

# **ON-Archive**

## **Quick Start Guide**

Version 9.2  
September 1999  
Part No. 000-6215

Published by Informix® Press

Informix Corporation  
4100 Bohannon Drive  
Menlo Park, CA 94025-1032

© 1999 Informix Corporation. All rights reserved. The following are trademarks of Informix Corporation or its affiliates, one or more of which may be registered in the United States or other jurisdictions:

Answers OnLine™; C-ISAM®; Client SDK™; DataBlade®; Data Director™; Decision Frontier™; Dynamic Scalable Architecture™; Dynamic Server™; Dynamic Server™, Developer Edition™; Dynamic Server™ with Advanced Decision Support Option™; Dynamic Server™ with Extended Parallel Option™; Dynamic Server™ with MetaCube®; Dynamic Server™ with Universal Data Option™; Dynamic Server™ with Web Integration Option™; Dynamic Server™, Workgroup Edition™; Dynamic Virtual Machine™; Enterprise Decision Server™; Formation™; Formation Architect™; Formation Flow Engine™; Gold Mine Data Access®; IIF.2000™; i.Reach™; i.Sell™; Illustra®; Informix®; Informix® 4GL; Informix® InquireSM; Informix® Internet Foundation.2000™; InformixLink®; Informix® Red Brick® Decision Server™; Informix Session Proxy™; Informix® Vista™; InfoShelf™; Interforum™; I-Spy™; Mediazation™; MetaCube®; NewEra™; ON-Bar™; OnLine Dynamic Server™; OnLine/Secure Dynamic Server™; OpenCase®; Orca™; PaVER™; Red Brick® and Design; Red Brick® Data Mine™; Red Brick® Mine Builder™; Red Brick® Decisionscape™; Red Brick® Ready™; Red Brick Systems®; Regency Support®; Rely on Red BrickSM; RISQL®; Solution DesignSM; STARindex™; STARjoin™; SuperView®; TARGETindex™; TARGETjoin™; The Data Warehouse Company®; The one with the smartest data wins.™; The world is being digitized. We're indexing it.SM; Universal Data Warehouse Blueprint™; Universal Database Components™; Universal Web Connect™; ViewPoint®; Visionary™; Web Integration Suite™. The Informix logo is registered with the United States Patent and Trademark Office. The DataBlade logo is registered with the United States Patent and Trademark Office.

Documentation Team: Evelyn Eldridge, Ann McKillop, Patrice O'Neill

#### GOVERNMENT LICENSE RIGHTS

Software and documentation acquired by or for the US Government are provided with rights as follows:

- (1) if for civilian agency use, with rights as restricted by vendor's standard license, as prescribed in FAR 12.212;
- (2) if for Dept. of Defense use, with rights as restricted by vendor's standard license, unless superseded by a negotiated vendor license, as prescribed in DFARS 227.7202. Any whole or partial reproduction of software or documentation marked with this legend must reproduce this legend.

# Table of Contents

About This Guide . . . . .	1
Overview . . . . .	1
Goals . . . . .	1
Assumptions . . . . .	2
Working on a Production System . . . . .	2
Terminology. . . . .	3
Setting Up Your Environment . . . . .	4
Verifying That the Tape Drives Are Ready . . . . .	4
Logging in as User informix . . . . .	6
Making Space for New dbspaces. . . . .	6
Setting Environment Variables . . . . .	6
Preparing the config.arc File . . . . .	7
Preparing the ONCONFIG File . . . . .	8
Starting the Database Server . . . . .	8
Creating dbspaces. . . . .	8
Creating Databases . . . . .	9
Archiving Data Using ON-Archive . . . . .	10
Starting the Cataloger . . . . .	10
Starting onarchive. . . . .	11
Defining a Volume Set . . . . .	11
Defining One or More Volumes . . . . .	13
Creating and Executing an Archive Request. . . . .	14
Backing Up Logical-Log Files Using ON-Archive . . . . .	16
Defining a Volume Set and a Volume . . . . .	17
Starting a Continuous Backup of the Logical-Log Files . . . . .	19
Ending the Continuous Backup of the Logical-Log Files. . . . .	21
Deleting Data from the Database. . . . .	22

Performing a Warm Restore . . . . .	23
Gathering the Appropriate Volumes . . . . .	23
Ensuring That Needed Devices Are Available . . . . .	24
Restoring Selected dbspaces . . . . .	24
Backing Up the Logical-Log File That Was Current When the Warm Restore Began . . . . .	26
Restoring the Logical-Log Files . . . . .	27
Exiting ON-Archive . . . . .	29
Preparing to Rerun the Exercise . . . . .	29
Problem Solving . . . . .	30
ON-Archive Error Message Format . . . . .	30
Typing Errors . . . . .	31
Tape Problems. . . . .	32
If a Request Fails . . . . .	33
Incorrect Permissions . . . . .	34
Cataloger Dead . . . . .	34
Problem Defining a Volume . . . . .	35
A Sample Session . . . . .	36

## About This Guide

This guide presents a complete example of how to set up, back up, and restore data for a sample Informix Dynamic Server 2000 database server. It introduces you to the concepts, terminology, and basic tasks involved in using ON-Archive as your recovery system. It includes a section on troubleshooting common problems and a scripted sample session.

