IBM Tivoli Software

# IBM Tivoli Directory Server 6.0 - Replication

Excerpt taken from presentation given on April 24, 2007
Implementing a replication topology: Preparing the System for Replica

**Support Technical Exchange Web site**
http://www-306.ibm.com/software/sysmgmt/products/support/supp_tech_exch.html

8/14/2007

# The boss walks in and says…

- **We need to have two ITDS 6.0 servers behind a Load Balancer on the intranet, but we also need a read only copy of this data on the internet. So we need 2 peers and one consumer for this example.**

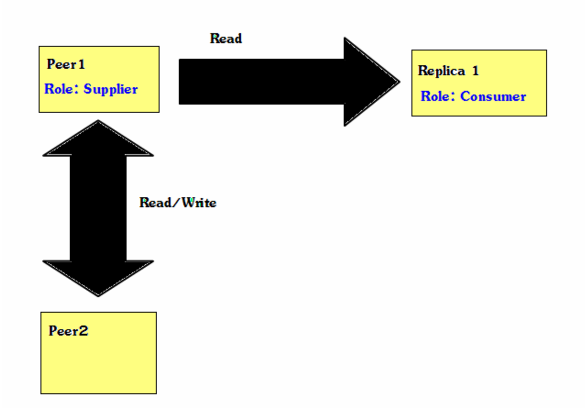

Peer1

Role: Supplier

Read

Replica 1

Role: Consumer

Read/Write

Peer2

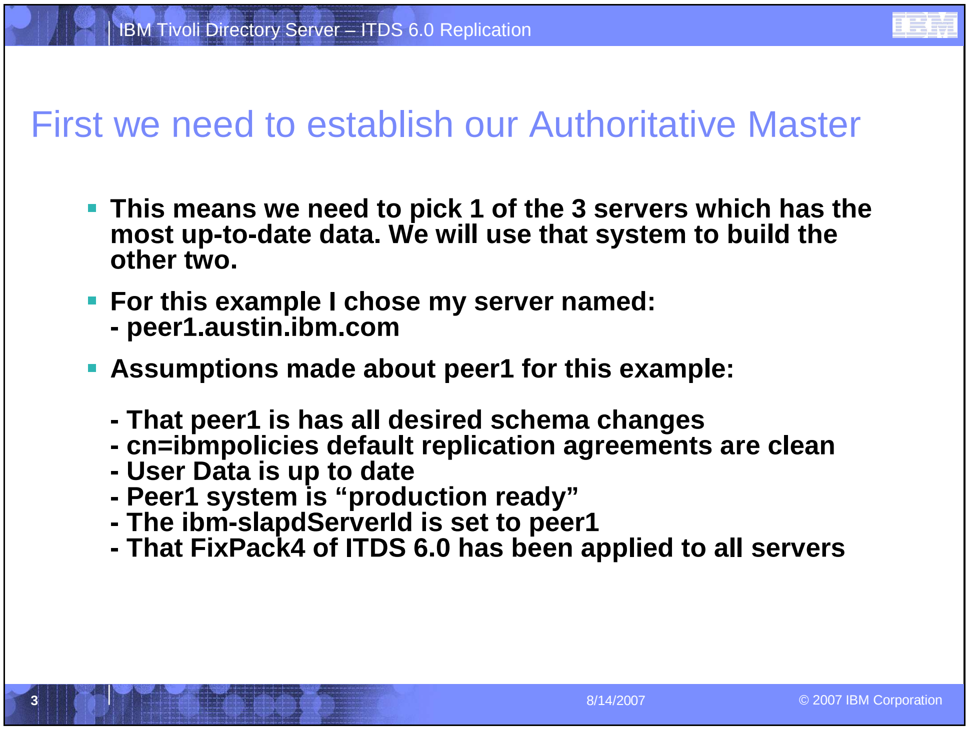

IBM Tivoli Directory Server – ITDS 6.0 Replication

## First we need to establish our Authoritative Master

- **This means we need to pick 1 of the 3 servers which has the most up-to-date data. We will use that system to build the other two.**
- **For this example I chose my server named:**
  **- peer1.austin.ibm.com**
- **Assumptions made about peer1 for this example:**

  **- That peer1 is has all desired schema changes**
  **- cn=ibmpolicies default replication agreements are clean**
  **- User Data is up to date**
  **- Peer1 system is “production ready”**
  **- The ibm-slapdServerId is set to peer1**
  **- That FixPack4 of ITDS 6.0 has been applied to all servers**

3 8/14/2007 © 2007 IBM Corporation

Changed:

This simply means we need to pick which of the 3 systems we will use for our configuration has the most current and up to date data.

To

This means we need to pick 1 or the 3 servers which has the most up-to-date data. We will use that system to build the other two.

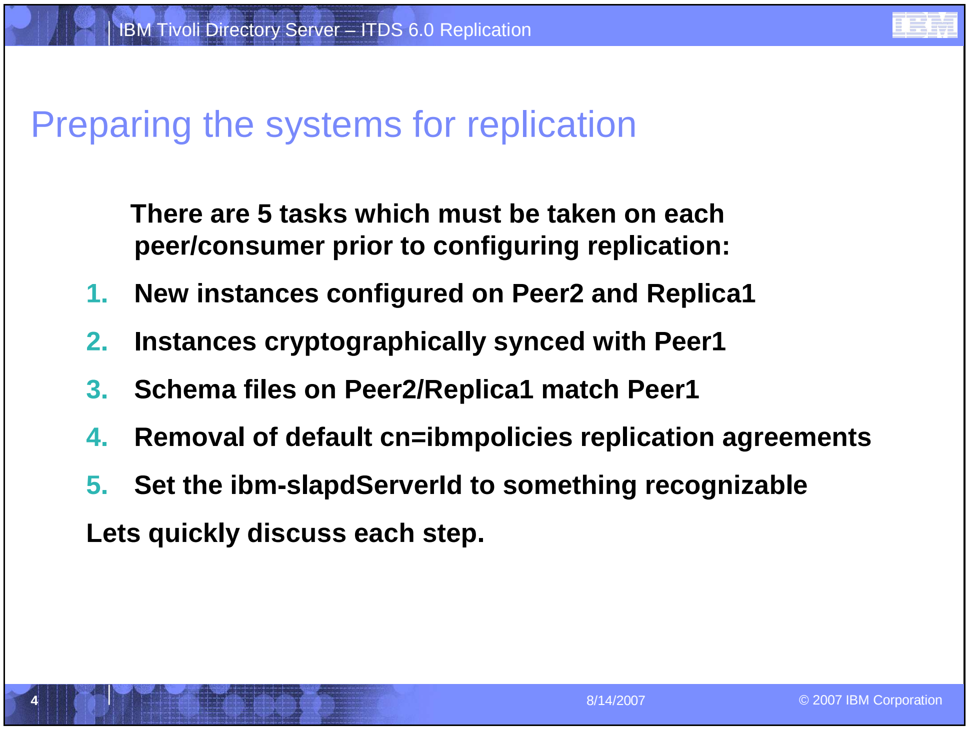

## Preparing the systems for replication

**There are 5 tasks which must be taken on each peer/consumer prior to configuring replication:**

1. **New instances configured on Peer2 and Replica1**
2. **Instances cryptographically synced with Peer1**
3. **Schema files on Peer2/Replica1 match Peer1**
4. **Removal of default cn=ibmpolicies replication agreements**
5. **Set the ibm-slapdServerId to something recognizable**

