

An Oracle White Paper

December 2010

ACFS File System Replication A How to Setup Guide

ORACLE

Introduction	2
Overview	3
Primary file system	3
Standby file system	3
Planning for ACFS Replication.....	4
Tagging Considerations	4
Space Management for Replication.....	5
Setting up Replication	5
Admin User Setup.....	6
File System Setup	7
Network Setup.....	7
Generate Oracle Network files	7
Set Network Tuanbles	10
Validating Network Configuration.....	10
Replication Configuration and Initiation	11
Initializing Standby File system.....	12
Initializing the Primary File System.....	12
Pause and Resume Replication	15
Termination of Replication.....	17

Introduction

In Oracle Release 11.2.0.2 the ACFS File System Replication feature was introduced. ACFS replication enables replication of an ACFS file system across a network to a remote site. This capability is useful for providing disaster recovery capability. Similarly to Data Guard, which replicates databases by capturing database redo, ACFS Replication captures ACFS file system changes on a primary file system and transmits these changes to a standby file system.

ACFS Replication leverages OracleNet and the NETWORK_FILE_TRANSFER PL/SQL package for transferring replicated data from a primary node to the standby file system node. ACFS replication is currently only supported on “Grid Infrastructure for Cluster” as selected on the Oracle Installer. ACFS Replication is not supported on “Grid Infrastructure for a Standalone Server.” However, Grid Infrastructure for a Cluster can be installed on a single node by supplying the necessary information for a single node during installation.

The combination of Oracle Real Application Clusters, Data Guard and ACFS Replication provides comprehensive site and Disaster Recovery policies for all files inside and outside the database.

This paper is designed to be a guide for System Administrators or DBAs that will be managing ACFS and ACFS replication. For background information on implementing ACFS please review the Storage Administrator’s Guide or MOS Doc ID 948187.1.

Overview

Primary file system

The source ACFS file system is referred to as a primary file system and the target ACFS file system as a standby file system. For every primary file system there can be only be one standby file system. ACFS Replication captures, in real-time, file system changes on the primary file system and saves them in files called replication logs (rlogs). Rlogs are stored in the `.ACFS/repl` directory of the file system that is being replicated. If the primary node is part of a multi-node cluster, then all rlogs (1 rlog per node) created at a specific instance are collectively called a “cord”. Rlogs combined into a cord are then transmitted to the standby node. The cord is then used to update the standby file system.

ACFS replicates all changes written to disk. However, unless data is written synchronously, data written to files is first buffered in a cache before being flushed, then at a later point in time to disk. ACFS guarantees that when data is committed to disk it will also be written to the standby file system.

Current restrictions (11.2.0.2)

- The minimum file system size that can be replicated is 4GB.
- ACFS currently supports a maximum of 8 node cluster for the primary file system
- ACFS cannot use encryption or security for replicated file systems
- ACFS replication is available only for Linux and Windows systems
- Cascading standbys is not currently supported
- ACFS standby file system must be empty before initiating

Standby file system


Replication logs are asynchronously transported to the node hosting the standby file system.

Replication logs are then read and applied to the standby file system. When the replication logs have been successfully applied to the standby file system, they are deleted on both the primary and standby file systems. Note, the standby file system is a read-only file system. One use-case for the standby file system is that it can be the source for backups.

Planning for ACFS Replication

Examples described in this white paper assume that Grid Infrastructure software has been installed on nodes hosting the ACFS filesystem, and that the ADVM volumes and ACFS file systems are mounted. Note that the primary and standby file systems must be of same OS, architecture and endianness. However, the primary and standby sites can have differing configurations; i.e., primary can be multi-node cluster and standby can be single node cluster. If a standby node is used for disaster recovery purposes, then it is recommended that the standby node have a similar configuration such as cluster configuration.

