

# 11g R1 Rac – Solaris – Installation Checklist

## Important Documents

[Clusterware installation guide for solaris operating system](#)  
[Real Application Clusters Installation Guide for Unix and Linux](#)  
[Release Notes for the solaris operating system](#)  
[Storage Administration Guide](#)

This checklist strives to document the minimum requirements to be met in order to accomplish a successful Oracle Rac Installation. This checklist can be used as a companion to the documents in the previous section. Please be aware that there could be requirements documented in the Official Oracle documentation that is not present in this document.

## Hardware requirements

### Procure Certified Hardware

Confirm with the hardware vendor (server and disk) that the combination of hardware you have procured, is certified to run with Oracle RAC (ie the vendor has tested the configuration and have successful production implementations). Components include but are not limited to

- Server make and model
- Type of HBA
- HBA driver to be used
- O/S kernel version to be used
- Multipathing driver to be used.

Make sure that you have redundant network interface cards to be used for the private nic and the public nic.

The private interconnect should be connected using a switch (Crossover cable not supported).

### [Technology compatibility matrix](#)

Review and familiarize yourself with the Rac technology compatibility matrix.

### [Memory Requirements](#)

Atleast 1gb of physical ram

### Solaris Containers

Oracle Real Application clusters is only certified to run in Global Containers (If you are using solaris containers) (Cannot run in local containers).

### Host bus adapters

Ensure that you have atleast 2 HBA's on each node (So that you have redundant access paths to the disk subsystem)

## Network Requirements

### Network interface requirements

Each node must have atleast 2 network

	<p>interfaces.</p> <p>We recommend having atleast 4 network interfaces. 2 for the private network and 2 for the public network.</p>
NIC Bonding	<p>If you have more than one network interface card for each type of network interfae (private and public), then bond a pair of interfaces together using IPMP.</p> <p>Refer to the following metalink notes for details on how to setup ipmp</p> <p><a href="#">368464.1</a></p> <p><a href="#">283107.1</a></p>
Interface Names	<p>The public interface names associated with the network adapters for each network must be the same on all nodes, and the private interface names associated with the network adaptors should be the same on all nodes.</p> <p>For example: With a two-node cluster, you cannot configure network adapters on node1 with eth0 as the public interface, but on node2 have eth1 as the public interface. Public interface names must be the same, so you must configure eth0 as public on both nodes. You should configure the private interfaces on the same network adapters as well. If eth1 is the private interface for node1, then eth1 should be the private interface for node2.</p>
Network Protocol (Public Nic)	<p>For the public network, each network adapter must support TCP/IP.</p>
Network Protocol (Private Nic)	<p>For the private network, the interconnect must support the user datagram protocol (UDP) using high-speed network adapters and switches that support TCP/IP (Gigabit Ethernet or better recommended).</p>
Public IP address	<p>Assign One IP address with an associated host name (or network name) registered in the DNS for the public interface. Record the host name and IP address in the system</p>

	<p>hosts file also, /etc/hosts.</p> <p>Each node needs a public IP address</p> <p>This can be an interface bonded using IPMP.</p>
Virtual IP address	<p>Assign One virtual IP (VIP) address with an associated host name registered in a DNS. Record the host name and VIP address in the system hosts file, /etc/hosts.</p> <p>Each node needs a virtual IP address.</p> <p>Select an address for your VIP that meets the following requirements:</p> <ul style="list-style-type: none"> <li>- The IP address and host name are currently unused (it can be registered in a DNS, but should not be accessible by a ping command)</li> <li>- The VIP is on the same subnet as your public interface</li> </ul>
Private IP address	<p>Configure the private interconnect network interface cards to have a private node name and a private IP address.</p> <p>This can be an interface bonded using IPMP.</p> <p>Choose a private IP address that is in the address range 10.*.* or 192.168.*.*</p> <p>The private IP address should be on a separate subnet than the public IP address.</p> <p>Oracle strongly recommends using a physically separate, private network (You can use a vlan on a shared switch if you must). You should ensure that the private IP addresses are reachable only by the cluster member nodes.</p>
<b>Operating System Requirements</b>	

