Alfred Christensen -- alfredch@us.ibm.com Scott Moonen – smoonen@us.ibm.com

z/OS® V1R13 Communications Server – IBM Software Group, Enterprise Networking Solutions

# V1R13 Communications Server Overview

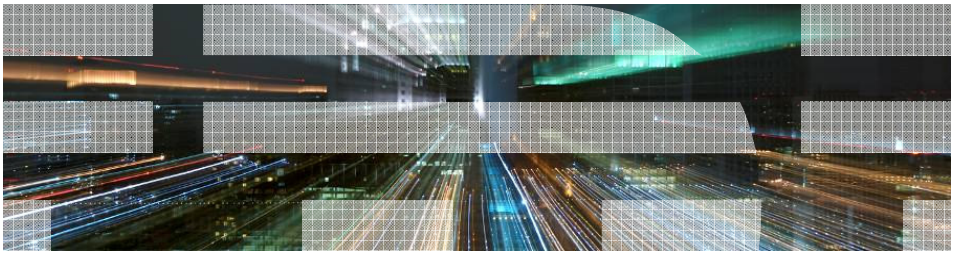



This presentation provides an overview of the new functions in z/OS V1R13 Communications Server.

# Agenda

**Introduction**

Security

Simplification

Economics and Platform Efficiency

Availability

Application/Middleware/Workload Enablement

EE/SNA

SOD and Reference

***z/OS V1R13 - planned GA: September 2011***

***Disclaimer: IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.***





Here is the Agenda. The first topic is a brief introduction.

## What will the z/OS community need from z/OS networking in 2011-2013?

- **System z technology is expected to continue to evolve**
  - Networking software needs to support new technologies such as zEnterprise
- **Access to System z system-level skills will continue to be an issue**
  - Existing skills are retiring
  - New people becoming responsible for the overall System z environment – including z/OS networking
- **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




In the near future, networking software needs to provide support for new System z technologies such as zEnterprise. In addition, System z skills will continue to be an issue as experienced people retire. Follow the IBM Academic Initiative https://www.ibm.com/developerworks/university/academicinitiative/. Security has been and continues to be a hot issue facing the z/OS community.

## What will the z/OS community need from z/OS networking in 2011-2013?

- **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 expected to follow 3Q2011**
  - IPv6 compliance (USGv6, IPv6-Ready, TAHI test suite, and so on)




There is a continued demand for reduced costs associated with System z and for autonomic integration. IPv6 utilization is expected to increase as the Internet Assigned Numbers Authority (IANA) has already run out of IPv4 addresses. Regional registries are expected to run out of IPv4 addresses in 2011 as well.

## Agenda

Introduction

**Security**

Simplification

Economics and Platform Efficiency

Availability

Application/Middleware/Workload Enablement

EE/SNA

SOD and Reference





The next topic is Security enhancements.

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

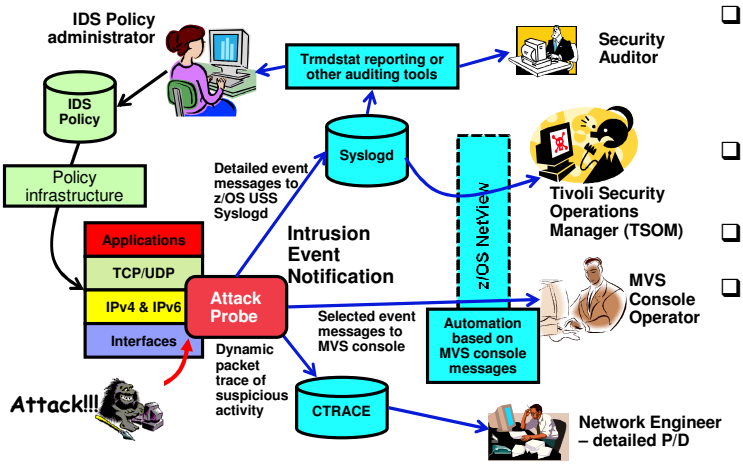


- **Events detected**
  - Scans
  - Attacks
  - Flooding
- **Defensive actions**
- **Reporting**
- **IDS Policy**
  - z/OS Comm Server Configuration Assistant

**The current IDS support is for IPv4 traffic only!**




Intrusion detection services (IDS) supports detection of scans, attacks and flooding. IDS can be configured to discard packets or limit connections. Events can be logged to syslogd, to the MVS console, to IDS packet trace and to Tivoli Security Operations Manager (TSOM). IDS is configured using policy and is supported by the Configuration Assistant. Currently IDS is only supported for IPv4 traffic.

## Intrusion Detection Services enhanced to include IPv6 traffic

- **Attack types supported for both IPv4 and IPv6**
  - Scan
  - Traffic regulation (TR)
  - Attack types extended to IPv6
  - Flood attacks
- **Plus, new attack types for IPv6-specific vulnerabilities**
  - Restricted IPv6 Next Header
  - Restricted IPv6 Hop Options
  - Restricted IPv6 Destination options
- **Note:** Defense Manager daemon already supports IPv6




The existing event detection for scans, traffic regulation, attacks and flood attacks is enhanced to include IPv6 traffic. TCP and UDP scan event rules now support IPv6 traffic. The ICMP scan event rule is unchanged and a corresponding ICMPv6 scan event rule is added. Scan exclusion lists can now include IPv6 addresses. TCP Traffic Regulation (TR) is enhanced to monitor IPv4 and IPv6 connection requests and UDP TR is enhanced to monitor IPv4 and IPv6 packets. The malformed packet event, UDP perpetual echo, and ICMP redirect restrictions attacks are extended to IPv6. IPv6 packets can be dropped due to malformed headers, options, or values. The SYN flood and interface flood attacks are extended to IPv6 as well. New attacks added for IPv6-specific vulnerabilities include Restricted IPv6 Next Header, Restricted IPv6 Hop Options, and Restricted IPv4 Destination options. The Defense Manager daemon already supports IPv6.

