


# z/OS Communications Server Technical Update

Gus Kassimis - [kassimis@us.ibm.com](mailto:kassimis@us.ibm.com)  
Sam Reynolds - [samr@us.ibm.com](mailto:samr@us.ibm.com)  
IBM Enterprise Networking Solutions  
Raleigh, NC, USA

Sessions: 11327 and 11328  
Monday, August 6 – 9:30 AM to 10:30 AM &  
11:00 AM to 12:00 PM



## z/OS Communications Server Technical Update

<b>Session number:</b>	11327 and 11328
<b>Date and time:</b>	Monday, August 6, 2012 - 9:30 AM – 10:30 and 11:00 AM – 12:00 PM
<b>Location:</b>	Platinum Ballroom Salon 9
<b>Program:</b>	Communications Infrastructure
<b>Project:</b>	Communications Server
<b>Track:</b>	Network Support and Management
<b>Classification:</b>	Technical
<b>Speaker:</b>	Gus Kassimis, IBM Sam Reynolds, IBM
<b>Abstract:</b>	<p>z/OS Communications Server combines TCP/IP and VTAM support to better address the needs of today's complex networks. This two-part session provides an overview of features in the most recent releases of z/OS Communications Server, and focuses on a more detailed exploration of selected functions, such as:</p> <ul style="list-style-type: none"> <li>• IEDN-enabled Hipersockets</li> <li>• Resolver autonomics</li> <li>• EE and TN3270 enhancements</li> <li>• ... and more!</li> </ul>

## 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® | • GDDM®                                     | • Language Environment® | • Rational Suite® | • zEnterprise    |
| • AIIX®                             | • GDPS®                                     | • MQSeries®             | • Rational®       | • zSeries®       |
| • alphaWorks®                       | • Geographically Dispersed Parallel Sysplex | • MVS                   | • Redbooks        | • z/Architecture |
| • AnyNet®                           | • HiperSockets                              | • NetView®              | • Redbooks (logo) | • z/OS®          |
| • AS/400®                           | • HPR Channel Connectivity                  | • OMEGAMON®             | • Sysplex Timer®  | • z/VM®          |
| • BladeCenter®                      | • HyperSwap                                 | • Open Power            | • System i5       | • z/VSE          |
| • Candle®                           | • i5/OS (logo)                              | • OpenPower             | • System p5       |                  |
| • CICS®                             | • i5/OS®                                    | • Operating System/2®   | • System x®       |                  |
| • DataPower®                        | • IBM eServer                               | • Operating System/400® | • System z®       |                  |
| • DB2 Connect                       | • IBM (logo)®                               | • OS/2®                 | • System z9®      |                  |
| • DB2®                              | • IBM®                                      | • OS/390®               | • System z10      |                  |
| • DRDA®                             | • IBM zEnterprise™ System                   | • OS/400®               | • Tivoli (logo)®  |                  |
| • e-business on demand®             | • IMS                                       | • Parallel Sysplex®     | • Tivoli®         |                  |
| • e-business (logo)                 | • InfiniBand®                               | • POWER®                | • VTAM®           |                  |
| • e business (logo)®                | • IP PrintWay                               | • POWER7®               | • WebSphere®      |                  |
| • ESCON®                            | • IPDS                                      | • PowerVM               | • xSeries®        |                  |
| • FICON®                            | • iSeries                                   | • PR/SM                 | • z9®             |                  |
|                                     | • LANDP®                                    | • pSeries®              | • z10 BC          |                  |
|                                     |   | • RACF®                 | • z10 EC          |                  |
- \* All other products may be trademarks or registered trademarks of their respective companies.

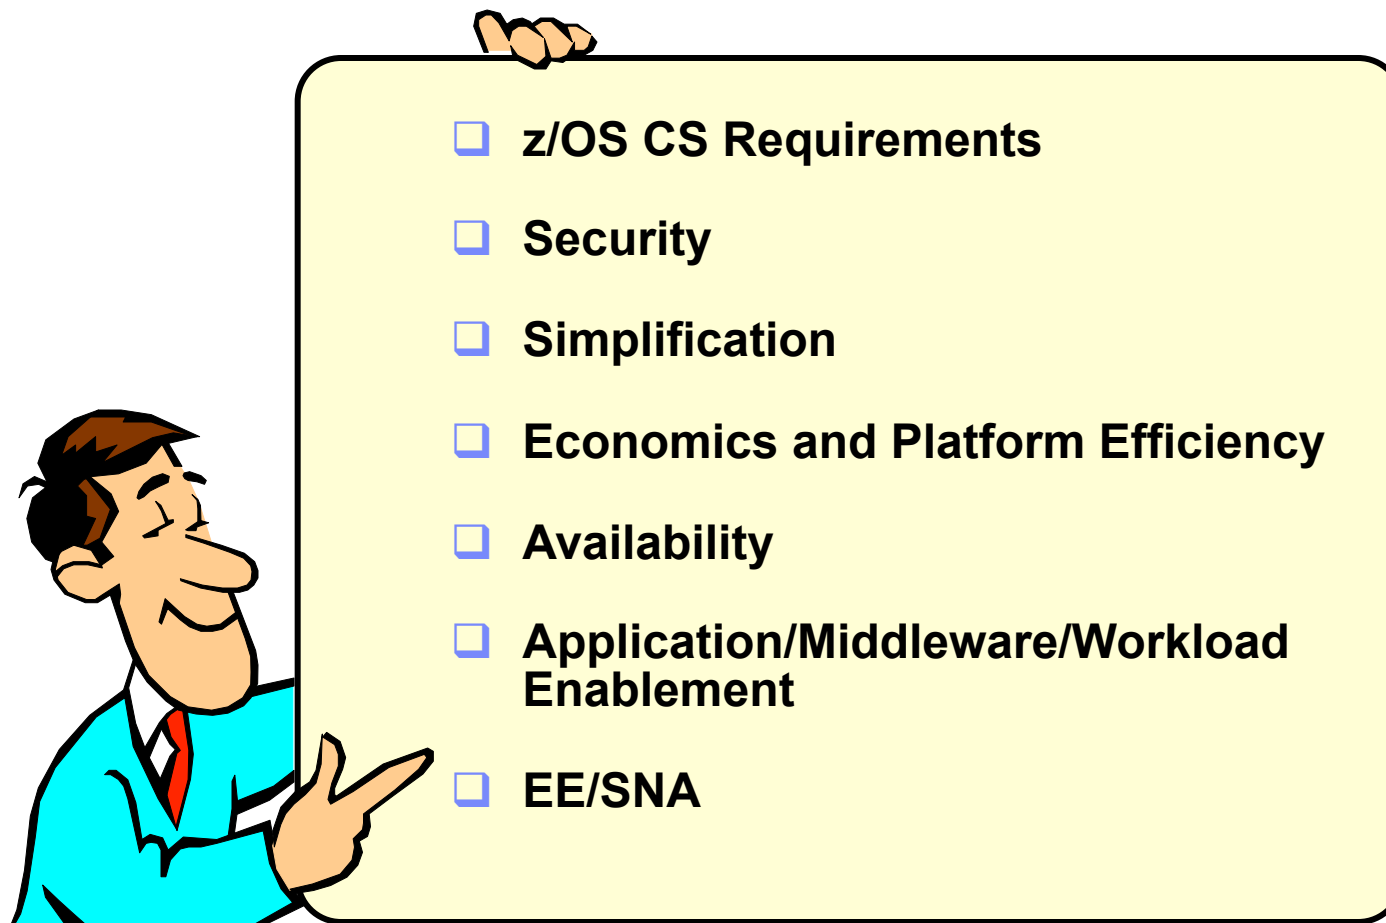
The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries or both:

- 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](http://www.ibm.com/legal/us) for further legal information.

## Agenda



***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 2013-2015?

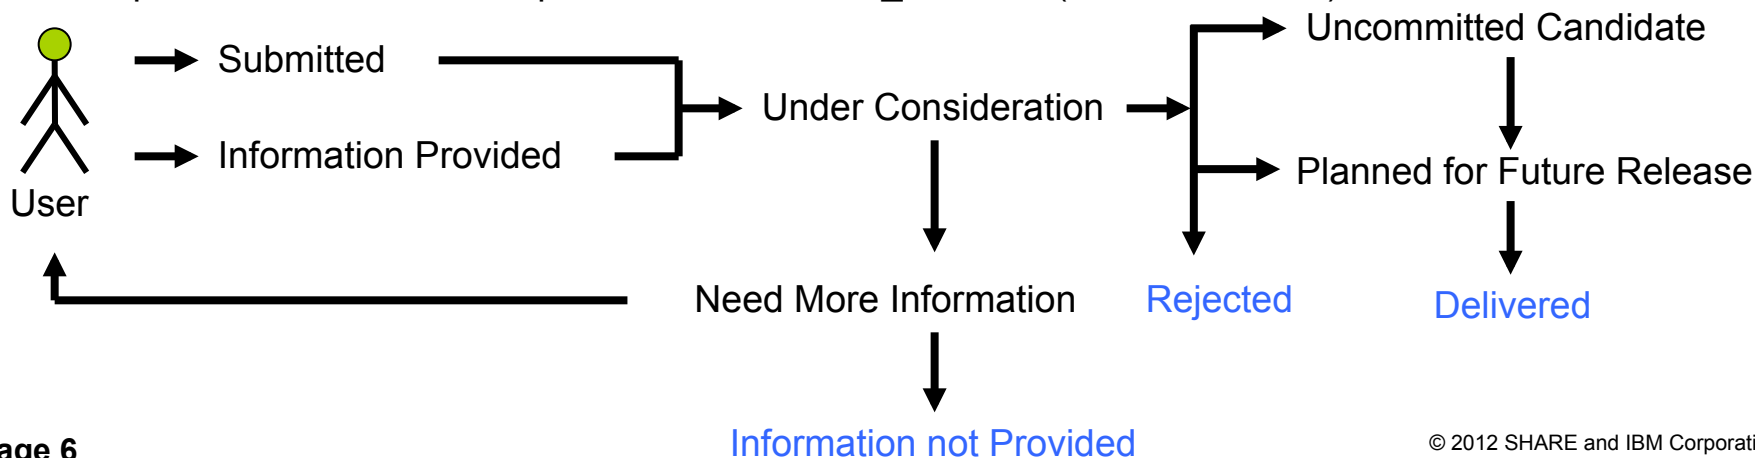


- **System z technology is expected to continue to evolve**
  - Networking software need to support new technologies such as zEnterprise
- **Access to System z system-level skills will continue to be an issue**
  - Retiring existing people, who grew up with system z
  - New people becoming responsible 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, TAMI 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](http://www.ibm.com/developerworks/rfe/?PROD_ID=498) (QR code above)



# RFE: New Requirements Process for z/OS Communications Server ...

IBM
English
Sign in (or register)

developerWorks
Technical topics
Evaluation software
Community
Events

Search developerWorks

developerWorks > RFE Community > WebSphere >

## WebSphere RFE Community

Welcome WebSphere users! Here you have an opportunity to collaborate directly with the WebSphere product development teams and other product users.

- [→ Search for RFEs](#) (view, comment, vote, and watch)
- [→ Submit RFEs](#)
- [→ Track your RFEs](#) (My RFEs)

Customize this page for your favorite product:

→

**Welcome z/OS Communications Server users**

The RFE Community does not currently contain any RFEs for z/OS Communications Server . Be the first to [submit a new RFE](#) for z/OS Communications Server.

All product RFEs are not included in the RFE Community database. Refer to the [FAQs](#) to see how to add existing RATLC RFEs to the RFE Community.

**Spotlight**

- [→ Announcements](#)
- [→ Give us your feedback](#)

**Brands**

- All brands
- Information Management
- Rational
- Tivoli
- **WebSphere**

**Latest RFE submitted**

No submitted RFEs for z/OS Communications Server product.

# RFE: New Requirements Process for z/OS Communications Server ...

## Submit a request for enhancement (RFE)

Use this form to submit an idea for a new product feature, also called a request for enhancement (RFE). The product development team will review your input and provide status updates as decisions are made regarding the RFE. Before you submit a new RFE, please [view RFEs that have already been submitted](#). If your idea has already been submitted, you can add comments to the existing RFE, thereby indicating your agreement with the idea. We may use this information to help prioritize development of new features.

**Note:** The company and business justification will not be visible on the Jazz.net site for RFEs submitted for Jazz products.

The fields indicated with an asterisk (\*) are required to complete the transaction. If you do not want to provide us with the required information, please use the Back button on your browser to return to the previous page.

A key icon indicates that the field is displayed only to the original submitter. The key icon next to an RFE indicates that the RFE is a private RFE.

<b>Submitter:*</b>	SamReynolds
<b>Company:*</b>	The Company field is visible to you and IBM only, as shown by the key icon (40 characters or less): <input type="text" value="none"/> (You have 36 characters left)
<b>Headline:*</b>	Please enter a summary of your request (125 characters or less): <input type="text"/> (You have 125 characters left)
<b>Submitter's ranking of priority:*</b>	What impact does this request have on your ability to use the product? <input type="button" value="Select a priority"/> <a href="#">Priority definitions</a>
<b>Brand:*</b>	<input type="text" value="WebSphere"/>
<b>Product family:*</b>	<input type="text" value="Enterprise Networking"/>

**Spotlight**

- [Announcements](#)
- [Give us your feedback](#)

**Brands**

- [All brands](#)
- [Information Management](#)
- [Rational](#)
- [Tivoli](#)
- [WebSphere](#)

**RFE activities**

- [Search RFEs](#)
- [Submit RFEs](#)

**My stuff**


- [My watchlist](#)
- [My votes](#)
- [My RFEs](#)
- [My group memberships](#)
- [My saved searches](#)
- [My RSS feeds](#)
- [My notifications](#)

**Groups**

- [Group directory](#)



## RFE: New Requirements Process for z/OS Communications Server ...

<b>Description:*</b>	Please enter a detailed description of the new feature that you want (5000 characters or less):
	<input type="text"/>
	(You have 5000 characters left)
<b>Use Case:*</b>	Please describe the scenario (use case) this feature would be used in (5000 characters or less):
	<input type="checkbox"/> <a href="#">Use case example</a>
	<input type="text"/>
	(You have 5000 characters left)
<b> Business justification:</b>	Please explain your business justification for why IBM should add this feature. Include information such as extent of individuals affected, impact on your business or project, and so forth. This information will not be publicly visible, as shown by the key icon. You can use this field to provide information that you only want to share with IBM (5000 characters or less):
	<input type="text"/>
	(You have 5000 characters left)
<b>Watch this RFE:</b>	<input checked="" type="checkbox"/> Add this RFE to my watchlist
<b>Attachments</b>	

# RFE: New Requirements Process for z/OS Communications Server ...

developerWorks > RFE Community > WebSphere >

## Submitted request for enhancement (RFE)

↓ RFE details	↓ Vote	↓ Reconsideration
↓ Attachments	↓ Comments	

The RFE shown below was submitted for a new feature of an existing product. The product development team will review your input and provide status updates as decisions are made regarding the RFE. Use this form to comment on the submission, to attach a file to the submission, to add the submission to a watchlist, or to add this submission to your most wanted features.

