IBM Unified V7000 1.3.2.3 PTF Feb 2013

Contents

Chapter 1. Overview of Information Center changes

The following topics are either new or changed. The "LDAP bind user requirements" section is new. Minor changes occurred to the rest of the topics shown in this PDF.

Chapter 2. LDAP bind user requirements

When configuring an Storwize V7000 Unified with LDAP as the authentication mechanism, the Storwize V7000 Unified system needs to connect to the LDAP server using an administrative user ID and password, referred to as bind user.

The bind user information and password are stored on the management node and are accessible by the root user. It is recommended, assuming that the root user is aware of internal details of the system, that bind user be given just enough privileges required by the storage system, to mitigate any security concerns.

This bind user must at least have permission to query users and groups that are defined in the LDAP server to allow storage system to authenticate these users to allow data access. The bind user information (binddn) is also used by Samba (CIFS server) while making LDAP queries to retrieve required information from the LDAP server.

Note: In the following sections it is assumed that the user account for the bind user already exists in the LDAP directory server. The bind user distinguished name (also known as dn) used in the following examples is `uid=ibmbinduser,ou=people,dc=ldapserver,dc=com` and this name needs to be updated accordingly based on the bind user used with the Storwize V7000 Unified system.

OpenLDAP Server ACLs

The following example describes the privileges required for the bind user. Note that the example uses ACLs needed for other users as well for the sake of completeness. It is likely that a corporate directory server has those ACLs configured already and only the entries for bind user need to be merged correctly in the slapd configuration file (generally `/etc/openldap/slapd.conf` file on Linux systems), as shown in the following example. Follow the ACL ordering rules to ensure correct ACLs are applied.

```
### some attributes need to be readable so that commands like 'id user','getent' etc can answer co
access to attrs=cn,objectClass,entry,homeDirectory,uid,uidNumber,gidNumber,memberUid
by dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com" read

###The following will not list userPassword when ldapsearch is performed with bind user.
### Anonymous is needed to allow bind to succeed and users to authenticate, should be a pre-existi
access to attrs=userPassword
    by dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com"  auth
    by self write
    by anonymous auth
    by * none


### Storage system needs to be able to find samba domain account specified on the cfgldap command.
###It is strongly recommended that domain account is pre-created to ensure
###consistent access to multiple storage systems.
###Uncomment ONLY if you want storage systems to create domain account when it does not exist.
#access to dn.base="dc=ldapserver,dc=com"
#    by dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com " write
#    by * none

access to dn.regex="sambadomainname=[^,]+,dc=ldapserver,dc=com"
    by dn=" uid=ibmbinduser,ou=people,dc=ldapserver,dc=com" read
```

```
    by * none


### all samba attributes need to be readable for samba access
access to attrs=cn,sambaLMPassword,sambaNTPassword,sambaPwdLastSet,sambaLogonTime,sambaLogoffTime,sam
    by dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com" read
    by self read
    by * none

### Need write access to record bad failed login attempt
access to attrs=cn,sambaBadPasswordCount,sambaBadPasswordTime,sambaAcctFlags
by dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com" write

### Required to check samba schema
access to dn.base="cn=Subschema" by  dn="uid=ibmbinduser,ou=people,dc=ldapserver,dc=com" read
```

Tivoli Directory Server ACLs

The following example describes the privileges required for the bind user, when using Tivoli Directory Server. These ACLs are provided in the LDIF format and can be applied by submitting the **ldapmodify** command.

```
dn: dc=ldapserver,dc=com
changetype: modify
add: ibm-filterAclEntry
ibm-filterAclEntry:access-id:uid=ibmbinduser,ou=people,dc=ldapserver,dc=com:(objectClass=sambaSamAcc
-
add:ibm-filterAclEntry
ibm-filterAclEntry:access-id:uid=ibmbinduser,ou=people,dc=ldapserver,dc=com:(objectclass=sambaDomain
 dn:uid=ibmbinduser,ou=people,dc=ldapserver,dc=com
add:aclEntry
aclentry: access-id:uid=ibmbinduser,ou=people,dc=ldapserver,dc=com:at.cn:r:at.objectClass:r:at.homeD
### Storage system needs to be able to find samba domain account specified on the cfgldap command.
###It is strongly recommended that domain account is pre-created to ensure
###consistent access to multiple storage systems.
###Uncomment ONLY if you want storage systems to create domain account when it does not exist.
 dn: dc=ldapserver,dc=com
changetype: modify
add:ibm-filterAclEntry
ibm-filterAclEntry:access-id:uid=ibmbinduser,ou=people,dc=ldapserver,dc=com:(objectclass=domain):obj
```

