

Administering Rational ClearQuest MultiSite Beta Release

IMPORTANT NOTICE

COPYRIGHT

Copyright ©1977-2001, Rational Software Corporation. All rights reserved.

Part Number: 800-024551-000

Revision Number: 2001A.04.00

PERMITTED USAGE

This Document Contains Proprietary Information Which Is The Property Of Rational Software Corporation ("Rational") And Is Furnished For The Sole Purpose Of The Operation And The Maintenance Of Products Of Rational. No Part Of This Publication Is To Be Used For Any Other Purpose, And Is Not To Be Reproduced, Copied, Adapted, Disclosed, Distributed, Transmitted, Stored In A Retrieval System Or Translated Into Any Human Or Computer Language, In Any Form, By Any Means, In Whole Or In Part, Without The Prior Express Written Consent Of Rational.

TRADEMARKS

Rational, Rational Software Corporation, the Rational logo, Rational the e-development company, ClearCase, ClearCase Attache, ClearCase MultiSite, ClearDDTS, ClearQuest, ClearQuest MultiSite, DDTS, Object Testing, Object-Oriented Recording, Objectory, PerformanceStudio, PureCoverage, PureDDTS, PureLink, Purify, Purify'd, Quantify, Rational Apex, Rational CRC, Rational Rose, Rational Summit, Rational Visual Test, Requisite, RequisitePro, SiteCheck, SoDA, TestMate, and The Rational Watch are trademarks or registered trademarks of Rational Software Corporation in the United States and in other countries. All other names are used for identification purposes only, and are trademarks or registered trademarks of their respective companies.

Microsoft, the Microsoft logo, ActiveX, BackOffice, the BackOffice logo, Developer Studio, the Microsoft Internet Explorer logo, the Microsoft Office Compatible logo, Microsoft Press, the Microsoft Press logo, Microsoft QuickBasic, MS-DOS, MSDN, the Office logo, SourceSafe, Visual C++, Visual J++, Visual SourceSafe, Visual Studio, the Visual Studio logo, Win32, Win32s, Win64, Windows, the Windows CE logo, the Windows logo, Windows NT, the Windows Start logo, and XENIX are trademarks or registered trademarks of Microsoft Corporation in the United States and other countries.

FLEXIm and GLOBEtrotter are trademarks or registered trademarks of GLOBEtrotter Software, Inc. Licensee shall not incorporate any GLOBEtrotter software (FLEXIm libraries and utilities) into any product or application the primary purpose of which is software license management.

Portions Copyright ©1992-2001, Summit Software Company. All rights reserved.

PATENT

U.S. Patent Nos. 5,193,180 and 5,335,344 and 5,535,329 and 5,835,701. Additional patents pending.

Purify is licensed under Sun Microsystems, Inc., U.S. Patent No. 5,404,499.

GOVERNMENT RIGHTS LEGEND

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the applicable Rational Software Corporation license agreement and as provided in DFARS 277.7202-1(a) and 277.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1988), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 227-14, as applicable.

WARRANTY DISCLAIMER

This document and its associated software may be used as stated in the underlying license agreement. Rational Software Corporation expressly disclaims all other warranties, express or implied, with respect to the media and software product and its documentation, including without limitation, the warranties of merchantability or fitness for a particular purpose or arising from a course of dealing, usage, or trade practice.

Printed in the U.S.A.

Part number: 800-024551-000

Contents

Preface	iii
About This Manual	iii
ClearQuest Documentation Roadmap	iv
Typographical Conventions	iv
Online Documentation	v
Contacting Rational Software	v
1 Introducing ClearQuest MultiSite	1
Understanding the architecture of ClearQuest MultiSite	2
Using clans	2
Using families	3
Using sites	4
Understanding mastership	5
Synchronizing replicas in a family	6
ClearQuest MultiSite, time, and time zones	6
Conflict resolution	7
Database operations and the oplog	8
Tracking Operations for Each Replica	9
Epoch Numbers	11
Optimization and the Epoch Number Matrix	12
2 Deploying ClearQuest MultiSite	17
Basic overview of deploying ClearQuest MultiSite	18
Planning your replication strategy	18
Synchronization Pattern	18
Synchronization Schedule	23
Defining a transport mechanism	24
Setting up mastership for records	24
Creating database replicas	25
3 Transport mechanisms with ClearQuest MultiSite	27
Transferring Packets with Store-and-Forward	28
Packet Storage During the Export and Import Phases	29
Packet Transport	29
Disk space needed for ClearQuest MultiSite	30
Configuring the MultiSite Control Panel	31

	Differentiating Packets with Storage Classes	31
	Sending Files That Are Not Packets	32
	Setting Up an Indirect Shipping Route	33
	Retries, Expirations, and Returned Data	33
	Automated Synchronization	35
	Scheduling on windows	35
	Using ClearQuest MultiSite through a Firewall	36
	Using E-Mail	36
	Using FTP	37
	Using Custom Software	38
	Installing Store-and-Forward on a Firewall Host	38
	Firewall Issues	38
	Controlling Ports Used by albd_server and shipping_server	39
	Specifying Port Values on Windows	40
4	ClearQuest MultiSite Command Set	41
	Location of ClearQuest MultiSite programs	41
	Multiutil use	41
	multiutil Subcommands	42
	Specifying Replicas in Commands	44
	ClearQuest API Functions Related to ClearQuest MultiSite	44
5	Creating Replicas	45
	Understanding the replication process	45
	Activating the database	46
	Creating a replica	47
	Importing a replica	48
	Creating empty vendor databases	48
	Importing the replica	49
	Replica-Creation Scenario	49
	Prerequisites	50
	Activating the database set	50
	Export Phase	50
	Transport Phase	51
	Import Phase	51
	Creating a replica using store-and-forward	52
	Communication Between Replica Hosts	53
	Limiting the Size of a Packet	53

6	Synchronizing replicas	55
	Understanding synchronization	56
	Planning synchronization	57
	Designing an Update Strategy	59
	Assumption of Successful Synchronization	60
	Applying packets that include schema updates	61
	Synchronizing a replica using store-and-forward	61
	Communication Between Replica Hosts	62
	Limiting the Size of a Packet	63
	Transport Options	63
	Automating synchronization	63
	Synchronization scenario	64
	Export Phase	64
	Transport Phase	64
	Import Phase	64
7	Managing mastership	65
	Understanding mastership	66
	Changing mastership	67
	Using mastership with records	68
	Understanding MultiSite system fields and record types	69
	Allowing users to access the ratl_mastership field	70
	Using the ratl_mastership field to change record mastership	71
	Using mastership with user administration	71
	Using the chmaster command to change object mastership	71
	Transferring mastership of a record	72
	Transferring mastership of a Workspace item	72
	Transferring Mastership of User Administration	72
	Changing mastership of a working schema repository	73
	Forcing a transfer of mastership	73
	Fixing an accidental mastership change	74
8	Managing replicas	75
	Understanding replica management	76
	Displaying properties of a replica	76
	Changing the host name for a replica	77
	Renaming a replica	77
	Moving a replica	78
	Moving a shipping server host	78

Moving a user database replica	79
Moving a replicated schema repository	79
Removing a replica	80
Remove replica scenario	81
Scrubbing Parameters for Database Replicas	83
9 Troubleshooting MultiSite operations	85
Replica-creation problems	85
Export Phase	85
Synchronization export problems	86
Cannot Find Oplog	86
Packets accumulate in outgoing storage bay	86
Replica cannot update itself	87
Transport problems	88
Error Messages	88
Invalid destination	89
Delivery fails	89
Shipping Order Expires	90
Synchronization Import Problems	90
Packets accumulate in incoming storage bay	90
Packet is not applicable to any local replicas	91
Miscellaneous problems	92
Recovering from lost packets	92
Restoring replicas	94
Restoring a replica from backup	94
Cleaning up from accidental deletion of a replica	96
1 MultiSite Command Line Reference	97
Command summary	97
activate	99
chepoch	101
chmaster	104
chreplica	108
describe	110
dumpoplog	112
lsepoch	115
lspacket	117
lsreplica	120
mkorder	124

mkreplica	128
MultiSite Control Panel	143
recoverpacket	148
rename	151
restorereplica	152
rmreplica	156
scruboplog	158
shipping_server	160
syncreplica	163

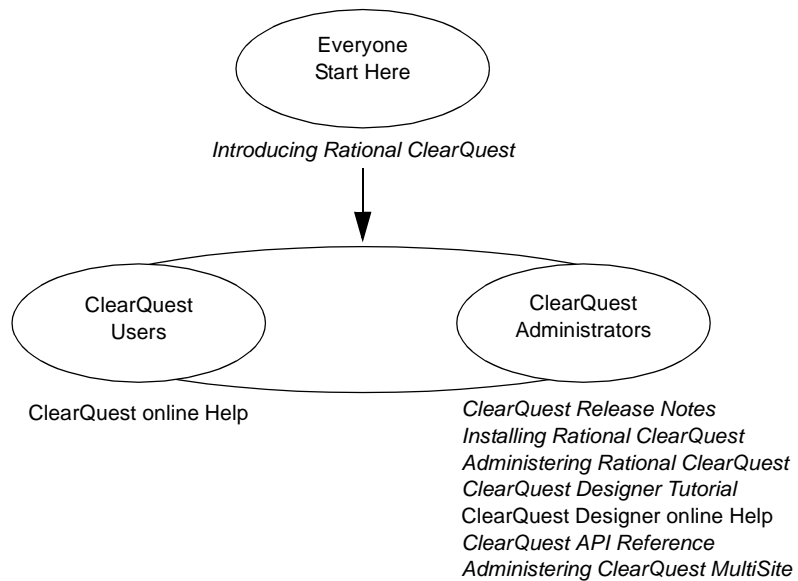
Preface

ClearQuest® MultiSite® (abbreviated to “MultiSite” in this manual) is a layered product option to Rational ClearQuest. It supports parallel software development and software reuse across project teams distributed geographically and provides automated, error-free replication of databases and transparent access to all database records.

About This Manual

This manual is for all MultiSite administrators. It assumes you have experience with ClearQuest. The manual provides an overview of MultiSite, describes how to set up and use it, and gives troubleshooting suggestions.

ClearQuest Documentation Roadmap



Typographical Conventions

This manual uses the following typographical conventions:

- *clearquest-home-dir* represents the directory into which the ClearQuest Product Family has been installed. By default, this directory is `C:\Program Files\Rational\ClearQuest` on Windows.
- **Bold** is used for names the user can enter; for example, all command names, file names, and branch names.
- *Italic* is used for variables, document titles, glossary terms, and emphasis.
- A monospaced font is used for examples. Where user input needs to be distinguished from program output, **bold** is used for user input.
- [] Brackets enclose optional items in format and syntax descriptions.

- {} Braces enclose a list from which you must choose an item in format and syntax descriptions.
- | A vertical bar separates items in a list of choices.
- ... In a syntax description, an ellipsis indicates you can repeat the preceding item or line one or more times. Otherwise, it can indicate omitted information.
- If a command or option name has a short form, a “medial dot” (·) character indicates the shortest legal abbreviation. For example:

scl·ass

This means that you can truncate the command name to **scl** or any of its intermediate spellings (**scla** and so on).

Online Documentation

MultiSite provides a help system that includes an online version of this manual and context-sensitive help for the MultiSite Control Panel and the MultiSite graphical interfaces on Windows.

MultiSite provides access to reference pages (detailed descriptions of MultiSite commands, utilities, and data structures) with the **multiutil man** command.

The **multiutil help** command displays individual subcommand syntax:

multiutil help lspacket

Usage: `lspacket [-long | -short] [pname ...]`

Contacting Rational Software

If you have a technical problem and you cannot find the solution in this guide or in the online Help, contact Rational Software Technical Support. Select Reference > Contacting Technical Support in the ClearQuest Help for addresses and phone numbers of technical support centers.

Introducing ClearQuest MultiSite

1

ClearQuest MultiSite[®] adds a powerful capability to Rational ClearQuest. With ClearQuest MultiSite, developers at different locations can use the same schema and user database. Each location (*site*) has synchronized copies (*replicas*) of the user database and its respective schema repository. At any time, a site can propagate the changes made in its particular replicas to other sites, by sending update packets. The update process can be scheduled by any number of methods to be automatic or can be started manually.

An organization can use ClearQuest MultiSite to distribute independent, but related development efforts across multiple cities, nations, or continents. For example, a company in the United States has development and testing sites in India, Argentina, Japan, and Australia. Because it is impractical for all engineers to access the ClearQuest database in the United States, the company uses ClearQuest MultiSite to distribute the database.

Understanding the architecture of ClearQuest MultiSite

Before you start using ClearQuest MultiSite, you should be familiar with the architecture of ClearQuest.

In summary, ClearQuest uses database sets to organize data. A database set consists of a schema repository and each of its associated user databases. A schema repository holds the metadata that defines user databases (schemas), while a user database is where record information is stored and is what users access to submit and modify records. One schema repository can support multiple user databases. For more information about ClearQuest architecture, see the *Administering Rational ClearQuest* manual.

ClearQuest MultiSite replicates both a schema repository and an associated user database when it creates a replica site. Consequently, when you work with ClearQuest replicas, you are working with two database replicas, a replicated schema repository and a user database replica.

A replica is a named set of one replicated schema repository and one user database replica. Replicas reside at physical sites.

ClearQuest MultiSite organizes replicas three ways:

- clans
- families
- sites

Using clans

A clan represents all the replicas of a single schema repository and its associated databases. Within a clan, a working schema repository is used to propagate changes to schemas and to create additional user databases which can then be replicated.

Understanding the role of the working schema repository

Within a ClearQuest MultiSite clan, there are two types of sites: a working schema repository site and a replicated schema repository site. Each site can perform different tasks and has a different role within MultiSite.

The working repository site within a clan is used to propagate changes to schemas and to create additional user databases which can then be replicated. Only the working schema repository can make schema changes and a clan can have only one working schema repository.

Table 1 MultiSite site roles

Type of Site	Supported tasks
replicated schema repository site	Performs all MultiSite operations Can submit new records Modifies records of which it has mastership Administers users of which it has mastership
working schema repository site	All supported tasks of the replica site, in addition to the following: Modifies schemas Creates new ClearQuest databases

Using families

Within a clan, replicas are grouped into replica families. A replica family consists of all the replicas of a specific database.

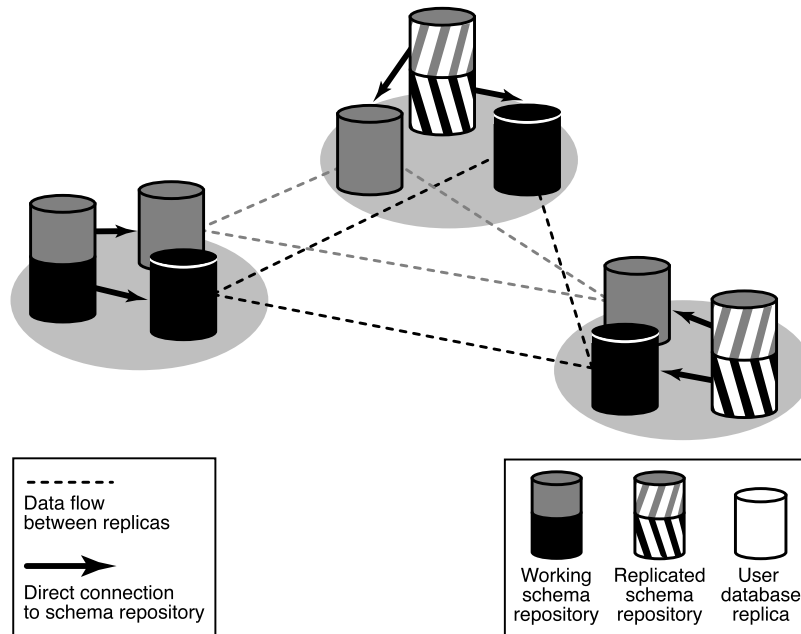
Both schema repositories and user databases participate in replica families. The family name is the same as the ClearQuest database name. The schema repository family name is always MASTR.

Also, when you receive update packets, those packets may contain updates to either the user database replica, the replicated schema repository or both.

Using sites

A site is a named collection of replicas in the same clan that reside at the same location. Each site has a replicated schema repository and at most one replica from each user database family. Each site is served by a host machine, which is responsible for receiving and sending update packets to replicas within its family.

Figure 1 .Diagram of a ClearQuest MultiSite clan



Understanding mastership

In a ClearQuest MultiSite environment, tracking changes and preventing data corruption is accomplished with an exclusive-right-to-modify scheme, called *mastership*. Mastership determines when a user of a database replica is allowed to modify data. For example, without using mastership, users could modify records in different replicas independently and simultaneously, the result would be chaos. Suppose a record SAMPL00001 was modified in three replicas at the same time. Which is the real record SAMPL00001, and what ought to happen to the other two versions?

With mastership, database records and other ClearQuest objects (schemas, queries, etc.) are assigned a *mastering replica*. The initial mastering replica of a ClearQuest object is the site where the object is created. The mastering replica can be changed subsequently (see Chapter 7, *Managing mastership*).

In general, an object can be modified or deleted only at its mastering replica. With database records, mastership information is stored as a field value in a record. To change the mastership of a record (and to allow modifying) users can change the value of the mastership field to the current site. Mastership of the record is then changed during the next synchronizations cycle. All other ClearQuest objects require an administrator to change the mastership.

The following ClearQuest objects have a mastering replica:

- Records
- Users and groups
- Workspace items (queries, reports, charts and folders)
- Schema repository

Synchronizing replicas in a family

Because the data in a replicated database is modified by multiple sites, the contents of each database replica in a *family* tend to diverge. In fact, a particular replica is rarely—and need never be—in the same state as any other replica. To keep the replicas from diverging too much, each site sends updates to one or more other sites. Updating a replica may change both its database and its schema repository to reflect the activity that has taken place in one or more other replicas.

Replica information is exported from a database in packets. A *logical packet* includes all the information required to create a new database replica (replica-creation packet) or to update one or more existing database replicas (update packet). For flexibility, and to accommodate limitations of data-transport facilities, each logical packet can be created as a set of *physical packets* to be copied to disk.

After a logical packet is created and sent to a site, it is processed at that site. The changes that occurred originally at the sending site (and perhaps some other sites, too) are added to the database and schema repository (if needed) of the replica at the receiving site. If the logical packet includes several physical packets, the import commands always process the physical packets in the correct order. No error occurs if the same packet is imported two or more times at a site.

You can match the synchronization strategy for each database family to its particular use patterns, your organization's needs, and the level of connectivity among the sites. For one database, you can update replicas every hour, using a high-speed network; for another database, you can send updates only once or twice a month, using electronic mail, magnetic tape or disk files as the delivery mechanism. See *Planning your replication strategy* on page 18 for information on planning synchronization. Chapter 6, *Synchronizing replicas*, discusses how to create and synchronize database replicas. *Database operations and the oplog* on page 8 describes the underlying mechanism that supports synchronization and replication.

ClearQuest MultiSite, time, and time zones

In ClearQuest MultiSite, time stamps are stored in Universal Coordinated Time (UTC) and are printed to reflect the local time. For example, if a developer in Bangalore modifies a record at 14:33 Bangalore time, the

modification time is stored as 09:03UTC. When a developer in San Francisco looks at the description of the version, the modification time is displayed as 01:03 San Francisco time.

When you automate synchronization, you must adjust schedules for time zone differences. For an example, see *Planning your replication strategy* on page 18.

Conflict resolution

Mastership restrictions prevent most inconsistent changes in different database replicas, but some inconsistent changes are unavoidable, particularly in the naming of new users, groups and other stateless record types. For example, a new user named `jsmith` can be created at two or more replicas during the same time period between synchronizations.

To avoid such naming conflicts, the ClearQuest MultiSite administrators for a family need to create and enforce some naming conventions for database records. A ClearQuest use model that is used consistently across sites reduces the potential for conflicts. For example, the administrators for a family agree that all site-specific records (such as user names) must include a site identifier, or that stateless record types such as user and e-mail rules that will be used at multiple sites are created only at a certain site, which forces the name to be unique.

When naming conflicts are found, ClearQuest MultiSite internally renames the conflicting records. If this happens, you should rename the conflicting records as soon as possible, .

TONIA'S QUESTION: Perhaps the details of detecting naming conflicts should go in a white paper for the first release?

Database operations and the oplog

This section describes the database replica mechanism that supports replica synchronization. This information is not required to use ClearQuest MultiSite, but is helpful when you want to deepen your understanding of the error-recovery facilities. See Chapter 9, *Troubleshooting MultiSite operations*.

For a replicated database, operation log entries (*oplogs*) are created. These entries store all the information required to replay the changes in another replica:

- The identity of the replica where the change originally took place.
- The specific changes to a database record or to a schema in the schema repository made during a single checkout; for example, submission of a new record, schema updates, and so on..
- An integer sequence number: 1 for the first change originating at a particular replica, 2 for the next change, and so on. This is called the *epoch number* of the oplog entry.


The exact kind and amount of information varies with the specific operation. For example, an oplog entry for the submission of a new record has different, and more, information than an oplog entry for modifying an existing record.

Note: Oplog entries are created only for replicated databases. You can delete a replica's oplog entries after they have been used to update other replicas. For more information, see *Scrubbing Parameters for Database Replicas* on page 83.

Tracking Operations for Each Replica

The history of an unreplicated database is a linear sequence of operations, as seen in Figure 2.

Figure 2 History of Changes to an Unreplicated database




For a replicated database, changes are tracked separately for each replica. (That is why an oplog entry includes the identity of the replica where the operation originated.) Thus, the history of a replica can be viewed as several stacks of oplog entries. Each stack is represented by a linear sequence of epoch numbers for the operations originating in that replica.

Figure 3 shows the state of two replicas in a family:

Operations with epoch numbers 1–950 have occurred at replica **lex_hub**.

Operations 1–702 have occurred at replica **sanfran_hub**.


Figure 3 State of a Family



A replica has accurate data only about its own operations. Until it receives update packets, its information about other replicas is out of date. For example, replica **lex_hub** records 950 local operations, but has received update packets for only 504 **sanfran_hub** operations. Similarly, replica **sanfran_hub** records 702 local operations, but has no current data about the **lex_hub** replica's state.

Figure 4 illustrates this scenario, in which each replica is out of date with respect to the operations originating at the other replica.

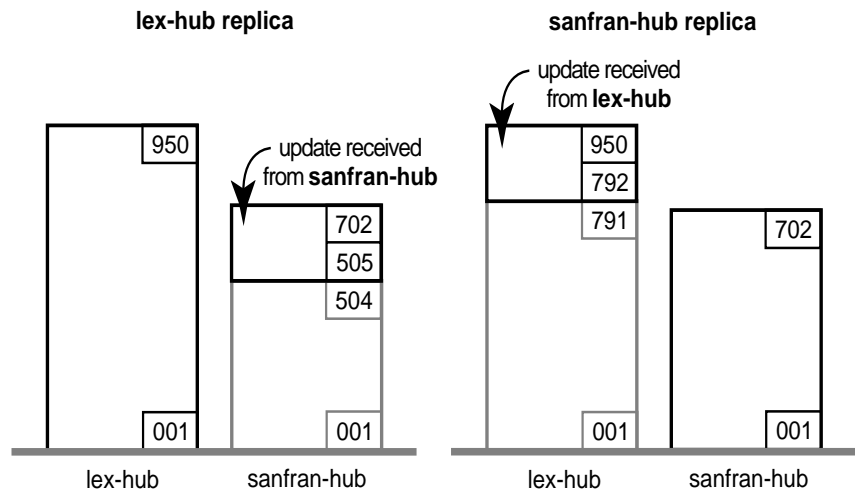
Figure 4 Oplog differences between replicas



Epoch Numbers

Picturing a replicated database as a set of oplog stacks, shown in Figure 4, makes it easy to understand the synchronization process. For example, an update packet sent from replica **lex_hub** to replica **sanfran_hub** consists of increments to the stack for replica **lex_hub** (operations 792–950). Figure 5 shows the two increments. Because **sanfran_hub** knows its own state, it needs updates only for other replicas. (In certain error-recovery situations, you must reset a replica’s data about its own operations. See Chapter 9, *Troubleshooting MultiSite operations*.)

Figure 5 Updates Between Two Replicas



Note: By the time the packet is imported at **sanfran_hub**, additional database changes may have been made at **lex_hub**. These changes are not included in the update packet.

Optimization and the Epoch Number Matrix

The ClearQuest MultiSite synchronization scheme attempts to minimize the amount of data transmitted among sites. Each replica keeps track of the following epoch numbers:

- 1 **Changes made in the current replica.** The epoch number that indicates how many operations originated at the current replica.
- 2 **Changes at sibling replicas.** When **syncreplica** writes an operation from an update packet to the current replica, it increments the epoch number for the sibling replica at which the operation originally occurred. This epoch number is the number of operations originating at the sibling replica that have been imported at the current replica.
- 3 **Current knowledge of the states of other replicas.** For each other replica, an estimate of its own changes and other replicas' changes.

Figure 6 shows how these epoch numbers fall into an *epoch number matrix*. Each replica maintains its own such matrix, revising its rows as work occurs locally and as it exchanges update packets with other replicas:

When work occurs in the **lex_hub** replica, its own number of oplog IDs is incremented.

When the **lex_hub** replica generates an update packet to be sent to **sanfran_hub**, it revises the **sanfran_hub** row in its epoch number matrix.

Note that a **syncreplica -export** command updates epoch numbers immediately. It does not wait for acknowledgment from the receiving site that the packet has been received and applied correctly. During normal ClearQuest and ClearQuest MultiSite processing, no manual intervention is required to maintain the accuracy of the epoch number matrices for the various replicas. However, failure to apply a packet may require manual intervention, as described in *Lost Update Packet* on page 93.