See the Comments section below for a complete list of comments and to add your own comment.

You can inform others of this [RFE via email](#).

### Spotlight

- Announcements
- Give us your feedback

---

### Brands

- All brands
- Information Management
- Rational
- Tivoli
- **WebSphere**

RFE details	
<b>Headline:</b>	Implement Policy Agent PBR support for IPv6
<b>ID:</b>	24239
<b>RTC ID:</b>	7066
<b>OTHER ID:</b>	1413
<b>State:</b>	Open
<b>Status:</b>	<a href="#">Uncommitted Candidate</a>
<b>Created on:</b>	06 Jul 2012, 09:40 AM Eastern Time (ET)
<b>Updated on:</b>	13 Jul 2012, 02:44 PM Eastern Time (ET)

### RFE stats

- 4 vote(s)
- 0 comment(s)
- 1 user watchlist(s)
- 0 attachment(s)

### RFE actions

- Add vote
- Add to My watchlist
- Email this RFE

---

## RFE: New Requirements Process for z/OS Communications Server ...

### ▪ Requirement Tips:

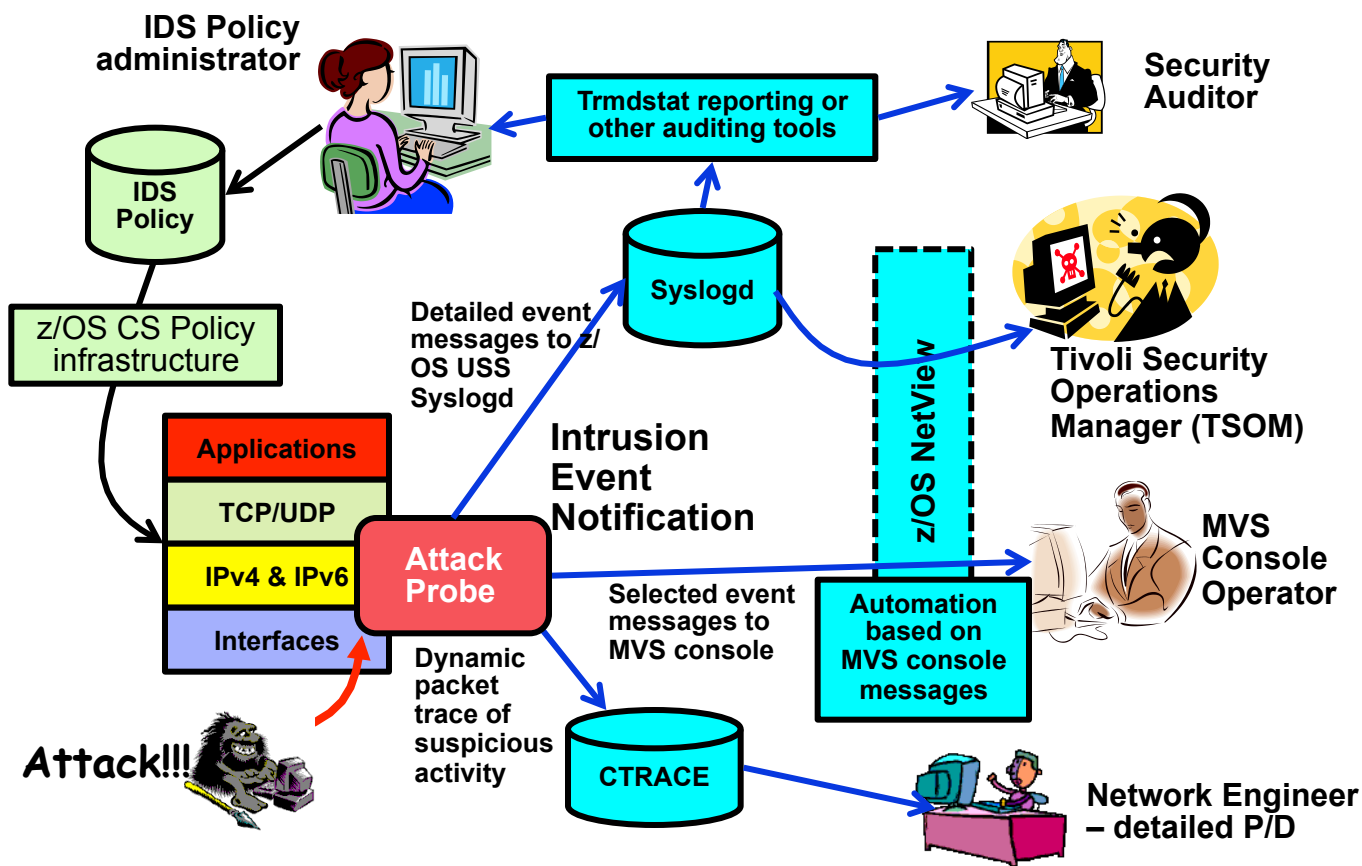
- Explain the problem you need solved, not just the requested solution
  - A request for a particular solution may not be feasible, or might take a very long time to deliver
  - By describing the problem to be addressed, we may be able to suggest alternatives that will be immediately beneficial, or a solution that we will be more likely able to implement in the reasonable future
  - Please understand that even in the best case, the delivery of a new function will likely be 2 to 3 years off due to our 2-year release cycle
- We try to disposition a requirement within a few weeks (typically marking it as an “uncommitted candidate” or closing it). During the time it is “Under Consideration” please monitor for our updates. We sometimes need to request more information, recommend alternate solutions, etc., before we can disposition it.
- If you see other z/OS Communications Server requirements that would be beneficial to your organization, please consider voting for them (as shown on the previous chart).

## **z/OS Communications Server Technical Update**

# **Security**



# Review: Intrusion Detection and Prevention services on z/OS



- ❑ **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!**

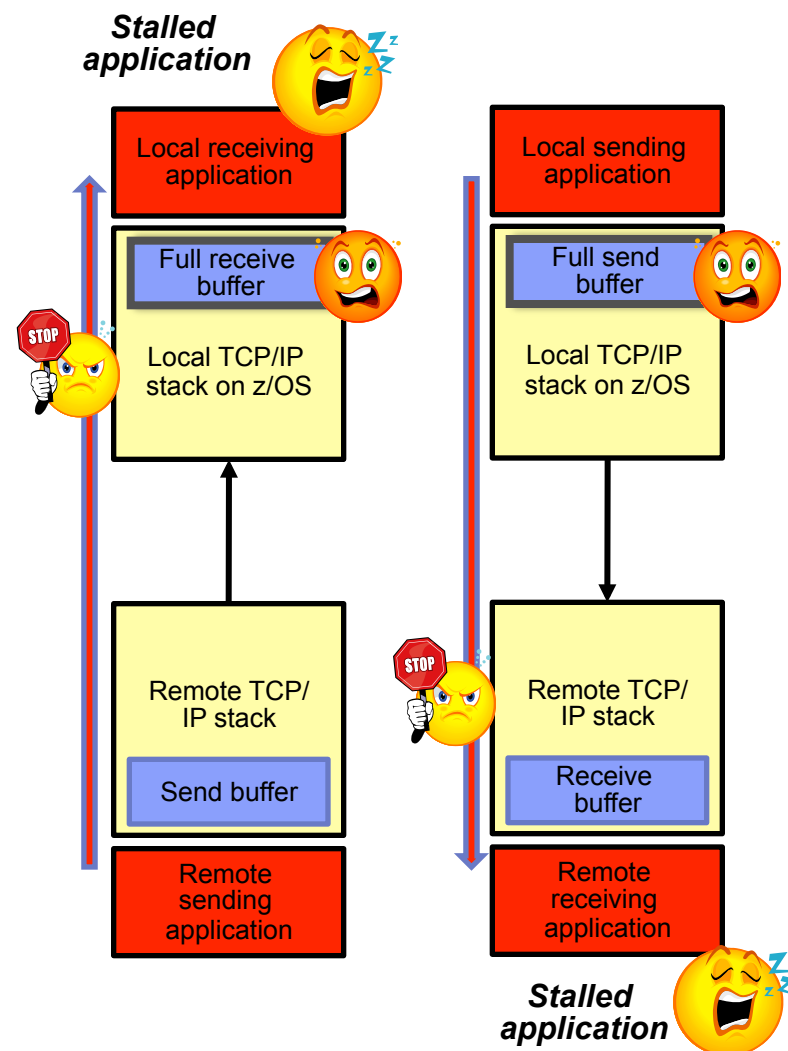
## Intrusion Detection Services enhanced to include IPv6 traffic

- **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



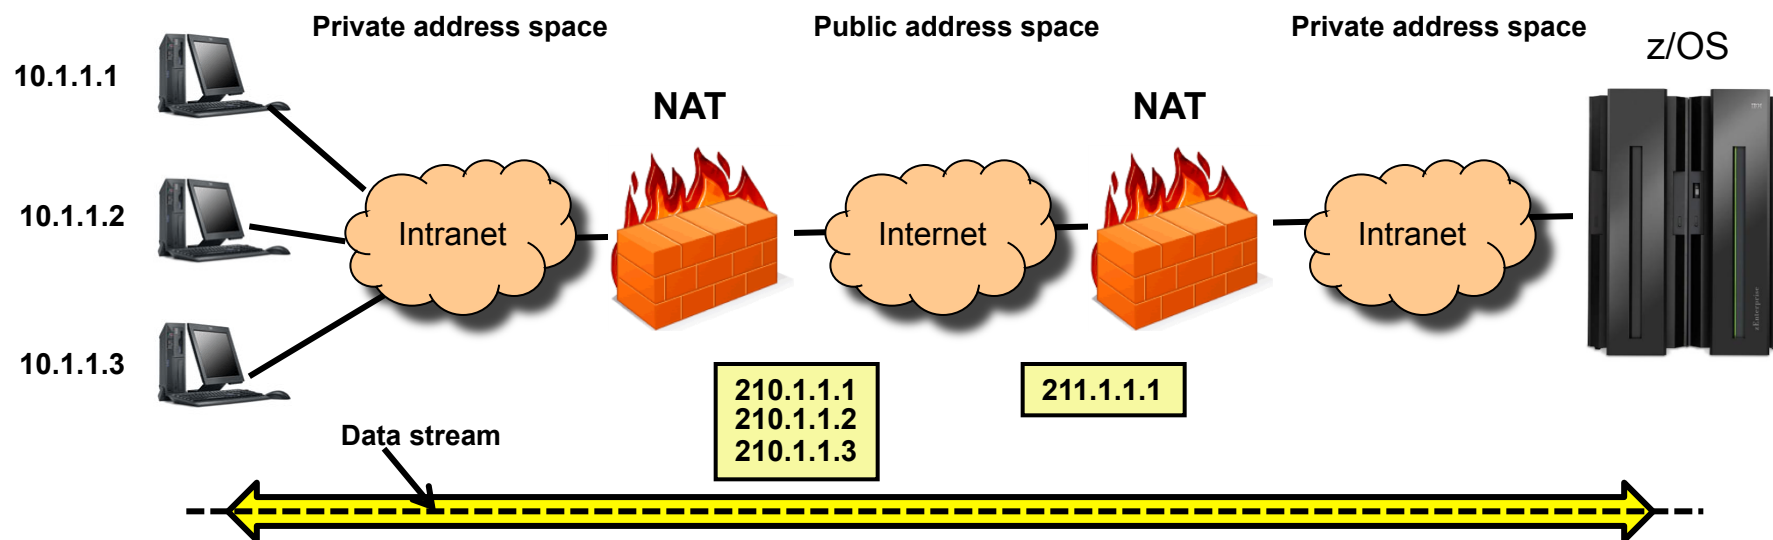
## New IDS attack types implemented for both IPv4 and IPv6

- **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



## 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!





## 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
    - APAR PM62213 added an [FTP.DATA](#) statement to indicate whether the FTP server supports logging in with passphrases.
  - 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