This guide does not describe full-scale archiving for a database server. It does not cover the more advanced features of ON-Archive, such as partial archives and parallel archives. Those topics are covered in the [Archive and Backup Guide](#).

---

## Overview

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 guide, Informix made the following assumptions about your environment:

- You are already familiar with your database server.
- You are archiving only one instance of your database server.
- Your database server 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 you sample commands that you use to perform various operations. For 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 your database server
- Two tape devices, named **tape** and **ltape**
- Two dbspaces, named **dbspace1** and **dbspace2**
- Two copies of the **stores\_demo** database, named **databaseOne** and **databaseTwo**

## Working on a Production System

When you perform 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 database server instances
- Performing a cold restore
- Restoring while users are active
- Using multiple tapes for the archive
- Restoring from multiple logical-log backup tapes

## Terminology

This guide uses the following terminology:

<i>archive</i>	refers to the saving of database data.
<i>back up</i>	refers to the saving of logical logs.
<i>restore</i>	refers to the copying of data from an archive storage medium to the active storage medium that your database server uses.

---

## Setting Up Your Environment

The next sections describe the steps that 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 your database server.
8. Create dbspaces.
9. Create databases.

## Verifying 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 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 archive or back up 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. One option on many UNIX systems is the **mt** command. The **mt** command has the following syntax:

```
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 receive 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.

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

## Logging in as User informix



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

## Making 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.

To prepare the files, enter the following commands:



```
# 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
```

## Setting Environment Variables



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

## Preparing the config.arc File

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

When you install a new instance of your database server, 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 syntax is for 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 that 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 and then 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**.

## Preparing the ONCONFIG File

Start an instance of your database server specifically for this exercise. Before you start your database server, you must prepare an ONCONFIG file. You can follow the basic procedure given for the learning environment in your [Administrator's Guide](#). However, set the TAPEDEV and LTAPEDEV parameters to something other than **/dev/null**, as the following example illustrates:

```
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:** For information about the undesirable effects of setting the tape-device parameters to **/dev/null**, refer to the “[Archive and Backup Guide](#).”

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

## Starting the Database Server

Before you start your database server, make sure you that have an entry for the **dbservername** in the **sqlhosts** file. You can then start your database server with the following command:

```
oninit -i
```



## Creating dbspaces

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

Enter the following commands:



```
% onspaces -c -d dbspace1 -p /tmp/space1 -o 0 -s 1000
% onspaces -c -d dbspace2 -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.

## Creating Databases

Put a database that uses logging in each of the dbspaces that you just created. An easy way to perform this task is to make copies of the **stores\_demo** database.

Enter the following commands:



```
% dbaccessdemo -dbspace dbspace1 -log databaseOne
% dbaccessdemo -dbspace dbspace2 -log databaseTwo
```

For more information on the **dbaccessdemo** command, see your [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.

The following sections describe each of these steps.

### Starting the Cataloger

To start the **oncatlgr** utility, your database server 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 7.03
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-JAN-1998:11:10:07.
```

## Starting onarchive

Once you start the cataloger, you can start the **onarchive** utility. 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 executes, ON-Archive displays the following ON-Archive prompt.

```
Onarchive>
```

Commands typed at the `Onarchive>` prompt can be in either 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:** For an explanation of how to abbreviate ON-Archive commands, see [“Typing Long Commands” on page 32](#).


## Defining 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 you start or continue 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, issue the `DEFINE/VSET` command at the ON-Archive prompt. To define a volume set named **arcvset**, enter the following command:



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

The text that follows `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>	automatically selects a vset depending on the access 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 that you defined in the <b>config.arc</b> file.



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.
```

## Defining One or More Volumes

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



```
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 that is associated with the correct physical device. In the preceding example, the user wants to mount the volume on the device **/dev/rst0** and types **1**.

ON-Archive then prompts you to mount a volume (insert a tape in 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 you are using this tape with ON-Archive for the first time, 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):
```



To respond to this message, type **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.

## Creating and Executing 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 that your database server manages, enter the following command:



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

The asterisk that follows the DBSPACESET qualifier indicates that you want to perform a full-system archive. The VSET qualifier is set to `arcvset`, which indicates 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, which is 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 a specific user owns, or that have a particular status.

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

```
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 [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 your database server 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.

