IBM XIV Hyper-Scale Consistency

Demonstrated with Oracle

Eran Orgad, Patrick Pollard
IBM Systems and Technology Group ISV Enablement

March 2014

Table of contents

Abstract

This white paper on IBM XIV Hyper-Scale Consistency explores an Oracle database with a configuration that spans multiple IBM XIV Storage Systems.

The IBM XIV grid architecture extends its capabilities for databases that require highly available storage systems with sustained high performance to easily use Hyper-Scale Mobility to move volumes from one storage systems to another while the database is up and running.

This paper discusses the consistency groups and provides some sample scripts to demonstrate such capabilities that enable better flexibility and ease of management of critical applications.

About IBM XIV Hyper-Scale Consistency

IBM® Hyper-Scale Consistency enables an administrator to establish a volume consistency group spanning multiple IBM XIV® Storage Systems. This feature is introduced in the IBM XIV storage software version 11.4 enables new sets of XCLI commands to coordinate the creation of snapshots across XIV frames.

IBM Hyper-Scale Consistency is another step forward in the Hyper-Scale vision, taking the XIV grid concept beyond the dimensions of a single XIV Storage System allowing better flexibility and easier control on critical data.

Executive summary

This paper demonstrates an implementation of the IBM XIV Hyper-Scale Consistency to multiple XIV Storage Systems.

Enterprise application databases typically need to be highly available. Yet the trend is to minimize maintenance time and at the same time optimize storage placement, such that highly available content moves to faster hosts and storage systems. This dynamic movement of critical data requires smart storage array capabilities that allow seamless volume migration. XIV Storage System is delivering this functionality with IBM Hyper-Scale Mobility, allowing storage administrators to move critical parts or the entire database to other XIV systems while the database is up and running. Doing that, it is highly desirable to maintain the consistency and integrity of the database that is spread across multiple volumes and XIV systems.

This paper demonstrates a real life example, keeping the consistency of an Oracle database that resides on two different XIV storage arrays. The scripts used in this paper are helpful to understand the concept and to inspire similar solutions aimed to automate parts of the backup procedures. We demonstrate a capability to create Hyper-Scale Consistency snapshot of both Oracle backup mode as well as a Point-in Time copy of the database without entering backup mode, producing a crash consistent* copy.

* Applications that are not crash consistent might need additional tools to assure consistence as VSS for Exchange server.

Introduction

The Hyper-Scale Consistency group setup uses two XIV systems to host an Oracle database that spans two sets of three 1 TB data volumes and three 50 GB Oracle log volumes. The sets are configured with the consistency group definition to house a *Normal*[*] Oracle ASM setting. This allows for Oracle preferred read from the faster set assuming systems have different hardware setting such as solid-state drive (SSD) on one XIV Storage System only. In such cases, reads of cache hits will benefit from the faster array, while Oracle Automatic Storage Management (ASM) mirrored array allows us the high level of availability with redundancy that assures database uptime under various hardware failures.

Consistency groups on multiple XIV configuration

The XIV code level 11.4 introduces Hyper-Scale Consistency groups. The following configuration setup used one set of consistency group devices on XIV (XIV1) and another consistency group on another XIV (XIV2). In this example, the focus is on a single host XIV197, Linux® host running Red Hat Enterprise Linux 5.9, with multiple volumes on both XIV systems defined within a consistency group (marked as CG1 and CG2).

