

**Using Oracle® Compute Cloud Service -
Dedicated Compute Capacity - SPARC
Model 300**



Part No: E74671-08
April 2018

Using Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300

Part No: E74671-08

Copyright © 2016, 2018, 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, then 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 about 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 unless otherwise set forth in an applicable agreement between you and Oracle. 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, except as set forth in an applicable agreement between you and Oracle.

Access to Oracle Support

Oracle customers that have purchased support 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.

Référence: E74671-08

Copyright © 2016, 2018, 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 stipulation expresse de votre contrat de licence ou de la loi, vous ne pouvez pas copier, reproduire, traduire, modifier, accorder de licence, 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 livré sous licence au Gouvernement des Etats-Unis, ou à quiconque qui aurait souscrit la licence de ce logiciel 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 un risque de dommages corporels. Si vous utilisez ce logiciel ou ce 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 des applications dangereuses.

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 de 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle. 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, sauf mention contraire stipulée dans un contrat entre vous et Oracle.

Accès aux services de support Oracle

Les clients Oracle qui ont souscrit un contrat de support ont accès au support électronique via My Oracle Support. Pour plus d'informations, visitez le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> ou le site <http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> si vous êtes malentendant.

Contents

Get Started With Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300	7
▼ Step 1: Get a Subscription	7
Step 2: Learn About the Service	8
Service Features	8
Users and Roles	10
Terminology	11
▼ Step 3: Log in to the Service	12
▼ Step 4: Learn How to Access the CLIs and BUI	12
Step 5: Learn About Opening Service Requests	13
▼ Update Your Oracle Support Profile	13
▼ Create a Service Request	14
 Managing Instances	 17
Learn About Instance Configurations	17
Storage	18
Network Settings	18
Instance Life Cycle	18
▼ Restart an Instance	19
▼ Re-create an Instance	19
Manage SSH Keys for Secure Access	20
About SSH Keys	20
▼ Learn About Generating an SSH Key Pair	21
▼ Generate an SSH Key Pair on Oracle Solaris	21
▼ Add an SSH Public Key to an Instance	22
▼ Modify an SSH Public Key	22
▼ Delete an SSH Public Key	23
▼ Log In to an Instance	23

Managing Storage	25
Learn About Storage	25
Managing Storage Volumes	26
▼ Create a Volume	26
▼ Prepare to Mount the Volume	30
▼ Create Initiators and Initiator Groups	31
▼ Define Storage Projects	32
▼ Define Storage Volumes	33
▼ Attach a Volume to an Instance	35
▼ Prepare a Volume for Use With Zones	37
▼ Configure an iSCSI LUN as a Native or Branded Zone Root Pool	39
▼ Configure an iSCSI LUN as a Kernel Zone Root Pool	40
▼ Create a Volume Snapshot	41
▼ Increase the Size of a Volume	42
▼ Detach a Volume	43
▼ Destroy a Volume	43
Managing Storage NFS Shares	44
▼ Create a Share	44
▼ Create a Network Exception	46
 Configuring Networking	 49
Learn About Network Settings	49
Service Port Details	50
▼ Reserve and Manage Public IP Addresses	51
▼ Attach Additional Network Devices to Compute Instances	52
 Getting Answers to Your Questions	 55
What's New	56
Best Practices	56
General Best Practices	56
Instance Best Practices	57
Storage Best Practices	58
Performance Best Practices	58
Troubleshoot the Service	59

Get Started With Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300

Complete these steps to get started using Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300.

- [“Step 1: Get a Subscription” on page 7](#)
- [“Step 2: Learn About the Service” on page 8](#)
- [“Step 3: Log in to the Service” on page 12](#)
- [“Step 4: Learn How to Access the CLIs and BUI” on page 12](#)
- [“Step 5: Learn About Opening Service Requests” on page 13](#)

Related Information

- [“Managing Instances” on page 17](#)
- [“Managing Storage” on page 25](#)
- [“Configuring Networking” on page 49](#)
- [“Getting Answers to Your Questions” on page 55](#)
- cloud.oracle.com

▼ Step 1: Get a Subscription

To get started with Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300, you must request a paid subscription.

After you purchase the subscription, Oracle will notify you when the service subscription is ready, and you can log in to the service and begin managing instances.

Note - Trial subscriptions are not available for this release of the service.

1. Purchase a subscription.

You cannot order the Oracle SPARC Model 300 Service from the available online tools. Instead, contact an Oracle Sales representative to purchase a subscription. Your Sales representative can review pricing and other information with you. To locate contact information, visit the [Oracle Corporate Contact](#) page.

Note - You can also access pricing and other product information by selecting the Infrastructure tab at cloud.oracle.com and accessing the Compute page.

2. Prepare the information needed to order a subscription.

When you order a paid subscription, Oracle will provide you with a spreadsheet into which you must specify information about the person (yourself, or another person) who will be the account administrator for the service. You will also supply other network and administration information. Refer to the spreadsheet for all of the required information.

Next Steps

- [“Step 2: Learn About the Service” on page 8](#)

Step 2: Learn About the Service

After you subscribe to the service and while you are waiting to receive notification that the service has been provisioned, learn about these topics to prepare to manage the service and its features.

- [“Service Features” on page 8](#)
- [“Users and Roles” on page 10](#)
- [“Terminology” on page 11](#)

Service Features

Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300 is a secure, reliable, low-cost, standards-based infrastructure service. You can use it to rapidly access Oracle Solaris Zones on Oracle Cloud with all the necessary storage and networking resources. You can manage and scale your virtualized environments in the cloud easily, and migrate your Oracle and third-party applications to Oracle Cloud.

Let's take a look at some of the features of the service (you can also view a [video](#) that provides an overview of these features):

- Benefit from the performance advantages provided by the hardware used by the service.
Oracle SPARC Model 300 Service runs on the world's fastest processor, the SPARC M7. This processor dramatically outperforms x86 for Java, database, in-memory, and security workloads. Oracle SPARC Model 300 Service offers superior database performance for real-time analytics through on-chip query acceleration and memory decompression. The SPARC M7 processor also delivers vastly more efficient virtualization than x86, meaning that you can run more applications in your cloud infrastructure with less performance and resource overhead.
- Migrate your applications to the Oracle public cloud.
When you subscribe to the service, you are given access to a dedicated environment, called a site, that consists of high-performance SPARC servers reserved for your use. The configuration that you subscribe to for the Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300 gives you compute power equivalent to 300 OCPUs (a CPU capacity equivalent to ten 30 physical core Oracle SPARC M7 processors) of dedicated virtual computing power running Oracle Solaris 11 in the Oracle Public Cloud.
You can migrate your on-premises applications to Oracle Solaris Zones that you create on these dedicated sites to take advantage of the compute, storage, and network capabilities that Oracle Compute Cloud Service provides. Because you're the only tenant on the site, you get predictable performance in the cloud. You can extend your data center to Oracle Cloud by requesting Oracle Network Cloud Service - VPN for Dedicated Compute.
- Use existing licenses with Oracle SPARC Model 300 Service.
You can use existing Oracle Database, middleware, and application software licenses according to the terms of those licenses and in compliance with Oracle IaaS policies.
- Create zones using Oracle-provided and custom machine images.
Create custom Oracle Solaris Zones (kernel zones, native zones, and branded zones) within your Oracle SPARC Model 300 Service. Customizations include the disk size, Oracle Solaris platform version, memory size, and number of vCPUs.
Using Oracle Solaris Unified Archive tools, you can create, archive, and deploy zone images into native zones, branded zones, or kernel zones.
- Attach highly-available storage to instances.
You can use the full capacity of the ZFSSA storage (approximately 32 TB or 60 TB, depending on the revision of the service purchased) to enable all instances in the service to store data, applications, or boot volumes for your zones. If you delete a zone, the data stored in the storage volumes remains intact until you delete the volumes. Once this storage capacity has been exhausted, no additional block storage can be provided using the standard service.
- Exercise fine-grained control over network traffic.
You can control network traffic among individual instances and also between specific groups of instances and external hosts. You can also control traffic to and from instances over specific protocols and ports that you can define.

- Reserve and assign fixed public IP addresses.

The default SPARC 300 Model Service configuration provides private IPs, for which all traffic to the internet is blocked. Oracle can open ports for you, by request. For an instance that requires access to the Internet, you can reserve and use a static public IP address. See [“Reserve and Manage Public IP Addresses” on page 51](#).

- Ensure secure access to instances.

You can configure your Oracle SPARC Model 300 Service zones and instances to be accessed securely from remote hosts by using SSH. Oracle recommends using SSH keys. See [“Manage SSH Keys for Secure Access” on page 20](#).

Users and Roles

Become familiar with the user types and roles that are provided by default for the Oracle SPARC Model 300 Service so you understand which tasks are supported for them.

User Type and Roles	Description
Service Compute Administrator (opc)	This administrator of instances is a less powerful role than root, and typically manages file systems and software installations. The default password for this account is <code>solaris_opc</code> . SSH public key authentication is the only method to log into your instance.
Service Compute Root Administrator (root)	This superuser of instances has full administrative access. The <code>opc</code> user must assume this role to perform certain administrative tasks like creating additional users. The <code>root</code> user cannot login directly, by default. The initial password for the <code>root</code> role is <code>solaris_opc</code> . Upon first logging into the <code>root</code> role you will be asked to change its password.
Service Storage Administrator (opc)	This administrator performs storage-related tasks like creating, deleting, resizing, and cloning storage volumes. The default password for this account is <code>solaris_opc</code> .

The `opc` account that Oracle initially provides on the instances is a limited-capability user login account. You can transition to a full administrative-capability `root` account once you have logged in as the `opc` user if you need full administrative access to the instance.

