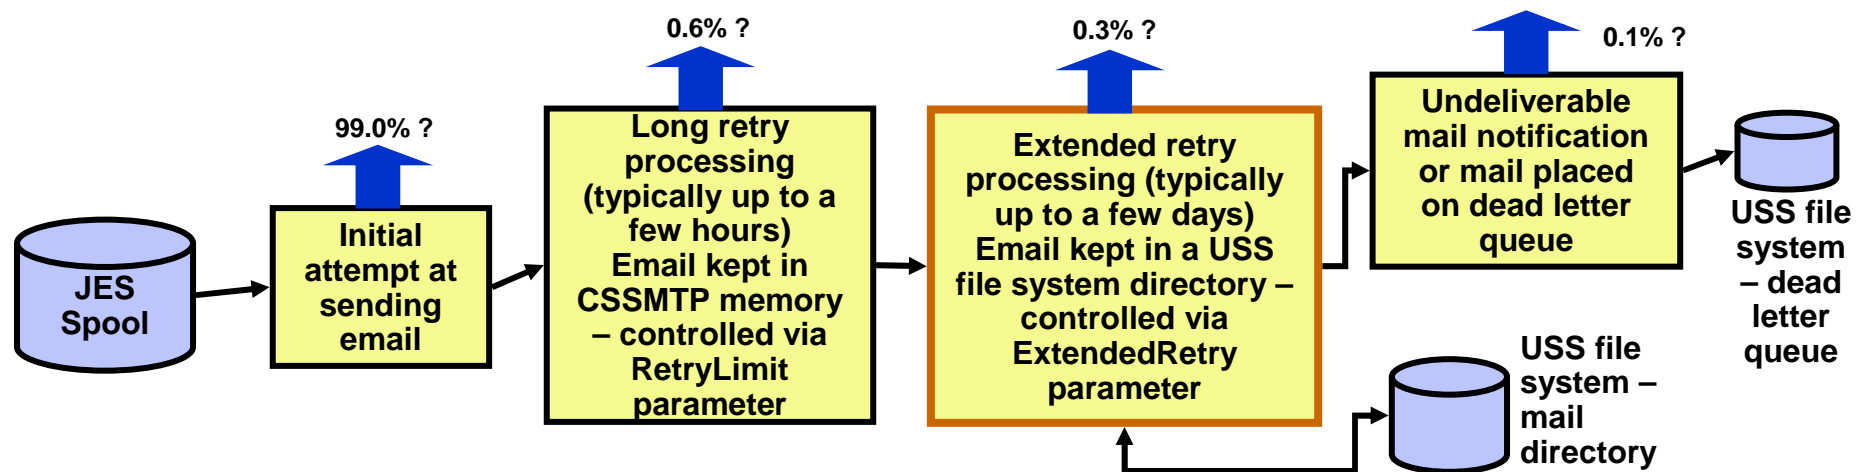
- If you are migrating to z/OS V1R13 Communications Server and you already have a user ID for the system resolver procedure, no action is required if you meet either of these conditions:
  - You explicitly defined an OMVS segment for the user ID.
  - An OMVS segment was created through the RACF automated assignment of unique z/OS UNIX identities support or through default OMVS segment support.

- If you are migrating to z/OS V1R13 Communications Server, and you *do not have* a user ID defined for the system resolver that has an associated OMVS segment, *you must take action*. If you do not take action, *the resolver address space initialization will fail and the initialization of all TCP/IP stacks will be delayed.*

  Steps to take:

  1. If you already have a resolver user ID but it does not have an OMVS segment, you must define an OMVS segment for the resolver user ID.
  2. If you do not have a resolver user ID, you must create one that includes an OMVS segment.

CSSMTP enhanced send error recovery

- CSSMTP sends batch email to the internet from z/OS JES spool files

- If target relays fail to acknowledge mail, will retry for configured interval up to 5 days (default 5 minutes) then drop the message, and return undeliverable notice, however:
  – Spool files cannot be deleted until all messages in the spool file are delivered
    • A spool file could contain thousands of messages but only a few are being retried
  – Messages being retried are retained in CSSMTP memory

- z/OS V1R13 provides file system storage of messages being retried for an extended interval (beyond initial retry limit), so that JES spool files and CSSMTP memory can be released.
  – Will continue to retry from memory and spool until initial retry limit reached
  – New parameter to indicate how long beyond existing interval to retry from file system

0.6% ? 0.3% ? 0.1% ?