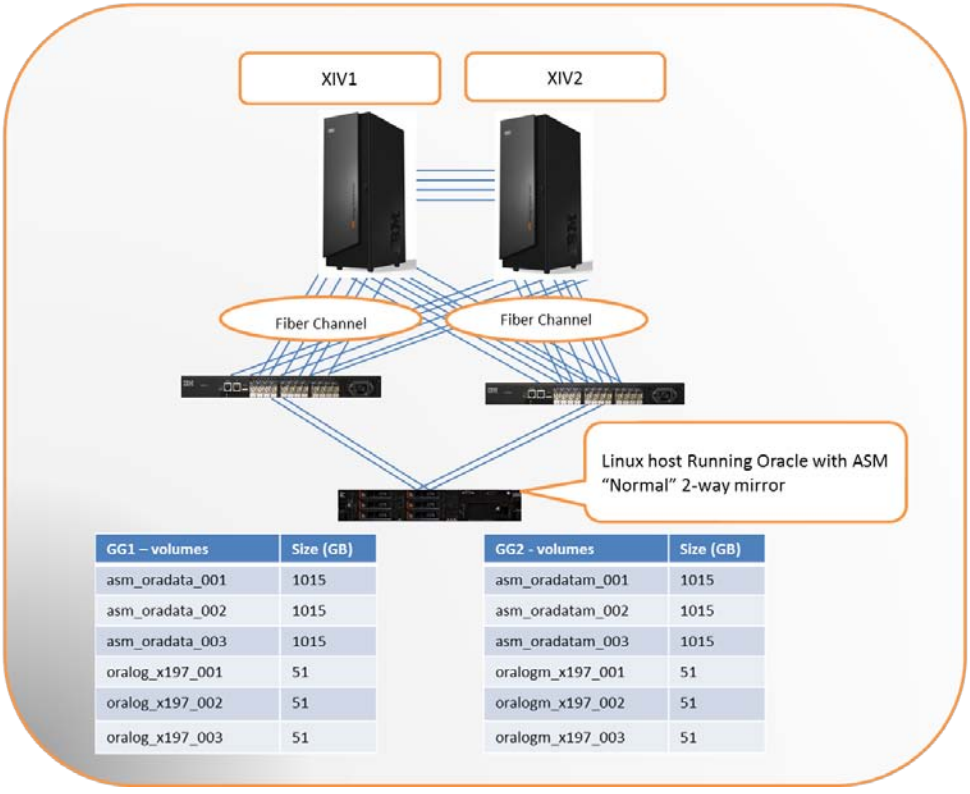

GG1 – volumes	Size (GB)
asm_oradata_001	1015
asm_oradata_002	1015
asm_oradata_003	1015
oralog_x197_001	51
oralog_x197_002	51
oralog_x197_003	51

GG2 - volumes	Size (GB)
asm_oradatam_001	1015
asm_oradatam_002	1015
asm_oradatam_003	1015
oralogm_x197_001	51
oralogm_x197_002	51
oralogm_x197_003	51

Figure 1*: Configuration of XIV consistency groups with Oracle ASM*

[*] Oracle ASM Normal = 2-way mirroring, High = 3-way mirroring, External = No mirroring

Hyper-Scale Consistency groups

Consistency groups can be used to take simultaneous snapshots of multiple volumes, thus ensuring consistent copies of a group of volumes. The Hyper-Scale vision takes this concept beyond a single XIV frame. This milestone shows how write ordering can be maintained using XCLI commands thus preserving the integrity of multiple CG on multiple XIV storage systems.

This concept provides a list of benefits as XIV snapshots are taken at sub seconds duration. The value for administrators is important as they have the flexibility to run multiple snapshots, not only to recover quickly from hardware failure, but also from a logical failure created by unintended delete or drop of tables from the database.

In the example scripts, the test team created two consistency groups that during the invocation of the wrapping script, the team has the Oracle database enter into backup mode, this step is not deemed necessary yet for the purpose of demonstration, this step is kept intact. Following that step, the team invoked a script that list the consistency groups and issues a `pause_io` XCLI command to multiple XIV frames. While in this mode, the team has 10 seconds to complete the consistent point in time snapshot before io_resume is issued automatically to both XIV arrays. Because Oracle is in the backup mode, the team completed the step to take Oracle out of the backup mode post snapshot completion. The same script supports that ability to invoke a point in time (PiT) snapshot only without entering that database into the backup mode.

XIV GUI ease of use

The XIV GUI provides the user with an intuitive interface to set and review that the consistency groups as well as any CG snapshots that are executed directly through the GUI or as a result of running a script that automates the regular backups and point in time copies created.

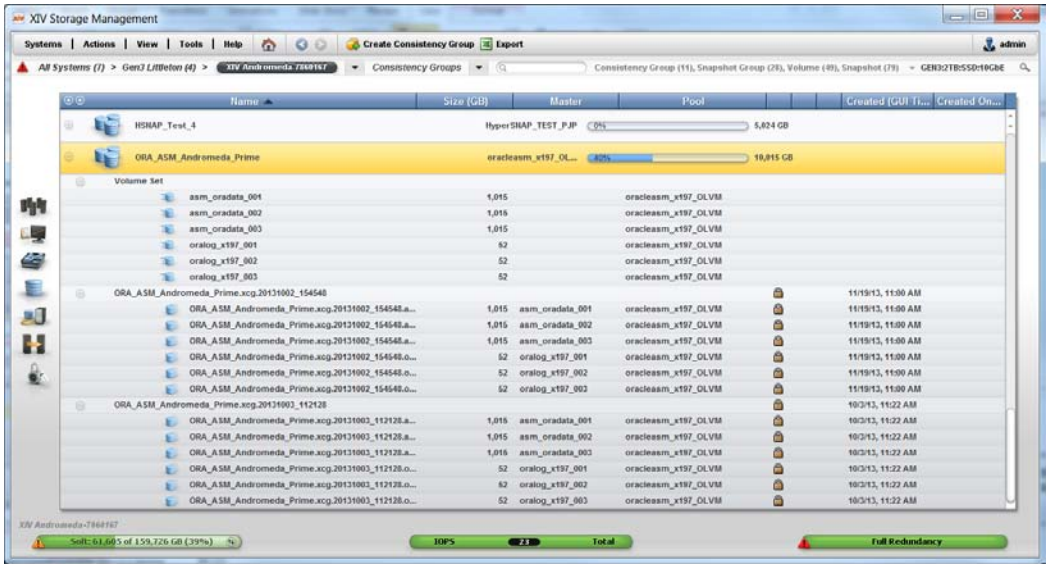

Figure 2: List Consistency of Groups and associated snapshots

At a high level, the XIV Connectivity view allows the user to view and monitor the state and connection of multiple arrays in a single screen, informing of a change in the connectivity and alerting the admin of any degradation (note the little triangle at the bottom of the view indicating alerts that should be reviewed).

