IBM eServer™

Systems management

Miscellaneous enhancements

ebusiness on demand software

Agenda - Systems management

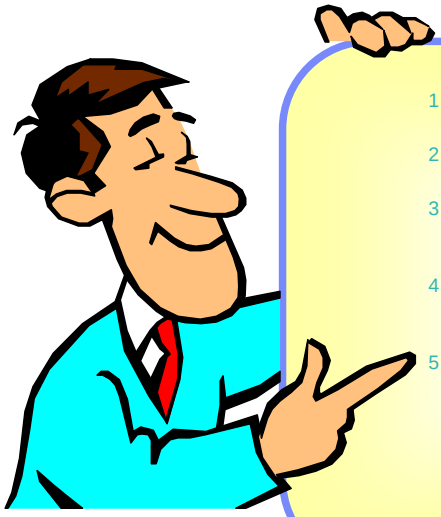

1. **SNMP enhancements**
2. **ARP/ND Takeover Visibility Enhancements**
3. **Allow wildcards on Enterprise Extender network management interface filters**
4. **Network management interface support for dropping socket endpoints**
5. **Provide reason code on TCP connection termination SMF record**

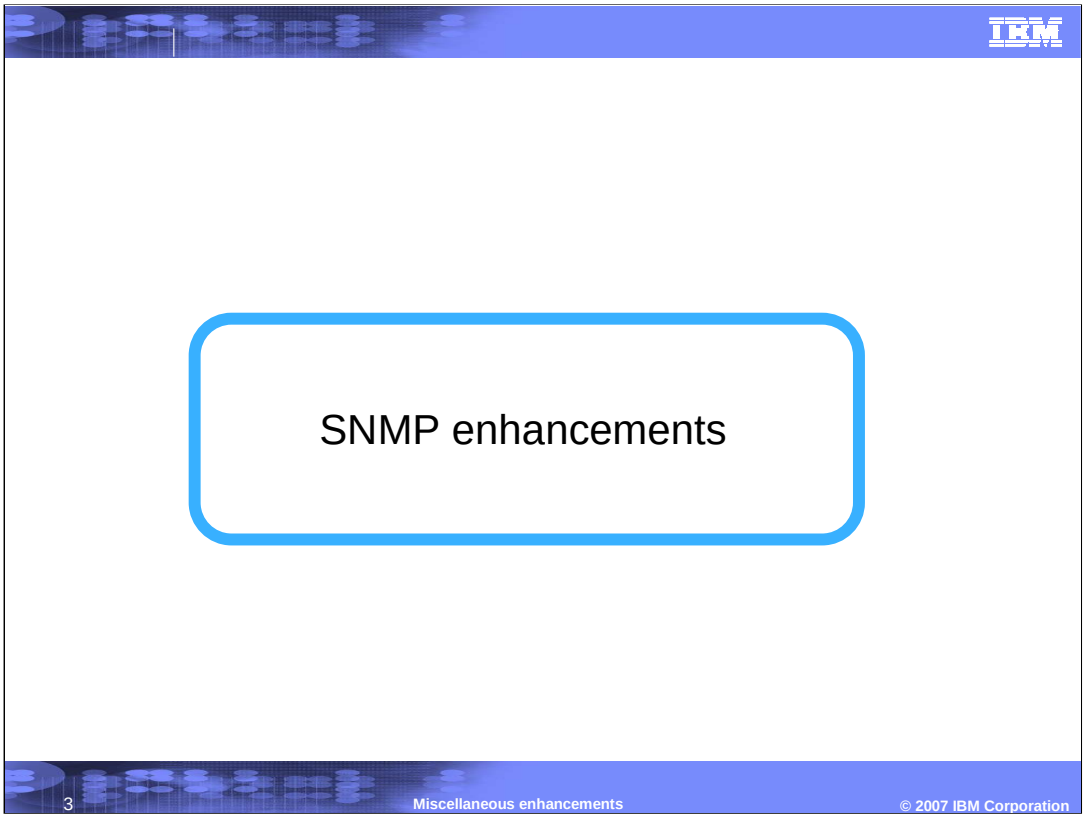
SNMP enhancements
3
Miscellaneous enhancements
© 2007 IBM Corporation

