This presentation will give you an overview of the enhancements to the Communications Server in z/OS V1R11 for system management and monitoring. The system management and monitoring theme includes various enhancements needed by network management vendors, including Tivoli®, and improved management functions that have been requested by customers.

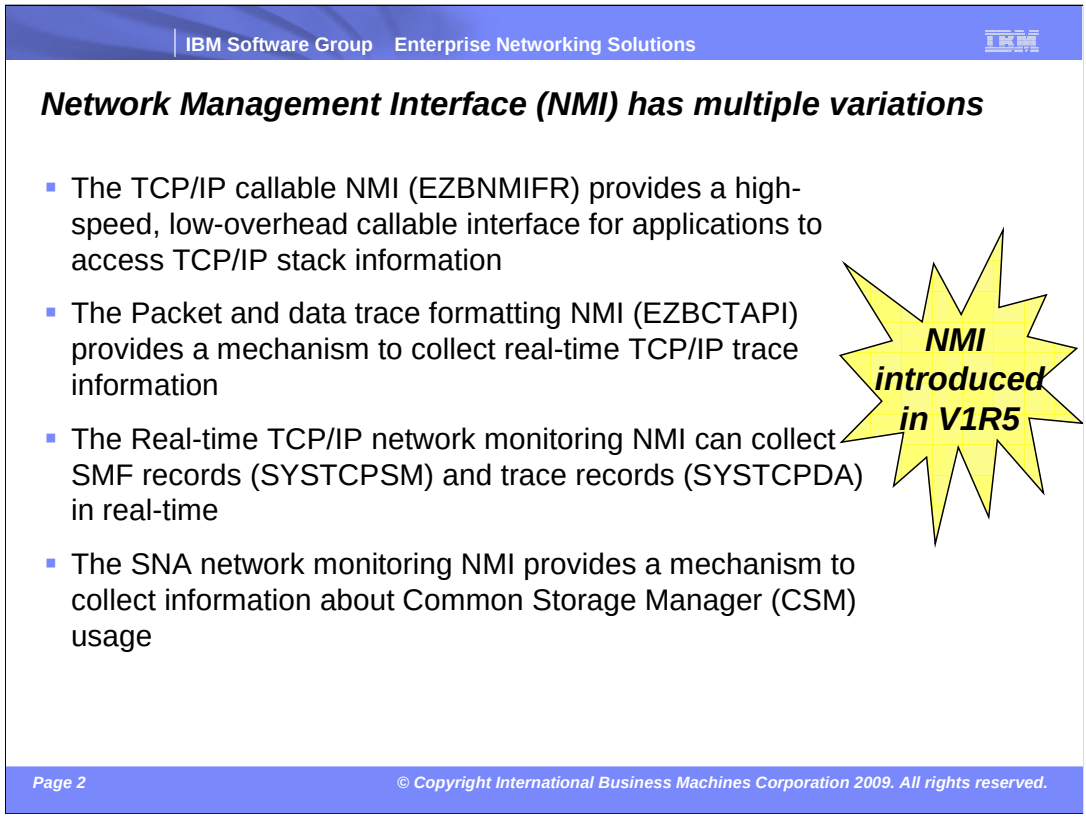

The Network Management Interface, or NMI, function was introduced in z/OS V1R5 Communications Server. It allows network management applications like IBM Tivoli NetView® for z/OS to obtain performance information from the TCP/IP stack. It also allows network management applications to manage TCP connections and UDP endpoints. There are many different forms of NMI functions available for use. Four of the different forms are listed on this slide.

EZBNMIFR, the TCP/IP callable NMI, provides a callable interface for acquiring TCP/IP stack information. This NMI is commonly used to monitor TCP connections, UDP endpoints, TCP/IP storage levels, TN3270E server performance, and to terminate TCP connections.

EZBCTAPI, the Packet and data trace formatting NMI, provides facilities for the collection and formatting of real time packet trace and data trace records from TCPIP.

The Real-time TCP/IP network monitoring NMI can be used to obtain certain types of information on a real-time basis. Two types of information available using this NMI are SMF records, and packet or data records.

The SNA network monitoring NMI can be used to obtain SNA information, for instance the amount of CSM storage used across the system.

