

ON-Archive

Database Utility

Quick Start Guide

Version 7.2
April 1996
Part No. 000-7890A

Published by INFORMIX® Press

Informix Software, Inc.
4100 Bohannon Drive
Menlo Park, CA 94025

The following are worldwide trademarks of Informix Software, Inc., or its subsidiaries, registered in the United States of America as indicated by “®,” and in numerous other countries worldwide:

INFORMIX®, C-ISAM®, INFORMIX®.OnLine Dynamic Server™

The following are worldwide trademarks of the indicated owners or their subsidiaries, registered in the United States of America as indicated by “®,” and in numerous other countries worldwide:

X/OpenCompany Ltd.: UNIX®, X/Open®

Some of the products or services mentioned in this document are provided by companies other than Informix. These products or services are identified by the trademark or servicemark of the appropriate company. If you have a question about one of those products or services, please call the company in question directly.

Documentation Team: Bob Berry, Susan Koehler, Katarina Stenstedt

Copyright © 1981-1996 by Informix Software, Inc. All rights reserved.

No part of this work covered by the copyright hereon may be reproduced or used in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems—without permission of the publisher.

To the extent that this software allows the user to store, display, and otherwise manipulate various forms of data, including, without limitation, multimedia content such as photographs, movies, music and other binary large objects (blobs), use of any single blob may potentially infringe upon numerous different third-party intellectual and/or proprietary rights. It is the user's responsibility to avoid infringements of any such third-party rights.

RESTRICTED RIGHTS LEGEND

Software and accompanying materials acquired with United States Federal Government funds or intended for use within or for any United States federal agency are provided with “Restricted Rights” as defined in DFARS 252.227-7013(c)(1)(ii) or FAR 52.227-19.

Table of Contents

Introduction	1
Goals	2
Assumptions	2
Working on a Production System	3
Terminology.	3
Setting Up Your Environment	4
Verify That the Tape Drives Are Ready	4
Log In as User informix	6
Make Space for New Dbspaces	6
Set Environment Variables	6
Prepare the config.arc File	7
Prepare the ONCONFIG File	8
Start OnLine.	8
Create Dbspaces	9
Create Databases	9
Archiving Data Using ON-Archive	10
Start the Cataloger	10
Start onarchive	11
Define a Volume Set	11
Define One or More Volumes	13
Create and Execute an Archive Request	14
Backing Up Logical-Log Files Using ON-Archive	16
Define a Volume Set and a Volume	17
Start a Continuous Backup of the Logical-Log Files	19
End the Continuous Backup of the Logical-Log Files	21
Delete Data from the Database?	22
Performing a Warm Restore	23
Gather the Appropriate Volumes	24
Ensure That Needed Devices Are Available	24
Restore Selected Dbspaces	24
Back Up the Logical-Log File That Was Current When the Warm Restore Began.	26
Restore the Logical-Log Files	28

Terminating ON-Archive	30
Preparing to Rerun the Exercise	30
Problem Solving	31
ON-Archive Error Message Format	31
Typing Errors	32
Tape Problems.	34
If a Request Fails	35
Incorrect Permissions	35
Cataloger Dead	36
Problem Defining a Volume	37
A Sample Session	38

The *ON-Archive Quick Start Guide* presents a complete example of setting up, backing up, and restoring data for a sample INFORMIX-OnLine Dynamic Server system. It introduces you to the concepts, terminology, and basic tasks involved in using ON-Archive as your OnLine recovery system. It includes a section on troubleshooting common problems and a scripted sample session.

This guide does not describe full-scale archiving for an INFORMIX-OnLine Dynamic Server production system. It does not cover the more advanced features of ON-Archive, such as partial archives and parallel archives. Those topics are covered in the [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#).

Introduction

The example in this guide provides the following benefits:

- It allows you to learn without endangering your production databases.
- It takes you through steps that are common to all ON-Archive procedures.
- It gives you a hands-on introduction to ON-Archive concepts and terminology.

Goals

The example in this guide familiarizes you with ON-Archive by requiring you to perform the following tasks:

- Start the cataloger and **onarchive**
- Create a volume set and volume
- Make a level-0 archive using **onarchive**
- Back up the logical-log files
- Restore from a level-0 archive and a logical-log backup using **onarchive**

Assumptions

For the example in this manual, Informix made the following assumptions about your environment:

- You are already familiar with OnLine.
- You are archiving only one instance of OnLine.
- OnLine uses continuous logical-log backups.
- Two tape drives are available.
- You can archive all the data on one tape.
- An operator is present during the archive and restore.
- You are working at a workstation where you can run multiple concurrent windows. (If you do not have multiple concurrent windows, you need two terminals so that you can devote one terminal to the continuous logical-log backups.)

The example gives explicit samples of the commands you should type to do various operations. If you want the sample commands to work, you must prepare the sample environment and name the elements to match the preceding assumptions. The instructions tell you to prepare the following items:

- A separate instance of OnLine
- Two tape devices, named **tape** and **ltape**

- Two dbspaces, named **dbspace1** and **dbspace2**
- Two copies of the **stores7** database, named **databaseOne** and **databaseTwo**

Working on a Production System

When you do archives on a full production system, you must consider the following factors, which are not treated in the example:

- Creating dbspace sets and save sets
- Using **onautovop** for unattended operations
- Using parallel restores
- Protecting storage
- Using temporary storage areas
- Interchanging data between two OnLine database servers
- Doing a cold restore
- Restoring while users are active
- Using multiple tapes for the archive
- Restoring from multiple logical-log backup tapes

Terminology

In this document we will use the following terminology:

<i>archive</i>	refers to the saving of database data.
<i>backup</i>	refers to the saving of logical logs.
<i>restore</i>	refers to the copying of data from an archive storage medium onto the active storage medium used by OnLine.

Setting Up Your Environment

The next sections describe the steps you must take to set up the practice environment. The arrow symbol (\Rightarrow) marks places where you need to take an action.

To set up the practice environment

1. Verify that the tape drives are ready.
2. Log in as user **informix**.
3. Make space for new dbspaces (see step 8).
4. Set environment variables.
5. Prepare the **config.arc** file.
6. Prepare the ONCONFIG file.
7. Start OnLine.
8. Create dbspaces.
9. Create databases.