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

- **Global TCP Stall**
  - Detects when a large number of TCP connections are stalled and unable to send data
- **Data Hiding**
  - Detects data hidden in reserved fields
- **TCP Queue Size**
  - Detects TCP send, receive, and out-of-order queues that are storage constrained

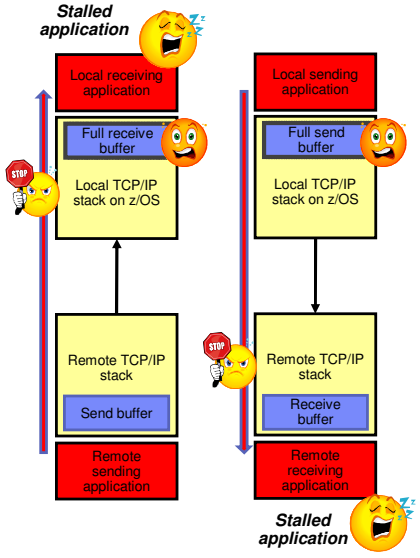






IDS implemented three new attack types for IPv4 and IPv6. The global TCP stall attack type prevents an attacker from creating connections with zero window sizes and keeping them open indefinitely. The data hiding attack type prevents an attacker from hiding data in reserved fields. This can be PadN options in IPv6 and reserved fields in IPv4 headers. The TCP queue size attack type helps you manage the amount of storage TCP can take up for the queues holding sent and received data. For example, out of order packets awaiting re-sequencing. It provides user control over storage constraint availability improvements added in z/OS V1R11. This helps avoid TCP storage constraint situations.

## 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, support was added for IKEv2, which was required for IPv6 currency
- Customers were encouraged to move to IKEv2, but for many, NAT is a requirement

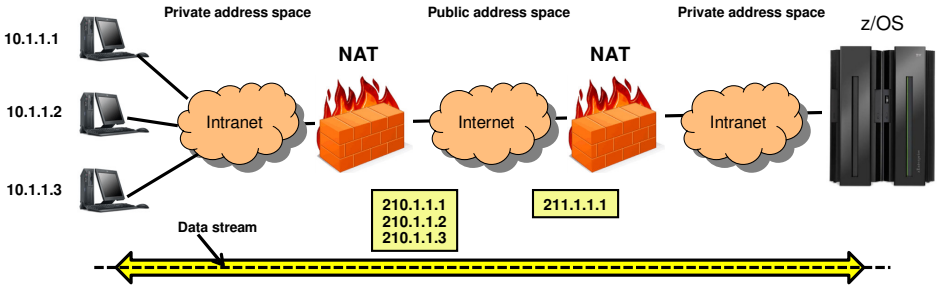





NAT is commonly used to conserve IPv4 addresses. IKEv2 support was added in V1R12 and supports both IPv4 and IPv6. Network address translation (NAT) is now supported when using IKEv2. You can now migrate from IKEv1 to IKEv2 if you are using NAT.

## Password phrase support in selected servers


- Password
  - One to eight characters
  - Limited range of characters allowed
- Password phrase
  - Nine to one hundred characters
  - Can contain most of the characters in the EBCDIC 1047code page
- Every user ID with a password phrase also has a password (since V1R10)
- Support for password phrases added to FTP and TN3270E in z/OS V1R13
  - TN3270E support is for solicitor screen only




The FTP and TN3270E servers have been updated to support password phrases. Passwords are one to eight characters in length and have a limited range of characters allowed. For example, a space is not allowed in a password. Password phrases extend the length to 100 characters and support most of the characters in the 1047 code page.

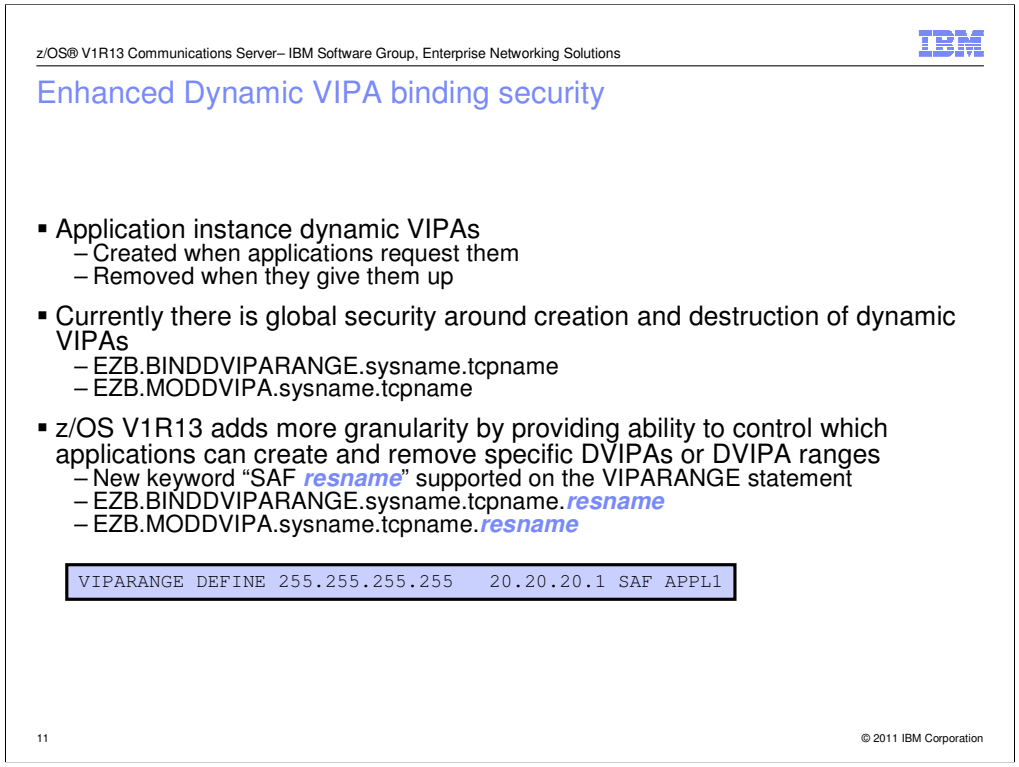


Application instance dynamic VIPAs are virtual IP addresses that are created when applications request them and removed when they give them up. They provide improved availability. For example a dynamic VIPA can move around in the sysplex, following the application when it moves, so clients are uninterrupted.

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

z/OS V1R13 adds more granularity by providing the ability to control which applications can create and remove specific DVIPAs or DVIPA ranges. This allows an application to create/remove its own DVIPAs but prevent it from interfering with other applications' ranges. A new keyword "SAF resname" is supported on the VIPARANGE statement. This identifies the resource profiles to use when creating or removing DVIPAs for the VIPARANGE statement. If the SAF keyword is not present, the existing profiles are used. In the example, to bind to 20.20.20.1, the application must be permitted to EZB.BINDDVIPARANGE.sysname.tcpname.APPL1. To issue MODDVIPA 20.20.20.1, the application must be permitted to EZB.MODDVIPA.sysname.tcpname.APPL1.

## Agenda

Introduction

Security

**Simplification**

Economics and Platform Efficiency

Availability

Application/Middleware/Workload Enablement

EE/SNA

SOD and Reference





The next topic is Simplification enhancements.

