

ZFS Replication

Manual

Revision 36

Armin Doerzbach

<http://www.doerzbach.com>

Table of Contents

1. Introduction.....	3
2. Disclaimer.....	3
3. Installation.....	3
4. Setup.....	4
5. Start and Stop Replication.....	9
6. Failover Szenarios.....	9
6.1. Failover for Tests.....	9
6.2. Disaster.....	10
6.2.1. Failover.....	10
6.2.2. Fallback.....	11
7. Troubleshooting.....	11
7.1. Replication service is in state maintenance.....	11
7.2. Errors during Fallback.....	12

1. Introduction

The ZFS replication is based on ZFS snapshots of ZFS filesystems, which are transferred using secure shell. Every ZFS filesystem can be replicated to a remote system. Multiple filesystems can be replicated by one replication service. There can be more than one replication service, and one source filesystem can be replicated by multiple replication services.

It is tested on Solaris 10 and Opensolaris and Solaris 11.4.

2. Disclaimer

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

3. Installation

Download the package from:

<http://www.doerzbach.com/downloads/protected/DOEzfsrepl.tar.gz>

The ZFS replication is provided as a Solaris Package DOEzfsrepl. It has to be installed on the source and the destination system first.

```
root@solaris10s rc:/tmp# gunzip -c /DOEzfsrepl.tar.gz |tar xvf -
DOEzfsrepl/
DOEzfsrepl/pkginfo
DOEzfsrepl/root/
DOEzfsrepl/root/opt/
DOEzfsrepl/root/opt/zfsrepl/
DOEzfsrepl/root/opt/zfsrepl/bin/
DOEzfsrepl/root/opt/zfsrepl/bin/sc_stop_svcs.bash
DOEzfsrepl/root/opt/zfsrepl/bin/sc_start_svcs.bash
DOEzfsrepl/root/opt/zfsrepl/bin/setup.bash
DOEzfsrepl/root/opt/zfsrepl/bin/zfs_replicate_service.bash
DOEzfsrepl/root/opt/zfsrepl/bin/zfs_replicate.bash
DOEzfsrepl/root/opt/zfsrepl/bin/functions.bash
DOEzfsrepl/pkgmap
root@solaris10src:/tmp#
root@solaris10src:/tmp# pkgadd -d .
```

```
The following packages are available:
 1  DOEzfsrepl      ZFS Replication Suite
      (all) 1.11
```

```
Select package(s) you wish to process (or 'all' to process
all packages). (default: all) [?,??,q]:
```

```
Processing package instance <DOEzfsrepl> from </tmp>
```

```
ZFS Replication Suite(all) 1.11
Doerzbach Engineering GmbH
## Processing package information.
## Processing system information.
  4 package pathnames are already properly installed.
## Verifying disk space requirements.
## Checking for conflicts with packages already installed.
## Checking for setuid/setgid programs.
```

```
Installing ZFS Replication Suite as <DOEzfsrepl>
```

```
## Installing part 1 of 1.
/opt/zfsrepl/bin/functions.bash
/opt/zfsrepl/bin/sc_start_svcs.bash
/opt/zfsrepl/bin/sc_stop_svcs.bash
/opt/zfsrepl/bin/setup.bash
/opt/zfsrepl/bin/zfs_replicate.bash
/opt/zfsrepl/bin/zfs_replicate_service.bash
[ verifying class <none> ]
```

```
Installation of <DOEzfsrepl> was successful.
```

4. Setup

First you have to get the following details for your service:

- User on the source system, which will do the replication, and the users password
- User on the destination system, which will receive the replication, and the users password
- DNS entry of the source and the destination system or add a entry to /etc/hosts on source and destination system
- The source ZFS filesystems which should be replicated with the same interval
- The related destination ZFS filesystems
- The root users passwords for both source and destination system
- If necessary the exact command to run before the snapshot is taken and after the snapshot is taken (e.g. Put the database into backup mode)

