

**Oracle® Solaris Cluster Data Service for
Oracle Traffic Director Guide**

ORACLE®

Part No: E35577
July 2014, E35577-01

Copyright © 2009, 2014, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Copyright © 2009, 2014, Oracle et/ou ses affiliés. Tous droits réservés.

Ce logiciel et la documentation qui l'accompagne sont protégés par les lois sur la propriété intellectuelle. Ils sont concédés sous licence et soumis à des restrictions d'utilisation et de divulgation. Sauf disposition de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, diffuser, modifier, breveter, transmettre, distribuer, exposer, exécuter, publier ou afficher le logiciel, même partiellement, sous quelque forme et par quelque procédé que ce soit. Par ailleurs, il est interdit de procéder à toute ingénierie inverse du logiciel, de le désassembler ou de le décompiler, excepté à des fins d'interopérabilité avec des logiciels tiers ou tel que prescrit par la loi.

Les informations fournies dans ce document sont susceptibles de modification sans préavis. Par ailleurs, Oracle Corporation ne garantit pas qu'elles soient exemptes d'erreurs et vous invite, le cas échéant, à lui en faire part par écrit.

Si ce logiciel, ou la documentation qui l'accompagne, est concédé sous licence au Gouvernement des Etats-Unis, ou à toute entité qui délivre la licence de ce logiciel ou l'utilise pour le compte du Gouvernement des Etats-Unis, la notice suivante s'applique:

U.S. GOVERNMENT END USERS. Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

Ce logiciel ou matériel a été développé pour un usage général dans le cadre d'applications de gestion des informations. Ce logiciel ou matériel n'est pas conçu ni n'est destiné à être utilisé dans des applications à risque, notamment dans des applications pouvant causer des dommages corporels. Si vous utilisez ce logiciel ou matériel dans le cadre d'applications dangereuses, il est de votre responsabilité de prendre toutes les mesures de secours, de sauvegarde, de redondance et autres mesures nécessaires à son utilisation dans des conditions optimales de sécurité. Oracle Corporation et ses affiliés déclinent toute responsabilité quant aux dommages causés par l'utilisation de ce logiciel ou matériel pour ce type d'applications.

Oracle et Java sont des marques déposées d'Oracle Corporation et/ou de ses affiliés. Tout autre nom mentionné peut correspondre à des marques appartenant à d'autres propriétaires qu'Oracle.

Intel et Intel Xeon sont des marques ou des marques déposées d'Intel Corporation. Toutes les marques SPARC sont utilisées sous licence et sont des marques ou des marques déposées de SPARC International, Inc. AMD, Opteron, le logo AMD et le logo AMD Opteron sont des marques ou des marques déposées d'Advanced Micro Devices. UNIX est une marque déposée d'The Open Group.

Ce logiciel ou matériel et la documentation qui l'accompagne peuvent fournir des informations ou des liens donnant accès à des contenus, des produits et des services émanant de tiers. Oracle Corporation et ses affiliés déclinent toute responsabilité ou garantie expresse quant aux contenus, produits ou services émanant de tiers. En aucun cas, Oracle Corporation et ses affiliés ne sauraient être tenus pour responsables des pertes subies, des coûts occasionnés ou des dommages causés par l'accès à des contenus, produits ou services tiers, ou à leur utilisation.

Contents

- Using This Documentation** 7

- 1 Installing and Configuring HA for Oracle Traffic Director** 9
 - HA for Oracle Traffic Director Overview 9
 - Overview of the Installation and Configuration Process for HA for Oracle Traffic Director 10
 - Planning the HA for Oracle Traffic Director Installation and Configuration 10
 - Installing and Configuring the Oracle Traffic Director Software 11
 - Installing the HA for Oracle Traffic Director Package 11
 - ▼ How to Install the HA for Oracle Traffic Director Package 12
 - Registering and Configuring HA for Oracle Traffic Director 13
 - Setting HA for Oracle Traffic Director Extension Properties 13
 - Setting IP Addresses and Port Numbers for the Oracle Traffic Director Application 14
 - ▼ How to Register and Configure HA for Oracle Traffic Director With a Failover Logical Host 14
 - ▼ How to Register and Configure HA for Oracle Traffic Director With a Scalable Shared Address 20
 - Tuning the HA for Oracle Traffic Director Fault Monitor 24
 - Operations by the Fault Monitor During a Probe 25

- A HA for Oracle Traffic Director Extension Properties** 29

- Index** 33

Using This Documentation

- **Overview** – Describes how to install and configure Oracle Solaris Cluster data service for Oracle Traffic Director.
- **Audience** – Experienced system administrators with extensive knowledge of Oracle software and hardware.
- **Required knowledge** – Knowledge of the Oracle Solaris operating system, of Oracle Solaris Cluster software, and expertise with the volume manager software that is used with Oracle Solaris Cluster software.

This document is not to be used as a planning or presales guide.

