

BEST PRACTICES DEPLOYMENT GUIDE

Oracle Validated Configuration with Cisco UCS, Nimble Storage, and Oracle Linux



Best Practices Deployment Guide: Oracle Validated Configuration with Cisco UCS, Nimble Storage, and Oracle Linux

This document will highlight end-to-end design best practices for deploying Oracle Database on SmartStack, and showcase what was validated jointly by Oracle, Cisco, and Nimble Storage. For more information on Oracle databases with Cisco UCS, Nimble Storage, and Oracle Linux with the Unbreakable Enterprise Kernel, please contact your sales rep or visit:


<http://www.nimblestorage.com/resources/SmartStack.php>

Table of Contents

3	Section 1: No Single Point of Failure
6	Section 2: Performance Optimizations
9	Section 3: Data Protection
10	Section 4: Operational Management
12	Appendix: Oracle Validated Configuration

1. No Single Point of Failure (SPOF)

When you deploy Oracle databases, you want to ensure the entire infrastructure has no single point of failure, for both hardware and software, across all layers (compute, network, storage, and applications). Here is a list of design considerations:

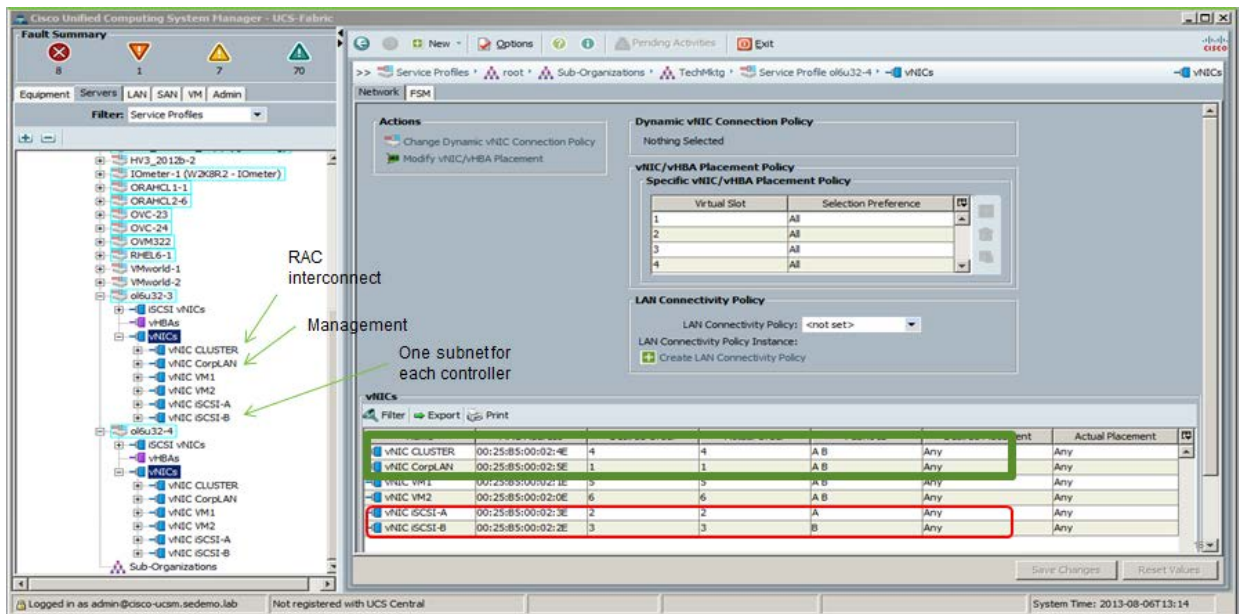


1. Compute

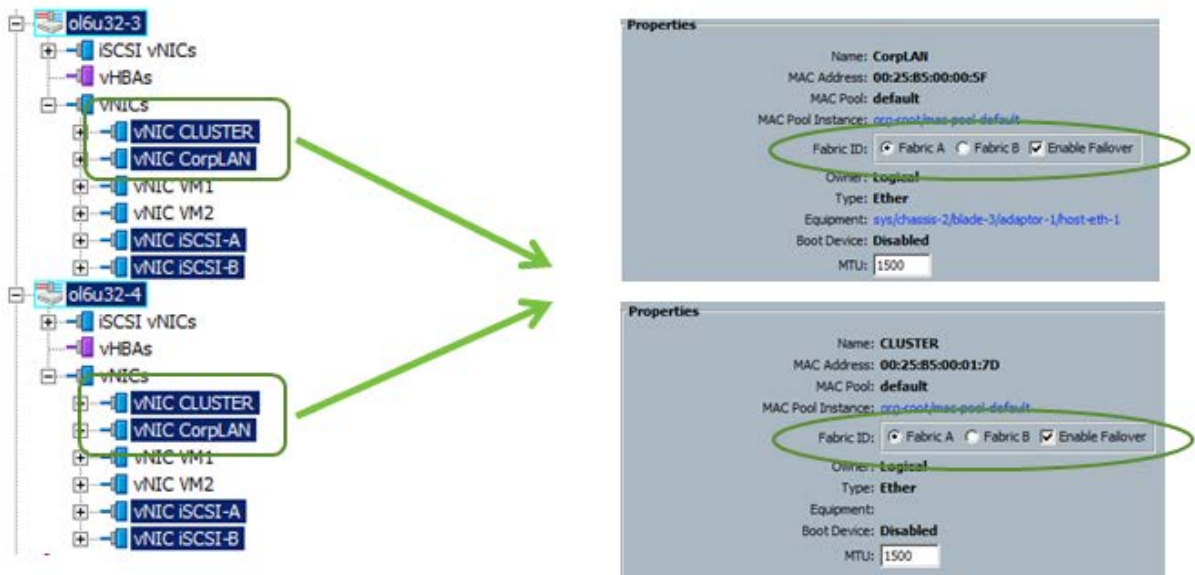
- Cisco UCS blade chassis has redundant everything
- Two UCS blade servers in case one of them fails
- Oracle Linux 6u4 x86_64 with the Unbreakable Enterprise Kernel
- Oracle Database 11gR2 RAC

2. Network

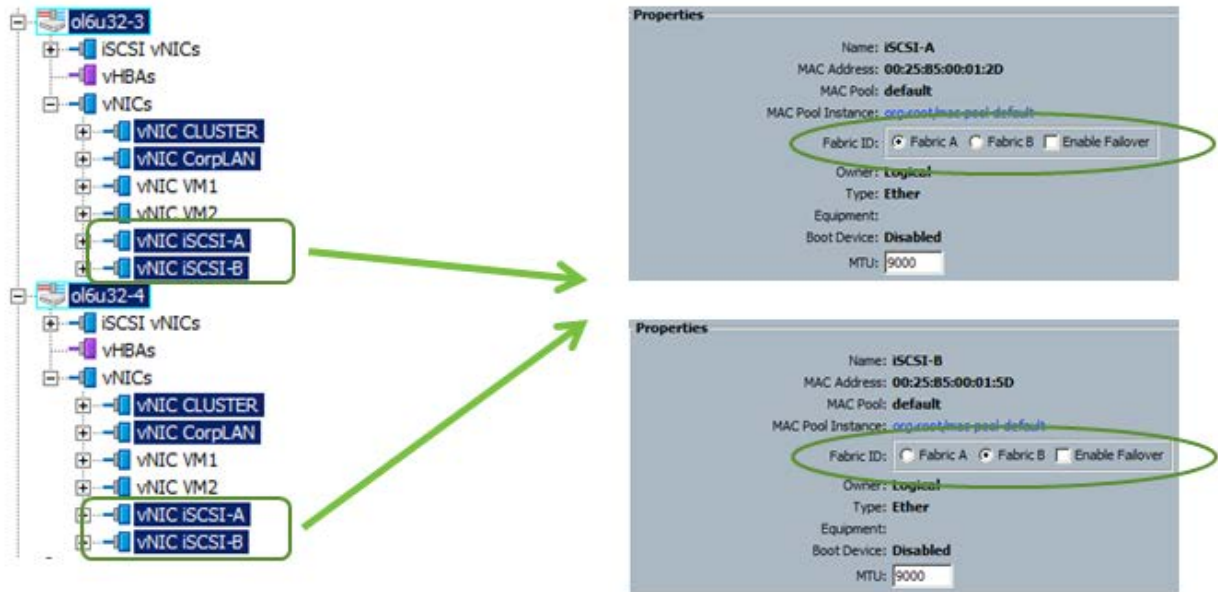
- Dual Fabric Interconnect
- Leverage UCS fabric NIC failover for management and cluster traffic only
- Use 10GigE for iSCSI traffic
- Leverage UCS Jumbo Frame (MTU 9000) for iSCSI traffic



Example of configuring the Cluster network and the Management network. Note the “Enable Failover” box is checked.