Verify That the Tape Drives Are Ready

Find out the device names of your tape drives so that you can substitute the proper device names in the commands. The tape-device names vary from one hardware platform to another, so you must find out the names for your system. In the examples, we use the following tape-device names:

- **/dev/rst0**
- **/dev/rst1**

You cannot do any archiving or backup to tape if your tape drive does not work. Check the status of both of your tape drives. Again, the technique for checking tape drives varies from hardware platform to platform. The UNIX **mt** command might be available. The syntax of the **mt** command follows:

```
mt -f tape_device_pathname status
```


Enter the following commands, substituting the device names of your tape drives:



```
# mt -f /dev/rst0 status
```

```
Archive QIC-150 tape drive:
sense key(0x0)= no sense   residual= 0   retries= 0
file no= 0    block no= 0
```



```
# mt -f /dev/rst1 status
```

```
Archive QIC-150 tape drive:
sense key(0x0)= no sense   residual= 0   retries= 0
file no= 0    block no= 0
```

If your tapes are properly connected and active, the **mt** commands return messages *similar* to (but not exactly the same as) those shown in the preceding example. If you do not get a description of the status, check for the following possible problems:

- No tape is loaded.
- Number on back of tape drive is set incorrectly.
- Tape drive is not connected to console.
- Tape drive is not turned on.

When the tape drive is not turned on, you notice a very heavy load on the system for a minute or two before you receive any error message.

Log In as User informix



When you use ON-Archive, you are usually logged in as user **informix**. Log in now as user **informix**.

Make Space for New Dbspaces

This step prepares two files, **space1** and **space2**, that are used for two dbspaces. It *does not* create the dbspaces. You create the dbspaces in a later step.

You can put the dbspaces in any storage location that has sufficient space. This example uses cooked files in the **/tmp** directory.

Enter the following commands to prepare the files:



```
# cd /tmp
# touch rootdbs
# touch space1
# touch space2
# chmod 660 rootdbs
# chmod 660 space1
# chmod 660 space2
# chgrp informix rootdbs
# chgrp informix space1
# chgrp informix space2
```

Set Environment Variables



Set your environment variables for a new instance of OnLine. See the [*INFORMIX-OnLine Dynamic Server Administrator's Guide*](#) for information on installing and configuring OnLine. For the procedures illustrated in this example, you do not need to set any environment variables for ON-Archive.

Prepare the config.arc File

The `$INFORMIXDIR/etc/config.arc` file specifies the device information and the file information used by ON-Archive. You do not need to change anything in this file except the device names.

When you install a new instance of OnLine, the **config.arc** file contains a series of entries similar to the following example:

```
! Device available for onarchive.
DEVICE HO = /home
DEVICE ho = /home
DEVICE DISK = /disk/path
DEVICE disk = /disk/path
DEVICE TAPE = /dev/tapedev, /dev/tapedev2
DEVICE tape = /dev/tapedev, /dev/tapedev2
DEVICE TAPE_VOP = /dev/tapedev
DEVICE tape_vop = /dev/tapedev
```

The following line shows the syntax of the DEVICE entry:

```
DEVICE logical_device_name = physical_device_pathname
```

The word that follows the DEVICE keyword is the name of the *logical device*. The preceding example defines the following logical-device names: **HO**, **ho**, **DISK**, **disk**, **TAPE**, **tape**, **TAPE_VOP**, and **tape_vop**. The equal sign is followed by the actual pathname that represents the *physical device*.

The definitions that are already in the **config.arc** file are sample definitions. You do not need to remove the sample definitions *unless* the lines you enter in the next step cause you to have duplicate logical-device names.

Before you put your definitions in **config.arc**, save a copy of the file. Enter the following commands:

```
% cd $INFORMIXDIR/etc
% cp config.arc config.arc.save
```

Edit the `$INFORMIXDIR/etc/config.arc` file. Enter the following lines, using the physical-device names of your tape drives:

```
DEVICE tape = pathname_for_device_1
DEVICE ltape = pathname_for_device_2
```

The two preceding lines in your **config.arc** file define two logical-device names for your tape drives: **tape** and **ltape**.

Prepare the ONCONFIG File

Start an instance of OnLine specifically for this exercise. Before you start OnLine you must prepare an ONCONFIG file. You can follow the basic procedure given for the Learning Environment in Chapter 3 of the [INFORMIX-OnLine Dynamic Server Administrator's Guide](#). However, set the TAPEDEV and LTAPEDEV parameters to something other than **/dev/null**, as illustrated in the following example:

```
TAPEDEV pathname_for_device_1
LTAPEDEV pathname_for_device_2
```

Even though ON-Archive does not use TAPEDEV and LTAPEDEV to specify devices, archives and logical-log backups will not work properly if these parameters are set to **/dev/null**.

Tip: Refer to the [“INFORMIX-OnLine Dynamic Server Archive and Backup Guide”](#) for information about the (undesirable) effects of setting the tape-device parameters to **/dev/null**.

Warning: Be sure that the ROOTNAME parameter is set to **rootdbs** and that ROOTPATH is set to **/tmp/rootdbs** or whatever you have chosen as the path to your example root dbspace. Otherwise you might inadvertently overwrite the **rootdbs** for a production database system when you start OnLine.

Start OnLine

Before starting OnLine, make sure you that have an entry for the **dbservername** in the **sqlhosts** file. Then you can start OnLine with the following command:

```
oninit -i
```



Create Dbspaces

For this example, you need two dbspaces in addition to the root dspace. You can use the **onspaces** utility to create dbspaces named **dbspace1** and **dbspace2** in the files that you created in [“Make Space for New Dbspaces” on page 6](#).

Enter the following commands:



```
% onspaces -c -d dbpace1 -p /tmp/space1 -o 0 -s 1000
% onspaces -c -d dbpace2 -p /tmp/space2 -o 0 -s 1000
```

After you execute the **onspaces** commands, you can use the **onstat** utility to verify that the dbspaces are present.

Enter the following command:



```
% onstat -d
```

The first few lines of the resulting message list the dbspaces in the rightmost column.

Create Databases

Put a database that uses logging in each of the dbspaces that you just created. An easy way to do this is to make copies of the **stores7** database.

Enter the following commands:



```
% dbaccessdemo7 -dbspace dbpace1 -log databaseOne
% dbaccessdemo7 -dbspace dbpace2 -log databaseTwo
```

The **dbaccessdemo7** command is documented in the Introduction of the [INFORMIX-OnLine Dynamic Server Administrator's Guide](#).

Archiving Data Using ON-Archive

To archive data using ON-Archive

1. Start the cataloger.
2. Start **onarchive**.
3. Define a volume set (vset).
4. Define one or more volumes.
5. Create and execute an archive request.

Each of these steps is described in detail in the following sections.