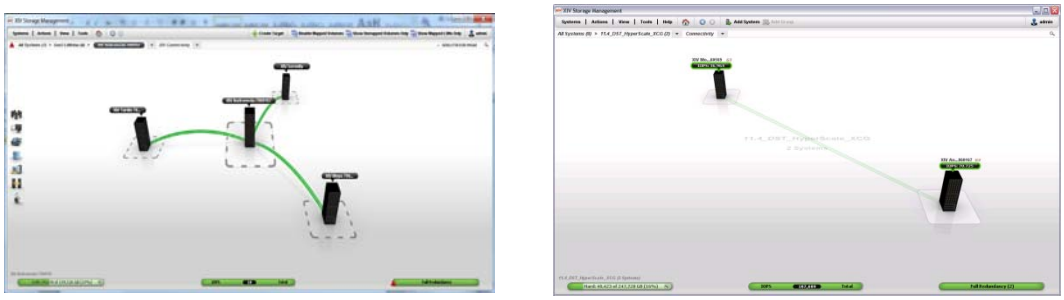

Figure 3: XIV connectivity view

Achieving optimal RPO

Multiple constraints and consideration impact decisions of how frequent users can use snapshot, and what mix of snapshot is best suited. When attempting to enter a database in the backup mode you are shifting operational resources that impact performance and latency, quiescence of database during operation might have other negative side effects on production.

Another consideration administrators need to take into account is the backup window as well as the toll on the host performance. Such constraints make the storage snapshot techniques more plausible to use when maintaining high performance in highly demanding environments.

Frequent snapshots might increase the usage of storage capacity. An ability to automate the creation and deletion of snapshots that gives the right mix of application-consistent backup as well as crash-consistent backups may provide an optimal way to protect applications to a desired Recovery Point Objective (RPO) for any given data set from an application maintaining write ordering. For deletion of snapshot one can take advantage of the auto-delete priority to determine the order in which snapshot are deleted should the pool snapshot space consume most available space.

Such tools allow also for a quick way to invoke an additional Application consistent snapshot right before determined event that are considered of higher risk to result in a failure that may require recovery, such as upgrades, or running a procedure that will modify tables, row type, indexing etc.

In such cases administrators would preferably have already the set of tools and scripts that will allow them to take snapshots in seconds right before invoking such procedures.

It is very reassuring when attempting such complex operations that a recovery path would be available in seconds should anything go wrong. By using either the XIV GUI XCLI, or a scripted consistency group restore, administrators can get the database or data sets to a healthy state. When a crash-consistent snapshot copy of the full database is restored and the instance is started, Oracle automatically performs and manages crash recovery. The database will then be transaction consistent up to the last redo commit

of the snapshot copy. After the database is opened, no further redo logs beyond the snapshot copy can be applied*.

Golden snapshot of a consistency group

To supplement the administrator's ability to easily control large subset of data scattered over multiple volumes (which is extended to multiple XIV), you can create a golden consistency group snapshot that will not be auto deleted with the mix of snapshots at higher priority.