Then run setup for the setup of the replication service:

```
root@solaris10src:/tmp # su - src
Oracle Corporation      SunOS 5.10      Generic Patch      January 2005
src@solaris10src:~ $ /opt/zfsrepl/bin/setup.bash -i
ZFSREPL: Copyright 2012 Doerzbach Engineering GmbH
          Written by Armin Doerzbach
          http://www.doerzbach.com
```

```
THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND
ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO,
THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR
A PARTICULAR PURPOSE ARE DISCLAIMED.
IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO,
PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS;
OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY,
WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
OF THE POSSIBILITY OF SUCH DAMAGE.
```

```
Please read the statement above carefully. Do you agree with it (YES/no)?
```

```
YES
```

```
Enter the filename of the responsefile you want to create: testreplication
```

Here you just enter a filename to store the inputs you give.

```
Enter the source system name (logical alias name for a cluster) which will send the snapshots
[default:]: solaris10src
```

Here you enter the hostname of the source host. If the source system is a cluster then you

enter the logical hostname of the resource group

Enter the destination system name (logical alias name for a cluster) which will receive the snapshots [default:]:**solaris10dest**

Here you enter the hostname of the dest host. If the source system is a cluster then you enter the logical hostname of the resource group

Enter the username on the destination system which will receive the snapshots [default:]:**dst**

Here you enter the username which will receive the snapshots (Normally it is an application user like oracle or so. It can also be root, but then the remote login for root must be possible)

The user id on the destination system will be checked now.
Please enter the password of the destination user if asked...
Warning: Identity file /opt/zfsrepl/etc/ssh_keys/src/ does not exist.
The authenticity of host 'solaris10dest (192.168.1.101)' can't be established.
RSA key fingerprint is 9b:7f:b3:0f:25:c4:fa:a8:04:3a:04:13:39:da:5f:e6.
Are you sure you want to continue connecting (yes/no)? **yes**
Warning: Permanently added 'solaris10dest,192.168.1.101' (RSA) to the list of known hosts.

```
-----  
| This system is for the use of authorized users only. |  
| Individuals using this computer system without authority, or in |  
| excess of their authority, are subject to having all of their |  
| activities on this system monitored and recorded by system |  
| personnel. |  
| |  
| In the course of monitoring individuals improperly using this |  
| system, or in the course of system maintenance, the activities |  
| of authorized users may also be monitored. |  
| |  
| Anyone using this system expressly consents to such monitoring |  
| and is advised that if such monitoring reveals possible |  
| evidence of criminal activity, system personnel may provide the |  
| evidence of such monitoring to law enforcement officials. |  
-----
```

Password:

Enter the password of the dst user here

How many filesystems do you want to replicate with this service [default:]:**2**
Enter source zfs filesystem[1] you want to replicate [default:]:**rpool/src/testfs1**
Enter destination zfs filesystem[1] [default:]:**rpool/dst/testfs1**
Going to check if the parent filesystem rpool/dst of the destination filesystem rpool/dst/testfs1 exists.

On the destination system the zfs filesystem rpool/dst must exist already. This will be checked here

Enter the password of the destination user if asked...
Password:
Enter source zfs filesystem[2] you want to replicate [default:]:**rpool/src/testfs2**
Enter destination zfs filesystem[2] [default:]:**rpool/dst/testfs2**
Going to check if the parent filesystem rpool/dst of the destination filesystem rpool/dst/testfs2 exists.
Enter the password of the destination user if asked...
Password:
Enter the replication interval in seconds [default:]:**60**
Enter the number of retries before failing the script or service [default:1]:**5**

In case of a network failure or other failures, this defines how many times to try before the service will fail.

Enter the number of snapshots to keep before deleting them [default:1]:**3**

How many snapshots should be kept? Do not keep too many as they will use space.

Enter the full path to the script to run before the snapshot should be taken (none if no script should be taken) [default:]:**none**

If you want to execute a script or command before the snapshot is taken, you have to enter the full path to it here. No arguments can be given to the script! If you enter "none" no script will be executed

Enter the full path to the script to run after the snapshot should be taken (none if no script should be taken) [default:]:**none**

If you want to execute a script or command after the snapshot has been taken, you have to enter the full path to it here. No arguments can be given to the script! If you enter "none" no script will be executed

Enter the service name which should be created [default:]:**testrepl**

Just enter a name for your zfsrepl instance. The service name in this case will be **svc:/application/zfsrepl:testrepl**

Do you want to make changes now (Y/[n])?

y

Please run the script **/export/home/src/doi_t_as_root.testrepl.solaris10src.bash** as root on the source system solaris10src now from another terminal.....