Start the Cataloger

To start the **oncatlgr** utility, OnLine must be in on-line mode. Enter the following command at the operating-system prompt:



```
% start_oncatlgr
```

Immediately after executing this command, ON-Archive displays a message similar to the following example:

```
Informix ONCATLGR Version 7.01  
Starting the cataloger.  
  
See file '/tmp/oncatlgr.out.3145' for results.
```

To check that **oncatlgr** started successfully, verify that the file mentioned in the preceding screen display contains a line similar to the following example:

```
Installation procedure successful on 13-APR-1994:11:10:07.
```

Start onarchive

Once the cataloger is started, you can start the **onarchive** utility. You use the **onarchive** utility to enter ON-Archive commands and execute requests. To start **onarchive**, enter the following command at the operating-system prompt:



```
% onarchive
```

Immediately after this command is executed, ON-Archive displays the following ON-Archive prompt:

```
Onarchive>
```

Commands typed at the `Onarchive>` prompt can be in uppercase or lowercase characters; ON-Archive is not case sensitive. This guide uses uppercase characters for ON-Archive commands and qualifiers and lowercase characters for values that the user supplies.



***Tip:** See [“Typing Long Commands” on page 33](#) for an explanation of how to abbreviate ON-Archive commands.*

Define a Volume Set

An ON-Archive *volume* is a single unit of physical media that ON-Archive uses to store archive or backup data. Because the media can be either disk or tape, the generic term volume is used to describe the media.

You must define a *volume set* before you define a volume. A volume set (or *vset*) is a logical grouping of volumes that you designate as the destination for an archive or logical-log backup. A volume set provides the following benefits:

- It allows you to cluster volumes that are used for the same purpose.
- It gives ON-Archive a pool of available volumes from which to choose when starting or continuing archive or backup operations.

For example, Figure 1 illustrates how you might name and group volumes into two such volume sets, one for an accounting system and one for a reservation system.




Figure 1
*Illustration of Two
Volume Sets:
account_set and
reserve_set*

To define a volume set, you issue the `DEFINE/VSET` command at the ON-Archive prompt. Enter the following command to define a vset named **arcvset**:



```
Onarchive>DEFINE/VSET=arcvset/CLASS=SYSTEM/DRIVER=TAPE/ACCESS=0/-  
_Onarchive>DEVICE_TYPE=tape
```

The text that follows the `DEFINE/VSET=arcvset` specifies values for the following qualifiers: `CLASS`, `DRIVER`, `ACCESS`, and `DEVICE_TYPE`. When you define a vset, you must provide values for these qualifiers. Each qualifier, together with a brief description of its function, is contained in the following list:

<code>CLASS</code>	Determines which users have access to a vset.
<code>DRIVER</code>	Is set to disk or tape.
<code>ACCESS</code>	ON-Archive automatically selects a vset depending on this value when you do not specify a vset at the command line.
<code>DEVICE_TYPE</code>	Is set to any one of the logical device drivers you defined in config.arc file.

If you wish to use a name other than **arcvset** for your vset, substitute that name and execute the preceding `DEFINE/VSET` command (with all qualifiers) now. Immediately after you enter a command to define your vset, ON-Archive responds with a message similar to the following example:

```
Vset arcvset created.
```


Define One or More Volumes

After you define a vset, you use the `DEFINE/VOLUME` command to define volumes for the vset. Enter the following `DEFINE/VOLUME` command to create a volume in a tape vset called **arcvset**:



```
Onarchive> DEFINE/VOLUME/VSET=arcvset
```

Immediately after you issue a `DEFINE/VOLUME` command, ON-Archive prompts you to select the device on which you plan to mount the volume. For example, suppose that you defined two logical devices with the following lines in **config.arc**:

```
DEVICE tape = /dev/rst0
DEVICE ltape = /dev/rst1
```

Given this configuration, ON-Archive prompts you as follows:

```
Devices available for ONARCHIVE
DEVICE #      DEVICE NAME
-----
      1      /dev/rst0

Enter the device numbers (1) or (C)ancel : 1
```



When ON-Archive displays this screen, enter the device number associated with the correct physical device. In the preceding example, the user wishes to mount the volume on the device **/dev/rst0**, and so types 1.

Then ON-Archive prompts you to mount a volume (insert a tape into your tape drive), as follows:

```
Please mount a volume on /dev/rst0 and press < RETURN > ** WRITE **
```



This message not only gives you time to mount a volume but also verifies the name of the physical device. Press RETURN, and ON-Archive initializes the volume. If this is the first time you are using this tape with ON-Archive, you receive a message similar to the following example:

```
ARC-I-01098, Not an ONARCHIVE volume.  
Proceed and overwrite contents, Cancel operation, or Retry with another  
volume?  
Enter (P/C/R):
```



Respond to this message by typing P for proceed. When the initialization is complete, ON-Archive displays a message similar to the following example:

```
Volume number 0001 created in arcvset.
```

The preceding message indicates that you successfully created a volume. ON-Archive is now prepared to perform an archive of the sample data.

Create and Execute an Archive Request

Many of the tasks that you perform with ON-Archive require a two-step process. First, you create a request, and then you execute the request. To create a request that performs a level-0 archive of all the data managed by your OnLine database server, enter the following command:



```
Onarchive> ARCHIVE/DBSPACESET=*/VSET=arcvset
```

The asterisk following the DBSPACESET qualifier indicates that you wish to perform a full-system archive. The VSET qualifier is set to `arcvset`, indicating to ON-Archive that the target volume set for this archive is `arcvset`. The default archive level is zero.

After you enter the ARCHIVE command, ON-Archive displays a message similar to the following example:

```
Request 00000001 registered in the catalog.
```

The number that occurs after the word `Request` is called a *request number*. (The request number displayed on your screen might be different from the number in the preceding example.) To execute the request, you must use the `EXECUTE/REQUEST` command, specifying the request number (leading zeroes are not required). Enter the following command, substituting the request number that ON-Archive displayed on your screen when you created the archive request:



```
Onarchive> EXECUTE/REQUEST=1
```

After you execute the request, ON-Archive prompts you to select a device number and mount the appropriate tapes as follows:

```
Onarchive> execute/request=1
Executing...
Devices available for ONARCHIVE
DEVICE #          DEVICE NAME
-----
      1              /dev/rst0

Enter the device number (1) or (C)ancel : 1
Please mount volume 0001 of arcvset and press < RETURN > ** WRITE **
```



Press RETURN to continue.

When the archive is complete, a message similar to the following example is displayed on your screen, and ON-Archive returns you to the `Onarchive>` prompt:

```
File rootdbs being processed.
File dbspace1 being processed.
File dbspace2 being processed.
Onarchive>
```

You are ready to proceed to the next step in this example, backing up your logical logs.

Getting Information on Requests

You can use the LIST/REQUEST command to get information on the requests that you create and ON-Archive stores in the catalog. You can use the LIST/REQUEST command to examine requests that were created before or after a particular date, that are owned by a specific user, or that have a particular status.