## Defining 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.

To create a vset for the backup, enter the following command at the Onarchive> prompt:



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

ON-Archive responds with the following message.

```
Vset logvset created.
```

To create a volume for the vset, enter the following command at the Onarchive> prompt:



```
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
```



To select **/dev/rst1**, enter 1. 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
```

To continue the operation and overwrite the data on the tape, enter **P** to proceed. To cancel the operation, enter **C**. 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 the following section explains.

## Starting 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 [Archive and Backup Guide](#) describes how to use **onautovop**.



You must create a second window or terminal to perform 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 end 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 Dynamic Server 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 similar to the following example.

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

To execute the request, enter the following command at the **Onarchive>** prompt, replacing 2 with the request number that ON-Archive just displayed:



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

## Starting 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 :
```



To select device **/dev/rst1**, enter 1.

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 the following example shows, 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
```





To select device **/dev/rst1**, enter 1.

ON-Archive prompts you again 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 end the continuous backup. The next section explains how to do end the continuous backup.

## Ending 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.

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

To end the continuous backup of the logical-log files, press the Interrupt key (CTRL-C).

ON-Archive displays the keystrokes that you entered and the message that the following example shows.

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

```
%
```

THE UNIX shell then ends *both the continuous backup and the **onarchive** process*, and returns you to the operating-system prompt.

When you end a continuous backup of the logical-log files, your backup is incomplete if ON-Archive is writing a log file to tape at the time. The [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.

## Deleting Data from the Database

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

You can, however, delete some data from your sample database—**databaseOne** or **databaseTwo**—and see how the operations that follow restore it.

You can use the DB-Access utility to delete data from your database. To use DB-Access, open a new window, 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 database without violating the referential constraints that the databases contain. For a description of the **stores\_demo** database, which is the source for the databases you are using, see the [Informix Guide to SQL: Reference](#). You can use DB-Access to enter and run one or more of the following statements. For information on how to use that product, see the [DB-Access User's Manual](#).



```
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. Restart ON-Archive and continue with the warm restore as described in the following section:



```
% onarchive
```

---

## Performing a Warm Restore

This section explains how to perform a *warm restore* by using ON-Archive. A warm restore restores selected noncritical dbspaces and blobspaces while your database server 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 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. For information on how to perform an *unattended* warm restore, see the [Archive and Backup Guide](#).

### To perform an attended warm restore by using ON-Archive

1. Gather the volumes that contain the archives and logical logs that you need to restore.
2. Ensure that the database server 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.

## Gathering the Appropriate Volumes

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

## Ensuring That Needed Devices Are Available

You must ensure that the raw disk spaces or files from which you archived your data are still available when you perform the restore. The following procedures assume that the dbspaces you archived earlier are still available.

## Restoring 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 your data. Dynamic Server prohibits access to the dbspaces that you select until the warm restore is complete, and any logical-log records for the selected spaces are rolled forward.

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



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

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

To execute the request, enter the following command at the `Onarchive>` prompt and replace 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 Dynamic Server dbspaces being retrieved will first be taken offline.
```

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



To proceed, enter **P**.

Dynamic Server displays the name of the first dbspace that it restores and then prompts you to enter the device number of the tape drive, as the following example shows.

```
File dbspace1 being retrieved.
Devices available for ONARCHIVE
```

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

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



To select device **/dev/rst0**, enter **1**.

ON-Archive prompts you to mount volume **0001** of the **arcvset** vset as the following example shows.

```
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 that you selected. As the following example shows, ON-Archive lists each dbspace it retrieves from the tape and returns to the **Onarchive>** prompt when it finishes.

```
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.

## Backing 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.

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



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

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

To begin the backup, enter the following command at the `Onarchive>` prompt and replace 4 with the request number that ON-Archive just displayed:



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

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

```
Devices available for ONARCHIVE
```

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

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



To select **/dev/rst1**, enter 1.

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 that you restored.

## Restoring 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.

Your database server 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 that contains the checkpoint that the archive generated.
2. The log file that contains the checkpoint that the warm restore generated.
3. All the log files between step 1 and step 2.

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

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

To request the logical restore, enter the following command at the `Onarchive>` prompt:



```
Onarchive> RETRIEVE/LOGFILE
```

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

To begin the logical restore, enter the following command at the `Onarchive>` prompt and replace 5 with the request number that ON-Archive displayed:



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

As the following example shows, ON-Archive retrieves the first log file that must be restored, in this case `LF00000003`, 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
```



To select **/dev/rst1**, enter **1**.

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 the prompts that you just entered for these first prompts.

When ON-Archive finishes, it displays the `Onarchive>` prompt as follows and your database server makes the restored dbspaces and blobspaces accessible to users.

```
Onarchive>
```

At this point, you are finished with the operations in this guide. You can exit ON-Archive as the next section describes.



