



IBM Software Group

# Communication Controller for Linux® on System z® V1.3

*DLSw dynamic peer keepalive*



@business on demand.

© 2008 IBM Corporation  
Updated December 30, 2008

## Background

- CCLDLS currently supports TCP keepalive for peer connections defined as either active or passive.
- There is no keepalive support for dynamic peer connections.

## Problem

- You could experience a catastrophic failure of your IP network that does not affect socket status.
- The failure could go undetected by ccdls, and therefore, CCL, NCP and VTAM®.
- Peer router could detect and try to connect sessions through a different path.
- VTAM on the ccdls end may reject the new connection requests believing there are still active sessions.
- Results in recovery failure and long network outages.



You could experience a catastrophic failure of your IP network that does not affect socket status. As a result, this failure could go undetected by CCLDLS, and as such go undetected by CCL, NCP and VTAM.

After a failure such as losing the IP network, the peer router could detect and try to re-drive sessions through a different path. However, the VTAM on the CCLDLS end may reject the new connection requests, since it believes there are still active sessions over the CCLDLS that did not detect a failure. This could lead to recovery failure and potentially long network outages.



## New global configuration parameter

- `dynamic_keepalive_support`



A new global configuration parameter, 'dynamic\_keepalive\_support' controls whether keepalive will be performed on dynamic peer connections.

Valid settings are 'enabled' or 'disabled'. Default is 'disabled'.

## New parameters in DLSw timers section

- `keepalive_idle`
- `keepalive_interval`
- `keepalive_count`

These timer values affect all active, passive and dynamic connections.



New parameters in the DLSw timers section are used for all connections supporting keepalive.

Use 'keepalive\_idle' to set `TCP_KEEPIDLE`, specifying the number of seconds the TCP connection is idle before sending keepalive probes. The default value is 7200 seconds.

Use `keepalive_interval` to set `TCP_KEEPINTVL`, specifying the number of seconds between keepalive attempts. The default value is 75 seconds.

Use `keepalive_count` to set `TCP_KEEPCNT`, specifying the number of keepalive attempts before declaring the connection inoperative. The default value is 9 attempts.

Even though `keepalive_count` is not a timer, it is included in the timer section because the three variables work together, and this is consistent with Linux.

## Console updated

- You can see the keepalive timers in the global DLSw timer section of the console display.



The console was modified to display the value of the keepalive timers in the global DLSw timer section. The use of dynamic keepalive support has always been displayed in the DLSw global panel as disabled. This was changed to reflect what is configured for dynamic keepalive.





## 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:

System z VTAM

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>

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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 2008. 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.