Product Documentation Library

Late-breaking information and known issues for this product are included in the documentation library at <http://www.oracle.com/pls/topic/lookup?ctx=E39579>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program web site at <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Feedback

Provide feedback about this documentation at <http://www.oracle.com/goto/docfeedback>.

Installing and Configuring HA for Oracle Traffic Director

This chapter provides the procedures to install and configure Oracle Solaris Cluster HA for Oracle Traffic Director (HA for Oracle Traffic Director).

Note - This data service is supported only on Oracle's engineered systems. For more information about Oracle's engineered systems, see [Oracle Engineered Systems](#). For updated information about supported configurations of this data service, see the [Oracle Solaris Cluster 4 Compatibility Guide](#).

Note - You can install and configure this data service to run in either the global zone or a zone cluster.

This chapter contains the following sections.

- [“Overview of the Installation and Configuration Process for HA for Oracle Traffic Director” on page 10](#)
- [“Planning the HA for Oracle Traffic Director Installation and Configuration” on page 10](#)
- [“Installing and Configuring the Oracle Traffic Director Software” on page 11](#)
- [“Installing the HA for Oracle Traffic Director Package” on page 11](#)
- [“Registering and Configuring HA for Oracle Traffic Director” on page 13](#)
- [“Tuning the HA for Oracle Traffic Director Fault Monitor” on page 24](#)

HA for Oracle Traffic Director Overview

HA for Oracle Traffic Director is a multi-master data service and can be configured with a failover logical hostname or a scalable shared address. See [Chapter 1, “Planning for Oracle Solaris Cluster Data Services,”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) and the [“Oracle Solaris Cluster Concepts Guide”](#) for general information about data services, resource groups, resources, and other related topics.

Note - You can use Oracle Solaris Cluster Manager to install and configure this data service. See the Oracle Solaris Cluster Manager online help for details.

Note - If you run multiple data services in your Oracle Solaris Cluster configuration, you can set up the data services in any order, with the following exception. If HA for Oracle Traffic Director depends on Oracle Solaris Cluster HA for DNS, you must set up DNS first. See [“Oracle Solaris Cluster Data Service for DNS Guide”](#) for details. The Oracle Solaris OS includes the DNS software. If the cluster is to obtain the DNS service from another server, then configure the cluster to be a DNS client first.

Note - After installation, do not manually start and stop Oracle Traffic Director except by using the `clresourcegroup` command. See the [clresourcegroup\(1CL\)](#) man page for details. After Oracle Traffic Director is started, the Oracle Solaris Cluster software controls it.

Overview of the Installation and Configuration Process for HA for Oracle Traffic Director

The following table lists the sections that describe the installation and configuration tasks.

TABLE 1-1 Task Map: Installing and Configuring HA for Oracle Traffic Director

Task	For Instructions
Install Oracle Traffic Director.	“Installing and Configuring the Oracle Traffic Director Software” on page 11
Install the HA for Oracle Traffic Director package.	“How to Install the HA for Oracle Traffic Director Package” on page 12
Register HA for Oracle Traffic Director and configure the cluster for the data service.	“Registering and Configuring HA for Oracle Traffic Director” on page 13
Tune the HA for Oracle Traffic Director fault monitor.	“Tuning the HA for Oracle Traffic Director Fault Monitor” on page 24

Planning the HA for Oracle Traffic Director Installation and Configuration

Answer the following questions before you start your installation.

- Will you run HA for Oracle Traffic Director with a logical hostname or a shared address? See the [“Oracle Solaris Cluster Concepts Guide”](#) document for information on the two types of network resources.

- Where will the installation and configuration directories and files reside?
- Do other Oracle Solaris Cluster resources have dependencies on the services that Oracle Traffic Director provides or does Oracle Traffic Director depend on any services provided by other Oracle Solaris Cluster resources? If so, resource dependencies might exist between the resources so that one starts or stops before the other. See the [r_properties\(5\)](#) man page for a description of the resource property `Resource_dependencies` that sets these dependencies.
- When using a logical hostname, HA-Oracle Traffic Director must be configured to bind to `INADDR_ANY` since a logical hostname can be present on only one node at a time.
- Determine the resource groups to use for network addresses and application resources and the dependencies between them. See the [rg_properties\(5\)](#) man page for a description of the resource group property `RG_dependencies` that sets these dependencies.
- Provide the logical hostname (for failover services) or shared address (for scalable services) for clients to use to access the data service.
- Because you can configure Oracle Traffic Director to bind to `INADDR_ANY`, if you plan to run multiple instances of the Oracle Traffic Director data service or multiple data services on the same node, each instance must bind to a unique network address and port number.
- Determine where to place logs and error files on the local file system.
- Determine where to place the contents on the cluster file system.

Installing and Configuring the Oracle Traffic Director Software

Before you can register HA for Oracle Traffic Director and configure the cluster for the data service, you must install the Oracle Traffic Director software. See the Oracle Traffic Director [documentation](#) for information on how to install the software.