There are no rigid primary or standby node roles; i.e. a primary node can provide the role of primary for one file system and also provide the role of standby for another file system. However, for simplicity, this paper will use the term ‘primary node’ to indicate the node hosting the primary filesystem, and ‘standby node’ for the node hosting the standby filesystem.



The above configuration represents the system used in the following examples. With respect to replication, some commands, such as “acfsutil” must be executed with root privileges. Other commands, such as ‘sqlplus’, are issued from the Oracle user id. In the examples, user id is shown with the command prompt.

Tagging Considerations

ACFS tagging is an important adjunct to ACFS replication. Rather than replicating an entire file system, ACFS tagging enables a user to select specific files and directories for replication. Using tagging with ACFS replication requires that a replication tag be specified when replication is first initiated on the primary node. Tagging with replication cannot be implemented after replication has been initiated. To begin tagging after replication has been initiated requires that replication is first terminated and then restarted with a tagname. ACFS implements tagging with extended attributes, thus some editing tools and backup utilities do not retain these extended attributes of the original file by default. Please review the ACFS Tagging section of the Storage Administrators guide for the list of common utilities and their respective switch settings, so that ACFS tag names are preserved on the original file.

Before implementing ACFS replication, it is important to determine how and what will be replicated; i.e., will all file system data be replicated, certain directories or only specific ACFS tagged files. This choice may impact file system sizing.

ACFS tagging assigns a common naming attribute to a group of files. ACFS Replication uses this tag to filter files with unique tag names for remote file system replication. Tagging enables data or attribute based replication. For more information on Tagging, please refer to the Oracle's Storage Administrator's Guide.

The following example illustrates recursively tagging all files of the /acfs directory with "reptag" tag.

```
[root@node1 ~]# /sbin/acfsutil tag set -r reptag /acfs
```

Keep in mind that the tags specified on the init command line need not be applied to files at the time of the initialization. For example, you can replicate files with tags Chicago and Boston, where at the time of replication only files with tags Chicago exist; i.e., however, no files with the Boston tags exist. Any subsequent files tagged with Boston will also begin to be replicated.

Space Management for Replication

It is critical that sufficient disk space is available on both the primary and the standby file systems for storing the replication logs. Please review the Storage Administrator's Guide or [Pause and Resume Replication](#) section of this paper for details on file system sizing when using replication.

It is recommended that ACFS administrators monitor and prevent both the primary file system and the standby file system from running out of space. Enterprise Manager can be used for this monitoring and sending EM alerts when the file system approaches more than 70% full.

If the primary file system runs out of space, then applications using that file system may fail because ACFS cannot create a new replication log. If the standby file system runs out of space, then it cannot accept new replication logs from the primary node and therefore, cannot apply changes to the standby file system, which cause replication logs to accumulate on the primary file system as well. In cases where ACFS file system space becomes depleted, ACFS administrators can expand the file system, remove unneeded ACFS snapshots, or remove files to reclaim space; although the latter is not recommended. If the primary file system runs out of space and the ACFS administrator intends to remove files to free space, then only files that are not currently being replicated (such as when ACFS tagging is used) should be removed, since the removal of a file that is replicated will itself be captured in a replication log.

Setting up Replication

Before initializing ACFS Replication, ensure that the `compatible.asm` and `compatible.advm` attributes for the diskgroup containing the ACFS file system is set to 11.2.0.2.0 on both the primary and standby nodes. This can be done with `sqlplus` as illustrated below. Notice that `sqlplus` is executed from user Oracle on node1.

```
[oracle@node1 ~]$ sqlplus / as sysasm
SQL> alter diskgroup data set attribute 'compatible.asm' =
'11.2.0.2.0';
Diskgroup altered.

SQL> alter diskgroup data set attribute 'compatible.advm' =
'11.2.0.2.0';
Diskgroup altered.
```

Admin User Setup

In most cases the SYS user in ASM instance can be used as the ACFS Replication Administrator. In which case, the SYS user will need to be granted SYSDBA privilege (on ASM instance). If there is a need to have separate roles for replication management (Replication admin) from daily ASM management, then a separate ASM user can be setup. This user must be granted SYSASM and SYDBA privileges. The following example shows how to setup a replication admin user with user id of *admin* and password of *admin1*.