## 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 rename" 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.rename

```
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

## Legacy SSL support for TN3270 and FTP

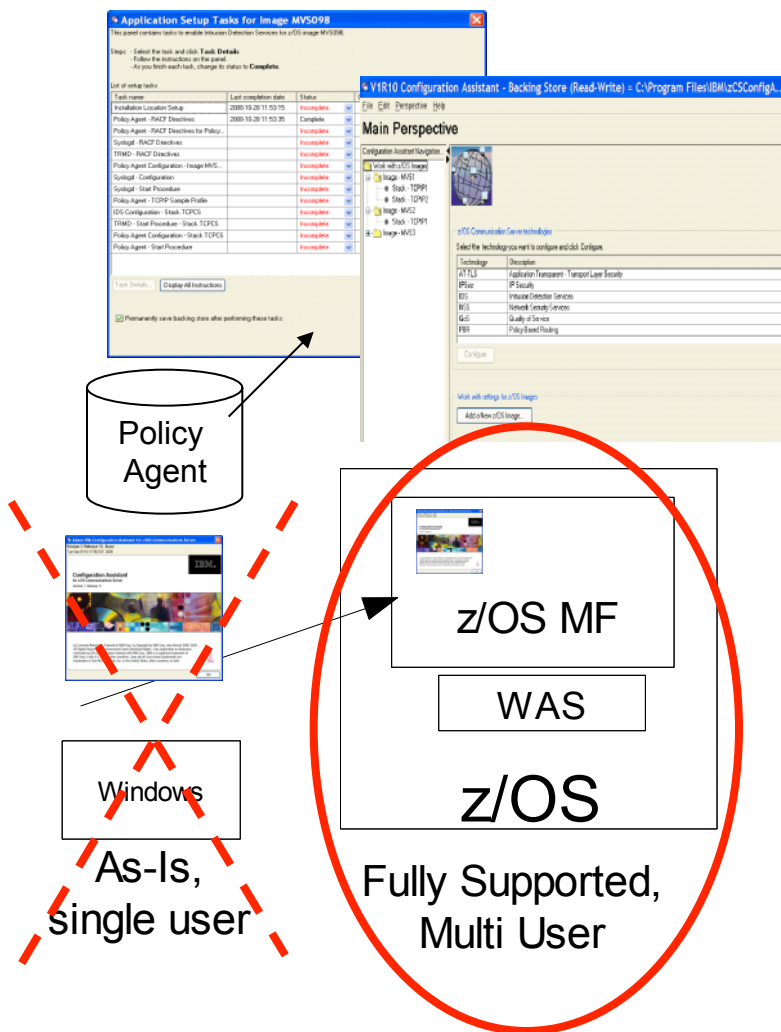
- TN3270 and FTP client/server were enabled for SSL before AT-TLS existed
- Both were enabled for AT-TLS in V1R9. This is the recommended approach
- Since then, all TLS/SSL enhancements come through AT-TLS only. Examples:
- It is our intent that all future TLS enhancements, including any new ciphers, will only be available through AT-TLS
  - Maximum “strength” of legacy SSL support: TLSv1.0, AES-CBC-256, SHA-1
  - Anything stronger/higher will require changover to AT-TLS-based protection
- We are considering removal of legacy SSL support from TN3270 and the FTP client/server in a future release
  - Will require some config changes in TN3270 / FTP configuration data sets
  - Will require use of policy agent and, ideally, configuration assistant

## **z/OS Communications Server Technical Update**

# **Simplification**



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



- 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](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>


---

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

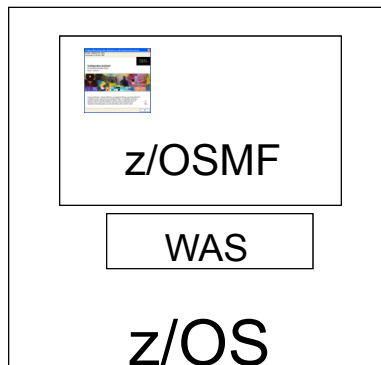
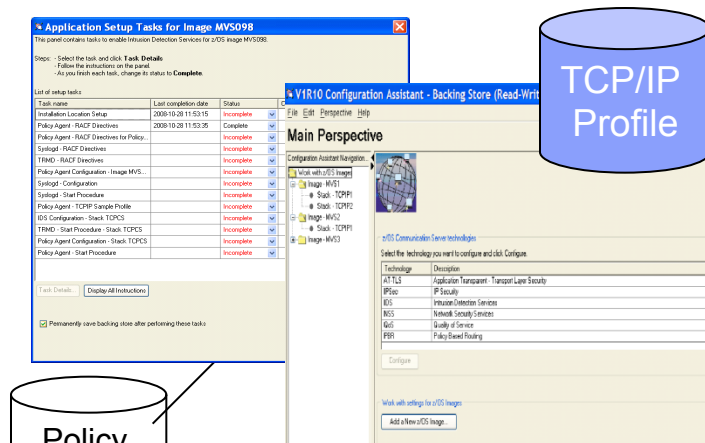
z/OS V1.13 is planned to be the final release for which the IBM Configuration Assistant for z/OS Communications Server tool that runs on Microsoft Windows will be provided by IBM. This tool is currently available as an as-is, nonwarranted web download. Customers who currently use Windows-based IBM Configuration 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 import of TCP/IP configuration information



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




## **z/OS Communications Server Technical Update**

# **Economics and platform efficiency**



## Customer Survey from Summer SHARE 2011: IPv6 responses


**IPv6 Testing on z/OS**



● Yes ● No ● Do not know

Winter, 2009 survey results:  
13% Yes, 80% No

**Date Need IPv6 in Production**



● 2011-2012 ● 2013-2014 ● 2015-2016  
● Later than 2016 ● Do not know

# IPv4 Address Exhaustion – When?

<http://www.potaroo.net/tools/ipv4/index.html>

**IPv4 Address Report**

This report generated at 17-May-2012 07:59 UTC.

---

IANA Unallocated Address Pool Exhaustion:

**03-Feb-2011**

Projected RIR Address Pool Exhaustion Dates:


RIR	Projected Exhaustion Date	Remaining Addresses in RIR Pool (/8s)
APNIC:	<b>19-Apr-2011</b>	1.1641
RIPENCC:	<b>14-Aug-2012</b>	2.2146
ARIN:	<b>20-Jun-2013</b>	4.9024
LACNIC:	<b>29-Jan-2014</b>	3.6292