See IBM Tivoli Directory Server Administration Guide for information about applying these ACLs on Tivoli Directory Server.

Updating LDAP user information with Samba attributes

To use Samba accounts, update LDAP user information with unique Samba attributes.

The following sample LDIF file shows the minimum required samba attributes:

```
dn: cn=cifsuser,ou=People,dc=ibm,dc=com
changetype: modify
add : objectClass
objectClass: sambaSamAccount
-
add: sambaSID
sambaSID: (S-1-0-41200)
-
add:sambaPasswordHistory
sambaPasswordHistory: (56 zeroes)
-
add:sambaNTPassword
sambaNTPassword: (valid samba password hash )
-
add:sambaPwdLastSet
sambaPwdLastSet: 1263386096
-
add:SambaAcctFlags
sambaAcctFlags: [U         ]
```

Note: Attributes must be separated with a dash as the first and only character on a separate line.

Perform the following steps to create the values for sambaNTPassword, sambaPwdLastSet and SambaAcctFlags, which must be generated from a perl module:

1. Download the module from http://search.cpan.org/~bjkuit/Crypt-SmbHash-0.12/SmbHash.pm. Create and install the module by following the README file.
2. Use the following perl script to generate the LM and NT password hashes:

```
# cat /tmp/Crypt-SmbHash-0.12/gen_hash.pl
#!/usr/local/bin/perl
use Crypt::SmbHash;
$username = $ARGV[0];
$password = $ARGV[1];
if ( !$password ) {
        print "Not enough arguments\n";
        print "Usage: $0 username password\n";
        exit 1;
}
$uid = (getpwnam($username))[2];
my ($login,undef,$uid) = getpwnam($ARGV[0]);
ntlmgen $password, $lm, $nt;
printf "%s:%d:%s:%s:[%-11s]:LCT-%08X\n", $login, $uid, $lm, $nt, "U", time;
```

3. Generate the password hashes for any user as in the following example for the user `test01`:

```
# perl gen_hash.pl cifsuser test01

:0:47F9DBCCD37D6B40AAD3B435B51404EE:82E6D500C194BA5B9716495691FB7DD6:

[U         ]:LCT-4C18B9FC
```

   Note: The line above is login name, uid, LM hash, NT hash, flags and time, respectively, with each field separated from the next by a colon. The login name and uid are omitted because the command was not run on the LDAP server.
4. Use the information from step 3 to update the LDIF file in the format provided by the example at the beginning of this topic.
   - To generate the `sambaPwdLastSet` value, use the hexadecimal time value from step 3 after the dash character and convert it into decimal.
   - A valid samba SID is required for a user to enable that user's access to an Storwize V7000 Unified share. To generate the samba SID, multiply the user's UID by 2 and add 1000. The users's SID should contain the samba SID from the `sambaDomainName` which is either generated or picked up from LDAP server, if it already exists. The following attributes for `sambaDomainName` LDIF entry are required:

```
dn: sambaDomainName=(Storwize V7000
Unified system name),dc=ibm,dc=com
sambaDomainName: (Storwize V7000
Unified system name)
sambaSID: S-1-5-21-1528920847-3529959213-2931869277
sambaPwdHistoryLength: 0
sambaMaxPwdAge: -1
sambaMinPwdAge: 0
```

   This entry can be created by the LDAP server administrator using either of the following two methods:
   - Write and run a bash script similar to the following example:

```
sambaSID=
for num in 1 2 3 ;do
    randNum=$(od -vAn -N4 -tu4 < /dev/urandom | sed -e 's/ //g')
    if [ -z       "$sambaSID" ];then
        sambaSID="S-1-5-21-$randNum"
    else
        sambaSID="${sambaSID}-$    {randNum}"
    fi
done
echo $sambaSID
```

     Then use the samba SID generated to create the LDIF file. The `sambaDomainName` must match the Storwize V7000 Unified system name.
   - Submit the **cfgldap** Storwize V7000 Unified CLI command, which creates the `sambaDomainName` if it does not exist.