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

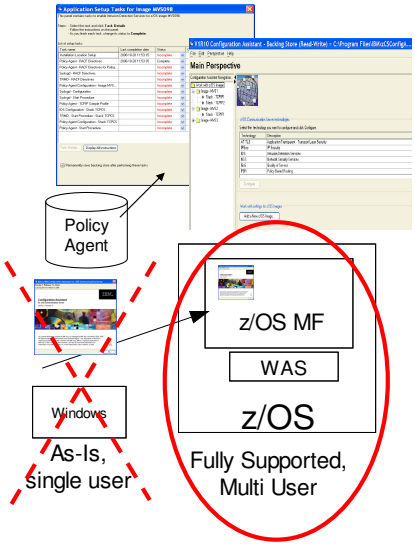


- IBM Configuration Assistant for z/OS Communications Server is integrated with z/OS Management Facility (z/OSMF)
  - Officially supported
- Stand-alone Windows version is still available
  - Available as-is - no official support




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

The stand-alone Windows version is still available, but is made available as-is, without any official support. It can be found at 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.

## Configuration Assistant support for multiple releases

- z/OSMF direction is for one instance in a sysplex
- Currently, Configuration Assistant only supports the release of Communications Server it ships with
- Makes life difficult for customers who have a mixed-release environment
- In V1R13, 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
  - You can now set the release of the operating system image
  - Release 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

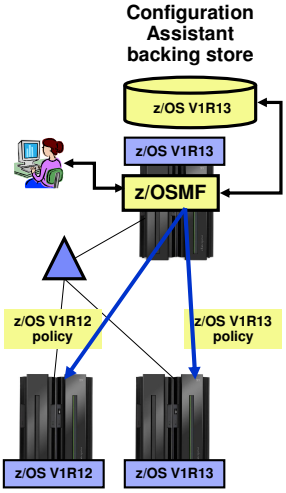





The z/OSMF direction is for a single instance of z/OSMF in a sysplex. For customers who have a mixed-release environment, this means that one instance has to manage multiple releases. Currently, the version of Configuration Assistant that ships with a release of Communications Server can only update 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 are supported in z/OS V1R13. When creating a new operating system image, you can now set the release of that image. The release of the operating system image can be changed at any time. This facilitates migration of images. Options that only apply to the higher-level release are ignored if the image is currently at the lower level. For example, new IDS attack types are ignored if the z/OS image is z/OS V1R12, but included if the image is z/OS V1R13.

## Configuration Assistant multiple stack policy configuration

- Currently security policy rules must be individually configured for each stack
- Often, a common security policy is needed throughout multiple stacks
  - For example, for a dynamic VIPA address which can move throughout the Sysplex
- Multiple stack policy configuration provides support for reusable rules
  - Reusable rules are created a single time and assigned to TCP/IP stacks
  - Single reusable rule updates are propagated to all stacks
  - Local address information in reusable rules defined using local address names
  - IKED identities in reusable rules defined using local names




Currently security policy rules must be individually configured for each stack in a Sysplex or other grouping. Many times, a common security policy is needed on each stack. For example, for a dynamic VIPA address which can move throughout the Sysplex. Currently, you must configure the rule for it on every host it can reside on. With z/OS V1R13, Sysplex-wide policy configuration allows for a rule to be configured once for a dynamic VIPA and pushed to every host it can reside on. This supports the Single System Image view for configuring Sysplex. Reusable rules are created a single time and assigned to TCP/IP stacks. If a reusable rule needs to be updated, only a single rule needs to be modified and the changes are propagated to all stacks. Local address information in a reusable rule is defined by way of local address names (a user-chosen symbol). Each stack that uses the reusable rule, defines the actual value of that symbol. IKED identities in reusable rules are also defined using local names.

## Reusable rules and named symbols - overview

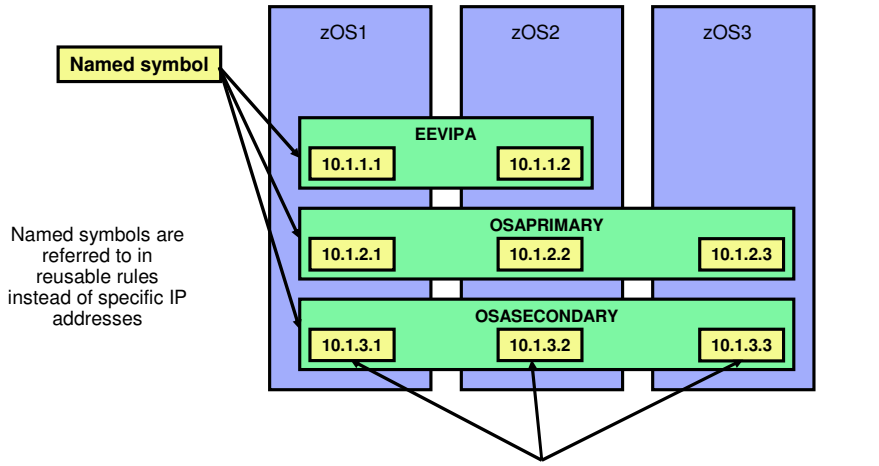





This diagram shows how reusable rules and named symbols can be used in the policy configuration for three different stacks. Reusable rules are created for all three stacks which contain named symbols. For each stack, the named symbols are configured with the values for the stack. Some of the named symbols can be dynamically configured by the Configuration Assistant.

## Configuration Assistant import of TCP/IP configuration information

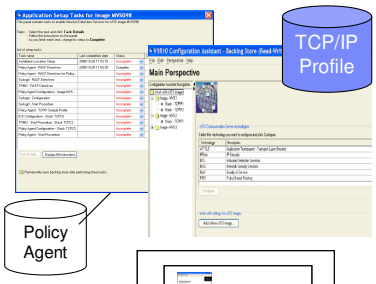


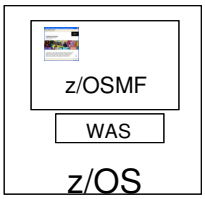


Fully Supported, Multi User

- In z/OS V1R13, Configuration Assistant will import profile information from running TCP/IP stacks
- Long-term goal is to be able to configure all of Communications Server with the Configuration Assistant
- z/OSMF version only





For z/OS V1R13, Configuration Assistant will import profile information from running TCP/IP stacks. This information can be used to help develop policy configuration. For example, you can discover home IP addresses and it will suggest address groups. The long-term goal is to be able to configure all of Communications Server with the Configuration Assistant. The ability to import TCP/IP configuration information is the first step in that direction. This support is only provided for the z/OSMF version of the Configuration Assistant – support is not provided in the as-is Windows version.