## Exiting ON-Archive

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

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



```
Onarchive> exit
```

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



```
%stop_oncatlgr
```

The **stop\_oncatlgr** script displays messages, such as the ones that follow, except that the process ID, in this case 4044, varies.

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



To end the **oncatlgr** process, enter `y`.

---

## Preparing to Rerun the Exercise

To repeat this exercise, you must return the system to its original state.

### To prepare to rerun the exercise

1. Set your environment variables for the instance of the database server 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. To remove the database server, execute the **onmode -k** command.
4. Return to [“Starting the Database Server” on page 8](#).

---

## Problem Solving

This section describes a few errors that you might make as you begin using ON-Archive. [“A Sample Session” on page 37](#) shows the screens as a typical user works through the example given in this guide.

### ON-Archive Error Message Format

For more information about ON-Archive error messages, see the [Archive and Backup Guide](#). A typical ON-Archive error message looks similar to the following example.

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

The error messages in the [Archive and Backup Guide](#) are sorted numerically. When you look up an error in the guide, 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 guide, 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 (/) (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 equals (=) 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

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

To avoid this problem, retype the line or make the window wider.

### *Typing Long Commands*

It is easy to make mistakes typing long commands. To reduce mistakes, shorten the length of the command. When you enter a command, 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

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 receive the following message.

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

You also receive 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. To re-execute the request, make a new request with a status of NEW. You can re-create the entire request, or you can use a MODIFY/COMMAND command. The following command creates a new request that is 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 you create the new request, you still need to re-execute it. Request execution is explained in [“Creating and Executing 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 the following example shows:

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

## Cataloger Dead

If you exit from ON-Archive and restart your database server, you must stop the cataloger and then restart it. If you do not follow this procedure, you might have an active but nonfunctional cataloger. In the following interaction, a cataloger from a previous session was 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.
```

To check that your current cataloger is still running, find its process number first in the **oncatlgr** output file. The following example illustrates how to find the **oncatlgr** process number:

```
% 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 Dynamic Server 2000.  
Installation procedure successful on 25-APR-1994:10:21:14.  
Process ID = 1597 Sernum = 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 an incorrect **DEVICE** definition in the **config.arc** file actually caused it.

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 guide. 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 user **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 your database server:

```
setenv INFORMIXDIR /infmxdm
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 your database server, including the disk space:

```
% oninit -i
```

```
This action will initialize Informix Dynamic Server 2000;
any existing Dynamic Server database will NOT be accessible.
Do you wish to continue (y/n)? y
```

Verify that your database server 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 ...
Dbospace/BLOBSpace successfully added.
```

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

```
Verifying physical disk space, please wait ...
Dbospace/BLOBSpace successfully added.
```

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

```
% onstat -d
```

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

```
Dbspaces
addressnumberflagsfchunknchunksflagsownername
80d9a01 1 1 1 N informixrootdb
80d9e42 1 2 1 N informixdbspace1
80da283 1 3 1 N informixdbspace2
. . .
```

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

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

The basic environment for testing ON-Archive is ready.

To start the archive process, start 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 Dynamic Server 2000.
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.
```

To begin the archive process, define 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 new tape and is not yet formatted as an ONARCHIVE volume. The response from ON-Archive is reasonable. 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 perform 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) cancel : 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 finishes writing to tape, end 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. Return 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 **
```

To restore the logical-log files, enter a request:

```
Onarchive> retrieve/logfile
```

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

To restore the logical-log files, execute the request:

```
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 restoring your data.

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

```
Onarchive> exit
```

To end the **oncatlgr** program, execute the **stop\_oncatlgr** script:

```
% stop_oncatlgr
```

```
Stopping ONCATLGR
```

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

Enter `y` to exit the **oncatlgr** program and end the sample session.



---

# Index

---

## C

CANCELLED status 33  
Continuous logical-log backup  
    ending 21

---

## D

DEFINE/VOLUME command,  
    example 13  
DEFINE/VSET command  
    description of 12  
Defining volumes 13  
Device number of physical  
    device 13

---

## F

FAILED status 33

---

## I

Information on requests 16

---

## L

LIST/REQUEST command  
    mentioned 16  
Logical-log file, and  
    ON-Archive 16

---

## M

Mounting a volume 14

---

## O

ON-Archive  
    and logical-log files 16  
    restoring data 23  
onarchive utility, starting 11  
oncatlgr utility  
    starting manually 10

---

## R

Requests  
    information about 16  
    status after failure 33  
Restoring data using ON-  
    Archive 23

---

## S

Starting  
    onarchive utility 11  
    oncatlgr utility 10

---

## T

Terminology 3

---

## V

Volume  
    defining 13  
    gather (warm restore) 23  
Volume set  
    defining 11  
vset 11

