z/OS V1R13

z/OS UNIX: IPV4 pktinfo support

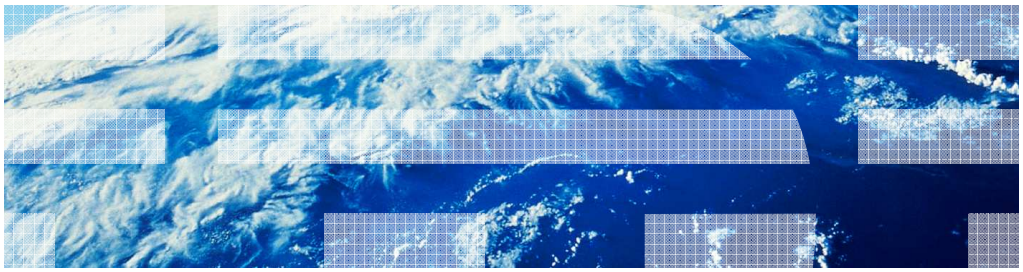

Session objectives

- Define the functional content and benefit
- Explain the external changes
- Explain the expected usage protocol
- Explain any migration issues or concerns
- Indicate list of Publications and References

Overview

- **Problem Statement / Need Addressed**
  - Currently under CINET environment with multiple TCPIP stack configuration, when a server system has multiple home addresses with multiple routes back to the client, the UDP reply packet might not flow on the same interface where the UDP request packet arrived.
- **Solution**
  - Provide a new z/OS IPv4 external interface to obtain request's inbound interface information and use it on the reply.
  - Similar to the existing IPv6 external interface
- **Benefit / Value**
  - Server can send the UDP reply packet to a client request out on the same inbound physical interface that the client's UDP request packet arrived

The server is receiving requests that were sent to all these local addresses and currently no correlation exists in AF_INET UDP datagram sockets between particular request packet and the subsequent reply packet. With multiple IP addresses for the server's UDP reply packet, the server system may not pick the interface used by the client to send the UDP request packet.

This is critical in security conscious environments where clients will not trust the replies that don't seem to come from the same server where the request was sent to.

Usage and invocation (1 of 2)

- **Expected IPv4 IP_PKTINFO Usage Protocol:**
  - Application requests inbound interface info to be part of recvmsg()
    - Use SetsockOpt() with new IP_RECVPKTINFO option
  - TCP/IP stack includes the inbound interface info in a new IN_PKTINFO structure as an ancillary data item on the recvmsg()
  - Returned inbound interface info is used untouched on the subsequent sendmsg()
- **Note:** The program must to call setsockopt() with IP_RECVPKTINFO option to have the TCPIP stack pass the client's return information in a IN_PKTINFO structure as ancillary data item on the recvmsg() call, and that IP_PKTINFO data is used, untouched, on the subsequent sendmsg() to have the reply flow out the same interface the request arrived.

Note that the z/OS unique IPv4 external interface is non-portable since it is not stated in the standard specification.

On the other hand, the IPv6 external interface is stated in the standard specification and the IPv4 external interface is modeled after the IPv6 specification. Other platforms have extended the IPv6 IP_PKTINFO external interface outside of the standard specification to IPv4 environment as well.

Usage and invocation (2 of 2)

- **IPv4 IP_PKTINFO external Interface:**
  - **Ancillary data type for the new IN_PKTINFO mapping structure:**
    - **IP_PKTINFO EQU 101**
  - **Setsockopt() call option to receive inbound interface information as ancillary data on recvmsg() call:**
    - **IP_RECVPKTINFO EQU 102 (IPPROTO_IP Level option)**
  - **Ancillary data mapping structure:**
    - **IN_PKTINFO DS 0F**
    - **IPI_ADDR DS CL4 IPv4 Address**
    - **IPI_IFINDEX DS F Interface index**

Following is defined in the netinet/in.h header:

Ancillary data type for the new IN_PKTINFO mapping structure:

```
#define IP_IPKTINFO        101
```

Setsockopt()/getsockopt() call option to receive inbound interface information as ancillary data on recvmsg() call:

```
#define IP_RECVPKTINFO     102
```

Ancillary data mapping structure:

```
struct in_pktinfo {
struct in_addr ipi_addr;        /* src/dst IPv4 address          */
unsigned int ipi_ifindex;       /* send/recv interface index */
};
```

Interactions and dependencies

- Software Dependencies
  - None.
- Hardware Dependencies
  - None.
- Exploiters
  - z/OS NFS Server

Migration and coexistence considerations

- None.

Installation

- None.

Session summary

- Provides the capability to obtain the AF_INET UDP request's inbound interface information, and allows to use it on the reply UDP packets.

Appendix - References

- **Publications References:**
  - SA22-7821-13 z/OS V1R13.0 XL C/C++ Run-Time Library Reference
  - SA22-7803-14 z/OS V1R13.0 UNIX System Services Programming: Assembler Callable Services Reference

Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "Copyright and trademark information" at http://www.ibm.com/legal/copytrade.shtml

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY.
Other company, product, or service names may be trademarks or service marks of others.

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.