99.0% ?

JES Spool → **Initial attempt at sending email** → **Long retry processing (typically up to a few hours) Email kept in CSSMTP memory – controlled via RetryLimit parameter** → **Extended retry processing (typically up to a few days) Email kept in a USS file system directory – controlled via ExtendedRetry parameter** → **Undeliverable mail notification or mail placed on dead letter queue** → **USS file system – dead letter queue**
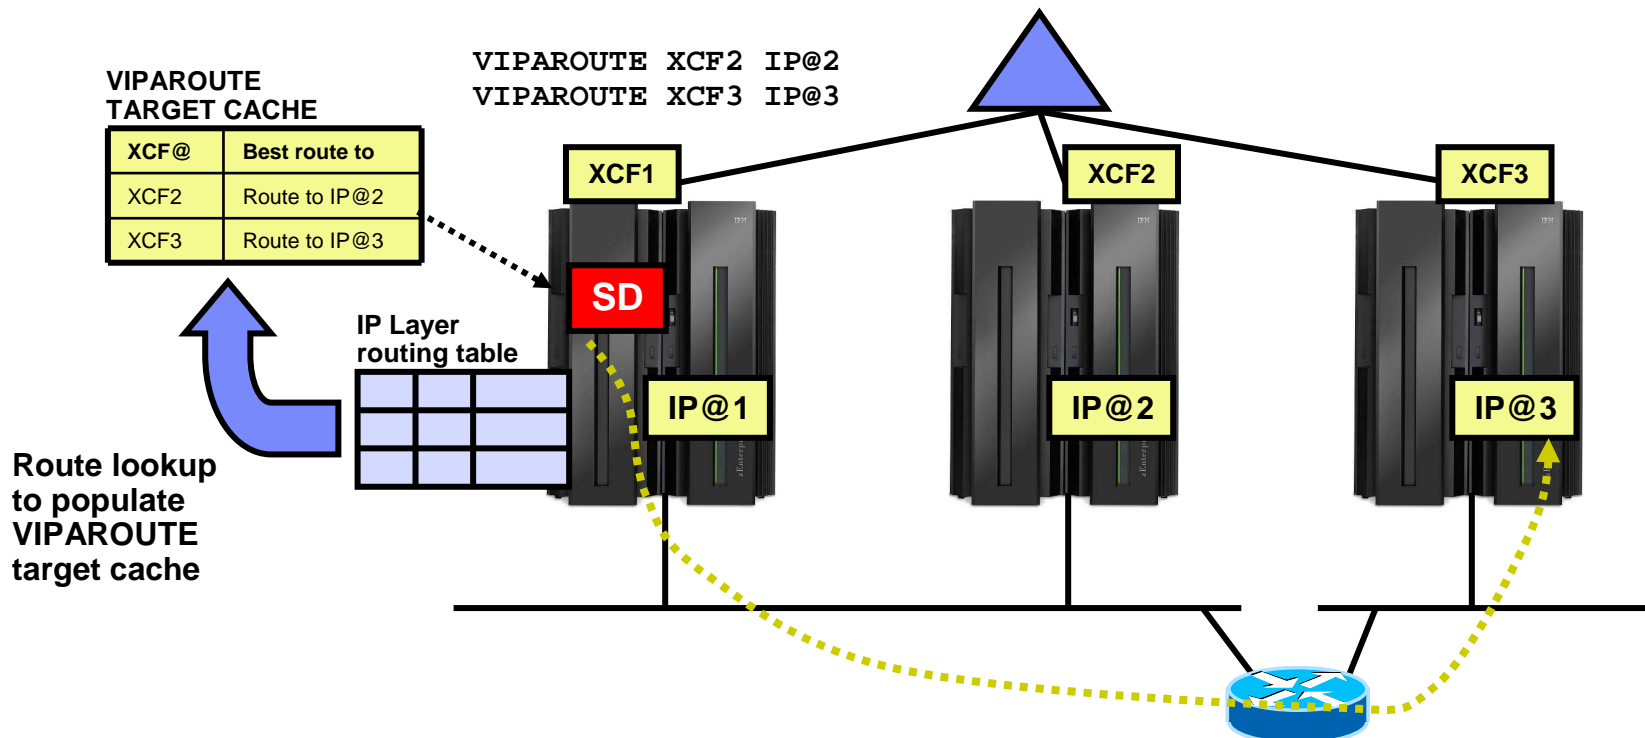
USS file system – mail directory