New SNMP MIB objects

- **TCP layer accepted connections counters**
  - Added the following 32-bit and 64-bit TCP layer accepted connection counters to the IBM MVS™ TCP/IP enterprise-specific MIB module:
    - **ibmMvsTcpAcceptCount** - 32-bit counter
    - **ibmMvsTcpHCAcceptCount** - 64-bit counter
  - These new counters represent the total number of connections accepted by the TCP/IP stack since the last time the stack was started.
  - Similar to other per stack, SNMP TCP counters (for example, tcpPassiveOpens, tcpCurrEstab)
- **CPC identifier**
  - Added new MIB object, **ibmMvsCpcNd**, to the IBM MVS TCP/IP enterprise-specific MIB module
    - The CPC is a physical collection of hardware (such as an System z9™ server) that consists of main storage, one or more central processors, timers, and channels.
    - Management applications wanted to be able to tell if applications from which they retrieved management data were running on the same CPC.
  - Value is CPC Node Descriptor, in hexadecimal, for the CPC on which the TCP/IP subagent is running, for example:
    - f0f0f2f0f8f4c4f3f2c9c2d4f0f2f0f0f0f0f0f0f0f9f9f9c1c1fff0

New SNMP MIB objects

NOTES

- Description of CPC Node Descriptor value
  - The CPC Node Descriptor value consists of the following information:
    - 002084 Type number
    - D32 Model number
    - IBM Manufacturer
    - 02 Plant of Manufacture
    - 0000000999AA Sequence Number
    - x'FFF0' Tag
  - The CPC Node Descriptor value is also displayed by the D M=CPU MVS operator's command

Message EZZ6317I - do not change sysObjectId

- **Message EZZ6317I**
  - Values for some MIB objects supported by the Agent can be configured by specifying OSNMPD.DATA information
    - sysObjectId is one of these MIB objects
  - The sysObjectId MIB object should not be set to other than its default value because it identifies the z/OS® CS Agent to network management applications
  - Message EZZ6317I was introduced in z/OS CS V1R4 to warn against configuring the sysObjectId value.
    - The ability to configure this value may be removed in a future release
  - Message EZZ6317I was only written to the syslog daemon, so it was not that visible to customers
  - Message EZZ6317I now written to the console and the syslog daemon
  - Warning more visible to those customers currently configuring this MIB object

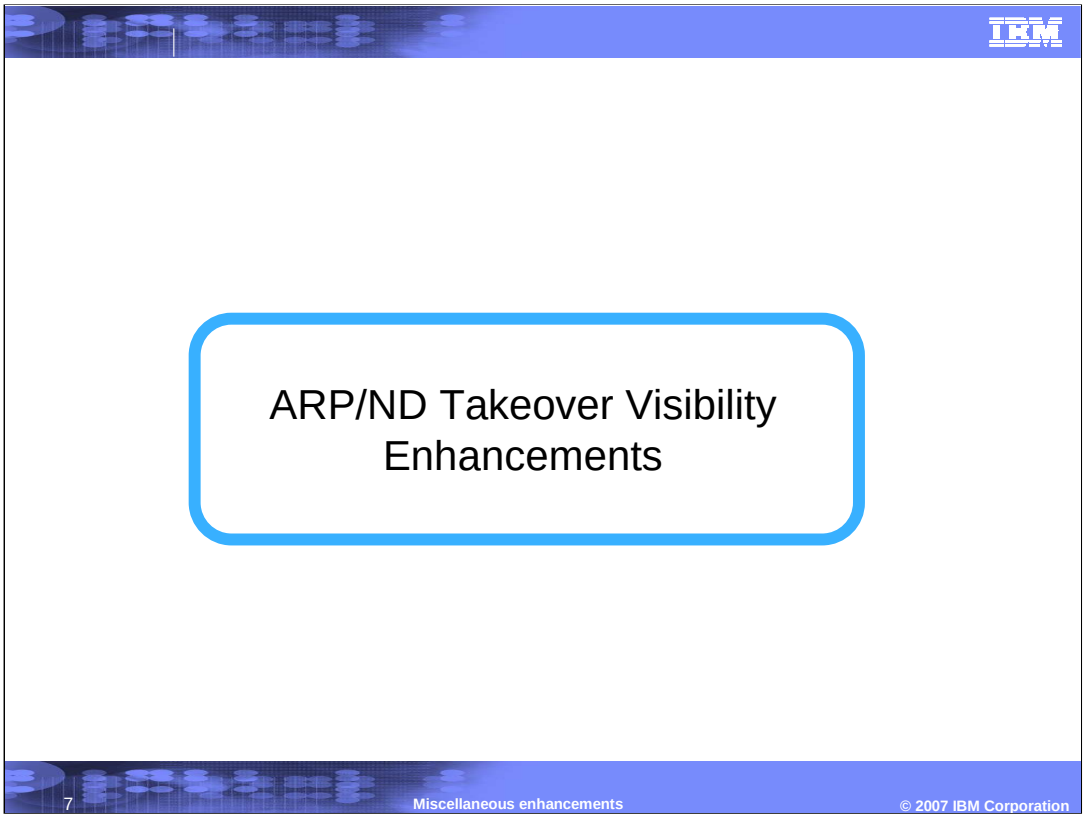
IBM
ARP/ND Takeover Visibility Enhancements
7
Miscellaneous enhancements
© 2007 IBM Corporation