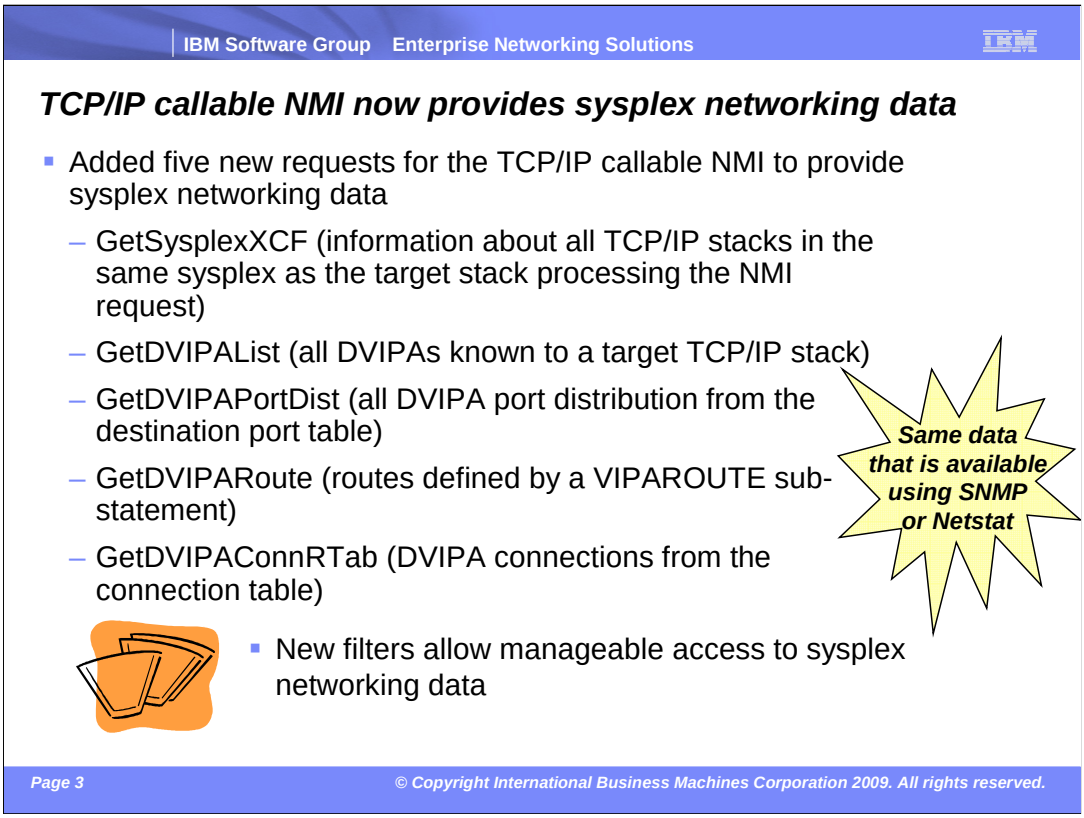

In z/OS V1R11 Communications Server, the TCP/IP callable NMI, EZBNMIFR, is enhanced to provide sysplex networking data. Five new requests are added to obtain different types of sysplex networking data.

The GetSysplexXCF request returns dynamic XCF information about all TCP/IP stacks running in the same sysplex group as the TCP/IP stack processing the NMI request. The GetDVIPAList request returns information about all dynamic VIPAs (DVIPAs) known to the target TCP/IP stack. The GetDVIPAPortDist request returns information about DVIPA port distribution from the destination port table. The GetDVIPARoute request returns information about routes defined by the VIPAROUTE sub-statement of the VIPADYNAMIC profile statement. The GetDVIPAConnRTab request returns information about DVIPAs connections from the connection routing table.

New NMI filters are also defined to allow network management applications to control the type and amount of sysplex networking data that is returned. The new filters include interface name, DVIPA address and port information, destination XCF address information, and target IP address information.

