



HIGH-AVAILABILITY

Acme Inc.



ORACLE®



Remedy Clustered Servers

High-Availability

Site Audit

Introduction

RSF-1 has been used successfully at Acme for a number of years to cluster the *Remedy AR System*, making it highly available to the organisation and its customers. The cluster originally consisted of an E450 and an E4500, later being upgraded to two E4500 nodes with associated shared storage. The AR System uses an Oracle back-end database to store data.

Recently the servers have been replaced by newer models of server and disk. The detail of the cluster will be covered in the following sections but a simplified layout of each server set is shown in Figure 1. The system comprises of a Sun V480 server with Sun S1 StorEDGE disk expansion unit and a 3510 fibre channel array.

Following a failure of some hardware on the 3510 arrays and subsequent action to identify and repair the system Acme have asked High-Availability.Com (HAC) to complete an audit of the systems identifying any areas for improvement and identifying single points of failure. This report is the result of that audit, which was completed during the evening of 22nd March 2005 by Giles Gamon of HAC accompanied by Joe Bloggs of Acme.

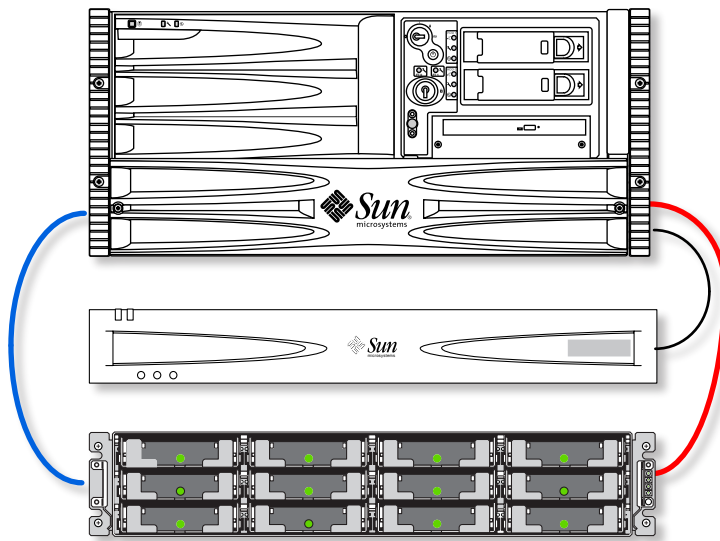


Figure 1.
Single Node Setup

The operating system is installed on the internal disks and mirrored on to the S1 StorEdge expansion unit. Oracle and AR System are installed on the 3510 array, which is mirrored using Solstice Disk Suite to another 3510 array attached to the other server.

A Netra t1 with L280 tape storage device is used to maintain backups of the system and data but is outside the scope of the audit conducted. There are two additional machines installed as part of the overall system. They provide for a development environment and a customer acceptance testing environment for pre-launch testing of any modifications etc. Again these systems were not included in the scope of the audit but point to good working practices.

Scope of Work

The audit was expected to include the following items;

- Examine the configuration of the Remedy cluster with a view to identifying single points of failure or other issues which may compromise the reliability or availability of the system.
- Review of the configuration of;
 - Hardware
 - Operating system
 - RSF-1 cluster software
- Carry out a failover operation in both directions. Successful restart and normal operation of the Remedy application needs to be tested during this phase.
- Make recommendations for relevant upgrades or additions to the RSF-1 cluster software or its modules.

Audit Process

The audit conducted was completed without making any changes to the system at the request of Acme. The audit began at approximately 17:00 with a set of non-intrusive tests and measurements. Failover testing was scheduled for and started at approximately 19:00.

Once all testing was complete the system was returned to exactly the same state that it was on in on HAC's arrival and functionality was tested and confirmed by the user base before HAC departed.