The `opc` account that Oracle provides on the ZFS Storage Appliance has capabilities limited to iSCSI volume creation, NFS share creation, and SNMP configuration and initialization. In addition, some limited administration capabilities are also allowed. No full administrative account is provided. Instead you must file a service request (SR) for more advanced administrative tasks to be carried out by Oracle personnel. Earlier releases of SPARC Model 300 service did not offer NFS capabilities. If your currently provisioned account does not have

NFS configuration capabilities, you can raise a Service Request (SR) requesting that the internal MOS note ID 2355735.1 be applied to your M300 ZFSSA to enable these capabilities.

Note - You can create additional compute instance users and roles as needed for your environment. The creation, modification, and deletion of users and roles requires root permissions.

Terminology

These key terms are relevant to the Oracle SPARC Model 300 Service.

- *Instance*. An Oracle VM Server for SPARC domain in the Oracle SPARC Model 300 Service. Oracle creates instances for your service with a specific machine image and set CPU, memory, and network resources. See [“Learn About Instance Configurations” on page 17](#).
- *IP reservation*. A public IP address that you can route to an Oracle Compute Cloud Service instance that requires access to or from the Internet or other Oracle services. See [About Public IP Addresses](#).
- *Oracle CPU (OCPU)*. For this service, an OCPU provides CPU capacity equivalent to one physical core of an Oracle SPARC M7 processor.
- *Storage volume snapshot*. A backup of all data currently stored on a storage volume. Use this snapshot to restore or clone a storage volume. See [“Create a Volume Snapshot” on page 41](#).
- *Virtual network device (VNIC)*. A pseudo network interface created on top of a network datalink device such as the provided LACP aggregation or native i40ge devices. VNICs can have specific IP, VLAN, and MTU definitions as required by the customer (within the restrictions imposed by the OPC service).
- *Volume*. A disk storage location implemented as a volume (iSCSI LUN) on the ZFS Storage Appliance.
- *VPN endpoint*. A VPN tunnel between your datacenter and your Oracle Compute Cloud Service site.
- *Zone*. A zone provides a virtual mapping from the application to the platform resources. Zones allow application components to be isolated from one another even though some zone types share a single instance of the Oracle Solaris operating system. Zones use resource management components to control how applications use available system resources.

In the Oracle SPARC Model 300 Service, you can create an Oracle Solaris Native Zone, Oracle Solaris Kernel Zone, or Oracle Solaris Legacy Zone in a specific instance to create additional virtual environments. See [Zone Brands](#) for a description of the different zone types.

▼ Step 3: Log in to the Service

When you order Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300 Service, Oracle will work with your System Administrator to collect needed information to deploy the service. This includes IP assignments and SSH public keys. A service request (SR) is also required to setup a secure VPN connection from your data center to the Oracle SPARC Model 300 Service.

Note - For this release, the Compute Cloud Service web console or REST APIs are not supported.

1. Provide required information to Oracle.

Provide all the information that Oracle requests after processing your order for the Oracle SPARC Model 300 Service, including reserved IP subnets, SSH keys, and some additional optional information.

You must also submit an SR to setup the secure VPN connection between your data center and the Oracle SPARC Model 300 Service. For information about opening an SR to setup a VPN connection, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

2. Leverage the public keys that have been installed on the instances.

From the client host, where you prepared the public keys installed by Oracle in the service's instances for the opc user, log in through a conventional Secure Shell (SSH) session as the opc user. Specify the previously provided the private/public keys pair.

Next Steps

- [“Step 4: Learn How to Access the CLIs and BUI” on page 12](#)

▼ Step 4: Learn How to Access the CLIs and BUI

- Access additional information about these resources:

If you want to...	See:
Use the Oracle Solaris CLI to interact with the instances in the service	Intro(1M)
Use the Oracle ZFS Storage Appliance RESTful API to interact with the service's storage	Oracle ZFS Storage Appliance RESTful API Guide

If you want to...	See:
Use the Oracle ZFSSA Command Line to interact with the services storage	Command Line Interface (CLI) chapter in the <i>ZFS Storage Appliance Administration Guide</i>
Use the Oracle ZFS Storage Appliance intuitive graphical interface for administration tasks, visualization concepts, and analyzing performance data the service's storage	Browser User Interface (BUI): Oracle ZFS Storage Appliance Administration Guide

Note - The Service Storage Administrator account provided to you by Oracle has capabilities limited to only administering storage volumes. Performing advanced management, networking, and maintenance operations for the storage appliance are restricted to Oracle personnel.

Next Steps

- [“Managing Instances” on page 17](#)
- [“Managing Storage” on page 25](#)
- [“Configuring Networking” on page 49](#)
- [“Getting Answers to Your Questions” on page 55](#)

Step 5: Learn About Opening Service Requests

Some of the tasks you might typically perform yourself in an on-premise data center must be performed by Oracle Support for this service. For example, you must file a Service Request (SR) so Oracle Support can perform the following actions for you:

- Set up a secure VPN connection from your data center to the Model 300 service.
- Set up NAT routing for public IP addresses.
- Recover or reinstall instances.

Refer to these topics for information about opening an SR:

- [“Update Your Oracle Support Profile” on page 13](#)
- [“Create a Service Request” on page 14](#)

▼ Update Your Oracle Support Profile

Before You Begin If you have not already configured your Oracle Support login profile for this service, you will need to do so before you can open a Service Request from My Oracle Support. To update your

profile, you will need your Customer Service Identifier (CSI) number, which you can find in the Welcome Letter you received when Oracle established your service.

1. **Access My Oracle Support at (<https://support.oracle.com>).**
2. **Sign in to My Oracle Support using your Oracle login.**
3. **Add a privilege to your login profile.**
 - a. **Click your name (Available) from the top right corner of the My Oracle Support screen.**
 - b. **Click *My Account* from the drop down menu.**
 - c. **Under the Support Identifiers heading, click *Request Access*.**
 - d. **In the Request Access to a Support Identifier window, provide your Support Identifier number and click *Request Access*.**

Your Customer Service Identifier (CSI) number was provided in the Welcome Letter you received from Oracle under the heading "HOW TO ACCESS MY ORACLE SUPPORT."
 - e. **In the Organization Name box, type *oracle* and click *Validate*.**

Within 24 hours, you will receive email from Oracle notifying you that the SI has been added to your profile. The next time you log in to My Oracle Support, you will see the *Create SR* button on the Service Requests tab when you log in to My Oracle Support.

Next Steps

- [“Create a Service Request” on page 14](#)

▼ Create a Service Request

Before You Begin If you have not already configured your Oracle Support login profile for this service, you will need to do so before you can open a Service Request from My Oracle Support. See [“Update Your Oracle Support Profile” on page 13](#) for more information.

1. **Access My Oracle Support at (<https://support.oracle.com>).**
2. **Sign in to My Oracle Support using your Oracle login.**
3. **Click *Create SR* on the Service Requests tab.**

4. **Under "What is the Problem," complete the *Problem Summary* and *Problem Description* fields.**

For example, if one of your instances has failed, you might provide the name of the failed instance that needs to be reinstalled and describe what caused the failure.

5. **Under "Where is the Problem," select the *Cloud* tab and provide the following information:**

- In the *Service Type* field, specify *Oracle SPARC Compute Service*.
- In the remaining fields, specify the values for the relevant service.

6. **Click *Next* at the top right side of the window.**

7. **Complete the required fields for each category and click *Next* until you have provided all required information.**

For more information, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

8. **After you have provided all required information, click *Submit*.**

You can monitor your Service Request from the Service Requests tab on the My Oracle Support dashboard tab.

Managing Instances

Read these topics to learn how to manage instances for Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300.

- [“Learn About Instance Configurations” on page 17](#)
- [“Restart an Instance” on page 19](#)
- [“Re-create an Instance” on page 19](#)
- [“Manage SSH Keys for Secure Access” on page 20](#)
- [“Log In to an Instance” on page 23](#)

Related Information

- [“Managing Storage” on page 25](#)
- [“Configuring Networking” on page 49](#)
- cloud.oracle.com

Learn About Instance Configurations

Each Oracle SPARC Model 300 Service includes 10 instances, or Oracle VM Server for SPARC logical domains (LDoms). Oracle creates these instances for you during the provisioning cycle with a specific operating system and set CPU, memory, and network resources.

Each instance is defined by its operating system and root domain configuration. Oracle provisions each instance to contain:

- Oracle Solaris 11.3 with Support Repository Updates applied, which is initially provisioned by Oracle.

You are responsible for patching and maintaining the operating system of the instance and zones, as well as backups and restorations of zones. Oracle releases a critical patch updates (CPUs) for Oracle Solaris 11 every 3 months in January, April, July and October. Keep all of your instances up to date with each of these CPUs to ensure that you receive the latest security fixes and fixes for other critical issues.

- 30 cores and 442 GB of memory. This is the only configuration option.
- Full superuser access to the instance for typical life cycle operations.

Storage

The Oracle SPARC Model 300 Service includes a dedicated ZFS Storage Appliance (ZFSSA) that is preconfigured and maintained by Oracle.

The provided storage consists of two highly available storage heads and associated storage trays that are clustered in an active-active configuration. Each head can be used to provide storage services for maximum performance and capacity. Heads and storage resources will failover in the event of a significant failure, which allows the service to be maintained (at a potentially reduced performance) while Oracle performs required maintenance operations.

Network access to each storage head is provided by a redundant highly available 40 GBps aggregated network connection. Separate network subnets have been provisioned for accessing storage or for carrying out management tasks on each ZFSSA head. A storage administration account is provided on each head to manage and provision storage volumes. For more advanced operations, open a service request (SR) to request action by Oracle support personnel who have full administrative capabilities.