If an ASM password file does not exist, then create the password file for ASM on all nodes (primary/standby and secondary nodes with multi-node clusters) as follows:

```
[oracle@node1 ~]$ orapwd file=$ORACLE_HOME/dbs/orapw+ASM
password=oracle entries=5
```

Please use a password appropriate for your installation.

Create ASM user on the primary node and assign appropriate roles:

```
[oracle@node1 ~]$ sqlplus / as sysasm
SQL> create user admin identified by admin1;
SQL> grant sysasm to admin;
SQL> grant sysdba to admin;
```

Create ASM user on the standby node and assign appropriate roles:

```
[oracle@node2 ~]$ sqlplus / as sysasm
SQL> create user admin identified by admin1;
SQL> grant sysasm to admin;
SQL> grant sysdba to admin;
```

Review changes to the password file by querying `v$pwfile_users`:

```
SQL> select * from v$pwfile_users;
USERNAME                               SYSDB SYSOP  SYSAS
-----
SYS                                     TRUE  TRUE  FALSE
ADMIN                                  TRUE  FALSE TRUE
```

Hereafter, the ACFS administrator “admin” will refer to the role that will manage the ACFS file system replication

File System Setup

Before initiating replication, the ACFS admin must ensure that the primary file system is mounted and the standby file system is only mounted on one node (in cluster configurations). It is recommended to have the same file system name for the standby and primary file system. Also, ensure that if replicating the entire file system; i.e., not using ACFS Tagging, that the standby file system is created with an equal or larger size as the primary file system. Please review the Storage Administrator’s Guide for details on file system sizing.