Note - This data service is supported only on Oracle's engineered systems. For more information about Oracle's engineered systems, see [Oracle Engineered Systems](#). For updated information about supported configurations of this data service, see the [Oracle Solaris Cluster 4 Compatibility Guide](#).

Installing the HA for Oracle Traffic Director Package

If you did not install the HA for Oracle Traffic Director package during your initial Oracle Solaris Cluster installation, perform this procedure to install the package.

▼ How to Install the HA for Oracle Traffic Director Package

Perform this procedure on each cluster node where you want the HA for Oracle Traffic Director software to run.

1. **On the cluster node where you are installing the data service package, assume the root role.**
2. **Ensure that the data service package is available from the configured publisher and that the `solaris` and `ha-cluster` publishers are valid.**

```
# pkg list -a ha-cluster/data-service/oracle-traffic-director
# pkg publisher
PUBLISHER                TYPE      STATUS  P  LOCATION
solaris                   origin   online  F  solaris-repository
ha-cluster                 origin   online  F  ha-cluster-repository
```

For information about setting the `solaris` publisher, see [“Adding, Modifying, or Removing Package Publishers”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

Tip - Use the `-nv` options whenever you install or update to see what changes will be made, such as which versions of which packages will be installed or updated and whether a new BE will be created.

If you do not get any error messages when you use the `-nv` options, run the command again without the `-n` option to actually perform the installation or update. If you do get error messages, run the command again with more `-v` options (for example, `-nvv`) or more of the package FMRI pattern to get more information to help you diagnose and fix the problem. For troubleshooting information, see [Appendix A, “Troubleshooting Package Installation and Update,”](#) in [“Adding and Updating Software in Oracle Solaris 11.2”](#).

3. **Install the HA for Oracle Traffic Director software package.**

```
# pkg install ha-cluster/data-service/oracle-traffic-director
```

4. **Verify that the package installed successfully.**

```
$ pkg info ha-cluster/data-service/oracle-traffic-director
```

Installation is successful if output shows that State is Installed.

5. **Perform any necessary updates to the Oracle Solaris Cluster software.**

For instructions on updating your software, see [Chapter 11, “Updating Your Software,”](#) in [“Oracle Solaris Cluster System Administration Guide”](#).

Registering and Configuring HA for Oracle Traffic Director

You can configure HA for Oracle Traffic Director to be used with a logical hostname or a shared address. In the first procedure in this section, these additional steps begin with a notation that they are required for scalable services only. Individual examples of a failover service and a scalable service follow the procedure.

- [“Setting HA for Oracle Traffic Director Extension Properties” on page 13](#)
- [“Setting IP Addresses and Port Numbers for the Oracle Traffic Director Application” on page 14](#)
- [“How to Register and Configure HA for Oracle Traffic Director With a Failover Logical Host” on page 14](#)
- [“How to Register and Configure HA for Oracle Traffic Director With a Scalable Shared Address” on page 20](#)

Setting HA for Oracle Traffic Director Extension Properties

The sections that follow contain instructions for registering and configuring HA for Oracle Traffic Director resources. For information about the extension properties, see [Appendix A, “HA for Oracle Traffic Director Extension Properties”](#). The Tunable entry indicates when you can update a property.

See the [r_properties\(5\)](#), [rg_properties\(5\)](#), and [rt_properties\(5\)](#) man pages for details on all the Oracle Solaris Cluster properties.

Also, see the [clresource\(1CL\)](#), [clresourcegroup\(1CL\)](#), and [clresourcetype\(1CL\)](#) man pages.

To set an extension property of a resource, include the following option in the `clresource` command that creates or modifies the resource:

`-p property=value`

`-p property`

Identifies the extension property that you are setting.

`value`

Specifies the value to which you are setting the extension property.

You can also use the procedures in [Chapter 2, “Administering Data Service Resources,” in “Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) to configure resources after the resources are created.

Setting IP Addresses and Port Numbers for the Oracle Traffic Director Application

The `Resource_dependencies` setting on the Oracle Traffic Director application resource determines the set of IP addresses that Oracle Traffic Director uses. The `Port_list` setting on the resource determines the list of port numbers that Oracle Traffic Director uses. The fault monitor monitors the Oracle Traffic Director instance through the first port in `Port_list`. All the ports used by Oracle Traffic Director must be configured in `Port_list` in order for the shared address load balancing feature to work correctly.

▼ How to Register and Configure HA for Oracle Traffic Director With a Failover Logical Host

This procedure describes how to use the Oracle Solaris Cluster maintenance commands to register and configure HA for Oracle Traffic Director with one or more failover logical hostnames.

To perform this procedure, you must have the following information.

- The name of the resource type for HA for Oracle Traffic Director. This name is `ORCL.otd`.
- The names of the cluster nodes on which Oracle Traffic Director is installed and running.
- The logical hostname (for failover services) that clients use to access the data service.
- The path to the Oracle Traffic Director instances. You can install the Oracle Traffic Director binaries on the local disks or the cluster file system. See [“Configuration Guidelines for Oracle Solaris Cluster Data Services”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#) for a discussion of the advantages and disadvantages of each location.
- When using a logical hostname, you must configure Oracle Traffic Director to bind to `INADDR_ANY` since the logical hostnames are available on only one node at a time. The fault monitor probes the local hostname when used with a logical hostname unless overridden by the `Server_URL` extension property.