For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

Network Settings

By default, all of your instances can talk to other instances within the service over a redundant network. Instances cannot be accessed from public networks or external hosts, by default, but can be configured to route using public IP addresses. To setup network address translation (NAT) routing for public IP addresses, open a service request (SR). For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

To enable different network communication settings, see [“Configuring Networking” on page 49](#).

Instance Life Cycle

The life cycle management of your instance is straightforward, with a few limitations. Oracle creates the instance for you. If the instance needs to be rebuilt, this would need to be done by

Oracle. To have your instance reinstalled, open a service request (SR). For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

You are not able to resize your instance, however, you can reboot your instance, update the operating system, and create zones.

▼ Restart an Instance

Once your instance is running, you can restart it, if needed, from the Oracle Solaris CLI. When you restart an instance, data on persistent storage volumes is not lost. Your instance also retains all of the configuration information. For example, the IP address for the instance and storage volumes that are attached and mounted on the instance will be retained following a restart.

Complete these steps to restart an instance.

1. **Ensure that you have the root role on the instance.**

You must have the root role to restart an instance. If it isn't assigned to you, or you're unsure whether it is assigned to you, ask your System Administrator to ensure that you have this role.

2. **Within the instance, restart the instance using appropriate Oracle Solaris CLI commands.**

Refer to the Oracle Solaris `reboot(1M)` man page for information about restarting an instance.



Caution - Do not use the shutdown command, `init 0, 1, 2, 3, 4, 5, S`, or take any other action that would prevent the system booting into multi-user mode. Doing so will cause your instance to become inaccessible because there is no console access. If your instance is (or becomes) unavailable, submit a service request (SR) with Oracle Support. For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

▼ Re-create an Instance

- **Re-create the instance.**

If your instance becomes unavailable and you need to re-create it, submit a service request (SR) with Oracle Support to have the instance re-created for you. For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

Note - An instance is re-created to its initial state, which is the same state it was when you first accessed the service.

Manage SSH Keys for Secure Access

To access your Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300 service, you must create a secure connection to your instances. The two-part security setup that you perform after your service is ordered involves setting up a secure VPN connection from the Oracle SPARC Model 300 Service to your data center, and setting up secure SSH keys on each of the instances.

Read these topics for information about generating and using SSH keys to enable secure access to your instances.

- [“About SSH Keys” on page 20](#)
- [“Learn About Generating an SSH Key Pair” on page 21](#)
- [“Generate an SSH Key Pair on Oracle Solaris” on page 21](#)
- [“Add an SSH Public Key to an Instance” on page 22](#)
- [“Modify an SSH Public Key” on page 22](#)
- [“Delete an SSH Public Key” on page 23](#)

About SSH Keys

You can securely log in to Oracle SPARC Model 300 Service instances from a remote host using a secure shell (SSH) connection.

SSH is a cryptographic network protocol that can use a public key and a private key pair to provide secure communication between two computers. SSH uses network port 22 by default.

You can associate one or more SSH public keys with each instance, enabling secure communications from one or more hosts to one or more instances. Add each public key that is associated with an instance to the `opc` user accounts `authorized_keys` file located at `/export/home/opc/.ssh/authorized_keys` within an instance. SSH public keys are created as a text file (typically named `id_rsa.pub`), and you should append the contents of the public key file to the `authorized_keys` file and should not be edited in any way.

▼ Learn About Generating an SSH Key Pair

To access your instances using SSH, generate at least one SSH key pair just as you did during the initial service registration. You can use any SSH utility to generate SSH keys and to log in to your instances. For example, if you're logging in from a Windows host, you can use PuTTY. If you're using a Linux host, you can use OpenSSH.

1. **Generate an SSH key pair.**
See [“Generate an SSH Key Pair on Oracle Solaris” on page 21](#) for an example.
2. **Ensure that the private key is available on each host that you'll use to access instances.**
Use the private key to log in to the instances using SSH.
3. **Associate the public key with your instances.**



Caution - Keep your SSH keys secure. Establish policies to ensure that the keys aren't lost or compromised when employees leave the organization or move to other departments. If you lose your private key, you can't access your instances. For business continuity, ensure that the SSH keys of at least two IT system administrators are added to your instances.

▼ Generate an SSH Key Pair on Oracle Solaris

Complete these steps to generate an SSH key pair on UNIX and UNIX-like systems.

1. **Run the `ssh-keygen` command.**
You can use the `-t` option to specify the type of key to create. For example, to create an RSA key, type:


```
# ssh-keygen -t rsa
```


You can also use the `-b` option to specify the length (bit size) of the key. For example:


```
# ssh-keygen -b 2048 -t rsa
```
2. **When prompted, enter the path to the file in which you want to save the key.**
The prompt displays a suggested default path and file name in parentheses. For example, the prompt might display `/export/home/user_name/.ssh/id_rsa`.

■ **To accept the default path and file name and press Enter.**

- **To supply your own values, type the preferred path and file name, and then press Enter.**
- 3. **Provide a passphrase.**

A passphrase is not required, but you can use it to protect your private key against unauthorized use.

 - **When prompted, type a passphrase.**
 - **When prompted, confirm the passphrase by typing it a second time.**
- 4. **Make a note of the path and file names of the private and public keys.**

When you log in to an instance, you'll need to provide the path to the corresponding SSH private key and the passphrase, when prompted.



Caution - Keep your SSH keys secure so they aren't lost or compromised. If you lose your private key, you can't access your instances.

▼ Add an SSH Public Key to an Instance

Oracle will add an SSH public key to all instances when an instance is deployed. This is the SSH public key that you generated and provided to Oracle when you initially ordered the Oracle SPARC Model 300 Service.

- **Add SSH public keys to an instance so you can access the instance from other hosts.**

Add the remote system's public key file, `id_rsa.pub`, to the local system by appending the `id_rsa.pub` file to the instance user's `/export/home/user_name/.ssh/authorized_keys` file.

Note - Copy the full contents of the `id_rsa.pub` file as a single line. Do not insert line breaks or carriage returns.

▼ Modify an SSH Public Key

Before you make any changes to the `~/.ssh/authorized_keys` file of the `opc` user on an instance, back up the file. You will need this backup if the `authorized_keys` file gets corrupted,

or you inadvertently make changes that result in the `opc` user getting locked out of the instance. Remember, if there is no other user set up on an instance, and changes to the `~/.ssh/authorized_keys` file result in the `opc` user getting locked out, you might be unable to access your instance.

Note - To log in to an instance by using SSH, you need to provide the private key that matches a public key associated with the instance.

Complete these steps to backup the `authorized_keys` file.

1. **Start a second ssh session. This will be your back-up version of the file.**
Ensure that the second ssh session remains connected while you edit the `authorized_keys` file.
2. **Test the changes you made to the `~/.ssh/authorized_keys` file.**
Before closing the backup session, log in with the new or updated SSH key.

▼ Delete an SSH Public Key

You can stop the public key from accessing all hosts or an individual host.

- **Complete one of these steps to stop public key access.**
 - **Stop all hosts accessing an instance with its public key.**
Delete the corresponding private key file on the instance located at `/export/home/user_name/.ssh/id_rsa`.
 - **Stop an individual host from accessing a server with a public key.**
Delete the corresponding public key entry from the file on the remote host at `/export/home/user_name/.ssh/authorized_keys`.



Caution - Do not remove all private keys from the instance, Doing so will prevent all access to the instance rendering it inaccessible.

▼ Log In to an Instance

After you've configured the VPN tunnel to your Oracle SPARC Model 300 Service, you can use the private IP address of your instance to connect to the instance. If you've configured your

public inbound IP address, you can log in to the instance with the public IP assigned to that instance.

- **Connect to your Oracle Solaris instance using ssh.**

See [Accessing an Oracle Solaris Instance Using SSH](#).

Note - The Oracle- provisioned instance requires you to initially log in as the `opc` user using the SSH keys you provided during the ordering of the service. You can't log in as `root`.

Managing Storage

Read these topics to learn how to manage storage for Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300.

- [“Learn About Storage” on page 25](#)
- [“Managing Storage Volumes” on page 26](#)
- [“Managing Storage NFS Shares” on page 44](#)

Related Information

- [“Managing Instances” on page 17](#)
- [“Configuring Networking” on page 49](#)
- cloud.oracle.com

Learn About Storage

Storage is available as either shares or volumes. A storage volume is a virtual disk that provides persistent block storage space for instances, and a share is an NFS share that provides filesystem storage in the Oracle SPARC Model 300 Service. You can attach the full capacity of the ZFSSA storage (approximately 32 TB or 60 TB, depending on the revision of the service purchased) to your instances for storing data and applications in any mix of volumes and shares.

You can use storage volumes to:

- Store data and applications.
- Serve as a persistent boot disk for the Oracle Solaris Zones. To do this, create a zone, attach one or more storage volumes to it, and then specify the associated volume as the boot disk for the zone.
- Scale the block storage capacity up or down by attaching or detaching storage volumes.

You can use shares to:

- Easily share data between applications and instances.

When using storage with your instances, the full capacity of the ZFSSA is available. Once this capacity is exhausted, additional capacity is not available. You will need to expand into alternate data storage services available in the Oracle Cloud.

Managing Storage Volumes

You can create storage volumes on the ZFS Storage Appliance component of the Oracle SPARC Model 300 Service and can configure them to be accessible by one or more compute instances. You can administer this process using either the CLI or BUI tools. Use whichever method best addresses your needs. See [“Step 4: Learn How to Access the CLIs and BUI” on page 12](#) for more information.

When you create a storage volume, you can specify the capacity you need.