**Lets quickly discuss each step.**

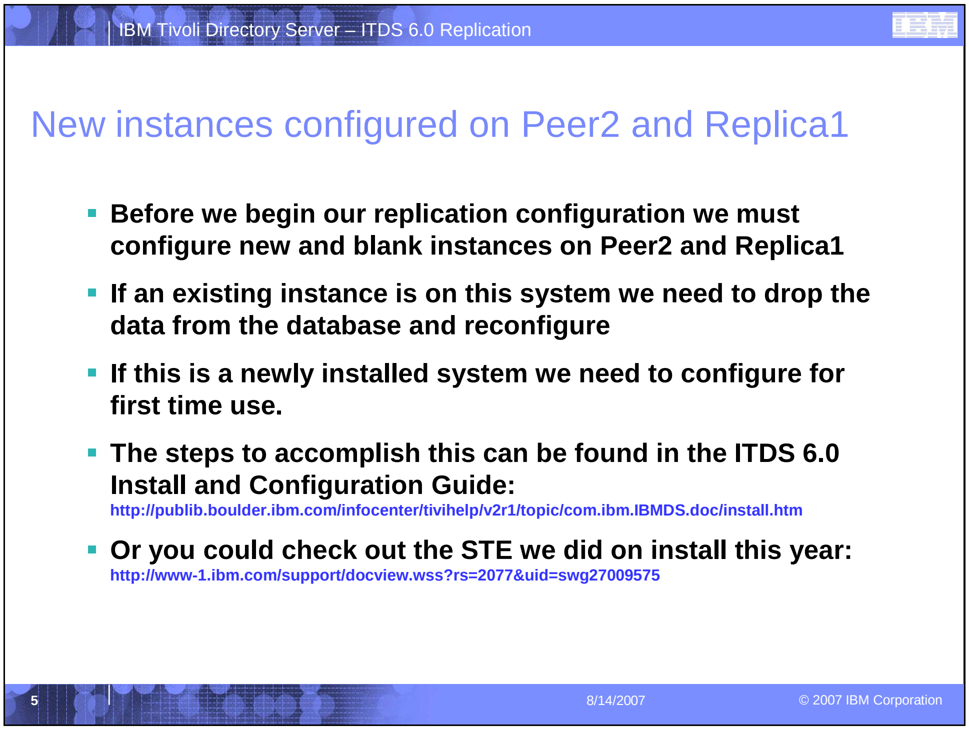

# New instances configured on Peer2 and Replica1

- **Before we begin our replication configuration we must configure new and blank instances on Peer2 and Replica1**
- **If an existing instance is on this system we need to drop the data from the database and reconfigure**
- **If this is a newly installed system we need to configure for first time use.**
- **The steps to accomplish this can be found in the ITDS 6.0 Install and Configuration Guide:**
  **http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc/install.htm**
- **Or you could check out the STE we did on install this year:**
  **http://www-1.ibm.com/support/docview.wss?rs=2077&uid=swg27009575**

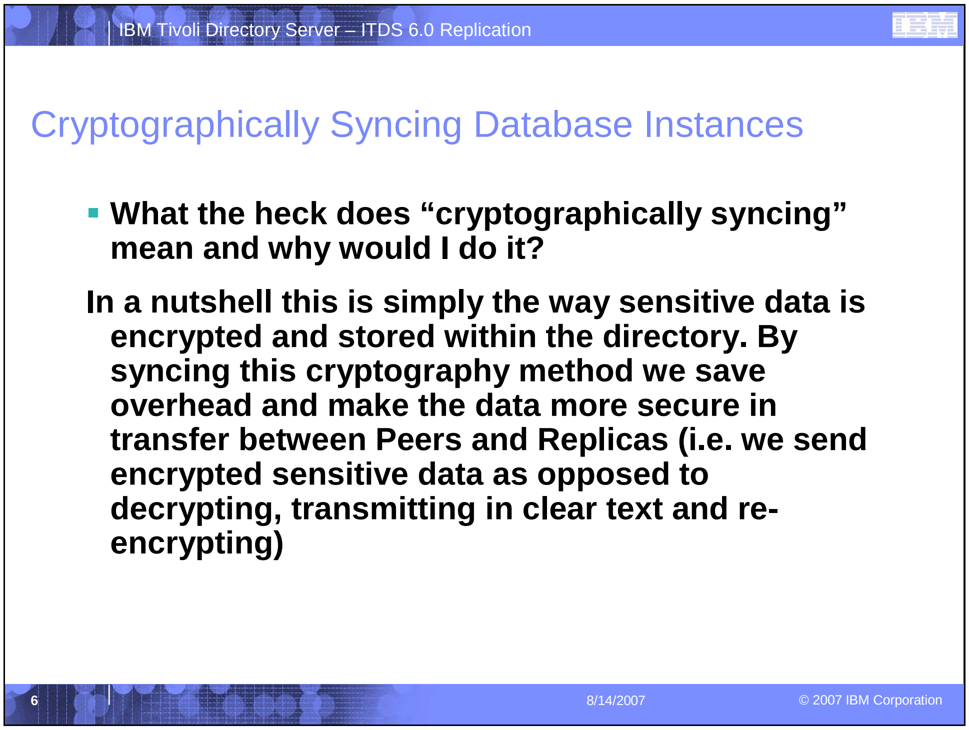

## Cryptographically Syncing Database Instances

- **What the heck does “cryptographically syncing” mean and why would I do it?**

**In a nutshell this is simply the way sensitive data is encrypted and stored within the directory. By syncing this cryptography method we save overhead and make the data more secure in transfer between Peers and Replicas (i.e. we send encrypted sensitive data as opposed to decrypting, transmitting in clear text and re-encrypting)**

# How do I cryptographically sync???

- **There are several good resources for instructions on this topic:**
  - http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc/install24.htm
  - http://publib.boulder.ibm.com/infocenter/tivihelp/v2r1/topic/com.ibm.IBMDS.doc/admin_gd40.htm
  - http://www-1.ibm.com/support/docview.wss?rs=767&context=SSVJJU&q1=corrections&uid=swg21210430&loc=en_US&cs=utf-8&lang=en
  - http://www-1.ibm.com/support/docview.wss?rs=767&context=SSVJJU&q1=ibmslapddir.ksf&uid=swg21248873&loc=en_US&cs=utf-8&lang=en
  - http://www-1.ibm.com/support/docview.wss?rs=767&context=SSVJJU&q1=ibmslapddir.ksf&uid=swg1IO03347&loc=en_US&cs=utf-8&lang=en