For example, the following command shows all the requests issued by the user **judith**:

```
Onarchive> EXECUTE/REQUEST=*/USER=judith
```

You can also use a LIST command to list the contents of a volume. For more information on the LIST commands, see the [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#).

Backing Up Logical-Log Files Using ON-Archive

This section explains how to start a continuous backup of the logical log using ON-Archive. A *continuous* backup of the logical log automatically backs up individual log files when OnLine indicates they are full and ready to be backed up. A continuous backup eliminates the need for you to monitor the logical log and determine when to back it up.

To start a continuous backup using ON-Archive

1. Define a volume set and volume to which ON-Archive will back up the logical-log files.
2. Start the continuous backup.

Define a Volume Set and a Volume

You need to define a second vset and volume to back up the logical-log files because the backup uses the second tape drive.

Enter the following command at the `Onarchive>` prompt to create a vset for the backup:



```
Onarchive> DEFINE/VSET=logvset/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_TYPE=ltape
```

ON-Archive responds with the following message:

```
Vset logvset created.
```

Enter the following command at the `Onarchive>` prompt to create a volume for the vset:



```
Onarchive> DEFINE/VOLUME/VSET=logvset
```

ON-Archive prompts you as follows to enter the device number of the tape drive:

```
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst1

```
Enter the device number (1) or (C)ancel : 1
```



Enter 1 to select **/dev/rst1**. This is the physical device that you assigned (in the **config.arc** file) to the logical device (**ltape**) for this vset.

ON-Archive prompts you as follows to mount a volume on **/dev/rst1**:

```
Please mount a volume on /dev/rst1 and press < RETURN > ** WRITE **
```



Insert a tape in the tape drive and press RETURN. If you insert a new tape that does not contain any data, ON-Archive begins to initialize the volume.

If you insert a tape that contains data, ON-Archive displays messages similar to those in the following example:

```
ON-Archive: Media Identification Function

Tape Information:

      Vset : logvset
      Volume : 000001

ARC-W-01845, Initializing this media may overwrite an ON-Archive volume.
Proceed and overwrite contents, Cancel operation, or Retry
with another volume?
Enter (P/C/R): P
```



Enter **P** to proceed if you wish to continue the operation and overwrite the data on the tape. Enter **C** if you wish to cancel the operation. You must define a **vset** and a **volume**, however, before you continue with these procedures. Otherwise, insert a new tape that does not contain data and press **R** to retry the operation.

ON-Archive displays the following message after it creates the volume:

```
Volume number 0001 created in logvset.
```

Once you create a volume, you are ready to start a continuous backup of the logical log, as explained in the following section.

Start a Continuous Backup of the Logical-Log Files

This section describes how to start a continuous backup of the logical-log files using the **onarchive** process. You can also use the **onautovop** process to start a continuous backup automatically. The [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#) describes how to use **onautovop**.



You must create a second window or terminal to do the continuous backup of the logical-log files because the **onarchive** process holds onto the window or terminal to interact with the operator as necessary. This feature makes the window unavailable for other purposes until you explicitly terminate the backup.



***Tip:** Perform the following steps in your second window or terminal before you start the backup.*

To start a continuous backup



1. Log in as user **informix**.



2. Set the same OnLine and ON-Archive environment variables that you set in the original window or terminal.

3. Enter the following command at the operating-system prompt to start a second **onarchive** process. (You *do not* need to start a second **oncatlgr** process.)



```
% onarchive
```

Enter the following command at the **Onarchive>** prompt to request a continuous backup of the logical log:



```
Onarchive> BACKUP/LOGFILE/CONTINUOUS/VSET=logvset
```

ON-Archive catalogs the request and displays the request number in a message like the following example:

```
Request 00000002 registered in the catalog.
```

Execute the request by entering the following command at the **Onarchive>** prompt, replacing 2 with the request number that ON-Archive just displayed:



```
Onarchive> EXECUTE/REQUEST=2
```

Start a Continuous Backup of the Logical-Log Files

ON-Archive prompts you as follows to enter the device number of the tape drive:

```
Executing...
Devices available for ONARCHIVE

DEVICE #          DEVICE NAME
-----          -
      1              /dev/rst1

Enter the device number (1) or (C)ancel :
```



Enter 1 to select device `/dev/rst1`.

ON-Archive prompts you as follows to mount volume 0001 of `vset logvset` in the tape drive:

```
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```



Insert the log backup tape in the tape drive and press RETURN. ON-Archive begins the backup.

If any log files are full at the time, ON-Archive copies them to the tape first. As shown in the following example, after ON-Archive backs up the log files that were full, it lists them and then prompts you again to enter the device number of the tape drive:

```
Logfile LF00000001 being processed.
Logfile LF00000002 being processed.
Logfile LF00000003 being processed.
Devices available for ONARCHIVE

DEVICE #          DEVICE NAME
-----          -
      1              /dev/rst1

Enter the device number (1) or (C)ancel : 1
```



Enter 1 to select device `/dev/rst1`.

ON-Archive again prompts you as follows to mount volume 0001 of vset logvset in the tape drive:

```
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```



Press RETURN. The tape is already in the tape drive.

At this point, or if no log files were full to begin with, you do not see any activity because ON-Archive is waiting for the current logical-log file to fill up before writing it to tape. The window or terminal you are using is now dedicated to the **onarchive** process until you terminate the continuous backup. The next section explains how to do this.

End the Continuous Backup of the Logical-Log Files

For this example, you must end the continuous backup of the logical-log files before you start the warm restore that is described in the following sections. You need the backup tape that you just created to restore the logical-log files after the warm restore.

You terminate a continuous backup of the logical-log files by pressing the Interrupt key. The Interrupt key is usually defined as the combination of the CONTROL key and the C key, or CTRL-C. You enter it by pressing both keys at the same time.



Enter the Interrupt key (CTRL-C) to terminate the continuous backup of the logical-log files.

ON-Archive displays the keystrokes you entered and the message shown in the following example:

```
^C
ARC-I-01483, No data to archive.
```

```
%
```

The UNIX shell then terminates *both the continuous backup and the **onarchive** process*, returning you to the operating-system prompt.

When you terminate a continuous backup of the logical-log files, your backup will be incomplete if ON-Archive is writing a log file to tape at the time. The [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#) describes the conditions that might exist when you interrupt a continuous backup of the logical-log files.



Return to your original window or terminal to perform the remaining procedures in this guide.

Delete Data from the Database?

This section is optional. It allows you to delete data from your databases prior to performing the warm restore described in “[Performing a Warm Restore](#)” on page 23. Typically, you restore dbspaces because you somehow have lost data that you need. This section simply provides justification for doing the warm restore. It is not a necessary step for either learning or performing the procedures that follow. If you wish, you can skip to “[Performing a Warm Restore](#)” on page 23 and begin the warm restore now.

If you like, however, you can delete some data from your sample databases—**databaseOne** or **databaseTwo**—to see it restored by the operations that follow.

You can use the DB-Access utility to delete data from your database. To use DB-Access, you should open a new window or use another terminal or exit ON-Archive temporarily to free up the current window. To exit ON-Archive, enter `exit` at the `Onarchive>` prompt as follows:



```
Onarchive> exit
```

The following DELETE statements are provided for your convenience. They delete data from either of your databases without violating the referential constraints that the databases contain. See the [Informix Guide to SQL: Reference](#) for a description of the **stores7** database, which is the source for the databases you are using. You can use DB-Access to enter and run one or more of the following statements. See the [DB-Access User Manual](#) for information on how to use that product.