When the **lex_hub** replica receives an update from **sanfran_hub**, it revises its own row (**lex_hub**) and the **sanfran_hub** row in its epoch number matrix.

Figure 6 Two-Row Epoch Number Matrix at Replica `lex_hub`

	VOB operations originated at lex-hub	VOB operations originated at sanfran-hub
this row is lex-hub's record of its own state	950	504
this row is lex-hub's estimate of sanfran-hub's state	912	504

The contents of this matrix are reported by the **multiutil lsepoch** command at the **lex_hub** replica:

```
multiutil lsepoch -clan telecomm -site lex_hub -family PRODA
-user lexadmin -password secret
Multiutil: Estimates of the epochs from each site replayed at
site 'lex_hub' (@host1):
lex_hub: 950
sanfran_hub: 504
Multiutil: Estimates of the epochs from each site replayed at
site 'sanfran_hub' (@host2):
lex_hub: 912
sanfran_hub: 504
```

A **syncreplica -export** command entered at **lex_hub** uses this matrix as follows to generate an update destined for **sanfran_hub**:

- 1 At **lex_hub**, the number of local operations is 950 (number in upper-left corner of matrix), and the estimate is that **sanfran_hub** has been updated only to epoch number 912 (number in lower-left corner).

- 2 The update packet that **lex_hub** sends to **sanfran_hub** includes **lex_hub** oplog entries 913-950. After the Lexington administrator invokes `syncreplica -export`, the **sanfran_hub** row is updated:
- ```
multiutil lsePOCH -clan telecom -site lex_hub -family PRODA
-user lexadmin -password secret
```
- Multiutil: Estimates of the epochs from each site replayed at site 'lex\_hub' (@host1):
- ```
lex_hub: 950
sanfran_hub: 504
```
- Multiutil: Estimates of the epochs from each site replayed at site 'sanfran_hub' (@host2):
- ```
lex_hub: 950
sanfran_hub: 504
```

## Indirect Synchronization

If there are more than two replicas in a database family, synchronization can occur indirectly. A replica can include nonlocal changes in update packets. For example, if **lex\_hub** exchanges updates with replicas **sanfran\_hub** and **bangalore**, it sends **bangalore** oplog entries that it has received previously from **sanfran\_hub**. These entries may or may not bring replica **bangalore** up to date on **sanfran\_hub**'s changes. (An update sent from **sanfran\_hub** to **bangalore** does bring **bangalore** up to date.)

**Note:** If a replica does not receive packets directly from some replicas in the database family, its rows for those replicas may contain zeros. This is expected behavior.

Figure 7 shows replica **lex\_hub**'s epoch number matrix.

**Figure 7** Epoch Number Matrix at Replica Lexington

|                                                            | Operations<br>originated at<br><b>lex-hub</b> | Operations<br>originated at<br><b>sanfran-hub</b> | Operations<br>originated at<br><b>bangalore</b> |
|------------------------------------------------------------|-----------------------------------------------|---------------------------------------------------|-------------------------------------------------|
| <b>lex-hub</b> 's record of<br>its own state               | 950                                           | 504                                               | 653                                             |
| <b>lex-hub</b> 's record of<br><b>sanfran-hub</b> 's state | 912                                           | 504                                               | 653                                             |
| <b>lex-hub</b> 's record of<br><b>bangalore</b> 's state   | 709                                           | 221                                               | 653                                             |

The contents of this matrix are reported by the **lsePOCH** command:

```
multiutil lsePOCH -clan telecom -site lex_hub -family PRODA
-user lexadmin -password secret
```

```
Multiutil: Estimates of the epochs from each site replayed at
site 'lex_hub' (@host1):
lex_hub: 950
sanfran_hub: 504
bangalore: 653
Multiutil: Estimates of the epochs from each site replayed at
site 'sanfran_hub' (@host2):
lex_hub: 912
sanfran_hub: 504
bangalore: 653
Multiutil: Estimates of the epochs from each site replayed at
site 'bangalore' (@host3):
lex_hub: 709
sanfran_hub: 221
bangalore: 653
```

A **syncreplica -export** command at Lexington uses this matrix to export an update for **bangalore**:

- 1 At Lexington, there are 950 local operations (number in upper-left corner of matrix), and the estimate is that **bangalore** has been updated only to epoch number 709 (lower-left corner).
- 2 For operations that originated at **sanfran\_hub**, **lex\_hub** has been updated to epoch number 504, and estimates that **bangalore** has been updated only to epoch number 221.
- 3 The update packet that **lex\_hub** sends to **bangalore** includes **lex\_hub** oplogs 710-950 and **sanfran\_hub** oplogs 222-504. The output of a **multiutil lsePOCH** command at Lexington now looks like this:

```
multiutil lsePOCH -clan telecom -site lex_hub -family PRODA
-user lexadmin -password secret
Multiutil: Estimates of the epochs from each site replayed at
site 'lex_hub' (@host1):
lex_hub: 950
sanfran_hub: 504
bangalore: 653
Multiutil: Estimates of the epochs from each site replayed at
site 'sanfran_hub' (@host2):
lex_hub: 912
sanfran_hub: 504
bangalore: 653
Multiutil: Estimates of the epochs from each site replayed at
site 'bangalore' (@host3):
lex_hub: 950
sanfran_hub: 504
bangalore: 653
```

# Deploying ClearQuest MultiSite

# 2

Deploying ClearQuest MultiSite requires setup tasks additional to installing Rational ClearQuest. You'll need to make decisions about how to transport update packets, what your synchronization strategy will be, and modify your schema to use the ClearQuest MultiSite mastership field. Deploying ClearQuest MultiSite involves the following tasks:

- Basic overview of deploying ClearQuest MultiSite
- Planning your replication strategy
- Defining a transport mechanism
- Setting up mastership for records
- Creating database replicas

## Basic overview of deploying ClearQuest MultiSite

---

Here's an overview of what tasks you'll need to do in order to deploy ClearQuest MultiSite.

- 1 Install ClearQuest MultiSite and configure licensing. For more information, see *Installing Rational ClearQuest* and *Administering Licenses for Rational Software*.
- 2 Plan out your replication strategy, including pattern and schedule.
- 3 Choose and configure your transport mechanism.
- 4 Create database replicas.

## Planning your replication strategy

---

When you plan your replication strategy, you should put thought into what type of synchronization pattern you want to use and what schedule you want to follow.

As you plan your synchronization pattern, remember that each site's update packet includes both its changes and all changes it receives from other sites. With the right strategy, there can be no need for each site to send update packets to every other site within the family.

Rational recommends that you document your plan in writing and implement your design decisions in a set of test replicas before implementing your ClearQuest MultiSite environment.

### Synchronization Pattern

The synchronization pattern for a database family defines which replicas exchange update packets and the direction of exchange. Your choice of pattern depends on the following factors:


- Bandwidth between sites
- Network topology
- Timing: how quickly do changes at one replica need to be received at another replica in the family?
- Failure tolerance

The following sections describe unidirectional and bidirectional exchanges and the most common synchronization patterns.

## Directions of Exchange

Synchronization can be unidirectional or bidirectional, as shown in Figure 8.

**Figure 8 Unidirectional and Bidirectional Updating**




In most cases, you will use bidirectional updating. Unidirectional updates are suitable in situations such as these:

- You use a replica as a backup.
- Your company provides read-only use of your database to another site (or company).
- A high-security development project uses the same files as a more open project. In this case, the open project sends updates to the high-security project, but no updates are sent in the other direction.


However, unidirectional updates carry some risk. For example, an accidental change of mastership cannot be fixed, and restoring from a replica that does not exchange updates directly with the broken replica involves extra work. Also, you must ensure that no work is done accidentally in a read-only replica.

## One-to-One and Ring Synchronization

**Figure 9** One-to-One Synchronization Pattern



**Figure 10** Ring Synchronization Pattern



The simple one-to-one and ring (or round-robin) patterns are simple patterns that are most suitable for small numbers of replicas. As the number of replicas grows larger, the amount of time increases for a change made at one replica to be received at a replica at the other side of the ring.



## One-to-Many Synchronization

Figure 11 Single-Hub Synchronization Pattern






Figure 12 Multiple-Hub Synchronization Pattern



**Figure 13 Tree Synchronization Pattern**



Advantages:


- More efficient for the spoke and branch replicas, which send to and receive from only one other replica.

Disadvantages:

- If the hub or root site goes down, all spoke sites must reconfigure their pattern to keep communication going.
- If you change the synchronization pattern so that replicas that did not synchronize directly now exchange packets, the first packets that are generated may be too large for the system.

## Many-to-Many Synchronization

Figure 14 Many-to-Many Synchronization Pattern



### Advantages:

- For companies with few sites, this pattern keeps each replica's epoch table the most accurate for all siblings.
- If one site is unavailable, the other sites do not have to change their patterns to continue synchronizing.

### Disadvantages:

- Each administrator must maintain more synchronization jobs and spend more time keeping track of packets.

## Synchronization Schedule

The synchronization schedule for a replica family defines when replicas in the family send and receive updates. The schedule is affected by many factors, including the rate of development at different sites and the connections among sites.

Consider the following issues when planning your synchronization strategy:

- Rate of development.  
If you schedule synchronizations frequently, updates take less time as fewer changes have taken place. Data loss is kept to a minimum if a replica is deleted accidentally and you must restore it from backup.
- Time zone differences.

## Defining a transport mechanism

---

There are multiple methods you can use to transport ClearQuest MultiSite packets. The method you choose depends on how your sites are connected, how quickly you must transfer packets, and how important security is. Table 2 lists the recommended methods for various situations.

For detailed information about how to set up your transport mechanism, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.

**Table 2** Choosing a Packet Transfer Method

| Your situation                            | Recommended methods                           | Source of more information                                                                                         |
|-------------------------------------------|-----------------------------------------------|--------------------------------------------------------------------------------------------------------------------|
| Sites are connected with high-speed lines | store-and-forward                             | <i>Transferring Packets with Store-and-Forward</i> on page 28<br><b>shipping_server</b> reference page             |
| One or more sites have firewalls          | E-mail or store-and-forward,                  | <i>Installing Store-and-Forward on a Firewall Host</i> on page 38<br><b>syncreplica</b> reference page             |
| Must transfer packets quickly             | E-mail, store-and-forward or ftp              | <b>shipping_server</b> reference page, <b>syncreplica</b> reference page<br>MultiSite Control Panel reference page |
| No electronic connection between sites    | external storage media such as diskette or CD | <b>syncreplica</b> reference page                                                                                  |

## Setting up mastership for records

---

ClearQuest MultiSite includes a system field, `ratl_mastership`, which tracks the mastership of a record. The value of this field is always the name of the replica which currently masters the respective record.

The mastership field can only be changed at the site of mastering replica. For more information about setting up mastership for records and detailed instructions on adding the field to the record form, see *Using mastership with records* on page 68.

## Creating database replicas

---

Replicating a database involves tasks on both the site that is replicating the database (exporting) and the site that will be receiving (importing) the replica.

At the replicating (exporting) site, the administrator must do the following:

- 1 If you are using the Rational Shipping Server, configure the MultiSite Control Panel to store, send, and receive replica packets, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.
- 2 Activate the respective database set to ClearQuest MultiSite status, if it isn't already activated. This is only necessary when creating the first replica of a clan.
- 3 Replicate the database using **mkreplica -export**.

At the receiving (importing) site, the administrator must do the following:

- 1 If you are using the Rational Shipping Server, configure the MultiSite Control Panel to store, send, and receive replica packets, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.
- 2 Create empty vendor databases for the schema repository and each respective user database replica. See *Installing Rational ClearQuest* for further instructions on creating vendor databases for ClearQuest to use.
- 3 Import the replica using **mkreplica -import**.

For more information about creating replicas, see Chapter 5, *Creating Replicas*.



# Transport mechanisms with ClearQuest MultiSite

# 3

ClearQuest MultiSite allows you to use transport replica creation and synchronization packets several ways. This chapter covers the following topics:

- Transferring Packets with Store-and-Forward
- When you use a file-based method for transport, you may need to use the `-maxsize` option to ensure the file is a manageable size.
- Using ClearQuest MultiSite through a Firewall
- Installing Store-and-Forward on a Firewall Host

## Transferring Packets with Store-and-Forward

The ClearQuest MultiSite *store-and-forward* facility is a file-transfer service that automates the *transport phase*. This facility can handle packets of any size, can route files through a series of ClearQuest MultiSite hosts, one hop at a time, and includes support for handling data-communications failures.

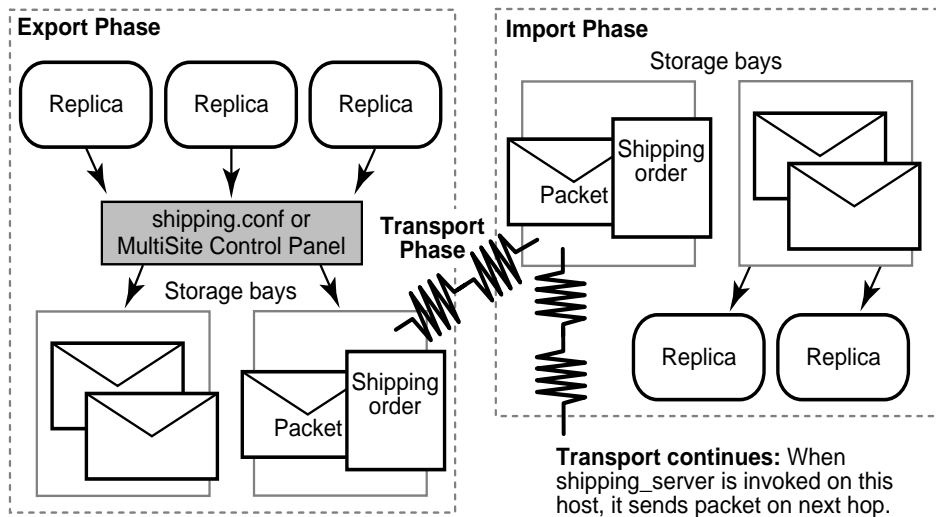
The major components of the store-and-forward facility include:

- MultiSite Control Panel
- Rational Shipping Server

The store-and-forward process is illustrated in Figure 15 and described in the following sections.

**Note:** To use store-and-forward, the sending host must be able to communicate with the receiving hosts. To determine whether the hosts can communicate, use the `rccp` command on the sending host to copy a file to the receiving host. If it fails, you may have to update the `hosts` file, `hosts` NIS map, or Domain Name Service before using store-and-forward.

**Figure 15** The Store-and-Forward Facility





## Packet Storage During the Export and Import Phases

When a physical packet file is exported from a database replica and submitted to the *store-and-forward* facility, it is accompanied by a *shipping order* file, which specifies delivery instructions. The packet is stored in one of the *storage bay* directories on the database replica host.

Each storage bay directory contains two subdirectories, **incoming** and **outgoing**, which hold the incoming and outgoing packets and their corresponding shipping order files. Shipping operations look in the **incoming** and **outgoing** directories for packets. A default storage bay, **ms\_ship**, is created on a host when the Rational Shipping Server is installed there.

**Note:** On Windows, the amount of available space on the disk partition where the shipping bays are located must be at least twice the size of the largest packet that will be stored in the shipping bays. There may be two copies of the same packet in the bay at one time: one on its way to another destination and another waiting to be applied to the replica on the host.

Return bays are similar to storage bays and provide “return-to-sender” storage for packets that could not be delivered successfully. A default return bay, **ms\_rtn**, is created on a host when ClearQuest MultiSite is installed there. This bay has two subdirectories, **incoming** and **outgoing**, which hold the incoming and outgoing packets. Shipping operations look in the subdirectories for packets.

## Packet Transport

The **shipping\_server** program transfers packet files from a storage bay (or return bay) at one site to the corresponding bay at another site.

An explicit command, manual or automated, invokes the **shipping\_server** on the sending host. The **shipping\_server** process contacts the **albd\_server** process on the receiving host, which in turn invokes the **shipping\_server** on the receiving host in receive mode. After a TCP/IP connection has been established between the sending and receiving invocations of **shipping\_server**, the file is transferred.

## Submitting Packets to Store-and-Forward

When you generate a replica-creation or update packet, you can specify that the *store-and-forward* facility must deliver it. Both **syncreplica** and **mkreplica** support the following options:

- The **-fship** option places the packet files and shipping order files in one of the host's storage bays, and runs **shipping\_server** to send the packet files to their destination host or route them to an intermediate host.

The **-ship** option places the packet files and shipping order files in a storage bay, but does not invoke **shipping\_server**. The packet files are sent the next time the **shipping\_server** polls the appropriate bay.

## Disk space needed for ClearQuest MultiSite

Regardless of which transport mechanism you use, you'll need to set up ClearQuest MultiSite storage bays to store and receive update packets from replicated databases. You'll need to allocate space for these directories on a machine. The amount of disk space varies according to the size of the update packets you will be sending.

If you are using the Rational Shipping Server to transport updates, the machine on which the storage bays reside needs to be accessible to the replicas that will be updated and any replicas that will be generating ClearQuest MultiSite packets.

The storage bays hold ClearQuest MultiSite packets, along with their corresponding *shipping order* files. Table 3 describes the amount of available disk space needed on the disk partition where the storage bay is located.

**Table 3 Disk Space Needed for Storage Bay**

| Type of packet   | Disk space needed                                                                                                                                                                                                                                                           |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| replica-creation | At least twice the size of database and schema repository. Packets can be up to four times as large as the databases they are created from.                                                                                                                                 |
| update           | On Windows, twice the size of the largest packet to be stored in bay. The reason is that there may be two instances of the same packet in the bay at one time: one on its way to another destination, and another waiting to be applied to the replica on the current host. |
|                  |                                                                                                                                                                                                                                                                             |

There is no specific formula for determining how large your update packets will be. The general rule is that more frequent synchronization results in smaller packets. However, even if you synchronize every hour, a large amount of database activity or release activity can occur in an hour and can

cause a large packet to be generated. If you are not sure that the available disk space can accommodate an unexpectedly large packet, you can configure ClearQuest MultiSite to limit the size of an update packet.

For more information on configuring the MultiSite Control Panel, see the **MultiSite Control Panel** (Windows) reference pages.

## Configuring the MultiSite Control Panel

You'll need to configure the MultiSite Control Panel before creating or receiving any replica packets. The MultiSite Control Panel controls the operation of the store-and-forward facility on the local host machine. It manages the locations of storage and return bays that store and receive update packets to and from other replicas.

The settings for the store-and-forward facility are host-specific. You can specify locations of storage and return bays, routing information to support multihop packet delivery, specifications to handle failure-to-deliver situations, receipt handlers, and so on.

For more information on specifying settings, see the *MultiSite Control Panel* on page 143 for specific instructions.

## Differentiating Packets with Storage Classes

You can configure the store-and-forward facility to handle updates for different replica families in different ways. Each packet can be assigned to a *storage class*, and each storage class can have its own *storage bay*, *return bay*, and *expiration period*.

You can use multiple storage classes to segregate the packets for replicas belonging to different families. By adjusting the operating system permissions on the storage bay directories, you can protect the packets from unauthorized use. You can also use a separate storage class when you use the store-and-forward facility to transfer non-ClearQuest MultiSite files between sites.

### Creating a storage class

There are a few guidelines you should use when creating a storage class.

- The name of the storage bay should be unique. You cannot use the same name or directory as you do for ClearCase MultiSite packets.
- Replicas within the same database family should use the same storage class name.

- The directory you specify should be on a drive that has sufficient room for the packet size you specify.
- Although you specify the administrator's e-mail address, you'll need to tell ClearQuest MultiSite which e-mail server to use. To do this, see the *Installing ClearQuest* manual or the release notes.
- Host name and info

To view the MultiSite Control Panel,

- 1 Click **Start > Settings > Control Panel**.
- 2 Double-click the MultiSite icon.

For more information on specifying settings, see the *MultiSite Control Panel* on page 143 for specific instructions.

## Receiving error notification

If a packet is delivered through a Windows host on which e-mail notification is not enabled, a failure on that Windows host means that no notification message is sent by electronic mail. Instead, a message is written to the event log; this message contains a request that the appropriate users be informed of the failure. For information on enabling e-mail notification, see the **MultiSite Control Panel** reference page.

## Sending Files That Are Not Packets

You can send any file using the store-and-forward facility if you create a shipping order for the file with the **mkorder** utility. You can send the file immediately or wait for the **shipping\_server** to send it.

- To send a file immediately, use the **-fship** option with **mkorder**:  
example here
- To store the file in a shipping bay so that **shipping\_server** will send the file the next time it runs, use the **-ship** option:  
example here

**Warning:** The shipping order must be located in the same directory as the file.

After you invoke the **mkorder** command, you can delete the original file.

If a file with the same name already exists on the receiving host, the file you send is renamed to *filename\_1*. If you transmit another file with the same name, it is renamed to *filename\_2*, and so on.

## Setting Up an Indirect Shipping Route

The shipping order for a packet includes the host name of the packet's final destination or several such host names. By default, the store-and-forward facility sends the packet directly to its destination host. You can specify that the packet must be sent to an intermediate host by associating it with a routing hop in the **MultiSite Control Panel** (Windows).

For example:

- On a Windows host, the Routing Information section in the MultiSite Control Panel specifies host **sydney-fw** in the **Next Routing Hop** box and hosts **sanfran\_hub**, **lex\_hub**, and **tokyo** in the **Destination Hostnames** box.

Any packet whose final destination is host **sanfran\_hub**, **lex\_hub**, or **tokyo** is forwarded to host **sydney-fw**. At this point, the local host has completed its task, and responsibility for delivering the packet now belongs to **sydney-fw**. Host **sydney-fw** can transmit the packet to its final destination directly, or send it to yet another intermediate host, depending on the settings in its **MultiSite Control Panel**.

**NOTE:** In a multihop transmission, using the **-fship** option on the original host causes the *first* hop to occur immediately. Subsequent hops occur when the Rational Shipping Server is invoked on the intermediate hosts, which may not be immediately after the packets are received.

## Retries, Expirations, and Returned Data

The **shipping\_server** makes one attempt to transmit a packet to another host. If the packet cannot be transmitted successfully (for example, because the receiving host is unavailable), **shipping\_server** generates an error message and log file entry, then exits. Administrators can set up a retry scheme to control its frequency:

- After successful transmission of a packet to another site, **shipping\_server** deletes both the packet and its shipping order. After a transmission failure, a packet and its shipping order remain in the storage bay.

- **shipping\_server -poll** transmits all packets it finds in one or more storage bays. Thus, any packets that remain after a transmission failure are sent (if possible) by the next invocation of **shipping\_server -poll**.

The following job definition performs this operation every hour:

```

Job.Begin
Job.Name: "Shipping Server Poll"
Job.Description.Begin:
Every hour, run the shipping server to send out any outstanding orders.
Job.Description.End:
Job.Schedule.Daily.Frequency: 1
Job.Schedule.FirstStartTime: 00:00:00
Job.Schedule.StartTimeRestartFrequency: 01:00:00
Job.Task: "MultiSite Sync yExport"
Job.Args: -quiet 1 -poll
Job.End

```

Attempts to transmit an undelivered packet can continue indefinitely, through repeated invocations of **shipping\_server**. However, administrators usually want to fix problems with failed transmissions instead of letting the attempts continue. Accordingly, each shipping order can include an expiration date-time, specified with one of the following:

- The command option **-pexpire**. A **Packet Expiration** value in the **MultiSite Control Panel** at the sending host.

By default, shipping orders expire 14 days after they are created.

When **shipping\_server** encounters a shipping order that has expired, it does not attempt to transmit the corresponding packet to its destination. Instead, it does the following:

- It modifies the shipping order to return the packet to the original sending host, where it is placed in a special *return bay*.
- It sends an electronic mail message to one or more addresses on the original sending host. (Another mail message is sent when the returned packet arrives at the original sending host.)

The return trip may involve multiple hops, as described in *Setting Up an Indirect Shipping Route* on page 33. During such a trip, a packet is placed in the return bay of each intermediate host. Each hop is handled by **shipping\_server -poll**, which processes a host's return bay in addition to its storage bays. The expiration time for a packet's return trip is 14 days; a packet that cannot be returned in that interval is deleted.

## Error Notification

# Automated Synchronization

---

You can write ClearQuest MultiSite scripts and utilities to automate all phases of synchronization:

**Export phase.** A ClearQuest MultiSite export script sends update packets from one or more replicas at the site to one or more siblings.

**Transport phase.** The *store-and-forward* facility handles packets of any size. You can invoke store-and-forward as part of the export phase, or automate packet transport separately.

**Import phase.** A ClearQuest MultiSite receipt handler runs whenever a packet is received at a replica.

## Scheduling on windows

You can define receipt handlers in the **MultiSite Control Panel** for different shipping classes.

For details about defining receipt handler entries, see the section *Receipt Handler Path* in the **MultiSite Control Panel** reference page.

## Using ClearQuest MultiSite through a Firewall

---

The ClearQuest MultiSite store-and-forward facility cannot operate through a firewall unless you configure ClearQuest MultiSite differently. Passing through a firewall is usually accomplished by granting access via specific ports and IP addresses. Because store-and-forward picks any available port number on each end to make the connection, there is no single port number (or even small range of port numbers) to which special access can be granted.

This section describes several ways to use ClearQuest MultiSite through a firewall:

- Use an existing electronic mail mechanism as the transport.
- Use the standard **ftp** utility to transport packets.
- Use a custom TCP application.
- Install the store-and-forward software on a host configured to communicate through the firewall.