Review the following topics for examples of how to perform these tasks:

- [“Create a Volume” on page 26](#)
- [“Prepare to Mount the Volume” on page 30](#)
- [“Create Initiators and Initiator Groups” on page 31](#)
- [“Define Storage Projects” on page 32](#)
- [“Define Storage Volumes” on page 33](#)
- [“Attach a Volume to an Instance” on page 35](#)
- [“Prepare a Volume for Use With Zones” on page 37](#)
- [“Configure an iSCSI LUN as a Native or Branded Zone Root Pool” on page 39](#)
- [“Configure an iSCSI LUN as a Kernel Zone Root Pool” on page 40](#)
- [“Create a Volume Snapshot” on page 41](#)
- [“Increase the Size of a Volume” on page 42](#)
- [“Detach a Volume” on page 43](#)
- [“Destroy a Volume” on page 43](#)

▼ Create a Volume

Perform these steps to create a volume (iSCSI LUN) using the Oracle ZFS Storage Appliance browser user interface (ZFSSA BUI).

1. **On a web UI, use the `opc` user account to log into the ZFSSA head where you want to create the storage volume (LUN).**

You can access the BUI at https://zfssa_head_ip:215.

2. **Identify and make note of the ZFSSA network interface information for the volume you want to create.**

Storage volumes (iSCSI LUNs) are provided over network devices. When configuring a volume, you must supply the appropriate network device over which the volume will be served. Oracle SPARC Model 300 Service provides storage access over `vnic5` (corresponding to head 1) and `vnic7` (corresponding to head 2) of the ZFSSA. You must identify which of these devices you wish to serve the volume over when configuring the volume.

Complete these steps to identify the network interface information for the volume (LUN):

- a. **Navigate from the Configuration > Network screen of the BUI, then select the Configuration tab.**

- b. **Locate the network for the volume from the 'Interfaces' list.**

The network you choose will vary depending on which of the two ZFSSA heads is used for the volume (LUN).

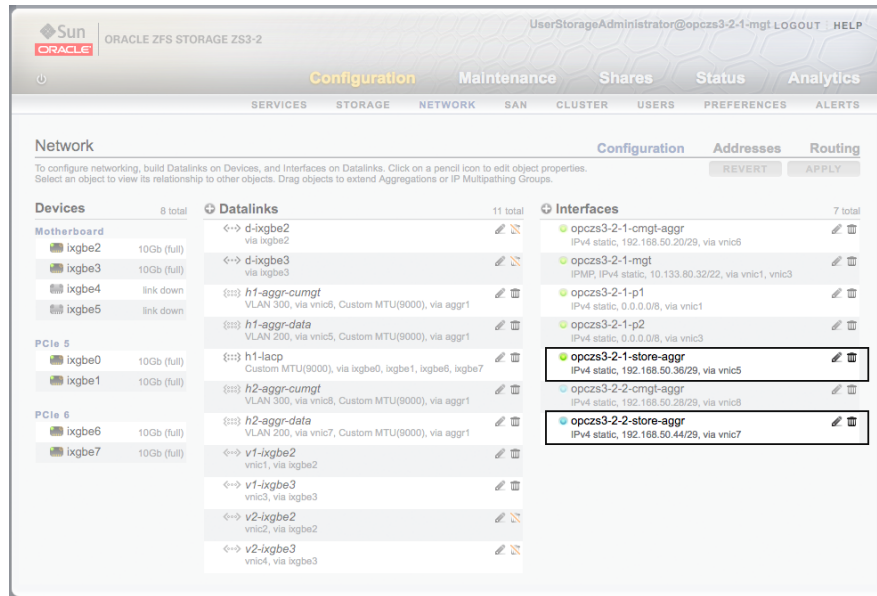
- Head 1 is associated with the interface named `opczs3-2-1-store-aggr`.
- Head 2 is associated with the interface named `opczs3-2-2-store-aggr`.

- c. **Make note of the following information and keep it accessible for upcoming steps.**

In this example, head 1 is used for the volume, so let's make note of these values:

- Interface name: `opczs3-2-1-store-aggr`
- VNIC shortname: `vnic5`

- IP address: 192.168.50.36

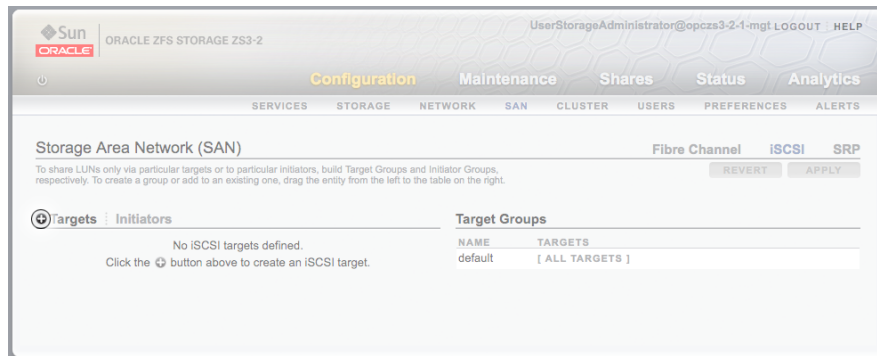


3. Define an iSCSI target.

Create a target group to enable the volumes (targets) to be shared with only specific initiators (one or more compute instances).

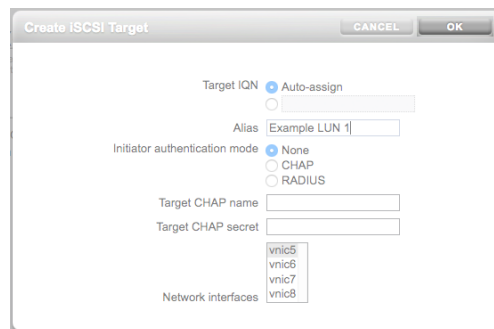
Complete these steps to create a target for the volume (LUN).

- a. Navigate from the Configuration > SAN screen of the BUI, select the iSCSI tab, and click the (+) button next to the Targets label, as shown here.



- b. When the 'Create iSCSI Target' window appears, enter values for the iSCSI target and click OK, as shown in this example.

- 'Alias' is an unambiguous reference name for administrators to use to refer to the disk.
- If additional security is required, select an 'Initiator authentication mode' and the associated target name and secret. This example does not enable CHAP or RADIUS.
- 'Network interface' is the name that was listed for the ZFSSA head for the volume (LUN) earlier in this procedure (vnic5).



4. Create a Target Group to which the target will belong.

Targets in the same target group will be visible to the same initiator(s). Complete these steps to create a target group:

- a. **Navigate from the Configuration > SAN screen of the BUI, select the iSCSI tab.**
- b. **Drag the target you want to add to a target group from the 'Targets' list to a group under the 'Target Groups' list.**

You can drag a target to an existing group or to the new group area to create a new group.

Note - Drag targets by floating the mouse over the target. When a drag handle appears to the left of the target, click and drag it to the desired target group.

In this example, the new group is created with the default name `targets-0`. You can change the default name, as we did in our example, by clicking the edit symbol that appears to the right of the name when you hover the mouse over the name. In the example, the group is renamed to `Example Target Group 1`.

- c. **Click Apply to create the volume (LUN).**

▼ Prepare to Mount the Volume

When you mount a new volume, you need to provide the iSCSI qualified name (IQN) for the initiators (compute instances) where the LUN will be mounted. This process should be repeated on each of the compute instances where an iSCSI LUN will be mounted.

Complete these steps to identify the IQN so you are prepared to mount the volume.

1. **Log in to the compute instance where you will mount the LUN with root privileges.**
2. **Use the `iscsiadm list initiator-node` command to identify the initiator node name.**

Ensure that you document the entire string (`iqn.1986-03.com.sun:01:opc-192-168-50-6`, in this example).

```
root@opc-192-168-50-6:~# iscsiadm list initiator-node
Initiator node name: iqn.1986-03.com.sun:01:opc-192-168-50-6
Initiator node alias: solaris.us.oracle.com
Login Parameters (Default/Configured):
```

```

Header Digest: NONE/-
Data Digest: NONE/-
Max Connections: 65535/-
Authentication Type: NONE
RADIUS Server: NONE
RADIUS Access: disabled
Tunable Parameters (Default/Configured):
Session Login Response Time: 60/-
Maximum Connection Retry Time: 180/-
Login Retry Time Interval: 60/-
Configured Sessions: 1
root@opc-192-168-50-6:~#

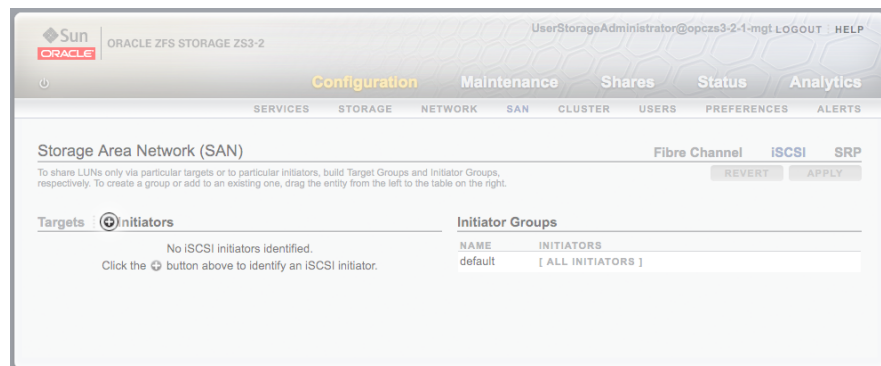
```

▼ Create Initiators and Initiator Groups

Initiators are compute instances that request storage. You can define groups of initiators such that the same target (LUN) is visible to multiple initiators (instances), which allows multiple instances to share access to one or more LUNs as is required for server clustering technologies like Oracle RAC or Oracle Solaris Cluster.

