HIGH-AVAILABILITY

steria

RSF-I

GE Access

Partnering for Success
STORM

High-Availability

Implementation Plan

Introduction

This document is intended for use during the implementation phase. It is also possible that components will be useful during pre-sales, post implementation and planning phases but that is not a primary objective of this document.

At High-Availability.Com we believe that good planning and documentation greatly improve the chances of a smooth implementation and good customer perception. Steria's STORM solution is being marketed and sold by Steria and the hardware, OS and clustering components are being supplied by GE Access. High-Availability.Com meet Steria's and GE Access's clustering needs by supplying RSF-1 through GE Access.

System for Tasking and Operational Resource Management (STORM)

Policing today is increasingly IT-centric. In the field, forces need greater ability to identify suspects, respond rapidly to incidents and deploy resources as effectively as possible. Behind the scenes, information needs to be shared between forces, and increasingly – in the fight against terrorism and other international crimes – across national boundaries, securely and swiftly.

Designed to give greater clarity in incident handling situations, STORM effectively marries information consolidation and resource allocation in one user-friendly interface. The result is that operational room staff can instantly see what type of incident is unfolding, what resources are needed to deal with the situation and the best way to allocate these resources to resolve the incident.

STORM Features and Benefits

STORM gives the Police the ability to manage and view all incidents, whether of a minor nature or a major disaster (a plane or rail crash, or terrorist attack, for example) with a greater degree of transparency. This means operatives can direct personnel and resources to a situation faster and more effectively. Typical implementations of the STORM system include:

- A dual screen interface displaying all relevant data
- Tight integration with mapping GIS
- Interface with other adjoining forces and national systems such as National Police Computer, Automatic Number Plate Recognition and the National Firearms System
- The ability to view skill sets, ensuring that all situations are attended by personnel with the necessary expertise
- Training of STORM operators, to ensure the system could function operationally at go-live
- A clear migration path to mobile capability, including tablet PC and PDA operation, and silent communication functionality

Key Steria Services to Police:

STORM (Command & Control)
Custody management
Crime management
Duty rostering
Imaging & biometrics
Electronic documents and records management
Rapid business process improvement
Demand management

High-Availability.Com

High-Availability.Com (HAC) was formed in 1995 and is well established as an application vendor providing clustering technology to businesses in every sector. The company operates 24x7 and has many high profile customers including the Tokyo Stock Exchange and the US Navy. The company's core product, RSF-1®, has been widely adopted in government and emergency services sectors and is very strong in the UK education market place.

RSF-1 is a middle-ware product that is invisible to the end user, who will benefit from its use but should never be aware of it. RSF-1 dramatically improves reliability by managing failure and automatically re-deploying an application to another location in the event of a failure.

RSF-1 has been tested at certification centres around the world including Sun, IBM, Check Point, Remedy and many others. The product is available on a wide variety of operating systems and integrations have been tested and certified by OS and application vendors (SunTone, StorageProven, Quality Partners, OPSEC etc.).

NetMon is a network interface fail-over facility, which ensures that the most common single point of failure, the NIC, is eliminated. Network interfaces are continually monitored and standby interfaces activated if a problem is detected. NetMon sends out packets on each interface it is monitoring and waits for a response from a defined target. When responses are received NetMon infers that the interface is up and running, even if not configured /plumbed in. If a response fails to be received when expected NetMon can infer that the interface has in some way failed.

IBM TotalStorage PROVEN™
SUNTONE™ CERTIFIED
ORACLE®
Remedy® a BMC Software company
Authorized MySQL™ Software Partner
Sun microsystems STRATEGIC DEVELOPER

Architecture

STORM has undergone significant development and can now be considered an 'off-the-shelf' solution which can be readily customised to the requirements of each force. At the 'heart' of STORM is a database server, this is integrated with systems such as the CAD server, PNC, HQ computer room facilities and CEDA. The system is typically deployed with clustering and a fail back site, for some forces the clustering also extends to the clustering site.