**Note:** When you use a file-based method for transport, you may need to use the **-maxsize** option to ensure the file is a manageable size.

sample code here

## Using E-Mail

You can use an existing e-mail mechanism as the transport. On the sending end, compress and encode the update packet; then send the resulting data to a specific mail alias at the receiving site. On the receiving end, redirect the mail alias to a script that decodes and decompresses the incoming information. To ensure that a mail message is not too large to be delivered, you can generate packets no larger than a specific size by using the **-maxsize** option or the **MultiSite Control Panel** (Windows).

Advantages:

- Transport mechanism is well understood and widely available.
- Little effort is required from the system administrator.

Disadvantages:

- No control over routing of data.
- Possibility that messages can be lost without notification.
- Messages can be intercepted easily.
- Less efficient than **ftp** or store-and-forward.



## Automating the e-mail process

You can write scripts to automate e-mail transport. The sending script creates the update packets, compresses and encodes them, and divides them into multiple small packets so they are not too big for the e-mail process. The script must mark the multiple packets with the correct sequencing. The script then sends the packets to an address at the target replica.

At the target replica, the account that receives the packets redirects or pipes the packets to a process that reassembles, decodes, and uncompresses the packets, and then places them in the replica's storage bay.

ClearQuest MultiSite import commands handle out-of-sequence and missing packet problems, so your scripts do not have to address these issues.

For security, you must encrypt the packets.

## Using FTP

The **ftp** utility can transport packets. On the sending end, the ClearQuest MultiSite administrator or a script creates and compresses the packet, and uses **ftp** to transfer the file to a location outside the firewall. This location, or dropsite, must be accessible by ClearQuest MultiSite administrators at other sites. Receiving sites poll the dropsite, looking for any new files. When new files arrive, the receiving sites retrieve them via **ftp**, decompress them, and process them as usual.

Advantages:

- Transport mechanism is well understood and widely available.
- More reliable and efficient than electronic mail.

Disadvantages:

- Use of a dropsite is required.
- Polling of the dropsite is required.
- More complicated to implement, due to the interactive nature of the **ftp** utility.
- More administration is required because a third system (the dropsite) is used.

## Using Custom Software

A custom TCP application can accept data and send it from one site to a waiting application at another site. Guidelines for simple applications that send data are often described in the network programming documentation provided by the vendor. If the sending and receiving applications use a fixed port number, the administrator can configure the firewall to permit access.

Advantages:

- Efficient and reliable.
- No dropsites required.
- Electronic mail-capable network is not required.
- Data interception is more difficult.

Disadvantages:

- Custom coding is required.
- Not as flexible as electronic mail or FTP solutions.

## Installing Store-and-Forward on a Firewall Host

---

**NOTE:** Because of security concerns, we recommend that you use this method only if other methods are unsuitable for your site.

An alternative to using mail, **ftp**, or custom software is to install the store-and-forward software on a “firewall host,” a host that can communicate through the firewall. ClearQuest MultiSite synchronization commands can forward data intended for systems on the other side of the firewall to this host. The software on this host then forwards packets through the firewall to the next hop. To specify the range of port numbers to be used on the host, you can use the environment variables `CLEARCASE_MIN_PORT` and `CLEARCASE_MAX_PORT`. In Figure 16, the hosts that communicate through the firewall are the firewall hosts; they have the MultiSite Control Panel software installed on them, but not ClearQuest software. The replica server hosts have ClearQuest MultiSite installed on them.

**Figure 16** Store-and-Forward Configuration

## Firewall Issues

This section describes issues you must consider before installing ClearQuest MultiSite on a firewall host and gives instructions for installation.

Before enabling **shipping\_server** on a firewall host, consider the following issues:

- Shipping bays can be overfilled.

Using **shipping\_server** on a firewall host enables anyone coming in from the network to fill shipping bays anywhere on the local network, on any machine where a **shipping\_server** is available.

To avoid full disks and the related problems:

Ensure that all shipping bays in the local network are on partitions of their own, so that filling the bays does not degrade system performance.

Install **shipping\_server** only on machines that need it: machines used by ClearQuest MultiSite administrators.

- Packets are susceptible to snooping.

In normal update packets, version information is not encoded. Therefore, anyone shipping packets across an unsecured network must encrypt the packets. Also, the format of a update packet is not very complicated; a dedicated programmer could figure out the format and create a packet with operations that damage a database. Encrypting the data makes this kind of attack much more difficult.

- Other servers can be accessible.

Allowing **shipping\_server** access also allows access to all servers created by the **albd\_server**. Because the **albd\_server** assigns port numbers in the allowed range to other servers running locally, programs from the outside network can connect to all of those servers. Therefore, the firewall host that runs the **shipping\_server** must not run other ClearCase servers.

If you can specify the ports to which programs can connect and the IP addresses that are allowed to connect, we recommend that you do so. It further limits the possibility that unauthorized machines can breach the firewall. (You specify ports during the firewall configuration process.)

## Controlling Ports Used by **albd\_server** and **shipping\_server**

The environment variables **CLEARCASE\_MIN\_PORT** and **CLEARCASE\_MAX\_PORT** specify the range of port numbers that the **albd\_server** and **shipping\_server** can allocate for communication purposes. When the server needs to assign a

port number, it starts with the value of `CLEARCASE_MIN_PORT` and continues through the range until it reaches `CLEARCASE_MAX_PORT`. If a port in the range cannot be allocated, the server sleeps and then tries the ports again.

When **shipping\_server** detects that the port environment variables are set, it tries to use TCP to make the connection with the **albd\_server** on the receiving host. If this connection fails, **shipping\_server** tries UDP. Therefore, if you have TCP connectivity, you do not have to enable UDP or open UDP ports on the firewall host.

On Windows, running an individual **shipping\_server** does not require more than two ports at a time.

## Specifying Port Values on Windows

To specify minimum and maximum port values:

On the firewall host, open Control Panel and click the **System** icon.

Create two system environment variables, `CLEARCASE_MIN_PORT` and `CLEARCASE_MAX_PORT`, and specify their values.

This chapter summarizes the commands available with ClearQuest MultiSite and the ClearQuest API functions that display or modify ClearQuest MultiSite information. Reference pages for the ClearQuest MultiSite commands are available in Chapter 1, *MultiSite Command Line Reference*, and are also available online:

- On Windows, the ClearQuest MultiSite **multiutil man** command displays reference pages in Windows Help.

## Location of ClearQuest MultiSite programs

---

The ClearQuest MultiSite installation places programs and configuration files in the ClearQuest installation directory .

## Multiutil use

---

multiutil has a set of subcommands that perform product functions, such as replica creation, synchronization, and management; change of mastership of objects stored in replicas; and failure recovery.

- Command options can always be abbreviated to three characters and sometimes fewer, as indicated in the reference pages.
- You can use **multiutil** in *single-command mode*. For example:  
sample here

Also in *interactive mode*:

- It has online help facilities. The **help** command displays syntax summaries, and the **man** command displays reference pages:

## multiutil Subcommands

The following sections describe the different kinds of **multiutil** subcommands.

**Table 4** Utility, Replica Creation, Synchronization, and Management Commands

| Command      | Description                                                                                     |
|--------------|-------------------------------------------------------------------------------------------------|
| <b>cd</b>    | Changes current working directory                                                               |
| <b>help</b>  | Displays <b>multiutil</b> command syntax                                                        |
| <b>man</b>   | Displays a ClearQuest MultiSite reference page                                                  |
| <b>pwd</b>   | Prints working directory                                                                        |
| <b>exit</b>  |                                                                                                 |
| <b>quit</b>  | Ends interactive <b>multiutil</b> session                                                       |
| <b>shell</b> | Creates subprocess to run shell or program                                                      |
| activate     | Prepares a database set to be used by ClearQuest MultiSite.                                     |
| chmaster     | Transfers mastership of ClearQuest objects including records, users, groups and Workspace items |
| chreplica    | Changes the properties of a replica.                                                            |
| describe     | Describes a replica                                                                             |
| dumpoplog    | Displays the content of a replica's oplog.                                                      |
| lspacket     | Describes contents of packet.                                                                   |
| mkorder      | Creates a shipping order for use by the store-and-forward facility.                             |
| mkreplica    | Creates a database replica.                                                                     |
| rename       | Renames a <i>replica</i>                                                                        |
| rmreplica    | Deletes a replica.                                                                              |
| scruboplog   | Deletes oplog entries.                                                                          |
| syncreplica  | Generates or applies update packets.                                                            |

## Object Mastership

To prevent conflicting changes from occurring at different replicas of a database, certain database objects are assigned a *mastering replica (master)*. The initial master of an object is the replica where the object is created. For more information on mastership, see *Managing mastership* on page 65.

**Table 5 Object Mastership Commands**

| Command  | Description                                 |
|----------|---------------------------------------------|
| chmaster | Transfers mastership of ClearQuest objects. |

## Failure Recovery

Each database uses an *epoch number matrix* to track its own level *and* the epoch level of all other replicas. (Because replicas are always changing, a replica knows what changes have been made to itself; but it can have only an estimate of the states of other replicas.) Each time a replica sends an update packet, it updates its own epoch number matrix, under the assumption that the packet will be delivered to its destinations and applied to the appropriate replicas. For more information, see *Database operations and the oplog* on page 8.

**multiutil** includes the following failure-recovery commands, for use when this assumption of successful delivery does not hold true:

**Table 6 Failure-Recovery Commands**

| Command        | Description                                                                                                                                                                                                                                                                                                    |
|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| chepoch        | Changes epoch numbers.                                                                                                                                                                                                                                                                                         |
| lsepoch        | Lists the current epoch estimates for specified replicas.                                                                                                                                                                                                                                                      |
| recoverpacket  | Resets epoch row table so changes in lost packets are resent (required when a packet is lost or unusable).                                                                                                                                                                                                     |
| restorereplica | Restores a <i>replica</i> from backup. This command places a replica in a special state, in which it sends epoch number matrix corrections to other replicas. The replica cannot be used for normal development work until it receives special updates that inform it of the current states of other replicas. |

## Specifying Replicas in Commands

---

mention how you have to use the family command if there's more than one user database replica at a site.

## ClearQuest API Functions Related to ClearQuest MultiSite

---

You can use the ClearQuest API in hooks and external applications to determine whether or not you are working with a replicated database and if you have mastership of the ClearQuest record or object you want to modify.

**Table 7**

| API | What it does |
|-----|--------------|
|     |              |
|     |              |
|     |              |



This chapter describes how to create database replicas. The topics covered include:

- Understanding the replication process
- Activating a database
- Exporting a replica
- Importing a replica
- Replica-creation scenario

## Understanding the replication process

---

Replicating a database involves tasks on both the site that is replicating the database (exporting) and the site that will be receiving (importing) the replica.

At the replicating (exporting) site, the administrator must do the following:

- 1 If using store-and-forward, configure the MultiSite Control Panel to store, send, and receive replica packets, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.
- 2 Activate the respective database set to MultiSite status using **activate**. This step is only required when making the initial replica of a clan.
- 3 Replicate the database using **mkreplica -export**.

At the receiving (importing) site, the administrator must do the following:

- 1 If using store-and-forward, configure the MultiSite Control Panel to store, send, and receive replica packets, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.
- 2 Create empty vendor databases for the schema repository and each respective user database replica.
- 3 Import the replica using **mkreplica -import**.

**Table 8     Replica creation tools**

| <b>Commands</b> | <b>When used:</b>                                                                                                               |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------|
| activate        | Before you can create the first replica of a clan. This command only needs to be done once,                                     |
| mkreplica       | To create and import replica packets. Administrators must first create a physical database to which to import a replica packet. |

## **Activating the database**

---

Before you can create the initial database replica of a clan, you must first activate the database set (a schema repository and its associated user databases). Once a database set has been activated, multiple families and sites can be replicated. A database set only needs to be activated once.

When you activate a database set to be replicated, you assign it a clan and site name and indicate the host where its update packets will be sent and received.

If you want to allow users to change the mastership of database records from the ClearQuest client, you need to make the appropriate schema changes and upgrade the respective user databases *before* you activate the database set (see Chapter 7, *Managing mastership*).

You activate a database to be replicated with the **activate** command. For more information about activating a database, See the activate reference page on page 99.

## Creating a replica

---

Replicas are made by creating replica packets and then sending them via the network or other transport method to the site which will host the database replica. This is done with the `mkreplica -export` command. For more details about using **mkreplica -export**, See the `mkreplica` reference page on page 128.

After a database set has been activated, you use this three-phase procedure to create new database replicas:

- 1 **Export phase**—At one site, enter a **mkreplica -export** command, which creates new replica objects and a replica-creation packet.
- 2 **Transport phase**—Send the packet to each of the new sites.
- 3 **Import phase**—At each new site, each administrator enters a **mkreplica -import** command, which creates a new database replica.

The basic procedure is the same for all methods of packet delivery and for all platforms.

During the export phase of replica creation, the replica creation command locks the database while copying it. The database is locked for the entire length of time the command runs and logins are not allowed.

Before running the **mkreplica -export** command on a database that has not been replicated, you **MUST** make sure that no users are currently logged into the database. If users are logged into the database during a **mkreplica -export**, any changes they make even after the database is unlocked will be LOST.

The time needed to create the packet depends on the size of the database and can take up to twice as long as it would to make a copy of the database or run a backup procedure.

Therefore, you need to schedule the export phase of replica creation during nonbusiness hours for your site. You must also cancel any scheduled backups for the duration of the export phase.

The following example creates the `paris` replica in the `telecommunications` clan:

```
multiutil mkreplica -export -telecommunications -SAMPL -user
admin -password secret -maxsize 50m -fship c:\temp\packets
-cq_default server3:paris
```

With this example, a replica of the SAMPL user database and its respective schema repository was made for the paris site which uses the host machine, server3. This example uses store-and-forward to ship the replica-creation packets

## Importing a replica

---

When you import a replica, you import raw data into existing vendor database(s). Replica packets do not contain databases, they contain the metadata and record data that form a database. Packets are not vendor-specific and can be used to create a ClearQuest database from any supported vendor database.

To receive (import) a replicated database at your site, you must do the following:

- 1 If using the Rational Shipping Server, configure the MultiSite Control Panel to store, send, and receive replica packets, see Chapter 3, *Transport mechanisms with ClearQuest MultiSite*.
- 2 Create at least two empty vendor database to which to import the replica data, see *Installing Rational ClearQuest*.
- 3 Import the replica by running the **mkreplica -import** command.

## Creating empty vendor databases

Before receiving and importing a database replica, you need to create an empty vendor database(s) in which to import the replica data. See *Installing Rational ClearQuest* for further instructions on creating vendor databases for ClearQuest to use.

**Note:** ClearQuest MultiSite does not support Microsoft Access or SQLAnywhere.

**Warning:** Do NOT create a ClearQuest database before receiving a database replica. A ClearQuest database is automatically created when you import a replica packet into an empty vendor database. If you have created a ClearQuest database with the Maintenance Tool or ClearQuest Designer, the replica import will fail.

## Importing the replica

After verifying that you've received a replica packet, you need to run the **mkreplica -import** command to import the replica data into the empty vendor database you've created. You'll need to enter the database parameters and login information for both the schema repository and user database you are importing.

## Adding additional replicas at the the same site

If you have an existing replica at your site and the replica you want to import is of the same clan, but a different family (it originates from the same working schema repository), you do not need to create a vendor database for the schema repository. If the incoming replica is of the same clan, the **mkreplica -import** command automatically associates the new replica with the existing replicated schema repository. See the **mkreplica** reference page on page 128.

## Replica-Creation Scenario

---

The replica-creation example in this section uses a fictional company whose software development takes place in Boston and in a new development office in San Francisco. Work is about to begin on a new release.

Relevant characteristics of the two replicas:

| <b>Boston Replica</b> |                   |
|-----------------------|-------------------|
| Host name:            | <b>minuteman</b>  |
| Replica name:         | <b>boston_hub</b> |

| <b>San Francisco Replica</b> |                    |
|------------------------------|--------------------|
| Host name:                   | <b>goldengate</b>  |
| Replica name:                | <b>sanfran_hub</b> |

## Prerequisites

Before you create a new replica, you must perform these steps at the original site:

- 1 Make sure ClearQuest MultiSite licenses are installed.

After you enter the **activate** command on a user database, users at the original site cannot access the database without a ClearQuest MultiSite license (in addition to a ClearQuest license).

- 2 When creating a replica of a specific database for the first time, all users must log off of the database.

The **mkreplica -export** command locks the database after you start to export the replica. All users should logoff before the procedure begins and log back on after it is finished. Data is lost if ClearQuest sessions are left open during the replication process.

- 3 Determine the size of the user database and schema repository.

Replica-creation packets can be up to five times larger than the respective databases. Be sure the working directory you use has enough free space. You must have write permission for the directory, and the directory you specify must not exist.

## Activating the database set

- 1 In Boston, activate the database set.

The following example activates the database set (schema repository and its associated databases) at the Boston site. It names the clan (telecomm) and the site (sanfran\_hub) and associates the host machine, minuteman. The host machine is where the storage bays for outgoing and incoming update packets are stored.

```
multiutil activate -user bostonadmin -password secret -clan telecomm -site boston -host minuteman
```

## Export Phase

At the Boston (originating) site, perform the following steps.

- 1 Use the **mkreplica -export** option to create the replica for the Boston site. See the **mkreplica** reference page on page 128 for information about restrictions on the command.

The following example creates the **sanfran\_hub** replica of the **PRODA** user database that participates in the telecommunications clan. The storage bays for this replica reside on the **goldengate** host machine.

Also, the administrator uses the **-fship** option which sends the packet immediately, using the Rational Shipping Server.

```
multiutil mkreplica -export -clan telecommunications -site
boston -family PRODA -user bostonadmin -password secret
-maxsize 50m -fship -workdir c:\temp\packets -sclass
cq_default goldengate:sanfran_hub
```

- 2 Back up the original database.

This backup records the fact that the database is replicated. If you have to restore a database replica from a backup copy that was made before the database was replicated, the ClearQuest MultiSite replica restoration procedure fails. (Although the **restorereplica** command may succeed, you will not be able to import update packets from other replicas because the original database is marked as unreplicated.)

## Transport Phase

- 1 Send the replica-creation packet to the new site. This process differs depending on the options you used during the export process:
  - If you used **-fship** in Step 1, the packet was sent to the new site immediately.
  - If you used **-ship**, you must run **shipping\_server** to send the packet to the new site. For example:

```
shipping_server -sclass cq_default
shipping_server -poll
```
  - If you used **-out** to write the packet to a file, you must transport the file to the new site using the media of your choice.

## Import Phase

These steps take place at the Boston (receiving) site. The Boston site has no existing replicas. The Boston site administrator needs to have created empty databases to which to import the incoming replicas of the schema repository and user database.

- 1 Verify the packet's arrival by entering the **lspacket** command on the receiving host.

By default, **lspacket** searches all the ClearQuest MultiSite storage bays for packets. For example, if host **goldengate** is the receiving host:

sample here

- 2 Create empty vendor databases. The Boston site administrator needs to have created empty databases to which to import the incoming replicas of the schema repository and user database.
- 3 Enter the import form of the replica-creation command.

In the **mkreplica -import** command, you must specify the pathname of the incoming packet as listed by the **lspacket** command. For example:

```
multiutil mkreplica -import -site boston -repository ORC1
-vendor ORACLE -dbologin orcadmin password -connectopts
host=boston_dbserver;SID=ORC1;server_ver:8.1;client_ver=8.
0;log_type=long -database ORC1 -vendor ORACLE ORC1
-dbologin orcuser password -connectopts
host=boston_dbserver;SID=ORC1;server_ver:8.1;client_ver=8.
0;log_type=long -comments ''Importing the initial replicas
of the SAMPL database and its schema repository for the
Boston site in the telecommunication clan''
<packet-file-path>
```

- 4 Delete the replica-creation packet. (Update packets are deleted automatically.)
- 5 Begin development.

Users in Paris can access the new replica in the same way they would access an unreplicated database.

## Creating a replica using store-and-forward

---

If your sites have a high-speed connection, you can take advantage of the ClearQuest MultiSite store-and-forward facility when you create a new replica. If your current site does not have IP connectivity to the site of the new replica, you can use magnetic tape or a file-based packet transfer method like ftp or email.

The following sections describe issues you must consider when you use the store-and-forward method.



## Communication Between Replica Hosts

The hosts must be able to communicate with each other. If your network uses host names, the sending host must be able to resolve the receiving host's name to an IP address. To accomplish this, you may have to update the **hosts** file, **hosts** NIS map, or Domain Name Service. Verify TCP/IP access by using **rcp** on each host to access the other hosts.

**Note:** If hosts in your network are known only by their IP addresses, you can use the IP addresses instead of host names, and no resolution is necessary.

## Limiting the Size of a Packet

The **mkreplica** command fails if it tries to create a packet larger than the size supported by your system. To prevent this problem and improve reliability, use the **-maxsize** option to divide the replica-creation packet into multiple packets:

```
multiutil mkreplica -export -maxsize 1g ...
```

For information on default packet size limits, see the **mkreplica** reference page.



# Synchronizing replicas

# 6

This chapter describes the process of synchronization. Synchronization uses the same export-transport-import procedure that is used during replica creation:

- Understanding sychronization
- Planning synchronization
- Applying packets that include schema updates
- Synchronizing a replica using store-and-forward
- Synchronization scenario

# Understanding synchronization

---

The synchronization process involves the same export-transport-import procedure that is used during replica creation:

- **Export phase**—At one site, a **syncreplica** (synchronize replica) command is invoked with the **-export** option. This creates a *packet* of data.
- **Transport phase**—The packet is sent to one or more other sites.
- **Import phase**—At the other sites, a **syncreplica** command is invoked with the **-import** option. This applies the changes in the packet to an existing replica.

Note that each synchronization is one-way. If two replicas update each other, two synchronizations are required.


The **syncreplica** command is optimized for performance; it creates a packet that contains only the information required to update the target replicas specified on the command line, based on its estimates of what the target replica has already been sent.

**Table 9 Synchronization tool summary**

| Command     | When used                                                                                                                                                                                                                                                                                                                                                    |
|-------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| syncreplica | To create update packets to send out and to import update packets from other replicas.                                                                                                                                                                                                                                                                       |
| chepoch     | Used when you need to recover from a lost packet. Allows you to re-create a packet that has already been sent.<br>Epoch numbers are automatically updated regardless of whether or not a packet is sent and applied successfully. If a packet fails, you'll need to reset the your replica's epoch estimate of the replica where the synchronization failed. |
| lsepoch     | Used when you need to recover from a lost packet. Allows you to list the current epoch estimates for the replicas within a family. You can then determine which epoch number to use if you need to reset an epoch estimate.                                                                                                                                  |

The following illustrates the ClearQuest ClearQuest MultiSite replica-synchronization scheme. At Site 1, a **syncreplica -export** command places database update information from **replica1** into a packet. The packet is sent to Site 2. At Site 2, a **syncreplica -import** command imports the contents of the packet into **replica2**.

**Figure 17 Replica Synchronization**



## Planning synchronization

---

Figure 17 shows a simple case, involving one point-to-point update. All updates need not be point to point, however, because they are cumulative. Suppose that the following updates take place among three replicas:

- Update 1: Replica 1 sends changes to Replica 2
- Update 2: Replica 2 sends changes to Replica 3

There is no need for Replica 1 to update Replica 3 directly, because the changes from Update 1 are included in Update 2. This feature gives administrators flexibility in devising update strategies and patterns. For efficiency, a single update can be targeted at multiple sites, for example, all other replicas in the database family.

In general, you can implement any update topology, as dictated by organizational structures, communications/transportation costs, and so on. Figure 18 shows a simple peer-to-peer synchronization update pattern and Figure 19 shows a double-hub hierarchical pattern.

**Figure 18 Peer-to-Peer Synchronization Pattern**  
peer-to-peer pattern



**Figure 19 Hierarchical Synchronization Pattern**  
hierarchical pattern



## Designing an Update Strategy


Site administrators must design a strategy for sending updates among the various replicas. They must specify an update pattern for the database family and an update frequency for each replica.

For example, the administrators for the database family in Figure 19 make the following decisions:

- The hub replicas, which undergo rapid development, synchronize every hour.
- Each hub replica synchronizes daily with its spoke replicas. Each spoke replica will send an update packet to the hub replica, and then the hub replica will send update packets back to the spoke replicas. Because these packets may be large and take a long time to import, the synchronization should not take place during working hours.
- All sites use receipt handlers to import packets as soon as they are received.

Figure 20 shows the synchronization timeline for the hub-spoke updates (but not the hub-to-hub updates). This timeline accounts for time zone differences and includes extra time to make sure that each synchronization phase completes before another begins.

**Figure 20 A Synchronization Export Schedule**



### Assumption of Successful Synchronization

The export and import phases of synchronization always occur at different times. A sending replica does not require acknowledgment from a sibling replica that a packet has been received and processed successfully. Instead, the sending replica assumes success. This enables an optimization: subsequent updates from a replica do not include the data sent in previous updates.



If a failure does occur (for example, a packet is lost in transit or a diskette is unreadable at the sibling replica's site), the sending site must adjust its records to enable the lost data to be resent.

For more on this topic, see Chapter 10, Troubleshooting ClearQuest MultiSite Operations.

## Applying packets that include schema updates

---

Packets originating from the working schema repository may include new schema revisions that require database upgrades at respective family sites. In this case, the **syncreplica -import** cannot finish its process until the user database replica is upgraded to the new schema version.