Network Setup

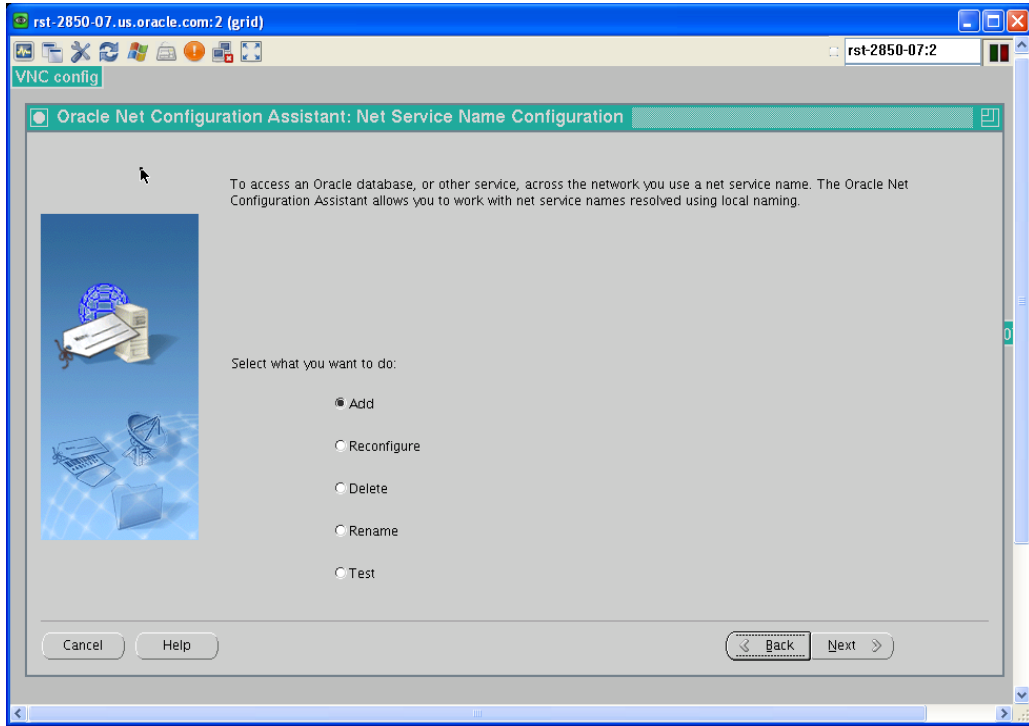
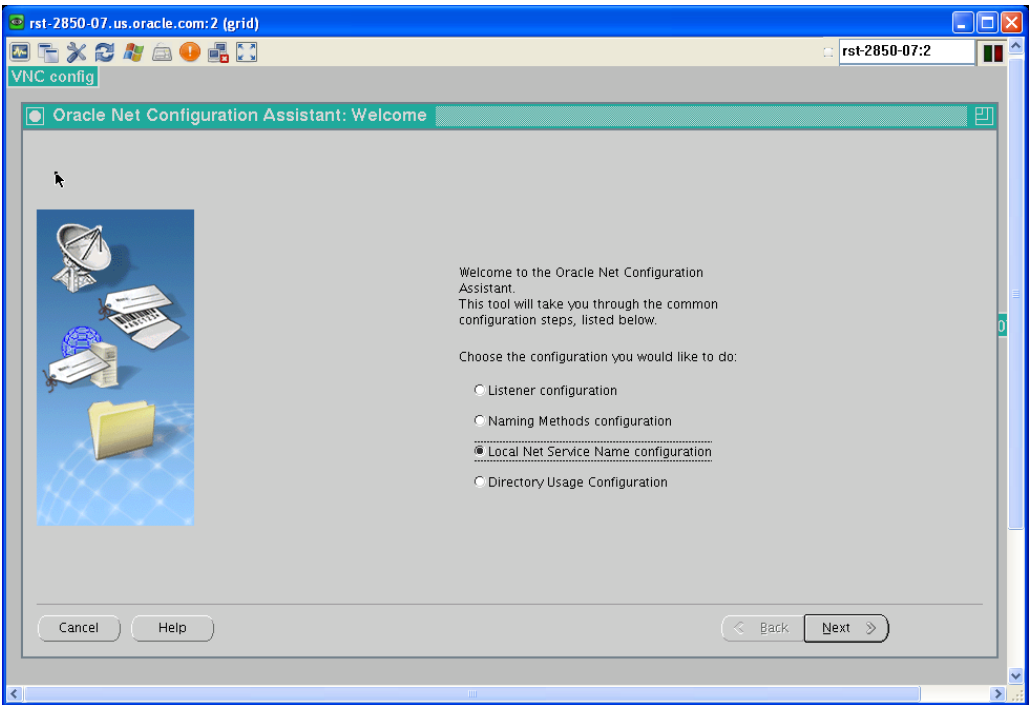
There two steps for configuring the network for ACFS replication:

1. Generate appropriate Oracle Network files. These files provide communication between the ASM instances and ACFS replication.
2. Set appropriate network parameters for network transmission. Since ACFS replication is heavily tied to network bandwidth, the appropriate settings need to be configured

Generate Oracle Network files

ACFS replication utilizes OracleNet for transmitting replication logs between primary and standby nodes. The principle OracleNet configuration is a file called “tnsnames.ora” and it resides at *\$ORACLE_HOME/network/admin/tnsnames.ora*. This file can be edited manually or through a configuration assistant called netca in the Grid Home. A tnsnames.ora file must be updated on each of the nodes participating in ACFS replication. The purpose of a tnsnames.ora file is to provides the Oracle environment the definition of remote endpoint used during replication. For example, there are tnsnames.ora files for both primary and standby nodes.

Once the filesystems are created, use *\$ORACLE_HOME/bin/netca* (from Grid Home) to create connect strings and network aliases for the primary/standby sites.



