

Installing Oracle GoldenGate for HP NonStop (Guardian)

After you decide how you want your data management environment to operate, you are ready to design its infrastructure and install Oracle GoldenGate for HP NonStop. This product is for the support of Enscribe and SQL/MP, it does not support SQL/MX. This document describes how to install Oracle GoldenGate data replication software for the HP NonStop platform using the Guardian operating system.

This chapter includes the following topics:

- [Overview to Installing Oracle GoldenGate for HP NonStop](#)
- [Preparing your Production Environment](#)
- [Preparing the Source and Target](#)
- [Documentation Accessibility](#)

Overview to Installing Oracle GoldenGate for HP NonStop

When you install Oracle GoldenGate for HP NonStop, you install Oracle GoldenGate components such as:

- Audserv
- GGSCI
- Manager
- Logger
- Extract
- Collector
- Replicat
- Coordinator
- Syncfile
- DEFGEN

- DDLGEN
- Supporting files, such as help

The install program auto-detects your Guardian operating system version and installs the native version of Oracle GoldenGate code where applicable. You can also tell Oracle GoldenGate to install TNS code during the installation routine. TNS objects are only required if you use TNS User Exits, otherwise its recommended not to install TNS code.

 **Note:**

HP NonStop Server installations use the `UNPAK` utility. Verify that you have a copy of this utility before proceeding.

- [Downloading Oracle GoldenGate for HP NonStop](#)
- [Installing Oracle GoldenGate for HP NonStop](#)

Downloading Oracle GoldenGate for HP NonStop

Go to the Oracle support site to look for the latest download of the release you want.

1. Navigate to <http://support.oracle.com>.
2. Log in with your Oracle ID and password.
3. Select the Patches and Upgrades tab.
4. On the Search tab, click Product or Family.
5. In the Product Field, enter **Oracle GoldenGate**.
6. From the Release drop-down list, select the release version that you want o download.
7. Make sure Platform is displayed as the default in the next field, and then select the platform from the drop-down list.
8. Leave the last field blank.
9. Click **Search**.
10. In the Advanced Patch Search Results list, select the available builds that satisfy the criteria that you supplied.
11. In the file Download dialog box, click the ZIP file to begin the download.

If patches are not available on the support site, go to the Oracle delivery site for the release download.

1. Navigate to <https://edelivery.oracle.com/>.
2. Sign in with your Oracle ID and password.
3. On the Terms and Restrictions page:
 - Accept the **Trial License Agreement** (even if you have a permanent license).

- Accept the **Export Restrictions**.
 - Click **Continue**.
4. On the Media Pack Search page:
 - Select the **Oracle Fusion Middleware** Product Pack.
 - Select the platform on which you will be installing the software.
 - Click **Go**.
 5. In the Results List:
 - Select the Oracle GoldenGate Media Pack that you want.
 - Click **Continue**.
 6. On the Download page:
 - View the `Readme` file.
 - Click **Download** for each component that you want. Follow the automatic download process to transfer the file to your system.

Installing Oracle GoldenGate for HP NonStop

Follow these steps to use the downloaded file to install Oracle GoldenGate for HP NonStop.

1. Save the downloaded file to a directory on your workstation and unzip it to a temporary folder.
2. Upload the `GGSPAK` and `GGSPAK_versionPK` files in *binary* mode to the volume and subvolume on your HP NonStop Server where you want to install and run Oracle GoldenGate. We recommend that you preface the subvolume name with `GG` and include the release of Oracle GoldenGate. For example, if you are installing Oracle GoldenGate release 12.2.0.1.4 then we recommend you install on `$VOL.GGS122014`.

Note:

If you are using Oracle GoldenGate against a TMF-enabled (or audited) database, then Oracle GoldenGate must be installed on a volume that is TMF-protected because the internal Oracle GoldenGate file must be audited. The product can only be installed on volume names of 6 digits or less.

