

Domino Replication Fundamentals and Troubleshooting

Author : Robert Steen

Index :

1. Introduction
2. Elements of Replication
3. Creating and Deleting Replicas
4. Running Replication
5. Replication Configurations
6. Other Areas of Replication
7. Logging and Debug options.
8. Troubleshooting

Chapter 1 : Introduction

Here we will be covering :

- 1.1 The aims of this document
- 1.2 What is Replication
- 1.3 Basic terminology

1.1 The Aims of this document

The aims of this document are to clearly explain the most commonly used aspects of Domino Replication in a Domino server environment. Then to continue the process by describing the most common issues seen with Domino Replication and detailing several troubleshooting techniques for resolving them.

1.2 What is replication

Replication is the process of exchanging modifications between certain Domino Databases or applications. Through replication, Domino can modify all of the replicas to be essentially identical over time.

This is an overall definition, of course there are settings that can limit the scope of the replication and schedule when it is triggered.

1.3 Basic terminology

Replication

The process used to exchange differences between databases.

Replica

A copy of a database that, because it shares a Replica ID with another database, can exchange information with it through replication.

More information on the Replica ID is below in Chapter 2 Elements of Replication.

Replicate

To update database replicas that are on different servers or on your workstation and a server. You can replicate the entire database, so that over time all database replicas are essentially identical, or select specific items or areas to replicate.

Replicator

The Domino task that performs the Replication.

Source

The database that is replicated from. Typically the server or Notes Client that initiates the replication.

Target

The database that is replicated to. Typically on the remote server.

Note

An individual element in a database, often referred to as a Document, but Note will also encompass design, ACL and other elements of the database.

Local

Usually refers to a database on a Notes Client. If a database is described as being Local it is not being accessed through a Domino Server but directly from a drive using only a Notes Client. This term can be used to differentiate between a Server based replica and a Notes Client based replica.

There are several other terms used throughout this document but these will be explained in context.

Chapter 2 : Elements of Replication

There are several elements that are used throughout Replication. Some of these need to be outlined to understand how replication functions overall.

Here we will be looking at :

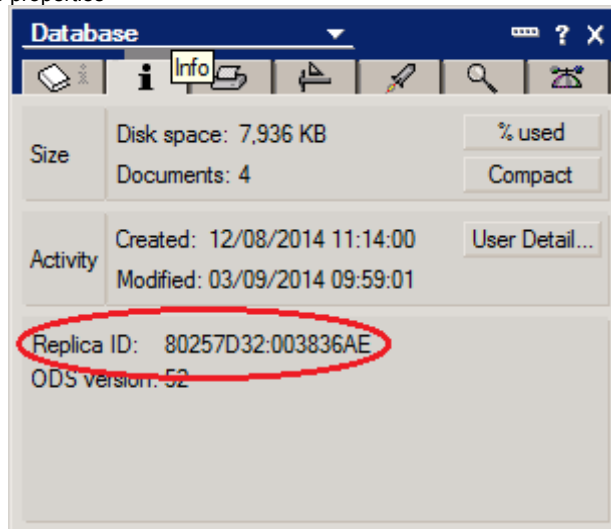
- 2.1 The Replica ID
- 2.2 Replication History
- 2.3 UNID and NoteID
- 2.4 Document Sequence Numbers
- 2.5 Connection Documents

2.1 The Replica ID

The Replica ID is a string assigned to a database that is used to identify the database for replication. This string needs to match between two databases for replication to be possible.

The Replica ID of a database can be seen in the Database properties on the 'info' tab :

Figure 2.1 - Replica ID in database properties



Note : Only the Replica ID needs to match between replicas for replication to identify them. The Filename, title and path for the databases do not need to be the same. This is important to understand as moving a database to a new path location on a server or renaming its filename will not prevent it from replicating.

2.2 Replication History

The Replication History is a log of the replication that has happened on a database.

It includes the Date and Time of the last replication, whether replication sent or received data, the servers involved in the replication and the filename of the database.

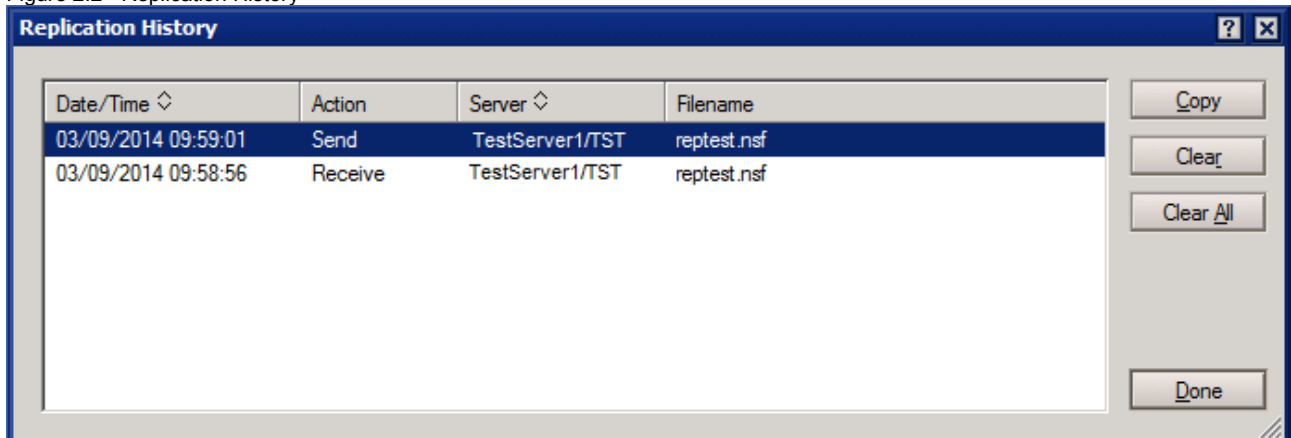
The Replication History is very important for replication. The Date and Time of the previous replication event is compared to the modified time for the database. Any documents that are listed as modified since the previous replication will then be listed to be replicated.

Using this process the Replicator task limits the amount of documents that are reviewed for replication and allows the Replication process to be more efficient reducing the amount of data needed to be processed and transmitted through a network.