AFRINIC:	<b>06-Nov-2014</b>	4.3216

---



**This is no longer a future long-term concern!!!!**

# IPv6 – State of z/OS and z/OS Communications Server



**A few applications and add-on functions still need IPv6-enablement: Sysplex-wide security associations, policy-based routing, remote commands, IPsec NAT traversal, and some less frequently used applications and functions.**

Important z/OS applications and subsystems are already IPv6 enabled

**z/OS Communications Server applications and z/OS-unique functions are not defined in any compliance criteria, but many are already IPv6 enabled:**

- High-availability functions IPv6-enabled: DVIPA, Sysplex, etc.
- Add-ons such as IP Security, AT-TLS, etc.
- Applications (TN3270, EE, FTP, CSSMTP, etc.)
- Management functions (SNMP, SMF records, NMI, OSPF, etc.)
- Subsystems are picking up (WAS, CICS, MQ, etc.)

Good for real, full-function, reliable "production" use

Good for US government use

z/OS V1R10 CS certified by DoD in 2008  
z/OS V1R12 CS certified By USGv6 in 2010

**US Government compliance criteria**

1. Department of Defense (DoD)
2. All other agencies via NIST (National Institute of Standards and Technology)



z/OS V1R8 and V1R11 CS certified as IPv6 Phase 2 Ready

**IPv6 Ready Logo compliance based on "Tahi" test**

Good for "commercial" use




Started in z/OS V1R4 CS – continually updating

**IPv6 Base RFC compliance based on standards bodies specifications**

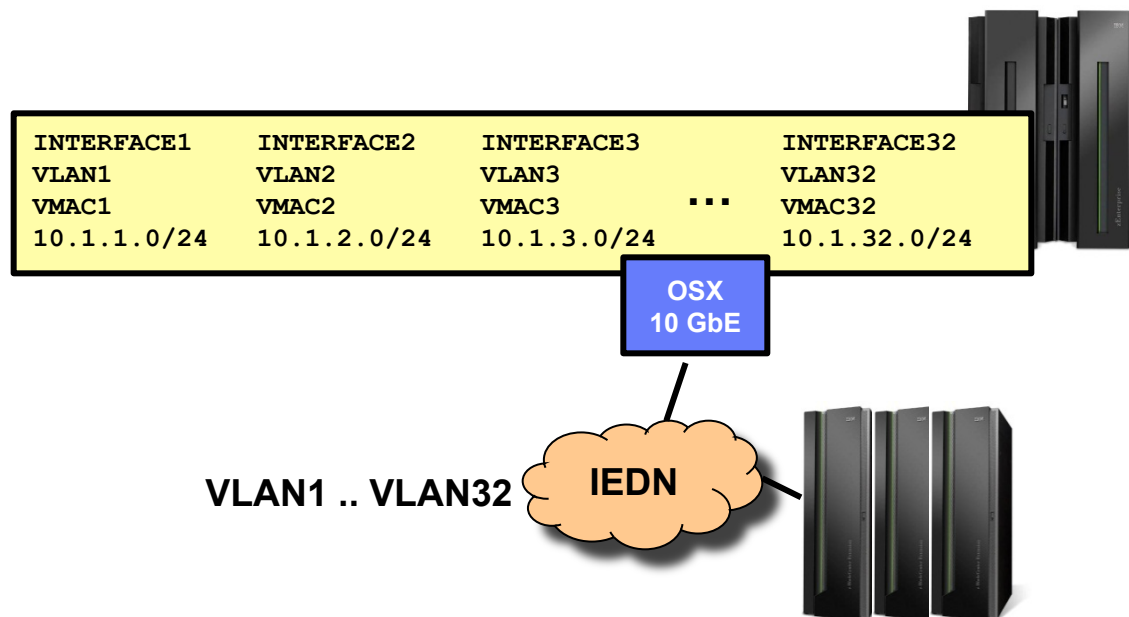
## Resolver support for IPv6 connections to DNS name servers (V1R12)

- Allows the system resolver to send requests to DNS name servers using IPv6 communication
  - Specify IPv6 addresses on the NSINTERADDR and NAMESERVER configuration statements
  - Resolver sends queries using IPv4, IPv6 or both based on the configuration
  
- Applications cannot manipulate IPv6 addresses using low-level resolver API calls, such as res\_query and res\_search
  - Only IPv4 addresses are supported on these APIs
  - The entire list, containing IPv4 and IPv6 addresses, is used for searching
    - Unless the application modifies the list, in which case only the returned IPv4 addresses are used
  
- The type of address returned (IPv4/IPv6) is not tied to the transport between the resolver and the name server. IPv6 addresses can be returned before z/OS V1R12



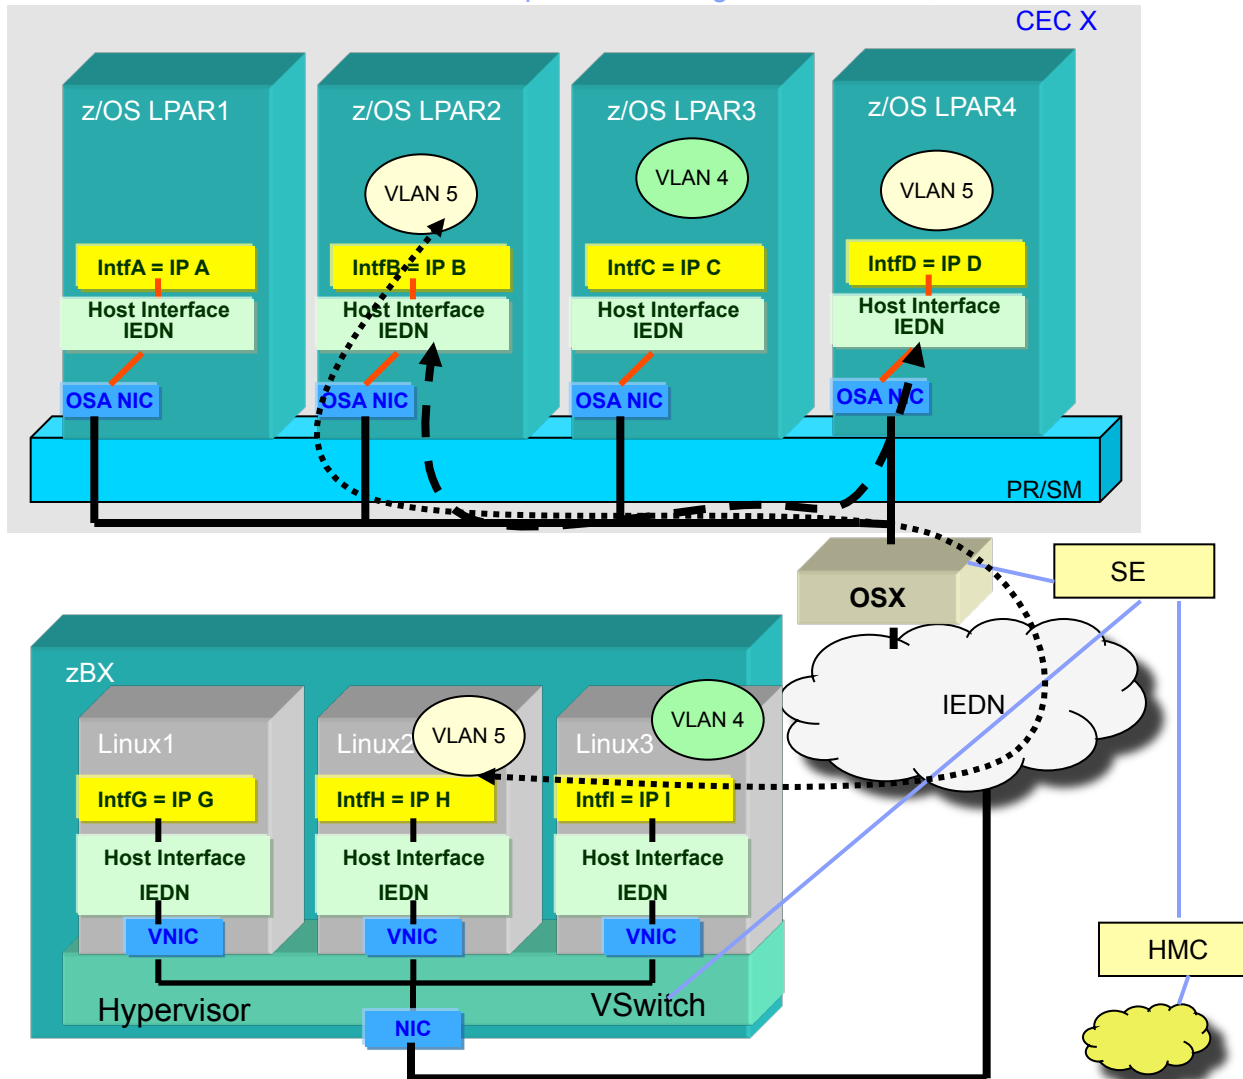
## Enhanced support for OSA VLANs

- 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



## zEnterprise IEDN without Hipersockets

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

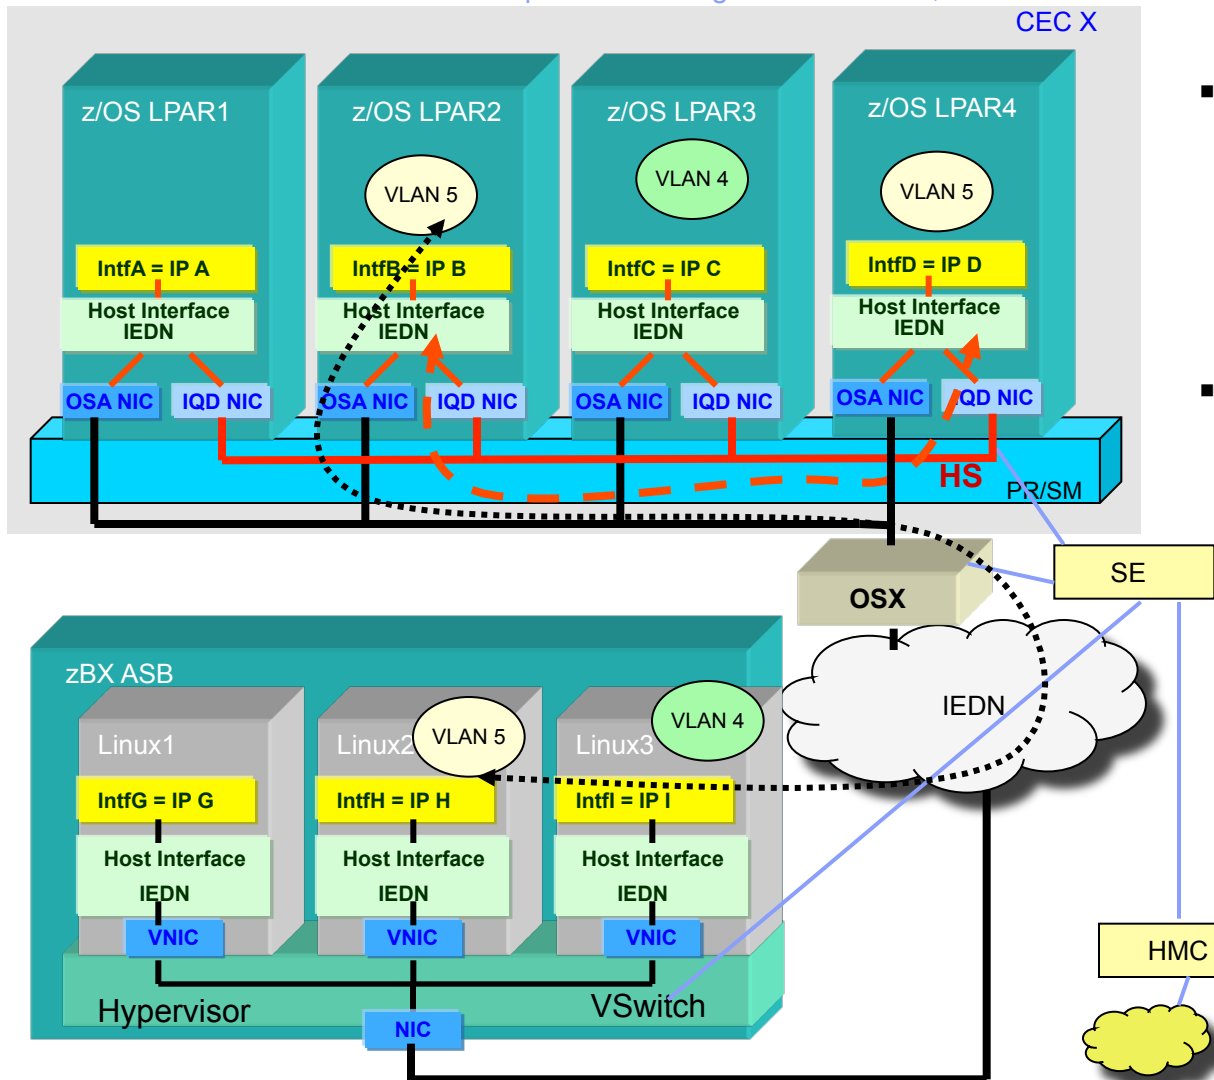


HiperSockets is another type of System z internal network that is a System z differentiator!

