



Table of Contents

Introduction.....	2
Installing Solaris Cluster manually.....	3
Using Auto Install with Solaris Cluster 4.....	4
Installing Solaris Cluster 4 using JetSC4.....	5
Solution Benefits.....	10
Summary.....	10
Appendix 1 - Quorum Servers.....	11



Introduction

In large organisations with many servers and Solaris Cluster configurations, it is important that the installation process is repeatable and human errors are minimised. Installing and configuring the same software by hand can lead to variances, which in some cases, require restarting the entire installation process or at worst, rebuilding the OS. To help reduce the requirement for manual installations, many organisations chose to utilise the Sun Microsystems (now owned by Oracle) JumpStart Enterprise Toolkit (JET) framework to automate the installation process. This toolkit, along with the JetSC3 module, allowed automated Cluster installations to be performed, providing consistent deployment techniques.

With the introduction of Solaris 11, Oracle changed the operating system installation process from being JumpStart based, to a new Auto Installer tool. While the overall end result of delivering an OS over the network is the same, the process of setting up the installation image has not; there are new commands and a reliance on XML files to facilitate the install. Combined with this, the Solaris Cluster 4 software uses a similar approach for installation and many of the guides assume an Automated Installer environment is available. This move to Automated Installer means organisations need to adopt a whole new approach to operating system installs, that requires retraining and implementation of a new methodology. As such, the increase of investment time in setting up such an environment has meant that the installation and configuration of Solaris Cluster has once more become a manual task.



Installing Solaris Cluster manually

Installing Solaris Cluster requires that all nodes have the same operating system level, patches and ideally hardware configuration. The exact hardware requirements for Solaris Cluster are available from Oracle; this paper assumes that a suitable set of hardware is available.

The process normally begins with the operating system installation; it's generally considered best practice to install the cluster framework onto 'clean' systems, before the applications are loaded.

Once the OS has been installed, any IPMP configuration for network redundancy is added. At this point, the Solaris Cluster software can be added and configured. Once installed, further SRUs (the new name for patches) may need applying.

Installing the Cluster framework is achieved through the use of the new Solaris 11 packaging system 'IPS'. The new packaging system needs to be told where to find the cluster software 'repository'.

If the clients are internet facing then they can utilise the Oracle Cluster repository, <https://pkg.oracle.com/ha-cluster/release/ha-cluster>, and the patch SRU, <https://pkg.oracle.com/ha-cluster/support/ha-cluster>

Access to the online Oracle repositories require an appropriate service contract; through which, you can generate the required SSL keys and certificates.

However on the majority of sites, servers do not normally have direct connection to the internet due to firewall and security policies. In these cases, the online Oracle repository can be copied to a local location, and the client machine can then be directed to the local http[s] repository or filesystem containing the software.

Framework Installation

Once the correct repo(s) are accessible the cluster framework can be installed:

```
# pkg install ha-cluster-framework-full
```

Any Cluster SRU should be applied at this point. SRUs can be accessed direct from the internet or from a local repository in the same manner as the main cluster repository.

```
# pkg update
```



Framework Configuration

Once the framework is installed, the software can be configured by running the command `/usr/cluster/bin/scinstall` on each node. Each node requires a series of questions answered correctly before they are rebooted and able to form a cluster.

Traditionally with Solaris Cluster 3 any errors at this point were often easier to rectify with a node rebuild. With Solaris Cluster 4 the package removal technology is more mature and should not necessitate this.

Once the framework has been installed and the nodes rebooted, the cluster should form, allowing the further configuration of applications and data services.

Using Auto Install with Solaris Cluster 4

Oracle have published a guide to incorporating a Sun Cluster 4 installation into an Automated Installer deployment at <http://www.oracle.com/technetwork/articles/servers-storage-admin/howto-automate-install-cluster4-1841534.html> . By using this guide, it is possible to have Solaris Cluster 4 installed during the initial OS install phase.

By amending the XML generated for an AI installation of a node the Solaris Cluster software can be installed at build time from a local repository.

