



IBM Tivoli Training
IBM Tivoli Directory Server Version 6.1
Blocked replication



© 2008 IBM Corporation
Updated October 24, 2008

<no sound>

Objectives

- Upon completion of this module, you should be able to:
 - ▶ Determine when it is safe to skip blocked replication entries
 - ▶ Save the Lightweight Data Interchange Format (LDIF) data for the blocked entries
 - ▶ Correct entries on a replica
 - ▶ Describe why it is important to maintain the same schema on all instances in a replication topology
 - ▶ Determine the status of a replication queue

This module will provide you with some insight into the management and determination of blocked replication entries.

When is it safe to skip a blocking entry

- Know your data
 - ▶ Understand the dependencies
- Be able to predict what will happen if entry is skipped
- Preserve the changed data
- Determine if the problem can be corrected so the blocked entry will not need to be skipped
- Document the problem
 - ▶ What caused it?
 - ▶ What was the data that will be skipped?
 - ▶ What will the predicted behavior be if skipped?

When determining if it is safe to skip a blocked replication entry, the most important thing you can do is to know your data. Understanding the relationship between each entity is the only means to determine what the consequences of skipping data might be. It is useless to skip an entry if the remaining entries in the queue depend on it.

Ask yourself the question: If I skip this entry, what will happen next? How will I eventually correct the problem so that all servers are back in sync?

Always capture a copy of the entry that you are skipping. An example of this process will be shown on one of the following slides.

Often you do not have to skip the entry. If you know the data and know what caused the problem, you can probably fix the issue so that the blocked entry will become unblocked.

Proper documentation of the problem is essential. Make sure that you document what caused the problem, the LDIF of the entry that was skipped, and what is to be expected after skipping. Also document how the data is to be synchronized.

Save the failing LDIF

- You need to know the replication agreement that has the failure

Search objectclass = ibm-replicationAgreement for the following attributes:

- ▶ ibm-replicationFailedChanges
- ▶ ibm-replicationFailedChangeCount

```
idsldapsearch -D <adminDN> -w <adminPW> -h <servername> -p  
<portnumber> -b " " -s base objectclass=ibm*nt ibm-  
replicationfailedchanges ibm-replicationfailedchangeount
```

To save the failing entry, you first must know the replication agreement that has the failure. You can search all of your replication agreements to determine the distinguished names (DNs) that have failures.

Save the failing LDIF (Continued)

- Save the data
- Use the controlreplerr extended operation
 - ▶ Show the error

```
ldapexop -D <adminDN> -w <adminPW> -op controlreplerr -show 1 -ra cn=<server>-  
1389,ibm-replicaServerId=<server>-389, ibm-replicaGroup=default,o=sample
```

Might return:

```
dn: entry-85,o=sample  
cn: entry-85  
objectclass: person  
objectclass: top  
userpassword: {AES256}tDO9yQT540xpp7ZMlg95mA==  
sn: user  
ibm-entryuuid: bf201fcb-758e-41dc-bdea-1855fe0b860b
```

Use the command `ldapexop` with the operation type of `controlreplerr` to show the LDIF for an error.

Correcting entries on the replica

- Use `idsldapdiff` command to compare the supplier and consumer server instances to confirm all differences
- Develop the correct LDIF file using operational LDIF format to perform the modify on the consumer
- Use the credentials that the supplier uses to bind to the consumer
- Make the changes to the consumer
- If the change corrects the problem, the blocked entry should now flow to the consumer

There are several ways to correct an error, depending on its nature. In many cases you can use the `idsldapdiff` command to compare and synchronize the servers. If you need to make a modification on the consumer, you must create an LDIF file suitable for the `idsldapmodify` command. Because the change needs to be made on the consumer, which is normally read-only, you must use the credentials that the supplier uses to bind to the consumer, not the administrator's credentials.

Keeping servers synchronized

- Never change data directly on a consumer unless it is to fix a problem
- Always document all changes to both suppliers and consumers
- The schema of all servers in a replication topology must be identical
- Know you data

Keeping servers synchronized is simple. Make sure that all changes made to a consumer are fully documented and understood. If you change schema, change it everywhere, or even better, have it replicated.

And above all else, know your data.

Training roadmap for IBM Tivoli Directory Server Version 6.1

http://www.ibm.com/software/tivoli/education/edu_prd.html



Summary

- You should now be able to:
 - ▶ Determine when it is safe to skip blocked replication entries
 - ▶ Save the Lightweight Data Interchange Format (LDIF) data for the blocked entries
 - ▶ Correct entries on a replica
 - ▶ Describe why it is important to maintain the same schema on all instances in a replication topology
 - ▶ Determine the status of a replication queue



Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send e-mail feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_replication.ppt

This module is also available in PDF format at: [../replication.pdf](..../replication.pdf)



You can help improve the quality of IBM Education Assistant content by providing feedback.

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 Tivoli

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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 NON-INFRINGEMENT. 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.