## Agenda

Introduction

Security

Simplification

**Economics and Platform Efficiency**

Availability

Application/Middleware/Workload Enablement

EE/SNA

SOD and Reference





The next topic is Economics and Platform Efficiency enhancements.

## 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
- Raise limit from 8 to 32 VLANs per stack per OSA port
  - No impact to OSA
  - VLANs used for network isolation in the zEnterprise IEDN

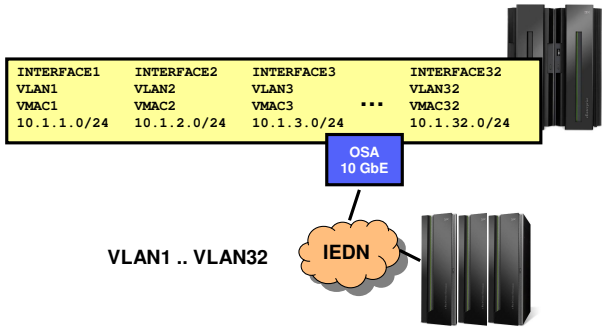






In z/OS V1R10, Communications Server added multiple VLAN support. Up to eight IPv4 and eight IPv6 VLANs are allowed per OSA port. A separate INTERFACE statement and data device are required per VLAN. The value of eight is a z/OS Communications Server software limitation. In z/OS V1R13, the limit is raised to 32 VLANs per stack per OSA port. There is no impact to OSA. This enhancement is in support of the emphasis on VLANs for network isolation in the zEnterprise IEDN.

## Additional 64-bit exploitation

- Traces relocated to 64 bit common storage
  - VTAM internal trace (VIT) moved from ECSA and data space to 64 bit common
  - Multiple CTRACE components moved from data spaces to 64 bit common

| CTRACE Component | Current location | z/IOS V1R13 change | User |
|---|---|---|---|
| SYSTCPIP | TCPIPDS1 Dataspace | 64 bit common | Stack |
| SYSTCPIP | TN3270E Private storage | 64 bit private | TN3270E |
| 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 |




Multiple trace buffers have been relocated to take advantage of 64 bit common storage. The VTAM internal trace (VIT) is moved from ECSA to 64 bit common storage. In addition, the VIT data space is no longer used. This move is transparent to you if you use the external VIT to obtain trace records.

For TCP/IP, multiple CTRACE components have been moved from data spaces to 64 bit common storage. The table summarizes the changes. These moves are transparent to you as long as you use the NMI interface to obtain trace data.

## z/OS virtual storage map

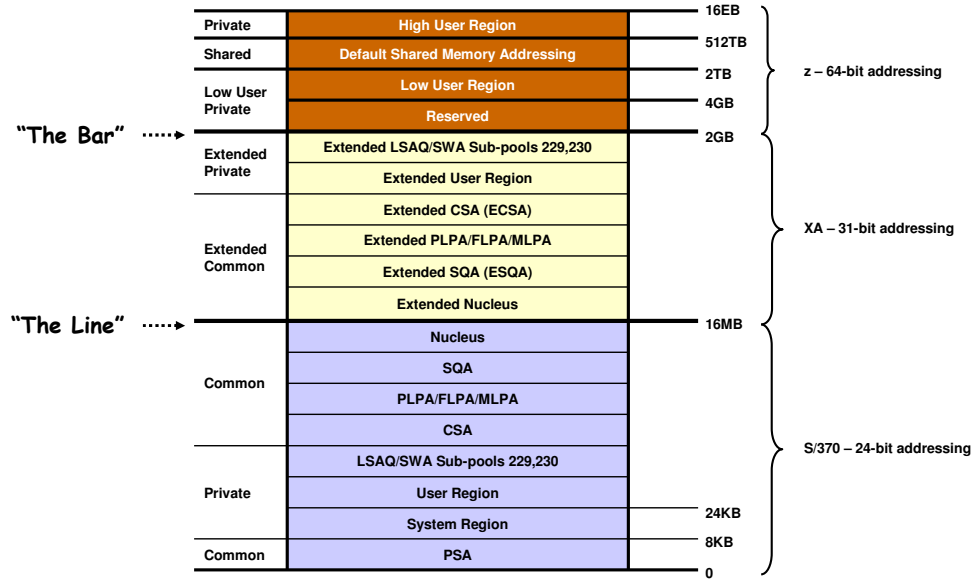






This slide shows a diagram of z/OS virtual storage. About 20 years ago, the focus was to move z/OS Communications Server storage utilization above the 31-bit “line” (16 MB). Now the attention is on moving storage utilization above the 64-bit “bar” (2 GB). In particular, reducing z/OS Communications Server’s ECSA footprint has been an area of focus for several releases.

## Miscellaneous TCP/IP performance improvements

- Communications Server improves performance and throughput in every release by focusing on software pathlength
- Goal is 5% performance improvement in TCP/IP per release
- Examples of improvements for z/OS V1R13:
  - Increased use of compiler optimization
  - Modified layout of internal data structures for better caching
  - Use of separate IPv4 and IPv6 modules
  - Improved TN3270 performance for long data streams (> 90 bytes)

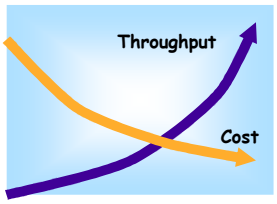





Communications Server development strives to improve performance and throughput in every release by focusing on software pathlength. The goal is at least a 5% performance improvement in TCP/IP per release. Some of the improvements for z/OS V1R13 are listed above.

## Agenda

Introduction

Security

Simplification

Economics and Platform Efficiency

**Availability**

Application/Middleware/Workload Enablement

EE/SNA

SOD and Reference





The next topic is Availability enhancements.