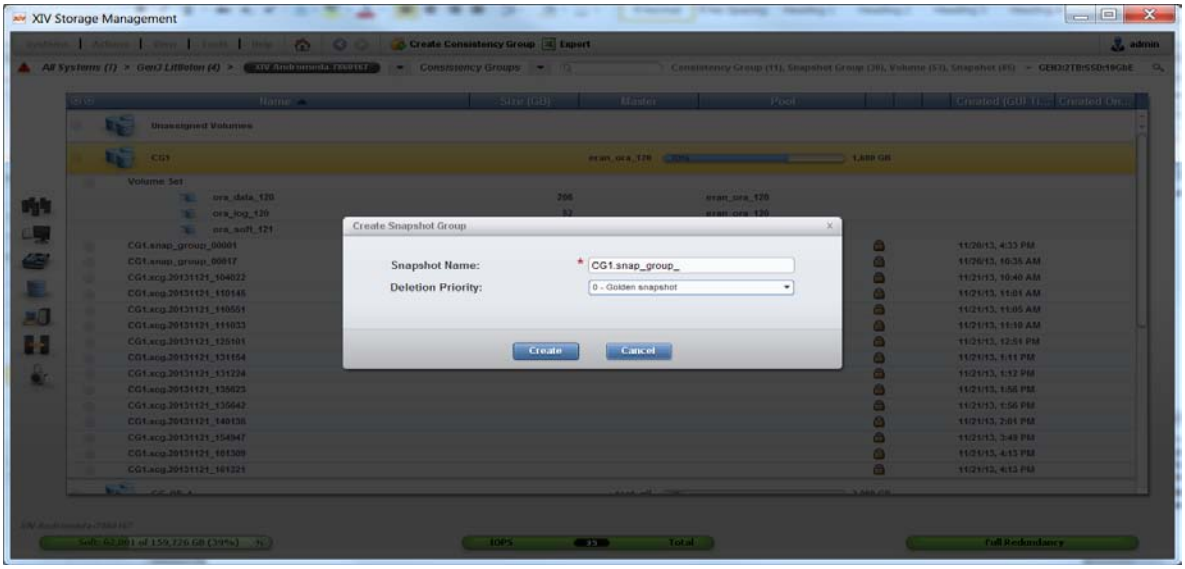

Figure 4: Golden consistency group snapshot

Using XCLI to script and automate operations

The XIV system comes with extensive sets of API and XCLI interfaces that allow users to customize, integrate, and automate critical or repeated tasks. For web interface integration into other systems, refer to the XIV RESTful API.

In this paper, the test team has also provided two example scripts showing how to use and automate tasks, particularly the ability to take a cross XIV consistent snapshot of a database (refer to Figure 1). The script allow the user some options such as take a point-in-time copy, which will invoke the io_pause, take a consistent snapshot, and auto resume as described earlier. Another option is to run the script with the backup option. This will enter Oracle into a backup mode, invoke the consistent snapshot, and take Oracle out of the backup mode.

Backup procedure

Performing a full database snapshot without entering the backup mode is regarded as a point-in-time copy of the database. The snapshot must contain all the data files, control files, and redo logs. If you need

* Oracle document ID 604683.1 – Supported Backup, Restore and Recovery operations using Third Party Snapshot Technologies

to restore this point-in-time copy, you will have a valid crash-consistent copy (refer to the Oracle document ID 604683.1 – *Supported Backup, Restore and Recovery operations using Third Party Snapshot Technologies*). This copy of the database will be consistent to the last redo commit of the snapshot copy. The following steps occur when you invoke the point-in-time copy as detailed in the example script.

The wrapper scripts send the consistency group list that include all the data files, control files, and redo logs to the xiv_xcg_backup.py script which parses the content of the cg_file.csv file producing pairs of consistency groups and XIV systems they reside on.

Next, the script would issue the `io_pause XCLi` command with the appropriate consistency group name to each pair. This command is issued with the auto resume, allowing for 10 seconds to take a consistent snapshot on all participating consistency groups on the various XIV systems. If the snapshot competes, I/O is resumed and you can have a valid point-in-time copy of the database. The scripts traps for some error condition if XCLI commands fail on a variety of error conditions, or fail to take a snapshot within the allowed duration of 10 seconds.

Restore procedure

Considering the RPO and the recovery time objective (RTO), it is apparent that having easy UI that allows you to invoke your recovery quickly is important, equally important, the work performed by XIV to restore is yet again almost instantaneous.

You can select the consistency group and recover from the UI as show in the following figure.