On netca exit, the following message should be displayed if the services were setup correctly:

```
Oracle Net Services Configuration:
Oracle Net Configuration Assistant is launched from Grid
Infrastructure home. Network configuration will be clusterwide.
Default local naming configuration complete.
  Created net service name: PRIMARY_DATA
Default local naming configuration complete.
  Created net service name: STANDBY_DATA
Oracle Net Services configuration successful. The exit code is 0
```

In our example we created a PRIMARY_DATA service and STANDBY_DATA service, for the primary filesystem and standby filesystem; respectively. In the example contained here, the tnsnames.ora file used for the primary node is:

- Primary Node tnsnames.ora

```
[oracle@node1 ~]$ cat $ORACLE_HOME/network/admin/tnsnames.ora
STANDBY =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = node2) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = acfs_fs)
    )
  )
```

The important elements are the alias name of “STANDBY”, the host name of “node2”, the default port of ‘1521’, and the service name of “acfs_fs”. This tnsnames.ora defines the remote endpoint for replication, in this case, standby, to node1, which is the primary node.

- Standby Node tnsnames.ora

The standby node requires a tnsnames.ora file that defines the primary endpoint. It contains:

```
[oracle@node2 ~]$ cat $ORACLE_HOME/network/admin/tnsnames.ora
PRIMARY =
  (DESCRIPTION =
    (ADDRESS_LIST =
      (ADDRESS = (PROTOCOL = TCP) (HOST = node1) (PORT = 1521))
    )
    (CONNECT_DATA =
      (SERVICE_NAME = acfs_fs)
    )
  )
```

Notice the symmetry between the two tnsnames.ora. For simplicity the service names are the same.

Set Network Tunables

The cornerstone of any successful replication deployment is network efficiency and bandwidth management, therefore the appropriate network tuning must be performed. For ACFS replication, first determine if Data Guard (DG) is already configured on the hosts. If DG is setup appropriately with the appropriate network tunable parameters, then ACFS replication can leverage the same settings. If DG is not enabled, then use the Data Guard best practices guide for network setup. The following document describes these best practices. Please see the *Redo Transport Best Practices* section of this paper.

<http://www.oracle.com/technetwork/database/features/availability/maa-wp-10gr2-dataguardnetworkbestpr-134557.pdf>

Validating Network Configuration

Use the `tnsping` utility and `SQL*PLUS` to test and ensure that basic connectivity exists between both sites and that the `tnsnames.ora` files are setup correctly.

- Primary Node

```
[oracle@node1 ~]$ tns ping standby
TNS Ping Utility for Linux: Version 11.2.0.2.0 - Production on 01-
DEC-2010 12:58:13
Copyright (c) 1997, 2010, Oracle. All rights reserved.
Used parameter files:
/u01/app/oracle/product/11.2.0/asm/network/admin/sqlnet.ora
Used TNSNAMES adapter to resolve the alias
Attempting to contact (DESCRIPTION = (ADDRESS_LIST = (ADDRESS =
(PROTOCOL = TCP) (HOST = node2) (PORT = 1521))) (CONNECT_DATA =
(SERVICE_NAME = acfs_fs)))
OK (0 msec)

[oracle@node1 ~]$ sqlplus admin/admin1 as sysasm@primary
SQL*Plus: Release 11.2.0.2.0 Production on Wed Dec 1 12:58:39 2010
Copyright (c) 1982, 2010, Oracle. All rights reserved.
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.2.0 - 64bit
Production
With the Real Application Clusters and Automatic Storage Management
options
SQL>
```