...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

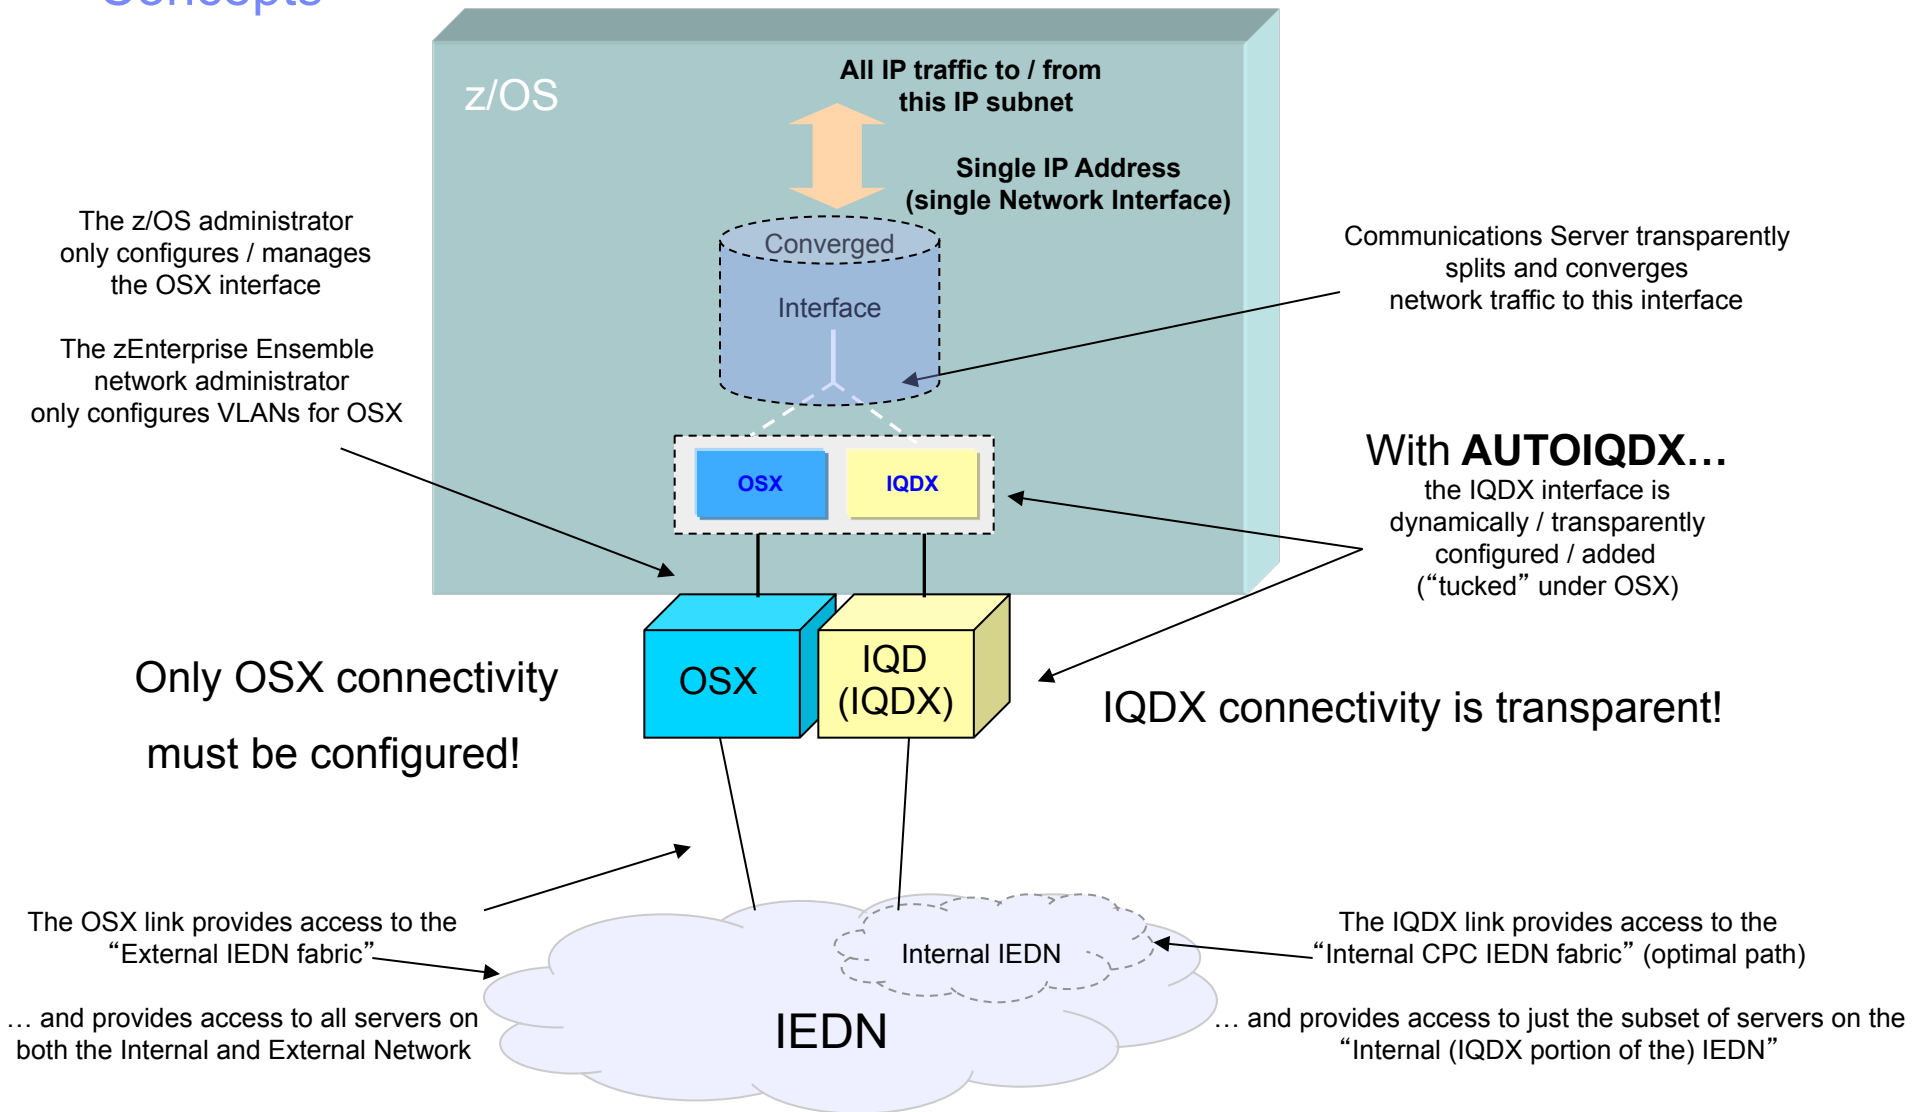
.... Intra Ensemble Data Network with platform managed virtualization, isolation and access controls



- HiperSockets becomes part of the IEDN
  - z/OS support in V1R13
  - zVM support in zVM 6.2
  - zEnterprise support required as well (see announcement)
  
- 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.

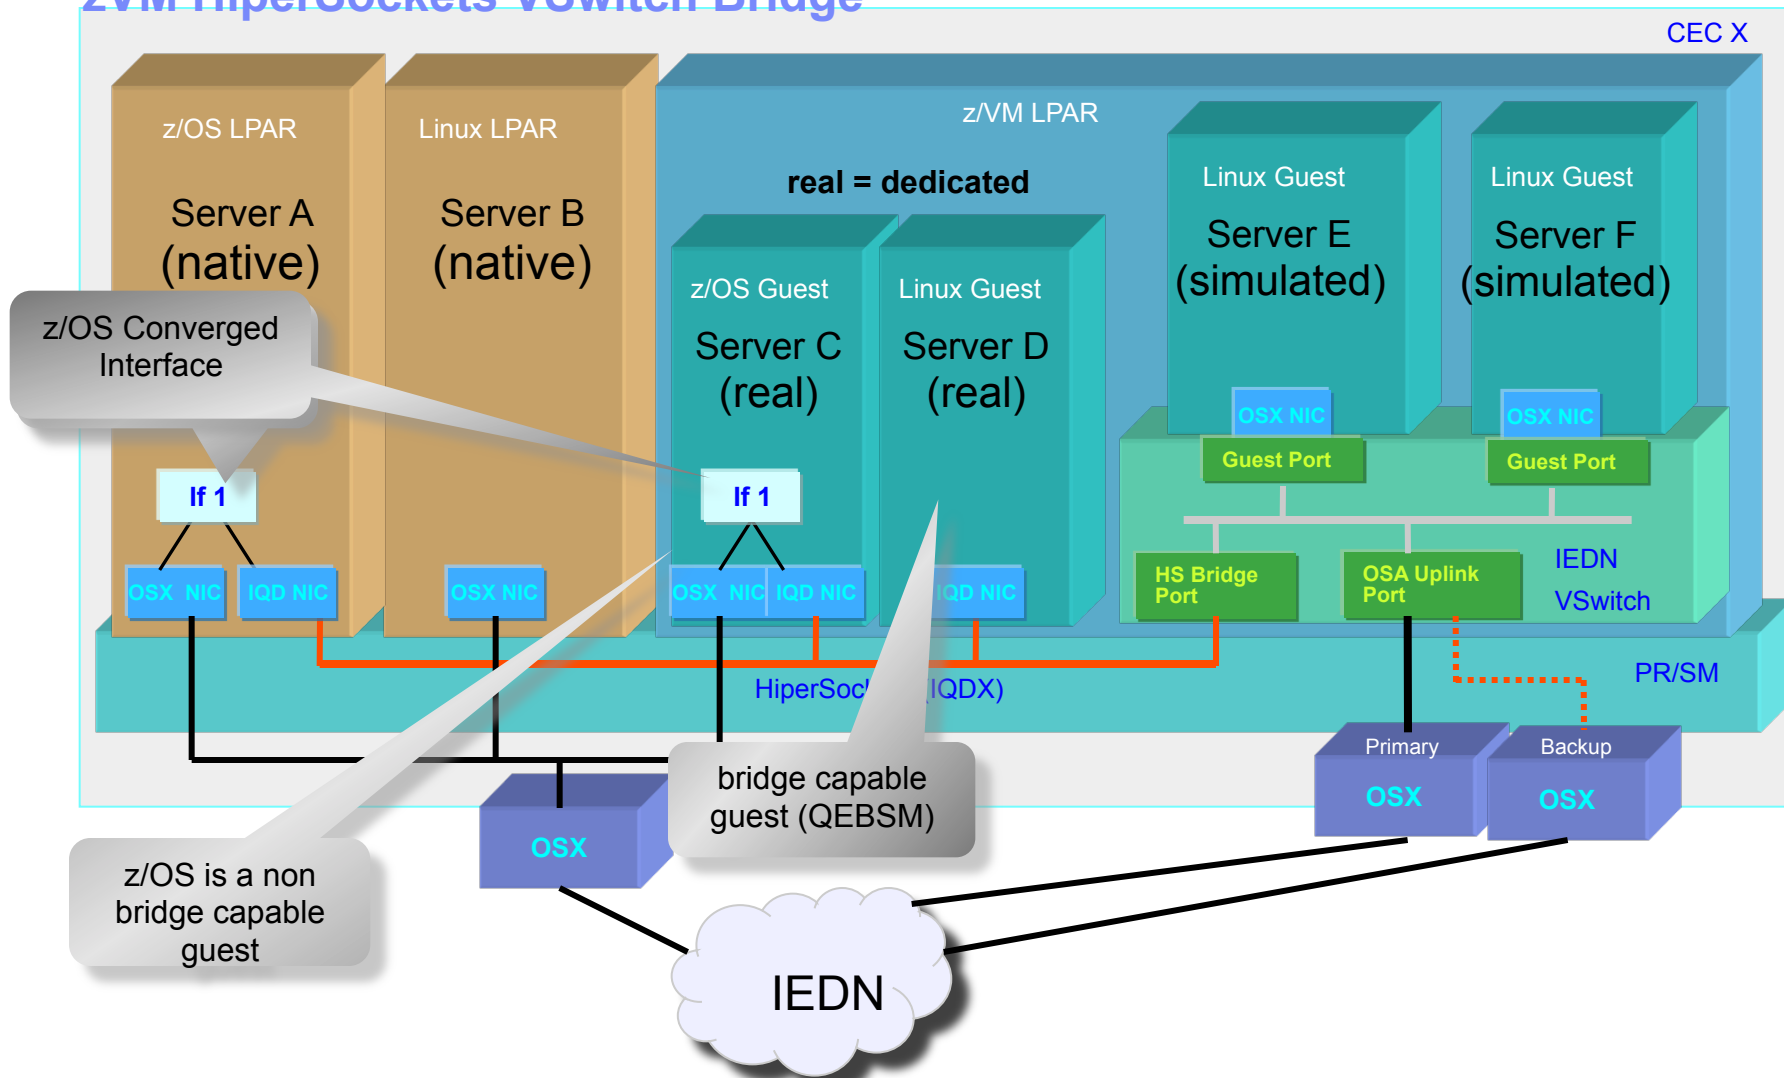


# IEDN enabled HiperSockets - z/OS “Converged IQDX Link” Concepts



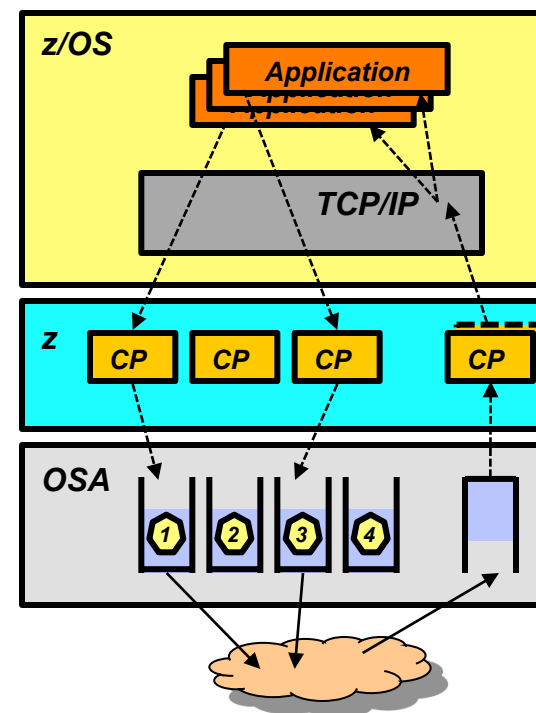
# HiperSockets IEDN Access IQDX Configuration

## zVM HiperSockets VSwitch Bridge



## 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](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

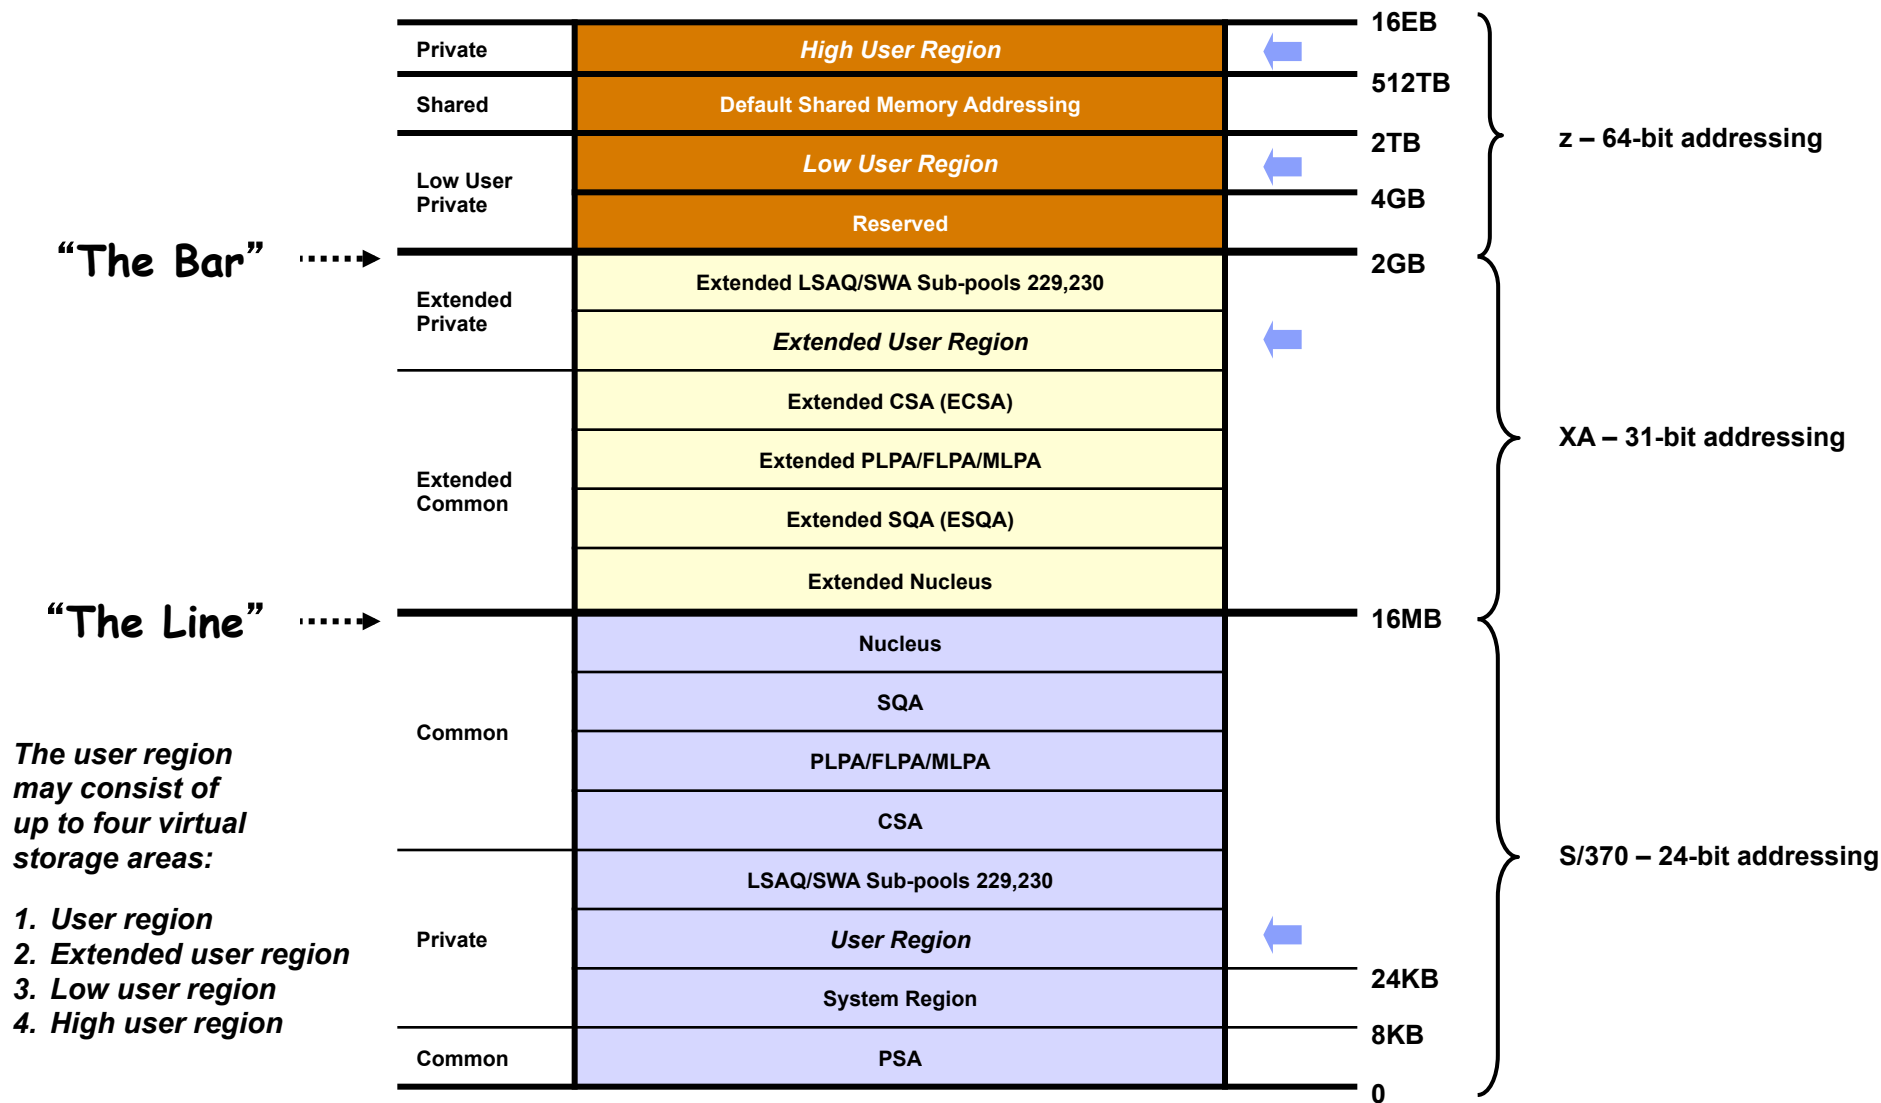


- Prior to V1R13, for IPv4 traffic from z/OS 1 to z/OS 3, checksum processing can be offloaded to the OSA NIC. TCP segmentation processing for IPv4 traffic can also be offloaded to OSA on this traffic pattern (using OSA-Express Large Send support). IPv6 traffic could not exploit either of these features.

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

# z/OS virtual storage map



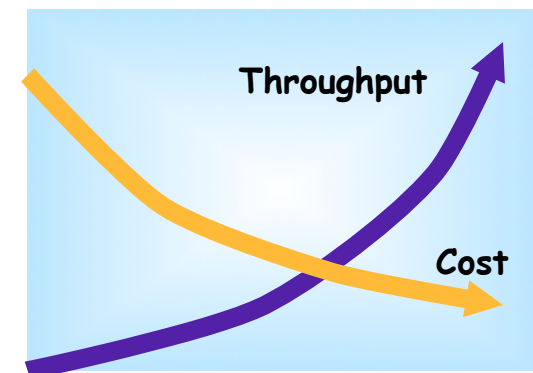
## Additional 64-bit exploitation

- 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

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

## 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)





---

## Please fill out your session evaluation

- z/OS Communications Server Technical Update, Part 1
- Session # 11327
- QR Code:

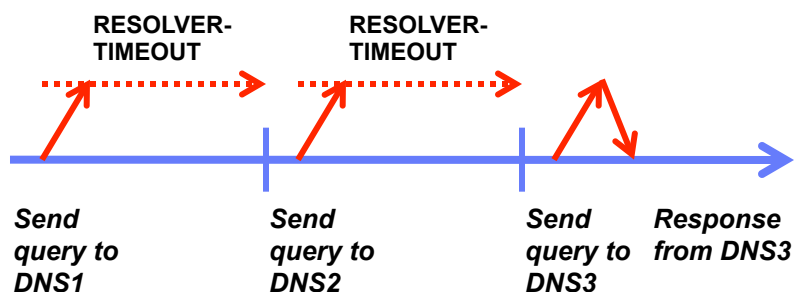


## **z/OS Communications Server Technical Update**

# **Availability**



## Improved resolver reaction to unresponsive name servers (V1R12)



**Assume:**

- 3 name servers in TCPIP.DATA
- 2 first are un-responsive
- RESOLVERTIMEOUT 30 seconds

**It takes 60+ seconds to get a response, and it will do so for every query made to the resolver**

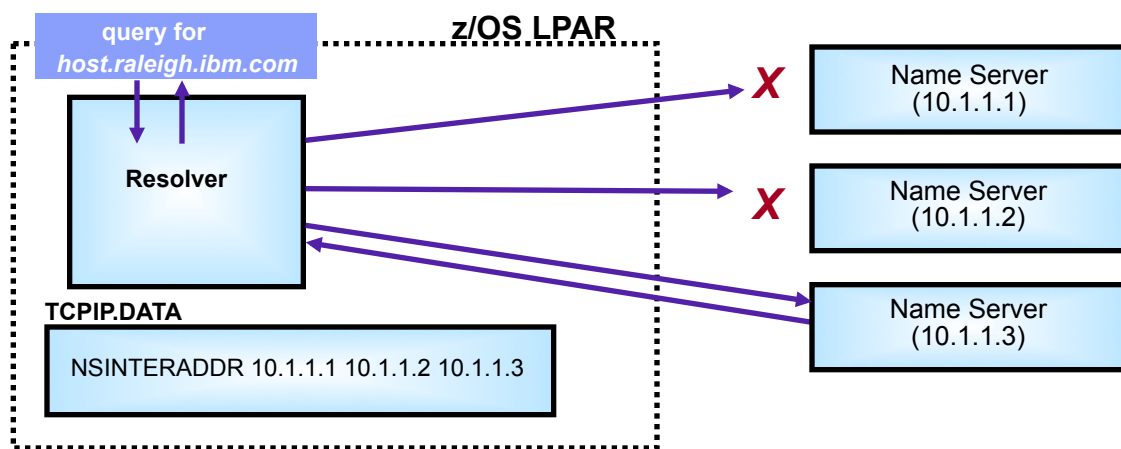
- Un-responsive name servers can impact performance significantly
  - Based on the setting of number of name servers, timeout, and retry limit in TCPIP.DATA
    - Beware that default RESOLVERTIMEOUT used to be 30 seconds – should be lowered to seconds or sub-seconds!
    - Default changed to 5 seconds in z/OS V1R12
- So far, no warning messages have been issued when name servers repetitively time out
- z/OS V1R12 adds messages to the console when name servers are un-responsive
- Configurable un-responsiveness threshold: percentage of failed queries over a 5-minute period
  - Default 25%
- A message will also be issued when a name server is deemed to have become responsive again