## Resolver autonomics for unresponsive name servers

- z/OS V1R12, resolver monitors name servers for responsiveness to queries
  - Network operator notification
  - New queries still sent to unresponsive name servers
  - No configuration required
- z/OS V1R13, can stop sending queries to unresponsive name servers
  - Network operator notification
  - Polls unresponsive name server to detect when responsive again
  - Configuration required

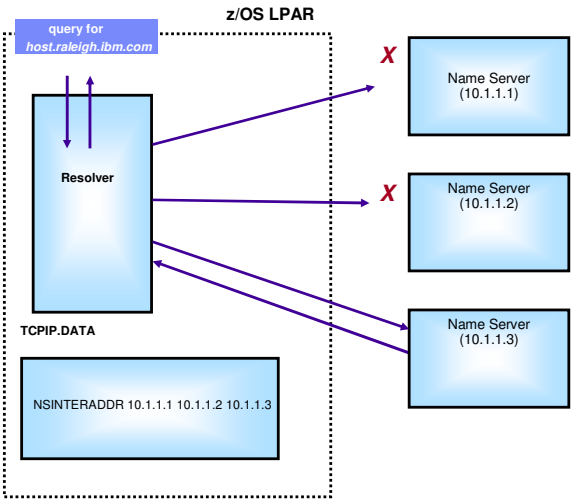






In z/OS V1R12, the resolver monitors name servers for responsiveness to queries. The network operator receives a notification when a name server becomes unresponsive. Although the resolver detected the unresponsive name server, new queries were still sent to that name server. This function is enabled by default.

In z/OS V1R13, the resolver can be configured to stop sending queries to unresponsive name servers. The resolver polls the unresponsive name server to detect when it becomes responsive again. The operator is notified of the condition using messages similar to those used in V1R12. The autonomic quiescing function must be explicitly enabled in the resolver setup file.

## CSSMTP enhanced send error recovery

- CSSMTP sends batch email to the internet from z/OS JES spool files
  - For example, paperless customer statements
- Messages retried if target relays fail to acknowledge mail
  - JES spool files and CSSMTP memory retained until retries completed
- z/OS V1R13 provides file system storage of messages being retried for an extended interval
  - JES spool files and CSSMTP memory can now be released

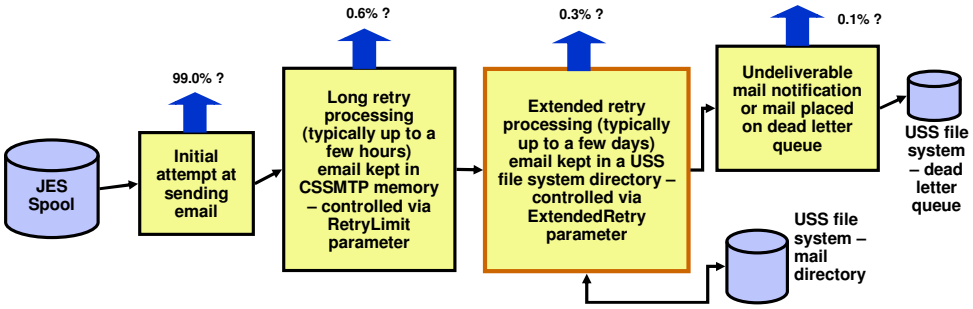





CSSMTP sends batch email to the internet from z/OS JES spool files. If the target relays fail to acknowledge mail, CSSMTP will retry for the configured interval up to five days then drop the message, and return an undeliverable notice. However, spool files cannot be deleted until all messages in the spool are delivered. A spool file can 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. It will continue to retry from memory and spool until the initial retry limit is reached. A new parameter can be configured to indicate how long beyond the existing interval to retry from the file system.

## CSSMTP translate code page support enhanced

- '@' character in SMTP mail messages MUST be translated correctly!
- CSSMTP currently supports only a fixed set of single byte code pages
- In V1R13, TRANSLATE statement allows additional code pages
  - Must be an EBCDIC code page
  - Code page must support roundtrip translations between IBM-1047 and ISO-8859-1 code pages

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




The '@' character has special meaning in SMTP mail messages. It 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. In z/OS V1R13, CSSMTP has been updated to allow the TRANSLATE statement to define additional code pages. Also, CSSMTP will allow the specification of a code page by its CCSID. Before z/OS V1R13, code pages must be a character string “IBM-XXXX”, where XXXX is a subset of possible code pages. In z/OS V1R13, CSSMTP has been updated to expand the list of supported code pages. Also, CSSMTP will 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), which are 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

- VIPAROUTE target cache is used to route a Sysplex Distributor packet
  - Updated every 60 seconds
- z/OS V1R13 shortens the interval for VIPAROUTE route lookups when the stack joins a Sysplex, or OMPROUTE is restarted
  - Will now start with five seconds, and gradually increase to 60 seconds

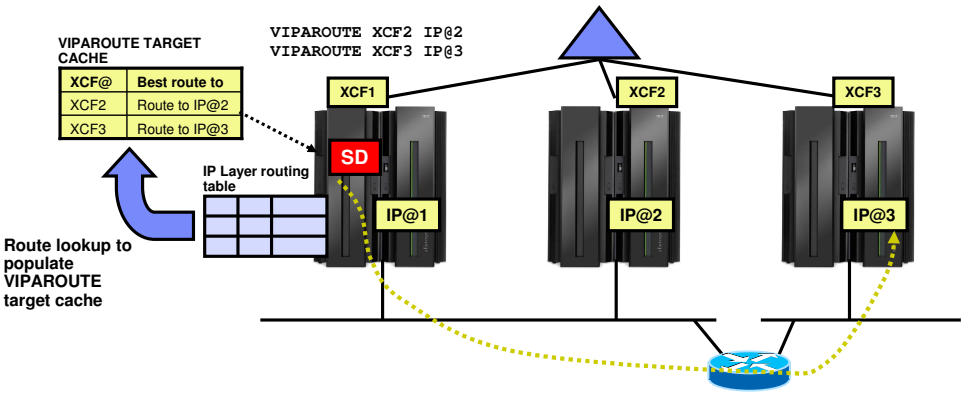





With VIPAROUTE support, distributed DVIPA traffic forwarded to a target stack can use any IP routing interface. An IP destination address is configured for each target XCF address. After the configuration profile is processed, normal IP routing rules are used to pick the best available route to that destination (IP routing table or policy-based routing). The IP route selected is cached for each target. Once a route to a target is cached, new connection requests distributed to that target use the cached target route. Periodically (every 60 seconds), a target cache refresh routine checks to see if the routing table has changed. If it has changed, a new route lookup is done to update each target's routing cache. If no route exists, the target is no longer considered as a distribution target; there is no fall back to use an XCF interface.

During a takeover that occurs because the primary routing stack is restarted, a route is unavailable until its interface finishes activating. If the initial target cache refresh occurs before this route is available, then there is a delay of close to 60 seconds until the next refresh. During this time, packets for distributed connections using VIPAROUTE might use non-optimal routes or might not be able to reach their targets.