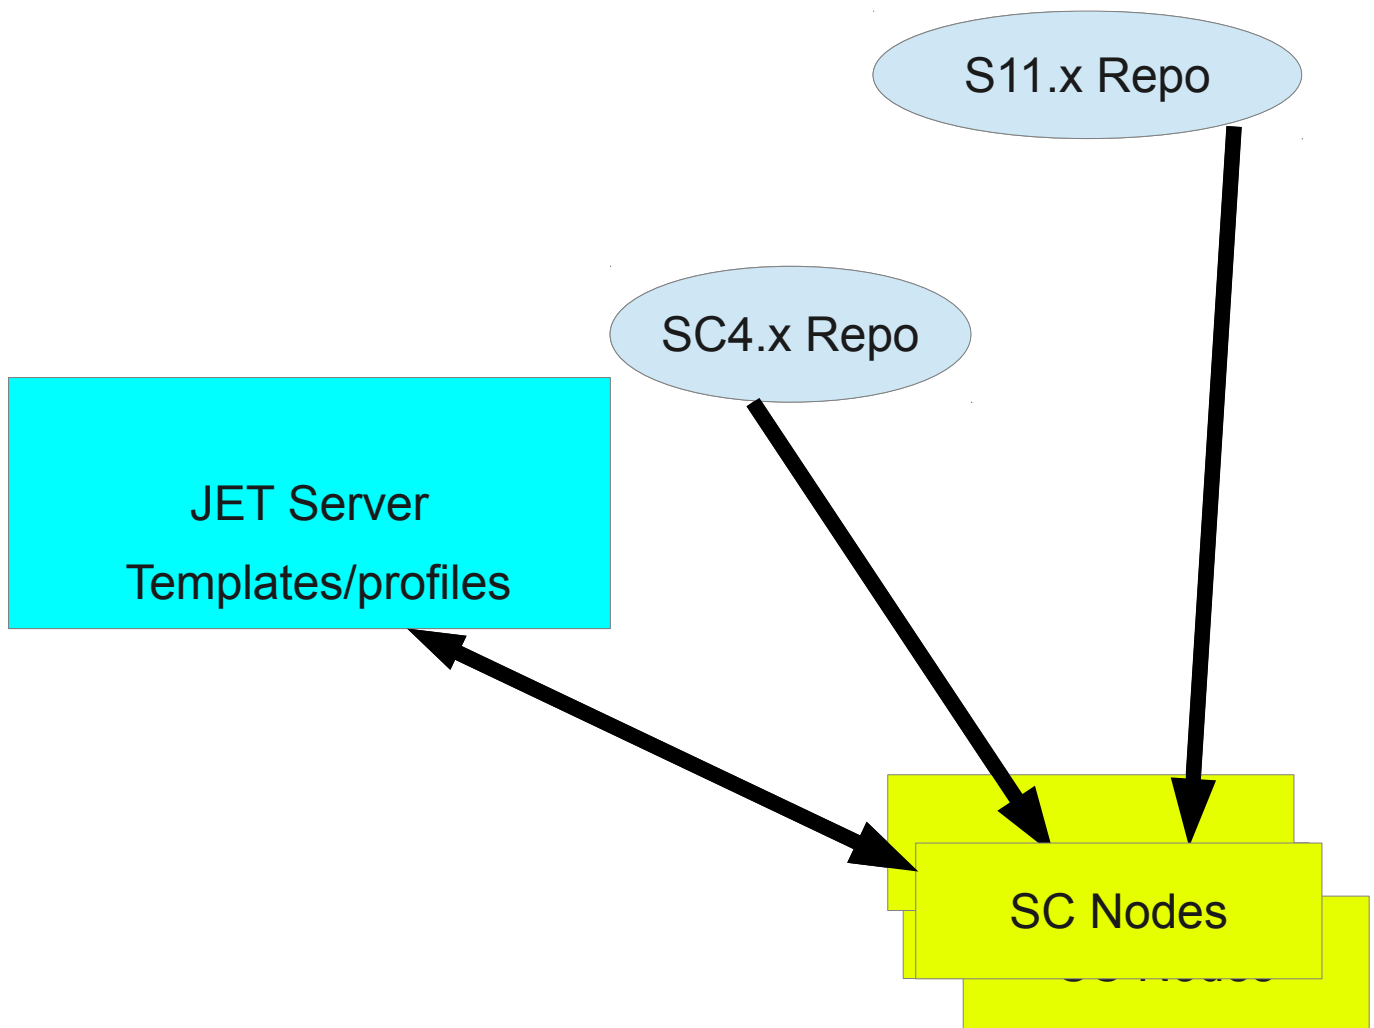
The Automated Installer approach provides a mechanism to install a base operating system and cluster framework, but still requires many manual steps to set up the build. The build is a standard Solaris install, and does not offer any easy ways to customise the system installation as would normally be required for corporate standard installs.



Installing Solaris Cluster 4 using JetSC4

Sun Microsystems developed, and subsequently released, the JumpStart Enterprise Toolkit (JET) in order to simplify the process for automated deployments of Solaris. Although Oracle have not released any support for the Automated Installer with JET, the functionality has been developed by the 3rd party Maui Systems Ltd. With JET now capable of deploying Solaris through AI, it was an obvious step to create a JetSC4 module.