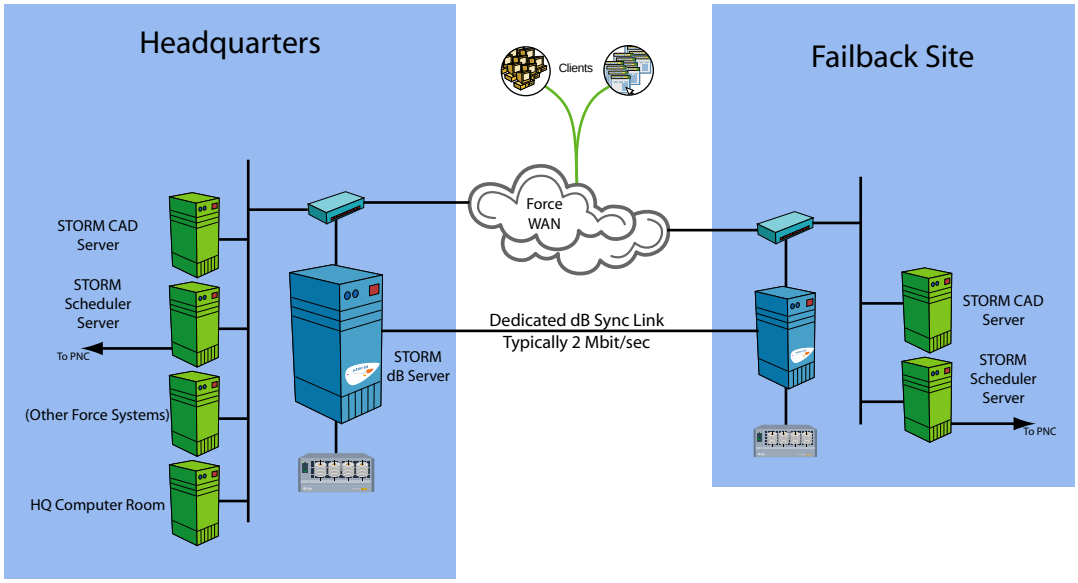

Figure 1.
High-Level STORM Architecture

This high level (Figure 1.) view of an installation can be usefully expanded to show the underlying components in a more resilient solution. Figure 2. Shows the logical connections between the replicated databases. Figure 2. also shows the clustered dB servers. This architecture is resilient to multiple simultaneous failures.