CSSMTP translate code page support enhanced

- The '@' character has special meaning in SMTP mail messages
    - MUST be translated correctly!
- CSSMTP currently supports only a fixed set of single byte code pages.
    - Some installations use a code page that is not supported by the CSSMTP TRANSLATE statement

Code Page	CCSID	@	§
IBM-1047 (EBCDIC)	1047	0x7C	0xB5
ISO-8859-1 (ASCII)	819	0x40	0xA7
IBM-273	273	0xB5	0x7C

- Allow the TRANSLATE statement to define additional code pages
    - Allow the specification of a code page by its CCSID.
        - Currently code pages must be a character string "IBM-XXXX", where XXXX is a subset of possible code pages.
    - Expand the list of supported code pages
        - Allow a user defined code page to be used
- The code page must be an EBCDIC code page.
    - The target ASCII code page is always ISO-8859-1 (or in reality, the US-ASCII subset of that code page)
- The code page must support roundtrip translations between the IBM-1047(EBCDIC) and the ISO-8859-1 (ASCII) code pages.
- The carriage return and line feed characters (CRLF) used to end the lines of commands and mail messages must translate properly to ISO-8859-1 (x'0D0A')
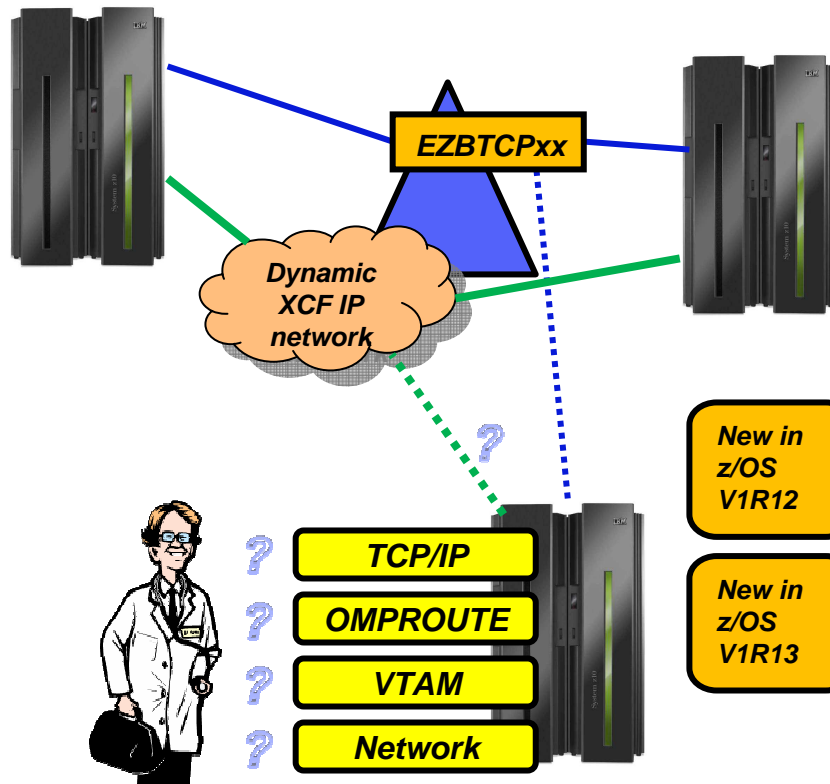
VIPAROUTE target cache update during initialization

- When using VIPAROUTE, a VIPAROUTE target cache is used to minimize the time it takes to route a Sysplex Distributor packet

- The target cache is updated every 60 seconds, which in some cases have caused delays during a primary stack's take-back of a distributed DVIPA

- z/OS V1R13 shortens the interval for VIPAROUTE route lookups in situations where the stack joins a Sysplex, or OMPROUTE is restarted
  - Will now start with 5 seconds, and gradually increase to 60 seconds

```
VIPAROUTE XCF2 IP@2
VIPAROUTE XCF3 IP@3
```