- Standby Node

```
[oracle@node2 ~]$ tnsping primary
TNS Ping Utility for Linux: Version 11.2.0.2.0 - Production on 01-DEC-2010 13:01:26
Copyright (c) 1997, 2010, Oracle. All rights reserved.
Used parameter files:
/u01/app/oracle/product/11.2.0/asm/network/admin/sqlnet.ora
Used TNSNAMES adapter to resolve the alias
Attempting to contact (DESCRIPTION = (ADDRESS_LIST = (ADDRESS =
(PROTOCOL = TCP) (HOST = node1) (PORT = 1521))) (CONNECT_DATA =
(SERVICE_NAME = acfs_fs)))
OK (0 msec)
```

```
[oracle@node2 ~]$ sqlplus admin/admin1 as sysasm@primary
SQL*Plus: Release 11.2.0.2.0 Production on Wed Dec 1 13:01:46 2010
Copyright (c) 1982, 2010, Oracle. All rights reserved.
Connected to:
Oracle Database 11g Enterprise Edition Release 11.2.0.2.0 - 64bit
Production
With the Real Application Clusters and Automatic Storage Management
options
```

Replication Configuration and Initiation

As a good measure, the following ensures that file systems are mounted on each node.

- Primary Node

```
[oracle@node1 ~]$ df -h
File system          Size  Used Avail Use% Mounted on
/dev/mapper/VolGroup00-LogVol100
                    25G   11G   13G   45% /
/dev/sda1            99M   16M   78M   17% /boot
tmpfs                1.5G  192M   1.3G   13% /dev/shm
.host:/              923G  295G   628G   32% /mnt/hgfs
/dev/asm/acfsvol-269 5.0G  1.2G  3.9G  23% /acfs
```

- Standby Node

```
[oracle@node2 ~]$ df -h
File system          Size  Used Avail Use% Mounted on
/dev/mapper/VolGroup00-LogVol100
                    25G   11G   13G   46% /
/dev/sda1            99M   16M   78M   17% /boot
tmpfs                1.5G  192M   1.3G   13% /dev/shm
.host:/              923G  295G   628G   32% /mnt/hgfs
/dev/asm/acfsvol-343 5.0G  101M  5.0G   2% /acfs
```

Initializing Standby File system

Replication is first initiated on the standby node, followed by initiating on the primary¹. Replication on the standby is initiated the using `/sbin/acfsutil repl init standby` command by the root user.

```
[root@node2 ~]# /sbin/acfsutil repl init standby -p \
  admin/admin1@primary /acfs
```

This command uses the following configuration information:

- `-p` option indicates the username connection to primary file system site, and the service name to be used to connect as ASMADMIN on the primary file system node
- The file system listed is standby file system (`/acfs`)
- Optional (Not shown): If the standby site is using a different service name than the primary file system site then the `-c service_name` is required.

Verify that the Standby File System is initiated

```
[root@node2 ~]# /sbin/acfsutil repl info -c /acfs
Site:                               Standby
Standby status:                   Online
Standby mount point:                 /acfs
Standby Oracle Net service name:     acfs_fs
Primary mount point:
Primary Oracle Net service name:     acfs_fs
Primary Oracle Net alias:            admin/****@primary
Replicated tags:
Log compression:                     Off
Debug log level:                      0
```

Initializing the Primary File System

Once the standby node has been enabled, the ACFS admin can initialize replication on the primary file system by running the `acfsutil repl init primary` command.

¹ If this command is interrupted for any reason, the user must recreate the standby file system, mount it on one node only of the site hosting the standby file system, and rerun the command.

```
[root@node1 ~]# /sbin/acfsutil repl init primary \
-s admin/admin1@standby /acfs
remote connection has been established
waiting for the standby replication site to initialize
waiting for the standby replication site to initialize
The standby replication site is initialized. ACFS replication will begin.
```

This command allows for the following configuration information:

- `-s` followed by connect string that is used to connect as ASMADMIN on the standby node
- ACFS file system that is to be replicated
- Optional (Not shown): `-m mountp`, the mount point on the standby node. If not specified, it is assumed that this mount point path is the same on the standby node as it is on the primary file system node.
- Optional (Not shown): `-c` option indicates the primary service name.