This function substantially reduces this initial delay before acquiring a preferred VIPAROUTE by starting with smaller target cache refresh intervals. The initial refresh pattern is five, five, 15, 35, and 60 seconds. After reaching 60 seconds, the refresh interval will remain at 60 seconds. The initial refresh pattern is used whenever the stack joins the sysplex group or OMPROUTE is recycled.

## Sysplex autonomics extended with CSM storage constrained monitoring

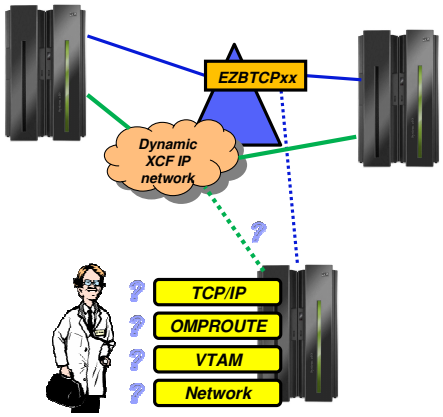


- **Monitoring:**
  - Comm Server health indicators
  - Dependent networking functions
  - Abends in Sysplex-related stack components
  - Selected network interface availability and routing
  - Repetitive internal abends in non-Sysplex related stack components
  - ***Detect when CSM FIXED or CSM ECSA has been constrained for multiple monitoring intervals***
- **Actions:**
  - Remove the stack from the IP Sysplex (manual or automatic)
  - Reactivate the currently inactive Sysplex configuration when a stack rejoins the Sysplex (manual or automatic)

***Sick? Better remove myself from the IP Sysplex!***

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





In V1R13, z/OS Communications Server more closely monitors common storage (CSM) usage and performs the recovery action if CSM is constrained (higher than 80% of defined CSM limit). Previous support only performs the recovery action when CSM is critical (higher than 90% of CSM limit). This reduces the impact of sysplex processing on other z/OS processes.

## Sysplex-wide Security Associations (SWSA) for IKEv2

- Sysplex-wide Security Associations (SWSA) allows IPSec protected traffic to be distributed through a sysplex and maintain full security to all sysplex endpoints
  - Security associations and characteristics (for example encryption) are moved and distributed through the sysplex with the target programs
- Support is extended for IKEv2
- IPv4 only

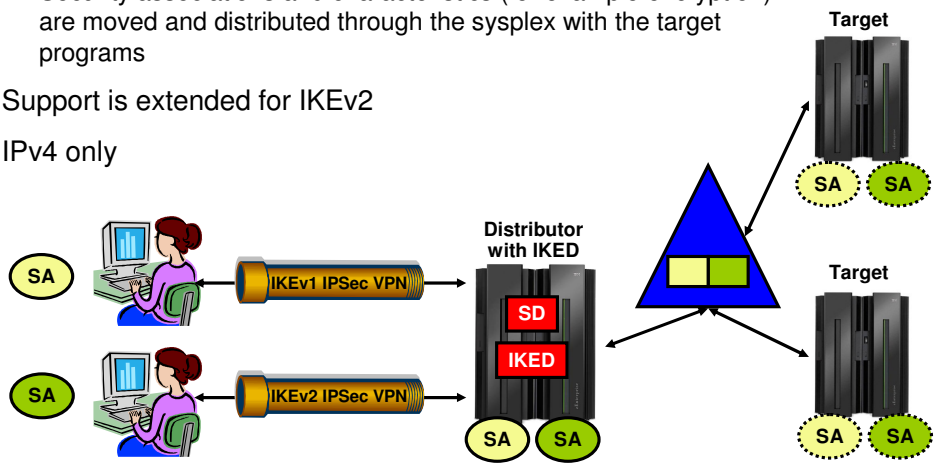





Sysplex Distributor negotiates security associations (SA) with the remote Client using the Internet Key Exchange protocol, IKE version 1 or IKE version 2. It sends copies of the SAs (shadows) to the targets for VPNs negotiated with either version of IKE. Targets use the SAs negotiated with either IKEv1 or IKEv2 to encrypt and decrypt data. The backup can recover SAs negotiated with either IKEv1 or IKEv2 in case of DVIPA takeover. The coupling facility stores shared data for SAs negotiated with either IKEv1 or IKEv2.

## Agenda

Introduction

Security

Simplification

Economics and Platform Efficiency

Availability

**Application/Middleware/Workload Enablement**

EE/SNA

SOD and Reference





The next topic is Application, Middleware and Workload Enablement enhancements.