VIPAROUTE TARGET CACHE

XCF@	Best route to
XCF2	Route to IP@2
XCF3	Route to IP@3

IP Layer routing table

Route lookup to populate VIPAROUTE target cache

XCF1   XCF2   XCF3

SD

IP@1   IP@2   IP@3

Sysplex autonomics extended with CSM storage constrained monitoring

- **Monitoring:**
  - Monitor CS health indicators
    - Storage usage critical condition (>90%) - CSM, TCPIP Private & ECSA
      - For more than TIMERSECS seconds
  - Monitor dependent networking functions
    - OMPROUTE availability
    - VTAM availability
    - XCF links available
  - Monitor for abends in Sysplex-related stack components
    - Selected internal components that are vital to Sysplex processing
      - Does not include "all" components
  - Selected network interface availability and routing
  - Monitor for repetitive internal abends in non-Sysplex related stack components
    - 5 times in less than 1 minute
  - *Detect when CSM FIXED or CSM ECSA has been constrained (>80% utilization) for multiple monitoring intervals*
    - *For 3 times the TIMERSECS value*

- **Actions:**
  - Remove the stack from the IP Sysplex (manual or automatic)
    - Retain the current Sysplex configuration data in an inactive state when a stack leaves the Sysplex
  - Reactivate the currently inactive Sysplex configuration when a stack rejoins the Sysplex (manual or automatic)
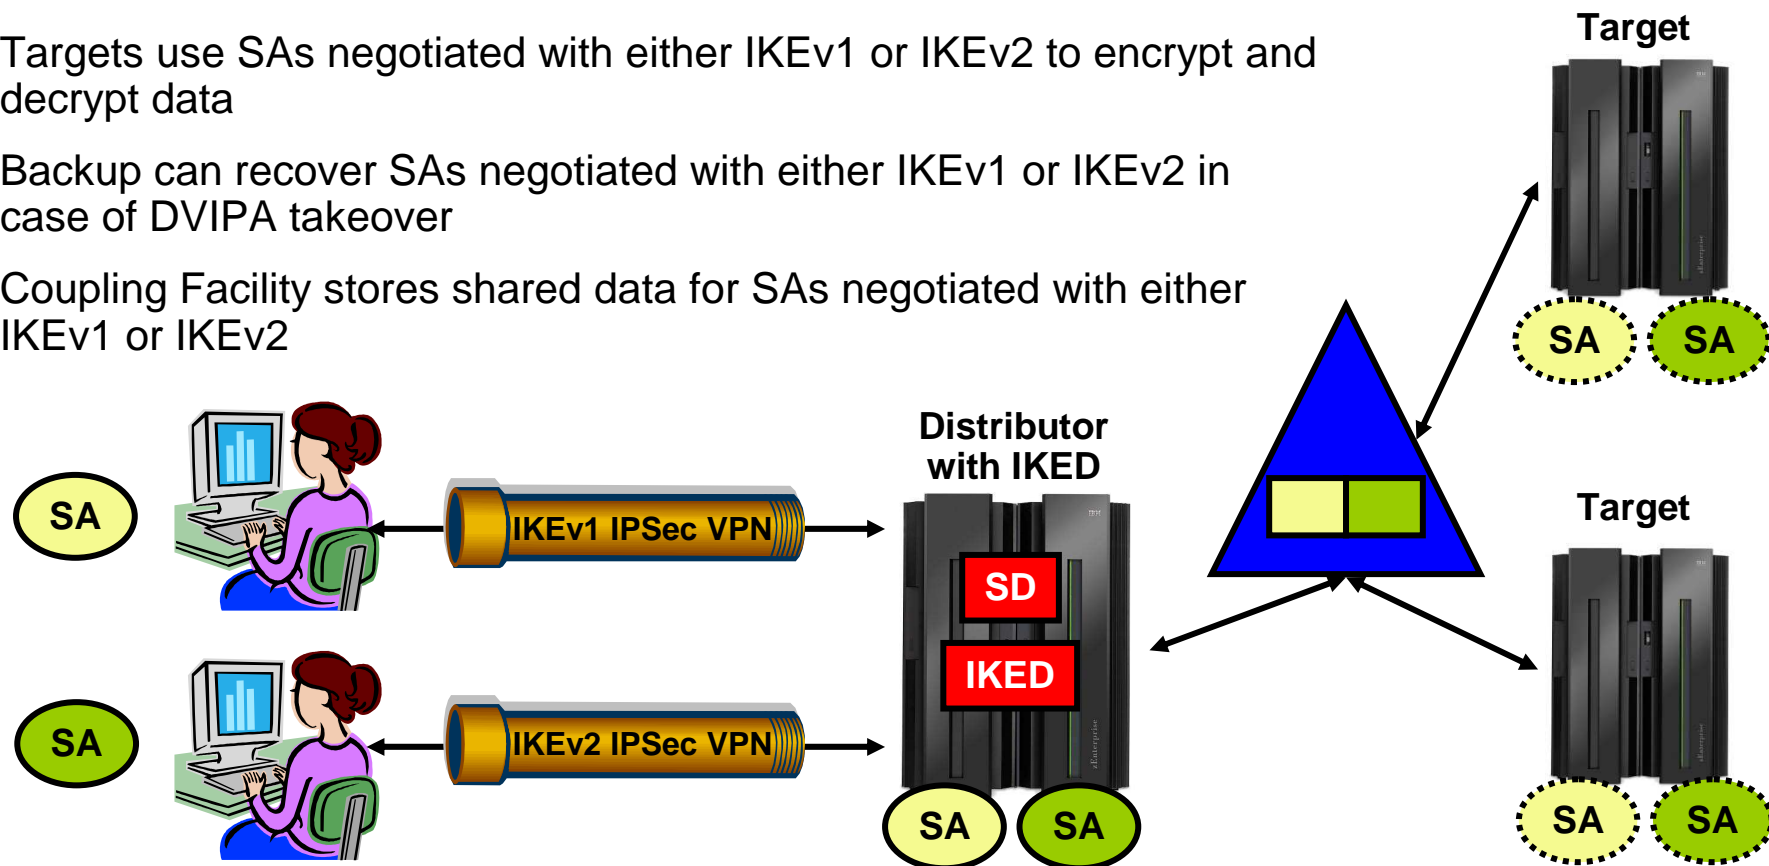
New in z/OS V1R12