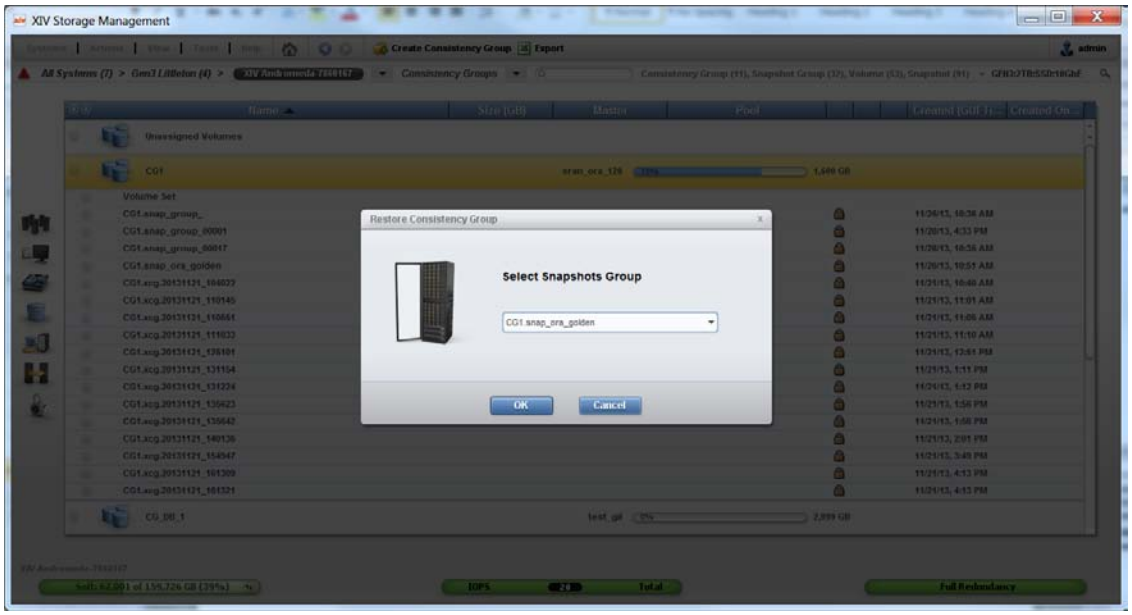

Figure 5: Consistency group restore

Other required steps might be more specific to the application you are using, however, if you look at the example followed in this paper, here are the steps at a high level. In this case, the RTO is in single minutes. The test team did not make any assumption on the RPO, however, when setting up automated interval to take snapshots, it is best to consider how much data is put at risk. In case a snapshot is taken every hour, you might lose up to one hour worth of data should you need to recover, such decisions should take into account your RPO and if necessary increasing the frequency of the snapshot should be considered.

You need to perform the following high-level steps for recovery.

1. Shut down the application database (DB).
2. Unmount the application DB disk from Oracle ASM.
3. Restore the application DB from cross-consistency snapshot through script or XIV GUI.
4. Mount the application DB disks from ASM.
5. Start the application DB.
6. DB will enter recovery because snapshots were taken in the backup mode.
7. Restart the DB with the MOUNT option
8. Alter the DB to allow to start normally.

Summary

IBM Hyper-Scale Consistency extends the Hyper-Scale vision with the ability to create consistency groups spanning multiple XIV frames. Hyper-Scale Consistency groups is one of three pillars in the Hyper-Scale expanding vision. Together with the other two pillars of this vision, Hyper-Scale Manager and Hyper-Scale Mobility, users of the IBM XIV system can enjoy better flexibility, ease of management, high performance, Scalability and Reliability leading to phenomenally low TCO.

Hyper-Scale Consistency makes it possible to use XIV frames that include SSDs in conjunction with frames that do not have SSD allowing the database hosts to take advantage of the faster frame, contributing to overall better performance.

Appendix A: How to use the example scripts

Along with this paper, you can use an example package of scripts and text files to help you in the process of deploying your own Hyper-Scale Consistency group, and automating the creation of consistent point-in-time snapshots and application snapshots when the database is in the backup mode.

List of items:

- README.txt – A readme file listing what is needed to run xiv_xcg_backup.py.
- xiv_xcg_backup.py – The script that allows you to take a Hyper-Scale Consistent snapshot on multiple consistency groups residing on multiple XIV systems.
- MyDB.csv – A simple comma separated file containing a pair of {XIV IP, CG}. This file will be read by the xiv_xcg_backup.py script and the list of consistency groups in this file will be considered as the cross system consistency group.
- Backup_xiv_ora.sh – A wrapper script example that allows for the multiple choices as follows:

  Usage: `./Backup_xiv_ora.sh [-pit|-backup|-control|-snapshot|-listsnapshot|-startdb]`

  - `-pit` is to take a consistent point-in-time snapshot of the consistency groups listed in the mydb.csv file.
  - `- backup` is to enter the Oracle database in the backup mode before taking a consistent snapshot of the consistency groups listed in the mydb.csv file, and when complete, exit the backup mode.
  - `-control` is to take backup of only the control files.
  - `-snapshot` is to use XCLI to create a snapshot using volume names.
  - `-listsnapshot` is to use XCLI to just list all the snapshots for a given list of volumes.
  - `-startdb` is a call to start your Oracle database from the script.