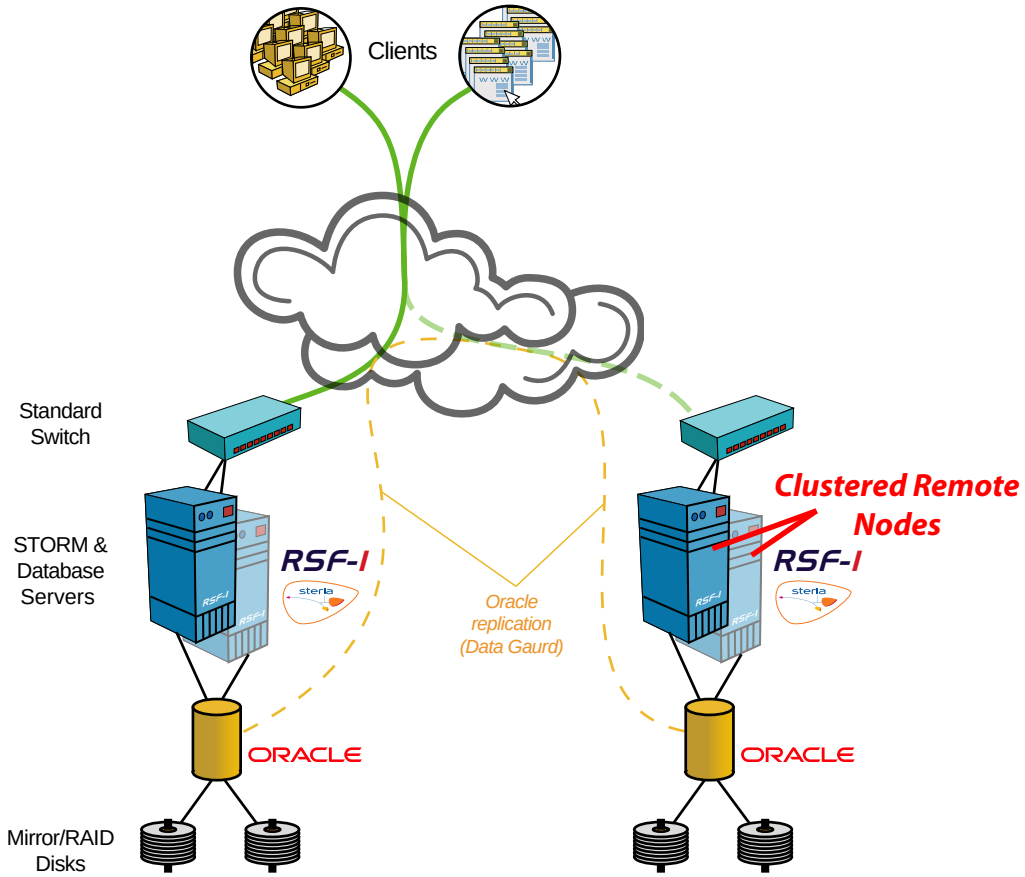

Figure 2.
STORM dB Replication & Clustering

Installation Process

The delivery of STORM to the customer in a smooth fashion requires a little planning. Sample questionnaires are supplied in Appendix A and should be completed and reviewed by the appropriate organisation prior to equipment being sent to site.

The clustering software can be installed once GE Access have finished but testing can not be completed until STORM is installed. Therefore this document sets out a process that provides enough detail to allow Steria to install STORM in such a way that it should not require substantial adjustment during the installation of the clustering. The Cluster installation will normally be done with Steria still on site but having completed the majority of their work.

Purely as an example the following Gantt chart shows the tasks. Normally there should be gaps between the processes to allow for delays but this example is intended to illustrate which processes can overlap and which can not.

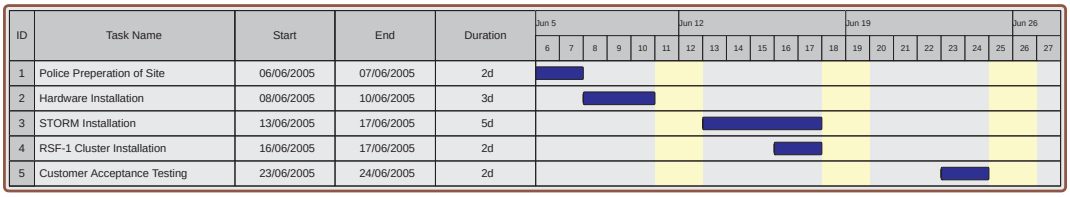

ID	Task Name	Start	End	Duration
1	Police Preperation of Site	06/06/2005	07/06/2005	2d
2	Hardware Installation	08/06/2005	10/06/2005	3d
3	STORM Installation	13/06/2005	17/06/2005	5d
4	RSF-1 Cluster Installation	16/06/2005	17/06/2005	2d
5	Customer Acceptance Testing	23/06/2005	24/06/2005	2d

Figure 3.
Example Gantt chart for STORM installation

Facilities to be provided by the Police

The details of facilities that are expected to be provided by the police vary depending on the customer but often these will include;

- Power for 19" racks and for engineer's lap tops etc.
- Networking (hubs) - there must be a network port for every NIC expected to be connected plus two for engineers. Each server will probably have two connections and each disk array one connection and the terminal concentrator will have one connection - so at least **8 Ethernet ports** should usually be available.
- Acceptance that engineers will need to connect lap tops etc. to the network to configure and administer the installation, this could be done with 'force' equipment rather than private but must include the ability to ftp and to use 'Putty'.
- IP Address information for each NIC, including DNS, subnet & router information. In addition a 'floating' IP address per active service (usually only one for STORM) must be supplied. Temporary addresses for lap tops will be required.