New in z/OS V1R13

EZBTCPxx

Dynamic XCF IP network

TCP/IP

OMPROUTE

VTAM

Network

Sick? Better remove myself from the IP Sysplex!

Feeling better? Maybe it's time to rejoin the IP Sysplex

Sysplex-wide Security Associations (SWSA) for IKEv2

- Sysplex Distributor
  - Negotiates SAs with remote Client using the Internet Key Exchange protocol, IKE version 1 or IKE version 2
  - Sends copies of SAs (shadows) to Targets for VPNs negotiated with either version of IKE

- Targets use SAs negotiated with either IKEv1 or IKEv2 to encrypt and decrypt data

- Backup can recover SAs negotiated with either IKEv1 or IKEv2 in case of DVIPA takeover

- Coupling Facility stores shared data for SAs negotiated with either IKEv1 or IKEv2

Target

SA **SA**

SA

Distributor with IKED

IKEv1 IPSec VPN

SD

IKED

Target

SA

IKEv2 IPSec VPN

SA **SA**

SA **SA**

What's Coming in z/OS Communication Server

Application / Middleware / Workload enablement

IBM®

z/OS FTP's journey to extended address volumes

- **z/OS V1R10**
  - DFSMS added support for VSAM data sets in Extended Addressing Space (EAS)
  - FTP doesn't support VSAM data set, so no impact

- **z/OS V1R11**
  - DFSMS added support for physical sequential extended format datasets to reside in the EAS
  - FTP added support for reading/writing to/from existing EAS data sets, but not creating them (toleration mode)
    - FTP qdisk option for SITE/LOCSITE output format will change

- **z/OS V1R13**
  - Full FTP support for extended address volumes without requiring use of SMS data class
    - PDS, PDSE, Physical Sequential Basic and Large
    - Prior to V1R13, can use EAV for SMS managed datasets only

- *VSAM Data Sets*
- *21-cylinder allocation units*
- *Extended Format Data Sets*
- *etc.*

EAS

All Data Set Types
3390-A
EAV

2314-1	3330-1	3350	3390-3	3390-9	3390-9	3390-9		
29 MB	101MB	317MB	3GB	9GB	27GB	54GB	223GB*	Architectural Limit:
~300 cyl	404 cyl	555 cyl	3,339 cyl	10,017 cyl	32,760 cyl	65,520 cyl	262,668 cyl	100s of TB*

Physical sequential data set formats

- Large format physical sequential datasets
  - Can have more than 64K tracks per volume
    - But don't have to. Can have fewer than 64K tracks per volume and still be large format

- DFSMS has offered large format physical sequential datasets since z/OS V1R7
  - Access methods supported: BSAM, QSAM, EXCP
  - Language Environment (LE) Runtime Library Large format dataset support completed in z/OS V1R13 now enables z/OS FTP support for these datasets

Sequential data sets	BASIC FORMAT	LARGE FORMAT	EXTENDED FORMAT
DSNTYPE	BASIC	LARGE	EXTPREF, EXTREQ
Max tracks per volume	65,535	16,777,215	2,147,483,647 (theoretical)
Max extents per volume	16	16	123
Why choose this format?	Maximum compatibility	Can be larger than basic	Can be larger than basic. Can be compressed or striped.