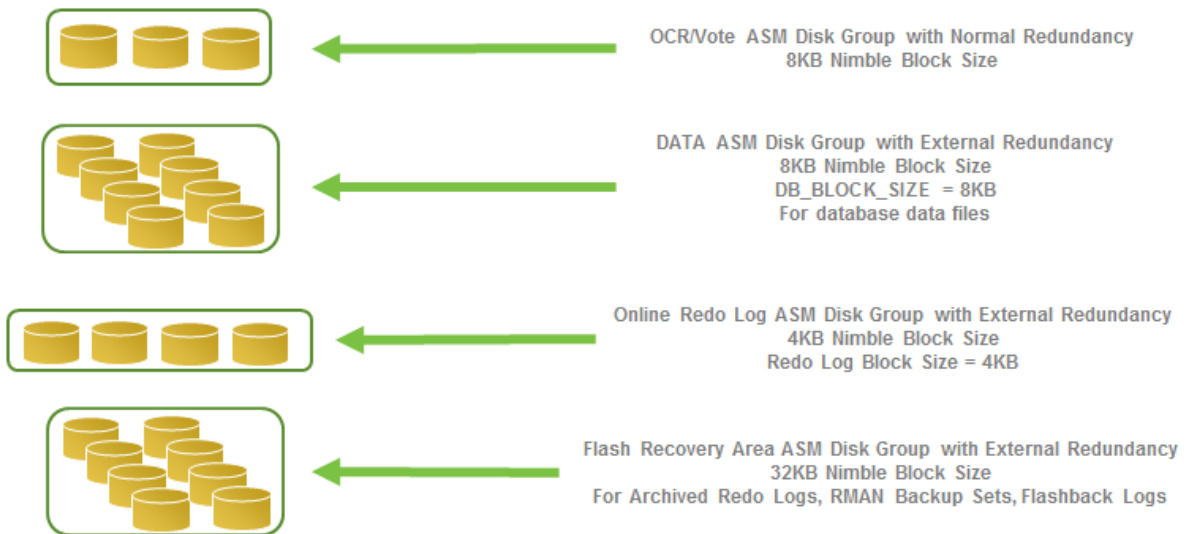


Example of configuring the iSCSI networks. Note the “Enable Failover” box is not checked.



3. Storage

- Nimble Storage CS-Series array has redundant everything
- All volumes provisioned use Linux DM-Multipath



2. Performance Optimizations

Server Operating System:

- Oracle Linux 6.x x86_64
 - Boot to Unbreakable Enterprise Kernel (default installation)
 - Use internal drive or Boot From SAN (Configure multipath for boot LUN if BFS)
- iSCSI Settings (/etc/iscsi/iscsi.conf)

```
node.session.timeo.replacement_timeout = 10
node.conn[0].timeo.noop_out_interval = 5
node.conn[0].timeo.noop_out_timeout = 10
```

- DM-Multipath Settings (/etc/multipath.conf)

```
defaults {
    user_friendly_names    yes
    path_grouping_policy   group_by_serial
    find_multipaths        yes
    features                "1 queue_if_no_path"
    path_selector          "round-robin 0"
}

devices {
    device {
        vendor            "NIMBLE"
        product           "Server"
        path_checker      tur
        rr_min_io_rq      1000
        rr_weight         priorities
        failback          immediate
    }
}

blacklist {
    wwid 26353900f02796769
    devnode "^\\(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^hd[a-z]"
}

multipaths {
    multipath {
        wwid            234eb4150c30140656c9ce900251fa9df
        alias           DATA1
    }
    multipath {
        wwid            2636311704885db8b6c9ce900251fa9df
        alias           DATA2
    }
    multipath {
        wwid            2636311704885db8b6c9ce900251fa9df
        alias           LOG1
    }
    multipath {
```

```

wwid      2636311704885db8b6c9ce900251fa9df
alias     LOG2
}
multipath {
  wwid      2636311704885db8b6c9ce900251fa9df
  alias     Add more LUNs.....
}
}