NMI application impacts for sysplex data collection

- NMI applications created using the z/OS V1R11 version of EZBNMRHA or EZBNMRHC will not run on releases before z/OS V1R11
  - The increased filter entry size is not supported by previous releases
  - "Version and release recognition" logic is needed within the NMI application to support multiple releases

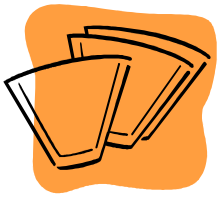

- Existing reason code **JrMustBeSysplex** is now set by EZBNMIFR
  - Indicates the target TCP/IP stack has not yet joined a sysplex group
  - Can be set for all of the new NMI requests except GetDVIPAList

The NMI mappings are available for use in IBM-provided assembler macro (EZBNMHRA) or C mapping file (EZBNMRHC). Network management application that are assembled or compiled using the z/OS V1R11 version of EZBNMRHA or EZBNMRHC typically will not run successfully on releases before z/OS V1R11. This is because the version of EZBNMIFR that exists on those earlier releases does not recognize the increased filter entry size introduced in this release.

For a network management application to run successfully on different versions of z/OS Communication Server, the application must recognize on which release of z/OS it is running and provide filters whose size is appropriate for that release.

The existing JrMustBeSysplex reason code can now be set by the TCP/IP callable NMI, EZBNMIFR. This reason code is set by all of the new sysplex networking data requests (except GetDVIPAList) when the target TCP/IP stack has not joined a sysplex group. If the new GetDVIPAList request is invoked while the target TCP/IP stack is not a member of a sysplex group, then information is only returned for deactivated dynamic VIPAs.

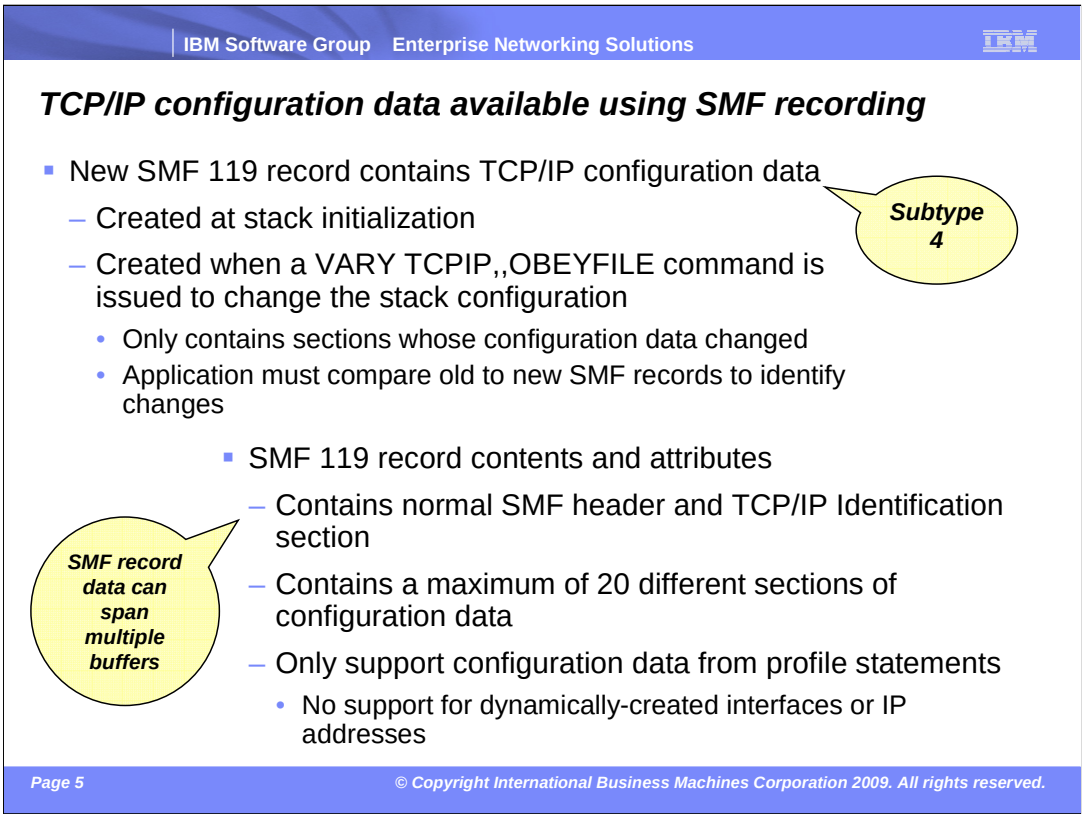

z/OS V1R11 Communications Server defines a new SMF 119 record to report TCP/IP configuration information. The new SMF 119 record has a subtype of 4. This SMF 119 record is available from the MVS SMF data sets and from the real-time SMF data NMI. If requested, the new SMF record is created at stack initialization and, whenever a VARY TCPIP,,OBEYFILE command is issued to change the stack configuration. When created at stack initialization, the SMF record contains all sections for which configuration data exists. When created due to a VARY TCPIP,,OBEYFILE command, then the SMF record only contains those sections whose configuration data changed. Only configuration data from profile statements is present in the SMF records. No dynamically-created IP addresses, run-time status, or counters are supported.