Physical Environment

All of the machines are situated in the same room at Acme's Big House, in High Street. The room is kept at a suitable temperature and the CPU reported temperatures between 52 and 56 degrees centigrade, which is within the acceptable working temperature. rsfnode2 *felt* as though the surrounding area was hotter, which may be a result of the fairly dense racking around it. This was not actually measured so is simply an observation.

No investigation into the power, air conditioning or the network infrastructure was conducted but they are believed, by Joe Bloggs, to offer adequate resilience. It should be noted that there is some exposure to risk from the fact that the systems are located in the same building and on the same floor.

The wiring inside the rack is of a reasonable standard but there is room for improvement, the fibre cabling in hung over rack supports and the effective bend radius is substantially less than recommended. **HAC recommend** that the cabling be tidied up and specifically the fibre cabling be dealt with, perhaps with some strengthened cable guides. The fibre cabling is of a light weight type, in that it is not fitted with loose jacket filled with kevlar strands, it is fitted with a fairly tight jacket. The fact that the fibre is of the light weight variety is not a problem but extra care must be taken with this fibre and any ties should be loosely fitted.

rsfnode2

The audit began with this machine as it was not running the live service. Further, this system was not set to automatically failover in the event of a problem.

Operating System

The server consists of a V480R server running a full install of Solaris 8, quad 900MHz processors with 8Gb RAM. The system is expanded with a quad network PCI card (qfe) and two fibre channel PCI cards, for the 3510 arrays, and a PCI SCSI card for connecting to the S1 StorEDGE unit. The OS patch level is fairly recent (Dec '04).

The 'cron' jobs are setup to run some components that will not run if the AR System application/Oracle is not active on this node. The errors that result are of little concern but they do generate an Email each time a cron job is run. There are also some cron jobs that when run successfully create Emails. The Emails are not being checked and date back to 1st November 2003. The root account mail account is currently 32Meg in size.

HAC recommend that the cron jobs be revised to include '*is cluster active?*' test prior to running each job and that the jobs which create output (stdout/stderr) be revised where appropriate to reduce unwarranted Emails, which should probably only be raised in the event of attention being required. Further, **HAC recommend** that the Emails generated are checked regularly, or even better forwarded to a group of users responsible for administering the system.

The /etc/hosts file and /etc/nsswitch.conf is generally well configured and does not result in RSF-1 depending on DNS, which is good. However, the *loghost* is specified as the floating cluster identity, *remedy-svc*, which means that logs will not always be retained, as the floating identity could be *down*. **HAC recommend** that the loghost be changed to be the local machine (i.e. rsfnode2 in this instance, or localhost.)

The '*remedy*' identity is set on each node as the local node identity (i.e. primary IP address). This has probably been done to accommodate AR System, which by default will use the primary machine address as the 'end point' and name for client connections. However, AR System can be configured to change this default setting, which will allow the remedy identity to be separated from the local machine identity. **HAC recommend** that the following optional setting, not currently set, be understood and Acme consider maintaining a separate 'remedy' identity;

```
/etc/ar.conf
AR server Server-Name: remedy
```

The 'remedy' identity is that used inside AR System schemas and active links etc. it is also the identity used by client connections into the system. There is a 'unix' identity which is currently configured in the same way as the remedy identity but the reason was unknown. This should be discussed and **possibly revised**.

The operating system disks, internal to the V480R, are properly laid out and all mirrored with the external S1 StorEDGE. There are no reported errors with the mirrors.

The V480R comes with two 'ce' interfaces and four more have been provisioned by use of a PCI 'qfe' card. The network connections and diagrams available did not fully explain the layout of what are thought to be backup and heartbeat networks. The 3510 array is connected for administrative purposes to the V480R via qfe2, however, while the physical connections appeared OK and the V480R configuration appeared to be correct, the status lights indicated a fault and the 3510 was not accessible, so could not be checked. **HAC recommend** that the network connection to the 3510 be repaired and the status of the 3510 confirmed to be OK. It should be noted that some issues were found, to be covered later, on the 3510 located with the rsfnode1 and may also exist on this unit.