Note - Perform this procedure on any cluster member.

Note - Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. **On a cluster member, become an administrator that provides `solaris.cluster.modify` and `solaris.cluster.admin` authorization.**

2. Register the ORCL.otd resource type for HA for Oracle Traffic Director.

```
# clresourcetype register ORCL.otd
```

3. Create a failover resource group for the logical hostname.

You can optionally select the set of nodes on which the data service can run with the `-n` option.

```
# clresourcegroup create [-n node[,...]] resource-group
```

resource-group

Specifies the name of the failover logical hostname resource group. This name can be your choice but must be unique for resource groups within the cluster.

`-n node[,...]`

Specifies a comma-separated, ordered list of nodes that can master this resource group.

This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

4. Verify that all of the network addresses that you use have been added to your name service database.

You should have performed this verification during the Oracle Solaris Cluster installation. See the planning chapter in the [“Oracle Solaris Cluster Software Installation Guide”](#) for details.

Note - To avoid any failures because of name service lookup, ensure that all logical hostnames and shared addresses are present in the server's and client's `/etc/inet/hosts` file. Configure name service mapping in `/etc/nsswitch.conf` on the servers to first check the local files before trying to access NIS or NIS+.

5. Add a logical hostname to the failover resource group.

```
# clreslogicalhostname create -g resource-group \  
-h logical-hostname[,...] \  
[-N netiflist] \  
resource
```

`-g resource-group`

Specifies the name of the failover resource group.

`-h logical-hostname[,...]`

Specifies a comma-separated list of logical hostnames that this resource is to make available.

-N *netiflist*

Specifies an optional, comma-separated list that identifies the IPMP groups that are on each node. The format of each entry in the list is *netif@node*. The replaceable items in this format are as follows:

netif Specifies an IPMP group name, such as *sc_ipmp0*, or a public network interface card (NIC). If you specify a public NIC, Oracle Solaris Cluster attempts to create the required IPMP groups.

node Specifies the name or ID of a node.

This list is optional. If you omit this list, Oracle Solaris Cluster attempts to create the required IPMP groups.

Note - If you require a fully qualified hostname, you must specify the fully qualified name with the **-h** option and you cannot use the fully qualified form in the resource name.

Note - Oracle Solaris Cluster does not currently support using the adapter name for *netif*.

resource

Specifies the name of the resource.

6. Bring the failover logical host resource group online.

```
# clresourcegroup online -emM resource-group
```

7. Configure the Oracle Traffic Director server instance that uses the failover logical host.

8. Create a scalable application resource group.

```
# clresourcegroup create -p Maximum primaries=m \  
-p Desired primaries=n \  
scalable-resource-group
```

-p *Maximum primaries=m*

Specifies the maximum number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

-p *Desired primaries=n*

Specifies the desired number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

scalable-resource-group

Specifies the scalable resource group.

Note - If the value of the `Maximum primaries` is equal to the value of the `Desired primaries`, and if that number is equal to the number of nodes in the nodes list, then you can substitute the `-p Maximum primaries=m \ -p Desired primaries=n` options with the `-S` option. For example:

```
# clresourcegroup create -S resource-group
```

9. Add an application resource to the application resource group.

You can repeat this step to add multiple application resources to the same resource group.

```
# clresource create -g resource-group \
-t resource-type \
-p ORACLE_HOME=oracle-traffic-director-installation-directory \
-p INSTANCE_HOME=instance-directory \
-p Port_list=port-number/protocol \
resource
```

-g resource-group

Specifies the name of the failover resource group into which the resources are to be placed.

-t resource-type

Specifies the type of the resource to add.

-p ORACLE_HOME=oracle-traffic-director-installation-directory

Specifies the directory where the Oracle Traffic Director software has been installed. This is a per-node setting and if the location is different on each node, each node must be qualified with the node name. For example:

```
-p ORACLE_HOME{node1}=oracle-traffic-director-installation-directory-node1 \
-p ORACLE_HOME{node2}=oracle-traffic-director-installation-directory-node2
```

-p INSTANCE_HOME=instance-directory

Specifies the directory where the Oracle Traffic Director instance configuration is located. This is a per-node setting and if the location is different on each node, each node must be qualified with the node name. For example:

```
-p INSTANCE_HOME{node1}=instance-directory-node1 \
-p INSTANCE_HOME{node2}=instance-directory-node2
```

-p Port_list=port-number/protocol

Specifies a comma-separated list of port numbers and protocol to use, for example, `80/tcp` and `81/tcp`.

resource

Specifies the name of the resource to add.

The resource is created in the enabled state.