If a packet contains updates to both the replicated schema repository and the user database replica, **syncreplica -import** stops the process and gives the following message:

```
text here TONIA's QUESTION
```

In this case, you'll need to upgrade the affected user database replica and then run **syncreplica -import** again.

If you've automated your synchronization process, the automation script will fail and additional packets dependent on the schema changes cannot be applied.

## Synchronizing a replica using store-and-forward

---

If your sites have a high-speed connection, you can take advantage of the ClearQuest MultiSite store-and-forward facility when you synchronizing replicas. If your current site does not have IP connectivity, you can use magnetic tape or a file-based packet transfer method like ftp or email.

The following sections describe issues you must consider when you use the store-and-forward method.

## Communication Between Replica Hosts

The hosts must be able to communicate with each other. If your network uses host names, the sending host must be able to resolve the receiving host's name to an IP address. To accomplish this, you may have to update the Domain Name Service.

**Note:** If hosts in your network are known only by their IP addresses, you can use the IP addresses instead of host names, and no resolution is necessary.

## Limiting the Size of a Packet

The **syncreplica** command fails if it tries to create a packet larger than the size supported by your system. To prevent this problem and improve reliability, use the **-maxsize** option to divide the update packet into multiple packets:

```
multiutil mkreplica -export -clan telecomm -site boston
-family PRODA -user bostonadmin -password secret -maxsize 50m
-fship -workdir c:\temp\packets -scl cq_default sanfran_hub
```

For information on default packet size limits, see the **syncreplica** reference page on page 163.

## Transport Options

When you enter the **syncreplica -export** command, you can use either the **-fship** option to send the packet immediately, or the **-ship** option to store the packet in the outgoing shipping bay. With **-ship**, you must invoke the **shipping\_server** to send the packet.

The outgoing packet is stored in the **outgoing** subdirectory of a storage bay. If you have not customized *store-and-forward*, the default storage bay is used.

The **incoming** and **outgoing** subdirectories of storage bays contain packets waiting for transport or processing. All shipping operations look for packets in these subdirectories.

- At the receiving site, the incoming packet is stored in the **incoming** subdirectory of a storage bay.

## Automating synchronization

You can automate all phases of synchronization with the use of third-party tools:

- **Export phase.** Create a script that runs at scheduled times that creates and sends update packets from one or more replicas at the site to one or more siblings.
- **Transport phase.** The *store-and-forward* facility handles packets of any size. You can invoke *store-and-forward* as part of the export phase, or automate packet transport separately.
- **Import phase.** A ClearQuest MultiSite receipt handler runs whenever a packet is received at a replica.

# Synchronization scenario

---

This section describes how to synchronize replicas by entering explicit **syncreplica** commands.

## Export Phase

- 1 **Create an update packet.** At the sending host, use the **syncreplica -export** command with the appropriate transport option.

If your sites are connected electronically, you can use store-and-forward to send the packet (**-fship**) or place it in a storage bay (**-ship**):

```
multiutil syncreplica -export -clan telecomm -site
sanfran_hub -family PRODA -user sfadmin -password secret
-maxsize 50m -workdir c:\temp\packets -fship -sclass
cq_default bangalore paris
```

## Transport Phase

**Send the packets.** If you did not use, **-fship**, Use electronic mail, regular mail, or your preferred delivery method. If you used **syncreplica -export -ship**, invoke **shipping\_server** in either of the following ways:

```
shipping_server -poll
shipping_server \\goldengate\shipping\outgoing\
```

## Import Phase

- 1 **(If you used diskettes or electronic mail) Copy the packet files into a directory.**
- 2 **Apply the packet.** At the receiving replica, use the **syncreplica -import** command to apply the changes in the packet to the replica.

This example specifies the **-receive** option; **syncreplica** imports any packets it finds in the incoming shipping directories.

```
multiutil syncreplica -import -family PRODA -user parisadmin
-password secret -receive -scl cq_default
\\baguette\shipping\incoming\
```

This example specifies a directory pathname as an argument. **syncreplica -import** looks in this directory for update packets and applies them to the replica.

```
multiutil syncreplica -import -family PRODA -user parisadmin
-password secret \\baguette\shipping\incoming\
```

This chapter describes how to manage the mastership of ClearQuest objects in a database replica, and covers the following topics

- Understanding mastership
- Using mastership with records
- Using mastership with user administration
- Using the chmaster command to change object mastership
- Changing mastership of a working schema repository
- Forcing a transfer of mastership
- Fixing an accidental mastership change

# Understanding mastership

---

In a MultiSite environment, tracking changes and preventing data corruption is accomplished with an exclusive-right-to-modify scheme, called *mastership*. Mastership determines when a user of a database replica is allowed to modify data. For example, without using mastership, users could modify records in different replicas independently and simultaneously, the result would be chaos. Suppose a record SAMPL00001 was modified in three replicas at the same time. Which is the real record SAMPL00001, and what ought to happen to the other two versions?

With mastership, database records and other ClearQuest objects (records, schemas, queries, etc.) are assigned a *mastering replica*. The initial mastering replica of a ClearQuest object is the site where the object is created. The mastering replica can be changed subsequently. When you create an object in a replicated database, your current replica is mastering replica of the new object.

The following ClearQuest MultiSite objects have mastership properties that can be transferred between replicas:

- Records
- Workspace items
- Users and groups
- Schema repositories

**Table 10 Mastership tools**

| Command  | When used:                                                                                                                                              |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------|
| chmaster | Used to change the mastering replica of a particular object. You can only change the mastership of an object from the mastering replica of that object. |
| describe | To find out which replica currently has mastership of the respective ClearQuest object.                                                                 |

## Changing mastership

You can transfer mastership of an object by an explicit **chmaster** command or in the case of records, you can use the `ratl_mastership` field, see *Using mastership with records* on page 68. Some examples of when mastership changes are appropriate:

- You want to make changes to user information that is mastered at a different site.
- You want to allow another site to modify a public query
- Before you decommission a replica, you must transfer mastership of each object mastered by that replica to one of the remaining replicas. (See *Removing a replica* on page 80.)

Mastership changes are communicated among replicas by the standard synchronization mechanism. The general procedure for changing mastership is as follows:

- 1 Notify the administrator of the mastering replica that you need mastership of a database object.
- 2 The administrator at the mastering replica changes mastership of one or more objects to another replica.
- 3 The administrator at the mastering replica site exports and sends an update packet from the old mastering replica to the new mastering replica.

#### 4 Import the update packet at the new mastering replica.

Until the update packet containing the mastership change is imported at the new mastering replica, mastership is “in the packet” and the replicas in the database family have different information about which replica masters the object.

For example, the administrator at the **sanfran\_hub** replica transfers mastership of the user group “QA\_ENGINEERING” to the **bangalore** replica, and then exports an update packet. At this point:

The **sanfran\_hub** replica considers the user group to be mastered by **bangalore**.

The **bangalore** replica considers the user group to be mastered by **sanfran\_hub**.

No one can modify the user group.

When you complete the mastership transfer by importing the update packet at **bangalore**, users at **bangalore** are able to modify the user group “QA\_ENGINEERING”.

## Using mastership with records

---

ClearQuest MultiSite ships with a system field called `ratl_mastership` whose value reflects the current mastering replica of a record. In other words, the value of the `ratl_mastership` field is always the name of the replica which currently masters the respective record.

To allow users to change the mastership of database records, you’ll need to add this field to the respective record form.

The `ratl_mastership` field provides great flexibility in controlling when users can change the mastership of a record. Record mastership scenarios include, but are not limited to:

- Unrestricted ability to change mastership. All users can modify the `ratl_mastership` field and thus change the mastership of any record.
- Restricted ability to change mastership. You can make the mastership field read-only to specific users and groups.
- Hook-driven mastership. You can write a field hook that automatically changes the value of the `ratl_mastership` field (thus the mastering replica) when certain criteria are met.



The mastership field can only be changed at the site of mastering replica.

## **Understanding MultiSite system fields and record types**

The `ratl_mastership` field is a reference field that points to the replica record type. The replica record type is a stateless record type that automatically tracks all replicas in a family. Whenever a new replica is added to a family, its information is stored as a replica record. The replica record type and all replica records are read-only. You can use ClearQuest to query on the replica record type to find out more about the replicas that participate in a given family.

For complete list of which system fields are used for ClearQuest MultiSite, see *Administering Rational ClearQuest, Appendix A*.

## Allowing users to access the ratl\_mastership field

The ratl\_mastership field automatically contains the value of the current mastering replica of a record. To allow users to change this value (thus giving mastership of a record to another replica), you must add the field to an existing form on the respective record type.

By allowing users to access this field (adding it to your record form), any ClearQuest users can change the mastership of any record, if they are at the mastering replica site of respective record.

To do this, you must add the field to the record form of each respective record type. Schema changes must be done at the working schema repository and are sent as update packets to members of the respective replica sites.

You can add the ratl\_mastership field to the schema used by a family at any time.

For each record type in your schema that you want to track mastership, use the ClearQuest Designer modify the schema of the respective replica family. You must do this at the site of the working schema repository.

- 1 Select a record type to which to add the field.
- 2 Choose a form to which to add the field.
- 3 Double-click the form to display it.
- 4 Select the tab to which you want to add the field. You can either add the field to an existing tab on the form or add a new tab to contain the field.
  - Select **Edit > Add Tab** to add a new tab.
- 5 Using the Field List, drag-and-drop the ratl\_mastership field to the appropriate tab.
- 6 Check-in the modified schema.
- 7 Upgrade the appropriate user database to use the new schema.
- 8 Generate an update packet. Update packets automatically contain schema updates, if any. Keep in mind that each site will need to upgrade their replica to use the new schema revision see *Applying packets that include schema updates* on page 61.

## Using the ratl\_mastership field to change record mastership

With the ratl\_mastership field, record mastership is changed with the following steps:

- 1 A user at Site A modifies SAMPL0001 and changes the value of the ratl\_mastership field from Site A to Site B.
- 2 During the next synchronization exchange, Site B's replica is updated and is given mastership of SAMPL0001. Remember that until the next synchronization exchange, mastership of the record is "in the packet" and neither site can modify the record.

## Using mastership with user administration

---

User administrators at replica sites can add users, assign database privileges and create user groups. However, administrators cannot modify users which are not mastered at their site. Although a user's database privileges apply to all sites, whether or not a user record is mastered at a particular site has some impact on what that user can do.

If a user's record is not mastered at a site, that user cannot:

- Edit their user profile
- Save defaults while submitting records
- Assign startup queries
- Add a workspace item to the query menu as a favorite

## Using the chmaster command to change object mastership

---

The following ClearQuest MultiSite objects have mastership properties that can be transferred between replicas:

- Records
- Workspace items
- Users and groups
- Schema repositories

## Transferring mastership of a record

When you use the `chmaster` command, only ClearQuest users with the Super User privileges and access to MultiSite administration tools can change the mastership of a record. In addition, mastership can only be changed at site of the mastering replica.

Because of the dynamic nature of database records, ClearQuest MultiSite provides two ways change the mastership of a record. You can change the mastership of database records by using the `chmaster` command or use the `ratl_mastership` field. For more information about record mastership, see *Using mastership with records* on page 68.

When you create a new record, it is mastered by the replica in which you create it. To transfer mastership of an record to another replica, follow these steps:

- 1 At the mastering replica, enter a **chmaster** command:  
`multiutil chmaster`
- 2 At the old mastering replica, export an update packet to the new mastering replica's site:  
`example here`
- 3 At the new mastering replica, import the packet:  
`example here`
- 4 At the new mastering replica, verify that mastership has been received:

## Transferring mastership of a Workspace item

This section describes how to change mastership of a Workspace item using the **chmaster** command.

## Transferring Mastership of User Administration

This section describes how to change mastership of a user or group using the **chmaster** command.

## Changing mastership of a working schema repository

---

This section describes how to change mastership of a schema repository using the **chmaster** command.

The working schema repository is responsible for modifying schemas and adding new replica families to a clan, see *Understanding the role of the working schema repository* on page 2. If you want to change the site where these tasks are done, you'll have to change mastership of the working schema repository.

When you create a new replica, its replica object is mastered by the replica at which you enter the **mkreplica -export** command. Mastership of the replica object affects replica-modification activities (renaming the replica, changing its properties, or deleting it). You must perform these activities at the replica that masters the replica object.

To transfer mastership of a replica object:

1 Determine where it is.

2 At the working schema repository replica, enter a **chmaster** command:

```
multiutil chmaster -clan telecomm -site paris -family
MASTR -user parisadmin secret boston -working master
```

3 At the old working schema repository, export an update packet to the new mastering replica:

```
multiutil syncreplica -export -family MASTR -u parisadmin
-p secret -workdir c:\temp\shipping -fship -scl cq_default
```

4 At the new working schema repository, import the packet:

```
multiutil syncreplica -import -family MASTR -u bostonadmin
-p secret -receive -scl cq_default
```

5 At the new working schema repository, verify that mastership has been received:

```
multiutil describe
```

## Forcing a transfer of mastership

---

**CAUTION:** Use the **-force** option with **chmaster** only when a replica is no longer available. Using **-force** to grab mastership from an available replica causes irreparable inconsistencies among the replicas in the replica family.

## Fixing an accidental mastership change

---

If a mastership change is made in your replica by mistake, follow these steps to undo the change:

- 1 At your replica, complete the transfer by sending an update packet to the new mastering replica.
- 2 At the new mastering replica, complete these steps:
  - a Import the packet.
- 3 Change mastership back to your replica.
- 4 Export an update packet to your replica.
- 5 At your replica, import the packet.

This chapter describes how to manage existing replicas, including how to delete a replica. For information on creating a replica, see Chapter 5, *Creating Replicas*. For information on enabling requests for mastership in a replica, see Chapter 7, *Managing mastership*.

The following topics are covered in this chapter:

- Understanding replica management
- Displaying properties of a replica
- Changing the host name for a replica
- Renaming a replica
- Moving a replica
- Removing a replica
- Scrubbing Parameters for Database Replicas

## Understanding replica management

---

Replica management involves replica administration tasks that are in addition to any database administration tasks such as creating backups. When working with a database replica, tasks such as moving a database become more complicated as MultiSite settings may need to change.

**Table 11** Replica management tools

| Commands       | When used:                                                                                                                                                                                                                                                                                                      |
|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| chreplica      | To change the location of the host that delivers update packets to a replica. Use the chreplica command to change the host properties.                                                                                                                                                                          |
| lsreplica      | List all replicas in a family or a clan.                                                                                                                                                                                                                                                                        |
| rename         | To change the name of a replica.                                                                                                                                                                                                                                                                                |
| restorereplica | Automatically brings a replica up-to-date after it has been restored from a vendor database backup.<br>You'll need to use this command to request packets from one or more sibling replicas that will include the data necessary to bring the restored replica current.is brought up-to-date.<br>resend packets |
| rmreplica      | When you want to decommission a replica and inform other replicas not to expect additional updates. Before you can remove a replica, you'll need to transfer the mastership of all objects that it masters to another replica(s).                                                                               |

## Displaying properties of a replica

---

Use the **lsreplica** command displays the properties of a replica. For more information see *lsepoch* on page 115.

For example, to display the name, and and host of a replica:

```
lsreplica -clan telecomm -site paris -family PRODA -user
parisadmin -password secret -short
```



## Changing the host name for a replica

---

When you move the Rational Shipping Server and the MultiSite Control Panel or when you rename a replica's host, you must update the host name in the Control Panel. The replica keeps track of the hosts on which the replicas in a database family reside so that the store-and-forward facility can determine how to route updates to the replicas.

To change the host name, use the **chreplica** command. The change is not propagated to other replicas in the database family until you export an update packet from the current replica and the packet is imported at the other replicas. For restrictions, see the **chreplica** reference page.

To change a host name using the **chreplica** command:

```
multiutil chreplica -clan telecomm -site paris -family PRODA
-user parisadmin -password secret -host pserver2 paris
```

## Renaming a replica

---

To change the name of a replica, use the **rename** command. When you rename a replica, the change is made immediately at the current replica. The change is not propagated to other replicas in the replica family until you export an update packet from the current replica and the packet is imported at the other replicas.

You can only change the name of the replica at the current site. For other restrictions, see *rename* on page 151.

To rename a replica using the **rename** command:

```
multiutil -rename -clan telecomm -site paris -family PRODA -user
parisadmin -password '''' paris bordeaux
```

## Moving a replica

---

You can move each of the three components of a replica site:

- the host
- the user database replica
- the schema repository replica

There are some special considerations when you move a replicated database or its associated shipping server host:

- Make sure ClearQuest MultiSite is installed on the new shipping server host.
- If you automated the synchronization process on the old host, you must set up synchronization export and import scripts on the new host.
- After physically moving a user database replica, use ClearQuest Designer to tell ClearQuest where to locate the database (**Database > Update User Database Properties** or **Database > Move User Database**).
- After moving the replica, export update packets to all sibling replicas.

After moving the database replica, change the host name associated with the replica by using **multiutil chreplica -host**. You must enter this command at the mastering replica of the replica you moved.

### Moving a shipping server host

There may be times where you want to use a different machine to send and receive update packets. For example, the host machine you were using has had a hardware failure and needs to be replaced.

In this case, use the following steps:

- 1 Install ClearQuest MultiSite and the Rational Shipping Server on the new machine.
- 2 Use the `chreplica` command to associate the new host name with the replica.

```
chreplica -clan telecomm -site bangalore -family PRODA -user
bangadmin -password secret server3 bangalore
```

## Moving a user database replica

There may be times where you want to move your user database replica to a different location on the network or switch it to use a different vendor database software.

In this case, use the following steps:

- 1 Follow the instructions in the Administrating Rational ClearQuest book when you move your user database. A user database replica can be moved just as a non-replicated database.
- 2 Any host info that changed should be updated using **chreplica**.

## Moving a replicated schema repository

TONIA'S QUESTION: will put this in release notes, waiting for decision from Sheryl.

## Removing a replica

---

This section describes how to remove a replica. You must complete all steps; if you do not, synchronization and mastership problems can occur in other replicas in the database family.

When you remove a replica using the `rmreplica` command, you effectively tell all replicas within its family not to expect any more update packets or to track epoch estimates for that replica. Removing a replica does not delete the database.

Removing a replica requires two synchronization cycles, one to transfer mastership of all of that replica's objects to another replica and one to inform all other replicas that the removed replica is no longer participating in the update process. Because this information can only be communicated using the synchronization process, a replica cannot remove itself because doing so prevents it from creating update packets.

Once a replica is removed from a family, it no longer participates in synchronization activities and ClearQuest MultiSite information is not tracked. It no longer keeps an oplog and it cannot transfer masterhip of any object, including records.

To remove a replica, you must do the following:

- 1 Transfer mastership of all objects to another replica.
- 2 Export and send an update packet from the replica to be removed.
- 3 Have any other replica run **rmreplica** command at their site to remove the respective replica
- 4 Export and send an update packet from the replica which removed the respective replica.
- 5 Run **rmreplica** at the site which is to be removed.

The replica is effectively removed after all replicas have been synchronized.

**Note:** If a replica is deleted mistakenly and you want to restore it from backup, see *Restoring replicas* on page 94. If a replica's storage directory is lost and there is no backup, see *Cleaning up from accidental deletion of a replica* on page 96.

## Remove replica scenario

In this scenario, the replica **tokyo** in the database family **PRODA** is being removed.

### 1 Transfer mastership of all objects to another replica.

- a At the site of the replica to be removed, transfer mastership of all objects mastered by the replica to another replica.

In this example, the administrator transfers mastership of all of its ClearQuest objects, including records to the Boston replica :

```
chmaster -clan telecomm -toyko -PRODA -user bostonadmin
-password secret boston -all
```

The replica that receives the mastership can later transfer mastership to other replicas, if necessary.

If mastership is not transferred for all objects, you must fix the problem and reenter the **chmaster -all -long** command. For an example, see *Transferring Mastership of User Administration* on page 72. If there are problems you cannot fix, another replica can recover from the error by assuming mastership of the objects. For a description of this procedure, see *Cleaning up from accidental deletion of a replica* on page 96.

### 2 Export and send an update packet from the replica to be removed.

The replica to be removed must send its final changes, including any new records and mastership changes, to the replica receiving mastership. The replica to be removed can broadcast its final changes to all other replicas, but it must update the new mastering replica (**boston** in this example).

```
syncreplica -export -clan telecomm -toyko -PRODA -user
toykoadmin -password secret -workdir e:\temp\syncstorage
-fship -sclass cq_default boston
```

### 3 Import the update packet at the replica that is (or will become) the mastering replica of all the ClearQuest objects that are mastered by the replica to be removed.

```
syncreplica -import -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -receive -sclass
cq_default
```

### 4 At another replica, remove the replica.

```
multiutil rmreplica -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret toyko
```

- 5 At the replica where you removed the respective replica, export and send an update packet to the remaining replicas in the family.

This update packet notifies the other replicas of the replica removal.

```
sync replica -export -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -workdir
e:\temp\syncstorage -fship -sclass cq_default sanfransisco
bangalore
```

- 6 At the replica that was removed, run **rmreplica**.

```
multiutil rmreplica -clan telecomm -site toyko -family PRODA
-user toykoadmin -password secret toyko
```

## Removing the last replica of a clan

If you decommission and remove all replicas, the one remaining replica can be switched to a regular ClearQuest database (a non-replicated database and schema repository), and developers no longer need a ClearQuest MultiSite license to access it.

In this case, you'll need to run the **rmreplica** command at the site you want to remove and use the **-dbset** argument to change the database set name associated with the schema repository and its user database(s). There are no mastership concerns because the database will no longer be replicated.

For example,

```
multiutil rmreplica -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -dbset maindatabase
```

## Scrubbing Parameters for Database Replicas

---

When a ClearQuest action or ClearQuest MultiSite command makes a change to a replica, an *oplog entry* is recorded in the replica's database. (See *Database operations and the oplog* on page 8 for more information on this mechanism.) Also, when you export an update packet, an `export_sync` record is created for each target replica. These records are stored in the database replica and are used by the **recoverpacket** command to reset a replica's epoch number matrix.

You can scrub `oplog` entries and `export_sync` records to reclaim disk space, but you must keep them long enough to ensure that you can recover from replica failures and packet losses. The following sections give guidelines for configuring scrubbing frequency.

### Oplog Scrubbing

Oplog entries must be kept in the database for a significant period. In the near term, they are required when the replica generates update packets to be sent to all other replicas. Beyond that, entries may be required to help other replicas recover from catastrophic failures. If no replica can supply these entries, the replica being restored must be re-created. Because of the need to use `oplog` entries during synchronization, your synchronization strategy determines how often `oplogs` can be scrubbed.

By default, an `oplog` entry is never scrubbed. When it is safe to delete `oplog` entries for a replica, follow these steps:

- 1 Coordinate with administrators at other sites to determine how long each site must keep `oplog` entries.

Each site must keep entries for as long as necessary to ensure that `rmreplica` operations complete successfully. The frequency with which you scrub `oplogs` depends on the following factors:

- The pattern of synchronization among replicas in the database family.
- How often the replicas are synchronized.

Frequency of synchronization refers both to how often updates are exported and how often they are imported at other sites. Also, consider setting up a verification scheme so you can ensure that packets are processed successfully at other replicas before any `oplog` entries are scrubbed.

- 2 Use the **scruboplog** command at the replica whose oplog you'd like to reduce. This example scrubs the oplog of the sanfrancisco replica.

```
scruboplog -clan telecommunications -sanfrancisco -family
PRODA -user sfadmin -password secret -before 10/31/01
```



This chapter describes common situations in which running a MultiSite command produces an unexpected result, sometimes accompanied by a warning or error message. The situations fall into these categories:

- **Expected conditions** occur because certain inconsistent changes at different replicas cannot be avoided. In many cases, a MultiSite operation resolves the inconsistency, so you need not take any action.
- **Recoverable errors** are user errors, hardware glitches, and other problems that you resolve by performing a recovery procedure.
- **Serious errors** are problems that may require assistance from Rational Technical Support.

The organization of the descriptions follows the general ClearQuest MultiSite data flow: from replica creation through the phases of replica synchronization—export, transport, and import. This chapter also describes replica restoration and replacement.

For information on changing mastership, see Chapter 7, *Managing mastership*.

## Replica-creation problems

---

Problems with replica creation can occur during the export phase or the import phase.

### Export Phase

If the **mkreplica -export** command finds that a replica with the specified name exists in the database family (`Replica replica-name already exists`), select another name for the new replica, and reenter the **mkreplica -export** command.

If **mkreplica -export -fship** fails while it is transporting the packet, it does not remove the new replica's replica object at the creating site. To complete the replica creation, use **shipping\_server** to transfer the replica-creation packet.

## Synchronization export problems

---

This section describes problems that can occur during the export phase of synchronization.

### Cannot Find Oplog

**syncreplica –export** can fail with the following warning message:

```
Can not find oplog from replica replica-name with id oplog-ID
Gap in oplog entries may indicate missing oplog entries
```

(For more information on oplog entries, see *Database operations and the oplog* on page 8 and *Scrubbing Parameters for Database Replicas* on page 83.)

This error occurs when the sending replica's epoch number matrix does not match its set of oplog entries. For example:

- Before sending an update from **sydney** to **buenosaires**, **syncreplica** checks the epoch number matrix for **sydney**. It determines that the last **sydney** operation sent to **buenosaires** was 3620.
- **syncreplica** finds that oplog scrubbing in the **sydney** database has removed some of the operations that follow 3620. The earliest **sydney** operation remaining in the oplog is 5755.