3. Alter `GGSPAK` to be an Edit file.

```
TACL> FUP ALTER GGSUNPAK, CODE 101
```


4. Run `GGSPAK` and confirm the installation location. If the correct location is not displayed, you are prompted for it, as shown here.

```
TACL 2> run ggsunpak
Installing Oracle GoldenGate at $DATA4.CPSTEMP
```

```
Is this correct?(Y/N) N
Enter the volume and subvolume to install Oracle GoldenGate: $VOL.GGS122014
Installing Oracle GoldenGate at $VOL.GGS122014
Is this correct?(Y/N) Y
```

- 5. Oracle GoldenGate recommends using the Native versions of these components, but if you have TNS mode user exits, enter Y at the prompts to install the TNS versions.

```
Since you are installing for J06.20.00, we recommend installing Native mode
obje cts
Do you use TNS mode User Exits in Extract (Y/N)? N
Do you use TNS mode User Exits in Replicat (Y/N)? N
```

 **Note:**
For the purposes of this documentation, it is assumed that Native mode objects are installed.

- 6. You are asked whether to build Oracle GoldenGate intercept libraries. Oracle GoldenGate recommends building a GGSLIB object, but it can be omitted if you don't need to capture non-audited Enscribe files or FUP-related activities. (You can build a GGSLIB object later.)

```
*****
*   If you need the TNS mode GGS Intercept       *
*   Library GGSLIB then answer Yes,            *
*   Otherwise you can skip this step.          *
*****
Build of new GGSLIB recommended. Build now (y/n)? Y
```

- 7. The BUILDMAC utility runs and searches for run-time libraries to bind with GGSLIB, and you may be prompted that one or more libraries is missing. Enter Y to continue if that occurs.

```
*** Running BUILDMAC to build GGSLIB ***
*****
Looking for runtime libraries to BIND with GGSLIB
*****
Adding $SYSTEM.ZCOBOLRT.CLIBOBJ (COBOL74)
Adding $SYSTEM.ZCOB85RT.C8LIB (COBOL85)
Adding $SYSTEM.SYSTEM.CRELIB (CRE)
Adding $SYSTEM.SYSTEM.COBOLLIB (COBOL85)
One or more libraries was missing. Continue anyway? (Y/N) Y
```

- 8. You are asked if you want to include a user library. Respond either Y or N.

```
Do you want to include your own User Library (y/n)? N
```

- 9. You are asked if you want to change the location of the AUDCFG.

```
Do you want to change the location for the AUDCFG segment (Y/N): N
```

If you respond with yes, it prompts you for the new default location (\$vol.subvol) of the AUDCFG segment.

Assuming you opted to build the intercept libraries, you will see the BIND steps being performed. You may see the following Binder messages that can be safely ignored. They are:

```
WARNING 60: Parameter mode mismatch on PXS_FILE_WRITE_parameter2. Parameter
mismatch between two external procedure declarations. The parameters are
declared as STRING and OTHER.
WARNING 61: Parameter type mismatch on PXFS_FILE_WRITE_parameter2. Parameter
mismatch between two external procedure declarations. The parameters are
declared as INT32 and OTHER.
WARNING 61 Parameter type mismatch on PXFS_FILE_WRITE_parameter3.
```

Other warnings may also be displayed and can be ignored.

10. You are asked if you want to build the native mode Oracle GoldenGate intercept libraries GGSSRL/GGSDLL. To install native code, enter Y and continue to the next step; otherwise, enter N to use TNS code and then skip to step 17.

```
*****
*   If you need the Native mode GGS Intercept   *
*   Library GGSSRL/GGSLIBR then answer Yes,    *
*   Otherwise you can skip this step.          *
*****
STOPPED: 2,33
CPU time: 0:00:00.945
1: Process terminated with warning diagnostics
STOPPED: 2,223
CPU time: 0:01:18.782
1: Process terminated with warning diagnostics
Build Native mode GGSLIBR & GGSSRL (Y/N?) Y
```