Here you have to run the script **/export/home/src/doi_t_as_root.testrepl.solaris10src.bash** as root on the source system:

```
root@solaris10src# /export/home/src/doi_t_as_root.testrepl.solaris10src.bash
Creating directory for the public and private keys /opt/zfsrepl/etc/ssh_keys/src
Running zfs allow -u src create,destroy,mount,send,snapshot,receive,compression,recordsize on ...
rpool/src
rpool/src
root@solaris10src#
```

The setup script will then automatically proceed:

```
Generating public/private rsa key pair.
Your identification has been saved in /opt/zfsrepl/etc/ssh_keys/src/testrepl.
Your public key has been saved in /opt/zfsrepl/etc/ssh_keys/src/testrepl.pub.
The key fingerprint is:
be:e7:22:28:d6:9a:73:7b:02:c3:4c:cb:35:2b:0a:89 zfs_replication:testrepl@solaris10src
Installing public key on remote system now...
Please enter the password if asked.
```

```
|-----|
| This system is for the use of authorized users only. |
| Individuals using this computer system without authority, or in |
| excess of their authority, are subject to having all of their |
| activities on this system monitored and recorded by system |
| personnel. |
| |
| In the course of monitoring individuals improperly using this |
| system, or in the course of system maintenance, the activities |
| of authorized users may also be monitored. |
| |
| Anyone using this system expressly consents to such monitoring |
| and is advised that if such monitoring reveals possible |
| evidence of criminal activity, system personnel may provide the |
| evidence of such monitoring to law enforcement officials. |
|-----|
```

Password:

Enter the password of dst user here

```
|-----|
| This system is for the use of authorized users only. |
| Individuals using this computer system without authority, or in |
```

```

| excess of their authority, are subject to having all of their
| activities on this system monitored and recorded by system
| personnel.
|
| In the course of monitoring individuals improperly using this
| system, or in the course of system maintenance, the activities
| of authorized users may also be monitored.
|
| Anyone using this system expressly consents to such monitoring
| and is advised that if such monitoring reveals possible
| evidence of criminal activity, system personnel may provide the
| evidence of such monitoring to law enforcement officials.
|-----|

```

```

doit_as_root.testrep 100% |
*****| 1053
00:00

```

```

|-----|
| This system is for the use of authorized users only.
| Individuals using this computer system without authority, or in
| excess of their authority, are subject to having all of their
| activities on this system monitored and recorded by system
| personnel.
|
| In the course of monitoring individuals improperly using this
| system, or in the course of system maintenance, the activities
| of authorized users may also be monitored.
|
| Anyone using this system expressly consents to such monitoring
| and is advised that if such monitoring reveals possible
| evidence of criminal activity, system personnel may provide the
| evidence of such monitoring to law enforcement officials.
|-----|

```

```

testreplication 100% |
*****| 301
00:00
Please run the script /tmp/doit_as_root.testrepl.dst.bash as root on the remote host solaris10dest
now....Thank you!!!

```

Now you have to run /tmp/doit_as_root.testrepl.dst.bash on the destination system:

```

root@solaris10dest:/tmp # /tmp/doit_as_root.testrepl.dst.bash
Creating directory for the public and private keys /opt/zfsrepl/etc/ssh_keys/dst
Running zfs allow -u dst create,destroy,mount,send,snapshot,receive,compression,recordsize and chown
dst on ...
rpool/dst
rpool/dst

```

The setup script will proceed automatically and finish the configuration:

```

|-----|
| This system is for the use of authorized users only.
| Individuals using this computer system without authority, or in
| excess of their authority, are subject to having all of their
| activities on this system monitored and recorded by system
| personnel.
|
| In the course of monitoring individuals improperly using this
| system, or in the course of system maintenance, the activities
| of authorized users may also be monitored.
|
| Anyone using this system expressly consents to such monitoring
| and is advised that if such monitoring reveals possible
| evidence of criminal activity, system personnel may provide the
| evidence of such monitoring to law enforcement officials.
|-----|

```

```

|-----|
| This system is for the use of authorized users only.
| Individuals using this computer system without authority, or in
| excess of their authority, are subject to having all of their