Complete these steps to create an initiator group:

1. **Log in to ZFSSA using the `opc` user account.**
2. **Navigate from the Configuration > SAN screen of the BUI.**
3. **Select the iSCSI tab and click the (+) button next to the Initiators label, as shown here.**



4. **When the 'Identify iSCSI Initiator' window appears, enter values for the iSCSI initiator and click OK, as shown in this example.**

- The initiator IQN that you noted in [“Prepare to Mount the Volume” on page 30](#) (iqn.1986-03.com.sun:01:opc-192-168-50-6).
- A required alias for the IQN (opc-192-168-50-6).
- If additional security is required, you can configure CHAP authentication for each initiator. This example does not use CHAP because the Oracle SPARC 300 Service storage network is entirely private, so can only be accessed by instances within your service.



5. **Create an initiator group.**

- a. **Drag the newly identified initiator to the 'Initiator Groups' list and edit the name of the group.**

This example uses compute instance group 1.

- b. **Click Apply to create the initiator group.**

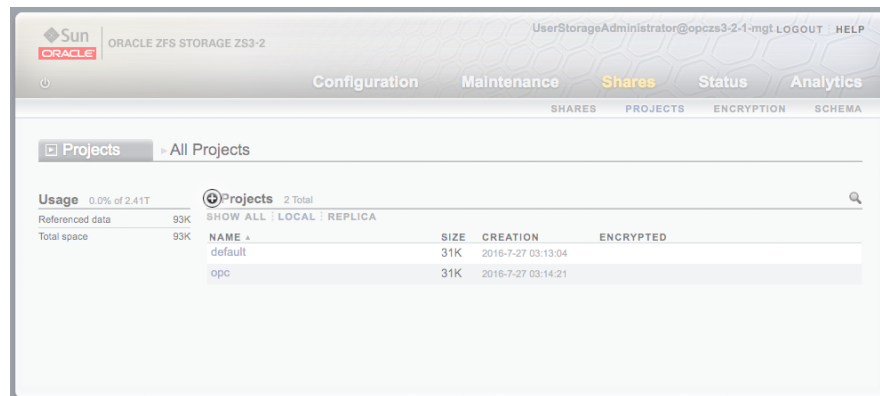
▼ Define Storage Projects

Projects are useful containers for specifying the scope of a snapshot of a group of volumes (LUNs). Oracle created a default opc project for you, but you can create additional projects, as needed.

Complete these steps to create a new project. This example uses the opc project.

1. **Log in to ZFSSA using the opc user account.**

2. **Navigate from the Shares > Projects screen of the BUI, and from the Projects tab, click the (+) button next to the Projects label, as shown here.**



3. **Enter a project name in the window that appears and click Apply.**
4. **Select the project you want to define (opc, in this example).**
Select the project by double clicking the name or by clicking the edit (pencil) button next to the project name.

▼ Define Storage Volumes

Oracle created a default opc project for you, which is the project where we'll add a storage volume (exampleLUN1) in this example.

Complete these steps to create a new storage volume (LUN).

1. **Log in to ZFSSA using the opc user account.**
2. **Navigate from the Shares > Projects screen of the BUI, and from the Projects tab, select the project you want to define (opc, in this example).**
Select the project by double clicking the name or by clicking the edit (pencil) button next to the project name.
3. **Select the 'LUNs' tab (next to the 'Filesystems'), and click the (+) button next to the LUNs label.**

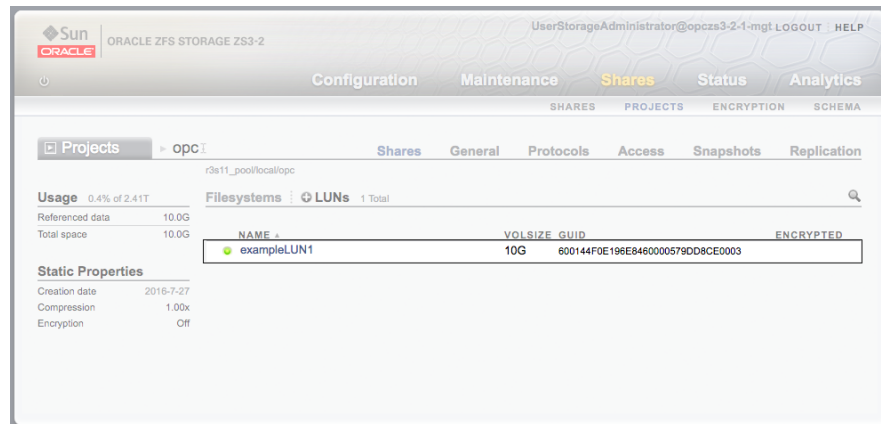
4. **When the 'Create LUN' window appears, enter values for the volume and click Apply, as shown in this example.**
 - Enter a name for the volume (exampleLUN1).
 - Enter a volume size for the LUN (10 G, for 10 gigabytes).
 - Select a target group (Example Target Group 1).
 - Select an initiator group (compute instance group 1).

The 'Create LUN' dialog box is shown with the following configuration:

- Project:** opc
- Name:** exampleLUN1
- Volume size:** 10 G
- Thin provisioned:** ☐
- Volume block size:** 8K
- Online:** ☒
- Target group:** Example Target Group 1
- Initiator group(s):** All initiators, compute instance group 1
- LU number:** 0, ☒ Auto-assign
- Encryption:** Off
- Inherit key:** ☐
- Key:** Local, OKM

5. **Verify that the volume was created by confirming its inclusion in the project.**

From the Shares > Projects screen of the BUI, select the `opc` project from the Projects tab, click the 'LUNs' tab (next to the 'Filesystems'). You can see the new volume (exampleLUN1), as shown here.



6. **Make note of the GUID number that is allocated to the share, as you will need it when you configure storage on the compute instance.**

In this example, the GUID is `600144F0E196E8460000579DD8CE0003`.

▼ Attach a Volume to an Instance

Complete these steps to attach a volume to an instance:

1. **Log into the compute instance with root privileges.**
2. **Configure 'Send target' iSCSI discovery.**

You only need to do this once on the instance.

```
root@opc-192-168-50-6:~# iscsiadm list discovery
Discovery:
Static: enabled
Send Targets: disabled
iSNS: disabled
root@opc-192-168-50-6:~# iscsiadm modify discovery -t enable
```

3. **Configure the IP address (192.168.50.36, in this example) of the ZFSSA head that is exporting the iSCSI LUNs.**

Remember that this address was noted in the Configurations tab during [“Create a Volume”](#) on page 26.



You only need to do this once on the instance.

```
root@opc-192-168-50-6:~# iscsiadm add discovery-address 192.168.50.36
Jul 31 12:12:58 opc-192-168-50-6 cmlb: WARNING:
/scsi_vhci/ssd@g600144f0e196e8460000579dd8ce0003 (ssd0):
Jul 31 12:12:58 opc-192-168-50-6 Corrupt label; wrong magic number
root@opc-192-168-50-6:~#
```

Note - The LUN should appear almost immediately upon its creation. You might receive warning messages on the instance console (for example, Corrupt label; wrong magic number, as shown in the preceding example). This simply indicates that the LUN is not yet formatted, and is expected behavior.

4. **To remove these warning messages, use the format command on the new disk.**

The format command will list all available disks in the compute instance and will offer to label any disks that appear to be missing labels, such as the newly created volume (LUN). For example,

```
root@opc-192-168-50-6:~# format
Searching for disks...
Jul 31 12:16:28 opc-192-168-50-6 cmlb: WARNING: /scsi_vhci/
ssd@g600144f0e196e8460000579dd8ce0003 (ssd0):
```

```

Jul 31 12:16:28 opc-192-168-50-6 Corrupt label; wrong magic number
Jul 31 12:16:28 opc-192-168-50-6 cmlb: WARNING: /scsi_vhci/
ssd@g600144f0e196e8460000579dd8ce0003 (ssd0):
Jul 31 12:16:28 opc-192-168-50-6 Corrupt label; wrong magic number done
c0t600144f0e196e8460000579DD8CE0003d0: configured with capacity of 9.94GB

AVAILABLE DISK SELECTIONS:
0. c0t600144f0e196e8460000579DD8CE0003d0 <SUN-ZFS Storage 7330-1.0 cyl 323 alt 2 hd 254
   sec 254>
/scsi_vhci/ssd@g600144f0e196e8460000579dd8ce0003
1. cld0 <SUN-DiskImage-20GB cyl 27902 alt 2 hd 96 sec 768> /virtual-devices@100/channel-
   devices@200/disk@0
Specify disk (enter its number): 0
selecting c0t600144f0e196e8460000579DD8CE0003d0 [disk formatted]
Jul 31 12:24:05 opc-192-168-50-6 cmlb: WARNING: /scsi_vhci/
ssd@g600144f0e196e8460000579dd8ce0003 (ssd0):
Jul 31 12:24:05 opc-192-168-50-6 Corrupt label; wrong magic number Disk not labeled.
Label it now? y
Jul 31 12:24:08 opc-192-168-50-6 cmlb: WARNING: /scsi_vhci/
ssd@g600144f0e196e8460000579dd8ce0003 (ssd0):
Jul 31 12:24:08 opc-192-168-50-6 Corrupt label; wrong magic number

FORMAT MENU:
disk - select a disk
type - select (define) a disk type
partition - select (define) a partition table
current - describe the current disk
format - format and analyze the disk
repair - repair a defective sector
label - write label to the disk
analyze - surface analysis
defect - defect list management
backup - search for backup labels
verify - read and display labels
save - save new disk/partition definitions
inquiry - show disk ID
volname - set 8-character volume name
!<cmd> - execute <cmd>, then return
quit
format> quit
root@opc-192-168-50-6:~#

```

▼ Prepare a Volume for Use With Zones

After you create a volume (LUN), you can prepare it for exclusive use by a zone.

By default, zone creation provides storage by automatically creating a ZFS filesystem in the Global Zone that contains all of the storage for the zone. This behavior is not optimal for zones created in the OPC environment since local storage is limited and use of local storage prevents migration of zones between instances. For OPC zones, configure zones to use volumes (iSCSI LUNs) provided by the ZFSSA as they can be shared and duplicated more easily.