   The sambaSID for every user should have the following format: `(samba SID for the domain)-(userID*2+1000)` For example: S-1-5-21-1528920847-3529959213-2931869277-1102

Note: To enable access to more than one Storwize V7000 Unified system, the domain SID prefix of all of the Storwize V7000 Unified systems must match. If you change the domain SID for an Storwize V7000 Unified system on the LDAP server, you must restart CTDB on that Storwize V7000 Unified system for the change to take effect.

5. Submit the ldapmodify command to update the user's information as shown in the following example:

```
# ldapmodify -h localhost -D cn=Manager,dc=ibm,dc=com -W -x -f /tmp/samba_user.ldif
```

Setting up external LDAP server prerequisites

Before configuring the Storwize V7000 Unified environment for external server LDAP integration, several external LDAP server prerequisites must be met.

- The external LDAP server must already be configured.
- Obtain in advance the administrative information for the external LDAP authentication server, such as the administration account, password, SSL certificate, and Kerberos keytab file.
- Ensure that the Storwize V7000 Unified file modules have proper connectivity to the external LDAP server, and vice versa.
- On each of the Storwize V7000 Unified file modules, you must synchronize the time with the external LDAP authentication server. Authentication does not work if the times on the Storwize V7000 Unified file modules and the external LDAP authentication server are not synchronized.
- Optionally, enable SSL or TLS encryption on the external LDAP server. Details on configuring SSL or TLS encryption on the server can be obtained from the OpenLDAP Administrator's Guide.
- To access CIFS shares, LDAP user information must be updated with unique Samba attributes in addition to the attributes that are stored for a normal LDAP user. Ensure that these required Samba attributes are present in the LDAP user entries. See “Updating LDAP user information with Samba attributes” on page 4.
- A special administrative user for the Storwize V7000 Unified system must be created on the external LDAP server. This user might not have permission to create users; however, this user must at a minimum have permission to query users and groups that are created in the external LDAP server. The user information is used as binddn by Samba while making LDAP queries, as shown in the following example LDAP query:

  ```
  ldapsearch -x -D "cn=Manager,dc=example,dc=com" -w <password> "<query-filter>"
  ```

  The example LDIF for this user will be similar to:

  ```
  dn: cn=Manager,dc=example,dc=com
  objectclass: organizationalRole
  cn: Manager
  ```

  For information on the privileges required by such special administrative users, see Chapter 2, “LDAP bind user requirements,” on page 3.
- Ensure not to have the same username for different organizational units of the LDAP server configured with the Storwize V7000 Unified system.

Chapter 3. Upgrade recovery

This section covers the recovery procedures that relate to upgrade.

Error codes and recommendations when running the applysoftware command

If any errors are posted after you issue the **applysoftware** command, see Table 1 and take the described course of action. Follow these guidelines:

1. Follow the actions in the order presented.
2. After each recommended fix, restart the upgrade by issuing the **applysoftware** command again. If the action fails, try the next recommended action.
3. If the recommended actions fail to resolve the issue, call the IBM Support Center.

Table 1. Upgrade error codes from using the **applysoftware** *command and recommended actions*

Error Code	The **applysoftware** command explanation	Action
EFSSG1000I	The command completed successfully.	None.
EFSSG4100	The command completed successfully.	None.
EFSSG4101	The required parameter was not specified.	Check the command and verify that the parameters are entered correctly.
EFSSG4101A	The applysoftware command returned required parameter not specified.	
EFSSG4102	The software package does not exist.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4102A	The applysoftware command returned software package does not exist	
EFSSG4103	The software package is not valid.	The package might be corrupt. If this problem persists, download a new package and try again.
EFSSG4103A	The applysoftware command returned invalid software package return code.	
EFSSG4104	An unexpected return code.	Call your next level of support.