11. The NLDLIB macro runs, and you are asked if you want to include a user library. Reply either Y or N.

```
*** Running NLDLIB to build GGSLIBR & GGSSRL ***
-----
NLDLIB builds the native relinkable libraries GGSLIBR and GGSSRL
A Native mode GGSLIBR needs the re-linkable versions of some
Tandem SRLs. ZCREREL and ZCOBREL. If Native cobol support is not
required then ZCOBREL may be omitted.
If Native Cobol support is required then you have to compile the
Cobol program and search in GGSLIBR
A Native mode GGSSRL will be built for a User Library for Dynamic
linking
Enter X at any prompt to EXIT
-----
Do you want to include your own User Library (Y/N) : N
```

12. You are asked if you want to change the location of the AUDCFG.

```
Do you want to change the location for the AUDCFG segment(Y/N) : N
```

If you respond with yes, it prompts you for the new default location (\$vol.subvol) of the AUDCFG segment.

13. You are asked if you want to include ZCREREL. It is recommended to enter Y to include it unless you know your applications do not use these libraries.

```
Do you want to include ZCREREL (Y/N) : Y
```

NLDLIB searches for ZCREREL locations and displays them.

```
Looking for ZCREREL ..  
$SYSTEM.SYS02.ZCREREL $SYSTEM.SYS03.ZCREREL $SYSTEM.SYS04.ZCREREL  
$DATA09.R1269G09.ZCREREL $DATA09.R8431G09.ZCREREL
```

14. You are asked to enter the location of the ZCREREL to use.

```
Enter location of ZCREREL : $SYSTEM.SYS03.ZCREREL
```

15. You are asked if you want to include ZCOBREL. It is recommended to enter 3 to include it unless you know your applications do not use these libraries.

```
Do you want to include ZCOBREL (Y/N) : Y
```

NLDLIB searches for ZCOBREL locations and displays them.

```
Looking for ZCOBREL ..  
$SYSTEM.SYS02.ZCOBREL $SYSTEM.SYS03.ZCOBREL $SYSTEM.SYS04.ZCOBREL  
$DATA09.R8107G09.ZCOBREL $DATA09.R8108D46.ZCOBREL
```

16. You are asked to enter the location of the ZCOBREL to use. If you plan to use COBOL, enter a location; otherwise, you can press **Enter** to bypass this prompt.

```
Enter location of ZCOBREL : $SYSTEM.SYS03.ZCOBREL
```

 **Note:**

If you omit a ZCOBREL location, you are prompted to confirm your choice:

```
Omit ZCOBREL (y/n):Y
```

NLDLIB builds the GGSLIBR and GGSSRL libraries, displaying a series of informational messages and the names of the files that were created.

17. You are asked if Oracle GoldenGate should use an existing SQL catalog or create a new catalog. Either enter X to specify no catalog or enter the name of a catalog to use. If you enter a catalog name and it does not exist, you are asked whether to create it.

```
SQL Catalog for Compilation (X for no catalog)? $VOL GGSCAT  
Catalog $VOL.GGSCAT does not exist. Create it (Y/N)? Y
```

The SQL compilation completes, and then libraries and programs are linked. This happens automatically and does not require action from you.

18. You are asked if you want to migrate information from another Oracle GoldenGate environment. If you are installing Oracle GoldenGate for the first time, enter N. A migration requires shutting down existing Oracle GoldenGate processes and, if Logger is being used, the user application. (The following assumes a new installation.)

```
Would you like to Migrate information  
from another GGS Environment (Y/N)? N
```

19. You are advised to license your programs as shown in the following prompt, and then the installation utility finishes.

Make sure to secure files appropriately. If EXTRACT is to be executed outside of the SUPER group, you must PROGID and LICENSE the AUDSERV program for SUPER.SUPER authority. PRIVLIB must also be licensed if executing Replicat outside SUPER group.