```

```

| activities on this system monitored and recorded by system
| personnel.
|
| In the course of monitoring individuals improperly using this
| system, or in the course of system maintenance, the activities
| of authorized users may also be monitored.
|
| Anyone using this system expressly consents to such monitoring
| and is advised that if such monitoring reveals possible
| evidence of criminal activity, system personnel may provide the
| evidence of such monitoring to law enforcement officials.
|-----

```

```

|-----
| This system is for the use of authorized users only.
| Individuals using this computer system without authority, or in
| excess of their authority, are subject to having all of their
| activities on this system monitored and recorded by system
| personnel.
|
| In the course of monitoring individuals improperly using this
| system, or in the course of system maintenance, the activities
| of authorized users may also be monitored.
|
| Anyone using this system expressly consents to such monitoring
| and is advised that if such monitoring reveals possible
| evidence of criminal activity, system personnel may provide the
| evidence of such monitoring to law enforcement officials.
|-----

```

```

Generating public/private rsa key pair.
Your identification has been saved in /opt/zfsrepl/etc/ssh_keys/dst/testrepl.
Your public key has been saved in /opt/zfsrepl/etc/ssh_keys/dst/testrepl.pub.
The key fingerprint is:
b4:71:c2:c8:a5:f4:88:55:6e:57:84:9f:84:3c:b2:e0 zfs_replication:testrepl@solaris10dest
check installation now...
The installation is correct.
The service application/zfsrepl:testrepl can be started now

```

The service `svc:/application/zfsrepl:testrepl` can now be started and the replication will run:

```

root@solaris10src# svcs zfsrepl
STATE      STIME      FMRI
disabled   19:06:37   svc:/application/zfsrepl:testrepl
root@solaris10src# svcadm enable zfsrepl:testrepl
src@solaris10src:~ $ svcs zfsrepl
STATE      STIME      FMRI
online     19:09:44   svc:/application/zfsrepl:testrepl

```

The initial snapshots will take a while to be sent to the destination system (up to many hours), as the whole filesystem is transferred. After that only incremental snapshots are sent....

After a while the zfs filesystems look like this on the source system:

```

root@solaris10src:/tmp # zfs list
NAME                                USED  AVAIL  REFER  MOUNTPOINT
rpool                                8.38G  31.2G  33.5K  /rpool
rpool/ROOT                           6.67G  31.2G   21K  legacy
rpool/ROOT/s10x_u9wos_14a            6.67G  31.2G  4.99G  /
rpool/ROOT/s10x_u9wos_14a/var        1.68G  8.32G  1.68G  /var
rpool/dump                            1.00G  31.2G  1.00G  -
rpool/export                          55.8M  31.2G   23K  /export
rpool/export/home                     55.8M  31.2G  55.8M  /export/home
rpool/src                              66K    31.2G   24K  /src
rpool/src/testfs1                     21K    31.2G   21K  /src/testfs1
rpool/src/testfs1@zfsrepl-src-dst-solaris10dest-15  0      -    21K  -
rpool/src/testfs1@zfsrepl-src-dst-solaris10dest-16  0      -    21K  -
rpool/src/testfs1@zfsrepl-src-dst-solaris10dest-17  0      -    21K  -
rpool/src/testfs2                     21K    31.2G   21K  /src/testfs2
rpool/src/testfs2@zfsrepl-src-dst-solaris10dest-15  0      -    21K  -
rpool/src/testfs2@zfsrepl-src-dst-solaris10dest-16  0      -    21K  -
rpool/src/testfs2@zfsrepl-src-dst-solaris10dest-17  0      -    21K  -
rpool/swap                            671M   31.9G  16.3M  -