Table 1. Upgrade error codes from using the **applysoftware** *command and recommended actions (continued)*

Error Code	The **applysoftware** command explanation	Action
EFSSG4105	Unable to mount the USB flash drive.	Run `umount /media/usb`, then remove the USB flash drive. Reinsert the USB flash drive. If the error persists, remove the USB flash drive and reboot. After the system reboots, reinsert the USB flash drive.
EFSSG4105C	The applysoftware command returned unable to mount USB.	
EFSSG4106A	The applysoftware command returned that there is insufficient system file system space.	
EFSSG4153	The required parameter was not specified.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4154	You must start on primary management node mgmt001st001.	Switch to the other node and try the command again.
EFSSG4154A	The applysoftware returned must start on primary management node mgmt001st001.	
EFSSG4155	Unable to mount USB flash drive.	Back up to a USB flash drive. Enter `# backupmanagementnode --unmount /media/usb`. Remove the USB flash drive and insert again. If the error persists, remove the USB flash drive and reboot. When the system is running, insert the USB flash drive again.
EFSSG4155I	The applysoftware command returned upgrade is already running.	
EFSSG4156	The specified International Organization for Standardization (ISO) does not exist.	Verify that the file actually exists where specified. Also verify that the command is passing the correct location parameters.
EFSSG4156A	The applysoftware command returned the specified ISO does not exist.	

Table 1. Upgrade error codes from using the **applysoftware** *command and recommended actions (continued)*

Error Code	The **applysoftware** command explanation	Action
EFSSG4157	The specific upgrade International Organization for Standardization (ISO) content is not valid.	The package might be corrupt. If this problem persists, download a new package and try again.
EFSSG4157I	The applysoftware command returned the specific upgrade ISO invalid content.	
EFSSG4158	The specific upgrade cannot be installed over the current version.	Check the upgrade documentation and verify that the level you are coming from is compatible with the level you are going to. If the upgrade level is not compatible, download the correct level and try again. If the upgrade level is compatible and the error persists, call the IBM Support Center.
EFSSG4158I	The applysoftware command returned the specific upgrade cannot be installed over the current version.	
EFSSG4159	The system is in an unhealthy state and the upgrade cannot start.	See Getting started troubleshooting. Determine if the system has issues.
EFSSG4159I	The applysoftware command returned that the system is in an unhealthy state and upgrade cannot start.	
EFSSG4160	The system has insufficient file system space.	At least 3 GB of space is required. Remove unneeded files from the /var file system.
EFSSA0201C	The license agreement has not been accepted.	

General upgrade error codes and recommended actions

If any errors are posted during the upgrade process, see Table 2 on page 10 and take the described course of action. If the error you see is not listed in this table, call the IBM Support Center. Follow these guidelines:

1. Follow the actions in the order presented.
2. After each recommended fix, restart the upgrade by issuing the **applysoftware** command again. If the action fails, try the next recommended action.
3. If the recommended actions fail to resolve the issue, call the IBM Support Center.

Table 2. Upgrade error codes and recommended actions

Error Code	Explanation	Action
019A	Yum update failed.	Contact IBM Remote Technical Support.
019B	Unable to remove StartBackupTSM task.	1. Check to see if management service is running on active node. If it is not, use **startmgtsrv** to start. 2. Contact IBM Remote Technical Support.
019C	Unable to determine active management node.	1. Check to see if management service is running on active node. If it is not use **startmgtsrv** to start. 2. Contact IBM Remote Technical Support.
019D	Check the system health.	1. Use **lnode** to determine what this node is showing unhealthy. (CTDB or GPFS). Possibly reboot unhealthy node and wait for node to come back up. Then check health of node with **lsnode**. 2. Contact IBM Remote Technical Support.
019E	Internal error - cluster or node not provided	Contact IBM Remote Technical Support.
019F	CIM restart failed.	Contact IBM Remote Technical Support.
01A0	Failed to reboot.	Determine the cause of the failed reboot: 1. Check console of system if able. See if the system is hung in BIOs or during boot. 2. Check cabling of system. 3. Check light path diagnostic for error indications. . 4. Reboot the system from console and restart upgrade. 5. Contact IBM Remote Technical Support.
01A1	Internal upgrade error.	Contact IBM Remote Technical Support.
01A3	Unable to uninstall CNCSM callbacks.	Contact IBM Remote Technical Support.
01A4	Unable to stop backup jobs.	1. Check the status of the backups by typing `lsjobstatus -j backup`. 2. Attempt to stop backups by typing `stopbackup --all`. 3. Contact IBM Remote Technical Support.