Install Correct O/S version	Install Solaris 10, with the latest patch cluster.
Swap space	Equal to physical memory if physical memory < 8gb. .75 x physical memory if physical memory > 8gb
Temp space	Atleast 400mb of space in /tmp, preferably atleast 2gb
Network Time protocol	Configure network time protocol (ntp), so that the times on the rac hosts are synchronized with the same clock. Ensure that each member node of the cluster is set as closely as possible to the same date and time.
Operating system packages required	SUNWarc SUNWbtool SUNWhea SUNWlibC SUNWlibm SUNWlibms SUNWsprot SUNWtoo SUNWi1of SUNWi1cs SUNWi15cs SUNWxwfmt  Make sure that the package <b>SUNWxcu4</b> is also installed on the machine.
If you are using sun cluster	SUNWscucm SUNWudlmr SUNWudlm
If you are using Clustered solaris volume manager	SUNWscmd
If you are using sun cluster and hardware raid	SUNWschwr
If you use any of the following oracle featuers Pro*C/C++, Oracle Call Interface, Oracle C++ Call Interface, Oracle XML Developer's Kit (XDK)	Install Sun ONE Studio 11 (C and C++ 5.8)
If you use the oracle ODBC driver	Install gcc 3.4.2
Operating system Patches Required	127111-02 SunOS 5.10: libc patch
<a href="#">Configure kernel parameters</a>	Configure the following kernel parameters using both /etc/system and resource controls.  Use the instructions in the section "Configuring kernel parameters on Solaris 10"

	<p>from the link  <a href="http://download.oracle.com/docs/cd/B28359_01/install.111/b28262/presolar.htm#BABCHAEED">http://download.oracle.com/docs/cd/B28359_01/install.111/b28262/presolar.htm#BABCHAEED</a></p> <pre>noexec_user_stack      1 semsys:seminfo_semmni 100 semsys:seminfo_semmns 1024 semsys:seminfo_semmsl 256 semsys:seminfo_semvmx 32767 shmsys:shminfo_shmmax 4294967295 shmsys:shminfo_shmmni 100</pre>
<a href="#">Set UDP parameters</a>	<pre>udp_xmit_hiwat=65536 udp_rcv_hiwat=65536</pre>
<a href="#">Configure SSH on all cluster nodes</a>	<p>Ssh should be configured on all nodes of the rac cluster, for the oracle user, and you should be able to ssh to every node without being prompted for the password.</p> <p>Make sure that there is no banner displayed when you login (or ssh) as oracle.</p> <p>Softlink ssh and scp to /usr/bin from /usr/local/bin</p>
Add all relevant node names and ip addresses in the /etc/hosts	<p>Check /etc/hosts to verify the following entries exist</p> <pre>127.0.0.1 localhost localhost.localdomain</pre> <p>The /etc/hosts should have the following entries for each database node</p> <ul style="list-style-type: none"> <li>– Your public node name, public node name.domainname</li> <li>– Your private node name,private node name.domainname</li> <li>– Your vip node name, vip node name.domainname.</li> </ul>
Create the dba group	Create the group named dba with the same gid on each server
Create Oracle Software Owner	Create the Unix user oracle, on each node and assign it to the group dba. Setup bash or ksh to be the default login shell
Set shell limits for the oracle user	<pre>rlim_fd_max    65536 maxuprc       16384</pre>

Set umask	Set the default umask for the oracle user to be 022
Allocate a local disk mount point for the CRS_HOME	<p>This mount point should be different from that allocated for ORACLE_HOME and ASM_HOME (eg: /crs).</p> <p>Allocate atleast 1gb of space for this mount point.</p>
Allocate a local disk mount point for the ASM_HOME	<p>This mount point should be different from that allocated for ORACLE_HOME and CRS_HOME (eg: /u01/asm). Allocate atleast 3gb of space for this mount point.</p>
Allocate a local disk mount point for the ORACLE_HOME	<p>This mount point should be different from that allocated for the CRS_HOME (eg: /u01/app). Allocate atleast 3Gb of space for this directory.</p>
Allocate three sharable raw devices for the voting disk.	<p>This device should preferably mirrored at the hardware level. Size 280Mb. Configure this device by creating a partition on the lun, that skips the first 1Mb on the LUN. For best availability, keep these three luns on physically separate disks.</p> <p>The devices can be emc powerpath devices.</p> <p>The device names that point to the same lun, should be exactly the same on both nodes.</p> <p>Eg: if /dev/rds/xxdisk1 points to lun1 on node 1 then /dev/rds/xxdisk1 has to point to lun1 on the second node.</p> <p>After creating the devices run the following command to make absolutely sure that you have 280Mb of free space.</p> <pre>dd if=/dev/zero of=/dev/raw/raw1 bs=+1M count=280 (Substitute raw1 with your rawdevice name)</pre> <p>Change the ownership and permissions of the disks</p> <pre>chown root:dba /dev/rds/c3t4d12s6 chmod 660 /dev/rds/c3t4d12s6</pre>

Allocate two sharable raw devices for the cluster registry.