This discrepancy may be an expected condition. For example, when a replica family changes its update topology, hosts that have not communicated with each other in the past start exchanging update packets. Synchronizing two replicas (**syncreplica –export** followed by **syncreplica –import**) updates epoch number matrix rows for the sending and receiving replicas, but it does not revise the row for any other replica. If two replicas rarely (or never) send updates to each other directly, the relevant rows in their epoch number matrices are quite out of date (possibly consisting of all zeros). This is not a problem, as long as the replicas receive each others' operations indirectly (for example, through a hub replica).

In this case, you must inform **sydney** about the true state of **buenosaires**, information that it has not received through the standard synchronization-update mechanism. This information enables **sydney** to determine which oplog entries must be sent to **buenosaires**.

use recover packet.

### Packets accumulate in outgoing storage bay

Problems with packet delivery are recoverable errors. In many cases, the MultiSite automatic-retry capability recovers from errors.

A replica-creation or update packet submitted to the *store-and-forward* facility for transport to one or more other hosts is accompanied by a *shipping order* file. (A logical packet can include multiple physical packets, each with its own shipping order.) The shipping order typically has an expiration time, determined by one of the following:

- A date-time specified with the **-pexpire** option in the **sync replica** or **mk replica** command that generated the packet (or the **mk order** command that submits an arbitrary file to the store-and-forward facility)
- The Packet Expiration value specified in the **MultiSite Control Panel** on the sending host.

Any number of delivery attempts may take place before the shipping order expires.

## Replica cannot update itself

You can receive the following message during export if you specify the sending replica as a destination:

```
A replica cannot update itself
```

If the sending replica is the only replica you specified, the **sync replica -export** command fails. If you specified other replicas, this message is printed as a warning, and the **sync replica -export** command continues its processing.

## Transport problems

---

This section describes problems that can occur during the transport phase of synchronization.

### Error Messages

The messages in Table 12 are generated by the **mkorder**, **mkreplica**, **shipping\_server**, and **syncreplica** commands.

**Table 12 Shipping Error Messages (Sheet 1 of 2)**

| Error Message                                                                                                                                                         | Meaning                                                                                                                                                                                              |
|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| cannot find a storage bay for class <i>class-name</i> : no such bay specified                                                                                         | No storage bay is assigned to storage class <i>class-name</i> in <b>MultiSite Control Panel</b> .                                                                                                    |
| cannot find a storage bay for class <i>class-name</i> : all applicable bays are either inaccessible or do not contain <i>byte-count</i> free bytes                    | Lack of permission or lack of free disk space prevents use of storage bays for class <i>class-name</i> .                                                                                             |
| cyclic delivery route detected to host <i>hostname</i> (via <i>next-hop-hostname</i> ) for order <i>shipping-order-pname</i>                                          | Sending the file to the <i>next-hop-hostname</i> specified in a <b>ROUTE</b> entry in the <b>Routing Information</b> section in the <b>MultiSite Control Panel</b> yields a circular delivery route. |
| file <i>file-pname</i> does not contain a valid shipping order                                                                                                        | <b>shipping_server</b> attempted to process a file that is not a shipping order.                                                                                                                     |
| for security reasons, shipping order <i>shipping-order-pname</i> cannot be processed: data file <i>file-pname</i> must be in the same directory as the shipping order | A shipping order and its associated packet file must be in the same directory. This security feature prevents transmission of arbitrary files.                                                       |
| giving up trying to return order <i>shipping-order-pname</i> to host <i>hostname</i> (original data file was <i>file-pname</i> )                                      | <b>shipping_server</b> cannot return a packet or other file to its original sending host (for example, because its shipping order expired) and has deleted the shipping order and data file.         |
| ignoring shipping bay <i>storage-bay-pname</i> : <i>reason</i>                                                                                                        | The storage bay directory specified in the <b>MultiSite Control Panel</b> is inaccessible, doesn't exist, and so on.                                                                                 |

**Table 12 Shipping Error Messages (Sheet 2 of 2)**

|                                                                                        |                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|----------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>shipping order <i>shipping-order-pname</i> not found (perhaps previously sent?)</p> | <p>During receipt handler processing, the <b>shipping_server</b> cannot find the shipping order of a packet that is to be forwarded to another host.</p> <p>A <b>shipping_server -poll</b> invocation may have sent the packet already. (If the packet is to be applied to replicas on the host, the imports occur before the packet is forwarded. This leaves a window of opportunity for a scheduled polling operation to send the packet.)</p> |
|----------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

## Invalid destination

The local host's Domain Name Service must list one of the following hosts:

- Destination host
- Next-hop host corresponding to the destination host (defined in the **Routing Information** section in the host's **MultiSite Control Panel**.)

**Note:** If hosts in your network are known only by their IP addresses, you can use the IP addresses instead of host names.

In the absence of such entries, **shipping\_server** fails, because it cannot determine where to deliver the packet. In this case, it writes error messages to its log file.

If the destination host name was misspelled, use the **mkorder** command to create a new shipping order with the correct host name. If a host name is misspelled in a **mkreplica -export** command, the incorrect host name is recorded in the replica . Verify the error with **lsreplica -long**, and correct the spelling with **chreplica**.

## Delivery fails

Each time **shipping\_server** cannot deliver a packet to a valid destination host, it logs error messages:

- In the Windows event viewer. It writes log messages to file `clearquest-home\var\log\shipping_server_log`.
- TONIA'S QUESTION: Is the above location accurate? Do store the same file?

If the problem is temporary (remote host is down, network connections are down, and so on), a subsequent invocation of **shipping\_server -poll** will transmit the packet successfully. If the problem is not temporary, the shipping order may expire eventually.

## Shipping Order Expires

If the **shipping\_server** finds that a shipping order has expired, it attempts to return the packet to the originating host. Also, it sends a mail message to one or more administrators on the original sending host, and sends another mail message when the packet is returned to the original sending host. On Windows, if e-mail notification is not enabled, **shipping\_server** writes a message to the Windows event viewer and records the error in the `clearquest-home\var\log\shipping_server_log` file.

TONIA'S QUESTION: Is the above location accurate? Do store the same file?

Use the **lspacket** command to check the *return bays* on your host. The packet files may have been returned by store-and-forward. If so, try again to deliver the packet:

- Fix the store-and-forward packet-delivery mechanism (for example, by fixing the network connection). Then, use **mkorder** to create a new shipping order for each physical packet file in the return bay.
- If you cannot fix the store-and-forward mechanism, deliver the packet by some other means. For example, copy the packet file to a diskette, and mail the diskette to the remote sites.

If the packet files are not in your host's return bays, they may be in transit. Search for the files immediately, because a packet that cannot be returned to its home host within 14 days is deleted.

## Synchronization Import Problems

---

This section describes problems that can occur during the import phase of synchronization.

### Packets accumulate in incoming storage bay

A recoverable error occurs when an update packet is lost and is not applied at your site. These are the symptoms:

- One or more replicas at your site are not being updated on their regular schedules.

- A **lspacket** command shows unprocessed packets accumulating in the storage bay. These packets depend on the missing packet and cannot be processed.

Verify that a packet is missing and determine which operations are needed:

- 1 Enter a **syncreplica -import -receive** command, which processes all incoming packets in the storage bay in the correct order. If **syncreplica** refuses to process any of them, a packet is missing.
- 2 Enter a **syncreplica -import** command that specifies the oldest packet in the storage bay:

```
multiutil syncreplica -import packet-pathname
Sync. packet packet-pathname was not applied to replica ...
- packet depends on changes not yet received
Packet requires changes up to 872; replica has only 756 from
replica: sanfran_hub
Packet requires changes up to 605; replica has only 500 from
replica: bangalore
```

In this example, one or more update packets are missing, containing operations 757–872 originally occurring at replica **sanfran\_hub** and operations 501–605 from **bangalore**. In general, a packet can contain operations from several replicas; the **syncreplica -import** command fails if operations are missing from *any* replica.

Locate the missing packets. They may be on a magnetic tape that you forgot to process or in packet files that were not processed because your store-and-forward configuration (the **MultiSite Control Panel**) specifies the wrong storage bay. If you locate the missing packets:

- 1 Process the missing packets by naming them in a **syncreplica -import** command. (Multiple packet files are imported in the correct order, regardless of the order of the command-line arguments.)
- 2 Process all the update packets that have accumulated in the storage bay by entering a single **syncreplica -import -receive** command.

If you cannot locate the missing packets, go to *Recovering from lost packets* on page 92.

## Packet is not applicable to any local replicas

Import can fail with the following message:

```
multiutil: Error: Sync. packet pathname is not applicable to
any local replicas
```

This error can occur when a replica has been moved and the hostname property has not been updated with the **chreplica** command.

- 1 Verify that the hostname property is wrong, use the **lsreplica** command.

For example, if the above error occurred at the bangalore replica:

```
multiutil lsreplica -site bangalore -user bangaloreadmin -p
secret -long bangalore
```

- 2 If the hostname is incorrect, use the **chreplica** command to change it. For this example, run **chreplica** at the **bangalore** replica:

```
multiutil chreplica -site bangalore -user bangaloreadmin -p
secret -host shiphost2 bangalore
Updated replica information for "bangalore".
```

- 3 Send an update packet to the other replicas in the family.

## Miscellaneous problems

Processing of an incoming replica-creation or update packet may fail because of these conditions:

- Disk partition is full.
- Receiving replica is locked.
- ClearQuest or MultiSite licensing failure.
- Multiple imports occurring simultaneously.

Make sure that multiple **syncreplica -import** commands do not run in the same replica simultaneously.

In such cases, fix the problem and reenter the **syncreplica -import** command.

## Recovering from lost packets

---

There are several circumstances in which a replica-creation or update packet is generated but is never applied at one or more of its destinations:

- The packet is stored on a magnetic tape or diskette that is destroyed or is not readable at the destination host.
- A packet file is lost when a hard disk fails.
- The packet is intact, but cannot be applied because another packet has been lost. (See *Packets accumulate in incoming storage bay* on page 90.)



## Lost Replica-Creation Packet

To recover a lost replica-creation packet:

- 1 At the replica where the **mkreplica -export** command was entered, remove the new replica with **rmreplica**.
- 2 Reenter the **mkreplica** command.

## Lost Update Packet

The **syncreplica -export** command assumes successful delivery of the update packet it generates. For example, when replica **lex\_hub** sends an update to replica **sanfran\_hub**, the **syncreplica** command assumes that the operations originating at **lex\_hub** are imported to the **sanfran\_hub** replica. For simplicity, this example does not reflect the fact that the update packet can also contain operations that originated at other replicas in the database family.

However, if the packet is lost, this assumption is invalid, and **lex\_hub** must reset its estimate of the state of replica **sanfran\_hub**. After this correction is made, the next update packet sent from **lex\_hub** to **sanfran\_hub** contains the operations **sanfran\_hub** needs.

To reset the epoch row, use the method described here.

- 1 At the receiving site, use the **lsepoch** command to display the replica's epoch number matrix:

```
multiutil lsepoch -clan telecom -site sanfran_hub -family
PRODA -user sfadmin -p secret
```

- 2 Use this output in a **chepoch** command at the sending site:

```
multiutil chepoch -family PRODA -user sfadmin -p secret
sanfran_hub bangalore=950 lex_hub=1300 sanfran_hub=2000
Change oplog ID in row "sanfran_hub", column "bangalore" to
950 [no] yes
Change oplog ID in row "sanfran_hub", column "lex_hub" to
1300 [no] yes
Change oplog ID in row "sanfran_hub", column "sanfran_hub" to
2000 [no] yes
Epoch row successfully set.
```

- 3 At the sending site, use this output in a **chepoch** command:

```
multiutil chepoch sanfran_hub bangalore=950 lex_hub=1300
sanfran_hub=2000
Change oplog ID in row "sanfran_hub", column "bangalore" to
950 [no] yes
Change oplog ID in row "sanfran_hub", column "lex_hub" to
```

```
1300 [no] yes
Change oplog ID in row "sanfran_hub", column "sanfran_hub" to
2000 [no] yes
Epoch row successfully set.
```

## Restoring replicas

---

Occasionally, a replica is lost. This can occur because of a hardware failure (for example, disk crash), a software failure (for example, OS-level file-system corruption), or a human error. If an unreplicated database is lost, you can restore a recent copy from backup and resume development work. The changes made between the time of the backup and the time of the failure are not recoverable.

Similarly, if you lose a replica, you can restore a recent copy from backup. But matters are more complicated:

- Some of the work done between the time of the backup and the time of the failure may be recoverable. If some of the operations were sent to other replicas in update packets, these operations must be retrieved and imported.
- The restored copy of the replica is out of date. You must make this replica consistent with the other replicas in the family before development can proceed at your site. Failure to reestablish consistency can lead to irreparable damage.

Because this procedure involves substantial effort, it is intended for situations where serious damage has occurred. (For example, the disk containing a replica is unusable.)

### Restoring a replica from backup

To restore a replica from backup:

- 1 Use your vendor database tools to restore a copy of the replicated database from backup.
- 2 Use the **restorereplica** command to start the restoration procedure. You must have Super User privileges.

```
multitutil restorereplica [-cl·an clan-name] [-site
site-name] [-fam·ily family-name]
```

```
-u·ser username -p·assword password [-force] [-completed]
[-replace][replica ...]
```

This places a special lock on the replica. Between this point and the completion of Step 6, the **syncreplica -import** command adjusts the ClearQuest lock temporarily to permit application of the update, then restores the full lock. During this time, only **syncreplica -import** can modify the replica.

- 3 Verify that all update packets have been processed at their destination replicas.
- 4 At the restored replica, generate update packets for all other replicas, and send the packets to the sibling replicas.

You can send the packets using your standard synchronization method. To recover the replica more quickly, create the packets with **syncreplica -export -fship**.

Because your replica is in the special restoration state, each outgoing update packet includes a special request for a return acknowledgment. It also includes your replica's old epoch numbers, which are now its current epoch numbers, by virtue of the restoration backup in Step 1. Each destination replica uses these numbers to roll back its row for your replica.

- 5 Wait for each other replica in the family to send an update packet to the restored replica. As in Step 4, you can accelerate the creation and delivery of the update packets.

Collectively, these update packets include all the operations that occurred between the time of the backup and the last update that your replica sent out before its storage was lost—even operations that originated at your replica. (The packets also include more recent operations that originated at other replicas.) In addition, each incoming packet includes the requested return acknowledgment from the sending host.

- 6 Process the incoming update packets with **syncreplica -import**. When your replica has received return acknowledgments from all other replicas in its family, **syncreplica -import** reports that restoration of the replica is complete:

```
Replica has completed restoration: ...
```

Development work in the replica can now resume.

## Cleaning up from accidental deletion of a replica

---

This situation is a more catastrophic variation of the problem described in *Restoring a replica from backup* on page 94: a replica is lost, and there is no backup to be restored. The procedure for handling this situation is similar to that in *Removing a replica* on page 80.

You can perform this procedure at any replica within the same family as the deleted replica.

- 1 Transfer mastership of all the objects that were mastered by the deleted replica. For example, if replica **tokyo** is deleted, enter this command at replica **sanfrancisco**:

```
multiutil chmaster -family PRODA -user toykoadmin -p secret
sanfrancisco -all -force tokyo -long sanfran_hub
```

**Warning:** Incorrect use of **-all -force** can lead to irreparable inconsistencies among the replicas in a family.

- 2 From the **sanfrancisco** site, remove the deleted replica.

```
multiutil rmreplica -family PRODA -user sfadmin -p secret
tokyo
```

- 3 Send an update packet to all other replicas in the family, to inform them of the mastership changes and the replica deletion.

```
multiutil syncreplica -export -clan telecomm -site
sanfrancisco -family PRODA -user sfadmin -p secret -workdir
c:\temp\shipping -fship -sclass cq_default bangalore paris
boston
```

## Command summary

---

### Creating replicas

- activate
- mkreplica

### Synchronizing replicas

- syncreplica
- lsepoch
- chepoch

### Changing mastership

- chmaster
- describe

### Managing update packets

- recoverpacket
- mkorder
- shipping\_server
- lspacket

### Maintaining replicas

- chreplica
- lsreplica
- rename
- restorereplica

- **rmreplica**
- **scruboplog**
- **dumpoplog**

## activate

---

Prepares a database set to be used by ClearQuest MultiSite.

**Synopsis**     **activate** [ **-dbset** *dbset-name*] **-u-ser** *user-name* [**-p-assword**] *password*  
                   **-cl-an** *clan-name* **-site** *site-name* **-host** *host-name*

**Description**     The **activate** command prepares a database set to be replicated. When you activate a database set, you give it a clan name and site name.

When a database is activated, its logical name is changed to CQMS.<clan-name>.<site-name>.

After you have activated a database set, you can replicated one or more user databases within that database set. To replicate a database that has been activated, use the **mkreplica** command.

**Note:** You only need to activate a database set once, before creating the first replica.

**Options and Arguments**     The following sections describe the options and arguments for use with **activate**.

**-dbset** *dbset-name*

The name of the database set you want to activate. You can omit this argument if your ClearQuest installation has only one database set (schema repository).

**-u-ser** *user-name*

The name of a user with ClearQuest Super User privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. You must have Super User privileges to administer ClearQuest MultiSite. You can omit this argument if the password is null.

**-cl-an** *clan-name*

Name by which the database set will be identified in subsequent multiutil commands. When you activate a database set, you also provide a clan name. The clan name must be unique and can be up to 50 characters long. The clan name cannot be changed after a database set has been activated.

## activate

### **-site** *site-name*

Name by which the replica will be identified in multiutil commands. The site name must be an identifier and can be up to 50 characters long.

### **-host** *host-name*

The host name refers to the <machine name> where the storage bays for the respective replica are located and must be consistent with the name of the host specified when you configured the MultiSite Control Panel. The Rational Shipping Server must also be installed on the host machine.

**Examples**    The following example activates the default database set, names the clan "telecomm" and names the site "boston". This replica will use server1 as its host.

```
activate -user bostonadmin -password secret -clan telecomm
-site boston -host server1
```

**See also**    **mkreplica**



## cheepoch

---

Changes epoch number estimates.

|                       |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Synopsis              | <b>cheepoch</b> [- <b>cl-an</b> <i>clan-name</i> ] [- <b>site</b> <i>site-name</i> ] <b>-fam-ily</b> <i>family-name</i><br><b>-u-ser</b> <i>username</i> [- <b>p-assword</b> ] <i>password</i><br><b>[ -f-orce ]</b> <i>replica</i> [ <i>replica=value ...</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Description           | <p>This command changes, in one replica, one or more of the <i>epoch numbers</i> that represent the replica's estimate of the epoch levels of other replicas. You cannot change a replica's own epoch numbers because these epoch numbers record the actual state of the replica.</p> <p>You use this command if you need to resend a packet that did not get applied.</p> <p>For more information about epoch numbers, see <i>Database operations and the oplog</i> on page 8.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Options and Arguments | <p><b>-cl-an</b> <i>clan-name</i></p> <p>The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.</p> <p><i>Default:</i> The first clan replicated at this site. If there is more than one clan at this site, this argument is required.</p> <p><b>-site</b> <i>site-name</i></p> <p>The site name of the replica to be modified.</p> <p><i>Default:</i> Current site. If there is more than replica on this machine, this argument is required.</p> <p><b>-fam-ily</b> <i>family-name</i></p> <ul style="list-style-type: none"> <li>• <b>user database family:</b> Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.</li> <li>• <b>replicated schema repository family:</b> To specify the family of the replicated schema repository use the family name, MASTR.</li> </ul> <p><i>Default:</i> None.</p> |

**-u-ser** *user*

The name of a user with Super User privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. You must have Super User privileges to administer ClearQuest MultiSite.

**-f-orce**

Suppresses confirmation steps.

**replica**

Specifies the replica whose epoch row is to be changed; that is, changes the current replica's estimate of the state of *replica*. Use the site name to specify the replica.

*[replica=value]*

One or more arguments,

where

*replica*      Column of the epoch number matrix. This argument, along with the preceding *replicar* argument, specifies a particular location in the matrix.

*value*        New epoch number to be entered at the specified matrix location.

**Default:** **chepoch** reads a set of *replica=value* pairs, one per line, from standard input. You can copy and paste **lsepoch** output, or type the data in the format described below. Extra white space is allowed. To terminate input, type a period character (.) and a carriage return (<CR>) at the beginning of a line.

- Examples**
- Changes two columns of epoch estimates in the current replica's (*boston*) row for the *bangalore* replica.

```
multiutil chepoch -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret bangalore paris=500
sanfrancisco=750
Change oplog ID in "bangalore", "paris" to 500 [no] yes
Change oplog ID in "bangalore", "sanfrancisco" to 500 [no]
yes
```

- Make the same change as in the proceeding example, but skip the confirmation step.

```
multiuitl chepoch -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -force bangalore paris=500
sanfrancisco=750
Epoch row successfully set.
```

**See Also**

- **lsepoch**
- **recoverpacket**
- **restorereplica**

## chmaster

---

Transfers mastership of a ClearQuest object from one site to another.

**Synopsis**    **chmaster** [-**cl-an** *clan-name*] [-**site** *site-name*] **-fam-ily** *family-name*  
                   **-u-ser** *username* [-**p-assword**] *password new-master-replica*  
                   {  
                   { *entity-selector...* | **-all** [ **-force** *old-master-replica* ] [ **-l-ong** ] }  
                   | **-workingmaster** }

**Description**    This command transfers the mastership of one or more objects from one replica to another. Only the current replica is affected immediately; other replicas are notified of the mastership transfers through the normal exchange of update packets.

Only the replica that masters an object can transfer that object, unless you use the **-force** argument.

Mastership restricts some of the operations you can perform on an object. For more information about mastership, see *Managing mastership* on page 65.

**Options and Arguments**

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-family** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR. Users and groups are mastered by the replicated schema repository. If you want to change the mastership of a user or group, specify the MASTR family.

You should also specify the MASTR family when you are changing the mastership of the working schema repository with the **-workingmaster** option.

*Default:* None.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

*new-master-replica*

The name of the replica to which you are transferring mastership.

*Default:* None.

## Specifying objects of which to change mastership

You can specify a particular object (*entity-selector*), that all objects for a particular family be transferred (**-all**) or that only the working master site be changed.

### entity-selector

This is used to indicate the object type of which you want to change mastership. The following ClearQuest objects use mastership.

| Object whose mastership is changing | Syntax                                     | Family used                         |
|-------------------------------------|--------------------------------------------|-------------------------------------|
| Record                              | entity:<record_id>                         | user database family                |
| Users and groups                    | user:<login_name> or<br>group:<group_name> | schema repository family<br>(MASTR) |
| Workspace item                      | workspace:<query_name>                     | user database family                |

*Default:* None.

**-a-ll** [ **-f-orce** *old-master-replica* ] [ **-l-ong** ]

**Warning:** Incorrect use of the **-force** form of the command can lead to irreparable divergence among the replicas in a database family.

Transfers to *new-master-replica* mastership of all objects that are located in and mastered by the current replica family you specify. If errors occur, the command continues. After finishing, it reports that not all mastership changes succeeded.

With **-long**, **chmaster** lists the objects whose mastership is changing.

With **-force**, **chmaster** transfers mastership of all objects in the replica family specified with *old-replica-master*. Use this form of **chmaster** only when replica *old-replica-master* is no longer available (for example, was deleted accidentally).

**Note:** You cannot change the working master site to a different site with the **-all** option. If you want to change mastership of schemas and other administrative tasks, use the **-workingmaster** option.

**-workingmaster**

Transfers mastership of the working schema repository to the site you specify. Only the working schema repository can change the mastership of the itself.

You need to have specified the MASTR family with the **-family** option to use this option.

**Examples**

- At the boston site, transfer mastership of user group "development" to the sanfrancisco site:

```
multiutil chmaster -clan telecomm -site boston -family MASTR
-user bostonadmin -password secret sanfrancisco
group:development
```

- At replica boston, transfer mastership of user "jsmith" to the sanfrancisco replica:

```
multiutil chmaster -clan telecomm -site boston -family MASTR
-user bostonadmin -password secret sanfrancisco user:jsmith
```

- At replica paris, transfer mastership of a report format called "Project\_report" to the bangalore replica:

```
multiutil chmaster -clan telecomm -site paris -family PRODA
-user parisadmin -password secret bangalore
workspace:project_report
```

- At replica paris which is the current working schema repository site, transfer mastership of all schemas and working schema repository tasks to boston:

```
multiutil chmaster -clan telecomm -site paris -family MASTR
-user parisadmin -password secret boston -working master
The working master has been changed from paris to boston.
```

- At replica paris, transfer mastership of all items mastered by paris to replica boston:

```
multiutil chmaster -clan telecomm -site paris -family PRODA
-user parisadmin -password secret boston -all
```

If you want to also transfer any users or groups, you'll need to run the command again and specify the MASTR family:

```
multiutil chmaster -clan telecomm -site paris -family MASTR
-user parisadmin -password secret boston -all
```

**See Also**

- **describe**
- **syncreplica**

## chreplica

---

Changes the host properties of a replica.

**Synopsis**    **chrep-lica** [-**cl-an** *clan-name*] [**-site** *site-name*]  
                   **-u-ser** *username* [**-p-assword**] *password* **-host** *hostname* *replica-selector*

**Description**    Use this command when you want to change the host information of a replica. If the host has been moved to a different location on the network, you'll need to change the host information that is stored with that replica.

**Options and Arguments**

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required..

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-host** *hostname*

Enter the name of the new shipping server host.

*replica-selector*

Specifies the replica to be changed. Use the site name to specify the replica you want to change.

*Default:* None.