If tagging was enabled for this directory, then a tagname “reptag” can be added in the initialization command such as;

```
[root@node1 ~]# /sbin/acfsutil repl init primary \
-s admin/admin1@standby reptag /acfs
validating the remote connection
remote connection has been established
waiting for the standby replication site to initialize
The standby replication site is initialized. ACFS replication will
begin.
```

Verify that the Primary File System is initiated

```
[root@node1 ~]# /sbin/acfsutil repl info -c /acfs
Site: Primary
Primary status: Online
Primary mount point: /acfs
Primary Oracle Net service name: acfs_fs
Standby mount point: /acfs
Standby Oracle Net service name: acfs_fs
Standby Oracle Net alias: admin/****@standby
Replicated tags:
Log compression: Off
Debug log level: 2
```

Once the `acfsutil repl init primary` command completes successfully, replication will begin transferring copies of all specified files to the standby file system.

The replication happens in two phases; The initial phase copies just the directory tree structure. The second phase copies the individual files. During this phase all updates or truncates to replicated files are blocked. Once a file is completely copied to the standby file system, replication logging for

that particular file are enabled. All changes to copied files, are logged, transported and applied to the standby file system.

Validate Replication Instantiation

```
[root@node1 ~]# /sbin/acfsutil repl info /acfs
Time: Wed Dec  1 13:24:58 2010
Event: Replication Synchronized
Sync Time: Wed Dec  1 13:24:58 2010
Logs were applied
Time: Wed Dec  1 13:24:53 2010
Event: Initialization Files Complete
Cord Number: 1
.
```

The rate of data change on the primary file system can be monitored using the *acfsutil info fs -s* command, with the *-s* flag indicating sample rate. The amount of change includes all user and metadata modifications to the file system. The following example illustrates its usage:

```
[root@node1 ~]# /sbin/acfsutil info fs -s 10 /acfs
/acfs
  amount of change since mount:          0.28 MB
  amount of change: 128.36 MB   rate of change: 13144 KB/s
  amount of change: 93.50 MB   rate of change: 9574 KB/s
```

This “amount” value approximates the size of replication logs generated when capturing changes to the file system. This command is useful for approximating the extra space required for storing replication logs in cases of planned or unplanned outages.

Pause and Resume Replication

The *acfsutil repl pause* command is used in instances when replication needs to be temporarily halted such as for planned downtime on either the primary or standby site. The ACFS pause command can be issued on either the primary or a standby file system node. However, there is a difference in behavior between the two scenarios.

A pause command issued on the standby node will continue to generate and propagate replication logs from primary to standby file system, but these rlogs will not be applied to the standby file system; i.e., it does not suspend transfer of rlogs from the primary node, only the application is deferred.

Consequently, rlogs will continue to accumulate at the node hosting the standby file system. As noted earlier, replication logs are deleted on the standby site only after they are successfully applied to the file system on the standby node, thus care should be taken to ensure that this does not cause the file system to run out of space on the primary and the standby file system.

A pause command issued on primary will generate replication logs but not propagate to standby; i.e., it will generate rlogs but suspends propagation. In this scenario, rlogs will continue to accumulate at the primary file system. This may cause the primary file system to run out of space.

In both cases, ACFS administrators should run *acfsutil repl resume* at earliest point, before the accumulated replication logs fill the file system. Note, the resume command should be executed at the same location where replication was paused.

In cases where there is a planned outage and the standby or primary file systems have to be unmounted, it is best to ensure that all the changes are propagated and applied on the standby file system. The *acfsutil repl sync* command is used for this purpose. It is used to synchronize the state of the primary and standby file system, and implicitly causes all outstanding replication data to be transferred to the standby file system. The *acfsutil repl sync* command returns success when this transfer is complete or, if the *apply* parameter is supplied, when all of these changes have been successfully applied to the standby file system. This command can only be run on the node hosting the primary file system.

For unplanned outages, if the cluster (or node) hosting the primary file system fails, the administrator of the standby file system should decide if the situation is a disaster or not. If it is not a disaster, then when the primary site recovers, replication will automatically restart. If it is a disaster, then issue an *acfsutil terminate* on the standby to convert the standby file system into a primary. Once the original primary is restarted and if replication needs to be re-instantiated, then replication initialization will need to be performed again.