Even if the profile statements referenced by the VARY TCPIP,,OBEYFILE command do not actually change the configuration, the SMF record will still be created. The sections in the record will contain all the configuration data, not just the changed data. Network management applications can compare a section in the current record with the same section in a previous record to determine which configuration values have actually changed.

The new SMF 119 subtype 4 is created with the normal SMF header and the first section is the TCP/IP identification section. The record can contain a maximum of 20 sections of different configuration data. When the data exceeds the maximum size of an SMF 119 record (32746 bytes), multiple SMF records are written to provide all the data.

TCP/IP configuration data available across NMI

- New GetProfile request added to TCP/IP callable NMI interface, EZBNMIFR
  - Returns a maximum of 20 different sections of configuration data
  - No filters are supported for this request

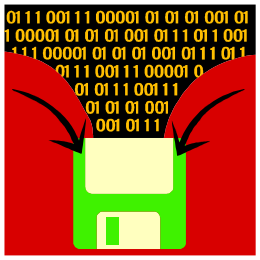

- SMF subtype 4 records can be collected using the real-time network monitoring NMI

In addition to the new SMF 119 subtype 4 record, z/OS V1R11 Communications Server also provides a new GetProfile for the TCP/IP callable NMI interface, EZBNMIFR. The GetProfile request can be used to retrieve the current settings of a stack's configuration data. As was the case for the SMF record, the NMI output can return a maximum of 20 different sections of different configuration data. No filters are supported for this request. Like the other TCP/IP callable NMI requests, if the output buffer is too small, the request returns an error and provides the total size needed for the data.

The SMF records can be acquired using the existing Real-time Network Monitoring NMI logic. The next slide discusses one way to use the SMF and NMI records in a comprehensive manner to monitor configuration changes to the TCP/IP stack.