Table 2. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01A5	Backup cron jobs are running.	1. Check the condition of tasks by typing `lstask -t cron`. 2. Attempt to remove the backup by typing `rmtask StartBackupTSM`. 3. Contact IBM Remote Technical Support.
01A6	Unable to install CNCSM callbacks.	Contact IBM Remote Technical Support.
01A7	Internal vital product data (VPD) error.	Contact IBM Remote Technical Support.
01A8	Check the health of management service.	1. Attempt to start the management service with **startmgtsrv** on active management node 2. Contact IBM Remote Technical Support.
01A9	Unable to stop performance collection daemon.	Contact IBM Remote Technical Support.
01AB	Internal upgrade error in node_setup_system.	Contact IBM Remote Technical Support.
01B1	Management node replication failed.	1. Follow the replication recovery procedure. See Resolving issues reported by lshealth for resolving the management node replication failure. 2. Contact IBM Remote Technical Support.
01B2	Unable to start performance collection daemon.	Contact IBM Remote Technical Support.
01B3	Failed to copy upgrade package to Storwize V7000.	This could be caused by a number of issues. Check **Monitoring > Events** under both the **block** tab and the **file** tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to Ethernet connectivity from file modules to the control enclosure.
01B4	Failed to start upgrade on Storwize V7000 with the **applysoftware** command.	This could be caused by a number of issues. Check **Monitoring > Events** under both the **block** tab and the **file** tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to Ethernet connectivity from file modules to the control enclosure.

Table 2. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01B5	Storwize V7000 multipaths are unhealthy.	Check the Fibre Channel connections to the system. Reseat Fibre Channel cables. For more information, see Fibre Channel connectivity between file modules and control enclosure.
01B6	Storwize V7000 vdisks are unhealthy as indicated by using the **lsvdisk** command.	See Control enclosure.
01B7	Failed to query status of Storwize V7000 upgrade by using the svcinfo **lssoftwareupgradestatus** command.	This could be caused by a number of issues. Check **Monitoring > Events** under both the **block** tab and the **file** tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to Ethernet connectivity from file modules to the control enclosure.
01B8	Failed to query status of Storwize V7000 nodes by using the **svcinfo lsnode** command.	See Control enclosure.
01B9	Failed to check the Storwize V7000 version.	This could be caused by a number of issues. Check **Monitoring > Events** under both the **block** tab and the **file** tab on the management GUI for an event that could have caused this error and follow the recommended action. If there is no obvious event that could have caused this error, refer to Ethernet connectivity from file modules to the control enclosure.
01BA	Unable to verify the correct software version.	1. Check the health of the storage controllers. 2. Contact IBM Remote Technical Support.
01BC	Check the health of storage controllers.	Contact IBM Remote Technical Support.
01BD	Unable to update software repository.	1. Ensure that the system is not under a heavy load. Restart the upgrade. 2. Contact IBM Remote Technical Support.
01BE	Unable to distribute upgrade callbacks.	1. Check on health of the cluster using **lshealth**. 2. Contact IBM Remote Technical Support.