```

EZZ9308E UNRESPONSIVE NAME SERVER DETECTED AT IP ADDRESS 9.43.25.200
EZZ9310I NAME SERVER 9.43.25.200
          TOTAL NUMBER OF QUERIES SENT          6000
          TOTAL NUMBER OF FAILURES              2100
          PERCENTAGE                             35%
    
```

## 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



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 (i.e. all name servers are unresponsive)
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

## Solution: when does resolver stop polling?

- The resolver polls a name server indefinitely until one of these situations occur:
  - Failure rate of queries to this name server drops under 1%
  - Name server is removed from NSINTERADDR statements in global TCPIP.DATA file
  - Autonomic quiescing function is stopped using MODIFY RESOLVER,REFRESH,SETUP= command
- Switching between network operator notification and autonomic quiescing functions dynamically is permitted
  - Current statistics are discarded
  - Current operator messages for unresponsive name servers are cleared from console



## Configuring the autonomic quiescing function

- You enable the autonomic quiescing function by specifying the new AUTOQUIESCE operand on the existing UNRESPONSIVETHRESHOLD resolver setup statement
  - AUTOQUIESCE is not set by default

```

.-UNRESPONSIVETHRESHOLD (25) -----
>>+-----+-----<<
' -UNRESPONSIVETHRESHOLD (percentage) ----- '
' -UNRESPONSIVETHRESHOLD (percentage,AUTOQUIESCE) - '

```

- You must specify GLOBALTCPIPDATA for the autonomic quiescing function to be active
  - AUTOQUIESCE ignored if no GLOBALTCPIPDATA specified

```

F RESOLVER,REFRESH,SETUP=USER.TCPPARMS (RESSETUP)
EZD2036I AUTOQUIESCE IGNORED - GLOBALTCPIPDATA REQUIRED

```

- You must provide a non-zero percentage value to enable the function
  - UNRESPONSIVETHRESHOLD(0) and UNRESPONSIVETHRESHOLD(0,AUTOQUIESCE) disable all monitoring functions

## Operator notifications (unresponsive server)

- The resolver issues EZZ9311E and EZZ9313I when a name server is considered to be unresponsive
  - One message pair per name server per instance of becoming unresponsive
- Statistics represent data from the last 30 or 60 seconds
  - 60 seconds (2 intervals) only used if insufficient sample size in the last 30 seconds, and resolver was polling the name server
  - “Total Number of Queries Sent” includes “Total Number of Resolver Polls Sent”
  - “Total Number of Failures” includes “Total Number of Failed Polls”

```
EZZ9311E STOPPED USING NAME SERVER AT IP ADDRESS 10.1.1.1
EZZ9313I NAME SERVER 10.1.1.1
      TOTAL NUMBER OF QUERIES SENT           510
      TOTAL NUMBER OF FAILURES              306
      TOTAL NUMBER OF RESOLVER POLLS SENT    10
      TOTAL NUMBER OF POLL FAILURES          6
      PERCENTAGE                             60%
```



## Operator notifications (responsive server)

- The resolver issues EZZ9312I when a previously unresponsive name server becomes responsive again
  - Failure rate during polling must be less than threshold for name server to be responsive
  - Resolver might continue polling, depending on the failure rate of the name server responding to the polls
- No EZZ9312I message is issued if autonomic quiescing is stopped dynamically and a name server was unresponsive at the time
  - Any EZZ9311E messages are cleared
- The resolver will wait at least 60 seconds after declaring a name server to be responsive before declaring it unresponsive a second time

```
EZZ9312I RESUMED USING NAME SERVER AT IP ADDRESS 10.1.1.1
```

## Displaying name server status


- When autonomic quiescing function is active ...
  - AUTOQUIESCE is included in the list of resolver setup statements
  - Name servers currently specified on NSINTERADDR statements in global TCPIP.DATA file are listed, along with status and failure rate at the last statistics checkpoint
- When autonomic quiescing function is not active ...
  - Only UNRESPONSIVETHRESHOLD setting is displayed
  - No name server status is displayed