- Desireable - external access using a browser and possibly Email or ISP access. The network may need to be firewall'd out from the main network or other arrangements made during installation.

Base Install - GE Access

The systems should be built and positioned in their proper places and all cabling completed. The serial heartbeat cable can be supplied in advance if required but will otherwise be installed by HAC engineers. The terminal server should be connected to the arrays and servers and once configured the connectivity must be tested.

The OS should be installed and patched with the appropriate patches. The system disks should use disk mirrors (RAID 0.)

The requirements and process for installation of the shared disks arrays is covered in HAC guides for the 3310 and extensively in the guide for the 3510, so the information will not be duplicated here. However, note the disks should be laid out in at least three partitions, with hardware RAID enabled. One very small (~10Mb) for disk heartbeats, one (/u02) for Oracle binaries and redo/logs. These two partitions should be logical volumes on the same physical disks (normally 2x36Gb, RAID 0) and the primary Oracle data area (mount point /u01) must not use any part of these disks. The remaining disks are for the primary Oracle data partition/logical volume (/u01) - RAID 5.

The disks should be added to Solaris Volume Manager (a.k.a. Solstice Disk Suite - SDS) or Veritas Volume Manager as shared disks and tests should be performed prior to leaving site. One machine will be designated 'primary' for the installation process and this should automount the shared file systems until the clustering (RSF-1) has been installed, the other machine should not automount the shared file systems. Once RSF-1 is installed these file systems will not be automounted by the OS.

Where NetMon is being used then two interfaces should be connected to each LAN for each server.

Each network interface (servers, disk arrays, terminal servers etc.) should be configured with the appropriate IP address and connectivity testing completed to verify that the machines are working correctly.

An installation sheet should be completed and left with the equipment so that other engineers can quickly identify the components and their configuration information. A sample configuration sheet is to be found in Appendix B.

The following lines should be added to **/etc/system** on each node;

```
* Primary Oracle dB - Max 1000 processes = 1000 * 2 + 10, but allow 500 extra semmn's contingency

set semsys:seminfo_semmni=100
set semsys:seminfo_semmns=2510
set semsys:seminfo_semmsl=512
set shmsys:shminfo_shmmax=4294967295
set shmsys:shminfo_shmmin=1
set shmsys:shminfo_shmmni=100
set shmsys:shminfo_shmseg=10

* Disable the console break / stop key - safety measure with clustering

set abort_enable=0"
```

STORM Install - Steria

The systems are now ready for Steria to complete the installation and configuration of Oracle and the STORM application suite. Steria provide detailed information for this process and often customisation for each site is required, however, the notes in this document are targeted at clustered environments specifically.

Oracle should be installed on the appropriate nodes; only one node in a cluster needs the installation if the binaries are being installed onto the shared disk - which is normal for Steria installations.

The Steria start/stop script for (Appendix C) is written in the korn shell (ksh) and will not work if placed in /etc/rc3.d/ directly. If the node is being clustered then when HAC complete their work this problem will be solved and meanwhile the script can be run manually when required. However, if the node is not being clustered, as with the example 'Single Remote Node' in Figure 4. below, then additional work must be completed now;