FTP support for large format physical sequential datasets

- FTP support means
  - Send from large format data sets
  - Allocate and store into large format data sets
  - configuration options to explicitly create large format data sets
  - In block mode, restart support for failed transfers to and from large format data sets

New client and server FTP.DATA keyword controls how FTP allocates physical sequential datasets. Displayable using SITE and LOCSITE

Value of SYSTEM uses SMS data class or system default value

```
    .-DSNTYPE SYSTEM----------.
>>-+-----------------------+----------------------------------><
    '-DSNTYPE----+-BASIC------+
                 +-LARGE------+
                 '-SYSTEM-----'
```

- Support also added for transfers to and from z/OS UNIX files larger than two gigabytes
  - Send from and allocate and store into z/OS UNIX files larger than two gigabytes
  - No additional FTP configuration needed, "just works"
  - Can restart failed transfers

Command to display all TN3270 servers

- A new D TCPIP,TELNET console command added to display the list of TN3270E servers that are or have been active on the system

- This can be a starting point for performing automation on TN3270E servers
  - After all, you have to know what's there before you can operate on it!

- Display example:

```
D TCPIP,TELNET
EZAOP60I TELNET STATUS REPORT
TELNET NAME   VERSION    STATUS
-----------   --------   ---------------------------------
TELNET        CS V1R13   ACTIVE
TELNET5       CS V1R13   INACTIVE (STOP CMD)
TELNET4       CS V1R13   INACTIVE (STOP CMD)
*** END TELNET STATUS REPORT ***
```

Retrieving System Resolver configuration data through NMI

- You can use the MODIFY RESOLVER,DISPLAY command to display the contents of the resolver setup file.

- However, there is no resolver command available to display the contents of the global TCPIP.DATA settings.
  - Currently, you must collect Trace Resolver output to see the global TCPIP.DATA settings dynamically, or alternatively, you can dump the resolver and examine the internal control blocks.

- The new resolver function in z/OS V1R13 to monitor and quiesce unresponsive name servers, depend on a global TCPIP.DATA
  - Only name servers specified in this global TCPIP.DATA will be monitored

- z/OS V1R13 Communications Server implements a resolver callable NMI (EZBREIFR)
  - Provides a high-speed, low-overhead callable programming interface for network management applications to access data related to the resolver
  - One request type:
    - GetResolverConfig
  - Returns:
    - The contents of the resolver configuration file and
    - The contents of the global TCPIPDATA file, if it's in use

TMI copy buffer interface improvement

- The TMI copy buffer interfaces currently require the caller to be APF authorized
  - EZBTMIC1, EZBTMIC4, and TMI_Copybuffer()
  - Causes problems for applications that want to fork().
    - APF authorization is not inherited
  - APF authorization gives application broad range of authority

- In z/OS V1R13, network management applications that use these TMI copy buffer interfaces will no longer have to be APF authorized
  - As an alternative to APF Authorization, the user ID that the applications are running under can be authorized to the appropriate SAF resource:
  - EZB.NETMGMT.systemname.tcpprocname.SYSTCPxx
    - SYSTCPxx can be:
      - SYSTCPDA, which provides access to packet trace data.
      - SYSTCPCN, which provides access to ongoing information about opening and closing TCP connections.
      - SYSTCPSM, which provides access to ongoing information about FTP and Telnet activity.
      - SYSTCPOT, which provides access to OSA Network Traffic Analyzer data.

PORTRANGE wildcard option

- A wildcard option was allowed for the jobname specified on a PORT statement since z/OS V1R5
    - PORT statement allows jobname to end in asterisk (*)
    - Characters before the asterisk define a prefix
    - Any applications with a jobname matching the prefix can access the specified port

- Similar wildcard support was not provided for the PORTRANGE statement

- PORTRANGE statement in z/OS V1R13 allows a wildcard jobname
    - Job name can end in an *
    - Characters before the * define a prefix
    - Only applications with job names that match the prefix have access to the specified port range