```

F RESOLVER,DISPLAY
EZZ9298I DEFAULTTCPIPDATA - None
EZZ9298I GLOBALTCPIPDATA - SYS1.TCPPARMS(TCPDATA)
EZZ9298I DEFAULTIPNODES - USER55.ETC.IPNODES
EZZ9298I GLOBALIPNODES - None
EZZ9304I CACHE
EZZ9298I CACHESIZE - 200M
EZZ9298I MAXTTL - 2147483647
EZZ9298I UNRESPONSIVETHRESHOLD - 25
EZZ9304I AUTOQUIESCE
EZD2305I NAME SERVER 10.1.1.1
                STATUS: ACTIVE           FAILURE RATE: 0%
EZD2305I NAME SERVER 10.1.1.2
                STATUS: QUIESCED        FAILURE RATE: 100%
EZD2305I NAME SERVER 10.1.1.3
                STATUS: ACTIVE           FAILURE RATE: *NA*
EZZ9293I DISPLAY COMMAND PROCESSED
  
```

## 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.

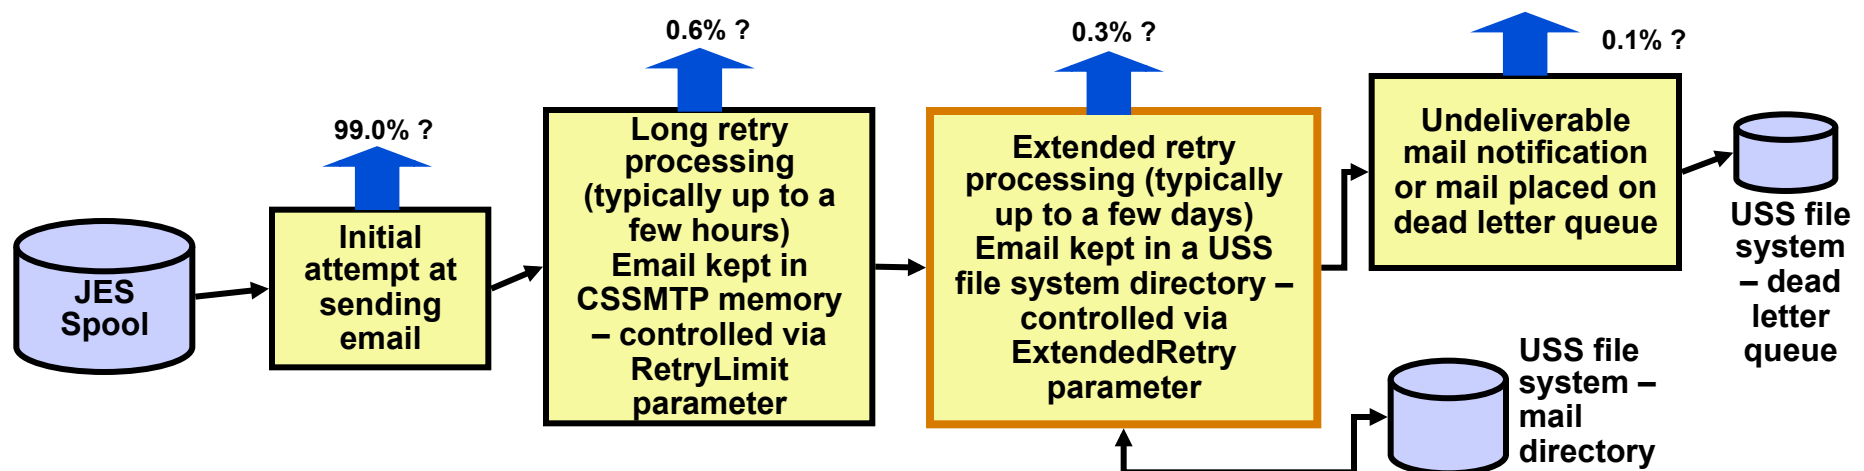
- If you are migrating to z/OS V1R13 Communications Server and you already have a user ID for the system resolver procedure, no action is required if you meet either of these conditions:
  - You explicitly defined an OMVS segment for the user ID.
  - An OMVS segment was created through the RACF automated 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



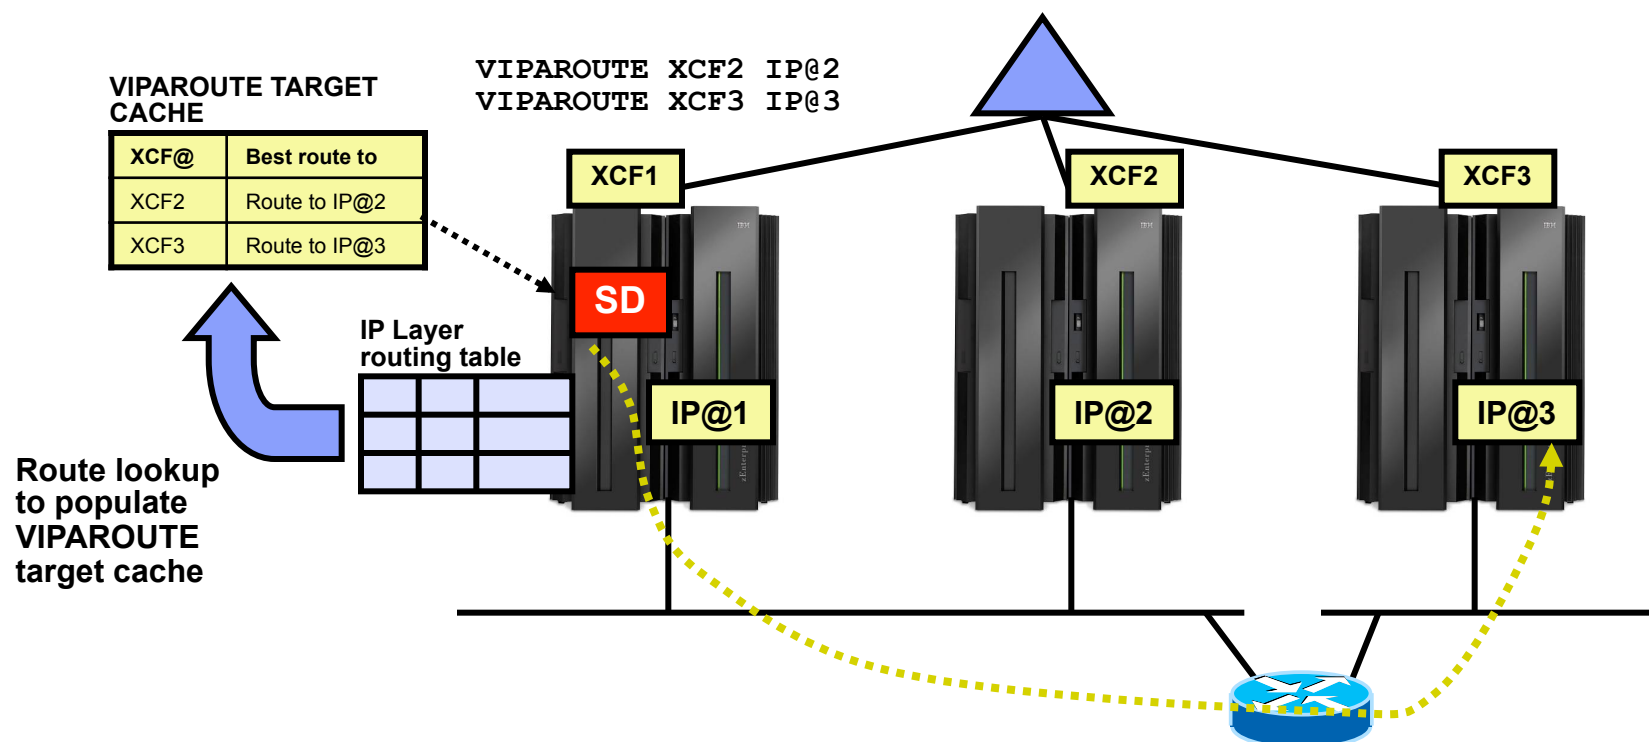
## CSSMTP translate code page support enhanced

- 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
- 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' )

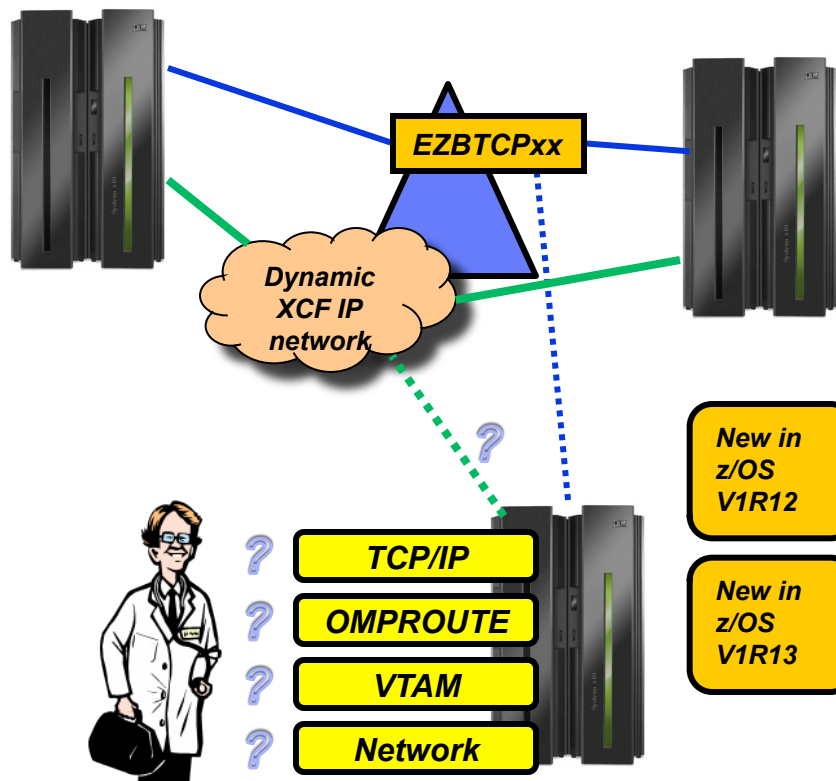
Code Page	CCSID	@	§
IBM-1047 (EBCDIC)	1047	0x7C	0xB5
ISO-8859-1 (ASCII)	819	0x40	0xA7
IBM-273	273	0xB5	0x7C

## 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



# Sysplex autonomies extended with CSM storage constrained monitoring



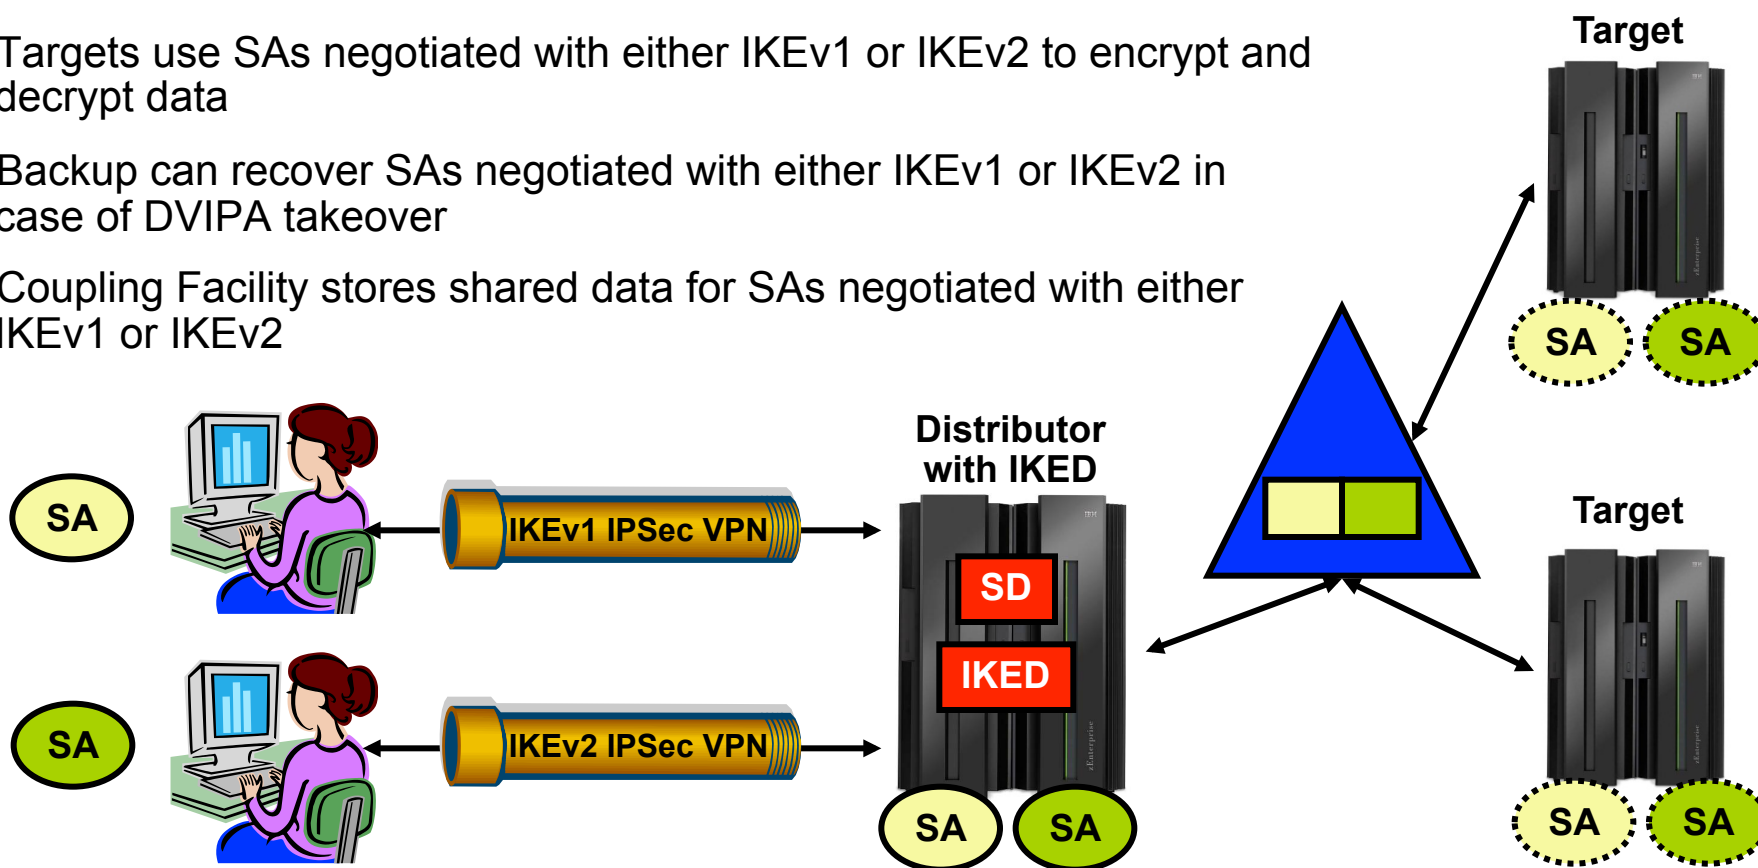
- **Monitoring:**
  - Monitor CS health indicators
    - Storage usage critical condition (>90%) - CSM, TCPIP Private & ECSA
      - For more than TIMERSECS seconds
  - Monitor dependent networking functions
    - OMPROUTE availability
    - VTAM availability
    - XCF links available
  - Monitor for 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)





## 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



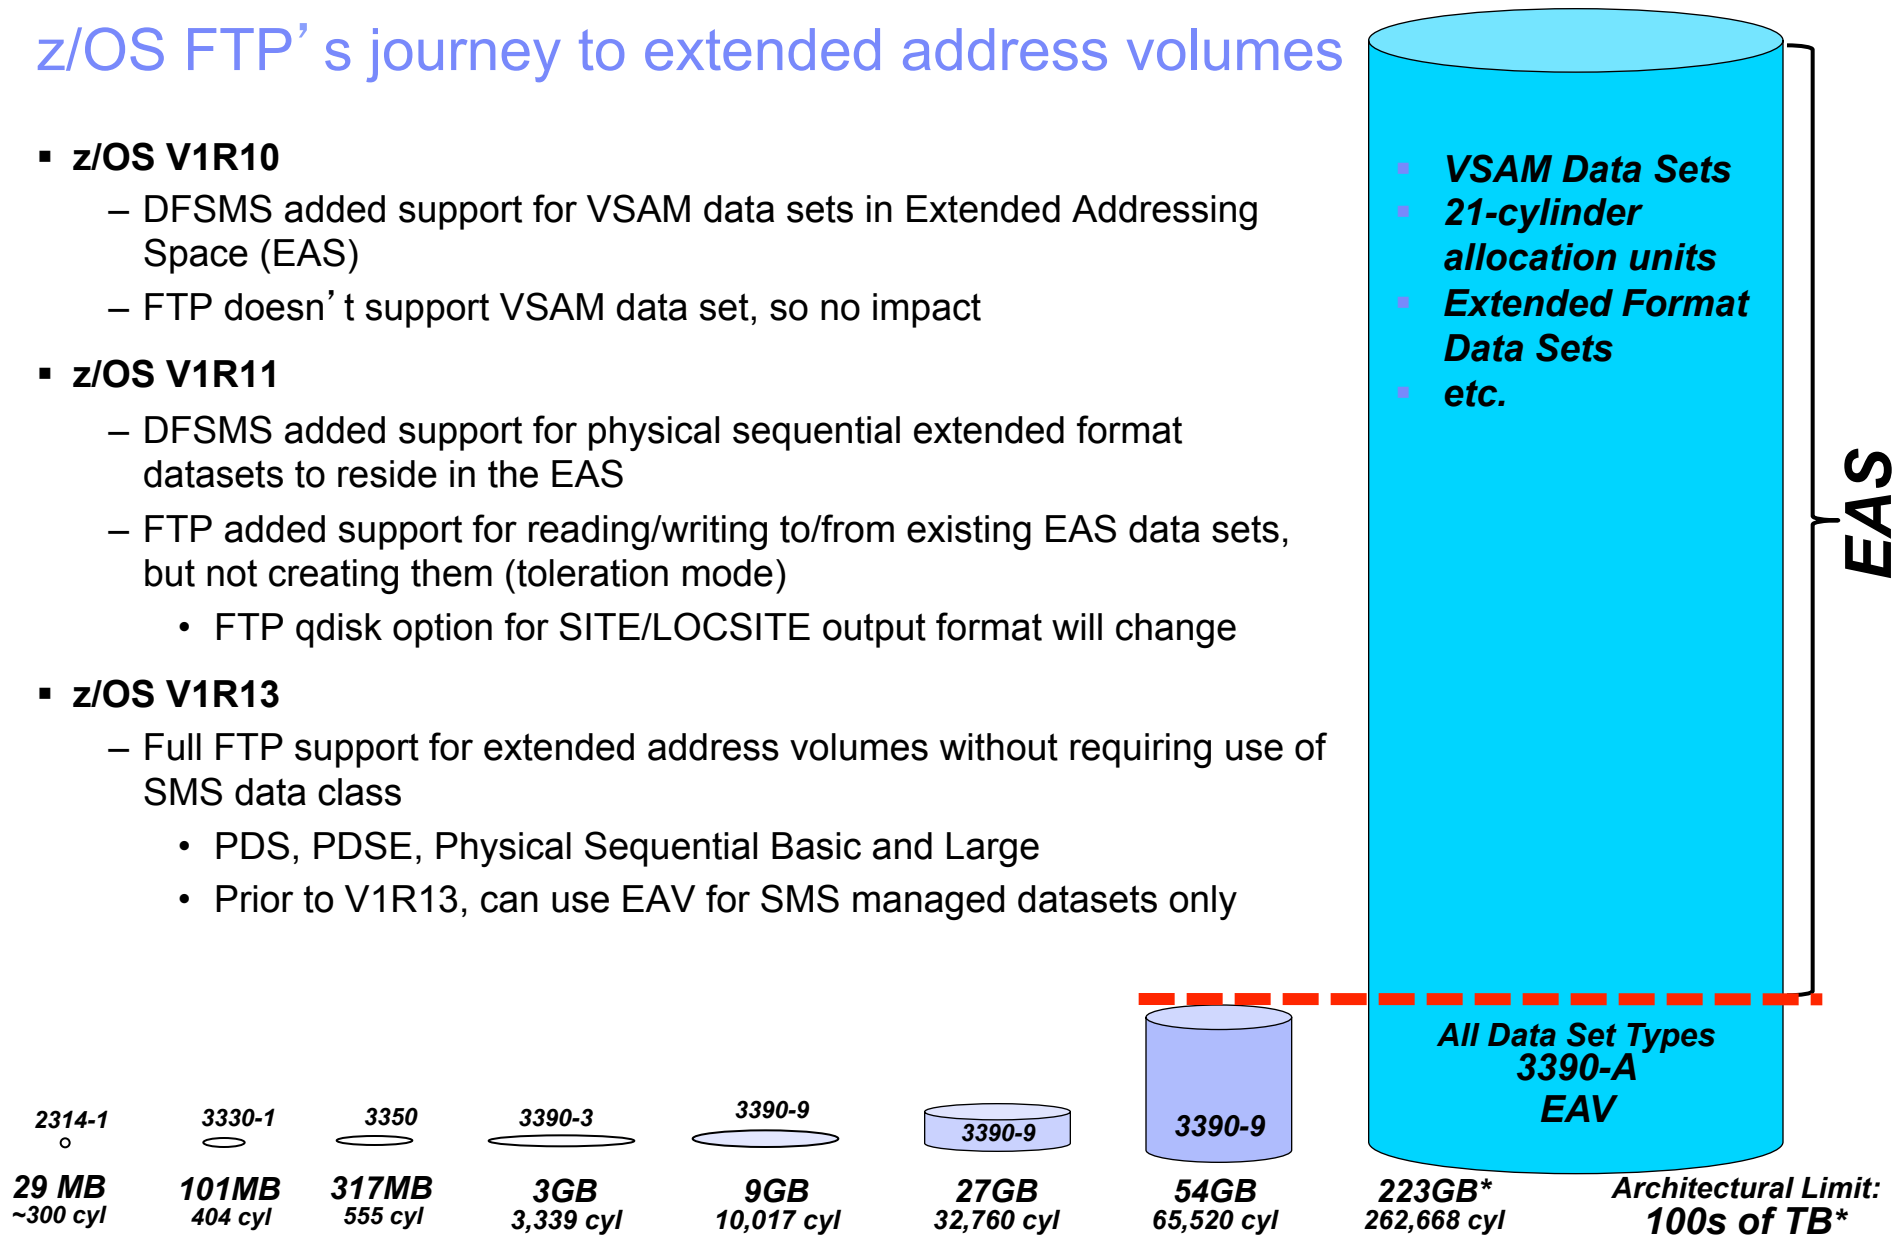
## **z/OS Communications Server Technical Update**

# **Application / Middleware / Workload enablement**



## 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



## 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

## 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*	



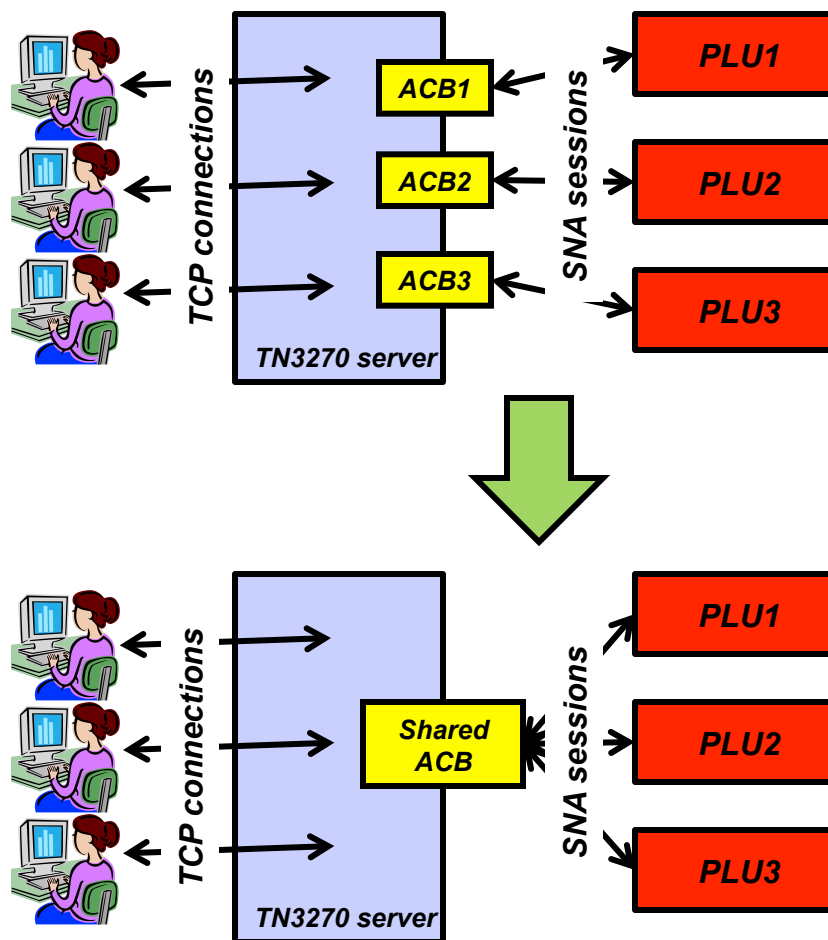
## **z/OS Communications Server Technical Update**

# **Enterprise Extender / TN3270 / SNA**



## TN3270 server improvements – shared ACB support for improved performance and reduced ECSA storage use (V1R12)


- Telnet shared ACB support can be turned on or off with a simple statement in TELNETGLOBALS section
- VTAM model statements must be used to define the Telnet LUs
- Shared ACBs remain open until the Telnet server is ended.
  - Improve path length for client logon by using an ACB which is already open
  - Improve path length for client logoff by avoiding CLOSE ACB
  - Improve path length for Telnet termination by having fewer ACBs to close
  - Reduce the likelihood of Telnet hangs due to CLOSE ACB
  - Reduce TN3270 server ECSA usage
- No change to VTAM display commands



## TN3270 server ECSA usage improvement up to and including z/OS V1R12 Communications Server (V1R12)

Release	ECSA for 256K TN3270 sessions
V1R7	798M
V1R8	708M
V1R9	480M
V1R10	440M
V1R11	347M
V1R12 <sup>(1)</sup>	249M

**ECSA for 256K TN3270 sessions**



The numbers are configuration dependent, but they should give you an idea of the magnitude of the savings achieved in the recent releases.

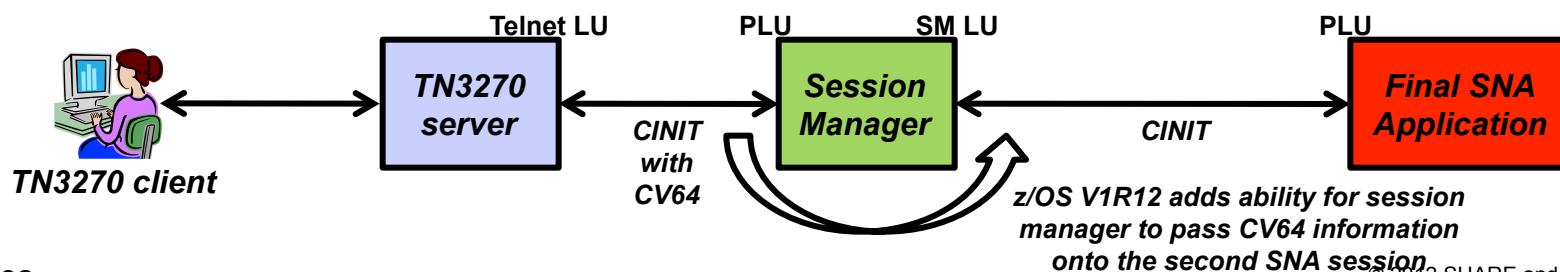
**Note (1):** The V1R12 number is a preliminary number based on use of shared ACBs.

## TN3270 server improvements – IP management information through a relay-mode session manager (V1R12)

- TN3270 server passes selected IP management information to the SNA side via a control vector known as a “CV64”
  - CV64 includes client IP address, port, optionally host name and secure connection status
  - A VTAM display of the Telnet LU includes some IP information