Table 2. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01BF	Upgrade callback failed	1. Contact your customer advocate. Upgrade callbacks are custom steps placed on a system before the start of upgrade. 2. Contact IBM Remote Technical Support.
01C0	Asynchronous replication is running. Stop asynchronous replication and continue with the upgrade.	1. Stop asynchronous replication by typing `stoprepl gpfs0 --kill`. Asynchronous replication is considered active if in RUNNING or KILLING state. 2. Contact IBM Remote Technical Support.
01C1	Asynchronous replication failed to stop. Stop asynchronous replication and continue with the upgrade.	1. Stop asynchronous replication by typing `stoprepl gpfs0 --kill`. Asynchronous replication is considered active if in RUNNING or KILLING state. 2. Contact IBM Remote Technical Support.
01C2	Failed while checking for current running asynchronous jobs.	1. Attempt to check status of lsrepl. If this command is working restart upgrade. 2. Contact IBM Remote Technical Support.
01C3	Could not stop CTDB.	Contact IBM Remote Technical Support.
01C4	Unable to remove callbacks	Contact IBM Remote Technical Support.
01C5	Could not reinstall Lib_Utils.	Contact IBM Remote Technical Support.
01C6	Failed while running sonas_update_yum.	Contact IBM Remote Technical Support.
01C7	Unable to get list of cluster nodes.	Contact IBM Remote Technical Support.
01C8	Failed while running `cnrsscconfig`.	Contact IBM Remote Technical Support.
01C9	Unable to install CIM configuration.	Contact IBM Remote Technical Support.
01CA	Unable to get name of cluster.	Contact IBM Remote Technical Support.
01CB	Unable to install GPFS packages.	Contact IBM Remote Technical Support.
01CC	Could not install platform. Upgrade on target system.	Contact IBM Remote Technical Support.
01CD	Unable to mount GPFS file systems.	1. See Checking the GPFS™ file system mount on each file module 2. Restart upgrade and see if this was a transient issue. 3. Follow SONAS GPFS troubleshooting documentation. 4. Contact IBM Remote Technical Support.

Table 2. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01CE	Unable to update system security.	1. Restart upgrade and see if this was a transient issue. 2. Contact IBM Remote Technical Support.
01CF	Unable to configure node.	1. Pull both power supply cables from subject node. Wait 10 seconds, then plug back in. After the system restarts, try again. 2. Contact IBM Remote Technical Support.
01D0	Unable to disable call home.	Contact IBM Remote Technical Support.
01D1	Unable to enable call home.	Contact IBM Remote Technical Support.
01D2	Failed to stop GPFS.	1. Follow SONAS GPFS troubleshooting documentation. 2. Contact IBM Remote Technical Support.
01D3	Could not determine if backups are running.	1. Attempt to stop backups. 2. Type `lsjobstatus -j backup;echo $?`. If the return code is 0, start the upgrade again. 3. If the return code is any other number, contact IBM Remote Technical Support.
01D5	Storwize V7000 stalled_non_redundant.	Refer to Storwize V7000 documentation.
01D6	Storwize V7000 system stalled.	Refer to Storwize V7000 documentation.
01D8	CTDB cluster is unhealthy.	1. See Checking CTDB health. 2. Use **`lshealth`** or RAS procedures to determine unhealthy components. 3. Contact IBM Remote Technical Support.
01DA	GPFS system is unhealthy.	1. See Checking the GPFS file system mount on each file module. 2. Use **`lsnode -r`** to confirm GPFS is unhealthy. If node GPFS is healthy restart the upgrade. 3. Contact IBM Remote Technical Support.
01DB	Failed to stop performance center.	Contact IBM Remote Technical Support.
01DC	Failed to configure performance center.	Contact IBM Remote Technical Support.
01DD	Failed to start performance center.	Contact IBM Remote Technical Support.

Table 2. Upgrade error codes and recommended actions (continued)

Error Code	Explanation	Action
01DE	Unable to communicate with passive management node.	1. Ensure that the active mgmt node can communicate with the passive management node before restarting the upgrade. 2. Contact IBM Remote Technical Support.
01DF	Upgrade must be resumed from the other management node.	Restart upgrade from other management node. This might require that a failover be issued first.
01E0	HSM upgrade failed.	Contact IBM Remote Technical Support.

Chapter 4. Installing and configuring the HBA on System z hosts

The host bus adapters (HBAs) for a System z® host must be ordered as features and they are either factory-installed when you order a new system or installed into an existing system by an IBM® service representative.