- As a result of this, the GetProfile request of the EZBNMIFR API can now include a wildcard value in the NMTP_PORTJobName field for entries that represent a PORTRANGE statement.

```
PORTRANGE
    111    1    UDP    PORTMAP
    111    1    TCP    PORTMAP
    500    5    UDP    USER1
    500    5    TCP    USER2
    700    4    TCP    ABCD*
```

What's Coming in z/OS Communication Server

Enterprise Extender / SNA

IBM®

Enterprise Extender firewall-friendly connectivity test

- DISPLAY EEDIAG,TEST=YES provides information about an Enterprise Extender partner and all the routers in between. But if a firewall in between is blocking ICMP messages, there can be a long delay before getting results
  - Because of timeouts waiting for ICMP messages that never come
  - Delay is (Number of router hops past the firewall) x (9 seconds)

- LIST=SUMMARY will provide quick test of partner reach-ability. Probe is sent to partner with TTL=255 so it doesn't probe any intermediate hops
  - No intermediate hop information provided
  - Hop count determination omitted

```
D NET,EEDIAG,TEST=YES,IPADDR=(9.67.1.1,9.67.1.5),LIST=SUMMARY
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2067I EEDIAG DISPLAY ISSUED ON 08/29/05 AT 15:41:22
"""""""""""""""""""""""""""""""""""""""""""""""""""""""""""
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.5
IST924I -------------------------------------------------------------
IST2133I INTFNAME: LTRLE1A                   INTFTYPE: MPCPTP
IST2134I    CONNECTIVITY SUCCESSFUL                        PORT: 12000
IST2137I      *NA 9.67.1.5                   RTT:      6
IST2134I    CONNECTIVITY SUCCESSFUL                        PORT: 12001
IST2137I      *NA 9.67.1.5                   RTT:      6
IST2134I    CONNECTIVITY SUCCESSFUL                        PORT: 12002
IST2137I      *NA 9.67.1.5                   RTT:      6
IST2134I    CONNECTIVITY SUCCESSFUL                        PORT: 12003
IST2137I      *NA 9.67.1.5                   RTT:      6
IST2134I    CONNECTIVITY SUCCESSFUL                        PORT: 12004
IST2137I      *NA 9.67.1.5                   RTT:      7
IST924I -------------------------------------------------------------
IST2139I CONNECTIVITY TEST RESULTS DISPLAYED FOR 1 OF 1 ROUTES
IST314I END
```
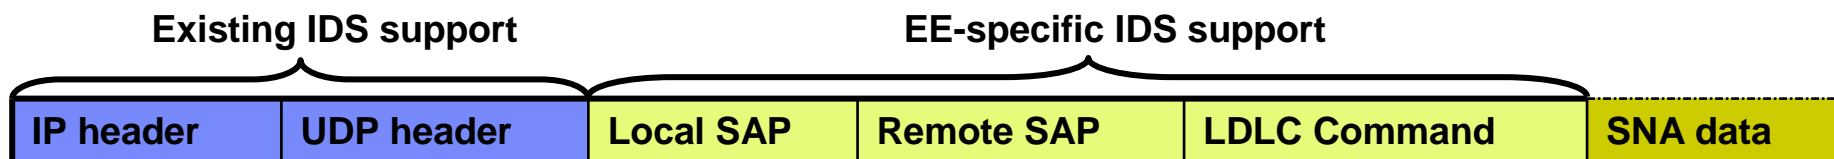
Improved APPN routing resilience

- APPN session routes are selected by APPN Topology and Routing Services (TRS)
  - Directed search routes used to locate resources
  - Routes for LU-LU sessions

- Optimal route determined by specified criteria
  - Line speed, Cost of data transmitted, and Security

- TRS builds complex routing trees to determine best path
  - Tree records represent nodes along preferred route
  - Transmission groups (TGs) between nodes are associated with tree records

- Recursive abends can occur if a pointer in a routing tree is compromised
  - Every time session route is requested using that tree, abends occur and session fails
  - Possible session setup failures with sense codes
    - 80140001, 087F0001, 087D0001, 08400007
  - Has so far required VTAM be restarted to recover

- Recovery routine for route selection abends is improved in z/OS V1R13 to prevent recursive abends – by removing the entire storage area that contains the routing tree and dynamically rebuild it
  - Existing sessions are not affected
  - Can temporarily impact VTAM performance in large APPN networks as routing tree is being rebuilt
  - Topology database is unaffected, so no network impact due to TDU flows

- For the rare case where the routing tree is corrupted, but no abends occur, a new command can be used to perform the same process as described above:
  - MODIFY TOPO,FUNCTION=CLRTREES