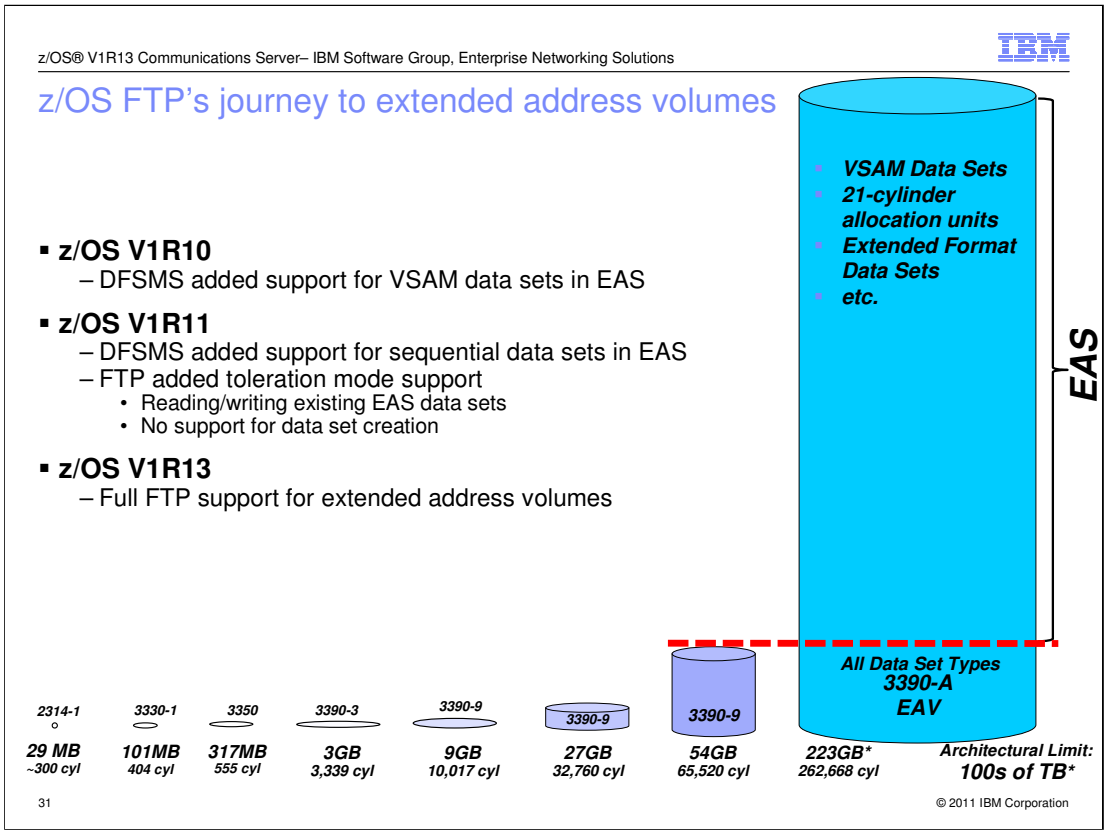


In z/OS V1R10, DFSMS added support for VSAM data sets in the extended addressing space (EAS). FTP does not support VSAM data sets, so there was no impact. In z/OS V1R11, DFSMS added support for extended format sequential data sets eligible to reside in the EAS. FTP added support for reading and writing existing EAS data sets, but not for creating them (toleration mode). In z/OS V1R13, FTP provides full support for extended address volumes (EAV) without requiring use of an SMS data class. Before z/OS V1R13, you can use an EAV for SMS managed datasets if set up properly in an SMS data class.

## FTP support for large format datasets

| | BASIC FORMAT | LARGE FORMAT | EXTENDED FORMAT |
|---|---|---|---|
| **DSNTYPE parameter value** | BASIC | LARGE | EXTPREF<br>EXTREQ |
| **Maximum tracks per volume** | 65,535 | 16,777,215 | Larger than 16,777,215 |
| **Maximum extents per volume** | 16 | 16 | 123 |
| **Why choose this format?** | Maximum compatibility. | Can be much larger than basic format. | Can be much larger than basic format.<br>Can be striped, compressed format or any combination. |

- **z/OS V1R7**
  - DFSMS added support for large format data sets
- **z/OS V1R13**
  - LE support for large format data sets completed
  - FTP added support for large format data sets
    - Reading/Writing/Creating/Restarting large format datasets
    - Reading/Writing/Creating/Restarting z/OS UNIX files > 2G





Large format physical sequential datasets can have more than 64K tracks per volume. They do not have to. They 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. The access methods supported are BSAM, QSAM, and EXCP. Language Environment (LE) Runtime Library large format dataset support completed in z/OS V1R13 now enables z/OS FTP support for these datasets. FTP support means sending from large format data sets and allocating and storing into large format data sets. In block mode, you can use the restart subcommand to resume the failed transfer to and from large format data sets. Support is also added for transfers to and from z/OS UNIX files larger than two gigabytes.

## Command to display all TN3270 servers

- New D TCPIP,TELNET console command
  - Displays the list of TN3270E servers that are or have been active on the system
- 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 ***
```




A new D TCPIP,TELNET console command has been 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!

## Retrieving System Resolver configuration data through NMI

- MODIFY RESOLVER,DISPLAY command displays contents of resolver setup file
- Currently no command to display contents of global TCPIP.DATA file
- New resolver autonomics for unresponsive name servers function in z/OS V1R13 requires global TCPIP.DATA file
- z/OS V1R13 Communications Server implements a resolver callable NMI (EZBREIFR)
  - Single request type:
    - GetResolverConfig
  - Returns:
    - The contents of the resolver configuration file and
    - The contents of the global TCPIPDATA file, if it's in use




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 file. Before z/OS V1R13, you must collect Trace Resolver output to see the global TCPIP.DATA settings dynamically. The new resolver function in z/OS V1R13 to monitor and quiesce unresponsive name servers, depends on a global TCPIP.DATA file. Only name servers specified in this global TCPIP.DATA file are monitored. z/OS V1R13 Communications Server implements a resolver callable NMI (EZBREIFR). The NMI provides a high-speed, low-overhead callable programming interface for network management applications to access data related to the resolver. There is a single request type (GetResolverConfig) which returns the contents of the resolver configuration file and the contents of the global TCPIP.DATA file, if it is in use.

## TMI copy buffer interface improvement

- Currently TMI copy buffer interfaces require the caller to be APF authorized
  - Causes problems for applications that want to fork()
  - APF authorization gives application broad range of authority
- In z/OS V1R13, applications that use the TMI copy buffer interfaces will no longer have to be APF authorized
  - Permit application to SAF resource EZB.NETMGMT.systemname.tcpprocname.*interface* where *interface* can be:
    - SYSTCPDA, Real-time TCP/IP packet and data trace NMI
    - SYSTCPCN, Real-time TCP connection SMF NMI
    - SYSTCPOT, Real-time OSAENTA packet trace NMI
    - SYSTCPSM, Real-time SMF NMI




The TCP/IP Management Interface (TMI) copy buffer interfaces currently require the caller to be APF authorized. The interfaces are EZBTMIC1, EZBTMIC4, and TMI_Copybuffer(). This causes problems for applications that want to fork() as APF authorization is not inherited. In addition, APF authorization gives applications a 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 application is running under can be authorized to the appropriate SAF resource EZB.NETMGMT.systemname.tcpprocname.interface. The interface can be SYSTCPDA, SYSTCPCN, SYSTCPOT, or SYSTCPSM.

## PORTRANGE wildcard option

- z/OS V1R5 added wildcard support for the jobname on a PORT statement
  - 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
- In z/OS V1R13, wildcard support is added for the jobname on a PORTRANGE statement
  - Only applications with job names that match the prefix have access to the specified port range
- TCP callable NMI GetProfile request can now include a wildcard value for PORTRANGE statement entries

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





A wildcard option has been allowed for the job name value specified on a PORT statement since z/OS V1R5. The PORT statement allows a job name to end in an asterisk. The characters before the asterisk define a prefix. Any applications with a job name matching the prefix can access the specified port. Similar wildcard support was not provided for the PORTRANGE statement. The PORTRANGE statement in z/OS V1R13 allows a wildcard job name. As with the PORT statement, the job name can end in an asterisk. The characters before the asterisk 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 TCP callable NMI (EZBNMIFR) can now include a wildcard value in the NMTP_PORTJobName field for entries that represent a PORTRANGE statement.

## Agenda

Introduction

Security

Simplification

Economics and Platform Efficiency

Availability

Application/Middleware/Workload Enablement

**EE/SNA**

SOD and Reference





The next topic is Enterprise Extender and SNA enhancements.

## Enterprise Extender firewall-friendly connectivity test

- DISPLAY EEDIAG,TEST=YES provides information about an EE partner and all routers in between
  - Firewalls blocking ICMP messages can cause long delays for results
- LIST=SUMMARY will now provide quick test of partner reach-ability
  - Hop count determination omitted
- LIST=DETAIL is unchanged