```

- /etc/security/limits.conf

```

# oracle-rdbms-server-11gR2-preinstall setting for nofile soft limit is 1024
oracle soft nofile 131072
# oracle-rdbms-server-11gR2-preinstall setting for nofile hard limit is 65536
oracle hard nofile 131072
# oracle-rdbms-server-11gR2-preinstall setting for nproc soft limit is 2047
oracle soft nproc 131072
# oracle-rdbms-server-11gR2-preinstall setting for nproc hard limit is 16384
oracle hard nproc 131072
# oracle-rdbms-server-11gR2-preinstall setting for stack soft limit is 10240KB
oracle soft stack 10240
# oracle-rdbms-server-11gR2-preinstall setting for stack hard limit is 10240KB
oracle hard stack 32768
oracle soft core unlimited
oracle hard core unlimited
oracle soft memlock 50000000
oracle hard memlock 50000000

```

- /etc/sysctl.conf

```

# oracle-rdbms-server-11gR2-preinstall setting for fs.file-max is 6815744
fs.file-max = 6815744
# oracle-rdbms-server-11gR2-preinstall setting for kernel.sem is '250 32000 100 128'
kernel.sem = 250 32000 100 128
# oracle-rdbms-server-11gR2-preinstall setting for kernel.shmmni is 4096
kernel.shmmni = 4096
# oracle-rdbms-server-11gR2-preinstall setting for kernel.shmall is 1073741824
kernel.shmall = 1073741824
# oracle-rdbms-server-11gR2-preinstall setting for kernel.shmmax is 4398046511104 on x86_64
kernel.shmmax = 4398046511104
# oracle-rdbms-server-11gR2-preinstall setting for net.core.rmem_default is 262144
net.core.rmem_default = 262144
# oracle-rdbms-server-11gR2-preinstall setting for net.core.rmem_max is 4194304
net.core.rmem_max = 4194304
# oracle-rdbms-server-11gR2-preinstall setting for net.core.wmem_default is 262144
net.core.wmem_default = 262144
# oracle-rdbms-server-11gR2-preinstall setting for net.core.wmem_max is 1048576
net.core.wmem_max = 1048576
# oracle-rdbms-server-11gR2-preinstall setting for fs.aio-max-nr is 1048576
fs.aio-max-nr = 3145728
# oracle-rdbms-server-11gR2-preinstall setting for net.ipv4.ip_local_port_range is 9000 65500
net.ipv4.ip_local_port_range = 9000 65500

# Increase these parameters
net.core.wmem_max=16780000
net.core.rmem_max=16780000
net.ipv4.tcp_rmem= 10240 87380 16780000
net.ipv4.tcp_wmem= 10240 87380 16780000

```

- /etc/profile (add to the end of the file)

```

if [ $USER = "oracle" ] || [ $USER = "grid" ]; then
  if [ $SHELL = "/bin/ksh" ]; then
    ulimit -p 16384
    ulimit -n 65536
  else
    ulimit -u 16384 -n 65536
  fi
fi