This device should preferably mirrored at the hardware level. Size 280 Mb. Configure this device by creating a partition on the lun, that skips the first 1Mb on the LUN. Device should preferably mirrored at the hardware level.

The devices can be emc powerpath devices.

Eg: if /dev/rdisk/xxdisk1 points to lun1 on node 1 then /dev/rdisk/xxdisk1 has to point to lun1 on the second node.

After creating the devices run the following command to make absolutely sure that you have 100Mb of free space.

```
dd if=/dev/zero of=/dev/raw/raw1 bs=+1M count=280 (Substitute raw1 with your rawdevice name)
```

Change the ownership and permissions on the devices

```
chown oracle:dba /dev/rdisk/c3t4d15s6  
chmod 660 /dev/rdisk/c3t4d15s6
```

Allocate devices for the ASM Data disk group

Take the total amount of space you require to place your database files (data, index, system, undo, temp, redo) and divide this up to create as many luns as possible. Rule of thumb is to present at least 4 luns and have the number of luns as close to the number of spindles you are allocating.

Present these luns to the operating system.

Eg: if you need to present 500gb of space then create and present 10 luns of 50gb each.

These luns can be emc powerpath devices.

Eg: if /dev/rdisk/xxdisk1 points to lun1 on node 1 then /dev/rdisk/xxdisk1 has to point to lun1 on the second node.

	<p>After the powerpath devices are configured, use fdisk to create a partition of the disk, skipping the first 1mb of the disk. This new partition is what ASM will use.</p> <p>Change the ownership and permissions on the devices</p> <pre>chown oracle:dba /dev/rdisk/c3t4d15s6 chmod 660 /dev/rdisk/c3t4d15s6</pre>
<p>Create the devices for the flash recovery area disk group</p>	<p>Take the total amount of space you require to place your Flash recovery area and divide this up to create as many luns as possible.</p> <p>Rule of thumb is to present at least 4 luns and have the number of luns as close to the number of spindles you are allocating.</p> <p>Present these luns to the operating system.</p> <p>Eg: if you need to present 500gb of space then create and present 10 luns of 50gb each.</p> <p>These luns can be emc powerpath devices.</p> <p>Eg: if /dev/rdisk/xxdisk1 points to lun1 on node 1 then /dev/rdisk/xxdisk1 has to point to lun1 on the second node.</p> <p>After the powerpath devices are configured, use fdisk to create a partition of the disk, skipping the first 1mb of the disk. This new partition is what ASM will use.</p> <p>Change the ownership and permissions on the devices</p> <pre>chown oracle:dba /dev/rdisk/c3t4d15s6 chmod 660 /dev/rdisk/c3t4d15s6</pre>
<p>Download latest cluster verification utility (cvu)</p>	<p><a href="#">Cluster Verification Utility – Solaris</a></p>
<p>Download Oracle Software</p>	<p>Download <a href="#">solaris.sparc64_11gR1_database_1013.zip</a>  <a href="#">solaris.sparc64_11gR1_clusterware.zip</a></p>
<p>Download Oracle Patchset</p>	<p>Download 11.1.0.7, patch number 6890831 from metalink</p>

<p><a href="#"><u>Run Cluster Verification Utility and send me the output.</u></a></p>	<p>As the oracle user ./runcluvfy.sh stage -pre crsinst -n node1,node2 ./runcluvfy.sh comp nodecon -n node1,node2 ./runcluvfy.sh stage -post hwos -n node1,node2 ./runcluvfy.sh comp sys -n node1,node2 -p crs -osdba dba -orainv dba</p> <p>Replace the node names with your actual public node names.</p>
<p>Xwindows client</p>	<p>Make sure that, on the desktop from which you are going to login to the servers and perform the installation, that you have a good xwindows client like exceed. (Open source clients like xming and/or xmanager will not work). If you are using reflections, then make sure that you have the latest version of reflections.</p>