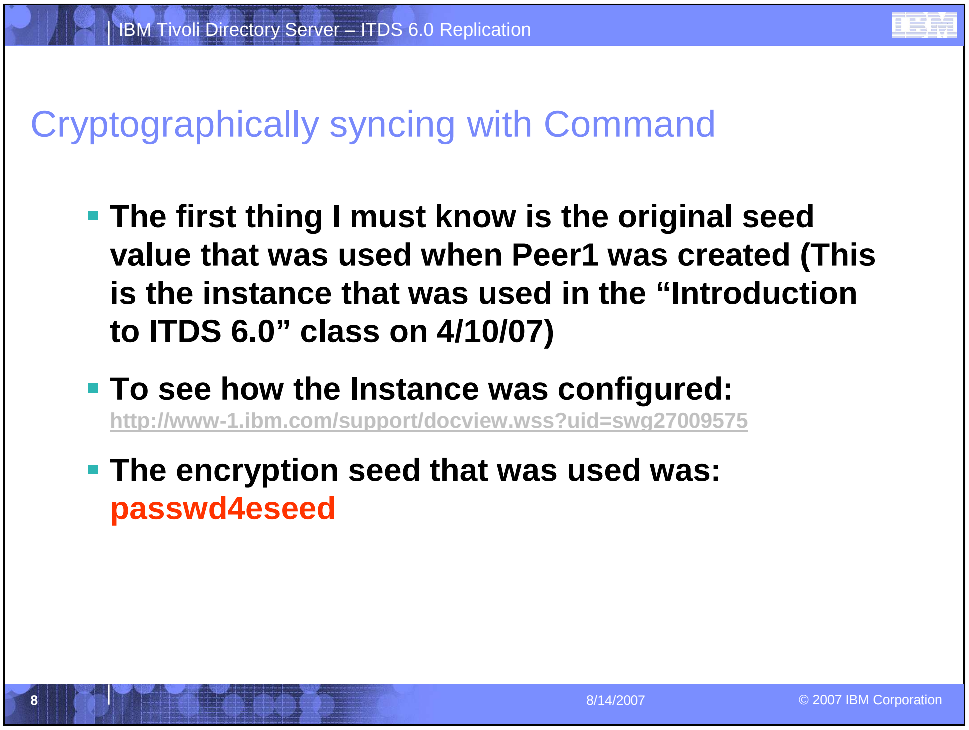

# Cryptographically syncing with Command

- **The first thing I must know is the original seed value that was used when Peer1 was created (This is the instance that was used in the “Introduction to ITDS 6.0” class on 4/10/07)**
- **To see how the Instance was configured:** **http://www-1.ibm.com/support/docview.wss?uid=swg27009575**
- **The encryption seed that was used was:** **passwd4eseed**

# Crypto Syncing… The Salt

- **The next thing I will require is the salt value used on my “authoritative master” or peer1.**
- **To find this out I run:**

```
ldapsearch -D cn=root -w secret -b cn=crypto,cn=localhost objectclass=*
cn=crypto,cn=localhost
cn=crypto
objectclass=ibm-cryptoConfig
objectclass=ibm-slapdConfigEntry
objectclass=top
ibm-slapdCryptoSync=Qm5rb4B9F+p2BvDd
ibm-slapdCryptoSalt=2&mX4AsaJ(|A
```

**So our salt value will be: 2&mX4AsaJ(|A**

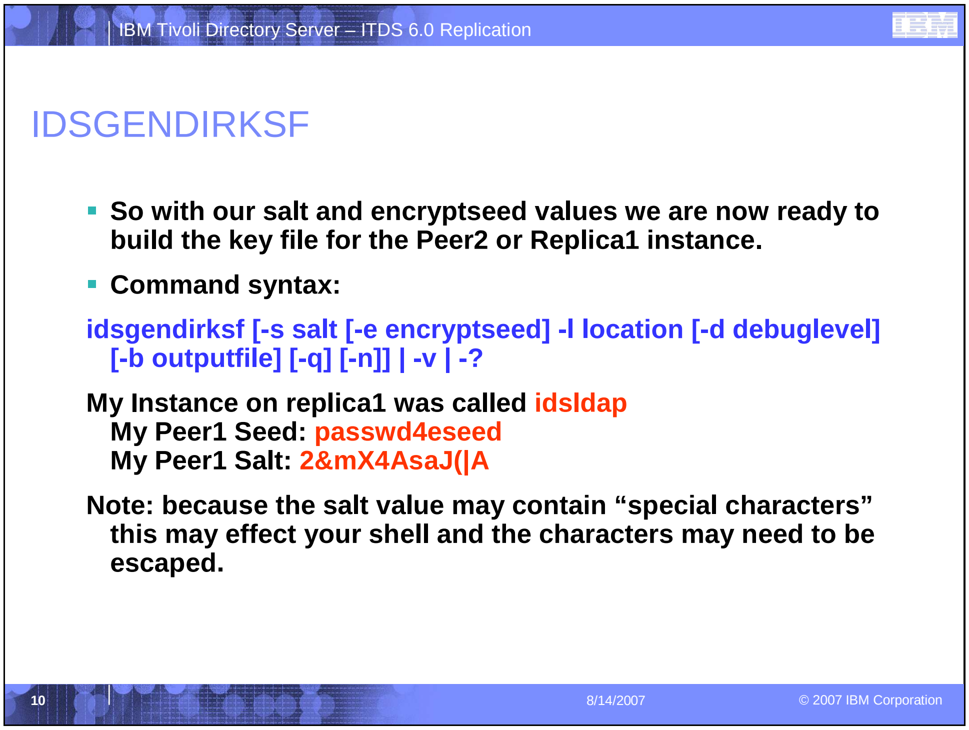

# IDSGENDIRKSF

- **So with our salt and encryptseed values we are now ready to build the key file for the Peer2 or Replica1 instance.**
- **Command syntax:**

**idsgendirksf [-s salt [-e encryptseed] -l location [-d debuglevel] [-b outputfile] [-q] [-n]] | -v | -?**

**My Instance on replica1 was called idsldap**
**My Peer1 Seed: passwd4eseed**
**My Peer1 Salt: 2&mX4AsaJ(|A**

**Note: because the salt value may contain “special characters” this may effect your shell and the characters may need to be escaped.**