```


swappool	15.6G	13.5M	21K	legacy
swappool/swap1	15.6G	15.6G	16K	-

After a while the zfs filesystems look like this on the destination system:

```
root@solaris10dest:/tmp # zfs list
NAME                                USED    AVAIL    REFER  MOUNTPOINT
rpool                                8.38G  31.2G   32.5K  /rpool
rpool/ROOT                           6.67G  31.2G    21K   legacy
rpool/ROOT/s10x_u9wos_14a            6.67G  31.2G   4.99G  /
rpool/ROOT/s10x_u9wos_14a/var       1.68G  8.32G   1.68G  /var
rpool/dst                             63K    31.2G    21K   legacy
rpool/dst/testfs1                    21K    31.2G    21K   legacy
rpool/dst/testfs1@zfsrepl-src-dst-solaris10dest-16  0      -        21K   -
rpool/dst/testfs1@zfsrepl-src-dst-solaris10dest-17  0      -        21K   -
rpool/dst/testfs1@zfsrepl-src-dst-solaris10dest-18  0      -        21K   -
rpool/dst/testfs2                    21K    31.2G    21K   legacy
rpool/dst/testfs2@zfsrepl-src-dst-solaris10dest-16  0      -        21K   -
rpool/dst/testfs2@zfsrepl-src-dst-solaris10dest-17  0      -        21K   -
rpool/dst/testfs2@zfsrepl-src-dst-solaris10dest-18  0      -        21K   -
rpool/dump                          1.00G  31.2G   1.00G  -
rpool/export                        55.8M  31.2G    23K   /export
rpool/export/home                   55.8M  31.2G   55.8M  /export/home
rpool/swap                          671M   31.9G   16.3M  -
swappool                            15.6G  13.5M    21K   legacy
swappool/swap1                      15.6G  15.6G    16K   -
```

5. Start and Stop Replication

You can just stop the SMF Service:

```
root@solaris10src:/tmp # svcadm disable zfsrepl:testrepl
root@solaris10src:/tmp # svcs zfsrepl
STATE      STIME      FMRI
disabled   19:31:23   svc:/application/zfsrepl:testrepl
```

You can just start the SMF Service:

```
root@solaris10src:/tmp # svcadm enable zfsrepl:testrepl
root@solaris10src:/tmp # svcs zfsrepl
STATE      STIME      FMRI
offline*   19:31:50   svc:/application/zfsrepl:testrepl
root@solaris10src:/tmp # svcs zfsrepl
STATE      STIME      FMRI
online     19:31:55   svc:/application/zfsrepl:testrepl
```

6. Failover Szenarios

6.1. Failover for Tests

If you want to take the disaster system online, for tests or maintenance reason you just have to stop the replication services on the source system first.

```
root@solaris10src:/tmp # svcadm disable zfsrepl:testrepl
```

Otherwise the replication will rollback the changes on the destination system everytime an incremental snapshot is replicated.

Then just modify the mountpoint of the destination filesystem using:

```
zfs set mountpoint=/wherevertomount/testfs1 rpool/dst/testfs1
```

Depending on the application it is necessary to stop the application before the last replication takes place.

The last replication can easily be invoked using the following command as src user on the source system:

```
src@solaris10src:~ $ /opt/zfsrepl/bin/zfs_replicate_service.bash runonce testrepl
```

At the end of the tests or the maintenance, you have two choices

- you can replicate back the changes from the destination system to the source system
- you just overwrite all the changes made on the destination

If you want to overwrite the changes you made on the destination system you just start the replication service on the source system again

```
root@solaris10src:/tmp # svcadm enable zfsrepl:testrepl
```

If you really want to replicate back the changes you made during the tests or the maintenance window, you will have to run the following command on the **destination** system as **destination user**:

```
dst@solaris10dest:~ $ /opt/zfsrepl/bin/zfs_replicate_service.bash runonce testrepl
```

But be sure, that the application was not running on the source system. Otherwise you will just overwrite all the changes made to the source system during maintenance window or tests, which you may not want!!!

You will be asked for the password of the source user on the source system.

6.2. Disaster

6.2.1. Failover

In a disaster case, the source system maybe temporarily not responsive or damaged. In this case we have to make sure, that no replication is taking place anymore before we start the application on the destination system. This can be done running the following command on the destination system as dst user:

```
dst@solaris10dest:~ $ /opt/zfsrepl/bin/zfs_replicate_service.bash lockoutkey testrepl
```

This command will remove the ssh-public key for the replication service from the `authorized_keys` file. Like this we make sure the replication can not take place anymore.

Then just modify the mountpoint of the destination filesystem(s) using:

```
zfs set mountpoint=/wherevertomount/testfs1 rpool/dst/testfs1
```

No we can startup the application on the destination system.