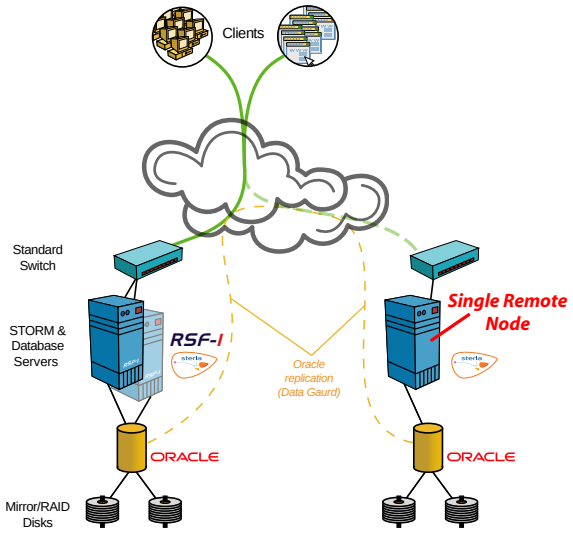

Figure 4.
Example of Two Clustered Nodes and a Single Remote Node

Install the korn script as /etc/rc3.d/oracle
Install the borne 'starter' script (Appendix C) as /etc/rc3.d/S30oracle
Link the S script to a K script. E.g.:- cd /etc/rc3.d/ ; ln -s S30oracle K70oracle

In all cases the script(s) must be made executable (chmod 555) and should be owned by root (chown root:sys).

Additionally the Steria scripts depend on RSF-1 scripts contained in HAC product packages (HACbase and HACrsf), which must be installed on all nodes. If clustering will not be used, as with 'Single Remote Node' in Figure 4. above, then /etc/rc2.d/S97rsfrc can be removed. The current packages can be found on the web http://www.high-availability.com/Downloads/.

The ORACLE_SID must be set to the correct value; e.g.:- STRM. Note that Data Guard nodes may have a different value, E.g.:- STRMSB. Check /var/opt/oracle/oratab & ~oracle/.profile & the scripts mentioned above.

Setup and check the Steria Data Guard scripts in $ORACLE_BASE/admin/dgd/; dbautostart.sh & dbautostop.sh.

Cluster Install - HAC

RSF-1 and NetMon[1] will be installed and licensed and will include the following tasks;

RSF-1 & OS Setup

Serial heartbeat setup and configuration
E.g.:- To use serial port /dev/ttyb - '/usr/sadm/bin/smserialport configure -p <root_passwd> -- -n b -x device=modemdialout -b 38400 -c "RSF-1 Heartbeat" -x service=y'
Disk heartbeat configuration
Network heartbeat configuration
Verify shared file system mount points exist
Ensure DNS is not relied upon for cluster component name resolution
Enable eeprom "local-mac-address?=true"
Create matching "rsfpasswd -u _rsfadmin" passwords on all nodes

RSF-1 Service Setup

Service configuration
Floating IP address setup
Set appropriate disks to be mounted by each service
Move any start service(application) scripts to /opt/HAC/RSF-1/etc/rc.<servicename>.d/
Run rsfklink to link the start scripts to stop scripts in the correct order

RSF-1 Oracle Service Specific - all nodes

Check /var/opt/oracle/oratab setup should end 'N' if Data Guard is in use, 'Y' if not.
Copy the following directories from the node on which Oracle was installed upon to the other node(s);
/usr/local/bin
/var/opt/oracle
Set the Oracle listener to listen on the floating IP address

Appendix A - Planning Questionnaires

Police Force

Primary Location - Normally HQ	
Location Name	
Primary Machine Name	
Primary Machine IP Address	
Secondary Machine Name	
Secondary Machine IP Address	
Domain Name	
Netmask	
Gateway IP Address	
DNS Server 1	
DNS Server 2	
DNS Domains to Search	
Steria Service (floating) IP Address	
Console Concentrator IP Address	
Disk Array IP Address	
Number of network ports available - total	
Access Restrictions?	

[1] - NetMon is an option, so may not be installed

Secondary Location - Disaster Recovery - Optional	
Location Name	
Primary Machine Name	
Primary Machine IP Address	
Secondary Machine Name	
Secondary Machine IP Address	
Domain Name	
Netmask	
Gateway IP Address	
DNS Server 1	
DNS Server 2	
DNS Domains to Search	
Steria Service (floating) IP Address	
Console Concentrator IP Address	
Disk Array IP Address	
Number of network ports available - total	
Access Restrictions?	