There is one free qfe network port and this could be used for a resilient network failover option called NetMon. **HAC recommend** that NetMon be configured and enabled to allow the public (user) interface to failover in case of failure. **HAC also recommend** that the public interfaces are connected to different switches to further reduce single points of failure.



Reviewing the logs it was found that the primary network appeared to have an intermittent fault, causing a reported outage of approximately 40 seconds each time. There is no discernable pattern to this but it appears to have started in February and has continued through March. This may be a result of known changes to the external components or a fault. However, problems were also found when ftp'ing files from rsfnode2 (and to a lesser extent from rsfnode1) to 96.100.3.23. This may be related to the fact that ftp logging is enabled. **HAC recommend** that this potential problem be investigated. HAC's view is that the two problems are unrelated and that the ftp issues are related to logging in some way. Again NetMon could be used here not only to detect the fault and log it accurately but also to fail over if required.

There were problems reported, in the RSF-1 logs, for the disk heartbeat (*/opt/HAC/RSF-1/log/rsfmon.log*) during November and December 2004 but these are assumed to be related to the known fault, now resolved, on the 3510 arrays. This is simply noted and as it is understood the fault has been identified and resolved no further action is recommended, beyond confirming this assumption. However, **HAC recommend** that regular review of the log files is conducted to catch any potential faults of a similar nature in future.

The Ethernet MAC addresses are set to be local (*local-mac-address?=true*) via the eeprom, which is correct.

Access could not be gained to the */dev/ttya* port via the Annex terminal server, which may have worked correctly but the information relating to the correct method of access was not available. **HAC recommend** that access is confirmed and the details recorded in a Standard Operating Procedures (SOP) type document for the system.

In the eeprom, the boot disk is defined as a sub mirror of the root disk, which is correct. However, it *may* be possible to add the other sub mirror as an alternative boot device. This would be done in case the first sub mirror had failed. It is not a significant issue but should be considered. Alternatively, the alternate mirror could be defined as an alias in the nvram (eeprom) to aide booting in the event of a failure.

Security on the systems could be improved. There are a number of processes that are not thought to be used (e.g.: sendmail) which offer a security risk. ssh is not installed and root can login via clear text with telnet and ftp. **HAC recommend** that, as a minimum, direct root login be removed from all but the *'console'* and administrators be compelled to login with a standard use account and *'su'* when required. Administrators should be encouraged to perform only tasks essential to be run as root as such and most work be run as themselves. The *'console'* port should be redefined, in */etc/default/login*, as */dev/ttya*.

The available free space on all file systems appeared adequate. The Oracle parameters were not checked in detail, sat for possible memory starvation, but there were no apparent faults in this regard and the configuration of swap, physical memory and Oracle parameters set in */etc/system* looked appropriate. An experienced DBA should be consulted if further re-assurance be required in this regard.

The file systems managed by RSF-1 are listed correctly in the RSF-1 configuration (*/opt/HAC/RSF-1/etc/fs/vfstab.remedy*), however, these files systems are still listed in */etc/vfstab*. They are, correctly, not mounted at boot time but many customers would also comment out, using *'#'*, the appropriate lines in */etc/vfstab*. **HAC recommend** that this 'commenting out' be considered to reduce confusion, or comments added to the */etc/vfstab* file about this installation and configuration of RSF-1.

rpc.metamedd has 'hung up' and **HAC recommend** it is restarted. This is probably due to a minor bug in Sun OS code and the occurrence is fairly rare. A quick search of Google shows other users have had similar issues and a restart solves the problem. During the audit HAC offered to perform this task but Acme preferred to have it documented here and actioned at a later date.

RSF-1 Cluster Software

RSF-1 v2.6.15 is installed on the server and is generally correctly configured.

There appears to have been an upgrade of the AR System software, currently running at 5.01.02 patch 1346, or associated dynamically linked library, *libthread*, files. This is evidenced by a rather confusing set of scripts in the remedy service directory (*/opt/HAC/RSF-1/etc/rc.remedy.d/*) and **HAC recommend** that these be 'cleaned up', particularly *S60arsystem* and other variants of this script.