- The multiple boot and node synchronising approach has been retained.
- The JET server co-ordinates and facilitates communication between the various components required, DHCP server, Repo(s) and nodes.
- The DHCP server can be remote or local to the JET server.
- Ability to specify local or remote repositories and any SRUs, including SSL key management.





Creating a cluster build (overview)

There are four steps required

- Create the templates
- Edit the template to populate all the variables required for your configuration.
- Run the `make_client` script to generate the AI XML code and JET configuration.
- Initiate the build on the nodes

Creating the JET template for Node 1

To create a node template run the JET utility 'make_ai_template'

```
# make_ai_template cl4node1 sc4
```

Edit the initial node template

Using your favourite editor, open the `/opt/SUNWjet/Templates/<nodename>` file, and populate with the cluster configuration information.

It is assumed at this point that the general Solaris configuration portion of the JET template is known and understood; this document concentrates on the SC4 implementation and proceeds from that section of the template. For further reading on JET and AI the white paper is listed in the references section.

The first item to populate is the location of the cluster software repository. If the software is to be installed from the Oracle Cluster repository then the appropriate SSL key and certificate files will need to be obtained from Oracle, and copied to the `/opt/SUNWjet/Clients/<clientname>` directory.

```
sc4_repo_key="Oracle_Solaris_Cluster_4.key.pem"  
sc4_repo_cert="Oracle_Solaris_Cluster_4.certificate.pem"  
sc4_repository="https://pkg.oracle.com/ha-cluster/release"
```



Alternatively, if the repository is local, you can leave the key & cert entries blank, and just populate the repository variable.

```
sc4_repo_key=""  
sc4_repo_cert=""  
sc4_repository="http://10.0.0.1:8081"
```

Select the Cluster package group to be installed; this can be one of :-

- ha-cluster-full
- ha-cluster-framework-full
- ha-cluster-data-services-full
- ha-cluster-minimal
- ha-cluster-framework-minimal

```
sc4_package="ha-cluster-framework-full"
```

If a software only install is required with no cluster configuration then no more variables are required.

Cluster configuration details

To configure a cluster simply fill in the relevant fields, if any are required fields then errors will be highlighted when *make_client* is run.

The maxnodes, maxprivatenets and netmask have been set to the Solaris Cluster defaults, if any of these are left blank the default will be used.



```
sc4_cluster_name="YTC_Cluster"  
sc4_auth_type="sys"  
sc4_private_net="172.16.0.0"  
sc4_maxnodes="64"  
sc4_maxprivatenets="10"  
sc4_netmask="255.255.240.0"  
sc4_node_names="sc4node1 sc4node2"  
sc4_node_interconnects="net0 net1"  
sc4_node_addresses="10.2.0.20 10.2.0.23"  
sc4_switch_names="switch1 switch2"
```

Notes:

- The 'node_names' and node_addresses should be entered in the same order;
- There should be the same number of private interconnect switch names as nodes.

Quorum Devices

A quorum device is required for 2 node clusters (and generally for clusters with an even number of nodes). This can either be a shared disk (capable of SCSI-3 reservations) or a 'Quorum Server'.

If the quorum_disks variable is left as "auto" then the Solaris Cluster install script will attempt to determine which disk to use, if none are available the install will fail !

When using a 'Quorum Server', this should be set up and available prior to building the cluster nodes; the JET check_client script will check if it is available.

A guide to installing a 'Quorum Server' can be found in the appendix.

```
sc4_quorum_disks="auto"  
sc4_quorum_server="YTCClusterQ"  
sc4_quorum_ip="10.2.0.99"  
sc4_quorum_port="9005"
```




Once the template for node1 has been completed, generate subsequent node templates using the `make_ai_template` command, but this time, specify a source template to take initial settings from:

```
# make_ai_template -T sc4node1 sc4node2
```

Unless you have pre-populated `/etc/hosts` and `/etc/ethers` with appropriate details, edit the new template and set the IP address and MAC address details accordingly.

Setting up the build environment for the nodes

Once you have created the templates, you can quickly set up the AI build environment by using the JET 'make_client' command.

```
# make_client sc4node1
```

If any errors are reported, edit the template and correct any errors; rerun 'make_client' as required.

Repeat the 'make_client' command for each cluster node template created.

```
# make_client sc4node2
```

Your cluster is now ready to be built

Note: If you have opted to use a quorum server, the make_client script will warn you if it cannot be contacted. You can proceed with the installation but the Solaris Cluster nodes will wait during the build, periodically checking for quorum server to be reachable.

Initiate the build on the nodes

On all nodes, initiate the appropriate PXE/DHCP boot commands to invoke an install.