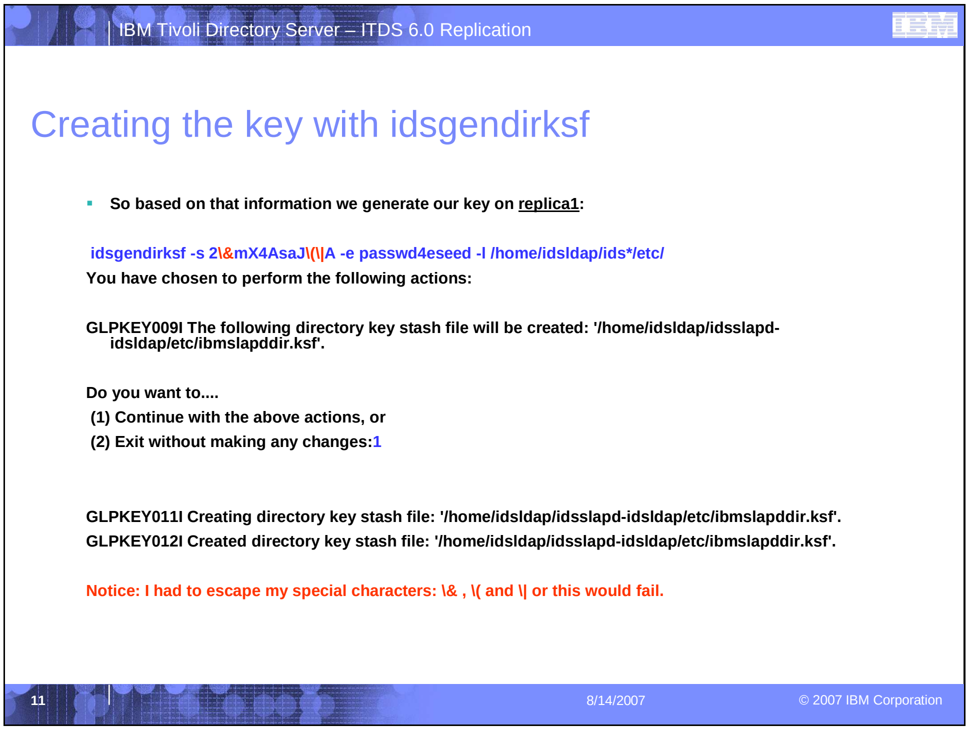

# Creating the key with idsgendirksf

- So based on that information we generate our key on replica1:

```
idsgendirksf -s 2\&mX4AsaJ\(\|A -e passwd4eseed -l /home/idsldap/ids*/etc/
You have chosen to perform the following actions:

GLPKEY009I The following directory key stash file will be created: '/home/idsldap/idsslapd-
     idsldap/etc/ibmslapddir.ksf'.

Do you want to....
 (1) Continue with the above actions, or
 (2) Exit without making any changes:1


GLPKEY011I Creating directory key stash file: '/home/idsldap/idsslapd-idsldap/etc/ibmslapddir.ksf'.
GLPKEY012I Created directory key stash file: '/home/idsldap/idsslapd-idsldap/etc/ibmslapddir.ksf'.
```

Notice: I had to escape my special characters: \& , \( and \| or this would fail.

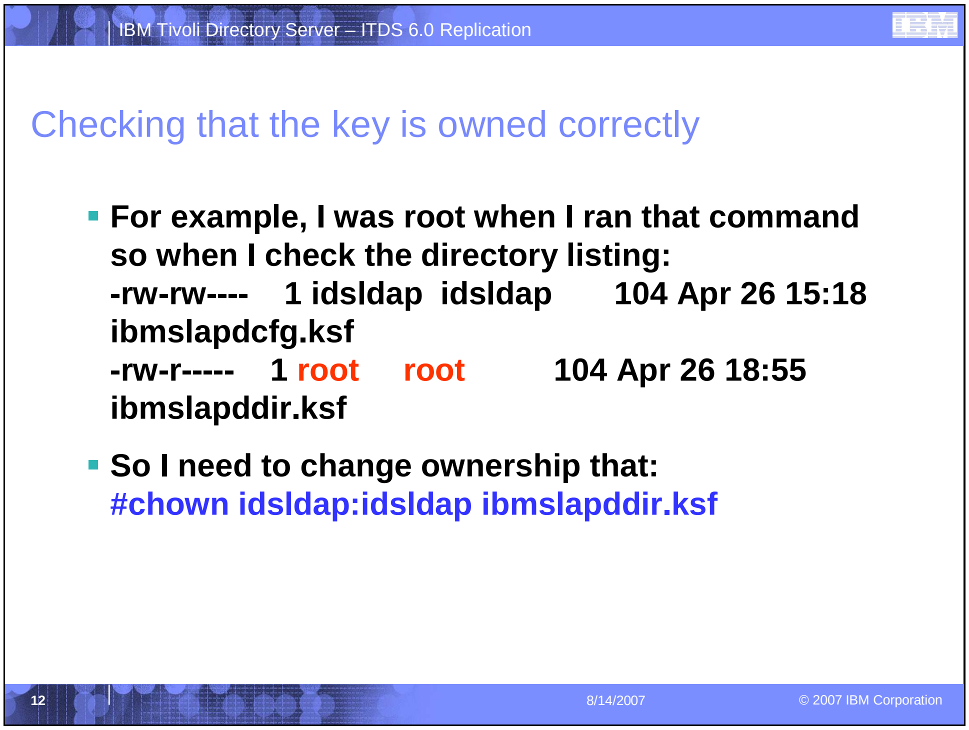

## Checking that the key is owned correctly

- For example, I was root when I ran that command so when I check the directory listing:

```
-rw-rw----    1 idsldap  idsldap       104 Apr 26 15:18
ibmslapdcfg.ksf
-rw-r-----    1 root     root          104 Apr 26 18:55
ibmslapddir.ksf
```

- So I need to change ownership that:

```
#chown idsldap:idsldap ibmslapddir.ksf
```

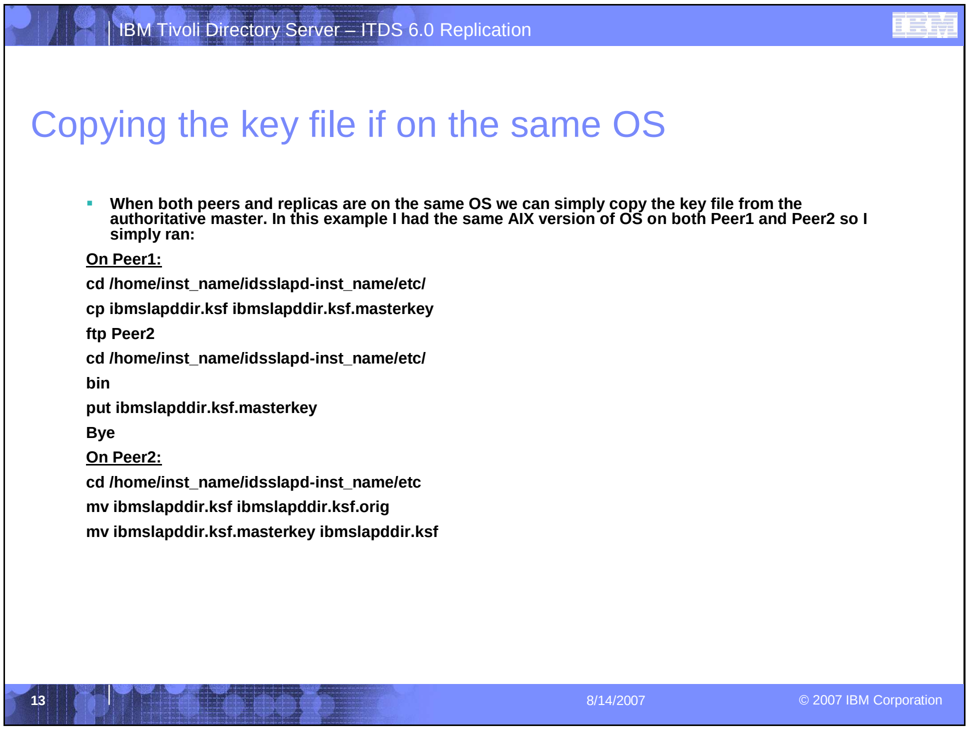

# Copying the key file if on the same OS

- When both peers and replicas are on the same OS we can simply copy the key file from the authoritative master. In this example I had the same AIX version of OS on both Peer1 and Peer2 so I simply ran:

On Peer1:

```
cd /home/inst_name/idsslapd-inst_name/etc/
cp ibmslapddir.ksf ibmslapddir.ksf.masterkey
ftp Peer2
cd /home/inst_name/idsslapd-inst_name/etc/
bin
put ibmslapddir.ksf.masterkey
Bye
```

On Peer2:

```
cd /home/inst_name/idsslapd-inst_name/etc
mv ibmslapddir.ksf ibmslapddir.ksf.orig
mv ibmslapddir.ksf.masterkey ibmslapddir.ksf
```

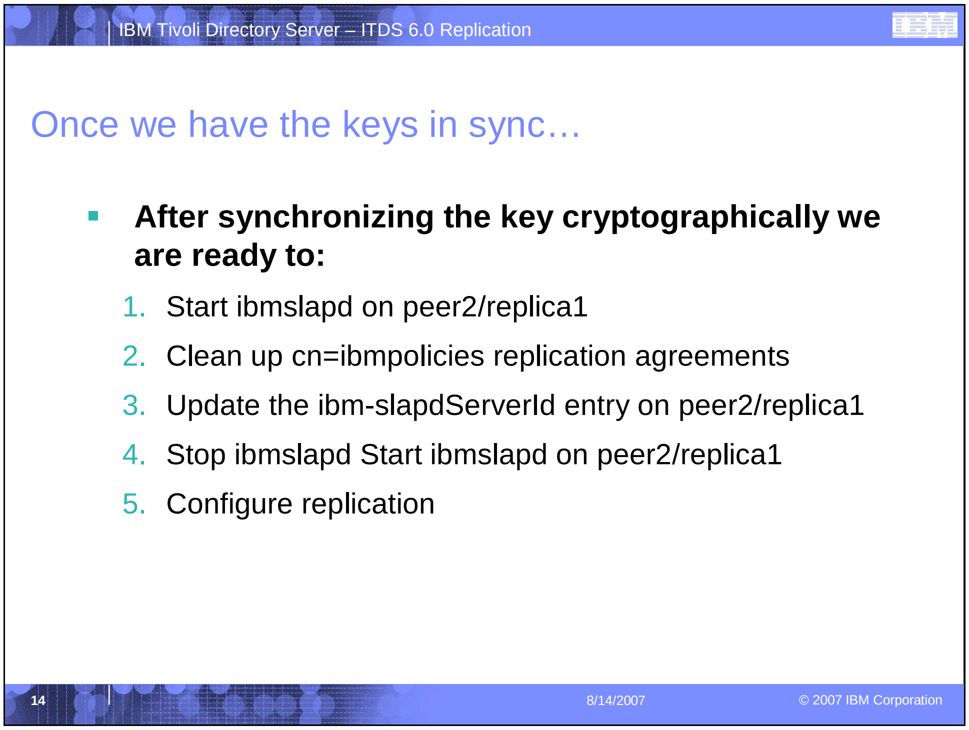

## Once we have the keys in sync…

- **After synchronizing the key cryptographically we are ready to:**

  1. Start ibmslapd on peer2/replica1
  2. Clean up cn=ibmpolicies replication agreements
  3. Update the ibm-slapdServerId entry on peer2/replica1
  4. Stop ibmslapd Start ibmslapd on peer2/replica1
  5. Configure replication

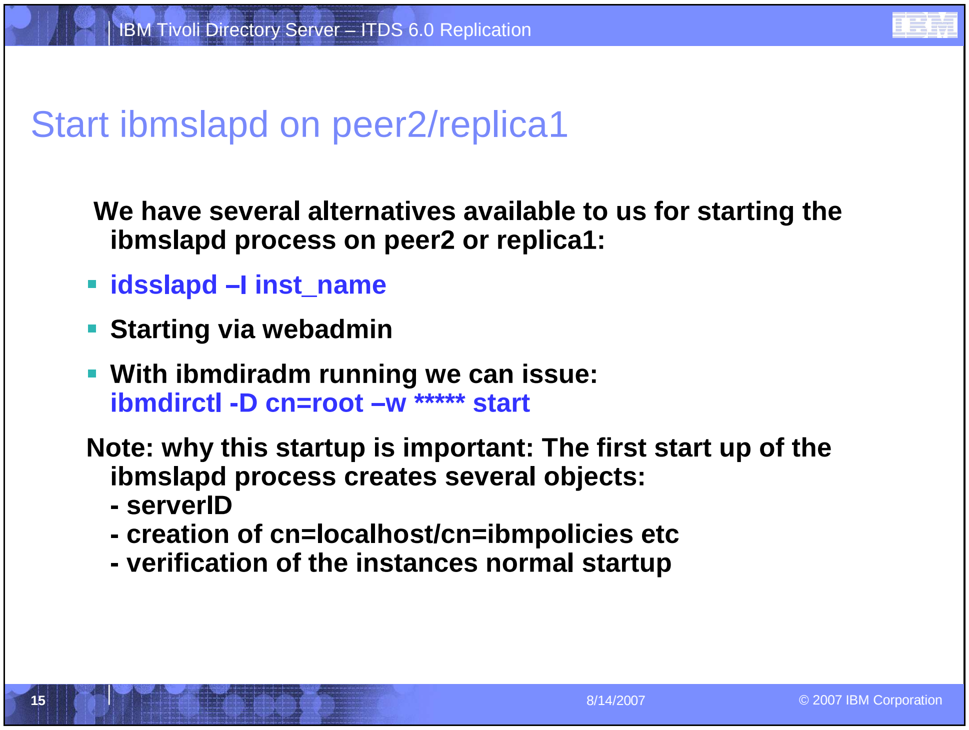

## Start ibmslapd on peer2/replica1

**We have several alternatives available to us for starting the ibmslapd process on peer2 or replica1:**

- **idsslapd –I inst_name**
- **Starting via webadmin**
- **With ibmdiradm running we can issue:**
  **ibmdirctl -D cn=root –w ***** start**

**Note: why this startup is important: The first start up of the ibmslapd process creates several objects:**
**- serverID**
**- creation of cn=localhost/cn=ibmpolicies etc**
**- verification of the instances normal startup**