```
DELETE FROM items WHERE order_num = 1021
DELETE FROM items WHERE stock_num = 6
DELETE FROM catalog WHERE manu_code = 'ANZ'
DELETE FROM cust_calls WHERE customer_num = 110
DELETE FROM cust_calls WHERE user_id = 'maryj'
```

After deleting data from your databases, either return to the window running **onarchive** or enter **onarchive** at the operating-system prompt in the current window to restart ON-Archive and continue with the warm restore in the following section:



```
% onarchive
```

Performing a Warm Restore

This section explains how to perform a *warm restore* using ON-Archive. A warm restore restores selected noncritical dbspaces and blobspaces *while OnLine is in on-line or quiescent mode*. A *noncritical dbspace* is any dbspace other than **rootdbs** or the dbspaces that contain the physical log and the logical-log files.

The warm restore presented here is an *attended* warm restore—that is, one that you execute and attend to as it runs. The [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#) explains how to perform an *unattended* warm restore.

To perform an attended warm restore using ON-Archive

1. Gather the volumes that contain the archives and logical logs you need to restore.
2. Ensure that the OnLine devices and disk spaces to which you will be restoring data are available.
3. Restore selected dbspaces (or blobspaces).

4. Back up the logical-log file that was current when you began the warm restore.
5. Restore the logical-log files.

The following sections describe each of these steps in detail.

Gather the Appropriate Volumes

ON-Archive tells you which vsets and volumes to mount during the warm restore. For the warm restore you are about to do, you need the tape you used for the level-0 archive and the tape you used for the continuous backup of the logical-log files.

Ensure That Needed Devices Are Available

In general, you must ensure that the raw disk spaces or files from which you archived OnLine data are still available when you perform the restore.

The following procedures assume that the dbspaces you archived earlier are still available.

Restore Selected Dbspaces

The RETRIEVE/DBSPACESET command enables you to restore either all or some of the dbspaces within a dbspace set. A *dbspace set* is a group of dbspaces that ON-Archive archives together.

When you do a warm restore, you restore selected noncritical dbspaces (or blobspaces) from an archive of OnLine data. OnLine prohibits access to the dbspaces you select until the warm restore is complete and any logical-log records for the selected spaces have been rolled forward.

Enter the following command at the `Onarchive>` prompt to restore **dbspace1** and **dbspace2** from the level-0 archive you performed earlier:



```
Onarchive> RETRIEVE/DBSPACESET=*/DBSPACE=(dbspace1,dbspace2)
```

ON-Archive catalogs the request and displays the request number.



Execute the request by entering the following command at the `Onarchive>` prompt, replacing 3 with the request number that ON-Archive just displayed:

```
Onarchive> EXECUTE/REQUEST=3
```

ON-Archive displays the following messages and asks if you want to proceed with the request or cancel it:

```
Executing...
```

```
Any online dbspaces being retrieved will first be taken offline.
```

```
Do you want to Proceed or Cancel this request? (P/C)
```



Enter **P** to proceed.

OnLine displays the name of the first dbspace it will restore and then prompts you to enter the device number of the tape drive, as shown in the following example:

```
File dbspacel being retrieved.  
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst0

```
Enter the device number (1) or (C)ancel :
```



Enter **1** to select device **/dev/rst0**.

ON-Archive prompts you to mount volume 0001 of the **arcvset** vset as shown in the following example:

```
Please mount volume 0001 of arcvset and press < RETURN > ** READ **
```



Insert the archive tape in the tape drive and press RETURN.

ON-Archive begins reading from the tape and restoring the dbspaces you selected. As shown in the following example, ON-Archive lists each dbspace it retrieves from the tape and returns to the `Onarchive>` prompt when it has finished:

```
File dbspace2 being retrieved.  
Onarchive>
```

After restoring the selected dbspaces and blobspaces, a warm restore requires you to roll forward any transactions in the logical log for these spaces. Before you can roll forward any transactions for these spaces, however, you must back up the logical-log file that was current when the warm restore began. The next section describes how to do this.

Back Up the Logical-Log File That Was Current When the Warm Restore Began

You must back up the logical-log file that was current when the warm restore began because this file contains the checkpoint that signals the end of transactions for the restored dbspaces and blobspaces.

Use the same tape for this backup that you used earlier for the continuous backup of the logical-log files.

Enter the following command at the `Onarchive>` prompt to request a backup of the logical-log file that was current when the warm restore began:



```
Onarchive> BACKUP/LOGFILE/AUTOMATIC/CURRENT/VSET=logvset
```

ON-Archive catalogs the request and displays the request number.

Enter the following command at the `Onarchive>` prompt to begin the backup, replacing 4 with the request number that ON-Archive just displayed:



```
Onarchive> EXECUTE/REQUEST=4
```

Back Up the Logical-Log File That Was Current When the Warm Restore Began

ON-Archive prompts you as follows to enter the device number of the tape drive you are using:

```
Devices available for ONARCHIVE
DEVICE #      DEVICE NAME
-----
      1      /dev/rst1
Enter the device number (1) or (C)ancel :
```



Enter 1 to select `/dev/rst1`.

ON-Archive prompts you as follows to place volume 0001 of the logvset vset in the tape drive:

```
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```



Insert the log backup tape in the tape drive and press RETURN.

ON-Archive backs up the log file that was current and returns to the `Onarchive >` prompt when it is finished. You are now ready to roll forward any transactions in the logical log for the dbspaces you restored.

Restore the Logical-Log Files

The last step in performing a warm restore is to roll forward any transactions in the logical log for the restored dbspaces. This procedure is called a *logical restore*. The logical restore restores any data that might have been written to the dbspaces since you performed the archive.

OnLine keeps a record of which log files are needed for a logical restore. In general, you supply the following logical-log files to roll forward any transactions for the restored spaces:

1. The log file containing the checkpoint generated by the archive.
2. The log file containing the checkpoint generated by the warm restore.
3. All the log files between step 1 and step 2.

If any of these log files is missing, the logical restore fails.

All the log files you need for the following logical restore are on the tape you used for the continuous backup of the logical-log files.

Enter the following command at the `Onarchive>` prompt to request the logical restore:



```
Onarchive> RETRIEVE/LOGFILE
```

ON-Archive catalogs the request and displays the request number.

Enter the following command at the `Onarchive>` prompt to begin the logical restore, replacing 5 with the request number that ON-Archive displayed:



```
Onarchive> EXECUTE/REQUEST=5
```


As seen in the following example, ON-Archive retrieves the first log file that must be restored—LF00000003 in this case—and prompts you to enter the device number of the tape drive.