There will be a data loss of all the data changed since the last replication took place!! Be sure that this is the best option to go, before you start the destination system.

6.2.2. Fallback

After such a case normally we would have to set up the source system again. Then we would have to recreate user for the replication:

```
useradd -g staff -d /export/home/src -m -s /bin/bash src
```

Then recreate the the parent zfs filesystems of the replicated zfs filesystems, with appropriate permissions for src user.

```
zfs create rpool/src
zfs allow -u src create,destroy,mount,send,snapshot,receive,compression,recordsize
zfs set mountpoint=/whereever/src rpool/src
chown src /whereever/src
```

Then we would invoke the following command on the **destination** system as **destination user** which replicates the data back:

```
dst@solaris10dest:~ $ /opt/zfsrepl/bin/zfs_replicate_service.bash runonce testrepl
```

You could run this command several times. Before you really switch back to the source system.

Switch back to the source system after you have made a final replication with the application stopped on the destination system.

After all start the application on the source system again. Then you should set up the replication from scratch from the src system.

To do this just copy the file `/opt/zfsrepl/etc/svcsdefinition/<service_name>.remote` back to the source system into the source users homedirectory. In our case the file would be named `/opt/zfsrepl/etc/svcsdefinition/testrepl.remote`

Login to the **source system** as **source user** and rename the file just copied to the `<service_name>` in our case to `testrepl`.

Then run setup as follows:

```
src@solaris10src:~ $ /opt/zfsrepl/bin/setup -r testrepl
```

and follow the dialogs you already have seen during the setup. You will just be asked for passwords and you have to run scripts as root as during setup. The definition of the replication is taken from the file `testrepl` you have copied from the destination system.

7. Troubleshooting

7.1. Replication service is in state maintenance

In rare cases the replication is not starting again:

```

root@testso111:~# svcs -xv
svc:/application/zfsrepl:test-sync (?)
  State: maintenance since Thu Jun 28 11:34:10 2018
  Reason: Restarting too quickly.
  See: http://support.oracle.com/msg/SMF-8000-L5
  See: /var/svc/log/application-zfsrepl:test-sync.log
Impact: This service is not running.
root@testso111:~# tail /var/svc/log/application-zfsrepl:test-sync.log
cannot receive: most recent snapshot of rpool/test/dst/fs1 does not
match incremental source
28.06.18 11:33:59 WARNING Sending the incremental snapshot rpool/test/src/fs1@zfsrepl-ado-dst-
localhost-124 failed with exit code 1
cannot receive: most recent snapshot of rpool/test/dst/fs1 does not
match incremental source
28.06.18 11:34:05 WARNING Sending the incremental snapshot rpool/test/src/fs1@zfsrepl-ado-dst-
localhost-124 failed with exit code 1
28.06.18 11:34:10 ALARM Could not send snapshots even after 4 retries. Stopping here now
[ 2018 Jun 28 11:34:10 Stopping because all processes in service exited. ]
[ 2018 Jun 28 11:34:10 Executing stop method (:kill). ]
[ 2018 Jun 28 11:34:10 Restarting too quickly, changing state to maintenance. ]

```

This is due to some changes of the destination filesystem and some missing snapshots:

```

root@testso111:~# zfs list |grep test/dst
rpool/test/dst                83K  17.6G   32K  /rpool/test/dst
rpool/test/dst/fs1            51K  17.6G   31K  /rpool/test/dst/fs1
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-120  1K    -      31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-121  1K    -      31K  -

```

```

root@testso111:~# zfs list |grep test/src
rpool/test/src                63K  17.6G   32K  /rpool/test/src
rpool/test/src/fs1            31K  17.6G   31K  /rpool/test/src/fs1
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-120  0    -      31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-121  0    -      31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-122  0    -      31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-123  0    -      31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-124  0    -      31K  -

```

Here we have only 121 snapshot on destination system and 122-124 on source system. The easiest is to destroy snapshots 122 through to 124 on source system. This does not destroy any data on the source system just the snapshots...

```

root@testso111:~# zfs destroy rpool/test/src/fs1@zfsrepl-ado-dst-localhost-122
root@testso111:~# zfs destroy rpool/test/src/fs1@zfsrepl-ado-dst-localhost-123
root@testso111:~# zfs destroy rpool/test/src/fs1@zfsrepl-ado-dst-localhost-124