The way in which the AR System is started and the way it runs results in many apparently spurious lines being logged. There is no discernable value to the log entries (in */opt/HAC/RSF-1/log/rsfmon.log*) from this and **HAC recommend** that the scripts be modified to direct these messages to */dev/null*. It is possible that this change is undesirable for reasons not apparent to HAC but this should be reviewed.

There is no serial heartbeat configured, although it may be possible to use */dev/ttyb* for this purpose. **HAC recommend** that this be considered to improve resilience. The specification for RS232 allows for 15m of cabling between nodes, which should be adequate, although most modern equipment can drive over 25m.

The disk heart beat is configured to write via one 3510 array and read from another. **HAC recommend** that this be modified to have a pair of disk heartbeats, double the existing, with each read/write operation for each heartbeat being confined to one disk. This is simply an RSF-1 configuration issue and does not require geometry changes to the disks etc.

ORACLE®

RSF-1



rsfnod1

The hardware and software configuration of this system is almost identical to rsfnod2. Differences will be noted but most of the comments and recommendations for rsfnod2, above, are applicable to rsfnod1.

Operating System

The *hostid* is set to be identical, for AR System licensing reasons, believed to be agreed in writing with Remedy, to that of rsfnod2 but the Ethernet MAC addresses are different. This is as expected.

The clock of this system appears to be closely synchronised with that of rsfnod2, which is good.

The crontab entries differ in small ways to those on rsfnod2. **HAC recommend** that the differences be reviewed and the crontab entries be revised where appropriate so they are identical. The service related cron jobs could be placed on the shared disk to reduce effort in this regard.

The 3510 **was** accessible via the private network connection and the status was checked. There are a number of alarms showing, that the on-board backup battery may have a problem and may need replacing. **HAC recommend** that Sun be consulted on whether the battery should be replaced.

There are repeated reports of an 'IOM SES Firmware Revision Mismatch' which **HAC recommend** is resolved. It is understood this is already being done and there is an outstanding call with Sun support for resolution.

RSF-1 Cluster Software

The RSF-1 configuration is very similar to that of rsfnod2. However, the `.../etc/rc.remedy.d/` directory is in greater need of a 'tidy up' operation.

rsfnod1 does not start the 'agents' that monitor Oracle and AR System, although rsfnod2 does. The reason for this is unknown and **HAC recommend** the reason is understood and the systems be configured to operate in the same way.

Failover Testing

Failover, testing began at 19:50 after the user community and help desks had been consulted. The systems were both set to automatic and then failed over in both directions.

Failover testing was carried out in a safe way, by stopping the service on one node and waiting for the other node to take over. Power drop tests were not carried out as there is very limited additional value in performing these tests and the period of interruption may be significantly greater.

In addition, as the agents were enabled on rsfnod2, these were tested and successfully restarted AR System as required before returning RSF-1 to normal operation. While concern has been raised, above, about the various versions of S60arsystem, AR System starts without error.

All fail over tests were completed successfully.

On completion of the testing, rsfnod2 was returned to the RSF-1 manual administrative state, as this was how the system was set before the audit began. **HAC recommend** that all nodes, not undergoing maintenance of some kind, should normally be set to automatic but there may be specific reasons for the manual state on this site.

Upgrade Recommendations

The currently installed version of RSF-1 is v2.6.15, which dates back to May 2003. The current version of RSF-1 is v2.7.8 and there are a number of valuable updates that are available without charge to Acme. For a comprehensive list please check here. **HAC recommend** that RSF-1 v2.7.7 or later be installed and the MOUNT_POINT option be used.

As noted elsewhere in this document, the use of NetMon is also recommended to reduce exposure to single points of failure.

General Comments

The system appears to be installed and maintained to a reasonably high standard and while documentation, equipment labelling and general SOPs could perhaps do with some work, HAC are comfortable that if the actions recommended in this report are acted upon the system will continue to provide high levels of availability.