It is possible to clear all the replication history data, this forces a 'Full Replication' where all documents are review and compared to their replica counterpart, this will take more time to complete, we will be looking at why this may be helpful in the Logging and Debug options and Troubleshooting chapters later on.

To view the Replication History select the database and click on the menu options File\Replication\History...

Figure 2.2 - Replication History



The Replication History can be cleared by clicking on “Clear All” and then Done.

2.3 UNID and NoteID

The UNID (Universal ID) is a string that is assigned to a note when the note is first created. The UNID uniquely identifies a note. UNIDs are used when replicating database notes and when replacing or refreshing database design notes.

The NoteID is a shorter string that is used to identify a document within a database. It is not used for Replication identification.

It is important to understand the difference between the two strings. The UNID must match for the document to replicate, if it doesn't match to a string on the target database then it will be considered to be a new document and will replicate creating a new document on the target replica.

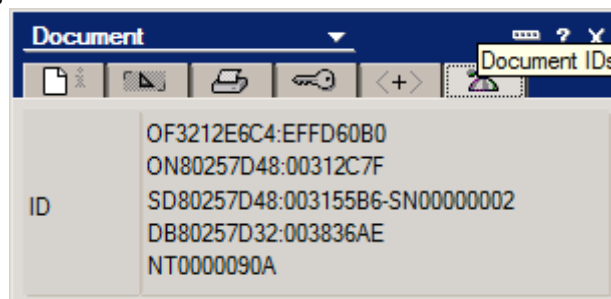
When a document is replicated the UNID will remain the same on the source and target replicas.

The NoteID does not need to match between replicas, although this can sometimes happen it is not required and should not be depended upon. A NoteID mis-match between replicas will not cause a new document to be created.

The NoteID is shorter and is often used in logging output - this can be useful as it can differentiate which replicating document has an issue as it can be different between the source and target replicas.

The UNID of a document can be seen in the “Document IDs” tab of the document properties :

Figure 2.3 - Document ID numbers



The first two lines make up the UNID.
Drop the OF and the ON – so the UNID is :

3212E6C4EFFD60B080257D4800312C7F

The NoteID is seen on the same tab of the Document properties – it is the last string on the bottom :

NT0000090A

Normally the NT and leading zeros are dropped to show the NoteID as 90A

Note : Some replication debug will refer to a document by its NoteID and not its UNID. This can be helpful as it can be used to differentiate which replica the debug output is referring to.

2.4 Document Sequence Numbers

This is a number logged on a document to track the number of updates to the document.

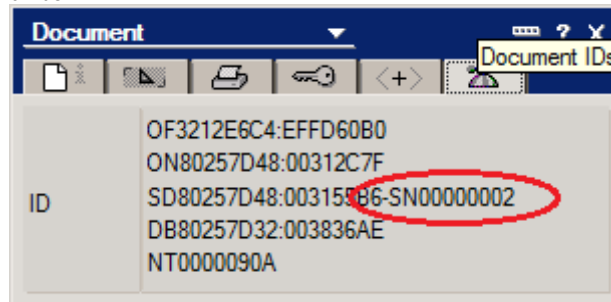
This is incremented whenever a document is updated.

Replication checks the Sequence number on a document with the sequence number of the document on the replica Database. If it is different from the last replication it triggers a replication of the document.

Replication will synchronize the sequence numbers between replicas – so if one replica has been updated several times and the other one not at all the sequence numbers will be matched.

As with the UNID and the NoteID the Sequence Number can be seen in the database properties on the “Document IDs” tab :

Figure 2.4 - Document Sequence Number



2.5 Connection Documents

In the Domino Directory, a Connection document enables communication between two servers and specifies how and when the information exchange occurs. In the Personal Address Book, it describes how a Notes Client accesses a certain server.

We are interested in the Domino Directory Connection documents as they can also be used to schedule and control Replication. Many of the features of Connection documents will not be covered here only those used for Replication.

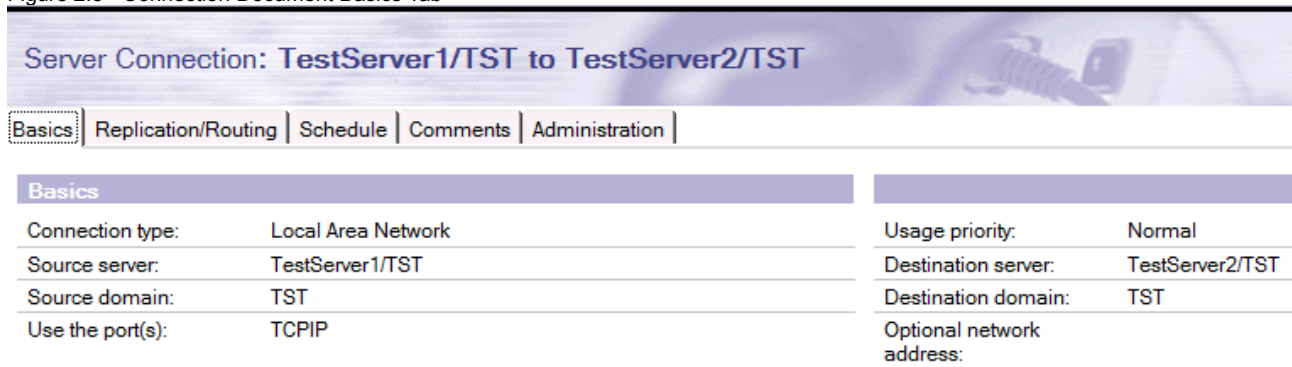
Connection Documents are not needed for replication to succeed between servers that are in the same Domino Domain but are used to automate replication tasks to a schedule. They are required for servers that are in different Domino Domains to allow traffic between the servers.

Connections Documents can be created and accessed using the Administrator Client from the Configuration tab. In the views select Server\Connections. From here all the existing Connection Documents will be displayed and the button 'Add Connection' will be available.

Connection Documents consist of several tabs :

Tab1 – Basics

Figure 2.5 - Connection Document Basics Tab



This section is used to define the source and destination servers and the type of connection between them.