```

Then the replication will replicate again...

```

root@testso111:~# svcadm clear test-sync
root@testso111:~# zfs list|grep test/dst
rpool/test/dst                67K  17.6G   32K  /rpool/test/dst
rpool/test/dst/fs1            35K  17.6G   31K  /rpool/test/dst/fs1
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-136  1K    -      31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-137  1K    -      31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-138  1K    -      31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-139  1K    -      31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-140  0    -      31K  -

```

7.2. Errors during Fallback

Sometimes during Fallback you will find it is not possible to apply incremental snapshot for inconsistency between source snapshot and destination snapshot.

```

root@testso111:~# su - dst
Oracle Corporation      SunOS 5.11      11.4.Beta      January 2018
dst@testso111:~$ /opt/zfsrepl/bin/zfs_replicate_service.bash runonce test-sync
You possibly destroy data on the source system localhost.

```

```

You will be asked to enter the password of the source user ado
Warning: Identity file /tmp/nokeyfile not accessible: No such file or directory.
Password:
Password:
Password:
cannot receive: most recent snapshot of rpool/test/src/fs1 does not
match incremental source
28.06.18 13:28:26 WARNING Sending the incremental snapshot rpool/test/dst/fs1@zfsrepl-ado-dst-
localhost-258 failed with exit code 1
cannot receive: most recent snapshot of rpool/test/src/fs1 does not
match incremental source
28.06.18 13:28:31 WARNING Sending the incremental snapshot rpool/test/dst/fs1@zfsrepl-ado-dst-
localhost-258 failed with exit code 1
cannot receive: most recent snapshot of rpool/test/src/fs1 does not
match incremental source
28.06.18 13:28:36 WARNING Sending the incremental snapshot rpool/test/dst/fs1@zfsrepl-ado-dst-
localhost-258 failed with exit code 1
cannot receive: most recent snapshot of rpool/test/src/fs1 does not
match incremental source
28.06.18 13:28:42 WARNING Sending the incremental snapshot rpool/test/dst/fs1@zfsrepl-ado-dst-
localhost-258 failed with exit code 1
28.06.18 13:28:47 ALARM Could not send snapshots even after 4 retries. Stopping here now

```

Then the easiest is to rollback snapshot on the source system (which is not running at the moment) to a more recent snapshot.

```

root@testso111:~# zfs list |grep test/src
rpool/test/src                82K  17.6G    32K  /rpool/test/src
rpool/test/src/fs1            50K  17.6G    31K  /rpool/test/src/fs1
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-256  19K   -       31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-257    0   -       31K  -
root@testso111:~# zfs rollback -r rpool/test/src/fs1@zfsrepl-ado-dst-localhost-256

```

Then to destroy the snapshots newer than the snapshot on the source system. Here 257 and 258

```

root@testso111:~# zfs list |grep test/dst
rpool/test/dst                83K  17.6G    32K  /rpool/test/dst
rpool/test/dst/fs1            51K  17.6G    31K  /rpool/test/dst/fs1
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-256    1K   -       31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-257    1K   -       31K  -
rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-258    0   -       31K  -
root@testso111:~# zfs destroy rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-257
root@testso111:~# zfs destroy rpool/test/dst/fs1@zfsrepl-ado-dst-localhost-258

```

Then try again to replicate back:

```

dst@testso111:~$ /opt/zfsrepl/bin/zfs_replicate_service.bash runonce test-sync
You possibly destroy data on the source system localhost.
You will be asked to enter the password of the source user ado
Warning: Identity file /tmp/nokeyfile not accessible: No such file or directory.
Password:

```

And check again on the source system if a new snapshot 257 was created:

```

root@testso111:~# zfs list |grep test/src
rpool/test/src                82K  17.6G    32K  /rpool/test/src
rpool/test/src/fs1            50K  17.6G    31K  /rpool/test/src/fs1
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-256  19K   -       31K  -
rpool/test/src/fs1@zfsrepl-ado-dst-localhost-257    0   -       31K  -

```

No the filesystems are in sync again.