Note - Optionally, you can set additional extension properties that belong to the HA for Oracle Traffic Director to override the default values of the properties. See [Appendix A, “HA for Oracle Traffic Director Extension Properties”](#) for a list of these properties.

10. Create a strong positive affinity with delgated failover from the logical hostname resource group to the Oracle Traffic Director instance resource group.

```
# clresourcegroup set \  
-p RG_affinities+==++instance-resource-group logical-hostname-resource-group
```

instance-resource-group

Specifies the name of the Oracle Traffic Director instance resource group.

logical-hostname-resource-group

Specifies the name of the failover logical hostname resource group.

11. Bring the scalable application resource group online.

```
# clresourcegroup online -eM resource-group
```

resource-group

Specifies the name of the scalable application resource group.

12. Create a local node offline restart dependency from the logical hostname resource to the Oracle Traffic Director instance resource.

```
# clresource set -p Resource_dependencies_offline_restart+=instance-  
resource{local_node} /  
logical-hostname-resource
```

instance-resource

Specifies the name of the Oracle Traffic Director instance resource.

logical-hostname-resource

Specifies the name of the failover logical hostname resource.

Example 1-1 Creating an Oracle Traffic Director Resource for Use With a Failover Logical Hostname

This example creates an Oracle Traffic Director resource, `otd-rs`, in a resource group, `otd-rg`. It is configured to run simultaneously on all the four nodes of a four-node cluster.

In an agent configuration where a logical hostname is being used, the Oracle Traffic Director instances must be configured to listen on all addresses, `INADDR_ANY`, which allows the fault monitor to connect to the `localhost` address of each node using the default `Port_List`, `80/tcp`. The clients use the IP addresses as configured in a logical hostname resource, `lh-rs`, which is contained in the resource group, `lh-rg`. The hostname `otd-a-lh` is configured in the naming service used by the cluster and in any of the clients that will be accessing the server instances.

Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

To create a logical hostname in this example:

```
# clresourcegroup create -p
  Nodelist="node1,node2,node3,node4"
  \
-p Failback=True lh-rg
# clreslogicalhostname create -g
lh-rg -h otd-a-lh lh-rs
# clresourcegroup online -eM lh-rg
```

To facilitate the automatic failover of the logical hostname to a node that has a running instance of Oracle Traffic Director:

- The logical hostname resource group must have a strong positive affinity with fail-over delegation to the Oracle Traffic Director resource group.
- The logical hostname resource must also have an offline-restart dependency on the Oracle Traffic Director resource with a local-node scope.

To create the Oracle Traffic Director resource group and resource, do the following:

```
# clresourcegroup create -S otd-rg
# clresourcetype register ORCL.otd
# clresource create -g otd-rg -t ORCL.otd
  \
-p ORACLE_HOME=/global/otd/otd-home
  \
-p INSTANCE_HOME{node1}=
/global/otd/otd-1/net-otd-a \
-p INSTANCE_HOME{node2}=
/global/otd/otd-2/net-otd-a \
-p INSTANCE_HOME{node3}=
/global/otd/otd-3/net-otd-a \
-p INSTANCE_HOME{node4}=
/global/otd/otd-4/net-otd-a \
-p Resource_dependencies_offline_restart=otd-gfs-rs
  \
-p Port_List=80/tcp
  \
otd-rs
```

```
# clresource set -p Resource_dependencies_offline_restart+=  
otd-rs{local_node} lh-rs  
# clresourcegroup online -eM otd-rg  
# clresourcegroup set -p RG_affinities+=+++otd-rg lh-rg
```

▼ How to Register and Configure HA for Oracle Traffic Director With a Scalable Shared Address

Before You Begin Ensure that the `/etc/netmasks` file has IP-address subnet and netmask entries for all logical hostnames. If necessary, edit the `/etc/netmasks` file to add any missing entries.

1. **On a cluster member, become an administrator that provides `solaris.cluster.modify` and `solaris.cluster.admin` authorization.**
2. **Register the `ORCL.otd` resource type for HA for Oracle Traffic Director.**

```
# clresourcetype register ORCL.otd
```

3. **Create a failover resource group for the scalable shared address resource.**

You can optionally select the set of nodes or zones on which the data service can run with the `-n` option.

```
# clresourcegroup create [-n node[,...]] resource-group
```

resource-group

Specifies the name of the failover resource group. This name can be your choice but must be unique for resource groups within the cluster.

`-n node[,...]`

Specifies a comma-separated, ordered list of nodes that can master this resource group.

This list is optional. If you omit this list, the global zone of each cluster node can master the resource group.

4. **Verify that all of the network addresses that you use have been added to your name service database.**

You should have performed this verification during the Oracle Solaris Cluster installation. See the planning chapter in the [“Oracle Solaris Cluster Software Installation Guide”](#) for details.

Note - To avoid any failures because of name service lookup, ensure that all logical hostnames and shared addresses are present in the server's and client's `/etc/inet/hosts` file. Configure name service mapping in `/etc/nsswitch.conf` on the servers to first check the local files before trying to access NIS or NIS+.