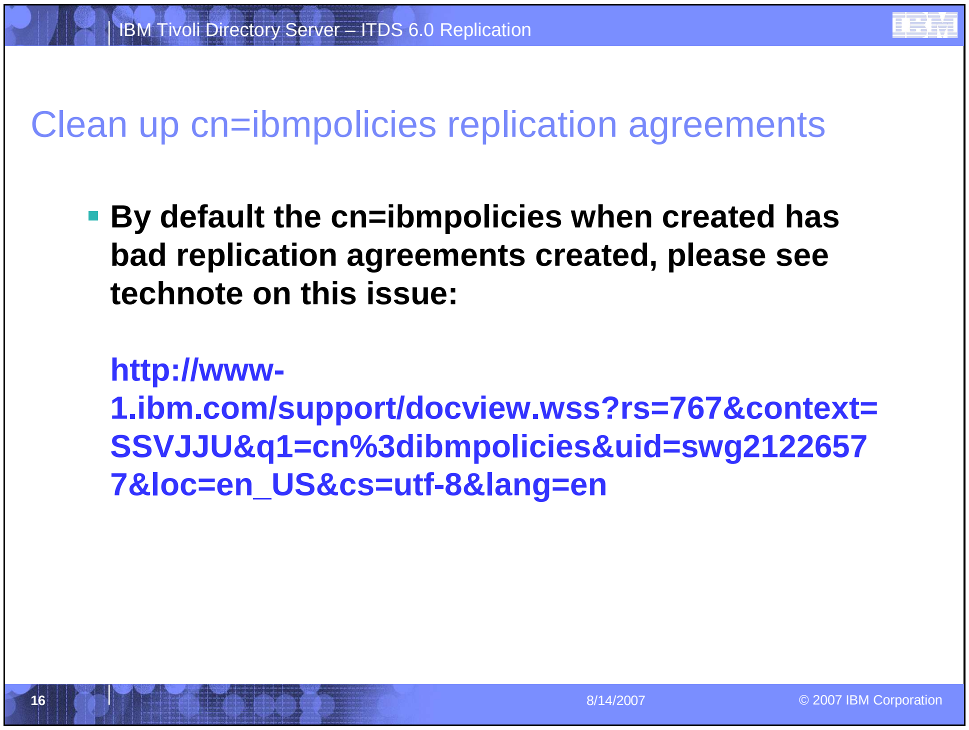

## Clean up cn=ibmpolicies replication agreements

- **By default the cn=ibmpolicies when created has bad replication agreements created, please see technote on this issue:**

  **http://www-1.ibm.com/support/docview.wss?rs=767&context=SSVJJU&q1=cn%3dibmpolicies&uid=swg21226577&loc=en_US&cs=utf-8&lang=en**

# Cleaning up cn=ibmpolicies replication agreements

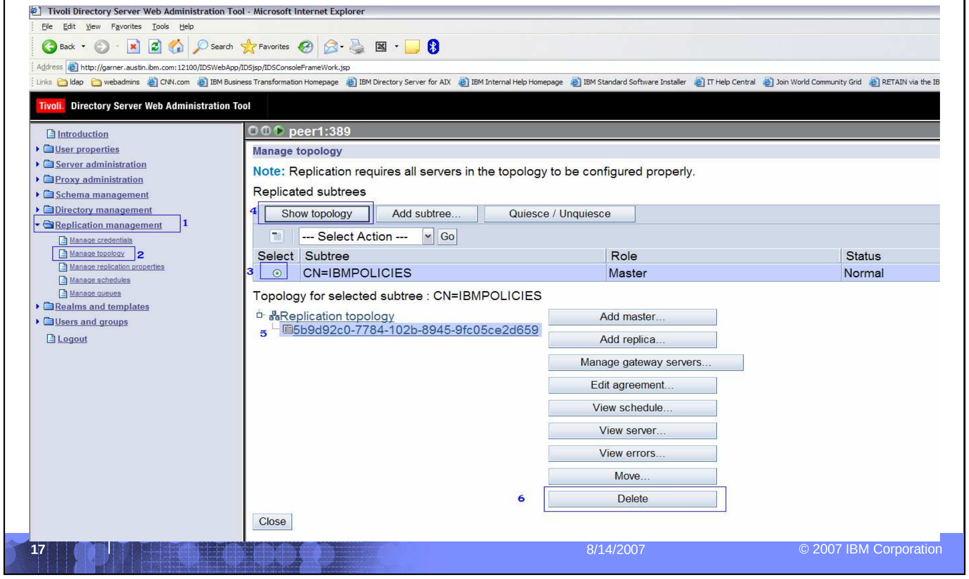
Tivoli Directory Server Web Administration Tool - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://garner.austin.ibm.com:12100/IDSWebApp/IDSjsp/IDSConsoleFrameWork.jsp
Links ldap webadmins CNN.com IBM Business Transformation Homepage IBM Directory Server for AIX IBM Internal Help Homepage IBM Standard Software Installer IT Help Central Join World Community Grid RETAIN via the IB
Tivoli Directory Server Web Administration Tool
Introduction
User properties
Server administration
Proxy administration
Schema management
Directory management
Replication management
1
Manage credentials
Manage topology
2
Manage replication properties
Manage schedules
Manage queues
Realms and templates
Users and groups
Logout
peer1:389
4
Show topology
Add subtree...
Quiesce / Unquiesce
--- Select Action ---
Go
3
5
Add master...
Add replica...
Manage gateway servers...
Edit agreement...
View schedule...
View server...
View errors...
Move...
6
Delete
Close

Manage topology

Note: Replication requires all servers in the topology to be configured properly.

Replicated subtrees

| Select | Subtree | Role | Status |
|---|---|---|---|
| ● | CN=IBMPOLICIES | Master | Normal |

Topology for selected subtree : CN=IBMPOLICIES

Replication topology
5b9d92c0-7784-102b-8945-9fc05ce2d659

# Cleaning up cn=ibmpolicies replication agreements

- **The webadmin will then prompt you to make sure you want to delete the agreement.**

  **Click ok**

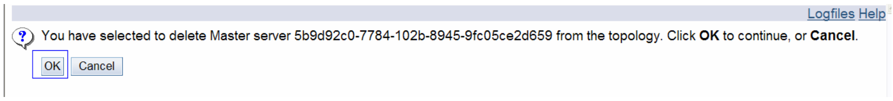

You have selected to delete Master server 5b9d92c0-7784-102b-8945-9fc05ce2d659 from the topology. Click **OK** to continue, or **Cancel**.