Complete these steps to remove the default storage that is allocated to a zone and instead configure use of a disk volume (iSCSI LUN).

1. Log in to the CLI with root privileges.

2. Identify the disks that are available on the instance for use by zones.

To exit the prompt after viewing the necessary information, press the ctrl+c key combination.

```
root@opc-192-168-50-6:~# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
0. c0t600144F0E196E8460000579DD8CE0003d0 <SUN-ZFS Storage 7330-1.0 cyl 323 alt 2 hd 254
   sec 254>
   /scsi_vhci/ssd@g600144f0e196e8460000579dd8ce0003
1. c1d0 <SUN-DiskImage-20GB cyl 27902 alt 2 hd 96 sec 768>
   /virtual-devices@100/channel-devices@200/disk@0
Specify disk (enter its number): ^C
root@opc-192-168-50-6:~#
```

In this example:

- Disk0 is c0t600144F0E196E8460000579DD8CE0003d0, which comes from the SUN-ZFS storage appliance.
- Disk1 is c1d0 and is the local boot disk.



Caution - Use extreme caution with this local boot disk. Damaging the data on this disk may render the whole compute instance unbootable, which would require you to open a service request (SR) with Oracle to repair/reinstall the instance.

For information about using Oracle Solaris Zones on shared storage, refer to [Getting Started With Oracle Solaris Zones on Shared Storage](#).

3. Determine the URI that can be used to provide this disk to a zone.

```
root@opc-192-168-50-6:~# suriadm lookup-uri /dev/dsk/
c0t600144F0E196E8460000579DD8CE0003d0
```

```
dev:dsk/c0t600144F0E196E8460000579DD8CE0003d0
iscsi://192.168.50.36/lunname.naa.600144f0e196e8460000579dd8ce0003
iscsi://192.168.50.36/target.iqn.1986-03.com.sun:02:303929e6-3928-4529-d4ae-
d2cb9a6922a8,lun.0
root@opc-192-168-50-6:~#
```

The list of available URIs that correspond to the volume includes several options, but the one that includes the volume GUID number (`iscsi://192.168.50.36/lunname.naa.600144f0e196e8460000579dd8ce0003`) is preferred.

4. Add the URI-described disk to the zone configuration.

The process of adding storage devices to native, legacy support (branded) or kernel zones does not modify the zones boot location it only adds additional storage on an iSCSI LUN. By default a dedicated ZFS filesystem is created in the Global Zone to host the zones root pool information, this filesystem is mounted at `/system/zones/zone-name` in the global zone. In order to have the zone boot from dedicated disks rather than space shared from the Global Zone, please follow the guidance in [“Configure an iSCSI LUN as a Native or Branded Zone Root Pool” on page 39](#) or [“Configure an iSCSI LUN as a Kernel Zone Root Pool” on page 40](#).

```
root@opc-192-168-50-6:~# zonecfg -z testzone
Use 'create' to begin configuring a new zone.
zonecfg:testzone> create
create: Using system default template 'SYSdefault'
zonecfg:testzone> add device
zonecfg:testzone:device> set storage=iscsi://192.168.50.36/lunname.naa.
600144f0e196e8460000579dd8ce0003
zonecfg:testzone:device> end
zonecfg:testzone> commit
zonecfg:testzone> exit
root@opc-192-168-50-6:~#
```

▼ Configure an iSCSI LUN as a Native or Branded Zone Root Pool

Complete these steps to configure a volume (LUN) as a native or branded root zone. See [Zone Brands](#) for a description of the different zone types.

- **Indicate a dedicated zpool for the zone installation that is managed by the zones framework.**

Use the `rootzpool` property, as follows:

```
root@opc-192-168-50-6:~# zonecfg -z testzone
```

```
Use 'create' to begin configuring a new zone.
zonecfg:testzone> create
create: Using system default template 'SYSdefault'
zonecfg:testzone> add rootzpool
zonecfg:testzone:rootzpool> set storage=iscsi://192.168.50.36/luname.naa.
600144f0e196e8460000579dd8ce0003
zonecfg:testzone:rootzpool> end
zonecfg:testzone> commit
zonecfg:testzone> exit
root@opc-192-168-50-6:~#
```

Note - This value of the storage property must be set to point to the URI path for the iSCSI LUN discovered in [“Prepare a Volume for Use With Zones” on page 37](#).

▼ Configure an iSCSI LUN as a Kernel Zone Root Pool

Complete these steps to configure a volume (LUN) as a kernel root zone. See [Zone Brands](#) for a description of the different zone types.

● Create a kernel zone.

Update the system-provided default device storage property (which defaults to a local zvol) to your new URI as follows:

```
root@opc-192-168-50-6:~# zonecfg -z kz
Use 'create' to begin configuring a new zone.
zonecfg:kz> create -t SYSsolaris-kz
zonecfg:kz> info device
device 0:
  storage.template: dev:/dev/zvol/dsk/{global-rootzpool}/VARSHARE/zones/{zonename}/disk
  %id}
  storage: dev:/dev/zvol/dsk/rpool/VARSHARE/zones/kz/disk0
  id: 0
  bootpri: 0
zonecfg:kz> select device 0
zonecfg:kz:device> set storage=iscsi://192.168.50.36/luname.naa.
600144f0e196e8460000579dd8ce0003
zonecfg:kz:device> end
zonecfg:kz> info device
device 0:
  storage: iscsi://192.168.50.36/luname.naa.600144f0e196e8460000579dd8ce0003
  id: 0
  bootpri: 0
```



```
zonecfg:kz> exit  
root@opc-192-168-50-6:~#
```

Note - See [“Prepare a Volume for Use With Zones” on page 37](#) for information about identifying the URI.

▼ Create a Volume Snapshot

You can use ZFSSA to take instantaneous snapshots of volumes (and projects). To take the fullest advantage of this feature, it is recommended that you group related LUNs into projects so you can take a consolidated snapshot of all related LUNs with a single snapshot.

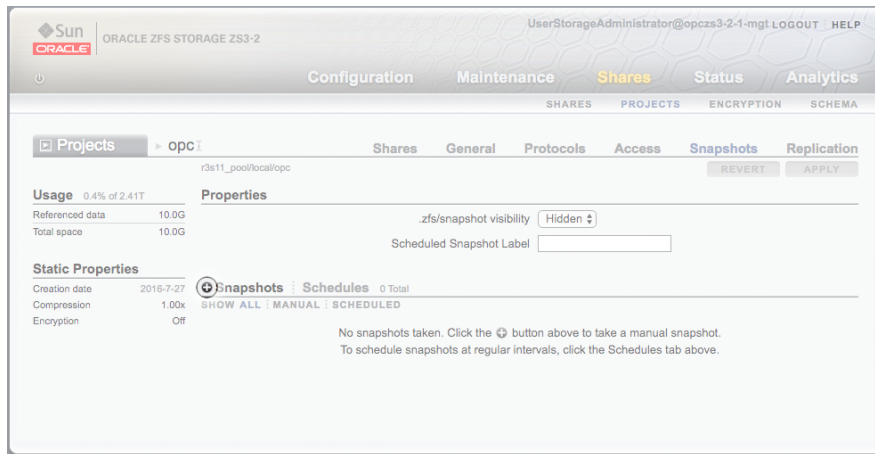
Note - Snapshots will consume space from the available storage capacity of the ZFSSA.

Complete these steps to create a snapshot:

1. **Log in to ZFSSA using the `opc` user account.**
2. **Navigate from the Shares > Projects screen of the BUI, and select the project of which you want to take a snapshot.**

In this example, the project name is `opc`.

3. Select the Snapshots tab and click the (+) button next to the Snapshots label, as shown here.



4. Enter the name for the snapshot in the window that appears and click OK.

▼ Increase the Size of a Volume

To resize a volume (LUN), modify the settings on ZFSSA.



Caution - Modifying the LUN size (especially reducing its size) has the potential to destroy data on the LUN. Before resizing any LUN that is already in use, ensure that you are aware of the potential risks and can proceed without impacting data.

1. **Log into ZFSSA using the `opc` user account.**
2. **Select the project that contains the volume (LUN).**
Navigate from the Shares > Projects screen of the BUI, and select the project that contains the LUN you want to resize (`opc`).
3. **Select the volume (LUN) to resize.**
From the LUNs screen on the Shares tab, click the edit (pencil) button next to the LUN you want to resize.

4. **Change the Volume Size setting and click Apply.**

▼ Detach a Volume

Complete these steps to detach a volume:

1. **Log in to ZFSSA using the `opc` user account.**
2. **Navigate from the Configuration > SAN screen of the BUI.**
3. **Detach the volume (LUN).**
Select the iSCSI tab and from the Initiators Group list, uncheck the box next to the instance name.
4. **Click Apply.**

Note - The presence of the volume in the compute instance is cached, but the ability to read or write data on the volume is removed. The volume will disappear completely on the next reboot of the compute instance.

▼ Destroy a Volume

You can permanently destroy a volume by deleting it on the ZFSSA. Complete these steps to destroy a volume:

1. **Log in to ZFSSA using the `opc` user account.**
2. **Select the project that contains the volume (LUN).**
Navigate from the Shares > Projects screen of the BUI, and select the project that contains the volume you want to destroy (`opc`, in this example).
3. **Select the volume (LUN) to destroy.**
Select the project that contains the volume (LUN). From the LUNs screen on the Shares tab, click the delete (trash bin) button next to the volume you want to destroy and click Apply.
4. **Select the project that contains the snapshot.**

Navigate from the Shares > Projects screen of the BUI, and select the project that contains the snapshot of this volume (opc, in this example).