About this task

Perform the following steps to check the installation of the HBA and to configure the HBA to work with the Storwize® V5000 system:

Procedure

1. Ensure that FICON®, FICON Express®, FICON Express2, or FICON Express4 features are installed on your System z host.
2. Configure the HBA to run in FCP mode.

What to do next

See the following IBM website for additional information about FCP connectivity:

www.ibm.com/systems/z/hardware/connectivity/

Chapter 5. Recovery procedure: Adding additional capacity for offline compressed file systems

In this situation, the storage pool has run out of capacity. As a result, the file system is unmounted and has gone offline, which makes all I/O to the file system fail.

To recover from this situation, you can either add available MDisks to the pool, or if free MDisks are not available, you can make spare drives available to build a new array (MDisk) to add to the pool. However, because spare drives are automatically used as backup drives when other drives fail on the system, using a spare drive to recover an offline file system can prevent an automated recovery if another drive fails on the system. After the file system is brought back online and capacity deficiencies have been addressed, return the drive to use as a spare or add another drive to replace it as a spare. If you add a new drive, new drives must be added to the system.

Increasing capacity to the storage pool

If MDisks are available to provide extra capacity to the storage pool that the compressed file system uses, you can add MDisks to the pool or create more MDisk (arrays).

Add any available MDisks: If an MDisk has already been created but not assigned to a pool, complete these steps:

1. In the management GUI, select **Pools** > **MDisk by Pools**.
2. Select **Not in pool** to display all the available MDisks that are not currently allocated to a storage pool.
3. Right-click the MDisks that you want to add to the storage pool and select **Add to Pool**.
4. On the **Add to Pool** dialog, select the pool and click **Add to Pool**.
5. Verify that the MDisk was added to the selected pool by expanding the pool and ensuring that the added MDisk is displayed.

Add any available drives: If MDisks have not been configured from available internal drives, you can provision the available drives into existing storage pools by completing these steps:

1. In the management GUI, select **Pools** > **Internal Storage**.
2. Select **Configure Storage**.
3. On the **Configure Internal Storage** dialog, select **Select a different configuration** and complete the following steps:
   a. In the **Drive Class** field, select the drive class that is available based on the installed storage on the system.
   b. In the **Preset** field, select the RAID configuration for the storage you are configuring.
   c. Select **Optimize for capacity** to configure all available capacity.
   d. Verify the configuration and click **Next**.
   e. Click **Expand an existing pool** and select the storage pool that is used for compression.

4. Click **Finish**.

Using spare drives to add capacity to the storage pool

If drives are not available, you need to make spare drives available to add capacity to the storage pool, bring the file system back online, ensure capacity for the storage pool does not run out again, and return spare drives to the system.

Note: If you are unfamiliar with managing spare goals and spare disks, contact IBM support for guidance. Increasing capacity in this way is meant only as a short term solution to this problem. Further provisioning to permanently resolve capacity constraints can be conducted with the help of IBM service personnel who might recommend that additional drives be added to your system.
To use spare drives to add capacity to the storage pool and bring file systems back online, complete these steps:

1. **Mark a spare drive as a candidate drive:** When block storage is configured on the system, available drives are categorized based on their drive class and drive type. To provide for drive redundancy, some drives are mark as spares, which provide backup drives in the event of a drive failure. Other drives are marked as candidates, which means they can be used as capacity for block storage pools. To mark a spare drive as a candidate and make it available to the block storage pool, complete these steps:
   a. In the management GUI, select **Pools** > **Internal Storage**.
   b. From the list of drives that display, right-click a drive that is marked as a spare drive and select **Mark as...** > **Candidate**.