Physical network interface recovery with multiple interfaces to the same network (subnet)

Requirement for this feature to function properly:

- At least two adapters attached to the same network (broadcast media).
- Adapters must use either LCS or QDIO
- The two adapters should be two physical adapters for real availability benefits

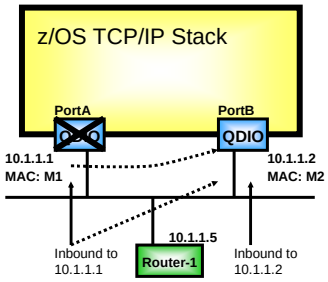

10.x.y.0/24

Router's initial ARP Cache

IP address	Mac address
10.1.1.1	M1
10.1.1.2	M2

Router's ARP Cache after movement of 10.1.1.1 to PortB

IP address	Mac address
10.1.1.1	M2
10.1.1.2	M2

Example: PortA fails or is shut down

1. The z/OS TCP/IP stack moves address 10.1.1.1 to the other QDIO adapter (PortB), which is on the same network (same network prefix) as PortA was.
2. The z/OS TCP/IP stack issues a gratuitous ARP for IP address 10.1.1.1 with the MAC address of PortB (M2) over the PortB adapter
3. Downstream TCP/IP nodes on the same subnet with that IP address in their ARP cache, will update their ARP caches to point to M2 for IP address 10.1.1.1 and will thereafter send inbound packets for both 10.1.1.1 and 10.1.1.2 to MAC address M2

Improved ARP/ND ownership visibility

- **New messages will be issued whenever a previously taken over link or interface has recovered, and takes back the ARP/ND responsibility.**
  - EZD0013I - LINK LOSAQDIO4 HAS TAKEN BACK ARP RESPONSIBILITY FROM LINK LOSAQDIO7
  - EZD0014I - INTERFACE OSAQDIO76 HAS TAKEN BACK ND RESPONSIBILITY FROM INTERFACE OSAQDIO46
- **Messages are already today issued when an interface takes over ARP/ND responsibility**
  - EZZ4329I for IPv4 links and EZZ4345I for IPv6 interfaces
- **Note that takeover/takeback processing does not change, only the visibility has been enhanced.**
- **Enhanced Netstat DEvlinks/-d report to better track the state of takeover:**
  - Displays ARP/ND information.

```
.........
IPv4 LAN Group Summary
LanGroup: 1     7F3DABD0
  LnkName            LnkStatus         ArpOwner          VipaOwner
  -------            ---------         --------          ---------
  LOSAQDIO4          Active            LOSAQDIO4         Yes
  LOSAQDIO7          Active            LOSAQDIO7         No
LanGroup: 2     7F320B90
  LnkName            LnkStatus         ArpOwner          VipaOwner
  -------            ---------         --------          ---------
  LOSAQDIO2          Active            LOSAQDIO2         Yes
```

Allow wildcards on Enterprise Extender network management interface filters

EE NMI filter support

- **The NMI client can supply 1 to 4 'filters' when making a SNA NMI HPR Connection request.**
  - Each filter can specify a CPName.
  - When a CPName is supplied as a filter, the response contains HPR connection records for each connection to the host with that CPName.
- **The CPName (and NetID) can now contain or be entirely comprised of the traditional '?' and '*' wildcards.**
  - '?' is a single character wildcard.
  - '*' is a multi-character wildcard.
- **Examples:**
  - NET*.SSCP*
  - *C.*9C
  - ?.?
  - *.*
- **This enhancement was requested by IBM Tivoli® OMEGAMON® XE for Mainframe Networks**