The settings are :

Connection type :

Select the type of connection you want to establish, for example Local Area Network

Usage priority :

This is used to specify when to use this connection information in a search for the destination server path. Default is Normal.

Source server :

This is where you would enter the Source Server's hierarchical name.

Destination server :

This is where you would enter the Target Server's hierarchical name.

Source domain :

This is where you would enter the Source Server's Domain name.

Destination domain :

This is where you would enter the Target Server's Domain name.

Use the port(s) :

Enter the port to use for this connection.

Optional network address :

Enter the destination server's IP address. This is optional and is only necessary if you use TCP/IP or another protocol that needs an explicit network address to make the connection.

Tab 2 – Replication/Routing

Figure 2.6 Connection Document - Replication/Routing Tab

Basics		Replication/Routing	Schedule	Comments	Administration
Replication			Routing		
Replication task:	Enabled		Routing task:	Mail Routing	
Replicate databases of:	Low & Medium & High priority		Route at once if:	1 messages pending	
Replication type:	Pull Push		Routing cost:	1	
Files/Directory paths to replicate:	(all if none specified)		Router type:	Push Only	
Files/Directory paths to NOT replicate:					
Replication time limit:	minutes				
AutoDialer					
Use AutoDialer to connect remote server to network:	Disabled				
AutoDialer connection name:					

This section is used to Enable or disable replication for this connection document, control the type of replication enabled with this document and list which databases are to be replicated.

The settings are :

Replication task:

Enable or Disable replication for this connection document.

Note : A connection document itself can still be needed and active for other tasks, even if the replication is disabled.

Replicate databases of:

Each database can be assigned a priority of Low, Medium or High. In this way a connection document can be defined to replicate High priority databases more frequently than Low priority databases. The priority of a database is defined in the Replication Options which we will be reviewing in Chapter 5 Replication Configurations later on.

Replication type

Here you can select the type of replication that will be performed with this connection. The options are :

Pull Push (the default)

Pull Pull

Pull only

Push only

We will be looking at the differences between these in Chapter 4 Running Replication.

Files/Directory paths to replicate:

Here an administrator can define a specific folder of databases to be replicated, or individual databases or if blank all databases on a server. Using this option could limit replication to databases in the Mail folder, or individual databases. This can be used to only replicate the required databases, or to replicate certain databases at different frequencies.

Files/Directory paths to NOT replicate:

This option allows the administrator to define folders or databases to NOT replicate. If the “Files/Directory paths to replicate” setting includes multiple databases then this option would allow the definition of databases or directories to not be replicated. It can often be easier to define a small number of databases to not replicate as opposed to defining all databases to be replicated individually.

Note : Defining a database or folder to not replicate here only prevents replication for this connection document, the databases can still be replicated through other connection documents or manually.

Replication time limit:

Here the administrator can enter the number of minutes replication can take before it is stopped. The default is blank, allowing the server to complete all defined replication. If the value is defined and the replication takes longer than the defined limit it could lead to unreplicated databases and documents. This feature is

less used now compared to when dial up connections were more common and needed to be limited in length.

We will not be covering the AutoDialer and Routing sections here as they do not relate to replication more information on these can be found in the Domino Administrator Help.

Tab 3 - Schedule

Figure 2.7 - Connection Document - Schedule Tab

Scheduled Connection	
Schedule:	Disabled
Connect at times:	09:00 - 18:00 each day
Repeat interval of:	120 minutes
Days of week:	Sun, Mon, Tue, Wed, Thu, Fri, Sat

Here the time and frequency of the replication is defined.

Schedule:

Specify the connection as enabled (the default) or disabled. If it is disabled the Connection Document is then only used for making connections to other servers.

Connect at times:

Define the times during the day that the connection should be active. If the Repeat interval is 0 (for one replication only) then this will be the time the replication starts.

Repeat interval of:

Define the interval to repeat replication. Enter 0 for only one replication to occur.

Days of week:

Define the days that this replication should occur on.

Tab 4 & 5 – Comments and Administration

Figure 2.8 - Connection Document - Comments Tab

The Comments tab allows the entry of comments regarding the connection document.

Figure 2.9 - Connection Document - Administration Tab

Administration	
Owners:	[dropdown menu]
Administrators:	[dropdown menu]

Here you can define :

Owners:

The hierarchical name of the user who created this document can be added here. If defined only this user or an administrator can modify the document, provided they have sufficient access to the NAMES.NSF database.

Administrators:

Enter the hierarchical name of users who are required to edit this document. Groups can be used as well.

More information on Connection Documents can be found in the Domino Administrator Help.

Chapter 3 : Creating and Deleting Replicas

For Replication to work there must be a replica of a database on a remote location.

As mentioned previously the Replica ID must match between replicas for Replication to work.

The replicator task itself will not create a replica on the target location.

So, here we will be reviewing the methods that can be used to create a replica of a database and how to delete them :

- 3.1 Operating System copy
- 3.2 Using a Notes Client
- 3.3 AdminP
- 3.4 Cluster Management
- 3.5 When Registering a User
- 3.6 Deleting Replicas

For many of these options (3.2, 3.3, 3.4, 3.5) the ID being used to create the Replica must be included in the "Server Access/Create new replicas" section of the Security Tab of the Server Document. If the ID is not included then the creation of the new replica will fail. The ID can be included individually or within a group.

Figure 3.1 - Server Document - Security Tab - Server Access section :

Server Access	Who can -
Access server:	All users can access this server
Not access server:	
Create databases & templates:	LocalDomainAdmins LocalDomainServers
Create new replicas:	LocalDomainAdmins LocalDomainServers
Create master templates:	LocalDomainAdmins LocalDomainServers

3.1 Operating System copy

A copy of a database made using the Copy function of the operating system can be used as a replica of the original database. This database will contain all the same documents, settings and designs from the original.