Page 54
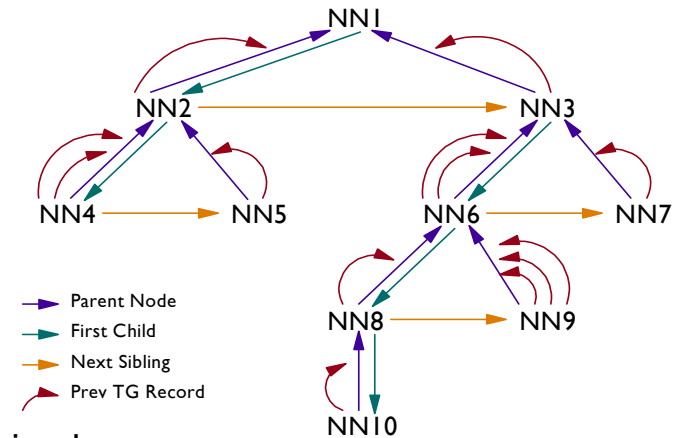
Intrusion Detection Services for Enterprise Extender traffic

- Implements four new IDS attack types:
  - EE Malformed Packet (Discard / Notify)
    - Checking for inbound LDLC packets with invalid lengths.
  - EE LDLC Check (Discard / Notify)
    - Checks that inbound LDLC control packets are received on the signaling port (12000)
  - EE Port Check (Discard / Notify)
    - Checks that the source port and destination port match in inbound EE packets.
  - EE XID Flood (Notify)
    - Checks if a threshold is met for inbound XIDs within one minute.

- Allow exclusion list for each attack type

- Notifications
  - System console message
  - Syslogd message
  - IDS Trace(SYSTCPIS)- NO IDS packet tracing done for EE XID flood
  - Statistics gathering

Existing IDS support   **EE-specific IDS support**

IP header	UDP header	Local SAP	Remote SAP	LDLC Command	SNA data

z/OS CS Configuration Assistant support for the new Intrusion Detection Services for Enterprise Extender

Statement of Direction:

z/OS V1.13 is planned to be the last release in which the BIND 9.2.0 function will be available. Customers who currently use or plan to use the z/OS BIND 9.2.0 function as a caching-only name server should use the resolver function, which became generally available in z/OS V1.11, to cache Domain Name Server (DNS) responses. Customers who currently use or plan to use the z/OS BIND 9.2.0 function as a primary or secondary authoritative name server should investigate using BIND on Linux for System z or BIND on an IBM blade in an IBM zEnterprise BladeCenter® Extension (zBX).
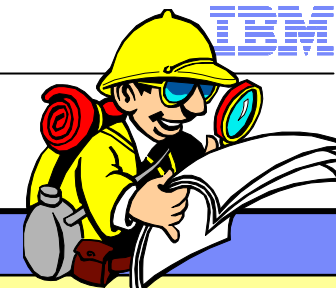
For more information

URL	Content
http://www.twitter.com/IBM_Commserver	IBM z/OS Communications Server Twitter Feed
http://www.facebook.com/IBMCommserver	IBM z/OS Communications Server Facebook Page
https://www.ibm.com/developerworks/mydeveloperworks/blogs/IBMCommserver/?lang=en	IBM z/OS Communications Server Blog
http://www.ibm.com/systems/z/	IBM System z in general
http://www.ibm.com/systems/z/hardware/networking/	IBM Mainframe System z networking
http://www.ibm.com/software/network/commserver/	IBM Software Communications Server products
http://www.ibm.com/software/network/commserver/zos/	IBM z/OS Communications Server
http://www.redbooks.ibm.com	ITSO Redbooks
http://www.ibm.com/software/network/commserver/zos/support/	IBM z/OS Communications Server technical Support – including TechNotes from service
http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/TechDocs	Technical support documentation from Washington Systems Center (techdocs, flashes, presentations, white papers, etc.)
http://www.rfc-editor.org/rfcsearch.html	Request For Comments (RFC)
http://www.ibm.com/systems/z/os/zos/bkserv/	IBM z/OS Internet library – PDF files of all z/OS manuals including Communications Server
http://www.ibm.com/developerworks/rfe/?PROD_ID=498	RFE Community for z/OS Communications Server
https://www.ibm.com/developerworks/rfe/execute?use_case=tutorials	RFE Community Tutorials

For pleasant reading ….