




IBM Software Group Enterprise Networking Solutions
z/OS® V1R11 Communications Server

z/OS V1R11 Communications Server – availability and business resilience

z/OS Communications Server Development, Raleigh, North Carolina



© Copyright International Business Machines Corporation 2009. All rights reserved.

This presentation will give you an overview of the enhancements to the Communications Server in z/OS V1R11 for availability and business resilience. This area covers enhancements that make the z/OS networking environment more resilient to various types of abnormal conditions that might impact the overall availability of the system.

Availability and business resilience


- *OMPROUTE detection of duplicate router ID*
- 🔍 Improved responsiveness to storage shortage conditions
- *Protection from run-away applications*
- *Disable moving DVIPA as source IP address*
- 🔍 Support for enhanced WLM routing algorithms

The main enhancements within this area are those listed on this slide.

Special emphasis in this release has been put into making the TCP/IP environment able to better cope with situations where storage resources are constrained – typically caused by extreme workload conditions.

Improved responsiveness to storage shortage conditions

- Capping amount of parallel resources (SRBs) used to process inbound QDIO data
- QDIO inbound “overrun” detection and recovery (drop packets at the DLC layer!)
- Improve OMPROUTE toleration of storage shortage conditions
- Provide alerts when excessive data builds up on TCP send or receive queues
- When fixed storage becomes constrained, provide relief by:
 - Mark unsent data on send queues as pageable (from fixed)
 - Mark all new data added to send queues pageable



- ✓ Stay up!
- ✓ Throttle workload at the source
- ✓ Prevent network spikes from monopolizing common z/OS storage

There have in the past been situations where OMPROUTE, in storage constrained situations, was not able to get storage for sending its periodic hello packets. This led to neighboring routers believing z/OS had gone down – even though it was still up-and-running, however in a very constrained state.

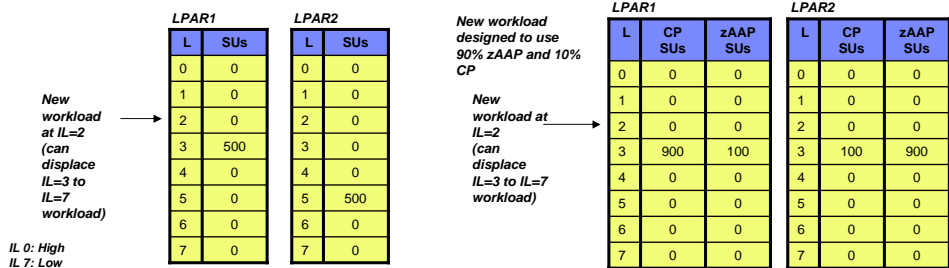
In this release z/OS Communications Server attempts to control situations that can lead to storage constraints by detecting early signs of extreme conditions for ECSA storage. If such signs are detected, z/OS Communications Server will attempt to limit inbound traffic at the source: the inbound QDIO network interfaces. This is done by capping the processor resources that are used for processing inbound QDIO buffers. z/OS Communications Server will further start dropping inbound packets (the normal IP-based recovery action of constrained nodes). The remote TCP will retransmit in an attempt to recover the lost packets.

OMPROUTE storage processing has been changed to ensure that OMPROUTE can get storage for sending its hello packets even in cases where the TCP/IP environment is severely storage constrained.

Another factor that can impact storage usage is applications that either do not read the inbound data that is queued for them in the receive buffers or accumulate large amounts of data in the send buffers.

Workload distribution algorithm enhancements by sysplex distributor and workload manager in z/OS V1R11

- New workload at Importance level 2
 - Which LPAR is best?
 - They both have 500 service units of displaceable workload
 - Before V1R11, they are equal
 - z/OS V1R11 takes importance level of displaceable workload into consideration
 - LPAR2 is preferred
- New workload at Importance level 2
 - Which LPAR is best?
 - They both have equal amount of displaceable service units
 - Before V1R11, they are equal
 - z/OS V1R11 takes amount of crossover to CP of displaceable workload into consideration
 - LPAR2 is preferred since it has the least amount of crossover



WLM supports a new method to calculate relative weights of the target servers, to which Sysplex Distributor distributes incoming connections.

The new method provides two enhancements.

The first enhancement is for WLM to consider the relative importance of existing work that is displaced on each target as new work is distributed. Target systems with more lower importance displaceable work is favored over systems with higher importance displaceable work.

The second enhancement is to factor in the use of specialty processors (zIIP and zAAP) and the amount of cross-over that occurs. If two target LPARs has the same amount of zIIP and zAAP processor capacity, WLM can be requested to prefer the target that has the least amount of cross-over. Cross-over means zIIP/zAAP eligible workload that is processed on a general CP.

These enhancements apply to the SERVERWLM distribution method.

This support is based on the new METHOD=EQUICPU function in the IWM4SRSC WLM service.

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:
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>

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.

© Copyright International Business Machines Corporation 2009. All rights reserved.

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.