5. Select the snapshot to delete.

From the Snapshots tab, click the delete (trash bin) button next to the Snapshots you want to delete.

Managing Storage NFS Shares

You can create storage shares on the ZFS Storage Appliance component of the SPARC Model 300 Service and configure them to be accessible by one or more of the compute instances. You can administer this process using either the CLI or BUI tools. Use whichever method best addresses your needs. See [“Step 4: Learn How to Access the CLIs and BUI” on page 12](#) for information.

When you create a storage share, you should carry out at least the following operations:

- [“Create a Share” on page 44](#)
- [“Create a Network Exception” on page 46](#)

▼ Create a Share

Perform these steps to create an NFS Share using the Oracle ZFS Storage Appliance browser user interface (ZFSSA BUI).

- 1. On a web UI, use the `opc` user account to log into the ZFSSA head where you want to create the NFS share.**

You can access the BUI at `https://zfssa_head_IP:215`

- 2. Identify and make note of the ZFSSA network interface information for the share you want to create.**

NFS shares are provided over network interfaces. When configuring a share, you must configure the appropriate network interface over which the share will be served. The Oracle SPARC Model 300 Service provides storage access from the ZFSSA over network devices `vnic5` (corresponding to head 1) and `vnic7` (corresponding to head 2). You must identify which of these devices you wish to serve the NFS share when configuring the volume.

Complete these steps to identify the network interface information for the share:

- a. **Navigate from the Configuration > Network screen of the BUI, then select the Configuration tab.**

- b. **Locate the network for the volume from the Interfaces list.**

The network you choose will vary depending on which of the two ZFSSA heads is used for the volume (LUN).

- Head 1 is associated with the interface named opczs3-2-1-store-aggr.
- Head 2 is associated with the interface named opczs3-2-2-store-aggr.

- c. **Make note of the following information and keep it accessible for upcoming steps.**

In this example, head 1 is used for the volume, so these values are noted:

- Interface name: opczs3-2-1-store-aggr
- VNIC short name: vnic5
- IP address: 192.168.50.36

The screenshot shows the Oracle ZFS Storage BUI Configuration > Network screen. The page has a top navigation bar with tabs: Configuration, Maintenance, Shares, Status, and Analytics. Below this is a sub-navigation bar with tabs: SERVICES, STORAGE, NETWORK, SAN, CLUSTER, USERS, PREFERENCES, and ALERTS. The main content area is titled 'Network' and has sub-tabs: Configuration, Addresses, and Routing. The Configuration tab is active, showing a table of network interfaces. The table has columns for Devices, Datalinks, and Interfaces. The Interfaces column lists several interfaces, including opczs3-2-1-cmgt-aggr, opczs3-2-1-mgt, opczs3-2-1-p1, opczs3-2-1-p2, opczs3-2-1-store-aggr, opczs3-2-2-cmgt-aggr, and opczs3-2-2-store-aggr. The opczs3-2-1-store-aggr interface is highlighted with a red box, showing its IP address as 192.168.50.36/29, via vnic5.

3. **Create an NFS share.**

NFS shares are created in the Shares page of the ZFSSA WebUI. A simple share creation is shown in the following steps. Detailed configuration is beyond the scope of this document. Refer to the [ZFS Storage Appliance Administration Guide](#) for detailed configuration information.

- a. **Navigate to the Shares > Shares screen of the BUI, then click on the + symbol next to the Filesystems heading to begin creating a new Filesystem Share.**
- b. **In the Pop-up window, configure at least these options:**
 - i. **Select the appropriate project where the share will reside.**
 - ii. **Define a name for the share.**
 - iii. **Select appropriate permissions and mountpoints.**
 - iv. **Apply any other options based on your use case.**

▼ Create a Network Exception

To limit the visibility of an NFS share, it is strongly advised to create appropriate network exceptions that apply to projects or individual shares to limit the clients that are able to mount or access the share. Network exceptions are defined based on IP addresses, Fully Qualified hostnames, or Subnets, and indicate that, for a specific share or project, only those clients fitting the exceptions definition can see or use the NFS share.

1. **Select a project or a file system, based on how you want the exception applied:**

Navigate to one of these options:

 - Project – Allows all shares within the project to inherit the same exception configuration.
 - Share – Configures a specific exception for just that share.
2. **Double click the Share or Project from the list in the window to start configuring the exception.**
3. **Navigate to the Protocols tab to start configuring an exception.**

4. **If applying an exception to one share and not all of a projects shares, you must uncheck the Inherit from Project box at the top of a shares Protocol page to enable editing of a specific exception for a share.**
5. **In the NFS field of the Protocols screen, set the Share Mode to None.**

This prevents the NFS share from being mounted by any client and allows only access to clients that match the exceptions that are being created.
6. **Click on the + symbol next to NFS Exceptions to create a new exception.**

Select a TYPE from the drop down menu option:

 - a. **Host(FQDN) or Netgroup – Enter an Entity IP address or fully Qualified Domain Name hostname.**
 - b. **DNS Domain – Enter a DNS Domain that will be able to access the share.**
 - c. **Network – Specify a network subnet in CIDR format (aaa.bbb.ccc.ddd/nn) that will be able to access the share**
7. **Complete the Exception information by choosing a suitable Share Mode, Charset, and whether Root Access is required.**
8. **Multiple exceptions can be added if required.**
9. **Click Apply to check and apply the exception.**

Configuring Networking

Read these topics to learn how to configure networking for Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300.

- [“Learn About Network Settings” on page 49](#)
- [“Service Port Details” on page 50](#)
- [“Reserve and Manage Public IP Addresses” on page 51](#)
- [“Attach Additional Network Devices to Compute Instances” on page 52](#)

Related Information

- [“Managing Instances” on page 17](#)
- [“Managing Storage” on page 25](#)
- cloud.oracle.com

Learn About Network Settings

By default, all instances in Oracle SPARC Model 300 Service can communicate with each other on the subnet that was allocated during the initial registration and deployment. Any traffic on the default VLAN (VLAN 100) can be routed outside of the service and can be accessible to the data center through a VPN service. All VLAN tags are removed from traffic leaving the Oracle SPARC Model 300 network, and all traffic entering the Oracle SPARC Model 300 network are tagged with VLAN 100. It is not possible to extend your VLAN into the Oracle SPARC Model 300 network from your data center.

You can further isolate network objects within the rack by using VLANs 101 through 199, which can be configured by the instance root role.

Storage in Oracle SPARC Model 300 Service can only be accessed by instances within the service.

Each compute instance (DomU) is preconfigured with a 2x10GbE LACP port aggregate datalink (aggr1), which is created using the net0 and net2 interfaces. This aggregate provides

a high-bandwidth redundant link on top of which you can configure virtual network devices (VNICS). By default, one VNIC (aggr1_vnic1) is preconfigured on the LACP aggregate aggr1 and is assigned to VLAN 100 to allow immediate connectivity to the Oracle SPARC Model 300 DomU. You can configure additional VNICS using this datalink, as needed. To create a VNIC link called aggr1_vnic2, for example, that is associated with VLAN 100 on the aggregate aggr1, run the following command as root user.

```
root@opc-192-168-50-6:~# dladm create-vnic -l aggr1 -v 100 aggr1_vnic2
```

where:

- -l selects the link aggregate on which the VNIC will be created.
- -v assigns the VNIC with a tagged VLAN ID.

VLANs 101 through 199 can be used for network isolation within the Oracle SPARC Model 300 and will not propagate out of the Oracle SPARC Model 300 network. In this case, you can create a VNIC using a command similar to the one in the preceding example, but you will use a tagged VLAN ID of 101 through 199 instead of VLAN ID 100.

Two additional 10GbE ports, net1 and net3, are available for use and are connected to the Oracle SPARC Model 300 networking, but are not configured by default. To use these interfaces, you must configure the ports as an additional LACP port aggregate. To create this aggregate, run the following command as root user.

```
root@opc-192-168-50-6:~# dladm create-aggr -m trunk -L active -P L3,L4 -l net1 -l net3 aggr2
```

where:

- -m sets the appropriate LACP aggregate mode.
- -L sets the LACP mode.
- -P sets the correct trunk aggregation policy.
- -l selects the network ports that will be added to the aggregate.

Any change to this command, other than to the name of the aggregate, is unsupported. The Oracle SPARC Model 300 network configuration expects an aggregate with specific characteristics, and changing them can produce unknown network performance or operation of the new aggregate. Do not use the net1 and net3 ports as independent ports due to existing SPARC Model 300 network requirements.

Service Port Details

Ports for services (like https and ssh) are required between network endpoints between your data center, the Oracle SPARC Model 300 Service, and the Internet. To ensure that you can

connect to your Oracle SPARC Model 300 Service from your data center, while maintaining a secure connection to the outside world, the service ports have been configured in the following manner, which can be modified through a request to Oracle Support.

- For connections between your data center and the Oracle SPARC Model 300 Service through the VPN service, all ports are set by default to permit, allowing all open communication. To restrict this port access, open a service request (SR) with Oracle Support, or deny the port on the customer VPN side. For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).
- For connections from the Oracle SPARC Model 300 Service to anything on the public IP space, all ports are set by default to deny, preventing any service port to connect across this public connection. To open port access, submit an SR with Oracle Support.

▼ Reserve and Manage Public IP Addresses

If you want to enable access to your instance over the public Internet, you must associate a public IP address with your instance. You can associate either a temporary or a persistent public IP address with an instance when you create the instance.

Temporary public IP addresses are assigned dynamically from a pool of public IP addresses. When you associate a temporary public IP address with an instance, if the instance is restarted or is deleted and created again later, its public IP address might change. If you want to assign a persistent public IP address to your instance, you must first create an IP reservation, and then associate the IP reservation with the instance.