Network management interface support for dropping socket endpoints

Dropping connections (socket end points)

- **The Netstat DRop/-d and VARY TCPIP,,NETSTAT,DROP intentionally only supports dropping one connection per request**
  - Concern for an end user who might reset multiple or all connections by mistake
    - These commands do not grant you a second chance
  - But this limits installations who want to have the ability to reset all connections under certain circumstances
- **z/OS V1R8 adds support for a network management application to drop TCP Connections and UDP Endpoints**
- **A new DropConnection request is available through the Network Management Interface API**
- **Support multiple connections drops from the API**
  - Applications that will support the DropConnection request should take special care to ensure that the connections input for termination have been examined carefully by programming logic that selects connections meeting a specific criteria, for example by state.

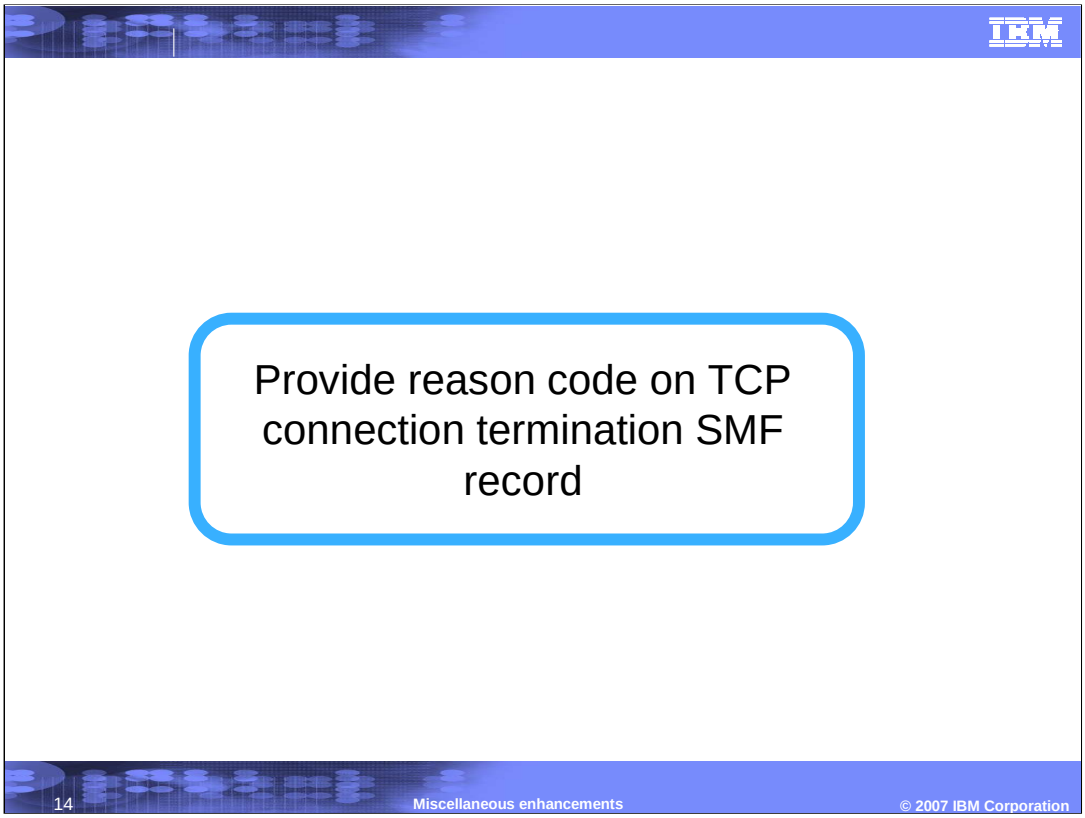
Provide reason code on TCP connection termination SMF record
14
Miscellaneous enhancements
© 2007 IBM Corporation