Steria

Solaris Version Required	Solaris 9
Total Disk Space Required	
Oracle primary partition mount point	/u01
Oracle redo / log partition size	32Gb RAID 0 less a small partition for RSF-1 Heartbeats * these two partitions together on own physical disks
Oracle redo / log partition mount point	/u02
Volume Management	Solaris Volume Management (Solstice Disk Suite)
Desired Disk Volume Group Name	STORM

GE Access

Primary Location - Normally HQ	
Solaris Patch Version	
Primary Machine hostid	
Primary Machine 'root' password	
Secondary Machine hostid	
Secondary Machine 'root' password	
Disk Array Password	

Secondary Location - Optional	
Solaris Patch Version	
Primary Machine hostid	
Primary Machine 'root' password	
Secondary Machine hostid	
Secondary Machine 'root' password	
Disk Array Password	

Appendix B - Sample Configuration Sheets

Primary Location	
Primary Machine RSF-1 License	
Primary Machine NetMon License	
Primary Machine Oracle ‘system’ Password	
Secondary Machine RSF-1 License	
Secondary Machine NetMon License	
Secondary Machine Oracle ‘system’ Password	
‘_rsfadmin’ Password	
Secondary Location	
Primary Machine RSF-1 License	
Primary Machine NetMon License	
Primary Machine Oracle ‘system’ Password	
Secondary Machine RSF-1 License	
Secondary Machine NetMon License	
Secondary Machine Oracle ‘system’ Password	
‘_rsfadmin’ Password	