A full diagram, Figure 2., of the system as known to HAC is attached overleaf. Figure 3. shows the Sun recommended connections for a 3510. The current detailed connections have not been checked against this because the labelling was unclear. This should be checked. The diagram includes scope for future expansion.



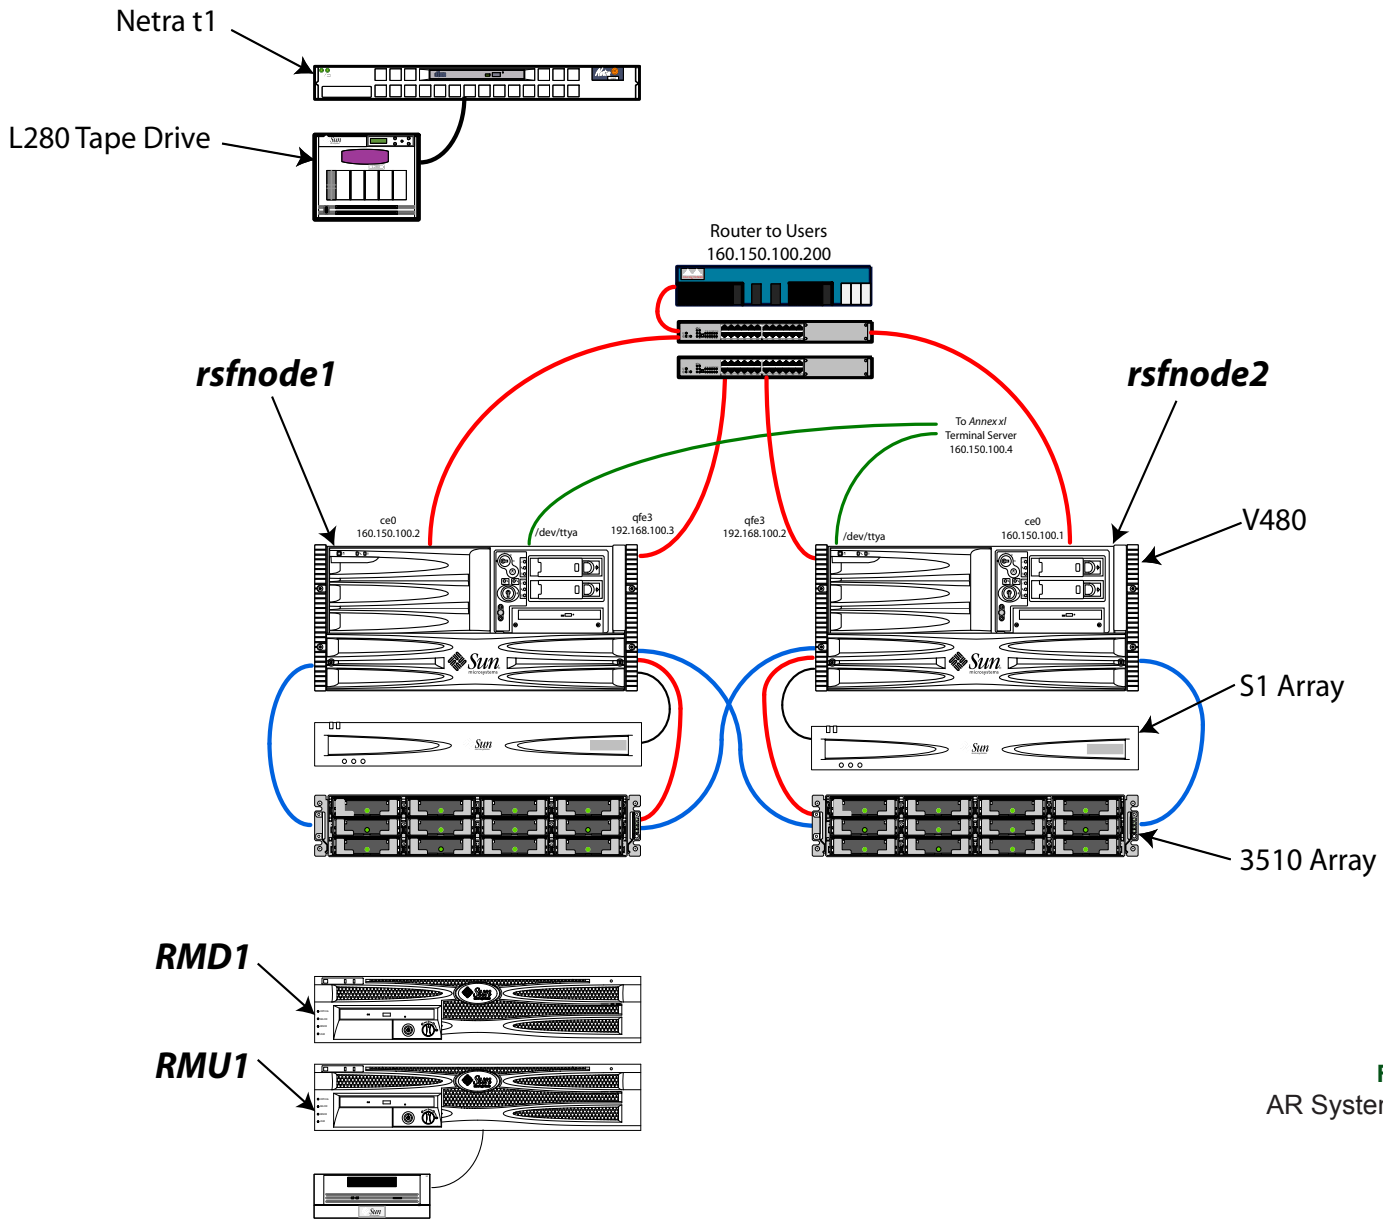


Figure 2.
AR System Setup

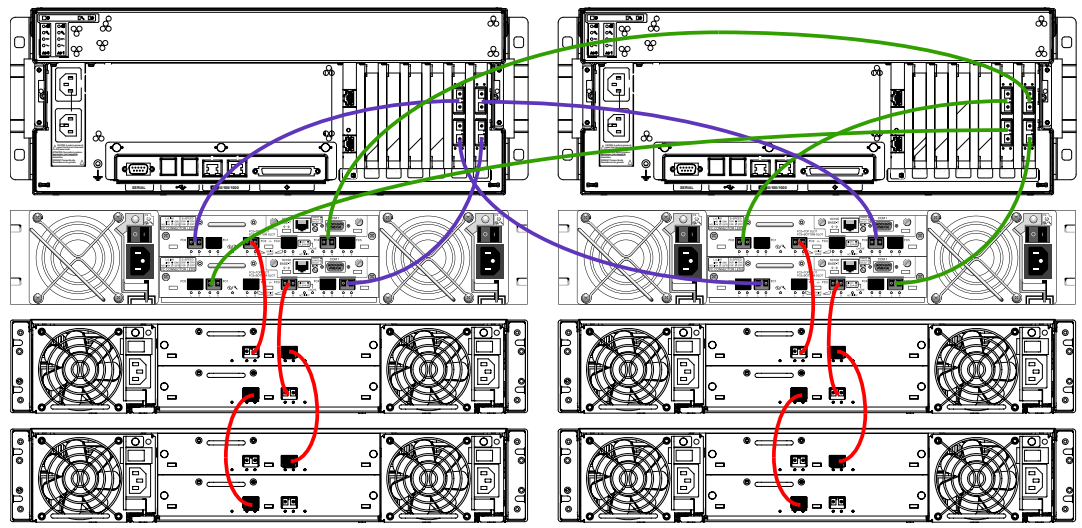


Figure 3.
Sun Recommended
Fibre Connections