Audserv/LOGDUMP licensing only applies to TMF installations. TMFARUL2 must be licensed

The Native Audserv must be licensed (A TMFARUL2 restriction)

The Native Logdump must be licensed if you are going to use it to read TMF trails (A TMFARUL2 restriction)

```
TACL> LOGON SUPER.SUPER
```

```
TACL> FUP
```

```
-GIVE Audserv, SUPER.SUPER
```

```
-SECURE Audserv, "NUNU", PROGID
```

```
-LICENSE Audserv
```

```
-LICENSE PRIV*
```

```
-LICENSE LOGDUMP
```

```
-LICENSE TMFARUL2
```

```
Installation Complete.
```

Licensing must be run by SUPER.SUPER because the following privileged access is required:

- Audserv and TMFARUL2 access the TMF trails through audit read routines.
- PRIVLIB contains privileged Replicat code that runs certain privileged FUP DUP or FUP LOAD operations.
- Logdump provides a version that can read the audit read routines. If you do not want Logdump to be able to read the audit trails, do not license it.

 **Note:**

If the version of your SQL catalog does not allow FUP to alter the program, you can PROGID and license the programs by using the following commands to license the components shown.

```
TACL> LOGON SUPER.SUPER
```

```
TACL> SQLCI
```

```
>> ALTER PROGRAM AUDSERV OWNER 255,255;
```

```
>> ALTER PROGRAM AUDSERV SECURE "NUNU" PROGID;
```

```
>> EXIT
```

```
TACL> FUP
```

```
- LICENSE (AUDSERV, PRIV*, LOGDUMP, TMFARUL2)
```

```
- EXIT
```

20. Once your product is installed, log on as the owner of your Oracle GoldenGate environment.

21. (Optional) You can purge the *PK files, if desired.

Preparing your Production Environment

This section outlines considerations that impact the overall architecture of your particular Oracle GoldenGate environment. These considerations include:

- Sizing
- CPU requirements
- Disk space requirements
- Data communications requirements
- [Sizing](#)
- [CPU Requirements](#)
- [Disk Space Requirements](#)
- [Data Communications Requirements](#)

Sizing

Before you begin running Oracle GoldenGate in your production environment, you should consider several factors regarding the size of your replication environment, including the following statistics.

- **Counts of inserts, updates and deletes generated on a per-file basis.** A common reference point is the application's transaction rate, from which these numbers can often be estimated.
- **The byte generation rate caused by inserts, updates and deletes.** Use this statistic to estimate disk space requirements and communications bandwidth. Byte generation totals during peak periods and over an extended time frame are also important.
- **For update-intensive applications, the number of bytes in a record that change during an update.** This is probably the most difficult statistic to extract (especially for Enscribe files) but can have a dramatic impact on bandwidth requirements. The number of bytes that change in a record during an update determines the compression rate for the record.
- **The cache hit rate and file busy percentages of the most heavily updated files.** This has a direct bearing on the throughput possible by individual Replicat processes on the target system.
- **Which programs are causing the heaviest amount of update activity.** It is often possible to avoid replicating certain types of activity (such as `FUP LOAD`) by excluding those programs from replication. Measuring individual programs is a concern primarily for non-TMF data processing.
- [Reducing Overall Requirements](#)

Reducing Overall Requirements

In some cases, due to bandwidth or application constraints, it can be desirable to eliminate certain files, programs or processes from Oracle GoldenGate.

To get a view of overall replication activity excluding certain files or programs, use the `EXCLUDEPROGRAM` parameters.

CPU Requirements

CPU requirements for Oracle GoldenGate components vary primarily according to transactions rates, record sizes, and the I/O mix (numbers of updates versus inserts and deletes).

 **Note:**

The following section details system performance statistics in a variety of situations. However, your operating environment, including operating system versions, hardware, and network topology will affect the performance you experience. The following figures should be taken only as guidelines, useful for reference when you are planning your Oracle GoldenGate implementation.