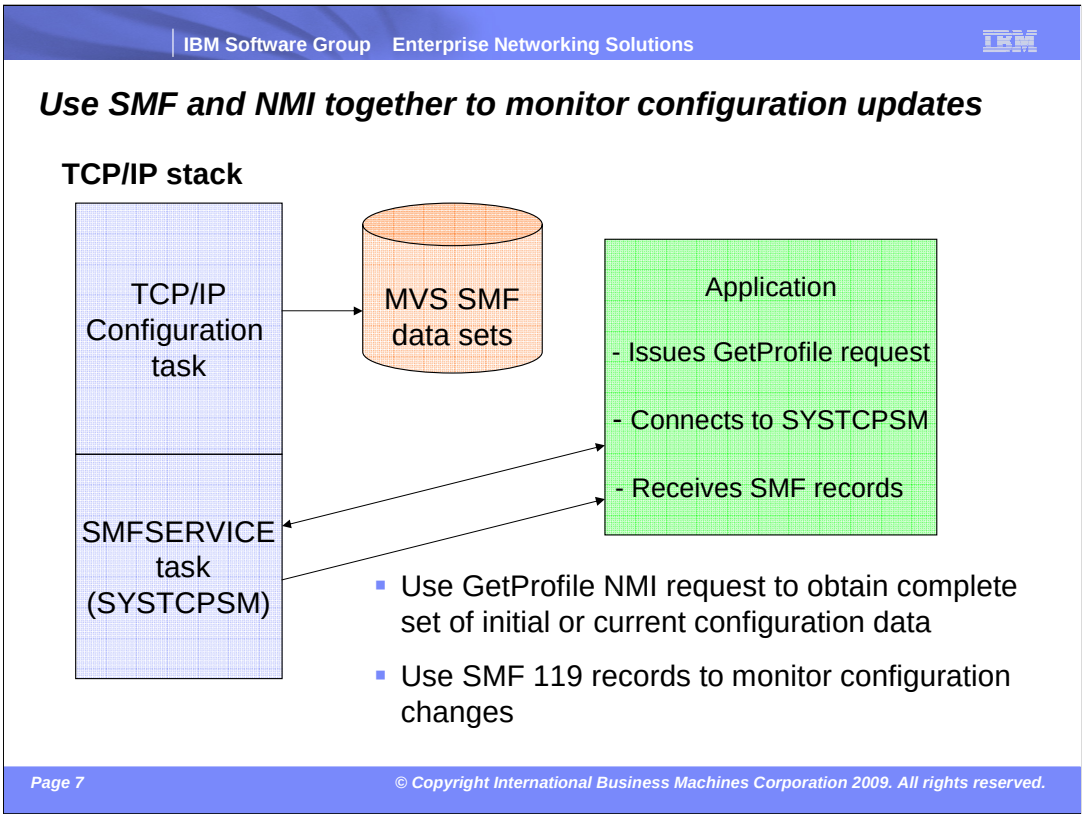

The TCP/IP stack's configuration task processes profile statements either from the initial profile data set or, the profile data set referenced on a VARY TCPIP,,OBEYFILE command. As part of the configuration task processing, the SMF 119 subtype 4 records are written to the MVS SMF data sets and to the real-time SMF data NMI task, based on configuration options. The new configuration statements for activating these functions are discussed on a subsequent slide.

Applications can use the new TCP/IP callable NMI request, GetProfile, to obtain the initial or the current TCP/IP stack configuration data. They can then connect to the real-time SMF data NMI task and receive SMF 119 subtype 4 records containing changes to the configuration.

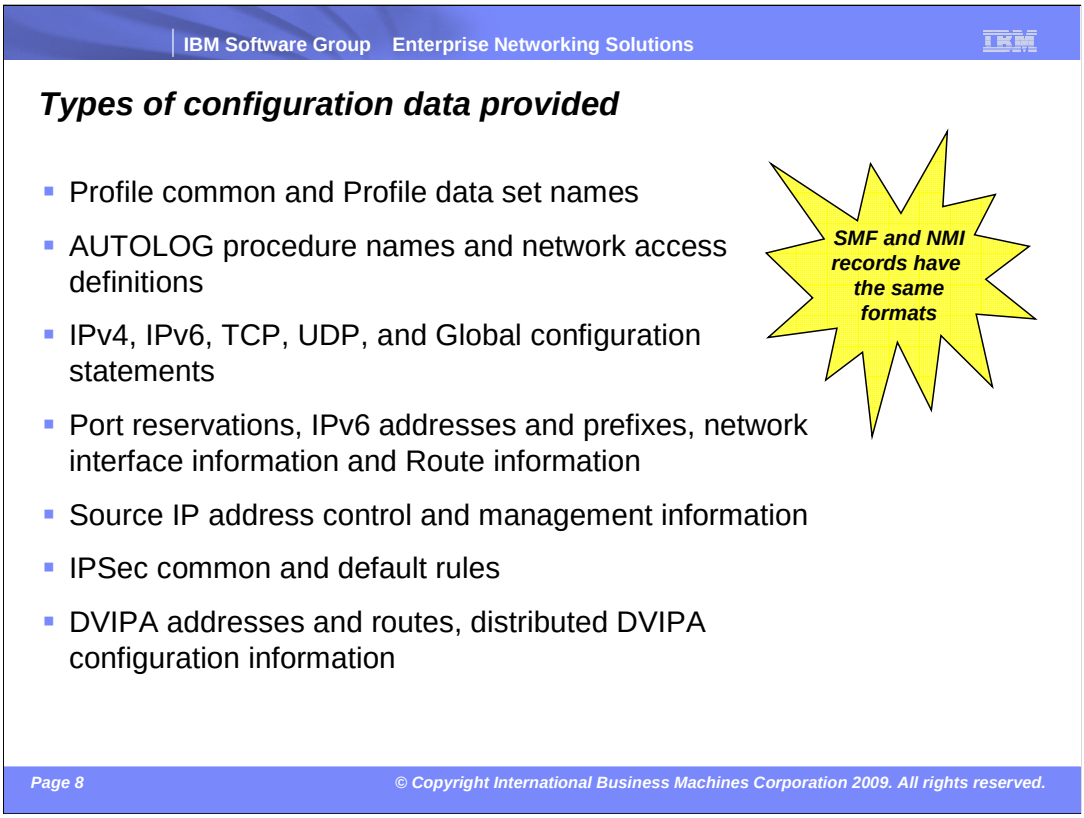

Both the SMF 119 record and the GetProfile NMI request provide the same sections of configuration data. As mentioned previously, a maximum of 20 different configuration sections can be reported. The range of configuration information provided by the SMF or NMI records are listed on the slide.