If the node hosting the standby file system fails, then a major concern is the amount of update activity that occurs on the primary file system relative to the amount of free space allocated to address standby file system outages. If the free space in the primary file system is exceeded because of the inability to transfer updates to the standby file system, then a file system out of space condition will occur and space will need to be made available; i.e. remove items no longer needed, particularly snapshots, do a file system resize to add space, etc. However, assuming the standby comes back, then as soon as primary file system space is available replication will continue. During this interval, where no space is available, the file system will return errors in response to update requests. If the standby file system is

going to be down for a long period, then it is recommended that the primary file system be unmounted to avoid update activity on the file system that could result in an out of space condition. When the standby file system becomes available the primary file system could be remounted and replication will restart automatically. Alternatively, the primary file system admin could elect to terminate and re-instantiate once the site hosting the standby file system is recovered.

To size the primary and standby file system appropriately for these planned and unplanned outages, the *acfsutil fs info* command, described earlier, can be used as a guide to determine the rate of replication log creation. First determine the approximate time interval when the primary file system is unable to send replication logs to the standby file system at its usual rate or when standby file systems are inaccessible while undergoing maintenance. Although it is not easy to determine how long an unplanned will last, this exercise helps in determining overall impact when an unplanned outage occurs.

As an aid, run `acfsutil info fs -s 1200` on the primary file system to collect the average rate of change over a 24 hour period with a 20 minute interval.

```
[root@node1 ~]# /sbin/acfsutil info fs -s 1200 /acfs
```

The output from this command helps determine the average rate of change, the peak rate of change, and how long the peaks last. Note, that this command only collects data on the node it is executed on, thus for clustered configurations, run the command and collect for all nodes in the cluster.

In the following scenario, assume $t = 60$ minutes is the time interval in that would adequately account for network problems or maintenance on site hosting the standby file system.

The following formula approximates the extra storage capacity needed for an outage of 60 minutes (assigned $t=60$):

N = Number of cluster nodes in the primary site generating rlogs

p_t = peak amount of change generated across all nodes for time t

$t = 60$ minutes

Extra storage capacity to hold replication logs = $(N * 1GB) + p_t$

In this use case example, assume a 4 node cluster on primary where all 4 are generating replication logs and during peak workload intervals, the total amount of change reported for 60 minutes is approximately 6 GB for all nodes. Using the storage capacity formula above, 10 GB of excess storage capacity on the site hosting the primary file system is required for storage of the replication logs.

Extra storage capacity to hold replication logs = $(4 * 1GB) + 6GB = 10GB$

Termination of Replication

The *acfsutil repl terminate* command is used in cases when administrators need to dissolve the instantiated replication. Note that this is performed at the file system level and not at the node level. To perform graceful ACFS Terminate, it is recommended to terminate the primary file system first, followed with by a terminate command at the standby file system. This will allow for the standby to apply all outstanding logs.

- Terminate Replication on Primary Node

```
[root@node1 ~]# /sbin/acfsutil repl terminate primary /acfs
```

- Terminate Replication on Standby Node

```
[root@node2 ~]# /sbin/acfsutil repl terminate standby /acfs  
acfsutilreplterminate: ACFS-05060: waiting for ACFS replication to terminate
```

Once file system termination has completed for a specific file system, no replication infrastructure exists between that primary and standby file system. Termination of replication is a permanent operation and requires a full re-initialization to instantiate replication. To restart replication, the *acfsutil repl init* command is used as previously illustrated.



ACFS Replication
December 2010
Author: Nitin Vengurlekar
Contributing Authors: Fred Glover, Barb Glover,
Diane Lebel, Jim Williams
Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark licensed through X/Open Company, Ltd. 1010

Hardware and Software, Engineered to Work Together