```

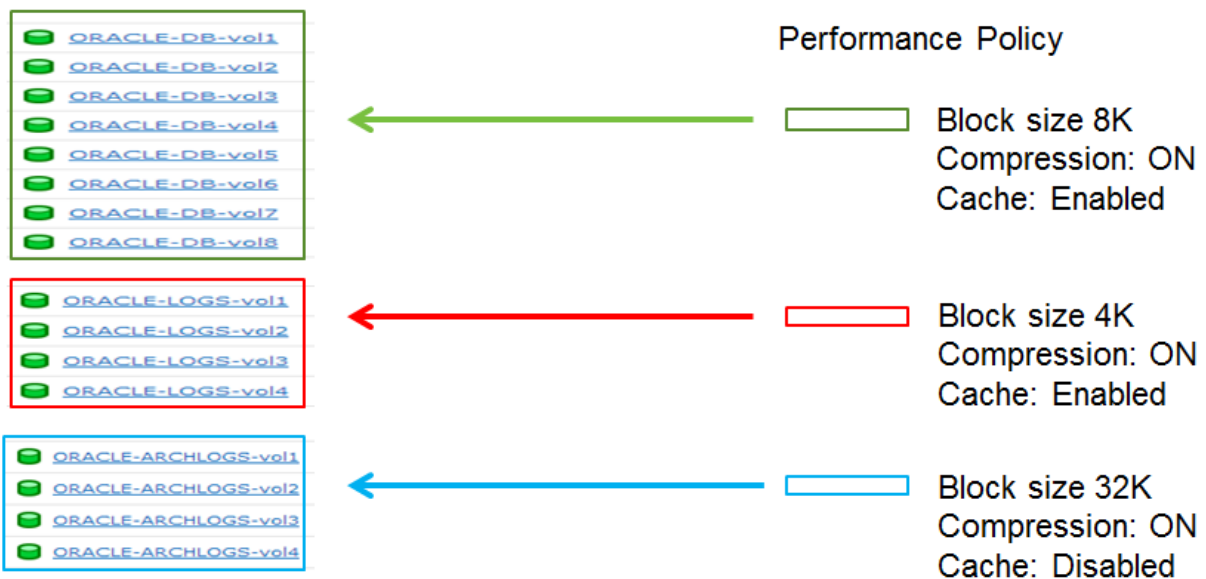
- IO Scheduler
 - noop

Oracle Database:

- Use Oracle Automatic Storage Management (ASM)
- Use External Redundancy for ASM disk groups
- Use default AU size (1MB)
- Configure the database to allow 4KB block size for Online Redo Logs
 - Change “_disk_sector_size_override=TRUE”
 - Re-create online redo logs with 4K option and remove the original

Storage:

Performance policy for database volumes:

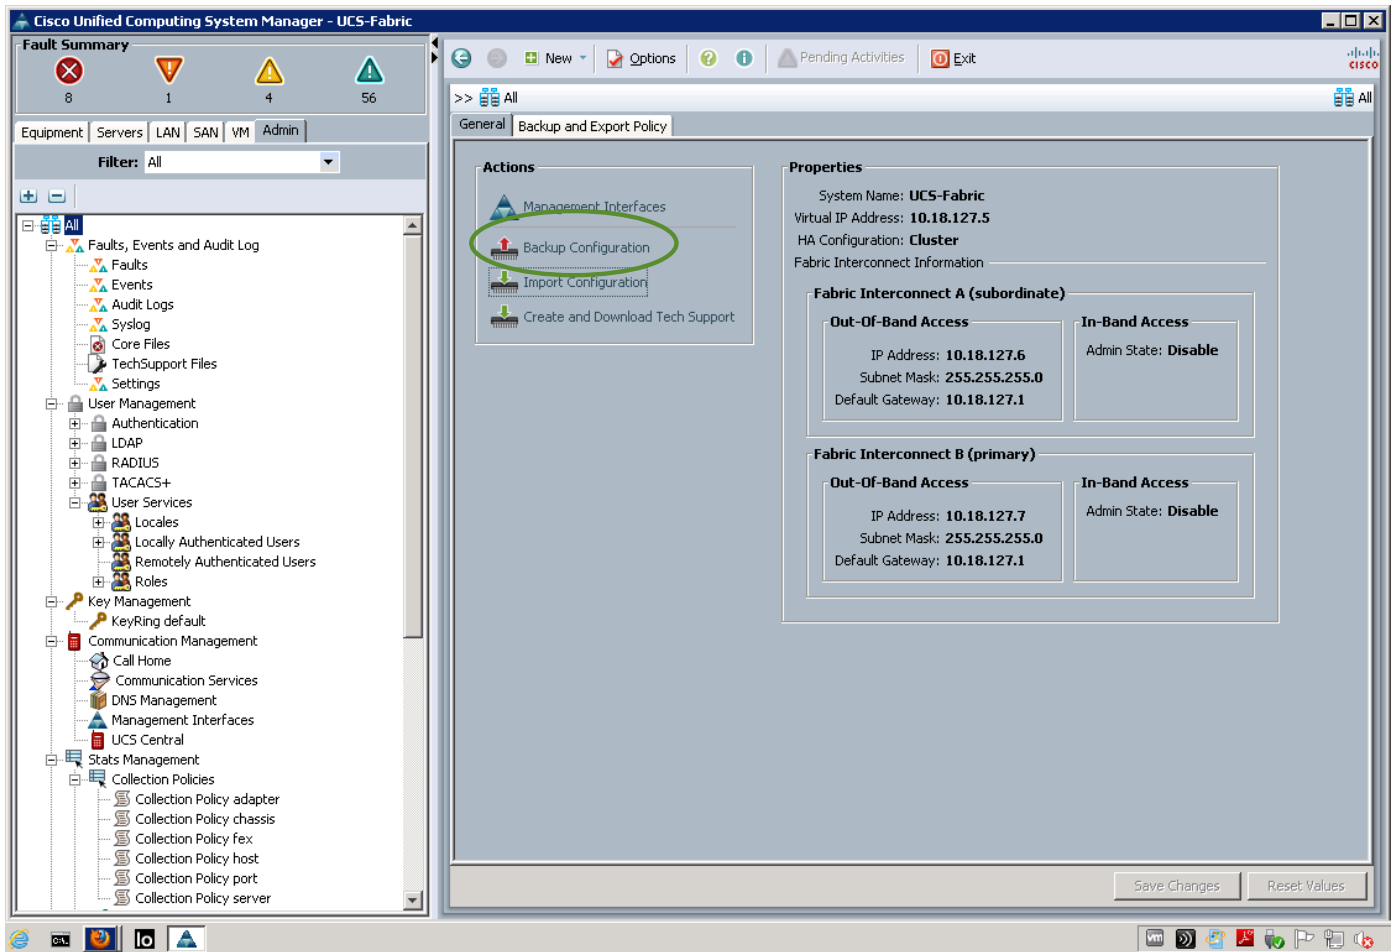


3: Data Protection

High Level Considerations

- Backup UCSM configuration
- All Oracle database volumes in same Volume Collection
- Volume Collection by application type + RPO

Infrastructure:



Application:

Oracle Database

-DB/Log/Archive/OCR-Vote volumes all in Oracle volume collection

Volume Collections > Oracle-ERP

The screenshot displays the Oracle Volume Manager interface for the Oracle-ERP volume collection. At the top, there are tabs for 'Status', 'Snapshots', and 'Replication'. Below the tabs, there are buttons for 'Edit...', 'Take Snapshot Collection...', 'Delete', 'Promote', 'Demote', 'Handover...', and 'Validate'. The status bar indicates 'Volumes: 19 Total Usage: 0 B Free: 18.85 TB'.

The interface is divided into three main sections:

- SYNCHRONIZATION:** Shows details for the volume collection, including Type (-), Server (N/A), Application (N/A), Username (N/A), and Password (N/A).
- PROTECTION STATUS:** Shows the last snapshot time (09/04 12:00 PM), next snapshot time (09/04 01:00 PM), and last complete replication (Unknown). It also includes a section for the "ORACLE-ERP" PROTECTION SCHEDULE, which details the snapshot frequency (1 hour), starting time (12:00 AM), repeat until time (11:59 PM), days of the week (Sun, Mon, Tue, Wed, Thu, Fri, Sat), and the number of snapshots to retain (5).
- ASSOCIATED VOLUMES:** Lists 19 volumes, including ORACLE-ARCHLOGS-vol1 through vol4, ORACLE-DB-vol1 through vol6, ORACLE-LOGS-vol1 through vol4, and ORACLE-OCR-VOTE-vol1 through vol3.