      Note: The **Use** column displays how a specific drive is used on the system.
   c. Click **OK**.
2. **Expand the storage pool:** After the spare drive has been marked as a candidate drive, you can expand the capacity of the block storage pool that is used for the offline file system.
   a. In the management GUI, select **Pools** > **Internal Storage**.
   b. Select **Configure Storage**.
   c. On the **Configure Internal Storage** dialog, select **Select a different configuration** and complete the following steps:
      1) In the **Drive Class** field, select the drive class of the candidate drive (former spare) that is available based on the installed storage on the system. Verify the correct number of disks is displayed.
      2) In the **Preset** field, select the RAID configuration for the storage you are configuring. If you are adding only one disk, the only RAID option is RAID0 which does not provide any data protection.
      3) Select **Optimize for capacity** to configure all available capacity.
      4) Verify the configuration and click **Next**.
      5) Click **Expand an existing pool** and select the storage pool that is used for compression.
3. **Check event logs to ensure all underlying volumes are back online.** Before bringing the file system back online ensure that all the errors for both the block volumes and the file system have been resolved by completing these steps:
   a. In the management GUI, select **Monitoring** > **Events** and select **Block**.
   b. Run the fix procedures in the recommended order for all events that are related to the block volume that is used by the file system.
   c. Select **File** and fix all errors that are related to the offline file systems.

4. **Bring file systems back online:** After the capacity has been added to the storage pool, bring the file system back online by completing these steps:
   a. In the management GUI, select **Files** > **File Systems**.
   b. Right-click the compressed file system that is offline and select **Mount**. If the file system does not come back online you may need to restart all of the disks that the file system uses Right-click the compressed file system that went offline and select **Start All Disks**.
5. **Prevent the file system from running out of capacity again:**

   First ensure that you have free capacity at least the size of the real capacity of the temporary drive that you are adding.

   To decrease the file system capacity, you can remove the disks (NSD) and the corresponding mapping to block volumes to force migration of the data to other NSDs, thus freeing up space on the file system. To remove an NSD, contact IBM Remote Technical Support.
6. **Return spare drives to the system:** To ensure that drive redundancy is not compromised, spare drives that were used to bring the offline file systems back online need be replaced either by returning the original drive back to its spare use or by adding a new drive to the system. Ensure that the file system capacity has been decreased accordingly before returning the spare drives to the systems. To return the drive back to its spare use, complete these steps:
   a. In the management GUI, select **Pools** > **Internal Storage**.
   b. From the list of drives that display, ensure that no MDisk are associated with the drive. If the drive is associated with MDisk, select **Pools** > **MDisks by Pool**. Right-click the MDisk and select **Remove from Pool**.
   c. In the management GUI, select **Pools** > **Internal Storage**.
   d. From the list of drives that display, right-click a drive you marked as a candidate in Step 1 and select **Mark as...** > **Spare** .
   e. Click **OK**.

   To add additional drives to the system, complete these steps:
   a. Acquire additional drives from IBM or vendor.
   b. Install drives into available drive slots on the enclosure. See Installing a hot-swap hard disk drive.
   c. After the drives are available, select **Pools** > **Internal Storage**.
   d. From the list of drives that display, right-click the new drive and select **Mark as...** > **Spare**.

Chapter 6. resumenode

Resume a list of nodes.

Name

resumenode - Resume a list of nodes.

Syntax

```
resumenode nodeNames [-c { clusterID | clusterName }]
```

Parameters

nodeName

Lists the nodes which will be resumed in a comma-separated list. Define them with the fully qualified domain name, short name or IP address.

Using unlisted arguments can lead to an error.

Options

-c, --cluster { *clusterID* | *clusterName* }

Selects the cluster for the operation. Use either the clusterID or the clusterName to identify the cluster. Optional.

Using unlisted options can lead to an error.

Description

The **resumenode** command resumes the node or nodes referred by the nodeNames argument. It enables a disabled node or unbans a banned node. The resumed node might be participating in the cluster and client records for the clustered trivial database (CTDB). It takes back its IP address and starts hosting services. The IP address reallocation might cause disruption to connected hosts.

Diagnostics

0 No error - The command was successful.

1 Syntax error - Unknown option.

8 Command error - An internal error occurred while executing the command.

9 Invalid object - The given cluster or node is unknown.

Copyright

See also

addnode, suspendnode, lsnode

Chapter 7. Adding File Modules to a Storwize V7000 System

You can now add Storwize V7000 Unified file modules to a Storwize V7000 system.

See the Adding Storwize V7000 Unified File modules to an Existing Storwize V7000 System PDF for more information.

Index