Primary Location Server Example Cabling

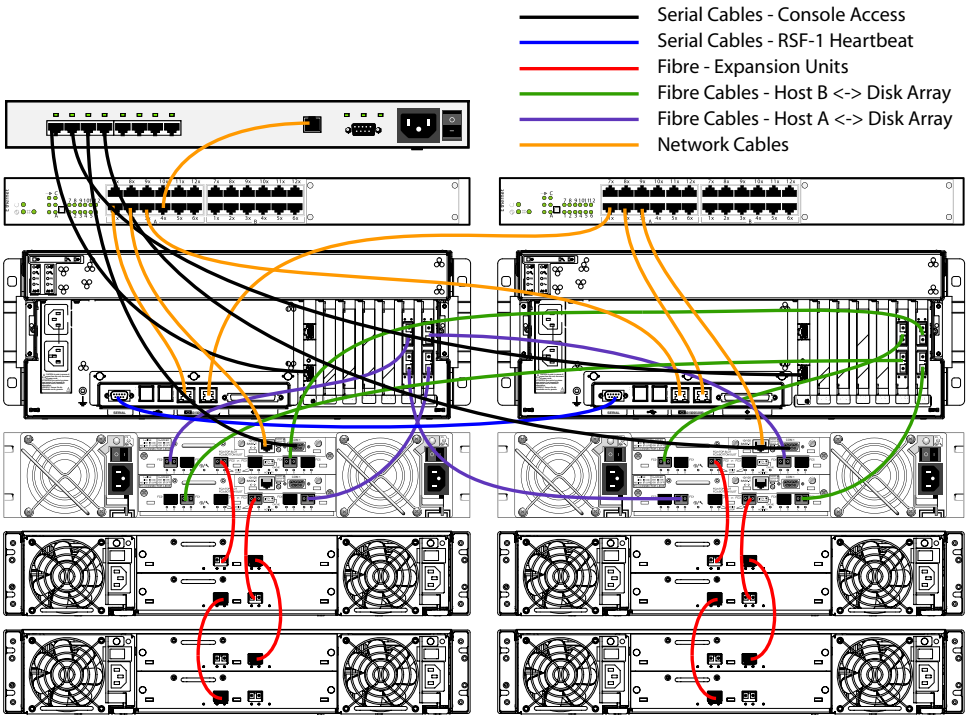

Figure 5.
Example Server Cabling Diagram

Secondary Location Server Example Cabling

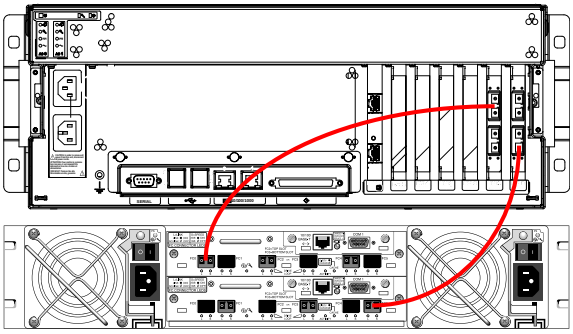

Figure 6.
Example Single Server Cabling Diagram

Appendix C - Steria's Start/Stop Scripts

The following scripts have been developed to allow RSF-1 to control STORM. They are included here for reference but may vary on site, depending on customer requirements etc.

Steria Korn Start Script

```
#!/usr/bin/ksh
#
# Script:                 oracle
#

. /opt/HAC/bin/rsf.sh

export PATH=$PATH:/usr/local/bin
export ORACLE_SID=STRM
export ORAENV_ASK=NO
service="oracle"

host=`uname -n`
script="$0"
SCRIPTS=/u01/app/oracle/admin/dgd
state=$1                # starting or stopping?
attempts=${2:-'1'}      # no of attempts at this action

# function for log mesg
dated_echo()
{
        date_now=`date +'%b %d %H:%M:%S'`
        echo "${date_now} ${host} ${script}[$$]: $*"
}

if [ ${attempts} -gt 3 ]; then
        # fail after three goes
        dated_echo "${service}: Too many retries on ${state}, aborting"
        exit ${COMPANY_ABORT}                   # abort code
fi

start_main ()
{
        . oraenv
        if              [ $? -ne 0 ];then
                        dated_echo "oraenv failed - SID $ORACLE_SID"
                        exit ${COMPANY_SAFE}
        fi

su - oracle "$SCRIPTS/dbautostart.sh"

        # uncomment next line if OEM required
        # agentctl start &

        su - oracle 'lsnrctl start STANDBY'

        # now start the STRM listener immediately for speed of getting node active,
        # but then wait to see if we have a standby database.  If so, shut down the STRM
        # listener again

        su - oracle 'lsnrctl start STRM'
        dated_echo "Started listeners, allow Data Gaurd to evaluate state"
        # see if database is mounted in standby mode
        dated_echo "Checking to make sure that $ORACLE_SID is not a standby database"
su - oracle "(export ORACLE_SID=$ORACLE_SID
export ORAENV_ASK=NO
. oraenv
sqlplus -s /nolog <<!EOF
connect / as sysdba
select 'DB=',name, open_mode, database_role from v\\\$database;
exit
!EOF
) "|grep "DB=" | read BLURB NAME OPENMODE SBSTATE
        if [ "$OPENMODE" != "READ" ]; then
                        if [ "$OPENMODE" != "MOUNTED" ]; then
                                        dated_echo "Database not mounted after 1 minute - manual intervention required!"
                                        dated_echo "leaving listener STRM up"
                                        exit ${COMPANY_RESTART}
                        fi
                        if [ "$SBSTATE" != "PHYSICAL STANDBY" ]; then
                                        dated_echo "$ORACLE_SID is not physical standby - it is $SBSTATE"
                                        exit ${COMPANY_RESTART}
                        else
                                        dated_echo "$ORACLE_SID is a Standby Database, so closing listener STRM"
                                        su - oracle 'lsnrctl stop STRM'
                                        exit ${COMPANY_OK}
                        fi
        else
                        dated_echo "$ORACLE_SID is OPEN, leaving listener STRM active"
                        exit ${COMPANY_OK}
        fi
        exit ${COMPANY_OK}
}

stop_main()
{
```

```
        . oraenv
        if  [ $? -ne 0 ];then
                dated_echo "oraenv failed - SID $ORACLE_SID"
                exit ${COMPANY_ABORT}
        fi
su - oracle "$SCRIPTS/dbautostop.sh"

        # uncomment next line if OEM required
        # agentctl stop &

        su - oracle 'lsnrctl stop STANDBY'
        su - oracle 'lsnrctl stop STRM'
        exit ${COMPANY_OK}
}

case "${state}" in

'start')
        dated_echo "${service}: Starting application"
        #
        # application startup commands here
        #

        start_main
        exit ${COMPANY_OK}
        ;;

'stop')
        dated_echo "${service}: Stopping application"

        stop_main
        exit ${COMPANY_OK}
        ;;

*)
        echo "Usage: $0 <start|stop>"
        exit ${COMPANY_WARN} # warning code
        ;;

esac

exit ${COMPANY_OK}              # OK code (default)
```