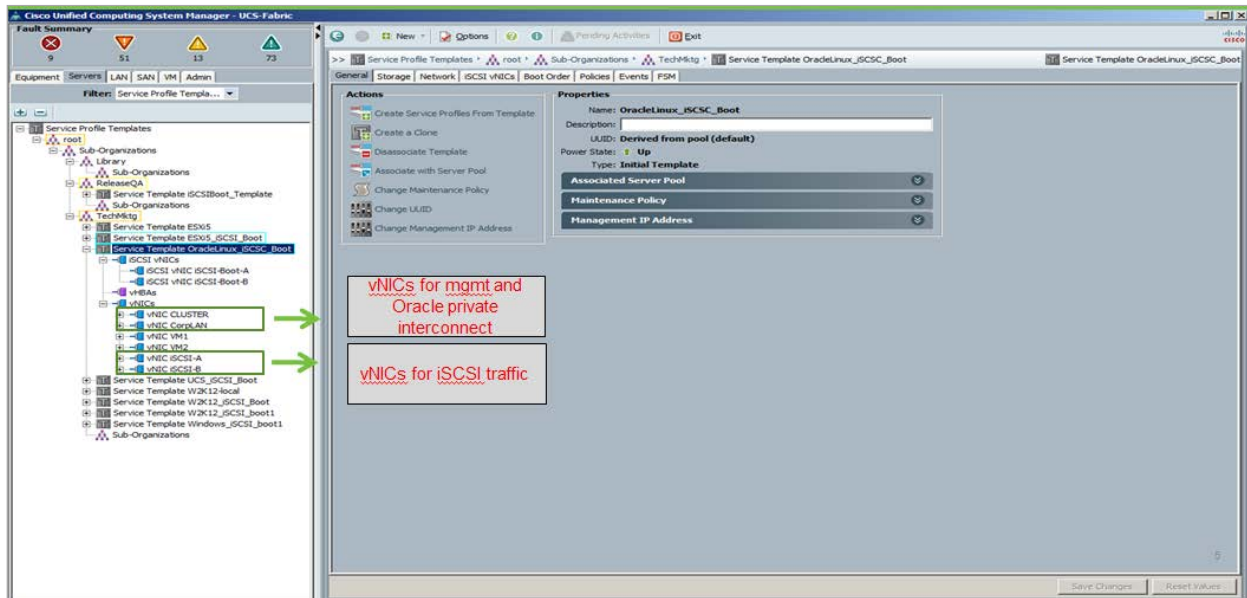
4: Operational Management

In this chapter, we will highlight tools and integrations that help making deployment and operational management simple and easy.

UCS Service Profile:

A custom UCS Service Profile template was created for Oracle RAC database. It creates a standard for deploying the Oracle environment serving business-critical applications, and simplifies scalability expansion down the line.

For vNIC



Nimble Storage InfoSight:

It is a good practice to regularly monitor Nimble InfoSight™ for storage health, availability, performance, and data protection reports based on heartbeats from the array:



Appendix: Oracle Validated Configuration

The following table includes details of the validation of SmartStack performance with real applications.

Configuration Summary	Oracle Database 11gR2 and 12c Single Instance and Oracle Real Application Cluster (RAC) for Linux x86_64 on Oracle Linux 6 Update 4 using ASMLib
Server Platform	Cisco UCS Blade B200 M3
Storage Model	Nimble Storage CS400 Series
Oracle Software	Oracle Database 11gR2 (11.2.0.3) and 12c (12.1.0.1.0) for Linux x86-64
Linux Distribution	Oracle Linux 6 Update 4 x86_64 with Unbreakable Enterprise Kernel

Server Model	2 x Cisco UCS Blade B200 M3
Processors	2 x 16 cores Intel® Xeon® CPU E5-2600
Memory	256GB
OnBoard Storage	300GB 6Gb SAS 10K RPM 2.5" HDD/hot plug/drive sled mounted
Network/Interconnect	2 x Cisco Fabric Interconnect
Multipath	Device mapper multipath 0.4.9-64.0.1.el6
Storage Model	Nimble Storage CS400 series 1.4.x
Kernel	2.6.39-400.17.1.el6uek.x86_64 or higher

NOTE: The Oracle Validated Configuration was conducted with Oracle 11gR2 and 12c Real Application Cluster (RAC).

Oracle Database 11g Release 2

http://linux.oracle.com/pls/apex/f?p=102:2:4213554191891752::NO::P2_VC_ID:627

Oracle Database 12c Release 1

http://linux.oracle.com/pls/apex/f?p=102:2:4213554191891752::NO::P2_VC_ID:631



Nimble Storage, Inc.

2740 Zanker Road., San Jose, CA 95134

Tel: 877-364-6253; 408-432-9600 | www.nimblestorage.com | info@nimblestorage.com

© 2013 Nimble Storage, Inc.. Nimble Storage, CASL, InfoSight, SmartStack, and NimbleConnect, are trademarks or registered trademarks of Nimble Storage, Inc. All other trademarks are the property of their respective owners. BPG-SMRT-ORACLE-0913