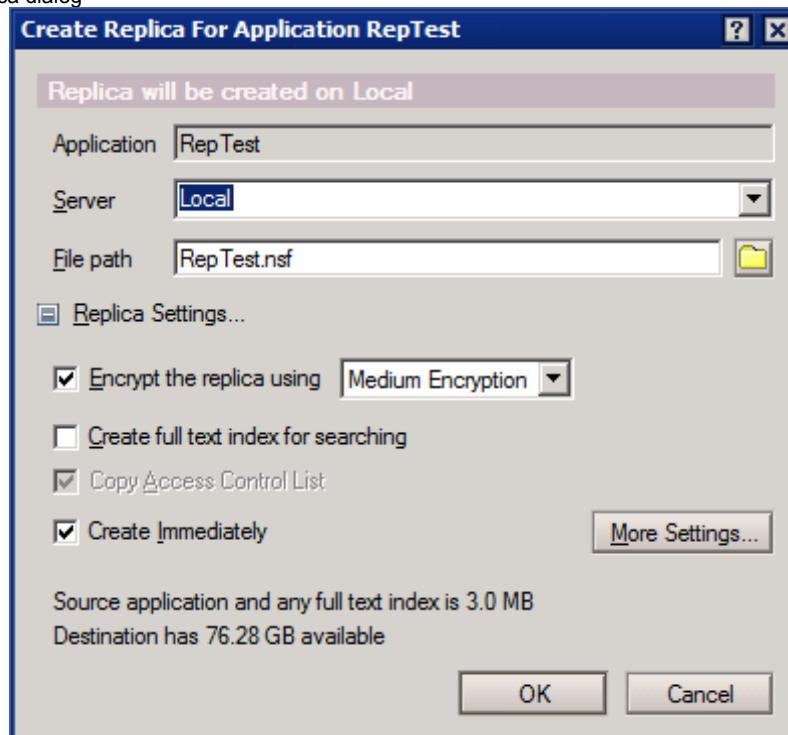
It is recommended that if the database is on a server that the server is shutdown before the copy is created, if the server is running the copy may be incomplete and corrupt.

This is a common method for creating a replica, it can often be the fastest option of creating multiple replicas on a target server.

3.2 Using a Notes Client

Select or open a database using a Notes Client. Then select File\Replication\New Replica...

Figure 3.2 - Create Replica dialog



Here you can define :

Server

The location the replica will be created on. This can be 'Local' to create a replica on the Notes Client itself or it can be a server to create the replica on a server.

File path

Define where the database will be created, on a server this will be relative to the Data folder, on a Client it can be anywhere on the hard drive. As mentioned above the replicas do not need to be stored in matching folders to replicate.

Note : The User ID being used by the Notes Client must be included in the "Create new replicas" field of the Security tab of the Server document.

Replica Settings :

Encryption : Define if the new replica will be encrypted. This will restrict access to the database if enabled. It will be enabled by default on all local replicas.

Create full text indexing for searching : Creates the Full Text Index for the database when the replica is created.

Copy Access control list : Replicates the ACL settings from the source database.

Create Immediately : Create the new replica immediately or using the background replication.

"More Settings..." button – this will open the Replication Options for this database. We will be looking at all these in Chapter 5 Replication Configurations.

Disadvantage of this process of creating a replica : In this option it is the Notes Client that does the replication. This means data is taken from the source server, through the Client and then to the target server, so the Notes Client can be a bottleneck. This can be inefficient compared to other methods.

Advantages : This can be a quick and easy step to create a replica on another server without having to have direct access to the server or to wait for the AdminP process (number 3 below) to complete.

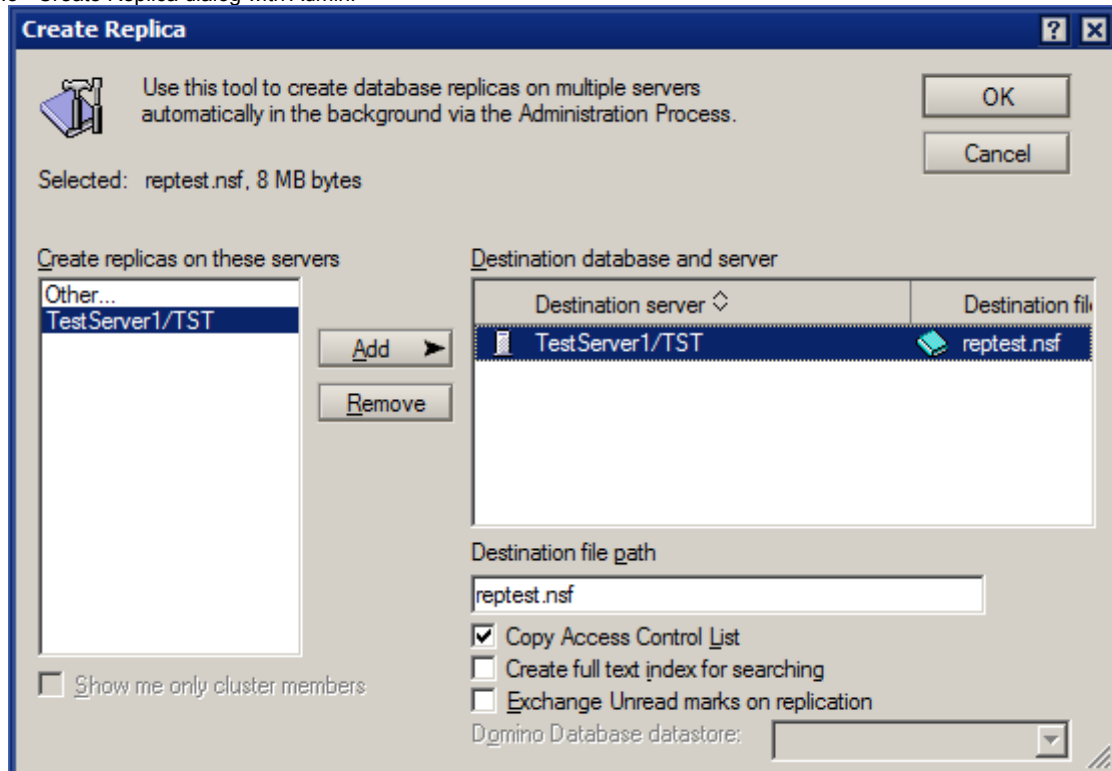
Note : It is possible to use the Notes Client to create a Notes Copy of a database. Using "File\Application\New Copy...". Although this can be used to create a copy of a database the resulting database will NOT be a replica. The Replica ID of the database will not match the source database and replication will not be possible between them.