Steria Borne OS 'Starter' Script

```
#!/sbin/sh
#
# Script:               S30oracle
#

/etc/rc3.d/oracle $*
```

HAC DataGuard State Script

```
#!/bin/ksh
#
# Script:               /opt/HAC/RSF-1/etc/rc.<service_name>.d/db_state.sh
#
# Determines if the database is the primary based on the host or whether the dB is in standby mode
#       used by /opt/HAC/RSF-1/bin/monitor.sh
#
#

export PATH=$PATH:/usr/local/bin
export ORAENV_ASK=NO

. oraenv

(sqlplus -s /nolog <<!EOF
set pages 0
connect / as sysdba
select 'DB= ',name,open_mode, database_role from v\$database;
exit
!EOF
) | grep 'DB='|read BURB PBNAME MODE MODE2 MODE3
OPENMODE="${MODE}${MODE2}"
if [ "$OPENMODE" != 'READWRITE' ]; then
 # oramon - looks like dB is not primary
 if [ "$MODE" = "MOUNTED" ]; then
   SBMODE="${MODE2}${MODE3}"
   if [ "$SBMODE" != "PHYSICAL STANDBY" ]; then
     # oramon - will not test dB as it's in standby mode
     exit 1
   fi
 fi
 exit
else
 # oramon - dB should work fine - do the tests
 exit 0
fi
```

Appendix D - Validation & Testing Configuration

The following diagram depicts the systems setup at GE Access for cluster integrated test, validation and demonstrations purposes.

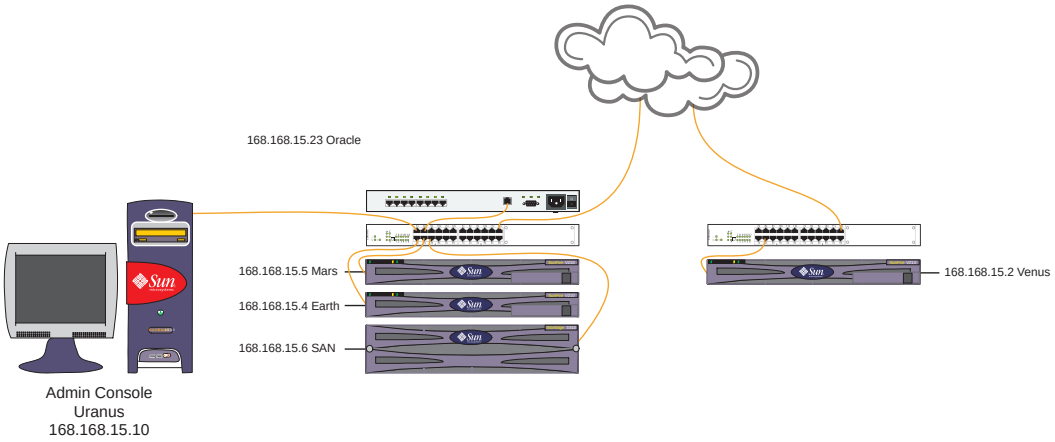

Figure 7.
GE Access Reference Model Setup

This is included here purely for refernce purposes. The compressed tar files of the setup and testing record is held by Giles Gamon and Kenneth (Ken) Scullion.