Packet and data trace NMI now collects OSANETA records

- Existing NMI support collecting packet trace records in real-time
- Function expanded to collect OSAENTA trace records now as well
- Function available to format the OSAENTA records in addition to collecting them

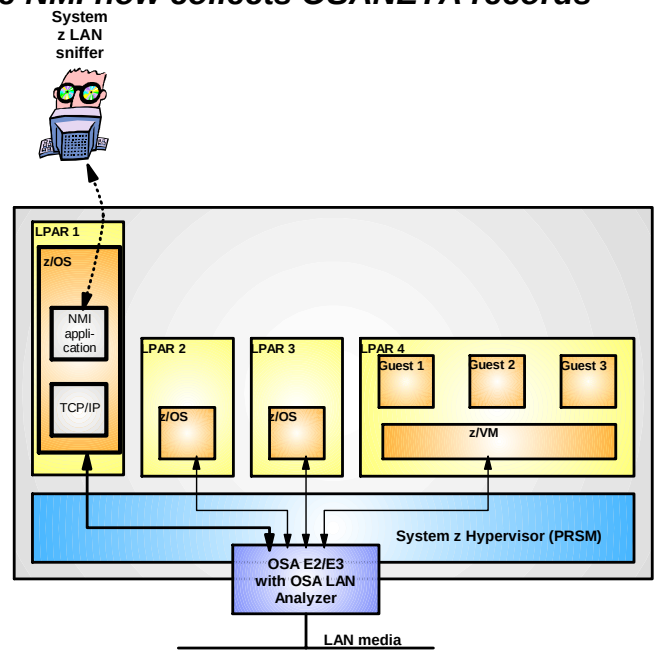

The Real-time TCP/IP network monitoring NMI currently can collect and format real time packet trace and data trace records from TCPIP. Using the same TCP/IP Management Interface (TMI) facilities, z/OS V1R11 Communications Server extends the function to collect and format OSAENTA trace records.

The TCP Monitor Interface will provide buffer tokens for each OSAENTA trace buffer to the application over a UNIX® sockets interface. The EZBCTAPI macro service used to format packet trace records can now be used to format OSAENTRA trace buffers. In addition, the EZBYPTO data area is updated to support both additional filters for packet trace and OSAENTA trace.