3.3 Using AdminP

This option is activated using the Administrator Client.

- ◆ Open the Administrator Client.
- ◆ Select the Files tab.
- ◆ Select the database or databases that you wish to make a replica of.
- ◆ Select Tools\Database\Create Replica(s)...

Figure 3.3 - Create Replica dialog with AdminP



Either select the destination servers in 'Create replicas on these servers' and click on Add or click on "Other..." and then Add to select further servers that may not be listed.

Once selected you can choose to :

Copy Access Control List – on by default to replicate the ACL.

Create full text index for searching – deselected by default, if enabled a Full Text Index will be created on for the replica on the target server

Exchange Unread marks on replication – off by default and also requires a Notes.INI setting (ADMINP_EXCHANGE_ALL_UNREAD_MARKS = 1).

This will initiate an Administration Process task that will create a Replication Stub on the target server. This is a replication place holder – basically a database with no design elements or documents. If the source and target servers are clustered then the Replication Stub will be populated automatically, if the source and target are not clustered then a manual replication will be needed to populate it.

The advantages to this process are that the Server will do all the replication – not the Client, also multiple databases can be selected and replicated in one process.

A disadvantage is that the Administration Process may take longer then expected to complete the creation of the replicas.

Note : The Administration server must be included in the "Create new replicas" field of the Security tab of the Server document for this process to complete.

For a successful replica creation several AdminP steps are generated these will need to replicate to and from the target server to complete the process :

Check Access for New Replica Creation

Accelerated Create Replica

3.4 Cluster Management

This is a Notes.ini setting that allows some cluster commands to be used at the Domino console. Regardless if the server is part of a cluster. There is no need to Restart the server after this command has been issued using "set config", and it becomes active immediately.

This enables this command with a few options:

CL COPY sourcedb targetdb [Options] Valid values of the Options are Empty -> aka no option used -> Creates a COPY of the sourcedatabase

REPLICA -> Creates a Full Replica of the Sourcedatabase

TEMPLATE -> Create a Template (as in design only) of the sourcedatabase

Examples for the CL COPY command:

CL COPY sourcedb targetdb

Examples for the CL COPY command:

> CL copy serverA!!db1.nsf serverB!!db2.nsf

This would create a regular non-replica copy of db1.nsf on serverA as db2.nsf on ServerB

> CL copy serverA!!db1.nsf serverB!!db2.nsf REPLICA

This would create a replica copy of db1.nsf on serverA as db2.nsf on ServerB

> CL copy serverA!!db1.nsf serverB!!db2.nsf TEMPLATE

This would create a template copy (only design, no data) of db1.nsf on serverA as db2.nsf on ServerB

> CL copy db1.nsf db2.nsf

This would create a regular non-replica copy of db1.nsf as db2.nsf on the same local server

Note : When creating new databases using the Cluster Manager the ID file of the source server is used not the ID of the user entering the requests. Because of this the Server ID in use must be included in the "Server Access/Create new replicas" section of the Security Tab of the Server Document.

3.5 When Registering a User

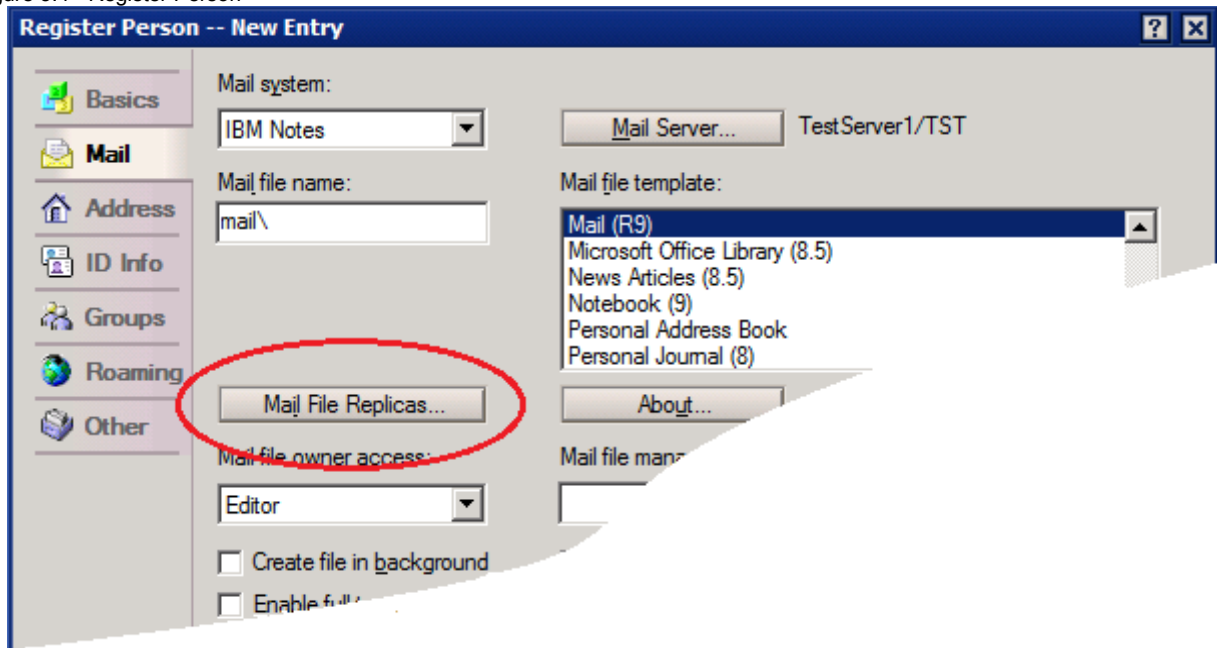
When registering a new user on a server a Mail file can be generated. If a replica of this database is required on other servers this can be initiated from the Register Person dialog box.

Open the Register Person dialog box.

Check the 'Advanced' checkbox on the lower left to enable the Advanced tabs.