As a percentage of overall application processing, Oracle GoldenGate as a general rule requires more CPU resources when:

- I/O is a larger percentage of overall application processing.
- Files are buffered.
- TMF auditing is off.
- Cache hits are high.
- When TMF is on, audit compression is used.
- All files are replicated.

More power is required because the application can produce relatively more I/O under these conditions.

- [TMF Data Extraction and Transfer](#)
- [Non-TMF Data Extraction and Transfer](#)
- [Intermediate Oracle GoldenGate Trails](#)
- [Replicat Performance](#)

TMF Data Extraction and Transfer

Oracle GoldenGate extracts data in TMF applications by reading TMF audit trails (using the Audserv program). Audserv passes the data in large blocks to Extract, which formats and moves the data to the target system.

Oracle GoldenGate has observed TMF extraction takes between 1% and 10% of overall system processing. Typically this figure is 2-3%, but can skew higher when:

- The database is SQL.
- As a percentage, there are a high number of updates in the application.
- The SQL update statements set a large number of columns.

Non-TMF Data Extraction and Transfer

Non-TMF data extraction takes more processing power than TMF because blocks of data are written to log trails on disk, then transferred to the target system by Extract in a separate step.

It has been observed that non-TMF extraction activities take between 2% and 14% of overall application processing. This means that, if the application is running with an average of 50% busy for the CPUs, after replication is implemented the CPUs will be between 51% and 57% busy. This represents a significant variance and can be traced to the environmental factors listed above.

Intermediate Oracle GoldenGate Trails

You can improve performance by writing data to intermediate trails. An intermediate trail is a trail written by Extract, extracted again, and transferred to Replicat. I/O to intermediate Oracle GoldenGate trails is serial, buffered and blocked at 28K bytes when the system is busy. This means that I/O rates upwards of three gigabytes per hour are possible to trails on the target (Higher rates are possible when you are running multiple instances of Extract, IP channels and controllers.)

Any impact resulting from intermediate trails is absorbed by the target system. The volume of data written to the target is likely to be far less than the corresponding audit generated on the source system for TMF applications.

To minimize contention with other disk activity, isolate trails on separate disk drives when possible.

Replicat Performance

Replicat performance on the target system is related closely to the I/O performance on the source system, as I/O is essentially the same on both. This means you can estimate the impact of I/O on the target system by determining the File Busy percentage of application files and tables on the source system.

Disk Space Requirements

Disk space requirements can vary depending on whether you are processing data from TMF-enabled applications or from non-TMF-enabled applications.

- [TMF Data Requirements](#)
- [Non-TMF Requirements](#)

TMF Data Requirements

TMF replication requires temporary storage for data read from audit trails. In most scenarios, extracted data is written directly to the target system. TMF audit trails generally act as safety storage if the target system or communication channel is unavailable or performing slowly.

In most cases, you must reserve enough space for TMF audit trails to tolerate an outage of one or more days. To determine your space requirements, check the rate at which audit is generated for a week or month and average the amount of space filled per day. Oracle GoldenGate can process tape dumps, so you can have a backlog as

old as the oldest tape dump. However, restoring a tape dump can be cumbersome and usually requires manual intervention, so disk is preferred whenever possible.

Another way to reduce source system disk requirements is to store data in a local trail, then move it across the network through a separate Extract process. In many cases, the amount of data written to a trail by Extract is far less than the corresponding amount of data in the audit trails, especially when a subset of data is being replicated. Filtering unnecessary data in this manner reduces the total disk requirement.

Non-TMF Requirements

Non-TMF database changes are stored in a log trail, which must be a local trail. Logged data is subsequently moved to the target system and replicated.

During typical activity, records in the log trail are replicated almost as soon as they become available, making the amount of disk required quite low. However, you must consider more extreme conditions when sizing disk space. For example, replication activity can fall behind due to:

- Target system maintenance (which might require suspension of Replicat processing)
- Slow data replication on the target due to other activity on the system
- Slow or downed communications links between the source and target

When any of these conditions arise, you must store backlog locally. Many organizations use the general rule of accommodating one full day's worth of database updates. Allocate the number of bytes generated during this period for log trails on the source system.