Appendix B: Scripts for XIV Hyper-Scale Consistency group

README.txt

```
###############################################################################################
'# xiv_xcg_backup:
'#
'# Licensed Materials - Property of IBM
'# IBM XIV Storage System Software
'# (c) Copyright IBM Corp. 2008. All Rights Reserved.
'# US Government Users Restricted Rights -
'# Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM
Corp.
'#
'# DISCLAIMER:
'# The following source code is sample code created by IBM Corporation.
'# This sample code is not currently part of any standard IBM product and is provided
to
'# you solely for the purpose of assisting you in the development of your applications.
'# The code is provided 'AS IS', without warranty or condition of any kind. IBM shall
'# not be liable for any damages arising out of your use of the sample code, even if
IBM
'# has been advised of the possibility of such damages.
'#
'# Author: Gil Sharon (gil@us.ibm.com)
'###############################################################################################

This script was written as an example script and does not cover all errors, inputs
atc...
It was tested (ah! you call that testing...) only on Windows and linux with Python
2.7.3

a few Notes:
- you need python to run this script... duh! (you can install the xpyv from the HAK)
- xcli should be in your PATH ( or you can edit the XCLI parameter in the
xiv_xcg_backup.py file...
- you MUST set the user/password using 'xcli --protect' before running the script....

Usage: python xiv_xcg_backup.py <my_ip_cg_file.csv>

an example csv file is attached..
the format is:
xiv1_ip,cg_name1
xiv1_ip,cg_name2
xiv2_ip,cg_name3
...

# changes:
*v1.3 - added stripping spaces from names in the csv file...
*v1.2 - made changes to run on linux...
*v1.1 - changed the 'cg_snapshots_create' to include 'auto_resume' so calls to
io_resume is not needed.
      - changed the name given to snapshots to include a constant snap_name and the
time.
```

*v1.0 - first version

xiv_xcg_backup.py

```
##############################################################################################
##
# xiv_xcg_backup: v1.3
#
# Licensed Materials - Property of IBM
# IBM XIV Storage System Software
# (c) Copyright IBM Corp. 2008. All Rights Reserved.
# US Government Users Restricted Rights -
# Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# DISCLAIMER:
# The following source code is sample code created by IBM Corporation.
# This sample code is not currently part of any standard IBM product and is provided to
# you solely for the purpose of assisting you in the development of your applications.
# The code is provided 'AS IS', without warranty or condition of any kind. IBM shall
# not be liable for any damages arising out of your use of the sample code, even if IBM
# has been advised of the possibility of such damages.
#
# Author: Gil Sharon (gil@us.ibm.com)
##############################################################################################
##

import os
import sys, traceback
import csv
import subprocess
import time
import xml.etree.ElementTree as ET

XCLI="xcli"
TIME_TO_STOP_IO="10000"
SNAP_NAME="xcg"

DEFAULT_POPEN_KWARGS = {
    "stdout": subprocess.PIPE,
    "stdin": subprocess.PIPE,
    "stderr": subprocess.PIPE}


def set_command_to_os(xcli_cmd, command):
    if 'win' in sys.platform:
        new_command = []
        for c in command:
            new_command.append(" " + c)
        xcli_cmd.append(new_command)

    if 'linux' in sys.platform:
        xcli_cmd.extend(command)

    return xcli_cmd

def run_xcli_shell_command(ip, command):
    xcli_cmd = [XCLI, "-m", ip, "-x"]
    xcli_cmd = set_command_to_os(xcli_cmd, command)
    print xcli_cmd
    try:
```

```
        additional_popen_kwargs = dict()
        additional_popen_kwargs.update(DEFAULT_POPEN_KWARGS)
        p = subprocess.Popen(xcli_cmd, **additional_popen_kwargs)

        out = ""
        # read line without blocking
        for line in iter(p.stdout.readline, ''):
#           print line
            out += line
            # if no /XCLIRETURN comes... we will be here forever...
            if "/XCLIRETURN" in line:
                return out
#       out, err = p.communicate()

    except (OSError, ValueError, subprocess.CalledProcessError), e:
        raise Exception ('Command %s had failed: %s' % (command, e))


class XivXcgBackup():
    def __init__(self, default=None):
        pass

    def read_input_file(self, argv):
        if(2 != len(argv)):
            raise Exception("Usage: xiv_xcg_backup <cg_file.csv>")

        records = []
        reader = csv.reader(open(argv[1], "rb"), delimiter=',', quoting=csv.QUOTE_NONE)
        for record in enumerate(reader):
            if 2 != len(record[1]) or record[1][0].startswith('#'):
                continue
            records.append([record[1][0].strip(), record[1][1].strip()])

        return records

    def pause_io(self, ip, cg):
        xcli_output = run_xcli_shell_command(ip, ["io_pause", "cg=" + cg,
"milli_seconds_to_resume=" + TIME_TO_STOP_IO])
        if "SUCCESS" not in xcli_output:
            print xcli_output
            raise Exception("could not pause io for %s on %s" % (cg, ip))

        # get token and return it..
        root = ET.fromstring(xcli_output)
        ch = root.getchildren()
#       print ch
        response = ch[0].find("token_id")
#       print response
#       response = root.find("/XCLIRETURN/OUTPUT")
        return str(response.get("value"))


    def snap_cg(self, ip, cg, tokens):
        new_snap_name = (cg + "." + SNAP_NAME + "." + time.strftime("%Y%m%d_%H%M%S",
time.localtime(self.start_time)))[-63:]
        xcli_output = run_xcli_shell_command(ip,
                                             ["cg_snapshots_create",
                                              "cg=" + cg,
                                              "snap_group=" + new_snap_name,
                                              "auto_resume=" + tokens[ip + "_" + cg]])
        if "SUCCESS" not in xcli_output:
            print xcli_output
```

```
            raise Exception("could not create snapshot for %s on %s" % (cg, ip))

    def resume_io(self, ip, token):
        xcli_output = run_xcli_shell_command(ip, ["io_resume", "token_id=" + token])
        if "SUCCESS" not in xcli_output:
            print xcli_output
            raise Exception("could not resume io for token %s on %s" % (token, ip))

    def release_ios(self, records, tokens):
        # for each cg:
        #   resume io token_id
        for r in records:
            if r[0] + "_" + r[1] in tokens:
                self.resume_io(r[0], tokens[r[0] + "_" + r[1]])

    def backup_start(self):
        records = {}
        try:
#            print run_xcli_shell_command([XCLI, "-x", "-m" ,sys.argv[1]
,"version_get"])


            tokens = {}
            # read input file
            records = self.read_input_file(sys.argv)

            # get start time
            self.start_time = time.time()
            try:
                # for each cg:
                #   io_pause cg=cg1 milli_seconds_to_resume=10000
                for r in records:
                    tokens[r[0] + "_" + r[1]] = self.pause_io(r[0], r[1])

                # for each cg:
                #   snap cg
                for r in records:
                    self.snap_cg(r[0], r[1], tokens) # TODO.. handle errors..
            except Exception, e:
                print "Oops: %s" % e
                traceback.print_exc()

                self.release_ios(records, tokens)

            end_time = time.time()
            print "exec time: " + str(end_time - self.start_time) + " sec."

        except Exception, e:
            print "Cannot run script: %s" % e
            traceback.print_exc()


if __name__ == "__main__":
    xcgObj = XivXcgBackup()
    xcgObj.backup_start()
```