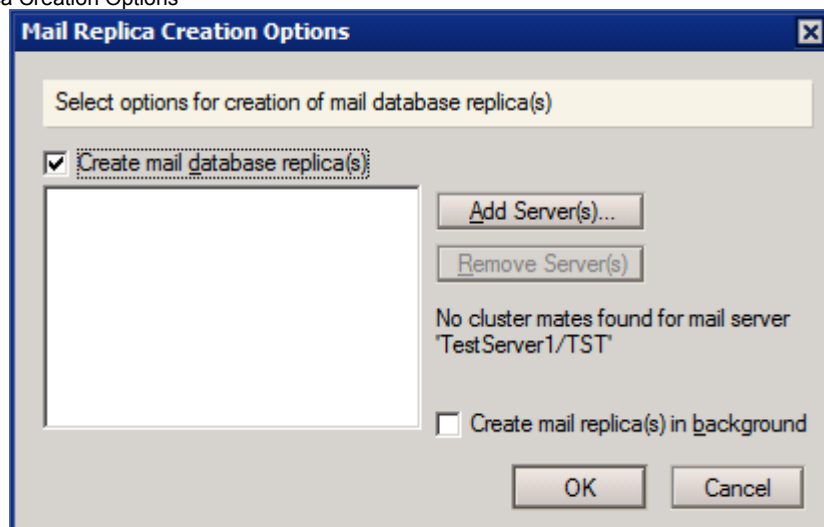
Open the 'Mail' tab :

Figure 3.4 - Register Person



Click on the button "Mail File Replicas" to open the "Mail Replica Creation Options" dialog :

Figure 3.5 - Mail Replica Creation Options



Here after checking "Create mail database replica(s)" you can enter the server that you wish to have replicas of the new mail database.

This process uses the AdminP process described above in "3.3 Using AdminP" to create the new replicas.

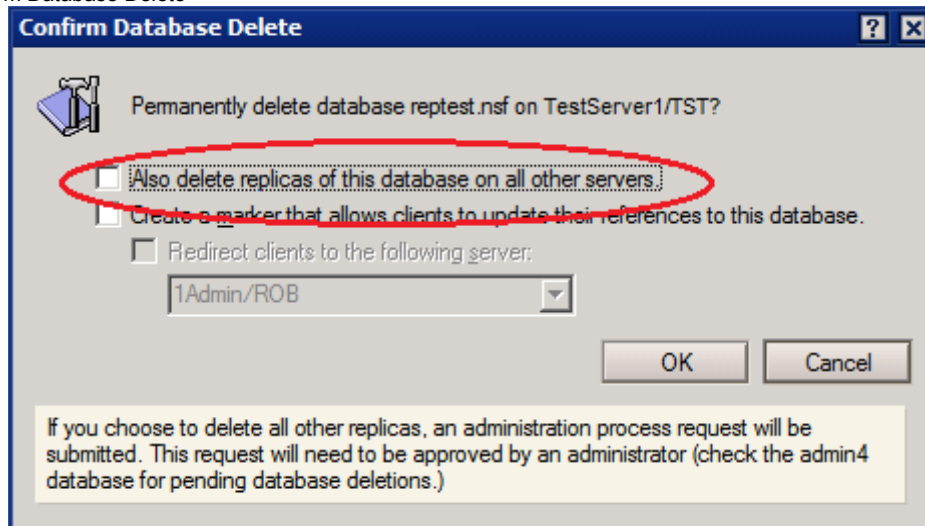
3.5 Deleting Replicas

Along with creating replicas it is often required to remove replicas from a Domino Server or from a Notes Client - individually this can be done as you would with any other database. However, removing all the replicas from a database off a entire Domain can be troublesome and a replica could be missed.

When using the Administrator Client in the files tab it is possible to Delete a database from a server with the option Tools\Database\Delete...

To open the Confirm Database Delete dialog :

Figure 3.6 - Confirm Database Delete



This initiates several Admin Process tasks that will replicate throughout the Domain :

Get Replica Information for Deletion - this will check each server for a replica of the database, on servers where a replica is found an "Approve Replica Deletion" request is generated.

The "Approve Replica Deletion" must be approved by an administrator with Manager access to the database and then the replicas will be removed.

One "Approve Replica Deletion" request will be generated by each server that hosts a replica of the database and each one will need to be approved for each server.

Chapter 4 : Running Replication

Once a replica of a database has been created you need to initiate the replication between the source and the target. This can be manually or on a schedule. Unless included in an existing Connection Document or a cluster simply creating a replica will not cause the databases to replicate further changes.

There are two types of replication that are used, all commands and options will use a combination of these.

The two types are Pull and Push.

Pull Replication is where changes from the target server are replicated to the source server. So, all changes are 'pulled' from the target server to the source server. With this option, on its own, changes from the source server are NOT replicated to the target server.

Push Replication is where changes from the source server are replicated to the target server. So, all changes are 'pushed' from the source server to the target server. With this option, on its own, changes from the target server are NOT replicated to the source server.

The default option in a Connection document is Pull Push – where changes are pulled from the target server to the source server and then from the source server to the target server.

Server to server

Manually.

Replication can be initiated on a Domino Server Console. The commands for this are :

Rep servername databasename (or wildcard)

Pull servername databasename (or wildcard)

Push servername databasename (or wildcard)

The 'Rep' command will initiate a Pull-Push replication :

Purpose : Performs replication of databases with other Domino servers.

For SCHEDULED replication: Do not specify any optional parameters.

For ONE TIME ONLY replication: You can specify optional parameters for replication, then the currently loaded replicator task will shut down.

Usage : Load REPLICA [[[server] database] mode]...

where :

[server] Name of destination server to replicate with.

[database] Name of database or folder to replicate with destination server.

[mode]

-p Perform pull-only replication.

-s Perform push-only replication.

(Default is 2-way replication)

Scheduled.

To allow replication to run between two servers on a schedule you will need to create a Connection Document and Enable replication on it. The schedule can then be defined in the connection document. Details on Connection Documents can be see above in Chapter2 Elements of Replication.

