# RP based Implementation of DR Solution for a Solaris Zone

Sangeeta
Department of Computer Science,
South Point Institute of Technology and Management,
Sonipat, Haryana-131001, India

Maneela
Department of Computer Science,
South Point Institute of Technology and Management,
Sonipat, Haryana-131001, India

## ABSTRACT
The main objective of this paper is to implement a Disaster recovery solution for a Solaris zone using emc recoverpoint technology. It is concerned with preventing any unexpected serious interruptions to IT services as a result of a natural disasters or other forms of force majeure having a catastrophic impact on the business. Today, almost all companies, large and small, depend to a greater or lesser extent on IT services. It is therefore to be expected that if these IT services are interrupted it will affect almost all aspects of the business. However, it is clear that there are strategic IT services on whose continuity the survival of the business may depend, and others that "simply" increase the productivity of the sales and work force.

Older technologies like sun guard and netbackup recovers a server using tapes which is very time consuming and a not good Disaster recovery solution in case of natural disaster where whole DC is impacted.

In this article, An RP based solution is deployed where data is in continuous sync between source site and target site which reduces RPO and RTO to a great extent.

## Keywords
Disaster Recovery, Recovery Point Objective, Recovery Time Objective, emc.

## 1. INTRODUCTION
Today, almost all companies, large and small, depend to a greater or lesser extent on IT services. It is therefore to be expected that if these IT services are interrupted it will affect almost all aspects of the business. The greater the impact associated with the interruption of a particular service, the greater the effort that needs to be devoted to prevention. Sooner or later, no matter how efficient our preventive activities have been, it will be necessary to bring recovery procedures into operation.

In general terms, there are three options for service recovery:

"Cold standby": which requires an alternative site where the live and service environments can be reproduced in not more than 72 hours.

"Warm standby": which requires an alternative site with active systems designed to allow recovery of critical services within 24 to 72 hours.

"Hot standby": which requires an alternative site with continuous replication of the data and all the systems active ready to substitute the live structure immediately.

And there are two key measures of Disaster Recovery capability:

Recovery Point Objective (RPO): The point in time to which (transaction) data must be recovered after an event (e.g. end of previous day's processing) in order to meet business needs as defined by your organization.

Recovery Time Objective (RTO): The period of time within which systems, applications, or functions must be recovered after an event (measured in minutes, hours or days). RTO is based on the business requirements and assessment from the Business Impact Analysis.

## 2. RP based DR Solution Framework
IT Service Continuity technologies to allow the replication and failover of servers and storage, from one DC to another DC. This paper details the process of setting up the Solaris systems to replicate a Solaris 10 zones using the EMC product Recover point. The intention is to fully replicate a Solaris 10 zone from one hosted on a global zone in the one DC to a second ITSC dedicated global zone hosted in the another DC. It does not cover the replication of third party applications installed within the zone.
In order to implement the DR solution for a remote datacenter, it is necessary to categorize the organization servers so that we can prioritize those servers first which are mission critical. Here are the various DR tiers:

**Mission Critical:** A server is defined as mission critical if it has a severe impact on one or more of the following; the ability of a customer to utilize equipment or services that an organization provides including installation or service of equipment, revenue generation, cash generation, customer contact, or legal and regulatory. For example, servers running share market applications are mission critical servers for a banking environment. If such a server gets crashed during market hours and there is no DR server against it, then there would be a huge loss to the bank on that day.

**Business Critical:** A server is defined as business critical if it slows down Organization's ability to achieve the business functions which may have one or more of the following

impacts; business unit performance analysis and reporting, internal reporting and internal accounting, or procurement.

**Non-Critical:** A Server is defined as non-critical if there is only a local impact and the scope impacts a work-group or individual employee's performance or productivity.

Now whenever a natural disaster occurs, our first priority would be to bring DR servers online against mission critical servers so that we can minimize the monitory loss associated with these servers, then we work for business critical servers. Also there is no need to have DR servers against Non-critical servers as there is no business loss associated with these servers. These are primarily used for application testing before it is deployed on a mission critical or business critical servers.