**Examples** Associate replica sanfrancisco with the host pacific1.  
`multitutil -clan telecomm -site sanfrancisco -family PRODA  
-user sfadmin -password secret -host pacific1 sanfrancisco`

**See Also**

- **chmaster**
- **syncreplica**

## describe

---

Describes the mastership properties of a ClearQuest object

**Synopsis**    **describe** [-**cl-an** *clan-name*] [-**site** *site-name*] **-fam-ily** *family-name*  
                   **-u-ser** *username* [-**p-assword**] *password entity-selector...*

**Description**    This command lists the mastering replica of one or more ClearQuest objects in a replica.

Mastership restricts some of the operations you can perform on an object. For more information about mastership, see *Managing mastership* on page 65.

**Options and Arguments**

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be described.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR. Use the MASTR family if you want to find out where respective users and groups are mastered.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**entity-selector**

This is used to indicate the object type of which you want to view the mastership properties. The following ClearQuest objects use mastership.

| Object whose mastership is changing | Syntax                                     | Family used                         |
|-------------------------------------|--------------------------------------------|-------------------------------------|
| Record                              | entity:<record_id>                         | user database family                |
| Users and groups                    | user:<login_name> or<br>group:<group_name> | schema repository family<br>(MASTR) |
| Workspace item                      | workspace:<query_name>                     | user database family                |

*Default:* None.

**Examples**

- At replica boston, find out which site is the mastering replica for the user group "development" :  

```
multiutil describe -clan telecomm -site boston -family MASTR
-user bostonadmin -password secret group:development
```
- At replica boston, find out which site is the mastering replica for the user "jsmith":  

```
multiutil describe -clan telecomm -site boston -family MASTR
-user bostonadmin -password secret user:jsmith
```
- At replica paris, find out which site is the mastering replica of a report format called "Project\_report"  

```
multiutil chmaster -clan telecomm -site paris -family PRODA
-user parisadmin -password secret workspace:project_report
```

**Note:** The easiest way to find out which site masters a records is to look at the value for the "ratl\_mastership" field for the record in question.

**See Also**

- **chmaster**
- **syncreplica**

## dumpoplog

---

Lists the contents of the specified replicas oplog.

- Synopsis** **dumpoplog** [-**cl-an** *clan-name*] [-**site** *site-name*] **-fam-ily** *family-name* **-u-ser** *username* [-**p-assword**] *password* [ **-l-ong** | **-s-hort**] [-**at** *replica...*] [*oplog-ent-order-num...* | [-**from** *oplog-ent-order-num*] [-**to** *oplog-ent-order-num*] [-**since** *date-time*] [-**reverse**]
- Description** Use **dumpoplog** to find out more about a replica's oplog. The oplog is used to track all database transactions including record changes and schema modifications. Each oplog entry is numbered. You can list (dump) an entire oplog or view just a range of entries. You can specify oplog entries by integer or by date-time or view only entries that originated from a specific replica site.
- Options** **-cl-an** *clan-name*
- The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.
- Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.
- site** *site-name*
- The site name of the respective replica.
- Default:* Current site. If there is more than replica on this machine, this argument is required.
- fam-ily** *family-name*
- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
  - **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR.
- Default:* None.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-l-ong** | **-s-hort**

Use -long to view all columns...

<more info here TONIA>

Use -short to view...<more info here>

*Default:* If no format is specified, -short is used.

**-at** *replica...*

Lists the oplog entries that originated from the replica(s) you specify. Use the site name to specify the replica.

*oplog-ent-order-num...*

Allows you to enumerate one or more oplog numbers. Oplog numbers are entered as integers.

**SPECIFYING THE OPLOG ENTRIES TO LIST:** You can list oplog entries according to number or use a date-time format with the **-since** argument.

**-from** *oplog-ent-order-num...*

Used in conjunction with the **-to-oplog-ent-order-num...** to indicate a range of oplog numbers. Oplog numbers are entered as integers.

**SPECIFYING THE OPLOG ENTRIES TO LIST:** You can list oplog entries according to number or use a date-time format with the **-since** argument.

**-to** *oplog-ent-order-num...*

Used in conjunction with the **-from oplog-ent-order-num...** to indicate a range of oplog numbers. Oplog numbers are entered as integers.

**SPECIFYING THE OPLOG ENTRIES TO LIST:** You can list oplog entries according to number or use a date-time format with the **-since** argument..

**-since** *date-time*

All oplog entries after the date-time you specify will be listed.

The *date-time* argument must be in the following format:

## dumpoplog

m[m]/d[d]/yyyy h[h]:m[m]:s[s]

For example:

10/7/2002 17:04:32

### **-reverse**

Allows you reverse the order of the list of the oplog entries.

**Examples** Dumps oplog of the  
TONIA example to be added.

**See Also**

- **lsepoch**
- **scruboplog**

## Isepoch

---

Lists the current epoch estimates for specified replicas.

|                              |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Synopsis</b>              | <b>Isepoch</b> [- <b>cl-an</b> <i>clan-name</i> ] [- <b>site</b> <i>site-name</i> ] - <b>family</b> <i>family-name</i> - <b>u-ser</b> <i>username</i> [- <b>p-assword</b> ] <i>password</i> [ <i>replica ...</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>Description</b>           | If a clan, site or family is not specified, the epoch estimates for all local instances of clans, sites, and families, respectively, are listed. The login must be valid for all instances. The replica's own epoch row in its matrix represents its actual level. The other rows represent the replica's best estimate of the other replicas' levels.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>Options and Arguments</b> | <p><b>-cl-an</b> <i>clan-name</i></p> <p>The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.</p> <p><i>Default:</i> The first clan replicated at this site. If there is more than one clan at this site, this argument is required.</p> <p><b>-site</b> <i>site-name</i></p> <p>The site name of the replica to be modified.</p> <p><i>Default:</i> Current site. If there is more than replica on this machine, this argument is required.</p> <p><b>-family</b> <i>family-name</i></p> <ul style="list-style-type: none"> <li>• <b>user database family:</b> Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.</li> <li>• <b>replicated schema repository family:</b> To specify the family of the replicated schema repository use the family name, MASTR.</li> </ul> <p><i>Default:</i> None.</p> <p><b>-u-ser</b> <i>user</i></p> <p>The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.</p> |

## lsepoch

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

*replica ...*

Specifies the replica for which you want to list epoch information. Use the site name(s) to specify the replica(s) you want to list.

*Default:* Lists the epoch estimate for each replica in the family.

### Examples

- Lists epoch estimate in the current replica's (*boston*) for the *bangalore* and *paris* replicas.

```
multiutil lsepoch -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret bangalore paris
```



## Ispacket

---

Describes contents of a packet.

**Synopsis** **Ispacket** [ **-l-ong** | **-s-hort** ] [ *pname ...* ]

**Description** This command lists a summary of the contents of one or more disk files that contain replica-creation or update packets. By default, the **Ispacket** output includes this information:

- Pathname of each packet
- Type of each packet (Replica Creation or Update)
- Date of generation
- The originating replica
- Clan and database family to which the packet applies
- Replicas for which the packet is intended.

An asterisk after a replica name indicates that the packet can be imported immediately because it does not depend on any other packet.

For example, if there are two packets waiting to be imported at a replica, the first packet has an asterisk and the second doesn't (because the second packet depends on the first).

- Packet sequence number (for a disk file storing one part of a logical packet that has been split into multiple physical packets)

**Options and Arguments** **-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-family** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR.

*Default:* None.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-l-ong |**

In addition to the default information, lists the following information:

- Name of the replica where the packet was created
- Oplug IDs (epoch numbers) that indicate the contents of the packet
- Packet revision number which is used by Rational Software

**-s-hort**

Lists only the pathname of a packet.

*pname ...*

One or more pathnames of files and/or directories (for example, pathnames of MultiSite storage bay directories).

Each file you specify is listed if it contains a physical packet (which may be one of several that make up a logical packet). For each directory you specify, **lspacket** lists packets stored in that directory.

*Default:* Lists all packets in all MultiSite storage bays on the current host.

- Examples**
- List a single update packet.  
`multiutil lspacket -long c:\temp\incoming\osakaupdate2`
  - List all packets in all of the local host's storage bays.  
`multiutil lspacket -short`
  - List all packets in a specific storage bay.  
`multiutil lspacket -long c:\temp\incoming\`
- See Also**
- **mkreplica**
  - **MultiSite Control Panel**
  - **syncreplica**

## Isreplica

---

Lists database replicas.

**Synopsis** **Isrep-lica** [**-cl-an** *clan-name*] [**-site** *site-name*] **-fam-ily** *family-name*  
**-u-ser** *username* [**-p-assword**] *password* [**-l-ong** | **-s-hort** | **-fmt** *format* ]  
 [ **-sib-lings** | [ **-sib-lings** ] **-infa-mily** *family* | *replica ...* ]

**Description** **Isreplica** lists information on all replicas recorded in the replica database of the current replica (except for deleted replicas). You can either list all replicas in a known clan or known family within a clan. Other replicas may exist, but the packets containing their creation information have not yet been imported at the current replica.

**Options and Arguments**

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the respective replica.

*Default:* Current site. If there is more than one replica on this machine, this argument is required.

**-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR.

You can use the **-siblings** argument or the **-siblings -infamily** arguments to list the replicas in a specific family within the clan, as known to you replicated schema repository.

*Default:* None.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-l-ong**

Includes each replica's creation information and host. If the current replica is in the process of restoration, this option annotates the listings of other replicas from which restoration updates are required. (See the **restorereplica** reference page.)

**LISTING FORMAT.** *Default:* Includes creation event information for each replica.

**-s-hort**

Lists only replica names.

**-fmt** *format-string*

With **-fmt** you can customize simple reports on replica information retrieved by **lsreplica**. Note that **-fmt** is a mutually exclusive alternative to the **-short** and **-long** options.

**-fmt** uses conversion specifications which identify particular data items to display and specifies its display format. The conversion specification format closely resembles that of the C-language function **printf()**:

- Percent sign (%)
- A key letter (lowercase), which indicates the kind of data to display

Unlike **printf()** specifiers, conversion specifications are not replaced by arguments supplied elsewhere on the command line; they are replaced automatically by **multiutil** with field values extracted from replica.

*format-string* is a character string, composed of alphanumeric characters, conversion specifications, and escape sequences. It must be enclosed in double quotes ( " ).

*Conversion specifications:*

|    |           |
|----|-----------|
| %h | Host name |
|----|-----------|

|    |                                 |
|----|---------------------------------|
| %n | Name of replica                 |
| %c | Clan name                       |
| %f | Family name                     |
| %d | Description of replica, if any. |
| %% | % character                     |

*Escape sequences:*

|      |                                   |
|------|-----------------------------------|
| \n   | <NL>                              |
| \t   | <TAB>                             |
| \'   | Single quote                      |
| \\   | Literal (uninterpreted) backslash |
| \nnn | Character specified by octal code |

**Example**

- Mimic the output from **lsreplica -long**. Note the use of single quotes to enclose the format string, which includes literal double quotes.
- TONIA's

**-sib-lings**

Lists the family members of the current replica, but does not list the current replica itself. This option is useful when you are writing scripts that process only sibling replicas.

**-infa-mily** *family*

Lists the all replicas in the same family of the specified replica. Use the site name to specify the replica. This option can only be used if you specified MASTR with the **-family** option.

*Default:* None.

*replica ...*

Specifies the replica for which you want to list information. Use the site name(s) to specify the replica(s) you want to list. You can only list replicas who are members of the same family.

- Examples**
- List the names of all the replicas participating in the family (PRODA) of the current replica (bangalore):  
`lsreplica -clan telecomm -site bangalore -family PRODA -user badmin -password secret -short`  
<more info here>
  - Display a long listing of replicas participating in the family (PRODA) of the current replica (bangalore):  
`lsreplica -clan telecomm -site bangalore -family PRODA -user badmin -password secret -long`  
<more info here>
  - List the names of all the replicas participating in the replicated schema repository family (MASTR) that participate in the PRODB family:  
`lsreplica -clan telecomm -site bangalore -family MASTR -user badmin -password secret -short -infamily PRODB`  
<more info here>
- See Also**
- **mkreplica**

## mkorder

---

Creates a shipping order for use by the store-and-forward facility.

**Synopsis**    **mkorder** **-dat-a** *packet-pname* [ **-sci-ass** *storage-class-name* ]  
                  [ **-pex-pire** *date-time* ] [ **-not-ify** *e-mail-address* ]  
                  [ **-shi-p** **-cop-y** | **-fsh-ip** [ **-cop-y** ] | **-out** *order-pname* ]  
                  *destination ...*

**Description**    This command creates a shipping order file for an existing packet or any other file. The shipping order is used by the **shipping\_server** to send the packet to one or more destinations.

**mkorder** submits to the store-and-forward facility a packet that was created with **mkreplica -out** or **syncreplica -out**. You can also use **mkorder** to resubmit store-and-forward packets whose shipping orders have expired, and to transmit arbitrary files among sites.

A shipping order must be located in the same directory as its associated packet or file.

**Note:** The store-and-forward facility deletes a packet after delivering it successfully (except when the destination is the local host). If you use this command to process a file that must be preserved at your site even after delivery to another site, you must specify the **-copy** option.

### Default Handling for Packet-delivery Failures.

*Default:* If a packet cannot be delivered, it is sent through the store-and-forward facility back to the administrator at the site of the originating replica. A mail message is sent to the store-and-forward administrator. This occurs after repeated attempts to deliver the packet have failed, and the allotted time has expired; it can also occur when the destination host is unknown or a data file does not exist. The store-and-forward configuration settings specify the expiration period and the e-mail address of the administrator.



## Options and Arguments

**-dat-a** *packet-pname*

The pathname of the file containing the packet.

*Default:* None.

**Note:** If *packet-pname* contains a colon character ( : ), **mkorder** changes the colon to a period character ( . ) during processing. This allows packets to be delivered to Windows machines, which do not allow colons within file names.

**-scl-ass** *class-name*

Specifies the storage class of the packet and shipping order. If you also use **-ship** or **-fship**, **mkorder** looks up the storage class in the store-and-forward configuration settings in the **MultiSite Control Panel** to determine the location of the storage bay to use.

*Default:* Creates a shipping order in the directory where the *packet-pname* file is located.

If you omit this option but use **-ship** or **-fship**, **mkorder** places the shipping order in the storage bay location specified for the **-cq\_default** class in the **MultiSite Control Panel**.

**-shi-p****-cop-y****-fsh-ip** [ **-cop-y** ]

Creates a shipping order for the *packet-pname* file. Using **-fship** (force ship) invokes **shipping\_server** to send the packet. Using **-ship** places the shipping order in a storage bay.

**-copy** is required with **-ship**, and optional with **-fship**:

- With **-copy**, **mkorder** copies the *packet-pname* file to one of the store-and-forward facility's storage bays, and places the shipping order in the bay. The copy is deleted after it is delivered successfully to all the destinations specified in the shipping order.
- Without **-copy**, **mkorder** does not copy *packet-pname*; **mkorder** places the shipping order in the directory where the file is located. *packet-pname* is deleted after it is delivered successfully to all the destinations specified in the shipping order.

**-out** *order-pname*

Places the shipping order in the specified file instead of in a storage bay. An error occurs if the file already exists.

**-pex-pire** *date-time*

Specifies the time at which the store-and-forward facility stops attempting to deliver the packet and generates a failure mail message instead.

This option overrides the storage class's Packet Expiration specification in the MultiSite Control Panel. See the **MultiSite Control Panel** reference page for a discussion of this specification, and of delivery retries in general.

The *date-time* argument can have any of the following formats:

*date.time* | *date* | *time* | **now**

where:

|                    |    |                                                                                                                                                    |
|--------------------|----|----------------------------------------------------------------------------------------------------------------------------------------------------|
| <i>date</i>        | := | <i>day-of-week</i>   <i>long-date</i>                                                                                                              |
| <i>time</i>        | := | <i>h</i> [ <i>h</i> ]: <i>m</i> [ <i>m</i> ][: <i>s</i> [ <i>s</i> ]] [ <b>UTC</b> [ [ +   - ] <i>h</i> [ <i>h</i> ][: <i>m</i> [ <i>m</i> ] ] ] ] |
| <i>day-of-week</i> | := | <b>today</b>   <b>yesterday</b>   <b>Sunday</b>   ...   <b>Saturday</b>   <b>Sun</b><br>  ...   <b>Sat</b>                                         |
| <i>long-date</i>   | := | <i>d</i> [ <i>d</i> ]- <i>month</i> [-[ <i>yy</i> ] <i>yy</i> ]                                                                                    |
| <i>month</i>       | := | <b>January</b>   ...   <b>December</b>   <b>Jan</b>   ...   <b>Dec</b>                                                                             |

Specify the *time* in 24-hour format, relative to the local time zone. If you omit the time, the default value is **00:00:00**. If you omit the *date*, the default value is **today**. If you omit the century, year, or a specific date, the most recent one is used. Specify **UTC** if you want the time to be resolved to the same moment in time regardless of time zone. Use the plus (+) or minus (-) operator to specify a positive or negative offset to the UTC time. If you specify **UTC** without hour or minute offsets, the default setting is Greenwich Mean Time (GMT). (Dates before January 1, 1970 Universal Coordinated Time (UTC) are invalid.)

Examples:

```
22-November-1998
sunday
yesterday.16:00
8-jun
13:00
today
9-Aug.10:00UTC
```

**-not-ify** *e-mail-address*

The delivery-failure message is sent to the specified e-mail address.

If a failure occurs on a Windows host that does not have e-mail notification enabled, a message appears in the Windows Event Viewer. The message includes the *e-mail-address* value specified with this option and a note requesting that this user be informed of the status of the operation. For information on enabling e-mail notification, see the **MultiSite Control Panel** reference page.

*destination ...*

One or more host names (which must be usable by hosts in different domains) or IP addresses. When sending a MultiSite packet, you must specify the host where the replica actually resides or is to be created.

*Default:* None.

**Examples**

- Create a shipping order for file **p1**, which is located in the `cq_default` storage bay. Store the shipping order in the same storage bay as **p1**, and specify that the file is to be sent to host **goldengate**. The lines are broken for readability. You must enter the command on a single physical line.  

```
mkorder -data -sclass cq_default "c:\temp\outgoing\p1" -out
"c:\temp\outgoing\p1_order" goldengate
Shipping order "c:\temp\outgoing\p1_order" generated
```
- Create a shipping order in the default storage bay for a specified file that is to be delivered to host **goldengate**. Specify that **admin** must be notified if the file is not delivered successfully.  

```
mkorder -data -sclass cq_default "c:\temp\outgoing\p1" -out
"c:\temp\outgoing\p1_order" goldengate
```
- Shipping order "c:\temp\outgoing\p1\_order" generated
- Create a shipping order for the same file, but place it in the storage bay for a particular storage class. Attempt immediate delivery (**-fship**), and allow delivery attempts to continue until the beginning of May 18.

**See Also**

- **mkreplica**,
- **MultiSite Control Panel**,
- **syncreplica**

## mkreplica

---

Creates a database replica.

### Synopsis

- Duplicate an existing database, generating a replica-creation packet:

```
mkrep-lica -exp-ort [-cl-an clan-name] [-site site-name] -fam-ily family-name
-u-ser username [-p-assword password] [-max-size size]
{
 {-sh-ip|-fsh-ip} -wor-kdir workpath [-scl-ass storage-class]
 [-pex-pire date-time] [-not-ify e-mail-addr] | -out packet-file-pname
} hostname:site-name ...
```

- Use a replica-creation packet to create a new replica, in addition to a new replicated schema repository:

```
mkrep-lica -imp-ort
{ [-site site-name] -repo-sitory db-info [-vendor <vendor-type>]
 <db-params>
}
{ [-data-base db-info [-vendor <vendor-type>] <db-params>]
-c-omments comments] { packet-file path | packet-dir-path }...
```

- Use a replica-creation packet to create a new replica within the same clan as the existing replicated schema repository at the current site:

```
mkrep-lica -imp-ort { [-cl-an clan-name] [-site site-name]
-u-ser username [-p-assword] password
}
{ -data-base db-info [-vendor <vendor-type>] <db-params>]
-c-omments comments] { packet-file path | packet-dir-path }...
```

where *<db-params>* :=

```
if <vendor> == DB2
 <db-info> := Database Alias
 <db-params> := -dbologin dbo-name [dbo-pwd]
```

```
if <vendor> == ORACLE
 <db-info> := SQL*Net Alias
 <db-params> := -dbologin dbo-name [dbo-pwd] [-connectopts
connect-options]
```

```
if <vendor> == SQL_SERVER
 <db-info> := Physical Database Name
 <db-params> := -server server-name -dbologin dbo-name [dbo-pwd]
-rwlogin rw-login [rw-pwd] [-rologin ro-login [ro-pwd]]
```

**Description** The creation of a new database replica is a two-phase process. Both phases require you to enter a **mkreplica** command.

**Note:** Before replicating a the first database of a clan, you must first **activate** the database set to which it belongs.

This command may take a long time. The database is locked while an export is in progress. For more information see *Creating a replica* on page 47.

**1** The **mkreplica –export** command duplicates the contents of the specified user database and its associated schema repository. This generates a single logical *replica-creation packet* for transmission to one or more other sites. As described in *Replica-creation packets* on page 130, it may be divided into multiple physical packets. (If you use **–fship** or **–ship**, **mkreplica** also generates a shipping order file for each physical packet file.)

**Note:** Creating multiple replicas in one **mkreplica –export** command is more efficient than using multiple **mkreplica –export** commands.

**2** At each new site, a **mkreplica –import** command uses the replica-creation packet to create a new replica. The new replica consists of two replicated databases, a schema repository and a user database. This command varies if you are adding a user database replica which belongs to a family within the same clan of an existing replicated schema repository at your site.

### Creating empty vendor databases

At each new site, the site administrator needs to have created empty vendor databases to receive the replica data. If this is the first replica at a site, you need at least two empty vendor databases, one for the replicated schema repository and one for the user database replica.

**Note:** If you are adding a user database replica family that is in the same clan as an already existing replicated schema repository at your site, you don't need to create a vendor database for the the replicated schema repository. You can associate the new user database replica with the existing replicated schema repository at your site.

## Oplog information

When a database is first replicated, the database's *oplog* (operation log) is enabled. All ClearQuest operations to be replicated are recorded in the oplog. Logging of operations continues until all but one of the database's replicas are deleted. Note that creation of additional replicas is recorded in oplog entries. Existing replicas learn about a newly created replica through the standard synchronization mechanism. (See the **syncreplica** reference page.)

**Note:** Before entering a **mkreplica -export** command, make sure ClearQuest MultiSite licenses are installed at the original site. After you enable replication of the original database, developers cannot access the database without a ClearQuest MultiSite license (in addition to a ClearQuest license). A ClearQuest MultiSite license is also required to run **mkreplica -export**.

## Replica-creation packets

Each invocation of **mkreplica -export** creates a single logical replica-creation packet. (This is true even if you create several new replicas with one **mkreplica** command.) The packet carries one or more replica specifications, each of which indicates the host on which a new replica is to be created, along with the new replica's name.

The user database and schema repository are locked during the **-export** phase.

The **-maxsize** option divides the single logical packet into multiple physical packets (disk files or tapes) to conform with limitations of the transfer medium.

## Cleaning up used packets

Replica-creation packets are not deleted after import. The MultiSite administrator at the new replica site must delete replica-creation packets after importing them with **mkreplica -import**.

## Error Handling for Packet Delivery Failures

If a packet cannot be delivered, it is sent through the store-and-forward facility back to the administrator at the site of the originating replica. A mail message is sent to the store-and-forward administrator. This occurs after repeated attempts to deliver the packet have all failed, and the allotted time has expired; it can also occur when the destination host is unknown or a data file does not exist. The store-and-forward configuration settings specify the expiration period and the e-mail address of the administrator.

**Restrictions** *Locks:* This command fails if the database is locked (i.e., during the upgrade process) or while another ClearQuest MultiSite operation is being performed.

*Other restrictions:*

- You cannot replicate a database to a host running a different version of ClearQuest MultiSite.

**Options and Arguments**

Export Phase

The following sections describe the options and arguments for use with **mkreplica –export**. The following arguments specify the database to be replicated.

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**–site** *site-name*

The site name of the replica to be created. The site name is specified when the database is replicated.

**–fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** Not applicable. When you run **mkreplica** the associated schema repository of the user database family you specify is automatically included in the replica-creation packet.

*Default:* None.

**–u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**–p-assword** *password*

The password associated with the user account being used to run this command. You must have Super User privileges to administer ClearQuest MultiSite.

**–max-size** *size*

The maximum size for a physical packet, expressed as a number followed by a single letter; for example:

|      |               |
|------|---------------|
| 500k | 500 kilobytes |
| 20m  | 20 megabytes  |
| 1.5g | 1.5 gigabytes |

**Specifying the replica-creation packet size:** *Default:* None. When you do not specify **–maxsize**, the default packet size depends on the shipping method you use:

- Packets created with **–ship** or **–fship** are no larger than the maximum packet size specified in the the **MultiSite Control Panel** .
- Packets created with **–out** are no larger than 2 GB.

The **mkreplica** command fails if it creates a packet larger than the size supported by your system.

**–shi-p****–fsh-ip**

Stores the replica-creation packet in one or more files in a store-and-forward storage bay. A separate shipping order file accompanies each physical packet, indicating how and where it is to be delivered.

**–fship** (force ship) invokes **shipping\_server** to send the replica-creation packet.

**–ship** places the packet in a storage bay. To send the packet, invoke **shipping\_server**.

*Default disposition of replica-creation packet:* None. You must specify how the replica-creation packet created by **mkreplica –export** is to be stored and/or transmitted to other sites.

**Note:** The disk partition where the storage bay is located (on the sending host and the receiving host) must have available space equal to or greater than the size of the replica-creation packet.