```
Executing...
File LF00000003 being retrieved.
Devices available for ONARCHIVE

DEVICE #          DEVICE NAME
-----          -
      1              /dev/rst1

Enter device number (1) or (C)ancel : 1
```



Enter 1 to select **/dev/rst1**.

ON-Archive prompts you as follows to mount volume 0001 of **vset** logvset:

```
Please mount volume 0001 of logvset and press < RETURN > ** READ **
```



Insert the log backup tape in the tape drive and press RETURN.



If ON-Archive restores multiple log files, it repeats the preceding prompts for each logical-log file that it rolls forward. Enter the same responses to those prompts that you just entered for these first prompts.

When ON-Archive finishes, it displays the `Onarchive>` prompt as follows and OnLine makes the restored dbspaces and blobspaces accessible to users:

```
Onarchive>
```

At this point, you are finished with the operations in this manual. If you wish, you can terminate ON-Archive as described in the next section.

Terminating ON-Archive

This section describes how to terminate ON-Archive. (You need to terminate ON-Archive only in the original window. You terminated ON-Archive in the second window when you terminated the **onarchive** process by entering CTRL-C.)

To terminate the **onarchive** process, enter `exit` at the `Onarchive>` prompt as follows:



```
Onarchive> exit
```

To terminate the **oncatlgr** process, enter the following command at the operating-system prompt:



```
%stop_oncatlgr
```

The **stop_oncatlgr** script displays messages like the ones that follow except that the process ID—4044 in this case—varies:

```
Stopping ONCATLGR  
Kill 4044? (y/n)
```



Enter `y` to terminate the **oncatlgr** process.

Preparing to Rerun the Exercise

If you or one of your colleagues wants to run through this exercise another time, you must put the system back into its original state.

To prepare to rerun the exercise

1. Set your environment variables for the instance of OnLine that was used in this exercise.
2. If you have not done so, exit **onarchive** and execute the **stop_oncatlgr** command to stop the cataloger.

3. Execute the **onmode -k** command to remove OnLine.
4. Return to [“Start OnLine” on page 8](#).

Problem Solving

This section describes a somewhat random collection of errors that you might make as you begin using ON-Archive. [“A Sample Session”](#) on page 38 shows the screens as a typical user works through the example given in this pamphlet.

ON-Archive Error Message Format

ON-Archive error messages are documented in Chapter 9 of the [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#). A typical ON-Archive error message looks like this:

```
ARC-E-01018      Request already completed or cancelled.
```

The error messages in the [INFORMIX-OnLine Dynamic Server Archive and Backup Guide](#) are sorted in numerical order. When you are looking up an error in the manual, disregard the letters.

The middle letter of the error number (as displayed on your console) is shown at the end of the error number in the manual, as follows:

```
ARC-01018E      Request already completed or cancelled.
```

Typing Errors

ON-Archive often requires long entries to create a request. It is easy to make a typing error and sometimes difficult to spot the error. The following examples illustrate some typing errors:

- **Omitting a slash mark (between DEFINE and VSET)**

```
Onarchive> DEFINE VSET=arcvset/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_TYPE=tape
```

ARC-E-01087, Missing slash (/) before qualifier name.

- **Typing a slash instead of an equal sign between VSET and arcvset**

```
Onarchive> DEFINE/VSET/arcvset/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_TYPE=tape
```

ARC-E-01083, Missing parameter for qualifier /VSET.

- **Mistaking a qualifier name (should be DRIVER, not DEVICE)**

```
Onarchive> DEFINE/VSET/arcvset/CLASS=SYSTEM/DEVICE=TAPE/ACCESS=0/DEVICE_T=tape
```

ARC-E-01012, Ambiguous qualifier abbreviation /DEVICE.

- Copying a command incorrectly

On the system that we use for testing, we can use mouse clicks to copy a line from the system prompt and enter the line at another system prompt. Unfortunately, if the original line wraps around, ON-Archive cannot interpret the resulting line. The following examples show two such errors:

```
Onarchive>DEFINE/VSET=aNewVolumeSet/CLASS=SYSTEM/DRIVER=TAPE/ACCESS=0-
_Onarchive>/DEVICE_TYPE=tape
```

ARC-E-01129, Unknown command

```
Onarchive>DEFINE/VSET=MyVolume3/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_TYPE=tape
```

ARC-E-01842, Invalid parameter format

You can avoid this problem by simply retyping the line or by making your window wider.

Typing Long Commands

It is easy to make mistakes typing long commands. You can reduce mistakes in typing by shortening the length of the command. When you enter a command, you need to type only enough letters of a qualifier to uniquely specify the qualifier name. For example, this request defines a vset named **zorch2**:

```
Onarchive> DEF/VSET=zorch2/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_T=tape
```

Notice that shortened forms were used for DEFINE (DEF), DRIVER (DRI) and DEVICE_TYPE (DEVICE_T).

Tape Problems

You might see the tape problems mentioned in [“Verify That the Tape Drives Are Ready” on page 4](#) in any tape application. The following tape problems are specific to ON-Archive.

If the tape in the tape drive has its write protection set, you cannot write to the tape to prepare a volume or to make an archive. You might get the following message:

```
ARC-E-01564, Cannot open tape.  
UNIX-E-00013, Permission denied
```

You also get the previous message if you use an incorrect type of tape; that is, the physical tape itself is incorrect. For example, you might have used an X-AB12 tape instead of an X-A88 tape.

An incorrectly formatted tape might cause the following interaction:

```
Please mount a volume on /dev/rst0 and press < RETURN > ** WRITE **  
  
ARC-E-01322, Error reading the ANSI volume label.  
Proceed and overwrite contents, Cancel operation, or Retry  
with another volume?  
Enter (P/C/R): p  
The volume 0001 has been deleted from the Catalog.  
  