Now let us build a Solaris zone and then replicate that using RP to a remote datacentre:

### Zone creation steps:

Before we create a zone, it is necessary to verify that MPXIO is enabled on the global server, so that new device is accessible through multiple paths. Then we sense the newly added LUN on the physical server so that we can create a zpool in case of ZFS or DG in case of Veritas FS. Here we will go with ZFS FS. We will further create a ZFS FS on which a new zone will be created.

```
#devfsadm –cv
#zpool create zone_pool  <disk_name>
#zfs  create zone_pool/root
#zfs set mountpoint=/export/zones/zonename  zone_pool/root
#chmod 700 /export/zones/zonename
```

Now we have a FS which can be used to create a zone on it.

```
# zonecfg -z prod_zone
prod_zone: No such zone configured

Use 'create' to begin configuring a new zone.

zonecfg:prod_zone> create -b

zonecfg:prod_zone>set zonepath=/export/zones/prod_zone

zonecfg:prod_zone> set autoboot=true

zonecfg:prod_zone> set ip-type=shared

zonecfg:prod_zone> add fs

zonecfg:prod_zone:fs> set dir=/data

zonecfg:prod_zone:fs> set special=prod_pool/apps

zonecfg:prod_zone:fs> set type=zfs

zonecfg:prod_zone:fs> end

zonecfg:prod_zone> verify

zonecfg:prod_zone> commit

zonecfg:prod_zone> exit

#zoneadm –c prod_zone install

# zoneadm –c prod_zone boot

# zlogin –C prod_zone
```

After this step, we have a Solaris zone fully configured and in running state.

At this point, a request is raised with storage team to allocate same size of storage in another DC on a physical server meant for DR and then setup an RP between the prod server and DR site. Here are the RP scripting steps for a successful replication:

1. Get the group status and bookmark an image

```
export DATE=`date '+%d%b%Y-%H:%M'`

export LOGFILE=/var/cg/rp.stat_$DATE.log

export CGLIST=/var/cg/s1-bos1.cglist

export RPA_ADDR=10.127.14.24

for i in `cat $CGLIST`; do

sshadmin@$$RPA_ADDR get_group_state
>/tmp/group/$i_grp_state

grep –i paused /tmp/group/$i_grp_state

if [ $? –eq 0 ] then;

echo "$i GROUP IS PAUSED- PLS CHECK GROUP "

fi

done
```

2. Bookmark the image

```
for i in `cat $CGLIST`; do

echo "Making a bookmark of $i"

ssh admin@$RPA_ADDR bookmark_image group=$i
bookmark=$i_$DATE
```

3. Verify if snapshots are successful

```
for i in `cat $CGLIST`; do

echo "Verification of image loaded on target"

ssh admin@$RPA_ADDR verify_group group=$i
bookmark_loaded=${i} ${DATE}

>/dev/null

while [ $? –ne 0 ]; do

echo "$i bookmark still not ready!"

sleep 60

ssh admin@$RPA_ADDR verify_group group=$i

bookmark_loaded=${i} ${DATE} >/dev/null

done
```

```
echo "$i bookmark is now ready"

done.
```

4. Enable image access

```
for i in `cat $CGLIST`; do

echo "Setting up image access for $i"

ssh admin@$RPA_ADDR enable_image_access group=$i
image=$i_$DATE

>/dev/null

sleep 30

done
```

5. Re-enable replication

```
for i in `cat $CGLIST`; do

echo "stop img access and return to replication for CG $i"

ssh admin@$RPA_ADDR disable_image_access group=$i
>/dev/null

done
```

Now we setup iDNS so that we can use same name for the zone at DR server as well, but with a different IP from different vlan. By doing so, application users will not even come to know whether they are on DR server during actual disaster. Here we have one zone name with two VIPs, one used at Prod site and another will be used at DR site during disaster. Therefore application team always have access to same name irrespective of DC.