**-workdir** *workpath*

A directory for use by **mkreplica** as a temporary workspace; it is deleted when **mkreplica** finishes. This directory must not already exist.

This directory is needed only when using **-ship** or **-fship**.

*Default Location for temporary Workspace:* None.

**-scl-ass** *storage-class*

Specifies the *storage class* of the packet and shipping order. **mkreplica** looks up the storage class in the store-and-forward configuration settings in the **MultiSite Control Panel** to determine the location of the storage bay to use.

If you omit this option, **mkreplica** places the packet in the storage bay location specified for the **cq\_default** class.

**-out** *packet-file-pname*

Places the first physical replica-creation packet in file *packet-file-pname*. Additional packets are placed in files named *packet-file-pname\_2*, *packet-file-pname\_3*, and so on.

The replica-creation packets are not delivered automatically; use an appropriate mechanism (for example, electronic mail, **ftp**, or postal service) to deliver them.

You can create a packet using **-out**, and subsequently deliver it using the store-and-forward facility. See the **mkorder** reference page for details.

**-pexpire** *date-time*

Specifies the time at which the store-and-forward facility stops trying to deliver the packet and generates a failure mail message instead.

This option overrides the storage class's Packet Expiration specification in the MultiSite Control Panel. See the **MultiSite Control Panel** reference page for a discussion of this specification and of delivery retries in general.

The *date-time* argument can have any of the following formats:

*date.time* | *date* | *time* | **now**

where:

|                    |    |                                                                                                                                               |
|--------------------|----|-----------------------------------------------------------------------------------------------------------------------------------------------|
| <i>date</i>        | := | <i>day-of-week</i>   <i>long-date</i>                                                                                                         |
| <i>time</i>        | := | <i>h</i> [ <i>h</i> ]: <i>m</i> [ <i>m</i> ]: <i>s</i> [ <i>s</i> ] [ <b>UTC</b> [ [ +   - ] <i>h</i> [ <i>h</i> ]: <i>m</i> [ <i>m</i> ] ] ] |
| <i>day-of-week</i> | := | <b>today</b>   <b>yesterday</b>   <b>Sunday</b>   ...   <b>Saturday</b>   <b>Sun</b><br>  ...   <b>Sat</b>                                    |
| <i>long-date</i>   | := | <i>d</i> [ <i>d</i> ]- <i>month</i> [ <i>-</i> [ <i>yy</i> ] <i>yy</i> ]                                                                      |
| <i>month</i>       | := | <b>January</b>   ...   <b>December</b>   <b>Jan</b>   ...   <b>Dec</b>                                                                        |

Specify the *time* in 24-hour format, relative to the local time zone. If you omit the time, the default value is **00:00:00**. If you omit the *date*, the default value is **today**. If you omit the century, year, or a specific date, the most recent one is used. Specify **UTC** if you want the time to be resolved to the same moment in time regardless of time zone. Use the plus (+) or minus (-) operator to specify a positive or negative offset to the UTC time. If you specify **UTC** without hour or minute offsets, the default setting is Greenwich Mean Time (GMT). (Dates before January 1, 1970 Universal Coordinated Time (UTC) are invalid.)

Examples:

```
22-November-1998
sunday
yesterday.16:00
8-jun
13:00
today
9-Aug.10:00UTC
```

**-not-ify** *e-mail-address*

The delivery-failure message is sent to the specified e-mail address.

If a failure occurs on a Windows host that does not have e-mail notification enabled, a message appears in the Windows Event Viewer. The message includes the *e-mail-address* value specified with this option and a note requesting that this user be informed of the status of the operation. For information on enabling e-mail notification, see the **MultiSite Control Panel** reference page.

**hostname:site-name...**

One or more arguments, each of which indicates one new replica to be created from this packet at another site.

*Default:* None.

*hostname* Names the machine where the new replica's storage directory will be created. *hostname* must be usable by hosts in different domains. It is used by the ClearQuest *store-and-forward* mechanism to determine how to route update packets to the replica. However, keep this information accurate even if your site does not use store-and-forward. (See the **chreplica** reference page.)

*hostname* can be either the IP address of the host or the computer name, for example, **minuteman**. You may have to append an IP domain name, for example, **minuteman.purpledoc.com**.

On Windows NT, the computer name is displayed in the **Network Settings** dialog box, which is accessible from the **Network** icon in the Control Panel. On Windows 2000, the computer name is displayed on the **Network Identification** tab in the **System Properties** dialog box, which is accessible from the **System** icon in the Control Panel.

*site-name* Name by which the replica will be identified in multiutil commands. The site name must be an identifier and can be up to 50 characters long. This name must be unique within the respective clan. In other words, there cannot be two sites with the same name participating in the same clan.

## Options and Arguments

### Import Phase

The following sections describe the options and arguments for use with **mkreplica -import**.

Before using **mkreplica -import**, the site administrator needs to have created empty vendor databases to receive the replica data.

### Importing both a replicated schema repository and a user database replica

If there is no existing replicated schema repository at the site to which you are importing the replica, you need to enter the repository name, site name and vendor database parameters.

#### **-site** *site-name*

The name of the site where the replica will be imported. The site name was given to the replica when it was exported. If you don't know the site name, contact your ClearQuest MultiSite administrator. This name must match the site name given to the replica when it was exported.

#### **-repo-sitory** *db-info*

Enter the database information respective to the vendor database you are using.

| Vendor Database | <dbinfo> value                |
|-----------------|-------------------------------|
| DB2             | <i>Database Alias</i>         |
| Oracle          | <i>SQL*Net Alias</i>          |
| SQL Server      | <i>Physical Database Name</i> |

#### **-vendor** *<vendor-type>*

Enter the database vendor you are using. Supported vendor types are: DB2, SQL\_SERVER, and ORACLE.

<db-params>

The required database parameters are the same parameters needed to connect to any ClearQuest database. Note these when you create the vendor database to which you import the replica. For more information about how to create empty vendor databases and the parameters you will need, see the *Installing Rational ClearQuest* documentation.

**Note:** When you import a replica, you'll need to specify the database parameters of both the vendor database that will hold the replicated schema repository and the vendor database that will hold the user database replica. You must have first created these databases before importing a replica packet.

here <db-params> :=

```
if <vendor> == DB2
 <db-info> := Database Alias
 <db-params> := -dbologin dbo-name [dbo-pwd]
```

```
if <vendor> == ORACLE
 <db-info> := SQL*Net Alias
 <db-params> := -dbologin dbo-name [dbo-pwd] [-connectopts
connect-options]
```

```
if <vendor> == SQL_SERVER
 <db-info> := Physical Database Name
 <db-params> := -server server-name -dbologin dbo-name [dbo-pwd]
-rlogin rw-login [rw-pwd] [-rologin ro-login [ro-pwd]]
```

**Note:** The read-only login (-rologin) for SQL\_SERVER is needed only when creating a replicated schema repository.

**-data-base** db-info

Enter the user database information respective to the vendor database you are using.

| Vendor Database | <dbinfo> value                |
|-----------------|-------------------------------|
| DB2             | <i>Database Alias</i>         |
| Oracle          | <i>SQL*Net Alias</i>          |
| SQL Server      | <i>Physical Database Name</i> |

**-vendor** <vendor-type>

Enter the database vendor you are using. Supported vendor types are: DB2, SQL\_SERVER, and ORACLE.

*Default:* Uses the same vendor as specified for the **-repository**.

## &lt;db-params&gt;

The required database parameters are the same parameters needed to connect to any ClearQuest database. Note these when you create the vendor database to which you import the replica. For more information about how to create empty vendor databases and the parameters you will need, see the *Installing Rational ClearQuest* documentation.

**Note:** When you import a replica, you'll need to specify the database parameters of both the vendor database that will hold the replicated schema repository and the vendor database that will hold the user database replica. You must have first created these databases before importing a replica packet.

here <db-params> :=

```
if <vendor> == DB2
 <db-info> := Database Alias
 <db-params> := -dbologin dbo-name [dbo-pwd]
```

```
if <vendor> == ORACLE
 <db-info> := SQL*Net Alias
 <db-params> := -dbologin dbo-name [dbo-pwd] [-connectopts
 connect-options]
```

```
if <vendor> == SQL_SERVER
 <db-info> := Physical Database Name
 <db-params> := -server server-name -dbologin dbo-name [dbo-pwd]
 -rwlogin rw-login [rw-pwd] [-rologin ro-login [ro-pwd]]
```

The read-only login (-rologin) for SQL\_SERVER is needed only when creating a replicated schema repository.

**--c-omments** comments

Enter any comments you want to include.

*packet-file-path* | *packet-dir-path* ...

Specifies a pathname of a replica-creation packet created by **mkreplica** **-export**. For a logical packet that spans multiple disk files, **mkreplica** scans the directory containing *packet-file-pname* for related physical packets.

If you also specify one or more *packet-dir-path* arguments, **mkreplica** searches for additional packets in these directories.

*Default:* None.

### **Adding an additional replica to an existing clan (replicated schema repository)**

If you are adding a replica family to an existing clan (replicated schema repository), you only need to create a vendor database for the user database replica. Instead of entering the database parameters for the a new replicated schema repository, you can use the clan name and indicate the following login information.

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be imported. The site name is specified when the database is replicated.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-data-base** *db-info*

Enter the user database information respective to the vendor database you are using.

| Vendor Database | <dbinfo> value                |
|-----------------|-------------------------------|
| DB2             | <i>Database Alias</i>         |
| Oracle          | <i>SQL*Net Alias</i>          |
| SQL Server      | <i>Physical Database Name</i> |

**-vendor** *<vendor-type>*

Enter the database vendor you are using. Supported vendor types are: DB2, SQL\_SERVER, and ORACLE.

*<db-params>*

The required database parameters are the same parameters needed to connect to any ClearQuest database. Note these when you create the vendor database to which you import the replica. For more information about how to create empty vendor databases and the parameters you will need, see the *Installing Rational ClearQuest* documentation.

**Note:** When you import a replica, you'll need to specify the database parameters of both the vendor database that will hold the replicated schema repository and the vendor database that will hold the user database replica. You must have first created these databases before importing a replica packet.

here *<db-params>* :=

```
if <vendor> == DB2
 <db-info> := Database Alias
 <db-params> := -dbologin dbo-name [dbo-pwd]
```

```
if <vendor> == ORACLE
 <db-info> := SQL*Net Alias
 <db-params> := -dbologin dbo-name [dbo-pwd] [-connectopts
connect-options]
```

```
if <vendor> == SQL_SERVER
```



*<db-info> := Physical Database Name*

*<db-params> := -server server-name -dbologin dbo-name [dbo-pwd]  
-rwlogin rw-login [rw-pwd] [-rologin ro-login [ro-pwd]]*

The read-only login (-rologin) for SQL\_SERVER is needed only when creating a replicated schema repository.

#### **--c-omments** *comments*

Enter any comments you want to include.

*packet-file-path | packet-dir-path ...*

Specifies a pathname of a replica-creation packet created by **mkreplica -export**. For a logical packet that spans multiple disk files, **mkreplica** scans the directory containing *packet-file-pname* for related physical packets.

If you also specify one or more *packet-dir-path* arguments, **mkreplica** searches for additional packets in these directories.

*Default:* None.

## Examples Exports

- Generate a replica-creation packet from the *boston* site of the *PRODA* database family, which will be used at remote host *sushi* to create a new replica named *osaka*. Place the packet in a file in directory *c:\temp*. The lines are broken for readability. You must enter the command on a single line.

```
multiutil mkreplica -export -clan telecomm -site boston
-family PRODA -user bostonadmin -password secret -c "make a
new replica for osaka" -out osakareplica sushi:osaka
```

- Similar to the preceding example but place the packet file in a storage bay, for shipping at some later time by the store-and-forward facility. The lines are broken for readability. You must enter the command on a single line.

```
multiutil mkreplica -export -clan telecomm -site boston
-family PRODA -user bostonadmin -password secret -c "make a
new replica for osaka" -ship -workdir c:\temp\working -sclass
cq_default -pexpire 22-November-2003 -notify
admin@company.com
```

- ## Imports
- Imports a new database replica *osaka* and its associated replicated schema repository into SQL Server databases.

```
multiutil mkreplica -import -osaka -repository cq_schemarepo
-vendor SQL_SERVER -server unagi -dbologin japanadmin secret
```

## mkreplica

```
-rwlogin japanadmin secret -rologin japanadmin secret
-database cq_userdb -vendor SQL_SERVER -dbologin juseradmin
secret -rwlogin juseradmin secret
```

- **Imports a new user database replica that is part of the telecomm clan at the osaka site. The new user database replica is being imported into a SQL Server database.**

```
multiutil mkreplica -import -clan telecomm -site osaka -user
cqjapanadmin -p secret
-database cq_userdb -vendor SQL_SERVER -dbologin juseradmin
secret -rwlogin juseradmin secret
```

**See Also**    **activate**

## MultiSite Control Panel

---

Configures store-and-forward facility

**Synopsis** %SystemRoot%\System32\ms.cpl

To open the MultiSite Control Panel, double-click the MultiSite icon in Windows Control Panel. TONIA'S QUESTION: enable e-mail notification info needs to go here.

**Description** The **MultiSite Control Panel** controls operation of the MultiSite store-and-forward facility on each host. MultiSite installation creates registry keys in which all these entries are defined. In some cases, the corresponding store-and-forward operation fails if a parameter is not defined, and in other cases there is a hard-coded default.

The **MultiSite Control Panel** provides controls for setting the configuration parameters described in the following sections.

**Maximum Packet Size** This setting controls the splitting of individual logical packets into multiple physical packets. This value specifies the maximum size for a physical packet file. Limiting the size of physical packets can improve the reliability of packet delivery in some networks. To specify no limit, use 0 (zero).

This value is used by the following commands (unless you also specify **-maxsize**):

- **mkreplica -fship**
- **mkreplica -ship**
- **syncreplica -fship**
- **syncreplica -ship**

When you invoke **mkreplica** or **syncreplica** with **-out** or **-tape**, this value is not used, and you must use **-maxsize** to limit the packet size.

*Default:* 2097151KB

**Administrator E-mail** TONIA'S QUESTION: In order to have e-mail work in cases where they aren't also running ClearCase MultiSite, the registry needs to be edited. That process is messing. Should I put it here anyway or release note it?

*Default:* None.

**Storage  
Classes**

**Storage Class Name**

This setting specifies the name of a storage class. You can associate values for packet expiration, the storage bay, the return bay, and the receipt handler with each storage class.

mention setting up the `cq_default` class

*Default:* None.

**Packet Expiration**

This setting specifies the expiration period (in days) for shipping orders generated in the associated storage class. This period begins at the time the shipping order is generated.

If a packet cannot be delivered to all its specified destinations in the specified number of days, the packet is returned to the original sending host and a message is sent to the address specified in the **Administrator Email** box.

If e-mail notification is not enabled, a message is written to the Windows Event Viewer.

A value of 0 (zero) specifies no expiration; delivery is reattempted indefinitely.

This setting is overridden by the `-pexpire` option to `syncreplica` or `mkreplica`.

**PACKET REDELIVERY.** The `shipping_server` program does not retry delivery of packets. The Packet Expiration specification is useful only if you set up a host to periodically attempt delivery of any undelivered packets.

*Default:* When the **Use Default Expiration** check box is selected, the storage class uses the packet expiration value associated with the `-cq_default` class. (This value is not shown in the **Packet Expiration** box; you must display the `-default` class to determine the value.) When MultiSite is installed for the first time, the Packet Expiration value for the `-cq_default` class is set to 14 days.

### Storage Bay Path

This setting defines the location of a storage bay, a directory that holds the outgoing and incoming update packets and shipping orders of a particular storage class.

Packets placed in a storage bay on an NTFS file system inherit the Windows ACL on the storage bay. Define ACLs on the storage and return bays to enable successful execution of MultiSite commands to process the packets, and to guard against unauthorized access.

Packets stored on FAT file systems have no protections.

Before using the store-and-forward facility, make sure that the disk partition where you will have your storage bays has sufficient free space for anticipated replica-creation and update packets. The amount of available space on the disk partition where the shipping and return bays are located must be at least twice as big as the largest packet that will be stored in the bays. This is because there may be two copies of the same packet in the bay at one time: one on its way to another destination, and another waiting to be applied to the replica on the host.

**Note:** When you create a new storage class, the storage bay and return bay that you specify are created, along with the *incoming* and *outgoing* directories within the bays.

### Return Bay Path

This setting defines a return bay (directory) to hold incoming or outgoing packets in the process of being returned to their origin because they could not be delivered to all specified destinations.

Packets placed in a return bay inherit the ACL on the directory.

## Receipt Handler Path

This setting specifies a batch file or program for the **shipping\_server** to run when a packet is received for the storage class. You can use this if you are using a scheduling program to automate the sending and applying of packets. By default, no file is specified.

For each packet that is received, **shipping\_server** does the following:

- 1 Reads the entries in the MultiSite Control Panel to find the appropriate **Receipt Handler** value for the packet.
  - If the packet is associated with a storage class and there is a **Receipt Handler** value for that storage class, **shipping\_server** uses the specified batch file or program.
  - If the packet is not associated with a storage class and there is a **Receipt Handler** value for the **-cq\_default** storage class, **shipping\_server** uses the batch file or program specified for **-cq\_default**.
- 2 Invokes the receipt handler, as follows:

```
script-pname [-d-ata packet-file-pname] [-a-ctual shipping-order-pname]
[-s-class storage-class] -o-rigin hostname
```

where

|                                             |                                                                                                                                              |
|---------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|
| <i>script-pname</i>                         | Script specified in the <b>RECEIPT-HANDLER</b> entry.                                                                                        |
| <b>-d-ata</b> <i>packet-file-pname</i>      | Location of the packet. This parameter is used only when the packet is destined for this host.                                               |
| <b>-a-ctual</b> <i>shipping-order-pname</i> | Location of the shipping order. This parameter is used only when the packet is destined for another host.                                    |
| <b>-s-class</b> <i>storage-class</i>        | Storage class associated with the packet. This parameter is used only if the packet was associated with a storage class when it was created. |
| <b>-o-rigin</b> <i>hostname</i>             | Host name of the machine from which the packet was first sent.                                                                               |

**Note:** If a packet is destined for both the local host and another host, both the **-data** and **-actual** parameters are used. The packet is imported at the replica on the host, then forwarded to its next destination.

*Default:* None.

## ROUTING INFORMATION

The Routing Information settings control the network routing of packets.

### Next Routing Hop

The host specified here is the next destination for packets whose final destination is any of the host names specified in the **Destination Hostnames** list. This host is responsible for delivery of the packet to its destinations. You can specify a host using either its host name (which must be usable by hosts in different domains) or its numeric IP address.

*Default:* None.

### Destination Host Names

Packets destined for any host listed in this field are sent to the host specified in the **Next Routing Hop** box. The value **-default** as the **Destination Hostname** accommodates all hosts that are not associated with a routing hop. You can specify a host using either its host name (which must be usable by hosts in different domains) or its numeric IP address.

*Default:* None.

## recoverpacket

---

Resets epoch row table so changes in lost packets are resent.

**Synopsis**    **recoverpacket** [-cl-an *clan-name*] [-site *site-name*] -fam-ily *family-name*  
                   -u-ser *username* [-p-assword] *password* [ -sin-ce *date-time* ] *replica* ...

**Description**    The **recoverpacket** command resets the epoch row at a sending replica to reflect the last synchronization sent to a receiving replica before a particular time. It scans through a list of epoch rows saved at the time of each export, looking for an entry prior to the time specified. When it finds this entry, it uses the associated row to reset the epoch row for just the specified receiving replica. The next time a packet is sent, it includes the changes that were in the lost packet.

### Resetting Epoch Numbers Automatically

When you send an update packet to another replica, success of the transport and import phases is assumed. Therefore, the sending replica's epoch number matrix is updated to reflect that the changes are made at the receiving replica. However, if the packet is lost before reaching the receiving replica, the sending replica's assumption that the receiving replica is up to date is incorrect.

The updated epoch numbers must be returned to the values they had before the packet was sent. Making these corrections to the sending replica's epoch number matrix causes it to include the same changes in the next update packet it sends to the receiving replica.

The administrator at the receiving replica needs to run an **dumpoplog** command to determine the time of the last successful import. The administrator at the sending replica uses this time in the **recoverpacket** command.

**Note:** If the two sites are not in the same time zone, or you do not send packets at the same time you generate them (for example, you generate packets at midnight and send them at 6:00 A.M.), you must adjust for the time.

### Resetting Epoch Numbers Manually

If there are no saved epoch rows for the replica that are as old as the specified time, the **recoverpacket** command fails. In this case, the administrator at the receiving site must use the **lsepoch** command to determine the correct epoch



number, and the administrator at the sending site must run **chepoch** on the sending replica to reset the epoch row. See the **chepoch** reference page for more information.

## Options and Arguments

### **-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

### **-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required..

### **-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** Use the MASTR family if you want to recover information about the working schema repository. If you've lost a packet, run **recoverpacket** on both the MASTR and user database families, then re-execute **syncreplica**.

### **-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

### **-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**–since** *date-time*

All oplog entries after the date-time you specify will be listed.

The *date-time* argument must be in the following format:

m[m]/d[d]/yyyy h[h]:m[m]:s[s]

For example:

10/7/2002 17:04:32

**replica...**

Specifies the replica for which the epoch row is reset. Use the site name to specify the replica.

**Examples**

- Reset the epoch row for replica **sanfran\_hub** so that changes sent since last Monday will be included in the next packet that is sent.

**example here**

- Reset the epoch row for replica **lex\_hub** so that the changes included in the most recent update packet will be included in the next packet that is sent.

**example here**

- Determine the last successful import at replica **bangalore**, reset the epoch row at replica **lex\_hub**, and send an update packet.

- At replica **bangalore**:

**example here**

- At replica **lex\_hub** (remember to adjust for the time zone difference):

**example here**

**See Also**

- **chepoch**,
- **lsepoch**
- **restorereplica**

## rename

---

Renames a replica

**Synopsis**    **rename** [-**cl-an** *clan-name*] [-**site** *site-name*]  
               -**u-ser** *username* [-**p-assword**] *password*  
               { *old-replica new-replica* }

**Description**    This option renames a replica. You can only rename your own site.

**Options and Arguments**

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

*old-replica new-replica*

Enter the current name of the replica followed by the new name of the replica. Use the site name for the replica name.

**See Also**    • **sync replica**

## restorereplica

---

Recovers packets generated since the restoration of a replica from backup.

**Synopsis**    **restorereplica** [-**cl-an** *clan-name*] [-**site** *site-name*] **-fam-ily** *family-name*  
**-u-ser** *username* [-**p-assword**] *password* [-**force**] [-**completed**]  
 [-**replace** ] [*replica ...*]

**Description**    *Execute this command IMMEDIATELY after restoring a replica from backup. Proceeding with normal development (and generating new changes) at a restored replica before executing this command can lead to IRREPARABLE inconsistencies among the replicas in a replica family.*

**restorereplica** replaces missing changes in a replica that has been restored from backup, as follows:

- 1 Causes the current replica to create special update packets that contain update requests to other replicas. These packets are sent to other replicas via with **syncreplica -export**.
- 2 Locks the current replica and marks the replica as being in the process of restoration.
- 3 Requests the recovery information for the replica.
- 4 Causes **lsreplica -long** to indicate which replicas must send restoration updates to the current replica.

The current replica remains in the restoration state until your site has received and applied (using **syncreplica -import**) all the restoration updates needed to bring the replica up to date with the state of the family. Collectively, these updates include all the changes to the family since the backup was made, including changes made in the current replica before its failure. During the process of restoration, the **lsreplica -long** command annotates its listing to indicate which replicas must send restoration updates to the replica.

For a description of the replica restoration procedure, see *Restoring replicas* on page 94.

## Locking Of The Replica

**restorereplica** locks the current replica. This ensures that while restoration proceeds through execution of **syncreplica –export** and **syncreplica –import** commands, no other changes are made to the current replica.

When **syncreplica** applies the final required update, it displays a message indicating that the restoration process is complete and unlocks the replica.

## Optimizing The Restoration Process

By default, **restorereplica** requires that the replica receive restoration updates from all other replicas in its family (either directly or indirectly). Only after all the updates are imported does the **syncreplica** command display the message indicating that restoration is complete.

In some cases, you can relax this requirement without compromising the correctness of the restoration process. The replica will be brought up to date if it receives a restoration update from only one replica—the last one to which the replica sent an update before it was restored from the backup version. You can specify the name of that last-updated replica (or a list of replicas, one of which must be the last-updated one) to **restorereplica**. **syncreplica** displays the restoration-completed message after receiving restoration updates from all the specified replicas.

The **restorereplica** command requires restoration updates from *all* other members of its family. The **syncreplica** command declares the replica to be restored completely only after all the updates have been processed.

**Warning:** Making a mistake in using this optimization can make the restored replica irreparably inconsistent with other replicas.

**Restrictions**     *Locks:* The database is locked during the restore operation. Read-only commands such as `lsepoch` and `lsreplica` can still access the database.

Options and Arguments

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** Not applicable. Restoring a member of user database family automatically requests updates and its associated replicated schema repository, if necessary.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-f-orce**

Suppresses the confirmation step.

*Default:* **restorereplica** prompts you for confirmation.

**Warning:** Incorrect use of this option allows new changes to be made to the replica before all missing changes are received from other replicas. This may place the entire replica family in an irreparably inconsistent state.

**–completed**

Use this option when you are finished restoring the replica. This option marks the replica as restored and unlocks the database. If this option is used, no more packet requests can be sent and no more restoration packets can be played at this replica.