In addition, you can move data to trails on the target system to minimize data loss caused by potential delays in Replicat processing. Size these trails for potential delays in replication.

When sizing for longer outages (for example, more than one day), you can use tape backups to reduce the amount of disk required.

Data Communications Requirements

Each record Oracle GoldenGate sends to the target system includes a header portion and a data portion. To compute total byte generation and transfer requirements in a given period, sum the byte counts of the headers and data for each record inserted, updated or deleted during the period.

The overall amount of activity can be calculated using the following statement. (Headers vary in size, however, the average is approximately 50 bytes. For this reason, the following example uses an 80-byte header.)

```
(number of inserts * (size of average insert + 80)) +  
(number of deletes * (size of average record + 80)) +  
(number of updates * (size of averaged compressed update + 80))
```

For Enscribe files, the size of insert and delete records can vary and may not be the same as the record size for the file (which represents the maximum size of each

record). For inserts and deletes, all bytes in the record are transferred, along with a header.

For SQL tables, all bytes are transferred for insert and delete records as well.

With Oracle GoldenGate, the following data is not transferred to the target system:

- Images of updates before the update occurs (Only after-images are sent for hot site backup.)
- Alternate key data
- SQL indexes
- Data generated by `FUP RELOAD`
- Assorted other TMF records

When sizing data communications bandwidth, peak byte generation is important. Bandwidth that can accommodate peak volumes helps guarantee low latency between the source and target databases at all times.

When moving data, Extract accumulates and sends records in 28K blocks across the network whenever database activity is high. Therefore the amount of data moved through Extract over a communication channel can be approximately 75-80% of channel capacity. For example, on a 10-megabit Ethernet network with approximately 4-megabits per second of true throughput, Extract can move approximately 3.5M bits per second (430KB/sec).

 **Note:**

All network performance statistics are suggestions. Your own network configuration, bandwidth demand, and operating environment will determine the kind of performance you can expect. You can boost actual throughput considerably by using router-based compression between the source and target systems (a 4 to 1 ratio is fairly typical). This is in addition to update compression used by Oracle GoldenGate modules. For more information, contact your network vendor or NonStop representative.

When sizing communications requirements, note that Oracle GoldenGate can send data over multiple TCP/IP channels simultaneously. You can achieve total bandwidth requirements over multiple lines or networks as an alternative to a single path.

Preparing the Source and Target

Once you have addressed Oracle GoldenGate planning considerations, you are ready to meet its prerequisites. These include:

- Downloading source data definitions
- Synchronizing source and target
- [Downloading Source Data Definitions](#)

- [Synchronizing Source and Target](#)

Downloading Source Data Definitions

With the enhancements of metadata in trails, using `DEFGEN` is no longer a requirement as long as the versions of source and target are of relative versions. This does not matter if Oracle GoldenGate replication is from NonStop to NonStop, NonStop to OpenSys, or Opens to NonStop. `DEFGEN` is still a means to obtain SQL/MP catalog or Enscribe dictionary info to use along with `DDLGEN` to create SQL table creates for OpenSys data bases.

Synchronizing Source and Target

Before running online synchronization for the first time, run an initial-load process to synchronize the source and target. This applies a copy of the source to the target. If Extract or Replicat stops running for any reason, you may need to re-synchronize the source and target and restart the Extract and Replicat groups to reset the checkpoints. For further assistance with your initial load, see the *Oracle GoldenGate for Mainframe Administering for HP NonStop (Guardian)* guide.

Documentation Accessibility

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

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.

Oracle® Fusion Middleware Installing Oracle GoldenGate for HP NonStop (Guardian), 19c (19.1.0.0.0)

F42363-01

Copyright © 2014, 2021, 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 embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. 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 Inside 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, Epyc, and the AMD 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.