By default, a /29 (8 IP addresses) for outbound source NAT and a /28 (16 IP addresses) for inbound NAT are available for each SPARC Model 300 Service. The /29 outbound IPs are a NAT pool for dynamic NAT of the instances to other Oracle Cloud services or to other public addresses. All network traffic from the compute instances on the outbound NAT are temporary public IPs.

The /28 inbound IPs use static mapping/persistent to instances or zones. You need to request these IPs (and additional public IP addresses) by submitting a service request (SR). Typically, hosts that require public facing services will need these.

● Request the required IP addresses.

For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

▼ Attach Additional Network Devices to Compute Instances

By default, each compute instance with a 20Gbps redundant LACP aggregate network device is connected on VLAN 100 (the default VLAN). You can map this device to your datacenter.

Oracle provides two additional network devices, which are not configured, but are attached to separate network paths so you can configure additional redundant network devices.

Complete these steps to add these devices.

1. **Identify the available devices in the compute instance.**
 - a. **Log in to the CLI with root privileges.**
 - b. **Run the following commands to display information for the additional network devices.**

```
root@opc-192-168-50-6:~# dladm show-phys
LINK  MEDIA    STATE  SPEED  DUPLEX  DEVICE
net0   Ethernet  up      10000  full    i40e0
net1   Ethernet  up      10000  full    i40e1
net3   Ethernet  up      10000  full    i40e3
net2   Ethernet  up      10000  full    i40e2

root@opc-192-168-50-6:~# dladm show-aggr -L
LINK  PORT  AGGREGATABLE  SYNC  COLL  DIST  DEFAULTED  EXPIRED
aggr1 net0   yes           yes   yes   yes   no         no
--    net2   yes           yes   yes   yes   no         no

root@opc-192-168-50-6:~# ipadm
NAME          CLASS/TYPE  STATE  UNDER  ADDR
aggr1_vnic1   ip          ok     --      --
aggr1_vnic1/v4static static    ok     --      192.168.50.6/28
root@opc-192-168-50-6:~#
```

From the preceding commands, you can see that four native network ports (net0, net1, net2, and net3) are provided, and that an LACP aggregate device (aggr1) had been created using ports net0 and net2. The aggregate port is then used by one virtual network interface (aggr1_vnic1), which holds the IP address of the instance, and through which the instance CLI is accessed.

2. **Follow these guidelines when you configure additional redundant network devices.**

- Place additional VNICs on top of the aggr1 device.
- Specify IP, VLAN, and MTU settings, as required (within OPC restrictions).
- Use VLANs 101-199.
- Use MTU sizes up to 9000.



Caution - Do not modify the Oracle-provided Aggregation device `aggr1` or the corresponding IP device `aggr1_vnic1` that is provided on top of it. Doing so can render the compute instance inaccessible and will require you to open a service request (SR) for Oracle to repair the broken network configuration.

Refer to [Strategies for Network Administration in Oracle Solaris 11.3](#) for procedures to configure the additional networks.

Getting Answers to Your Questions

To get answers to some of your questions about the Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300, refer to the [SPARC Model 300 Service FAQ](#). Then, read these topics for best practice and troubleshooting recommendations for using the service.

- [“What's New” on page 56](#)
- [“Best Practices” on page 56](#)
- [“Troubleshoot the Service” on page 59](#)

Related Information

- [“Managing Instances” on page 17](#)
- [“Managing Storage” on page 25](#)
- [“Configuring Networking” on page 49](#)
- [cloud.oracle.com](#)

What's New

Learn about the new and changed features for the Oracle Compute Cloud Service - Dedicated Compute Capacity - SPARC Model 300.

Date	Feature/Change	Description
April 2018	NFS shares on the ZFS storage appliance	The opc account that Oracle provides on the ZFS Storage Appliance supports NFS share creation, and the ability to modify SNMP services. In addition, some limited administration capabilities are also allowed. See “Managing Storage” on page 25 .
April 2018	Oracle Cloud Infrastructure Object Storage Classic	This storage feature can now be used to backup and restore data from the SPARC Model 300. Data can be uploaded and downloaded to the Oracle Cloud Infrastructure Object Storage Classic service.

Best Practices

Consider these recommendations when configuring your service:

- [“General Best Practices” on page 56](#)
- [“Instance Best Practices” on page 57](#)
- [“Storage Best Practices” on page 58](#)
- [“Performance Best Practices” on page 58](#)

General Best Practices

- Limit the amount of custom analytics reports that are configured and running on the ZFSSA to prevent overloading the storage device with analytic query processing rather than provisioning storage devices. This will ensure the best performance from the ZFSSA.
- Do not modify or destroy the default `aggr1` or `aggr1_vnic1` network datalink and interface provided in your Oracle-installed compute instance. These devices provide the only initial

means by which you can access an instance, and modifying or deleting them is likely to cause a compute instance to become inaccessible, which would require you to open a service request (SR) to Oracle to recover the network configuration.

Instance Best Practices

- Do not modify `/etc/system` settings without specific recommendations from Oracle. Each instance has been pre-configured to limit the ZFS arc caches based on current best practices.
- Do not delete the primary interface (data link or first VNIC interface) or modify the IP address that your instance was delivered with. In doing so, you risk losing connectivity to your instance, which would require you to file a service request (SR) to restore connectivity.
- Do not use the shutdown command, `init 0`, `1`, `2`, `3`, `4`, `5`, `S`, or take any other action that would prevent the system booting into multi-user mode. Doing so will cause your instance to become inaccessible because there is no console access. Remember to also factor in the CPU and memory resources that are necessary for the operating system if using kernel zones or legacy support zones.

If your instance is (or becomes) unavailable, submit a service request (SR) with Oracle Support. For information about opening an SR, see [“Step 5: Learn About Opening Service Requests” on page 13](#) and [Cloud Support Help: Service Requests](#).

- It is highly recommended that you make backups of each instance for your service. In the case of a catastrophic error that requires the instance to be rebuilt, the instance is re-created to its initial state, which is the same state it was when you first accessed the service. Any changes to the instance will be lost.
- Oracle recommends that you use ZFSSA iSCSI LUN storage volumes instead of using local storage.
- Change any default passwords that the system is delivered with.
- To ensure that Oracle SPARC Model 300 Service instances provide a resilient platform for your workloads, ensure that the latest security patches are applied to the operating system running on the instances and zones. In addition, before deploying applications on an instance, review the security configuration of the operating system and verify that it complies with your security policies and standards. For security and patching-related guidelines, see the documentation for your operating system (for example, refer to [Administering CVE Updates: Oracle Solaris 11.3 Security Compliance Guide](#) and [Oracle Solaris 11 Security and Hardening Guidelines](#)).

When you create zones and storage volumes, select the name of the object carefully. Pick a name that helps you quickly identify the key characteristics of the object later. For example, when creating a bootable storage volume, consider including the operating system name and the image disk size in the name of the storage volume.

- While configuring the size for a zone, consider the nature of the applications that you plan to deploy on the instance, the number of users that you expect to use the applications, and

also how you expect the load to scale in the future. Remember to also factor in the CPU and memory resources that are necessary for the operating system if using kernel zones or branded zones.

Configure the zone that meets the requirements of your workload with a sufficient buffer for intermittent spikes in the load. If you're not sure what shape is appropriate for an instance, then start small, experiment with a representative workload, and then settle on a zone configuration. This approach may help you achieve an optimal trade-off between resource allocation and performance.

- If you're creating an Oracle Solaris instance, try to determine, up front, how many users you expect to access the instance and plan for a separate SSH key pair for each user.
- Keep your SSH keys secure. Lay down policies to ensure that the keys aren't lost or compromised when employees leave the organization or move to other departments. If you lose your private key, then you can't access your instances. For business continuity, ensure that the SSH keys of at least two IT system administrators are added to your instances.

Storage Best Practices

- Limit the number of volumes (iSCSI LUNs) provided by each ZFSSA head to 500. This is not a hard restriction, but doing so will limit extended failover times seen by users in the event of a failure of one or other of the ZFSSA heads, and the resultant take over by the remaining head. This will help prevent unintended application failures that might result from extended storage wait times when trying to access storage that is being transitioned from the failed head to the remaining head.
- Storage volumes are iSCSI LUNs. Typically, you should mount them to the instance first, and then map them to the zones that require the dedicated or shared storage volume.

Performance Best Practices

You might be able to improve the transfer speeds, and thus the performance, of applications running on Oracle Solaris 11.3 by switching the mediator implementation you use to OpenSSH. For more information about using OpenSSH with Oracle Solaris 11.3, refer to [“How to Use the OpenSSH Implementation of Secure Shell” in *Managing Secure Shell Access in Oracle Solaris 11.3*](#).

Note - In some cases, where your Oracle Solaris 11.3 application depends on SunSSH-specific features (for example, x590 support or OpenSSL engine), switching to OpenSSH is not a viable option. In these cases, you might see some performance gain by increasing the value for the `tcp_recv_hiwat` tunable to 64MB.

Troubleshoot the Service

This section describes potential problems that you might encounter when using Oracle Compute Cloud Service and directs you to information to solve them. These include problems with:

- For help troubleshooting network configuration issues (for example, administration, naming services, and other network-related issues) in the Oracle Solaris OS, refer to [Troubleshooting Network Administration Issues in Oracle Solaris 11.3](#).
- For help troubleshooting Oracle Solaris system-related issues (for example, system crashes, hangs, reboot failures, and the like), refer to [Troubleshooting System Administration Issues in Oracle Solaris 11.3](#).

If these resources do not help resolve the issues you are facing, or for any other issues relating to storage or networking, open a service request (SR) with Oracle Support. See “[Step 5: Learn About Opening Service Requests](#)” on page 13 and [Cloud Support Help: Service Requests](#) for more information about opening an SR.