mydb2.csv

```
9.32.222.111,CG1
9.32.222.133,CG2
```

Backup_xiv_ora.sh

```
#!/bin/bash

################################################################################################
##
# Backup_xiv_ora.sh: v1.1
#
# Licensed Materials - Property of IBM
# IBM XIV Storage System Software
# (c) Copyright IBM Corp. 2008. All Rights Reserved.
# US Government Users Restricted Rights -
# Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# DISCLAIMER:
# The following source code is sample code created by IBM Corporation.
# This sample code is not currently part of any standard IBM product and is provided to
# you solely for the purpose of assisting you in the development of your applications.
# The code is provided 'AS IS', without warranty or condition of any kind. IBM shall
# not be liable for any damages arising out of your use of the sample code, even if IBM
# has been advised of the possibility of such damages.
#
# Author: Eran Orgad (eorgad@us.ibm.com)
################################################################################################
##

MY_IP_CG_FILE_CSV=/home/oracle/mydb2.csv
SYS_PASSWD=<OraclePasswd>
ADMIN_PASS=<adminPasswd>
SID=ora120
XIV1=9.32.248.171
XIV2=9.32.248.182
HOST=9.32.248.131
PORT=1521
USER1=SYSTEM
CONTROL_FILE1=+FLASHDG/xivoradb/controlfile/control_one
CONTROL_FILE2=+FLASHDG/xivoradb/controlfile/control_two
export PATH=$PATH:/opt/xiv/IBM_XIV_Storage_Management_GUI

function control_backup {
rman << EOF
rman TARGET sys/$SYS_PASSWD CATALOG rman/rman@CATALOG
run { allocate channel t1 type disk;
copy current controlfile to '$CONTROL_FILE1';
copy current controlfile to '$CONTROL_FILE2';
}
resync catalog;
exit
EOF
}

function end_backup {
echo "#### Executing end_backup ####";
echo sqlplus $USER1/$SYS_PASSWD@$HOST:$PORT/@$SID
sqlplus SYSTEM/oracle@$HOST:1521/ora120 <<EOF
alter database $SID end backup;
alter system archive log current;
alter system switch logfile;
```

```
exit
EOF
}


function begin_backup {
echo "#### Executing begin_backup ####";
sqlplus SYSTEM/oracle@$HOST:1521/ora120 << EOF
alter system archive log current;
alter database $SID begin backup;
exit
EOF
}

function select_count {
echo "#### Executing select_count ####";
sqlplus '/nolog' <<EOF
connect / as sysdba
select  FIRST_NAME, LAST_NAME, EMAIL from employees;
exit
EOF
}


function start_asm {
echo "#### Executing start_asm ####";
sqlplus sys/$SYS_PASSWD as sysdba << EOF
startup;
exit
EOF
}


function start_db {
echo "#### Executing start_db ####";
sqlplus sys/$SYS_PASSWD as sysdba << EOF
startup mount;
exit
EOF
}

### MAIN PROGRAN ###
echo "################# Script start ########################";

case "$1" in
'-backup')
## Get the oracle database in Backup mode
echo "performing a select statement";
#select_count
begin_backup
# Taking multiple XIV consistent snapshot using Gil's script
echo xiv_xcg_backup.py $MY_IP_CG_FILE_CSV
/opt/xiv/host_attach/xpyv/bin/xpyv xiv_xcg_backup.py $MY_IP_CG_FILE_CSV
## We have a consistent backup so we can end backup mode
end_backup
;;
'-control')
## To Backup control file uncomment the following
echo "control_backup"
;;
'-pit')
echo "## Point in time copy only ##"
```

```
# Taking multiple XIV consistent snapshot using XCG script
echo xiv_xcg_backup.py $MY_IP_CG_FILE_CSV
/opt/xiv/host_attach/xpyv/bin/xpyv xiv_xcg_backup.py $MY_IP_CG_FILE_CSV
;;
'-startdb')
echo "## start asm and start db ##"
start_asm;
start_db;
;;
'-snapshot')
echo "## perform volume snapshot via XCLI ##"
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_create vol=ora_data_120
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_create vol=ora_soft_121
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_create vol=ora_log_120
xcli -u admin -p $ADMIN_PASS -m $XIV2 snapshot_create vol=ora_log_120
;;
'-listsnapshot')
echo "## list volume snapshot via XCLI ##"
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_list vol=ora_data_120
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_list vol=ora_soft_121
xcli -u admin -p $ADMIN_PASS -m $XIV1 snapshot_list vol=ora_log_120
xcli -u admin -p $ADMIN_PASS -m $XIV2 snapshot_list vol=ora_log_120
;;

*)
echo "Usage: $0 [-pit|-backup|-control|-snapshot|-listsnapshot|-startdb]"
;;
esac

echo "################ Script end ######################";
```