```
IST1727I DNS NAME: CRUSET60P.RALEIGH.IBM.COM
IST1669I IPADDR..PORT 9.27.40.41..3907
```

- The CV64 is also passed to the SNA primary logical unit (PLU) via its logon exit
- When the SNA PLU is a session manager that relays the SNA session over another LU to the final SNA application PLU, the CV64 information is lost on that second session
  - The session manager has no SNA APIs available to propagate the CV64 information
- z/OS V1R12 adds such an API, allowing an enabled session manager to pass the CV64 information to the final SNA application



## Additional TN3270 server usability enhancements (V1R12)

- OMVS can be shutdown and restarted without re-IPLing z/OS
  - F OMVS,Shutdown
  - F OMVS,Restart
- Following the shutdown of OMVS, you are supposed to manually stop telnet
  - If Telnet stays up after OMVS is restarted, Telnet behavior is unpredictable.
- In z/OS V1R12 Telnet server address spaces register with OMVS and get notified when OMVS is being shut down
  - Telnet will shut down with OMVS
    - OMVS shutdown is delayed until Telnet has shut down
  - Must be restarted after OMVS has been restarted
- A new option is passed in the CV64 control vector to an SNA primary LU on the CINIT flow
  - The option informs the SNA application if the TN3270 connection is a secure connection or not
  - Can be used by the SNA application to determine requirements for additional security
- To prevent a change of TN3270 connection attributes during a takeover process, a new configuration option is added to the takeover definitions:
  - TKOGENLURECON and TKOSPECLURECON – SAMECONNTYPE
- TN3270 server messages will now indicate the name of the TN3270 server address space instead of just saying ‘TELNET’

```

F OMVS ,SHUTDOWN
BPXI055I OMVS SHUTDOWN REQUEST ACCEPTED
EZZ6008I TELNET STOPPING
EZZ6028I TELNET TRANSFORM HAS ENDED
EZZ6010I TELNET SERVER ENDED FOR PORT 3023
EZZ6010I TELNET SERVER ENDED FOR PORT 2023
EZZ6010I TELNET SERVER ENDED FOR PORT 1024
EZZ6010I TELNET SERVER ENDED FOR PORT 1023
EZZ6009I TELNET SERVER STOPPED
  
```

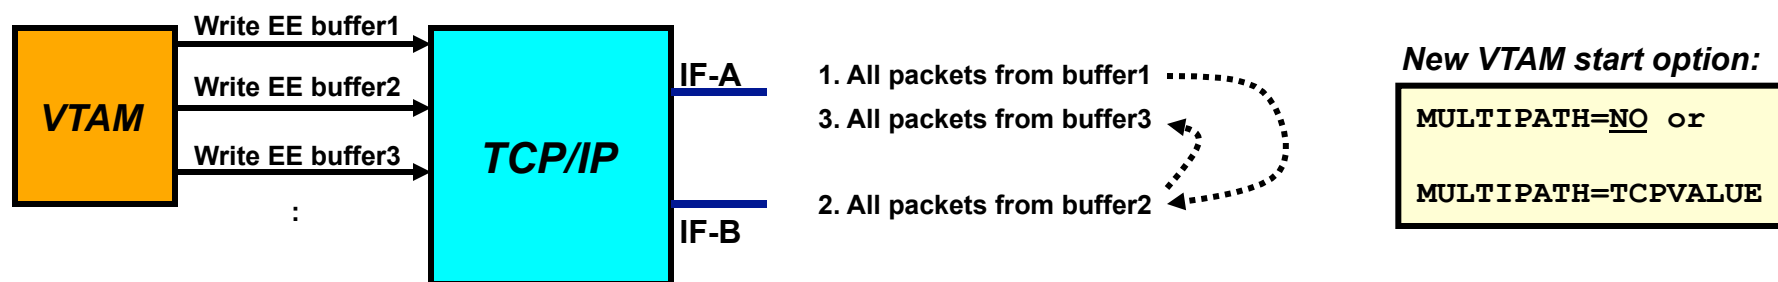
## Command to display all TN3270 servers

- 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 ***
```


## Multipath control for Enterprise Extender (V1R12)

- With multipath enabled in TCP/IP, all packets in one EE write buffer will be sent over one interface, and all packets in the next EE write buffer will be sent over another interface
  - A modified per-packet algorithm – really a per-EE-buffer algorithm
- Same behavior independent of PERCONNECTION / PERPACKET setting in TCP/IP
- EE traffic may incur performance issues if the different paths are not truly equal in terms of bandwidth and delay
- Per-connection multipath is generally beneficial for other TCP/IP traffic
- New support to allow TCP/IP to specify use of Multipath, but disable it by default for EE traffic



## Improved recovery from RTP pipe stalls (V1R12)

- z/OS V1R10 provided a version of Path MTU Discovery (PMTU) for Enterprise Extender.
  - However, MTU size changes over an active EE link are only communicated to the two nodes that act as the endpoint of that EE link (NN2 and NN3 below)

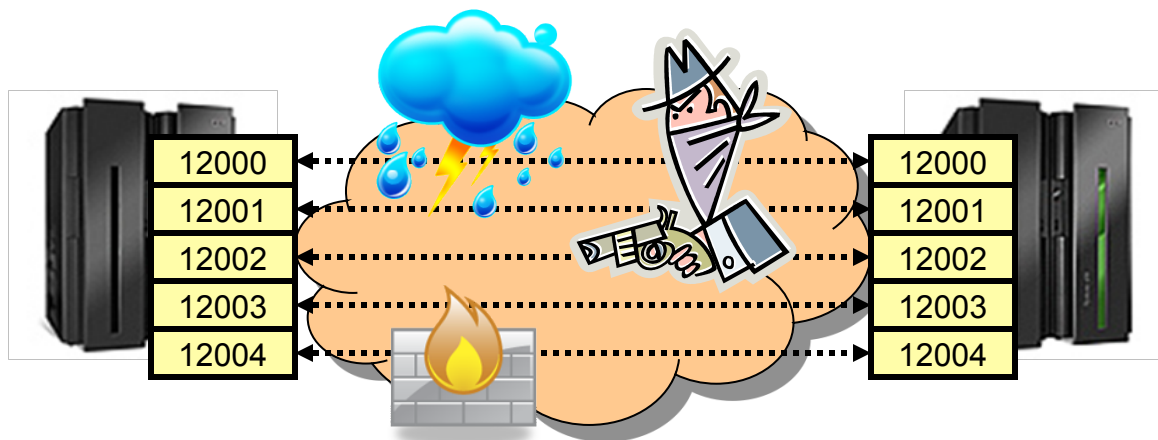


- If an existing RTP pipe begins on a node other than the EE link endpoint, it will not learn the PMTU-discovered MTU size, and will continue to send packets at a non-optimal size, potentially resulting in packet loss and transmission stalls.
- z/OS V1R12 adds logic for VTAM to drive the path switch logic if multiple retransmissions occur (stall detection)
  - Thereby letting NN1 above learn the new current MTU size and adapt

**IST2335I PATH SWITCH REASON: XMIT STALL RECOVERY**



## Enterprise Extender connection health verification (V1R12)



- Questions:
  - Are all five EE ports reachable at EE connection initialization point in time?
  - Do all five EE ports remain reachable?
- Apart from something not working correctly, you really do not know!
- z/OS V1R12 adds optional probing logic during EE connection initialization and during the lifetime of the EE connection to verify the health of the EE connection.
  - EEVERIFY=NEVER
    - Do not send any probes
  - EEVERIFY=ACTIVATE
    - Probe during connection initialization
  - EEVERIFY=timer-interval

## Enterprise Extender connection health verification – example (V1R12)

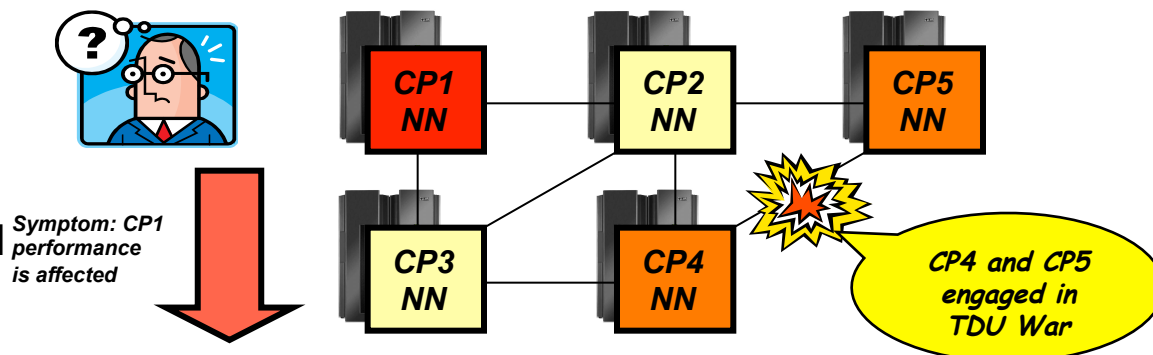
- To see all failed connections, issue the following command:

```
d net,ee,list=eeverify
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2000I ENTERPRISE EXTENDER GENERAL INFORMATION
IST1685I TCP/IP JOB NAME = TCPCS
IST2003I ENTERPRISE EXTENDER XCA MAJOR NODE NAME = XCAIP
IST2004I LIVTIME = (10,0)          SRQTIME =      15  SRQRETRY =      3
IST2005I IPRESOLV =      0
IST2231I CURRENT HPR CLOCK RATE = STANDARD
IST924I -----
IST2006I PORT PRIORITY =  SIGNAL          NETWORK          HIGH          MEDIUM          LOW
IST2007I IPPORT NUMBER =   12000          12001          12002          12003          12004
IST2008I IPTOS VALUE   =      C0          C0              80              40              20
IST924I -----
IST2324I EE HEALTH VERIFICATION: FAILED CONNECTION INFORMATION
IST2325I LINE LNIP1 PU SWIP2A1 ON 12/21/09 AT 15:56:39
IST2326I EE HEALTH VERIFICATION TOTAL CONNECTION FAILURES = 1
IST2017I TOTAL RTP PIPES =                1          LU-LU SESSIONS =                2
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS =                1
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS =                0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS =                0
IST2021I TOTAL ACTIVE EE CONNECTIONS =                1
IST314I END
```

## Enhancements to topology database diagnostics (V1R12)

- Enhancements in V1R11 defined a new control vector for TDU flows

- Topology Resource Sequence Number Update (x' 4E' ) control vector to identify node that set the RSN



- TDUDIAG start option available to control frequency of when new control vector is included
- Still required dumps and traces to diagnose TDU war
- z/OS V1R12 enhances various commands to improve the ability to better diagnose the TDU war scenario:
  - Enhance existing DISPLAY TOPO,LIST=TDUINFO output
  - New DISPLAY TOPO,LIST=TDUDIAG summary command
  - Diagnostic information from the Topology RSN Update control vector added in V1R11 is saved
  - New displays of diagnostic information from the x' 4E' control vector
    - DISPLAY TOPO,LIST=TDUDIAG command for a TG
    - DISPLAY TOPO,LIST=TDUDIAG command for a node

## 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



## 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



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

**New Requirement Map - Attacks**

Use this panel to indicate if you want attack protection

Enable attack protection

**Steps**

1. Select the action for each enabled attack type.
2. To disable protection for an attack type, select the row from the Enabled protection table and click the "Disable" button.
3. To enable protection for a specific attack type, select a row from the Attack type table and click the "Enable" button.

You will be prompted for additional details related to your attack type selection. Fill in the details and click OK.

**Attack type**

- Data Hiding Attack
- Flood Attack
- Global TCP Stall Attack
- ICMP Redirect Attack
- IPv4 Fragment Attack**
- IPv4 Options Attack
- IPv4 Outbound Raw Attack
- IPv4 Protocols Attack
- IPv6 Destination Options Attack
- IPv6 Hop-by-Hop Options Attack
- IPv6 Next Header Attack
- IPv6 Outbound Raw Attack

**Enabled protection**

Attack Type	Rule Name	Action
EE Malformed Packet Attack	EEMalformedPacket	Report Events
EE LDLC Check Attack	EELDLCCheck	Drop Packets or Connection
EE Port Check Attack	EESPortCheck	Both Drop and Report
<b>EE XID Flood Attack</b>	<b>EEXIDFlood</b>	<b>Report Events</b>

Buttons: Enable -->, <-- Disable, Modify..., Copy..., Advanced..., View Details...

Default Report Settings for Attacks...


Help ? < Back Next > Finish Cancel

---



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

**For pleasant reading ....**

## Please complete your session evaluation

- z/OS Communications Server Technical Update, Part 2
- Session # 11328
- QR Code:






Find us on Facebook at  
<http://www.facebook.com/IBMCommserver>

Follow us on Twitter at  
[http://www.twitter.com/IBM\\_Commserver](http://www.twitter.com/IBM_Commserver)

Read the z/OS Communications Server blog at  
<http://tinyurl.com/zoscsblog>

Visit the z/OS CS YouTube channel at  
<http://www.youtube.com/user/zOSCommServer>