UNIX-E-00013, Permission denied  
ARC-E-01324, Error writing the ANSI HDR1.  
ARC-E-01267, Tape not initialized.  
ARC-E-01269, Tape not defined.
```

If a Request Fails

If a request generates an error when it is executed, **onarchive** does not automatically retry the request. Instead, the request status changes to CANCELLED or FAILED. If you want to execute the request again, you need to make a new request with a status of NEW. You can re-create the entire request, or you can use a simple MODIFY/COMMAND command. The following command creates a new request that is exactly the same as request 47:

```
Onarchive> MODIFY/COMMAND=47
```

To delete the original FAILED or CANCELLED request when you create the new request, use the DELETE qualifier with the MODIFY/COMMAND command, as follows:

```
Onarchive> MODIFY/COMMAND=47/DELETE
```

Once the new request is created, you still need to execute it again. Request execution is explained in [“Create and Execute an Archive Request” on page 14](#).

Incorrect Permissions

If the permissions and group of the file used for a dbspace are not correct, the **onspaces** command fails. In the following example, the group of **/tmp/space3** is not **informix**:

```
% onspaces -c -d dbspace3 -p /tmp/space3 -o 0 -s 1000
```

```
Verifying physical disk space, please wait ...  
Bad GROUP-ID: '/tmp/space3'.
```

You can use the **ls -lg** command to check the permissions, the owner, and the group of a file, as shown in the following example:

```
%ls -lg /tmp/spaces3
```

Cataloger Dead

If you exit from ON-Archive and restart OnLine, you must stop the cataloger and then restart it. If you forget, you might have an active but nonfunctional cataloger. In the following interaction, a cataloger from a previous session has been left running.

```
Onarchive> define/vset=testdb/class=system/driver=tape/access=0/device_type=tape
```

```
The cataloger is dead.  
ARC-E-01211, The cataloger is not installed.
```

If you want to check that your current cataloger is still running, first find its process number in the **oncatlgr** output file as shown in the following example.

The following example illustrates how to find the **oncatlgr** process:

```
% start_oncatlgr
```

```
Informix ONCATLGR Version 7.01  
Starting the cataloger.  
  
See file '/tmp/oncatlgr.out.1597' for results.
```

```
% cat /tmp/oncatlgr.out.1597
```

```
...  
...  
ONCATLGR utility For INFORMIX-Online.  
Installation procedure successful on 25-APR-1994:10:21:14.  
Process ID = 1597 Servernum = 32.
```

```
% ps -e | grep 1597 | grep -v grep
```

```
informix 1597 0.0 2.9 280 860 p4 I 10:21 0:00 /infmtxdir
```


You might need to use a different form of the **ps** command, depending on what version of the UNIX operating system you are using.

Problem Defining a Volume

The following problem appears to relate to defining a volume, but it is actually caused by an incorrect DEVICE definition in the **config.arc** file.

The **config.arc** file includes the following definition. However, the device **/dev/rzz0** does not exist.

```
DEVICE xyz = /dev/rzz0
```

ON-Archive does not check the validity of the device when you define a vset. The following command creates a vset on device_type **xyz**:

```
Onarchive>define/vset=testdb/class=system/driver=tape/access=0/device_type=xyz
```

```
Vset testdb created.
```

However, when you attempt to define a volume, ON-Archive must search for the device to prepare the volume. The volume definition fails, as follows:

```
Onarchive> define/volume/vset=testdb:
```

```
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/disk/path

```
Enter the device number (1) or (C)ancel : 1
```

```
Please mount a volume on /disk/path and press < RETURN > ** WRITE **
```

```
ARC-E-01564, Cannot open tape.
```

```
UNIX-E-00002, No such file or directory
```

You can choose a different device or correct the definition of the device in the **config.arc** file.

A Sample Session

The following sample session illustrates the example from this manual. The first part of the session shows an archive and a logical-log backup. The second part shows a restore of **dbspace1** and **dbspace2**, followed by a restore of the logical-log files.

Log in as **informix**, create spaces, and set group and permissions:

```
> su informix
Password:
# cd /tmp
# touch rootdbs
# touch space1
# touch space2
# chmod 660 rootdbs
# chmod 660 space1
# chmod 660 space2
# chgrp informix rootdbs
# chgrp informix space1
# chgrp informix space2
```

List the new spaces to verify owner, group, and permissions:

```
# ls -lg spac*
```

```
-rw-rw---- 1 informix informix      0 Apr 21 16:09 space1
-rw-rw---- 1 informix informix      0 Apr 21 16:09 space2
```

Set environment variables appropriately for OnLine:

```
setenv INFORMIXDIR /infmtxdir
setenv PATH $INFORMIXDIR/bin:$PATH
setenv ONCONFIG onconfig.archtest
setenv INFORMIXSERVER arch_qstart
```

Verify environment variables (only a partial list is shown here):

```
% env
```

```
.....
PATH=/infmxdir/bin:/home/frame/bin::/usr/ucb:/bin:/usr/bin
INFORMIXDIR=/infmxdir
ONCONFIG=onconfig.archtest
INFORMIXSERVER=arch_qstart
...
```

Move to **\$INFORMIXDIR** and save a copy of **config.arc**. Edit **config.arc** and enter the **DEVICE** definitions as described on page 7. (The editing interaction in **vi** is not shown.)

```
% cd $INFORMIXDIR/etc
% cp config.arc config.arc.save
% vi config.arc
```

Prepare the **ONCONFIG** file as described on page 8. (The editing interaction in **vi** is not shown.)

```
% vi $ONCONFIG
```

Initialize OnLine, including the disk space:

```
% oninit -i
```

```
This action will initialize INFORMIX-OnLine;
any existing INFORMIX-OnLine database will NOT be accessible.
Do you wish to continue (y/n)? y
```

Verify that OnLine is active:

```
% onstat -b
```

```
RSAM Version 7.01.UC1  -- On-Line -- Up 15:36:43 -- 8856 Kbytes
. . .
```

Create two new dbspaces:

```
% onspaces -c -d dbspace1 -p /tmp/space1 -o 0 -s 1000
```

```
Verifying physical disk space, please wait ...  
Dbpace/BL0BSpace successfully added.
```

```
% onspaces -c -d dbspace2 -p /tmp/space2 -o 0 -s 1000
```

```
Verifying physical disk space, please wait ...  
Dbpace/BL0BSpace successfully added.
```

Verify that the dbspaces are active. Notice the rightmost column:

```
% onstat -d
```

```
RSAM Version 7.01.UC1 -- On-Line -- Up 00:01:08 -- 8856 Kbytes
```

Dbspaces

address	number	flags	sfchunk	nchunks	flags	ownername
80d9a0	1	1	1	1		Ninformixrootdb
80d9e4	2	1	2	1		Ninformixdbspace1
80da28	3	1	3	1		Ninformixdbspace2
. . .						