# Cleaning up cn=ibmpolicies replication agreements

- **We can now remove the cn=ibmpolicies replication topology totally:**

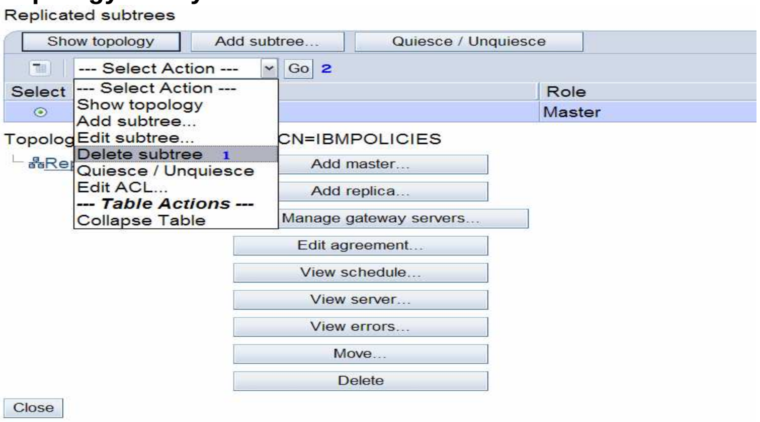

# The webadmin always wants to make sure…

- **Before most tasks will complete in ITDS client or the web admin there is usually a prompt making sure we want to accomplish the task. The same is true for the removal of this replication topology:**

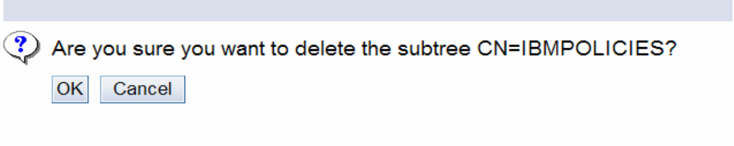
Are you sure you want to delete the subtree CN=IBMPOLICIES?

**Click on OK**

## What we are left with is a completely clean replication topology:

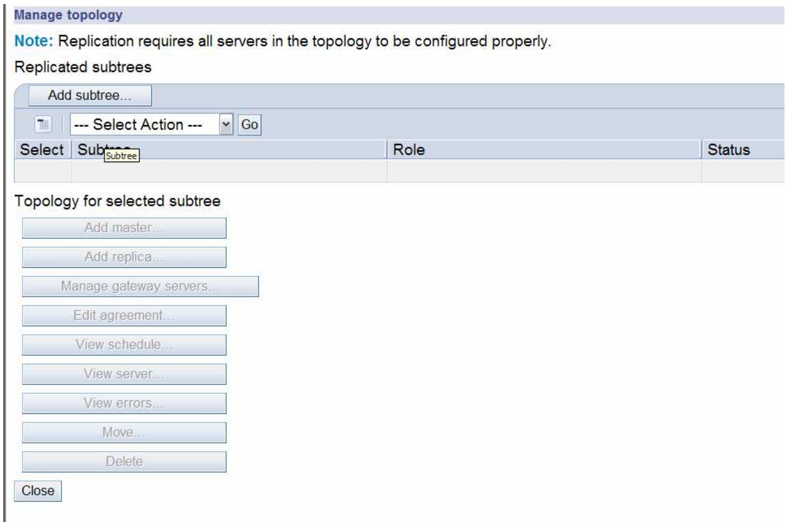

**Manage topology**

Note: Replication requires all servers in the topology to be configured properly.

Replicated subtrees

| Select | Subtree | Role | Status |
|---|---|---|---|

Topology for selected subtree

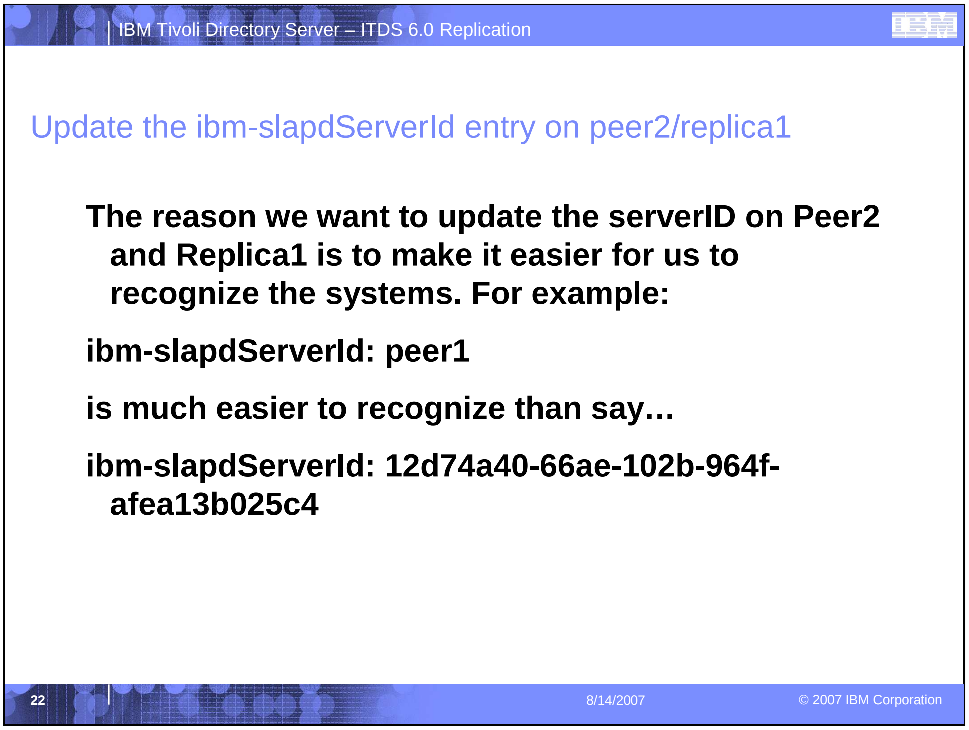

## Update the ibm-slapdServerId entry on peer2/replica1

**The reason we want to update the serverID on Peer2 and Replica1 is to make it easier for us to recognize the systems. For example:**