SMF TCP connection termination records

- **The TCP Connection Termination record is collected whenever a TCP connection is closed or aborted.**
- **This record contains all pertinent information about the connection, such as elapsed time, bytes transferred, and so on.**
- **If the TCP connection was opened by the TELNET server, an additional TCP Connection Termination Section is embedded at the end of the record.**
  - This section holds TELNET information like LU Name, Target Application Name and so on.
- **The TCP Connection Termination SMF record had various information about the stack**
  - Connection initiation time
  - Connection termination time
  - number of bytes inbound
  - number of bytes outbound
  - many more
- **But it does not hold any information on why the connection was terminated.**
  - was it a client action?
  - was it a server application action?
  - was it a protocol error?
- **Similarly, the TELNET section in the SMF record didnot show the reason why TELNET server closed a connection.**

SMF TCP connection termination records

- **z/OS V1R8 adds three new fields to the TCP connection termination smf record**
  - SMF119AP_TTTermCode
    - the termination reason for the connection
  - SMF119AP_TTDUPAcksRcvd
    - number of dup acks received over the life of the connection
  - SMF119AP_TTTelTermCode
    - Part of the TELNET section of the TCP connection termination record
    - If the connection was closed by the TELNET server, holds the reason why it was closed. Documented as part of message EZZ6034I.
    - If the connection was closed by the TCP layer, holds nothing.

TCP Connection termination reason codes

- **x'11'**
  - SMF119AP_TTTermCode_FRCA_ERROR
  - Error occurred during a send using FRCA(AFPA), possibly because the stack is terminating.
- **x'12'**
  - SMF119AP_TTTermCode_FRCA_FIN
  - A persistent socket used by FRCA (AFPA) was closed by a FIN.
- **x'21'**
  - SMF119AP_TTTermCode_STK_TERM
  - The connection was terminated because the stack was terminating.
- **x'22'**
  - SMF119AP_TTTermCode_NODVIPA_OWNER
  - Last stack that can own the dynamic VIPA bound to the socket is terminating.
- **x'31'**
  - SMF119AP_TTTermCode_IDS_TAGGED
  - Intrusion Detection found the connection to be malicious and closed the connection.
- **x'32'**
  - SMF119AP_TTTermCode_NETACCESS_DENY
  - Connection was denied because of a NetAccess rule.
- **x'33'**
  - SMF119AP_TTTermCode_INVACK
  - ACK received in lastack state.
- **x'41'**
  - SMF119AP_TTTermCode_ADMIN
  - The connection was terminated because of an administrator action (for example, using Netstat DRop/-D command or the NMI API).

TCP Connection termination reason codes (cont.)

- **x'42'**
  - SMF119AP_TTTermCode_BOUND_IP_DEL
  - the connection was terminated because the local IP Address bound by the application has been deleted from the stack.
- **x'51'**
  - SMF119AP_TTTermCode_NOT_ACCEPTED
  - The connection from a client was terminated because the application closed the socket before performing an accept().
- **x'52'**
  - SMF119AP_TTTermCode_APP_CLOSE
  - The application using the socket, closed the connection using a close().
- **x'53'**
  - SMF119AP_TTTermCode_PASCAL_CLOSE
  - A pascal routine issued an orderly close request.
- **x'54'**
  - SMF119AP_TTTermCode_PASCAL_DISCO
  - A pascal routine issued a disconnect.
- **x'55'**
  - SMF119AP_TTTermCode_ACCEPT_ERROR
  - An error occurred during a pascal accept.
- **x'61'**
  - SMF119AP_TTTermCode_RST_RCVD
  - The connection was terminated because the client sent a reset

TCP Connection termination reason codes (cont.)

- **x'71'**
  - SMF119AP_TTTermCode_MULTIPLE_REXMIT
  - The connection was closed because the same packet was being re-transmitted multiple times.
- **x'72'**
  - SMF119AP_TTTermCode_ZERO_WINDOW
  - The connection was closed because the TCP window the client got reduced to zero and multiple window probes were not acknowledged.
- **x'73'**
  - SMF119AP_TTTermCode_KEEPALIVE_FAIL
  - The connection was closed because multiple keepalive probes were not acknowledged.
- **x'74'**
  - SMF119AP_TTTermCode_NO_FIN
  - The connection was terminated because the stack timed out waiting for a fin in the finwait-2 state.

Trademarks, copyrights, and disclaimers

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

IBM MVS OMEGAMON Tivoli z/OS z9

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements or changes in the products or programs described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (for example, IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products.

IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance 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 or performance improvements equivalent to the ratios stated here.