Create two new databases, one in each of the new dbspaces (The messages output from **dbaccessdemo7** are not shown.):

```
% dbaccessdemo7 -dbspace dbspace1 -log databaseOne  
. . .  
% dbaccessdemo7 -dbspace dbspace2 -log databaseTwo  
. . .
```

The basic environment for testing ON-Archive is ready.

Start the archive process by starting the cataloger:

```
% start_oncatlgr
```

```
Informix ONCATLGR Version 7.01  
Starting the cataloger.  
  
See file '/tmp/oncatlgr.out.1597' for results.
```

Look at the last few lines of the `oncatlgr.out.1597` file:

```
% cat /tmp/oncatlgr.out.1597
```

```
ONCATLGR utility For INFORMIX-Online.  
Installation procedure successful on 25-APR-1994:10:21:14.  
Process ID = 1597 Servernum = 32.
```

Start ON-Archive:

```
% onarchive
```

```
ON-Archive Version 7.01  
Copyright (c) 1993 - Informix Software, Inc.
```

You begin the archive process by defining a volume set:

```
Onarchive>DEF/VSET=arcvset/CLASS=SYSTEM/DRI=TAPE/ACCESS=0/DEVICE_T=tape.
```

```
Vset arcvset created
```

Now you can create a volume:

```
Onarchive> DEFINE/VOLUME/VSET=arcvset
```

```
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst0

If you defined several devices in the **config.arc** file, you have a number of choices. In this case, only **/dev/rst0** was defined for device **tape**, so the only choice is 1:

```
Enter the device number (1) or (C)ancel : 1
Please mount a volume on /dev/rst0 and press < RETURN > ** WRITE **

ARC-I-01098, Not an ONARCHIVE volume.
```

The tape in the tape device is probably a brand new tape and is not yet formatted as an ONARCHIVE volume. The response from ON-Archive is reasonable. Just continue and let **onarchive** format the tape:

```
Proceed and overwrite contents, Cancel operation, or Retry
with another volume?
Enter (P/C/R): p
Volume number 0001 created in arcvset.
```

Now that a volume is defined, you can do an archive:

```
Onarchive> archive/dbspaceset=*/vset=arcvset
```

```
Request 00000001 registered in the catalog.
```

Execute the request:

```
Onarchive> execute/request=1
```

```
Executing...
Devices available for ONARCHIVE
DEVICE #      DEVICE NAME
-----
1             /dev/rst0
Enter one of the device numbers (1-2) or (C) ancel : 1
Please mount volume 0001 of arcvset and press < RETURN > ** WRITE **
File rootdbs being processed.
File dbspace1          being processed.
File dbspace2          being processed.
```

The data is archived.

Define a second volume set for a continuous backup of the logical-log files:

```
Onarchive>define/vset=logvset/class=system/driver=tape/access=0/-
_Onarchive>device_type=ltape
```

```
Vset logvset created.
```

Create a volume in the logvset volume:

```
Onarchive> define/volume/vset=logvset
```

Devices available for ONARCHIVE

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst1

Enter the device number (1) or (C)ancel : 1
Please mount a volume on /dev/rst0 and press < RETURN > ** WRITE **

ON-Archive: Media Identification Function

Tape Information:

Vset : logvset
Volume : 000001

RC-W-01845, Initializing this media may overwrite an ON-Archive volume.

Proceed and overwrite contents, Cancel operation, or Retry with another volume?

Enter (P/C/R): p
Volume number 0001 created in logvset.

At this point you must move to another window or terminal and log in as user **informix. Set your environment variables and then start ON-Archive:**

```
% onarchive
```

Start the continuous backup of the logical logs:

```
Onarchive> backup/logfile/continuous/vset=logvset
```

```
Request 00000002 registered in the catalog.
```


Execute the request:

```
Onarchive> execute/request=2
```

```
Executing...
Devices available for ONARCHIVE

DEVICE #          DEVICE NAME
-----          -
      1              /dev/rst1

Enter the device number (1) or (C)ancel : 1
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```

At this point ON-Archive might list the log files and prompt you to specify a device again:

```
LF00000001 being processed.
LF00000002 being processed.
LF00000003 being processed.
Devices available for ONARCHIVE

DEVICE #          DEVICE NAME
-----          -
      1              /dev/rst1

Enter the device number (1) or (C)ancel : 1
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```

Enter the device number 1 again:

After ON-Archive has finished writing to tape, terminate the continuous logical-log backup with CTRL-C:

```
^C
```

```
ARC-I-01483, No data to archive.
```

The backup procedure is finished and you are back at the UNIX prompt. Label the archive and backup tapes and store them safely.

Begin restoring the databases. Go back to the first window.

Enter the retrieve request:

```
Onarchive> retrieve/dbspaceset=*/dbspace=(dbspace1, dbspace2):
```

Request 00000003 registered in the catalog.

Execute the retrieve request:

```
Onarchive> execute/request=3
```

Executing...

Any online dbspaces being retrieved will first be taken offline.
Do you want to Proceed or Cancel this request? (P/C)p

ON-Archive retrieves the dbspaces:

File dbspace1 being retrieved.
Devices available for ONARCHIVE

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst0

Please mount volume 0001 of arcvset and press < RETURN > ** READ **
File dbspace2 being retrieved.

Enter a request to back up the log file that was current when you began the warm restore. Back up the log file to the **logvset volume set:**

```
Onarchive> backup/logfile/automatic/current/vset=logvset
```

Request 00000004 registered in the catalog.

Execute the backup request:

```
Onarchive> execute/request=4
```

```
Executing...
```

```
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst1

```
Enter the device number (1) or (C)ancel : 1
```

```
Please mount volume 0001 of logvset and press < RETURN > ** WRITE **
```

Enter a request to restore the logical-log files:

```
Onarchive> retrieve/logfile
```

```
Request 00000005 registered in the catalog.
```

Execute the request to restore the logical-log files:

```
Onarchive> execute/request=5
```

```
Executing...
```

```
File LF00000003 being retrieved.
```

```
Devices available for ONARCHIVE
```

DEVICE #	DEVICE NAME
-----	-----
1	/dev/rst1

```
Enter the device number (1) or (C)ancel : 1
```

```
Please mount volume 0001 of logvset and press < RETURN > ** READ **
```

When ON-Archive finishes restoring the logical-log files, you are finished with the operation of restoring your OnLine data.

To terminate the **onarchive** program, enter `exit` at the `Onarchive>` prompt as follows:

```
Onarchive> exit
```

To terminate the **oncatlgr** program, execute the **stop_oncatlgr** script:

```
% stop_oncatlgr
```

```
Stopping ONCATLGR  
Kill 7221? (y/n)
```

Enter `y` to terminate the **oncatlgr** program.

End of sample session.