```
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
```





The DISPLAY EE DIAG TEST=YES command provides information about an Enterprise Extender partner and all of the routers in between. If a firewall in between is blocking ICMP messages, there can be a long delay before getting results. The delay is because of timeouts waiting for ICMP messages that never come. The delay is the number of router hops past the firewall times nine seconds. The existing LIST=SUMMARY option now provides a quick test of partner reach-ability. The probe is sent to the partner with TTL=255 so it does not probe any intermediate hops. The output no longer includes the hop count determination. The LIST=DETAIL output is unchanged and includes intermediate hop information and the hop count determination.

## Intrusion Detection Services for Enterprise Extender traffic

- Implement new EE-specific IDS attack types
  - EE Malformed Packet
  - EE LDLC Check
  - EE Port Check
  - EE XID Flood
- Exclusion list allowed for each attack type
- Actions are discard and notify
- IDS policy - z/OS Comm Server Configuration Assistant
- IPv4 and IPv6

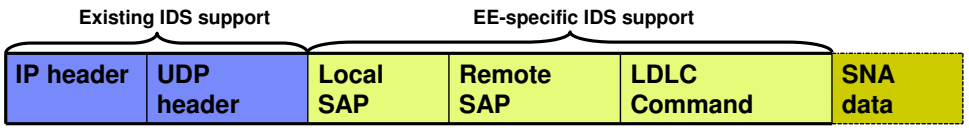





Intrusion detection services (IDS) is enhanced to implement four new IDS attack types for Enterprise Extender (EE). These attack types are supported for IPv4 and IPv6 EE traffic. The EE Malformed Packet attack type checks inbound EE packets for incorrect lengths. The EE LDLC Check attack type checks that inbound LDLC control commands are only received on the signaling port (12000). The EE Port Check attack type checks that inbound EE packets contain matching source and destination ports. The EE XID Flood attack type checks if a threshold is met for inbound XIDs within one minute. The actions allowed are to discard the packet and to provide a notification. The EE XID Flood attack only supports the notify action. An exclusion list can be created to exclude specific hosts from attack checking. Events notifications can be sent to syslogd, to the console, to IDS packet trace and to Tivoli Security Operations Manager (TSOM). IDS is configured using policies and is supported by the Configuration Assistant.

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

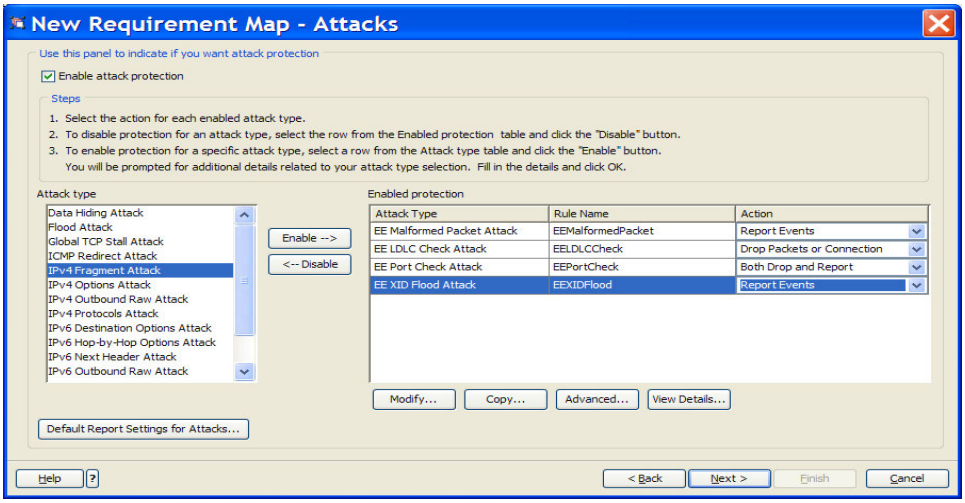





The Configuration Assistant was updated to support the four new EE attack types. The window here shows a new requirement map with the four EE attack types listed. You can disable an attack type by selecting it and clicking the disable button.

## Improved APPN routing resilience

- APPN session routes are selected by APPN Topology and Routing Services (TRS)
- Optimal route determined by specified criteria
- TRS builds routing trees to determine best route
- Recursive abends can occur if the routing tree is compromised
- Recovery routine for route selection abends is improved to prevent recursive abends
  - Routing tree is deleted
  - Dynamically rebuilt
- New FUNCTION=CLRTREES operand on MODIFY TOPO command

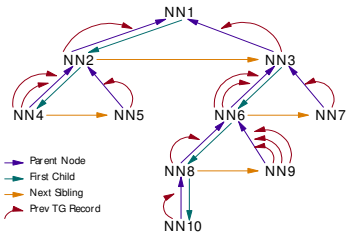





APPN session routes are selected by APPN Topology and Routing Services (TRS). These routes are used to locate resources (directed search routes) and for LU-LU sessions. The optimal route is determined by specified criteria such as line speed, cost of data transmitted, security, and user-defined values. TRS builds complex routing trees to determine the best path. The tree records represent nodes along preferred routes. Recursive abends can occur if a pointer in a routing tree is compromised. Every time a session route is requested using that tree, abends occur and the session fails. VTAM must be restarted to recover. The 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 rebuilding it. Existing sessions are not affected. It can temporarily impact VTAM performance in large APPN networks as the routing tree is being rebuilt. 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.

## Agenda

Introduction

Security

Simplification

Economics and Platform Efficiency

Availability

Application/Middleware/Workload Enablement

EE/SNA

**SOD and Reference**





The next topic is Enterprise Extender and SNA enhancements.

## 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).***




The next topic is Enterprise Extender and SNA enhancements.

## For more information

| URL | Content |
|---|---|
| http://www.twitter.com/IBM_Commserver | IBM Communications Server Twitter Feed |
| http://www.facebook.com/IBMCommserver | IBM Communications Server Facebook Page |
| 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.ibm.com/software/network/commserver/z_lin/ | IBM Communications Server for Linux on System z |
| http://www.ibm.com/software/network/ccl/ | IBM Communication Controller for Linux on System z |
| http://www.ibm.com/software/network/commserver/library/ | IBM Communications Server library |
| 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, and so on) |
| 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 |

***For pleasant reading ….***





This page provides a series of URLs for web pages where additional information on z/OS Communications Server can be found.