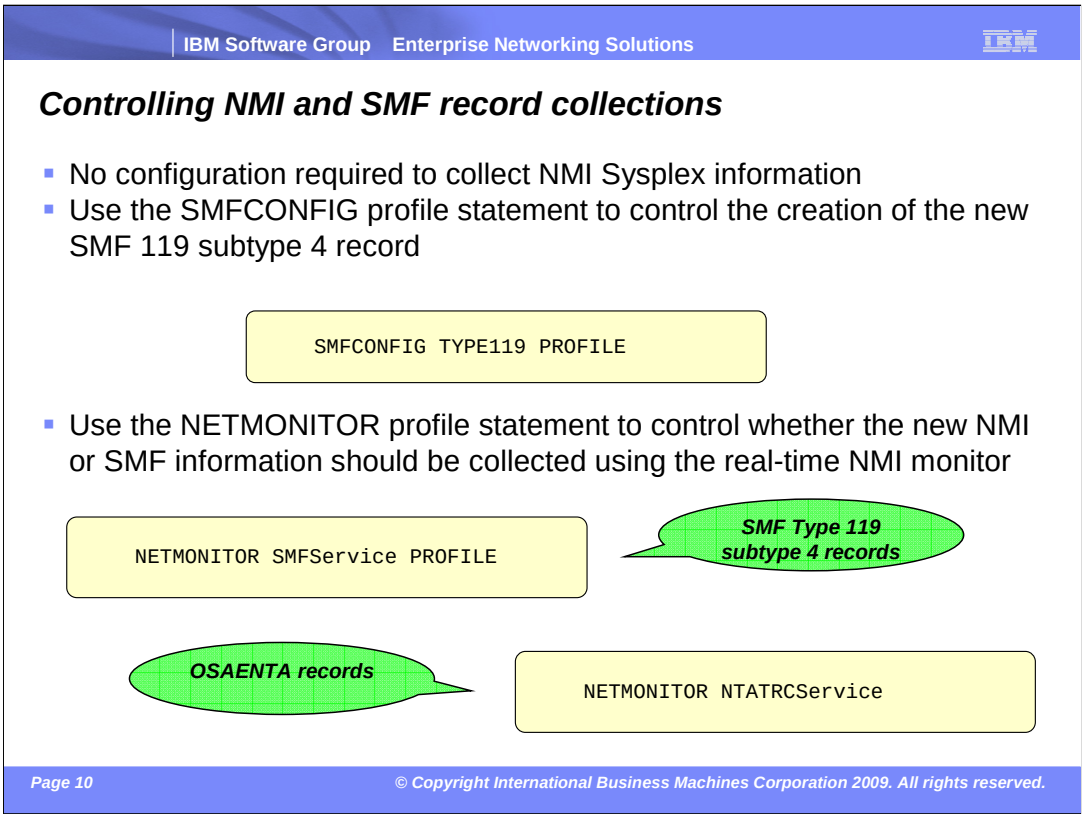

In some cases, new configuration settings are required to collect the TCP/IP NMI information that has been discussed. There is no configuration required to collect the sysplex NMI information, but there is for the other two functions.

A new sub-parameter keyword, PROFILE, is provided on the TYPE119 parameter of the SMFCONFIG profile statement to control the creation of the new SMF 119 subtype 4 records. This sub-statement controls only the collection of the SMF records in the MVS SMF data sets. A corresponding NOPROFILE sub-parameter is also provided to disable the collection of these records. By default, the NOPROFILE parameter setting is in effect, meaning that the new SMF 119 record is not written to the MVS SMF data sets. This setting can be changed by a VARY TCPIP,,OBEYFILE command.

A new sub-parameter, PROFILE, on the TCPIP Profile NETMONITOR SMFSERVICE statement is provided to collect the new SMF records using the Real-time Network Monitor NMI. A corresponding sub-parameter, NOPROFILE, can be used to disable collection of the records in this manner.

Similarly, a new sub-parameter, NTATRCService, on the TCPIP Profile NETMONITOR statement is provided to collect the OSAENTA trace record information using the Real-time Network Monitor NMI. The NONTATRCService sub-parameter can be used to disable the function.

SNA NMI request now can request detailed CSM usage

- Provide an optional ASID filter on the CSM statistics request to request detailed CSM usage data
  - Allow requests for detailed CSM usage data for all owners or a specific owner
  - Detailed CSM usage data is only returned when a filter is provided
  - Current CSM storage pool statistics and CSM summary information are always still returned

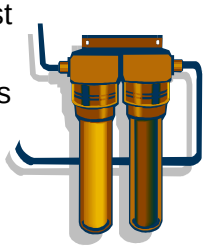

- Provide detailed CSM usage data on the CSM statistics response
  - Jobname and ASID
  - Amount of storage owned in each storage pool
  - Includes owners which have freed all CSM buffers

z/OS V1R11 Communications Server enhances the SNA network monitoring NMI to provide detailed CSM usage data on CSM statistics responses.

A new ASID filter is provided for the CSM statistics request. The ASID filter is optional, and either zero or a specific *owner_asid* value can be provided on the filter. When zero is provided, z/OS V1R11 will return the amount of storage owned by all CSM buffer owners. When a specific *owner_asid* value is provided, then only the amount of storage owned by that specific owner is returned. One to four ASID filters can be supplied on a single CSM statistics request. The new detailed CSM usage data is only returned when the ASID filter is supplied. The existing CSM storage pool statistics and CSM summary information continue to be provided on all CSM statistics requests.

For each owner ID reported, the data provided includes the jobname and ASID and the amount of storage owned in each storage pool. As is the case in the CSM storage pool statistics data, the 31 bit and 64 bit data space pools are not separately reported. The data returned includes owners which have freed all CSM buffers , which allows the NMI application to more accurately track when storage is freed by the system.

Trademarks, copyrights, and disclaimers

IBM, the IBM logo, ibm.com, and the following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

Current NetView Tivoli z/OS

If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. 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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

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.

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.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.