Appendix C: Consistency group XCLI commands

Category	Name	Description
consistency-group	cg_add_vol	Adds a volume to a consistency group
consistency-group	cg_create	Creates a consistency group
consistency-group	cg_delete	Deletes a consistency group
consistency-group	cg_list	Lists consistency groups
consistency-group	cg_remove_vol	Removes a volume from a consistency group
consistency-group	cg_rename	Renames consistency groups
consistency-group	cg_snapshots_create	Creates a snapshot group of a consistency group.
consistency-group	io_pause	Suspends I/O execution on a consistency group
consistency-group	io_pause_list	Lists the consistency group's io_pause state
consistency-group	io_resume	Resumes I/O execution on a consistency group previously suspended with resume_io

Appendix D: XCLI vol_list_cg

The following is a result of the vol_list XCLI command:

```
XIV Moya-7860169>>vol_list cg=ORA_ASM_Moya_Mirror
Name            Size (GB)  Consistency Group         Pool               Used Capacity (GB)
asm_oradatam_001 1015      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                    996
asm_oradatam_002 1015      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                    996
asm_oradatam_003 1015      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                    996
oralogm_x197_001   51      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                      1
oralogm_x197_002   51      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                      1
oralogm_x197_003   51      ORA_ASM_Moya_Mirror      oraasm_x197_olvm                      1

XIV Andromeda-7860167>>vol_list cg=ORA_ASM_Andromeda_Prime
Name            Size (GB)  Consistency Group         Pool                Used Capacity (GB)
asm_oradata_001 1015       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM               996
asm_oradata_002 1015       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM               996
asm_oradata_003 1015       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM               996
oralog_x197_001   51       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM                 1
oralog_x197_002   51       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM                 1
oralog_x197_003   51       ORA_ASM_Andromeda_Prime   oracleasm_x197_OLVM                 1
```

Appendix E: Resources

The following websites provide useful references to supplement the information contained in this paper:

- IBM Systems on PartnerWorld
  ibm.com/partnerworld/systems

- Power Development Platform
  ibm.com/partnerworld/pdp

- IBM Power Systems Information Center
  http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp

- IBM i 7.1 Information Center
  http://publib.boulder.ibm.com/infocenter/iseries/v7r1m0/index.jsp

- IBM AIX 7.1 Information Center
  http://publib.boulder.ibm.com/infocenter/aix/v7r1/index.jsp

- IBM Redbooks
  ibm.com/redbooks

- IBM Publications Center
  www.elink.ibmlink.ibm.com/public/applications/publications/cgibin/pbi.cgi?CTY=US

Acknowledgements

The authors would like to thank the following colleagues for the help, guidance, and technical reviews, production of scripts and setup and test of the entire configuration.

Glenn H Fujimoto, Gil Sharon, Eyal Abraham, Eliran Zada and Yossi Siles

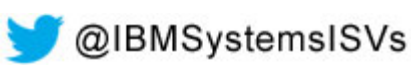

Trademarks and special notices

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local IBM office or IBM authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.