**–rep-lace** *replica...*

Specifies the subset of replicas from which restoration updates are required. Use the site name(s) to specify the replica(s) from which you want to receive restoration packets.

**Examples** For an example of restoring a replica, see *Restoring replicas* on page 94.

**See Also**

- **chepoch**
- **lsepoch**
- **lsreplica**
- **syncreplica**

## rmreplica

---

Deletes a replica.

**Synopsis** **rmreplica** [**-cl-an** *clan-name*] [**-site** *site-name*] **-fam-ily** *family-name*  
**-u-ser** *username* [**-p-assword**] *password* [**-dbset** *new-name*] *replica*

**Description** **Warning:** To delete a replica, you must complete all steps described in *Removing a replica* on page 80. If you do not complete all steps in the correct order, synchronization and mastership problems can occur in other replicas in the database family.

This command deletes from the current replica's database the database-replica record that records the existence and identity of another replica. Typically, you use this command at your site to record the fact that a replica at another site has been decommissioned and deleted.

### Options and Arguments

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** Not applicable. If there is only one user database family at the site specified, this command removes the replicated schema repository as well. If there is more than one user database family at the site, the schema repository will not be removed.

*Default:* None.



**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-dbset** *new-name*

This option is only used when removing the last replica of a clan. When you remove the last replica of a clan, you need to rename the dbset name so it does not include any ClearQuest MultiSite flags.

*replica*

Specifies the replica to be deleted from the current replica's database. Use the site name to specify the replica.

*Default:* Defaults to the current replica when running the command from the location of the replica to be removed. This argument is mandatory when running the command from any other replica location.

**Examples**

- Remove the replica record that records the existence of replica **paris** from the database of the current replica site, **Boston**.

```
multiutil rmreplica -clan telecomm -site Boston -family PRODA
-user admin -password "" paris
```

- Remove the last replica of a clan and return the existing replica to a non-replicated state.

```
multiutil rmreplica -user admin -password "" -dbset
cqdatabase
```

**See Also**

- **chmaster**
- **mkreplica**

## scruboplog

---

Deletes oplog entries associated with the respective replica.

**Synopsis** **scruboplog** [**-cl-an** *clan-name*] [**-site** *site-name*] **-fam-ily** *family-name* [**-u-ser** *username*] [**-p-assword** *password*] **-before** { *date-time* | *oplog-number* }

**Description** Oplog entries must be kept in the replica for a significant period. In the near term, they are required when the replica generates update packets to be sent to all other replicas. Beyond that, oplog entries may be required to help other replicas recover from failures.

However, you may want to delete (scrub) oplog entries occasionally to optimize hard drive space where the replica resides. You can also use the scruboplog command to delete the oplog of a replica which will no longer be used.

Although oplog entries only record the changes that have taken place in your database, overtime this information could require as much space as the data itself.

Before scrubbing (deleting) oplog entries for a replica, you should be sure that they are no longer needed and that each replica has the current information from the replica's oplog that you want to delete.

### Options and Arguments

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.

- **replicated schema repository family:** To specify the family of the replicated schema repository use the family name, MASTR.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-before** { *date-time* | *oplog-number* }

You can either enter the date-time of the oplog you want to delete or you can specify a number. All oplog entries previous to the date-time or oplog number you specify will be deleted.

The *date-time* argument must be in the following format:

m[m]/d[d]/yyyy h[h]:m[m]:s[s]

For example:

10/7/2002 17:04:32

The *oplog-number* argument should be entered as an integer. If you're unsure of what oplog number to specify, use **dumpoplog** to find out more information about the replica's oplog you want to delete.

- See Also
- **dumpoplog**
  - **syncreplica**

## shipping\_server

---

Store-and-forward packet transport server.

**Synopsis**    **shipping\_server** [ **-scl-*ass*** *storage-class-name* ] { **-pol-*l*** | *sources ...* }

**Description**    This command processes one or more shipping orders on the local host, causing the associated packets (or other files) to be sent to remote sites.

After delivering the data file specified in a shipping order to all its destinations, **shipping\_server** deletes the data file unless one of the destinations is the local host.

**Note:** When **shipping\_server** starts processing a shipping order, it locks the order. This prevents subsequent invocations of **shipping\_server** from processing the order.

### TCP/IP Connection

To transmit a file, **shipping\_server** uses UDP to contact the **albd\_server** process on the receiving host, and **albd\_server** invokes **shipping\_server** in receive mode on the receiving host.

After a TCP connection is established between the two **shipping\_server** processes, they transfer the file. The receiving **shipping\_server** selects a storage bay using the local store-and-forward configuration settings. See the **MultiSite Control Panel** (Windows). If a storage class is assigned multiple storage bays, available disk space determines the selection of a bay.

After a TCP connection is established between the two **shipping\_server** processes, they transfer the file. The receiving **shipping\_server** selects a storage bay using the local store-and-forward configuration settings. See the **MultiSite Control Panel** (Windows). If a storage class is assigned multiple storage bays, available disk space determines the selection of a bay.

**Windows:** The packet file inherits permissions from the Windows ACL on the storage bay directory.

### Colon Characters in Packet Names

If a packet name contains a colon ( : ), **shipping\_server** changes the colon to a period ( . ) during processing. This change allows packets to be delivered to Windows machines, which do not allow colons within file names.

## Handling of File Name Conflicts

You can use the **mkorder** and **shipping\_server** commands to transmit arbitrary files if the files are located in the same directory as their associated shipping orders. If a file with the same name already exists on the receiving host, the new file is renamed to *filename\_1* (if you send another file with the same name, it is renamed to *filename\_2*, and so on).

## Log File

**shipping\_server** writes records of all packets sent and received, notification messages, and all errors to the Windows event viewer. It writes log messages to file *ccase-home-dir\var\log\shipping\_server\_log*.

**TONIA's QUESTION: Where does CQMS write log messages to, if anywhere. The directory listed above doesn't exist in the CQMS install.**

**Restrictions**      The shipping order and the data file it specifies must be located in the same directory.

**Options and Arguments**      **-scl-*ass*** *storage-class-name*

Processes shipping orders for the specified storage class only.

*Default:* Processes all shipping orders specified or found in a search.

**-pol-*l***

Processes shipping orders located in some (if you use **-sclass**) or all MultiSite storage bays defined in the **MultiSite Control Panel** on Windows.

Processes shipping orders located in some (if you use **-sclass**) or all MultiSite storage bays defined in the **MultiSite Control Panel** on Windows.

*Default:* None.

**Note:** **shipping\_server** processes only shipping orders whose file names start with the characters "sh\_o\_". If you create shipping orders, name them according to this convention, or omit the **-poll** option and specify the shipping order pathnames.

**sources ...**

One or more pathnames of files and/or directories. Each file you specify is processed if it contains a valid shipping order. For each directory you specify, **shipping\_server** processes some (if you use **-sclass**) or all shipping orders stored in that directory.

## shipping\_server

- Examples**
- Process all shipping orders in all MultiSite storage bays.  
**shipping\_server -poll**
  - Process a particular shipping order. Note that the pathname argument specifies the shipping order file, not the data file to be transmitted.
  - Process all shipping order files in a specified directory.  
**shipping\_server**  
**"c:\Program Files\Rational\ClearQuest\var\shipping\ms\_ship"**
  - Process all shipping orders in the storage bays of a specified storage class.
- See Also**
- **mkorder, MultiSite Control Panel**

## syncreplica

---

Generates or applies update packets.

**Synopsis**    Export an update packet:

```
sync-replica --export [-cl-an clan-name] [-site site-name] -fam-ily
family-name
 -u-ser username [-p-assword] password [-max-size max-packet-size
 [-lim-it num-packets]] -wor-kdir directory
 { { -sh-ip | -fsh-ip }
 [-scl-ass storage-class] [-pex-pire date]
 [-not-ify email] | -out { packet-file-pname | staging-area-pname }
 } replica ...
```

Import an update packet:

```
sync-replica --imp-ort [-cl-an clan-name] [-site site-name] -fam-ily
family-name
 -u-ser username [-p-assword] password
 { -rec-eive [-scl-ass storage-class]
 | { packet-file-pname | staging-area-pname } ...
 }
```

**Description**    Synchronization of an existing replica with one or more replicas at other sites is a two-phase process:

- 1 At one site, a **syncreplica --export** command creates an update packet that contains changes that have occurred in the replica at that site (and perhaps other replicas, as well).
- 2 At another site, a **syncreplica --import** command applies the changes in the update packet to its replica of the same database.

Step #2 occurs at all sites that receive the packet.

Contents of an update packet:

- All changes that have occurred in the current database replica since the last update generated for the destination replicas. (Changes already sent to the destination replicas are excluded from the packet).
- Changes that have occurred in other replicas, which the current replica has received in previous update packets from those replicas, but has not already passed on to the destination replicas.

In all cases, **syncreplica –export** creates a single logical update packet for use at all the specified destinations; the packet can be used to update *only* those particular replicas.

### Notes On The Export Phase

MultiSite is designed for efficient updating of replicas. **syncreplica –export** attempts to exclude operations that have been sent previously. (However, there is no harm in sending the same update packet multiple times to the same replica; the first operation is imported and subsequent identical operations are ignored.)

### Specifying a Directory for Temporary Files

**syncreplica –export** stores temporary files in the directory specified by the TMP environment variable on Windows.

### Notes On The Import Phase

An update packet is applied to the appropriate replicas on the host on which you import it. You do not have to specify particular replicas or storage locations.

The import process applies update packets in the correct order; you can specify packets in any order on the command line.

The database replica is not locked for normal database operations during the import phase. But it is locked for all other MultiSite operations.

### Specifying a Directory for Temporary Files

**syncreplica –import** stores temporary files in the directory specified by the TMP environment variable on Windows.

### Skipping Packets

**syncreplica –import** refuses to process an update packet in the following situations:

- The update packet contains changes that depend on other changes that have not yet been applied to this replica. This usually means that an update packet destined for this replica has not been sent or was lost during transport.
- Problems were encountered processing an earlier physical packet in a multiple-part logical packet.



In any of these cases, **sync replica –import** displays an explanatory message.

### Update Failures / Replaying Packets

In some cases, **sync replica –import** begins to apply operations to a replica, but terminates with an error message. For example, another process locks the database, causing the import to fail. After the database is unlocked, you can have **sync replica –import** process the entire update packet again.

There is no harm in importing update packets that have already been processed successfully; the same change will not be made twice. Thus, even importing an entire update packet multiple times causes no error and does no harm.

See *Recovering from lost packets* on page 92, for more information on update failures.

### Deletion of Update Packets

If a single invocation of **sync replica –import** applies a packet successfully to all target replicas on the host, the update packet is deleted when the command completes its work. If the packet is processed with multiple **sync replica –import** commands, it is not deleted.

### Hooks Firing

ClearQuest hooks do not fire in response to changes made during packet import.

### Handling Naming Conflicts

**sync replica** resolves naming conflicts among type objects created at different replicas. For more information, see [.<more info later>](#)

### Delayed View Updates

**sync replica** does not inform any ClearQuest users of the updates to replicas. All active users see updates within a few seconds, through ClearQuest's normal database-polling routines.

### Error Handling for Packet-delivery Failures.

If a packet cannot be delivered, it is sent through the store-and-forward facility back to the administrator at the site of the originating replica. A mail message is sent to the store-and-forward administrator. This occurs after repeated attempts to deliver the packet have all failed, and the allotted time

has expired; it can also occur when the destination host is unknown or a data file does not exist. The store-and-forward configuration settings specify the expiration period and the e-mail address of the administrator.

## Options and Arguments

### Export Phase

The following sections describe the options and arguments for use with **syncreplica –export**.

#### **-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

#### **-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

#### **-fam-ily** *family-name*

- **user database family:** Use the name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.
- **replicated schema repository family:** Not applicable. Update packets automatically include updates to the replicated schema repository, if any.

*Default:* None.

#### **-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

#### **-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**–max-size** *max-packet-size* [ **–lim-it** *num-packets* ]

The maximum size for a physical packet, expressed as a number followed by a single letter. For example:

|      |               |
|------|---------------|
| 500k | 500 kilobytes |
| 20m  | 20 megabytes  |
| 1.5g | 1.5 gigabytes |

[ **–lim-it** *num-packets* ]

The **–limit** option limits the number of packets **syncreplica** generates; each packet is no larger than *max-packet-size*. Use this option when the disk space for your shipping bay or staging area is limited.

*Default:* When you do not specify **–maxsize**, the default packet size depends on the shipping method you use:

- Packets created with **–ship** or **–fship** are no larger than the maximum packet size specified in the **MultiSite Control Panel** (Windows).
- Packets created with **–out** are no larger than 2 GB.

**–shi-p**

**–fsh-ip**

Stores the update packet in one or more files in a store-and-forward storage bay; **syncreplica** creates a separate shipping order for each physical packet, indicating how and where it is to be delivered. The destinations are the host names associated in the replica database with the *replica-name* arguments. (Replica-name/host-name associations are created with **mkreplica –export** and can be changed with **chreplica**.)

Using **–fship** (force ship) invokes the **shipping\_server** to send the update packet immediately. Using **–ship** does not invoke this server.

*Default:* None. You must specify how the update packets created by **syncreplica –export** are to be stored and/or transmitted to other sites.

**–wor-kdir** *directory*

Specify a temporary working directory for **syncreplica** to use. This directory must not already exist and is deleted after the **syncreplica export** process.

**-scl-ass** *class-name*

Specifies the storage class of the packet and shipping order. **syncreplica** looks up the storage class in the **MultiSite Control Panel** on Windows to determine the location of the storage bay to use.

If you omit this option, **syncreplica** places the packet in the storage bay location specified for the **-cq\_default** class **MultiSite Control Panel**.

**-out** *packet-file-pname*

Places the first update packet in file *packet-file-pname*. Additional physical packets, if any, are placed in files named *packet-file-pname\_2*, *packet-file-pname\_3*, and so on.

Places the packets in the directory indicated with *staging-file-name*.

The update packets are not delivered automatically; use an appropriate mechanism (e-mail, ftp, postal service, and so on) to deliver them.

You can create a packet using **-out**, and deliver it using the store-and-forward facility. See the **mkorder** reference page.

*staging-area-path*

Name a directory name where packet files generated.

**-pex-pire** *date-time*

Specifies the time at which the store-and-forward facility stops attempting to deliver the packet and generates a failure mail message instead.

**Windows:** This option overrides the storage class's Packet Expiration specification in the MultiSite Control Panel. See the **MultiSite Control Panel** reference page for a description of this specification, and of delivery retries in general.

The *date-time* argument can have any of the following formats:

*date.time* | *date* | *time*

where:

*date* := *day-of-week* | *long-date*

*time* := *h[h]:m[m][:s[s]]* [UTC [ [+ | - ]*h[h]:m[m]* ] ]

*day-of-week* := **today** | **yesterday** | **Sunday** | ... | **Saturday** | **Sun** | ... | **Sat**

*long-date* := *d[d]-month[-[yy]yy]*

*month* := **January** | ... | **December** | **Jan** | ... | **Dec**

Specify the *time* in 24-hour format, relative to the local time zone. If you omit the time, the default value is **00:00:00**. If you omit the *date*, the default value is **today**. If you omit the century, year, or a specific date, the most recent one is used. Specify **UTC** if you want the time to be resolved to the same moment in time regardless of time zone. Use the plus (+) or minus (-) operator to specify a positive or negative offset to the UTC time. If you specify **UTC** without hour or minute offsets, the default setting is Greenwich Mean Time (GMT). (Dates before January 1, 1970 Universal Coordinated Time (UTC) are invalid.)

Examples:

```
22-November-1998
sunday
yesterday.16:00
8-jun
13:00
today
9-Aug.10:00UTC
```

**-not-ify** *e-mail-address*

The delivery-failure message is sent to the specified e-mail address.

If a failure occurs on a Windows host that does not have e-mail notification enabled, a message appears in the Windows Event Viewer. The message includes the *e-mail-address* value specified with this option and a note requesting that this user be informed of the status of the operation. For information on enabling e-mail notification, see the **MultiSite Control Panel** reference page.

**replica ...**

You can specify one or more destination replicas. Use the site name to refer to a replica. For example, *sitea* indicates that site a will receive the update packet, while *sitea siteb* indicates that both sites a and b will receive the update packet.

*Default:* None.

**Options and Arguments**

**Import Phase**

The following sections describe the options and arguments for use with **syncreplica -import**.

**-cl-an** *clan-name*

The name of the clan to which the replica belongs. All the replicas associated with a single schema repository form a clan. The clan name is specified when the replica is activated.

*Default:* The first clan replicated at this site. If there is more than one clan at this site, this argument is required.

**-site** *site-name*

The site name of the replica to be modified.

*Default:* Current site. If there is more than replica on this machine, this argument is required.

**-fam-ily** *family-name*

The name of the family to which this replica belongs. The replicas of a single user database form a family. The family name is five-character logical database name that was given to the user database when it was created within ClearQuest.

**-u-ser** *user*

The name of a user with the appropriate database privileges. You must have Super User privileges to administer ClearQuest MultiSite.

**-p-assword** *password*

The password associated with the user account being used to run this command. Note that you must have Super User privileges to administer ClearQuest MultiSite.

**-rec-eive** [ **-scl-ass** *storage-class* ]

Scans one or more of the current host's storage bays. Any unprocessed update packets intended for this host are applied to the appropriate replicas on the host. Using the **-sclass** option restricts processing to the storage bays of the specified storage class.

*Default:* If **-sclass** option is not specified, this command updates all replicas that are on the current host and are specified in the update packets.

If **syncreplica** finds any replica-creation packets, it sends mail to the store-and-forward administrator. (If the current host is a Windows host and there is no valid host specified in the **SMTP Host** box in the MultiSite Control Panel, a message appears in the Windows Event Viewer.) Use **mkreplica** to import these replica-creation packets.

*packet-file-pname | staging-area-pname ...*

Processes each *packet-file-pname* as an update packet. For each *staging-area-pname* specified, locates all previously unprocessed update packets in the directory and applies them to the appropriate replicas.

*Default:* There is no default update packet location.

## Examples Exports

- Generate an update packet to be sent to replica **lex\_hub**. Store the packet in a file in directory **d:\transfer\packet**  

```
syncreplica -export -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -maxsize 500mb -workdir
c:\work -out
```
- Similar to preceding example, but place the packet file in a storage bay, for shipping at some later time by the store-and-forward facility.

```
syncreplica -export -clan telecomm -site boston -family PRODA
-user bostonadmin -password secret -maxsize 500mb -workdir
c:\work -ship -sclass cq_default -pexpire
```

## syncreplica

- Similar to preceding example, but ship the packet immediately.  
sample code here

### Imports

- Process an incoming update packet in directory **/usr/tmp**.  
sample code here
- Process all incoming update packets in the current host's storage bays.  
sample code here

### See Also

- **mkorder**
- **mkreplica**
- **MultiSite Control Panel**



# Index

## A

- ACLs
  - storage bays 145
- administration
  - disk space for storage bays 30
  - scrubbing 83
- albd\_server, control of ports used 39
- asterisks in packet listings 117

## B

- bidirectional synchronization
  - about 19
- branch mastership
  - transfer procedure 72

## C

- ccase-home-dir* directory iv
- chepoch command 101
- chmaster command 104, 110, 112, 151
- chreplica command 108
- CLEARCASE\_MAX\_PORT environment variable 39
- CLEARCASE\_MIN\_PORT environment variable 39
- cleartool and multitool commands 42
- commands for MultiSite
  - about 41
  - ClearCase 44
  - multitool 41
- conventions, typographical iv
- creating replicas
  - command for 99, 128
  - common problems 85
  - export procedure 50
  - import procedure 52
  - scenario 49
  - with store-and-forward facility 52, 61

## D

- disk space
  - replica-creation directory 50
  - storage bays 30
- documentation
  - MultiSite online help description v

## E

- elements
  - transfer of mastership 73
- e-mail and firewalls 36
- encryption of update packets 39
- environment variables 39
- epoch number matrix
  - about 12
  - listing contents of 13, 43, 115
  - zeros in 14
- epoch numbers
  - about 8
  - changing, commands for 101, 148
  - changing, methods for 93
  - gaps in 86
  - role in updates 11
- error messages
  - Gap in oplog entries 86
  - Replica already exists 85
  - transport operations, list of 88
- event records
  - about 8
- export operations
  - packets accumulate in storage bay 86
  - replica creation 50, 53, 62
  - replica-creation packets, recovering lost 93
  - resending lost packets 148
  - synchronization problems 86
  - synchronization procedure, manual 64
  - update packet delivery patterns 57

## F

firewalls

- shipping\_server on 38

- synchronization and 36

ftp and firewalls 37

## H

help v

host name of replica, changing 77

## I

import operations

- assumption of success 60

- common synchronization problems 90

- failures, possible causes 92

- lost packets 90, 92

- replica creation 52

- replicasynchronization 63

- synchronization procedure, manual 64

installation and licensing 30

## L

licenses needed for ClearCase and MultiSite 30

lsepoch command 43, 115

lspacket command 117

lsreplica command 120

## M

make database replica 128

man command 41

mastership

- about 2

- chmaster command description 104, 110, 112, 151, 158

- elements, transferring 73

- fixing accidental change in 74

- management of 65

- objects in removed replicas 80

- transferring, replica removal 72

mkorder command 42, 124

mkreplica command 128

mkreplica command 99

MultiSite Control Panel 31, 32, 143

multitool commands

- about 41

- summary 42

## O

online help v

oplogs (operation logs)

- about 8

- gaps in epoch numbers 86

- scrubbing 83

ownership-preserving replicas

- changing properties of 108

## P

packets

- about 6

- listing contents of 117

- logical and physical 6

- processing imported 6

- redelivering 144

- routing 147

- splitting logical into physical 143

- submitting to store-and-forward facility 29

planning issues

- about 17

- design documentation 18

- firewalls 38

- licensing 30

- synchronization strategy 23

- time zones and synchronization strategy 59

ports, control of for servers 39

## R

receipt handlers, paths 146

recoverpacket command 148

- replica objects
  - deleting 156
  - transferring mastership of 73
- replica-creation packets
  - contents and cleanup 130
  - how to split 53, 63
- replicas
  - accidental deletion, recovery 96
  - changing hosts or host names 77, 108
  - displaying properties of 76
  - history of changes, how tracked 9
  - lsreplica command description 120
  - moving 78
  - removal procedure 80
  - renaming 77
  - resolving name conflicts 7
  - restoring from backup 94
  - scrubbing oplogs 83
  - self-mastering 73
  - transferring mastership of objects in 72
- restorereplica command 152
- return bays
  - about 29
  - ACLs 145
  - handling packets in 90
  - paths 145
- rmreplica command 156

## S

- scrubbing 83
- scruboplog command 158
- shipping orders
  - about 29
  - creating 42, 124
  - expiration date, specifying 144
  - expired 34, 90
  - processing 160
- shipping\_server
  - about 29, 160
  - error handling mechanisms 33
  - installing on firewall 38
  - log file 161
  - troubleshooting scenarios 86

- sites
  - about 1
  - documentation of design 18
- storage bays
  - about 29
  - ACLs 145
  - disk space requirements 30
  - packets in 86, 90
  - paths 145
- storage classes
  - naming 144
  - use in synchronization 31
- storage directories, restoring lost 94
- store-and-forward facility
  - about 28
  - creating replicas 52, 61
  - creating shipping orders 42, 124
  - customizing 143
  - deliveries attempted 33
  - firewalls 38
  - indirect shipping routes 33
  - reliability of and packet size 53, 63
  - sending files with 32
  - storage classes 31
  - submitting packets 29
  - use with firewalls 36
- synchronizing replicas
  - about 6, 27, 55
  - assumption of success 60
  - automating 35, 63
  - command for 163
  - common export problems 86
  - deliveries attempted 33
  - delivery patterns 57
  - firewalls, methods for handling 36
  - indirect routes 14
  - manual procedure 64
  - planning issues 23
  - risks of scrubbing oplogs 83
  - risks of unidirectional scheme 19
  - role of epoch numbers 11
  - schedule for 59
  - unidirectional vs. bidirectional 19
  - VOB database mechanism 8
- syncreplica command 163

examples 64

## T

TCP applications and firewalls 38

time stamps, interpretation of format 6

time zones, issues for synchronization strategies 59

topology for update packets 57

transport operations

- common problems 88

- delivery failure 89

- delivery mechanisms 6

- firewalls 36

- indirect routes 33

- invalid destinations 89

- recommended methods 24

- replica creation 51

- shipping order expired 90

- shipping\_server 160

- store-and-forward facility 28

- synchronization procedure, manual 64

triggers

- firing during synchronization 165

troubleshooting

- about 85

- accidental transfer of mastership 74

- deliveries, reattempting 33

- delivery failed 89

- expired shipping order 90

- export of update packets 86

- gap in oplog entries 86

- import failed 92

- import problems 90

- incoming packets accumulate 90

- invalid destinations 89

- lost packets 92

- packet size for store-and-forward facility 53, 63

- recovery from VOB server crash 94

- replica already exists 85

- replica deleted 96

- replica-creation problems 85

- shipping\_server log files 161

- shipping\_server problems 86

- success of synchronization 60

- synchronization and scrubbed oplogs 83

- transport problems 88

- update packet creation 86

typographical conventions iv

## U

unidirectional synchronization

- about 19

- risks 19

update packets

- contents of 163

- creating manually 64

- deleting 165

- encryption 39

- error notification in mixed environments 32

- storage classes 31

## V

views

- updating with replica changes 165

VOB database, mechanism for replica  
synchronization 8

VOB families

- listing members of 120