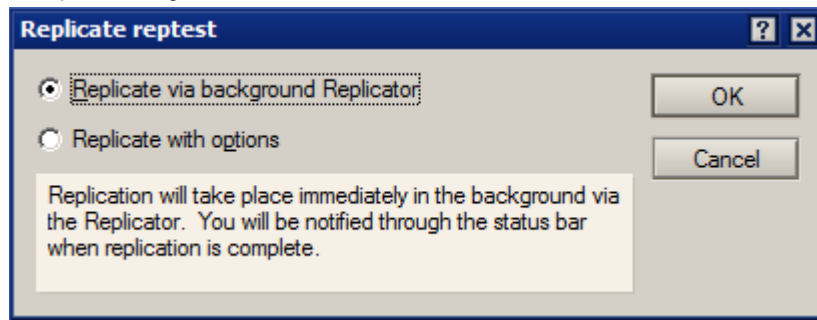
Client to Server

Manually

Once the replica is created on a Notes Client it can be replicated. Select the local replica and choose File\Replication\Replicate...

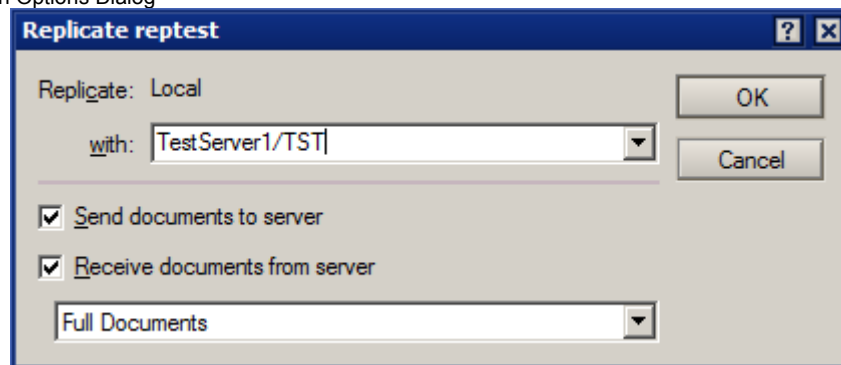
This gives you the option to “Replicate via background Replicator” or “Replicate with options” :

Figure 4.1 - Local Replicate Options dialog



Replicate with options gives you :

Figure 4.2 - Replication Options Dialog



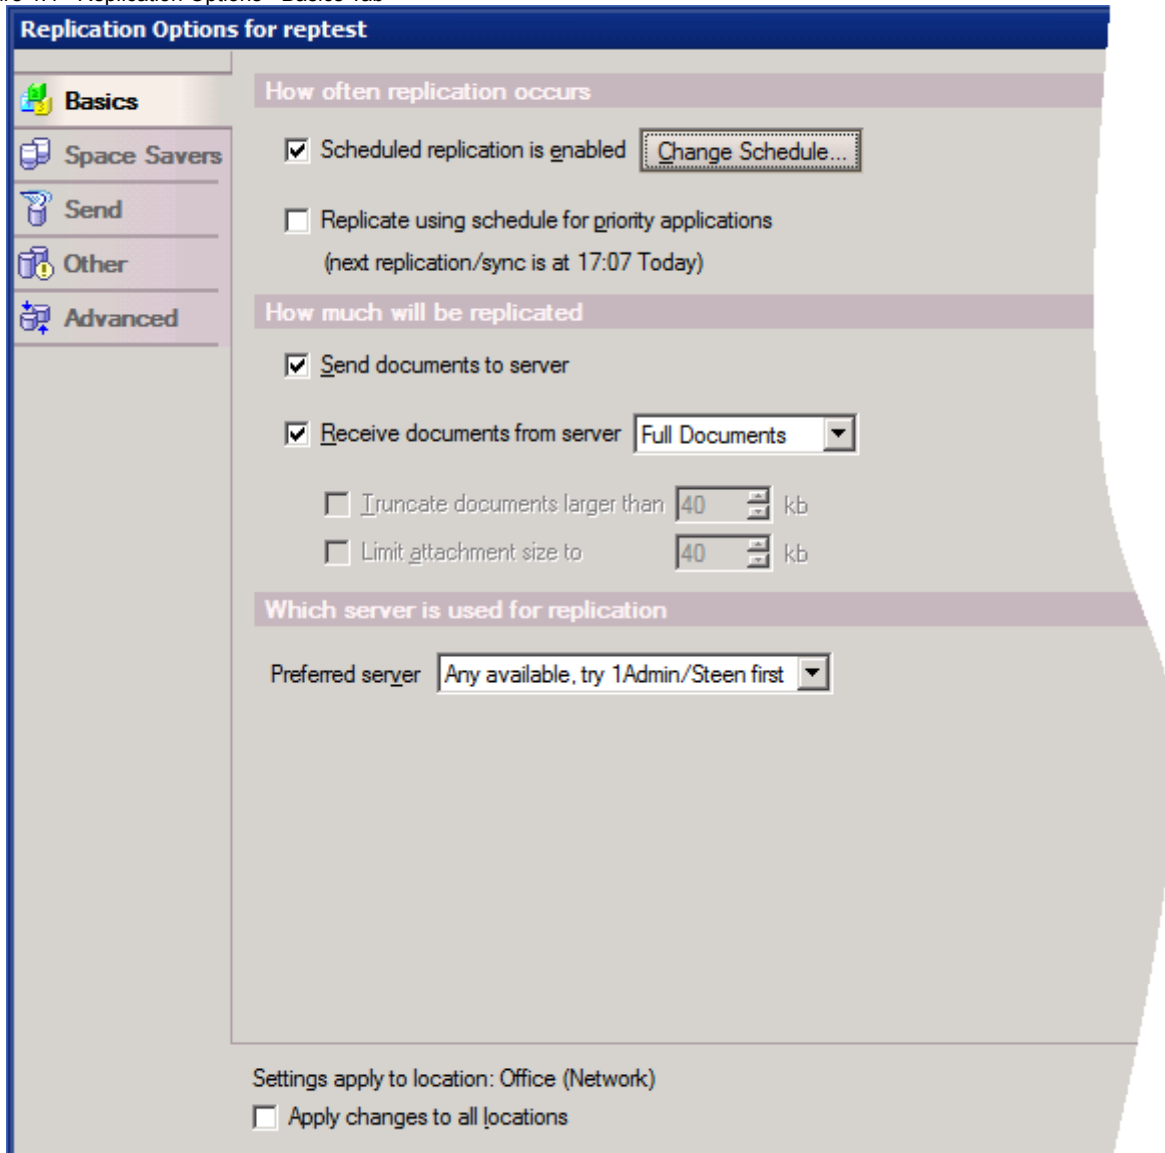
Allowing you to Pull, Push or Pull and Push to/from the server.