**Testing DR setup:**

A request is made to storage team to break the RP. Once it is confirmed replication is broken the zone can be un-mounted and deported.

```
#zoneadm –z prod_zone halt

#zoneadm –z prod_zone detach
```

Now we need to scan for the disk on the target system, import the zpool, and attach the zone. attaching the zone will upgrade it to the patch level on the global server.

```
#devfsadm  –cv
#zpool import zone_pool
#zoneadm –z prod_zone attach
```

Here we assume that physical servers at both DC are at same level and zone configuration file is created/copied manually at DR site.

Once zone pool is imported and zone gets attached successfully, we need to boot the zone so as to verify it

```
#zoneadm  –z  prod_zone boot
#zlogin  -C  prod_zone
```

## 3. Figure showing zone replication between remote sites

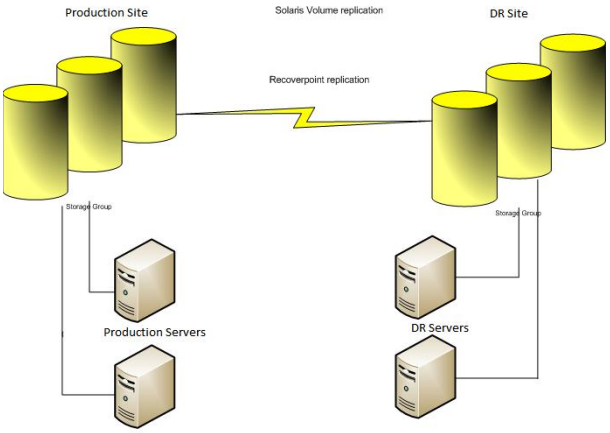


In the diagram, we have a Production site(source DC) where an organization's mission critical and business critical servers are racked along with the storage devices. These servers are replicated over the private WAN using various technologies like data domain for old servers, data guard for database FS etc. RP is one of the many techniques which is used here for Solaris zone replication. Here DR site(target DC) is the remote data center which hosts various DR servers. Recoverpoints are setup between the Solaris zones of Production site and that of DR site which continuously sync data over the WAN. Whenever there is a natural disaster at Production site, all of the Solaris zones at DR site will be brought online within very limited RPO and RTO.

## 4. CONCLUSIONS

Disaster recovery (DR) has taken on a new sense of urgency in recent years. Emerging issues like natural disasters, computer software and hardware failures, terrorism, hackers, computer viruses have all led to increase our focus on preparing for disasters. Therefore disaster recovery plan must be integrated with the overall organization ITSC approach and that must be tested through regular yearly drills as well as a test drill whenever an environmental change occurs. The plan should include documented and tested procedures with tested RPO and RTO. Here are the significance of RP based research on Solaris zones:

- Improved RTO and RPO
- Enable Alignment with Business Continuity
- Reduce multi-million dollar spend with 3rd parties like SunGard and Symantec
- Operational improvements
- Faster backup and restore times through data replication
- Data De-Duplication and Replication
- Reduces the dependency on tape technology
- Avoids storage cost / less disk required

## 5. REFERENCES

[1] John Dix, "cloud computing causing rethinking of disaster recovery", Network World: July 30, 2013.

[2] EMC, "Improving VMware Disaster Recovery with EMC RecoverPoint", emc.com, August 2012

[3] Laura DuBois, "Best practices in Business Continuity and Disaster Recovery", IDC, Feb 2011.

## 6. AUTHORS PROFILE

- Sangeeta obtained her B. Tech from Kurukshetra University, and M. Tech (CE) Pursuing from Deenbandhu Chhotu Ram University of Science and Technology, Murthal, India. She has attended various national seminars, conferences and presented research papers on ITSC.

- Maneela obtained her M. Tech from M. D. University, Rohtak. She started her teaching career three years back serving as a lecturer and head of the department of computer science, at South Point Institute of Technology and Management. She has attended various national seminars, conferences and presented research papers on ITSC.