5. Add a scalable shared address resource to the failover resource group.

```
# clressharedaddress create -g resource-group \  
-h shared-address[,...] \  
[-N netiflist] \  
resource
```

-g resource-group

Specifies the name of the failover resource group.

-h shared-address[,...]

Specifies a comma-separated list of shared addresses that this resource is to make available.

resource-group

Specifies the name of the resource group. This name can be your choice but must be unique for resource groups within the cluster.

-N netiflist

Specifies an optional, comma-separated list that identifies the IPMP groups that are on each node or zone. The format of each entry in the list is *netif@node*. The replaceable items in this format are as follows:

<i>netif</i>	Specifies an IPMP group name, such as <i>sc_ipmp0</i> , or a public network interface card (NIC). If you specify a public NIC, Oracle Solaris Cluster attempts to create the required IPMP groups.
--------------	--

<i>node</i>	Specifies the name or ID of a node. To specify the global zone, or to specify a node without non-global zones, specify only <i>node[,...]</i>
-------------	---

This list is optional. If you omit this list, Oracle Solaris Cluster attempts to create the required IPMP groups.

Note - Oracle Solaris Cluster does not currently support using the adapter name for *netif*.

resource

Specifies the name of the resource.

6. Bring the failover resource group online.

```
# clresourcegroup online -eM otd-rg
```

otd-rg

Specifies the name of the failover resource group.

7. Configure the Oracle Traffic Director server instance that uses the shared address.

8. Create a scalable resource group to run on all the desired cluster nodes.

Create a resource group to hold a data service application resource. You must specify the maximum and desired number of primary nodes as well as a dependency between this resource group and the failover resource group that you created in [Step 3](#). This dependency ensures that in the event of failover, the resource manager starts the network resource before starting any data services that depend on the network resource.

```
# clresourcegroup create -p Maximum primaries=m \  
-p Desired primaries=n \  
scalable-resource-group
```

-p Maximum primaries=m

Specifies the maximum number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

-p Desired primaries=n

Specifies the desired number of active primary nodes allowed for this resource group. If you do not assign a value to this property, the default is 1.

scalable-resource-group

Specifies the scalable resource group.

Note - If the value of the `Maximum primaries` is equal to the value of the `Desired primaries`, and if that number is equal to the number of nodes in the nodes list, then you can substitute the `-p Maximum primaries=m \ -p Desired primaries=n` options with the `-S` option. For example:

```
# clresourcegroup create -S resource-group
```

9. Create an application resource in the scalable resource group.

You can repeat this step to add multiple application resources, such as secure and insecure versions, to the same resource group.

To set load balancing for the data service, use the two standard resource properties `Load_balancing_policy` and `Load_balancing_weights`. See the [r_properties\(5\)](#) man page for a description of these properties. Additionally, see the examples that follow this section.

```
# clresource create -g scalable-resource-group \  
-t resource-type \  
-p ORACLE_HOME=oracle-traffic-director-installation-directory \  
-p INSTANCE_HOME=instance-directory \  
-p Resource_dependencies=shared-address[,...] \  
-p Port_list=port-number/protocol[,...] \  
-p Scalable=True \  
resource
```

`-g scalable-resource-group`

Specifies the name of the scalable resource group into which the resources are to be placed.

`-t resource-type`

Specifies the type of the resource to add.

`-p ORACLE_HOME=oracle-traffic-director-installation-directory`

Specifies the directory where the Oracle Traffic Director software has been installed. This is a per-node setting and if the location is different on each node, each node must be qualified with the node name. For example:

```
-p ORACLE_HOME{node1}=oracle-traffic-director-installation-directory-node1 \
-p ORACLE_HOME{node2}=oracle-traffic-director-installation-directory-node2
```

`-p INSTANCE_HOME=instance-directory`

Specifies the directory where the Oracle Traffic Director instance configuration is located. This is a per-node setting and if the location is different on each node, each node must be qualified with the node name. For example:

```
-p INSTANCE_HOME{node1}=instance-directory-node1 \
-p INSTANCE_HOME{node2}=instance-directory-node2
```

`-p Resource_dependencies=shared-address[,...]`

Specifies a comma-separated list of network resources that identify the shared addresses that the data service uses.

`-p Port_list=port-number/protocol[,...]`

Specifies a comma-separated list of port numbers and protocol to be used, for example, `80/tcp,81/tcp`.

`-p Scalable=True`

Specifies a Boolean that is required for scalable services.

`resource`

Specifies the name of the resource to add.

The resource is created in the enabled state.

10. Bring the scalable resource group online.

```
# clresourcegroup online -eM resource-group
```

`-eM resource-group`

Specifies the name of the scalable resource group.

Example 1-2 Creating an Oracle Traffic Director Resource for Use With a Scalable Shared Address

This example creates an Oracle Traffic Director `otd-rs` resource named `otd-rg` in a resource group named `web-rg`, which is configured to run simultaneously on all four nodes of a four-node cluster.