**ibm-slapdServerId: peer1**

**is much easier to recognize than say…**

**ibm-slapdServerId: 12d74a40-66ae-102b-964f-afea13b025c4**

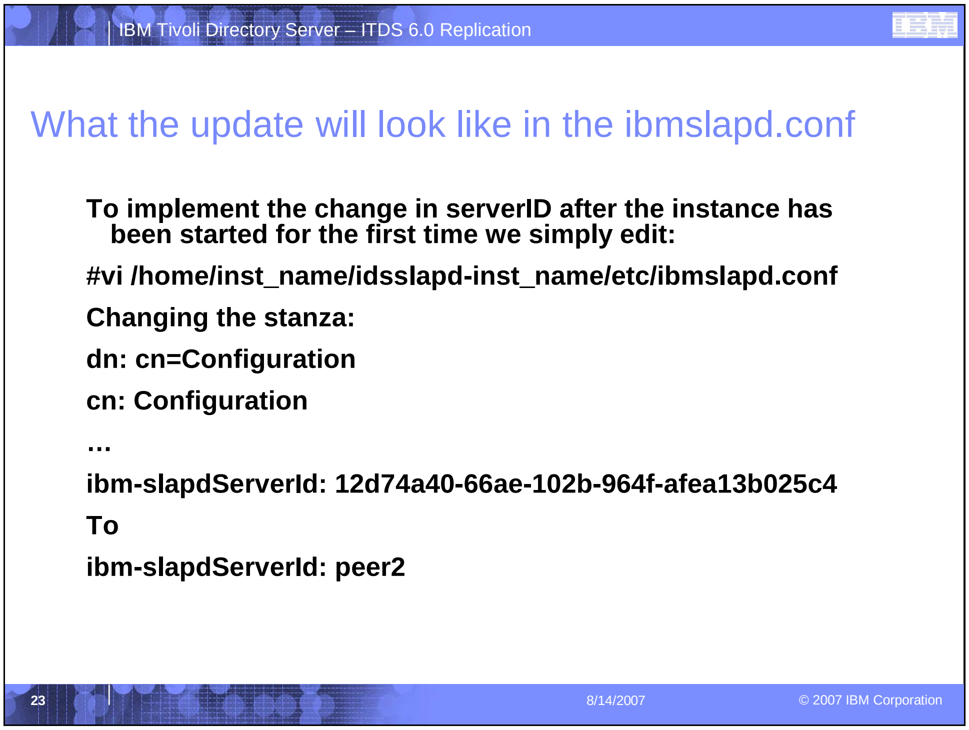

## What the update will look like in the ibmslapd.conf

**To implement the change in serverID after the instance has been started for the first time we simply edit:**

```
#vi /home/inst_name/idsslapd-inst_name/etc/ibmslapd.conf
```

**Changing the stanza:**

```
dn: cn=Configuration
cn: Configuration
…
ibm-slapdServerId: 12d74a40-66ae-102b-964f-afea13b025c4
```

**To**

```
ibm-slapdServerId: peer2
```

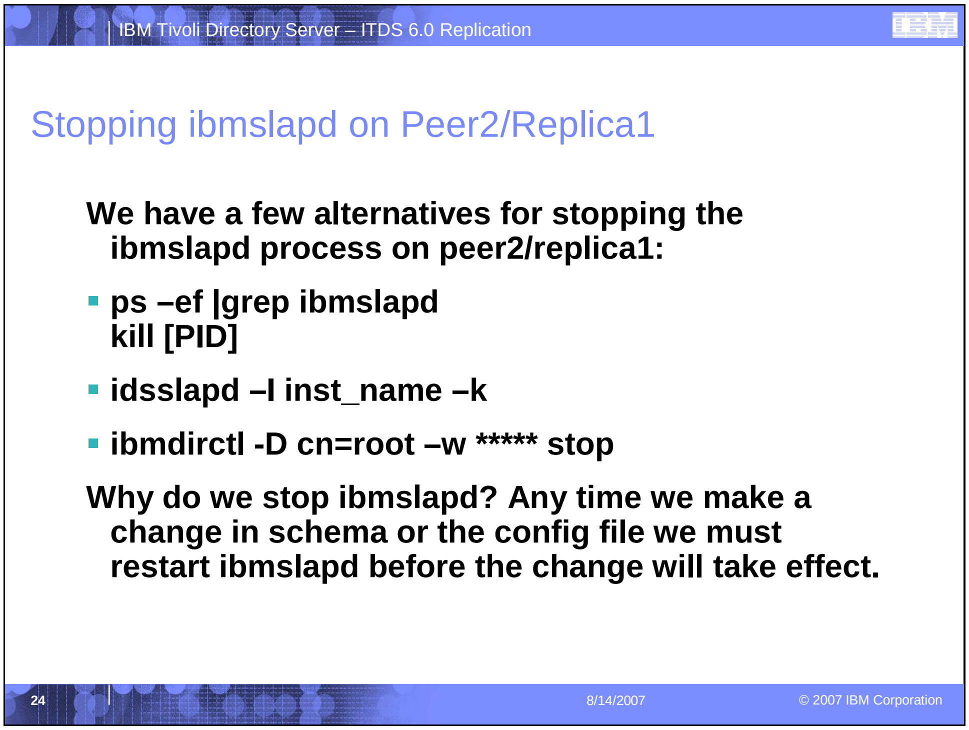

## Stopping ibmslapd on Peer2/Replica1

**We have a few alternatives for stopping the ibmslapd process on peer2/replica1:**

- **ps –ef |grep ibmslapd**
  **kill [PID]**
- **idsslapd –I inst_name –k**
- **ibmdirctl -D cn=root –w ***** stop**

**Why do we stop ibmslapd? Any time we make a change in schema or the config file we must restart ibmslapd before the change will take effect.**

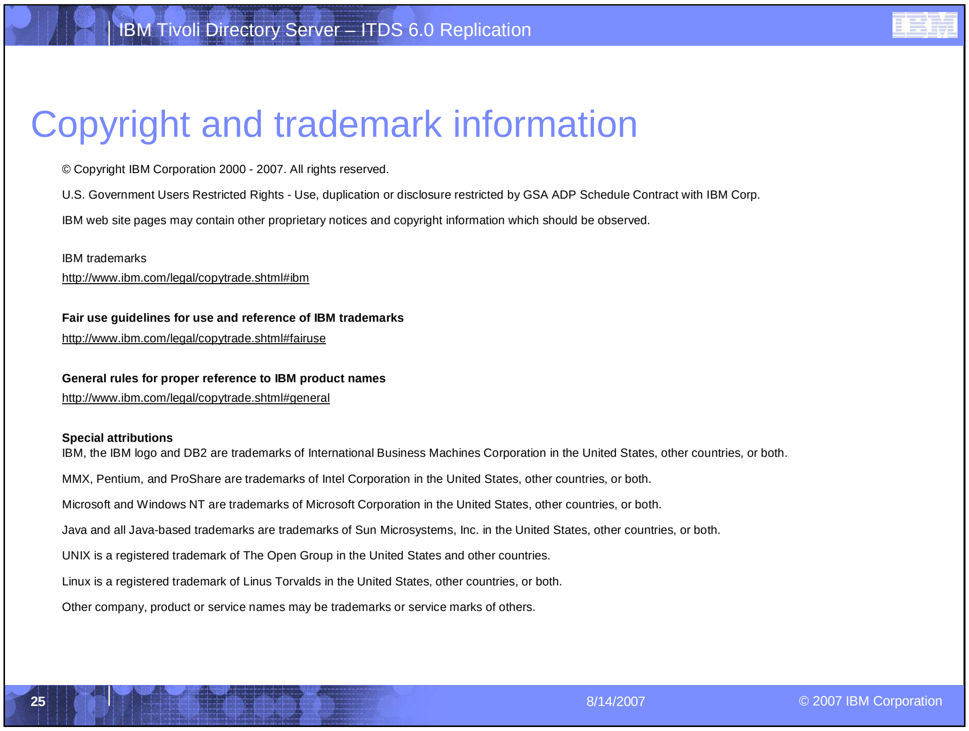

# Copyright and trademark information

IBM trademarks
http://www.ibm.com/legal/copytrade.shtml#ibm

**Fair use guidelines for use and reference of IBM trademarks**
http://www.ibm.com/legal/copytrade.shtml#fairuse

**General rules for proper reference to IBM product names**
http://www.ibm.com/legal/copytrade.shtml#general

**Special attributions**