Solution Benefits

By using JetAI and JetSC4, any investment already made in the JET infrastructure, knowledge and training can still be utilised. In organisations with large numbers of departments and servers that require constant, repetitive roll outs the transition to Solaris 11 and Solaris Cluster 4 becomes less time consuming. Builds can conform to a standard, company template; removing the per-installer variances added by individuals.

The same Solaris 11 based JET server can be used to build both Solaris 10 (including Solaris Cluster 3 support) clients and also Solaris 11 based installs, including Solaris Cluster 4.

Summary

The solution shows that even though the operating system delivery method has been changed by Oracle, any investment in using JET as an installation platform can be retained and expanded upon.

Setting up Solaris Cluster manually is prone to human typographic errors which can be time consuming to isolate and rectify.

Benefits from using JetAI and JetSC4 :-

- No need to manually edit XML files for AI & SC4; all handled by JetAI and JetSC4
- Utilising JetAI to deploy Solaris 11, files, scripts and customisations can be applied quickly, easily and consistently
- Multiple installations can be performed by one engineer, simultaneously
- Templates can be prepared, and tested, prior to installation
- Existing investment in JumpStart and JET infrastructure is retained



Appendix 1 - Quorum Servers

Using a quorum server releases the requirement from allocation SCSI3 storage to a Solaris Cluster and as the quorum hardware does not need to be on the Solaris Cluster 4.x HW support list it increases the number of available options.

A quorum server can be configured to act as quorum for multiple clusters and multiple cluster versions, so Solaris 10 Solaris Cluster 3.x quorum server can have a vote for a Solaris Cluster 4 configuration.

A JET quorum server module will probably be created in the future.

Configuring the quorum server

Depending on the Solaris version the software installation differs.

Installing the quorum server software on Solaris 10

Using the JES installer, select *Quorum Server* and then *Configure Later*

Installing the quorum server software on Solaris 11

Select the appropriate repository and install.

```
pkg install ha-cluster-quorum-server-full
```

Edit the quorum server configuration file

The same configuration file, `/etc/scqsd/scqsd.conf`, is used for all Solaris Clusters the server is utilised on.

First create an new, unique to the Solaris Cluster, directory for the quorum server files.

```
# mkdir /var/scqsd/YTCClusterQ
```

Edit the file `/etc/scqsd/scqsd.conf` and add one entry per Solaris Cluster

```
/usr/cluster/lib/sc/scqsd -i YTCClusterQ -p 9005 -d /var/scqsd/YTCClusterQ
```

Where :-

-i <Quorum Server>, the unique name for the quorum server, must match the template entry **sc4_quorum_server**

-p <port>, the port that will be listened to, this must match the template entry **sc4_quorum_port**

-d <quorum directory>, the unique directory for each quorum server configured



Start the quorum server

Once added the server can be started.

```
# clqs start YTCClusterQ
```

Quorum server tasks on Solaris Cluster nodes.

The Solaris Cluster nodes require the hostname and IP address of the quorum server to be set in the `/etc/inet/hosts` file. The JetSC4 module will add this during the post install phase.



References

“Oracle Solaris Cluster Software Installation Guide”

http://docs.oracle.com/cd/E29086_01/html/E29085/index.html

“Using JET to Deploy Solaris 11”, Marty Lee, Maui Systems Ltd.

<http://www.maui-systems.co.uk/white-papers/deploying-solaris-11-using-jet>

“How to Automate the Installation and Setup of Oracle Solaris Cluster 4.0”

<http://www.oracle.com/technetwork/articles/servers-storage-admin/howto-automate-install-cluster4-1841534.html>



About the Author

Bruce Porter is the Lead Technical Consultant of YTC Systems Ltd and has worked with JumpStart since its introduction in Solaris 2. While working for Sun Microsystems (Professional Services) he became part of the JET (JumpStart Enterprise Toolkit) core team working on many JET modules. He is currently still supplying Professional Services to many large organisations, including Oracle.



Disclaimers

Copyright © 2014, YTC Systems Ltd. All rights reserved. The information contained in this document is the proprietary and exclusive property of YTC Systems Ltd except as otherwise indicated. No part of this document, in whole or in part, may be reproduced, stored, transmitted, or used for design purposes without our prior written permission. The information contained in this document is subject to change without notice and is provided for informational purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document.

Oracle, Sun, Solaris, Solaris Cluster and JumpStart are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners

Installing Solaris Cluster 4.x
using JET

Author: Bruce D Porter

January 2014

YTC Systems Limited

30 Syston Avenue

St Helens

Merseyside

WA11 9JJ

<http://www.ytc-systems.co.uk>