The Oracle Traffic Director instances are configured to listen on port 80 and uses the IP addresses as configured in a `SharedAddress` resource named `sa-rs`, which is contained in the resource group `sa-rg`. The hostname `otd-a-sa`, is configured in the naming service used by the cluster and any of the clients that will be accessing the server instances.

To create the shared address resource group and resource for this example, do the following:

```
# clresourcegroup create sa-rg
# clressharedaddress create -g sa-rg -h
otd-a-sa sa-rs
# clresourcegroup online -eM sa-rg
```

To create the Oracle Traffic Director resource group and resource, do the following:

```
# clresourcegroup create -S otd-rg
# clresourcetype register ORCL.otd
# clresource create -g otd-rg -t ORCL.otd
\
-p ORACLE_HOME=/global/otd/otd-home
\
-p INSTANCE_HOME{node1}=
/global/otd/otd-1/net-otd-a \
-p INSTANCE_HOME{node2}=
/global/otd/otd-2/net-otd-a \
-p INSTANCE_HOME{node3}=
/global/otd/otd-3/net-otd-a \
-p INSTANCE_HOME{node4}=
/global/otd/otd-4/net-otd-a \
-p Resource_dependencies_offline_restart=otd-gfs-rs
\
-p Resource_dependencies=sa-rs
\
-p Port_List=80/tcp
\
-p Scalable=True \
otd-rs
# clresourcegroup online -eM otd-rg
```

Tuning the HA for Oracle Traffic Director Fault Monitor

The HA for Oracle Traffic Director fault monitor is contained in the resource that represents Oracle Traffic Director. You create this resource when you register and configure HA for Oracle

Traffic Director. For more information, see [“Registering and Configuring HA for Oracle Traffic Director” on page 13](#).

System properties and extension properties of this resource control the behavior of the fault monitor. The default values of these properties determine the preset behavior of the fault monitor. The preset behavior should be suitable for most Oracle Solaris Cluster installations. Therefore, you should tune the HA for Oracle Traffic Director fault monitor *only* if you need to modify this preset behavior.

For more information, see the following sections.

- [“Tuning Fault Monitors for Oracle Solaris Cluster Data Services”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide ”](#)
- [“Changing Resource Type, Resource Group, and Resource Properties”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide ”](#)
- [“Standard Properties”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide ”](#)

Operations by the Fault Monitor During a Probe

The probe for HA for Oracle Traffic Director uses a request to the Oracle Traffic Director instance to query the health of the Oracle Traffic Director instance. Before the probe actually queries the Oracle Traffic Director instance, a check is made to confirm that the network resources are configured for the Oracle Traffic Director instance resource. If no network resources are configured, an error message is logged, and the probe exits with a failure.

If the Oracle Traffic Director instance resource is used in conjunction with a failover logical hostname, the fault monitor will retrieve the `Port_list` extension property and connect to the Oracle Traffic Director instance through the localhost on that port. When used with a Shared Address, the `Resource_dependencies` resource-property setting on the Oracle Traffic Director resource determines the set of IP addresses that the Oracle Traffic Director instance uses. The `Port_list` resource-property setting determines the list of port numbers that the Oracle Traffic Director instance uses. The fault monitor monitors the Oracle Traffic Director server instance through the first port in `Port_list`. All the ports used by Oracle Traffic Director must be configured in `Port_list` in order for the shared address load balancing feature to work correctly.

If the probe fails to connect to the Oracle Traffic Director instance using a specified IP address and port combination, a complete failure occurs. The probe records the failure and takes appropriate action.

The probe sends an HTTP HEAD request to the Oracle Traffic Director instance and waits for the response. The request can be unsuccessful for various reasons, including heavy network traffic, heavy system load, and misconfiguration.

Misconfigurations can occur during the following conditions.

- The Oracle Traffic Director instance is not configured to listen on all ports (INADDR_ANY), when the Oracle Traffic Director resource is configured with a failover logical hostname.
- The Oracle Traffic Director instance is not configured to listen on the shared address when the Oracle Traffic Director resource is configured with a shared address.
- The `Resource_dependencies` and `Port_list` resource properties were not set correctly when you created the resource.

If the reply to the query is not received within the `Probe_timeout` resource time limit, the probe considers this probe a failure of HA for Oracle Traffic Director. The failure is recorded in the probe's history.

A probe failure can be a complete or partial failure. The following probe failures are considered complete failures.

- Failure to connect to the server. The following error message is sent, where `%s` indicates the host name and `%d` indicates the port number.

```
Failed to connect to %s port %d
```
- Timeout (exceeding the resource-property timeout `Probe_timeout`) after trying to connect to the server.
- Failure to successfully send the probe string to the server. The following error message is sent, where the first `%s` indicates the host name, `%d` indicates the port number, and the second `%s` indicates further details about the error.

```
Failed to communicate with server %s port %d: %s
```

The monitor accumulates two such partial failures within the resource-property interval `Retry_interval` and counts them as one failure.

The following probe failures are considered partial failures.

- Timeout (exceeding the resource-property timeout `Probe_timeout`) while trying to read the reply from the server to the probe's query.
- Failing to read data from the server for other reasons. The following error message is sent, where the first `%s` indicates the host name, `%d` indicates the port number, and the second `%s` indicates further details about the error.

```
Failed to communicate with server %s port %d: %s
```

The probe connects to the Oracle Traffic Director instance and performs an HTTP GET check by sending a HTTP request to the `Server_URL` property. If the HTTP server return code is 500 (Internal Server Error) or if the connect fails, the probe will take action.

The result of the HTTP requests is either a success or a failure. If all the requests from the Oracle Traffic Director instance receive a response that indicates a success, then the probe returns and continues the next cycle of probing.

Heavy network traffic, heavy system load, and misconfiguration can cause the HTTP GET probe to fail. Misconfiguration of the `Server_URL` property can cause a failure if a URI in the `Server_URL` property includes an incorrect port or hostname. For example, if the URI was specified to connect to a host other than the localhost when configured with a failover logical host, or to connect to a host other than the shared address when configured with a shared address.

Based on the history of failures, a failure can cause either a local restart or a failover of the data service. This action is further described in [“Tuning Fault Monitors for Oracle Solaris Cluster Data Services”](#) in [“Oracle Solaris Cluster Data Services Planning and Administration Guide”](#).

HA for Oracle Traffic Director Extension Properties

This section describes the extension properties for the resource type `ORCL.otd`. This resource type represents the Oracle Traffic Director application in an Oracle Solaris Cluster configuration.

For details about system-defined properties, see the [r_properties\(5\)](#) man page and the [rg_properties\(5\)](#) man page.

The extension properties of the `ORCL.otd` resource type are as follows:

ORACLE_HOME

This property is a per-node string of the Oracle Traffic Director installation directory path name. You must specify this property for each node at resource creation time.

Category	Required
Data type	String
Default	---not set---
Tunable	When disabled

INSTANCE_home

This property is a per-node string of the Oracle Traffic Director installation directory path name. You must specify this property for each node at resource creation time.

Category	Required
Data type	String
Default	---not set---
Tunable	When disabled

Server_URL

By default, the Oracle Traffic Director server instance is probed either through the `localhost:port` when a logical host is used, or the `shared_address:port` when a shared address is used.

This property allows you to specify an alternate URL to monitor the Oracle Traffic Director server instance. The fault monitor periodically runs the HTTP GET command for the URL specified and takes action if the HTTP request returns with response code 500 "Internal Server Error", or if the application server does not respond within the configured timeout period.

Category	Required
Data type	String
Default	Null
Tunable	Any time

Monitor_retry_count

Controls the restarts of the fault monitor. This property indicates the number of times the fault monitor is restarted by the process monitor facility and corresponds to the `-n` option passed to the `pmfadm` command. The number of restarts is counted in a specified time window. See the `Monitor_retry_interval` property for more information. Note that this property refers to the restarts of the fault monitor itself, not the Oracle Traffic Director instance. The restarts of the Oracle Traffic Director instance are controlled by the system-defined properties `Thorough_probe_interval`, `Retry_interval`, and `Retry_count`, as specified in their descriptions. See the [clresource\(1CL\)](#) man page for more information. You can modify the value for this property at any time.

Category	Required
Data type	Integer
Minimum	-1
Default	4
Tunable	Any time

Monitor_retry_interval

Indicates the time in minutes, over which the failures of the fault monitor are counted, and corresponds to the `-t` option passed to the `pmfadm` command. If the number of times the fault monitor fails exceeds the value of `Monitor_retry_count`, the fault monitor is not

restarted by the Process Monitor Facility. You can modify the value for this property at any time.

Category	Required
Data type	Integer
Minimum	-1
Default	2
Tunable	Any time

Probe_timeout

This property is the time-out value, in seconds, that is used by the fault monitor to probe an Oracle Traffic Director instance. You can modify the value for this property at any time.

Category	Required
Data type	Integer
Minimum	15
Default	90
Tunable	Any time

Num_probes

This property indicates the number of times that the fault monitor can timeout while probing before taking action on an Oracle Traffic Director server instance. You can modify the value for this property at any time.

Category	Required
Data type	Integer
Minimum	1
Default	2
Tunable	Any time

Index

E

extension properties
ORCL.otd resource type, 29

F

fault monitor
HA for Oracle Traffic Director, 24

H

HA for Oracle Traffic Director
configuration
 planning, 10
fault monitor, 24
installation
 planning, 10
 task map, 10
installing, 11
software package, installing, 11
task map, 10

I

installing
HA for Oracle Traffic Director, 11

O

Oracle Solaris Cluster software
 publisher, 12, 12
ORCL.otd resource type
 extension properties, 29

P

package, 11
publisher
 Oracle Solaris Cluster software, 12, 12

S

software package, 11

T

task map
 HA for Oracle Traffic Director, 10