You can also select to replicate Full Documents, Partial documents, Summary only, Smallest first. More on these options when we go through the Replication Options later.

Scheduled

Once a local replica of a database is created the Basics tab of the Replication Options is activated, from here you can define the schedule for the replication. To access this click on File\Replication\Options for this replication :

Figure 4.4 - Replication Options - Basics Tab



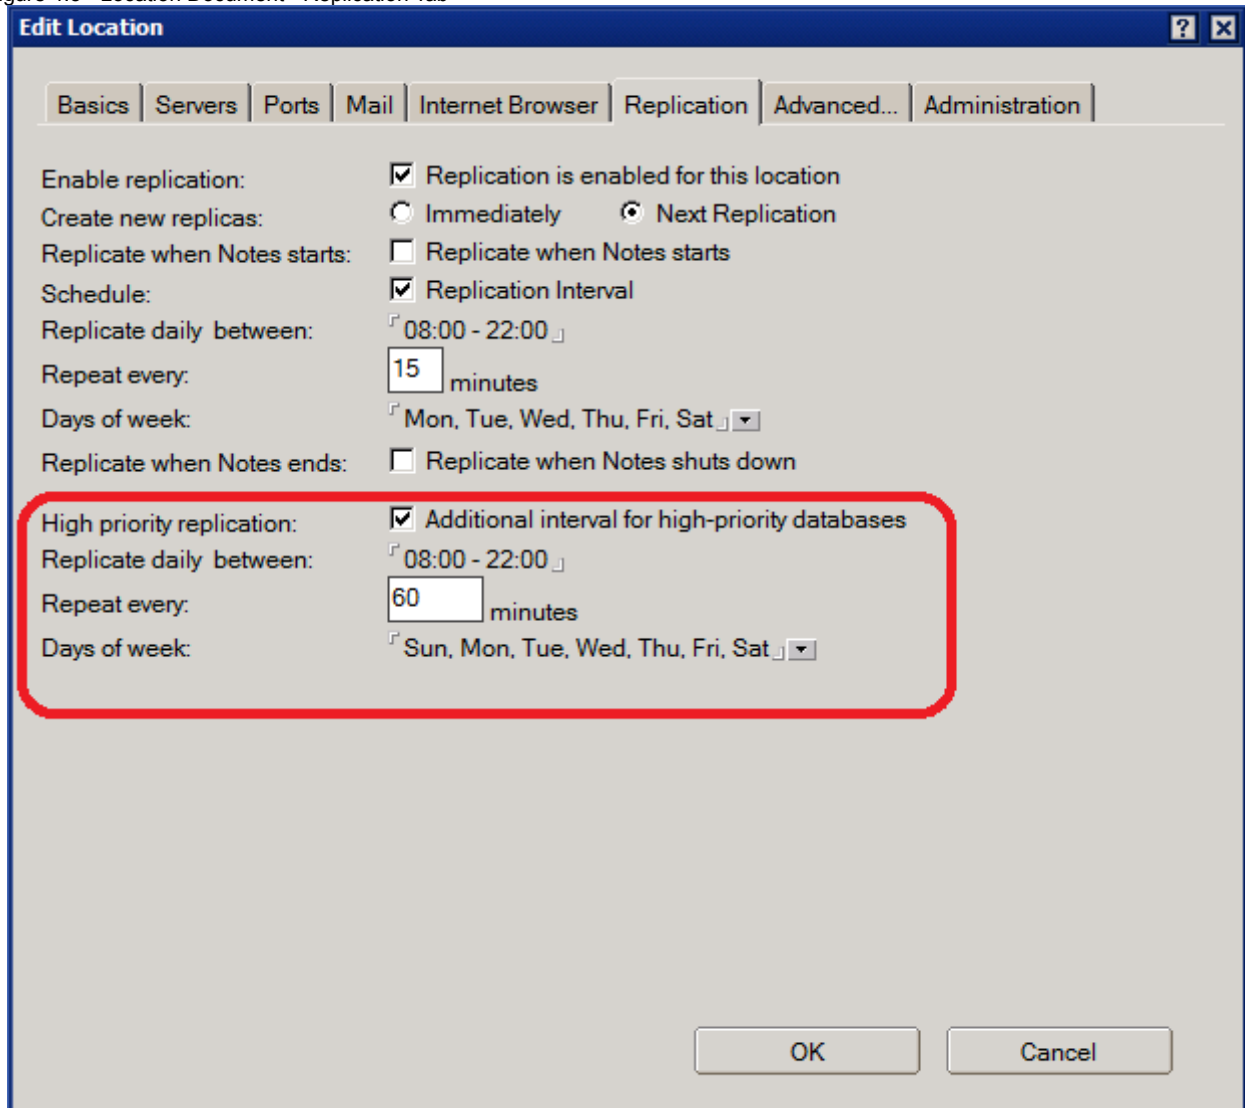
The settings here are :

Scheduled replication is enabled – once checked this activates the “Change Schedule” button where you can define the replication schedule, by interval time, repeat interval and day. This can be defined separately for normal or high priority documents.

Other schedule options are to “Replicate/Sync when I start the client” and “Replicate/Sync when I shut down the client”

Replicate using schedule for priority applications – Using the internal schedule for high priority databases. This is defined in the Replication tab of the Location document :

Figure 4.5 - Location Document - Replication Tab



Send documents to server – sets that changes to documents made locally are sent to the server.

Receive documents from server – sets that changes to documents made on the server are sent to the local replica with the options :

Full documents – complete documents are replicated.

Partial documents – documents are truncated or attachments limited to a size.

Summary only – only a summary of information about the document is replicated.

Smallest first – the order of replication is set so the smallest documents replicate first.

Preferred server – allows the selection of which server is replicated with first, in the case where there are multiple replicas available it may be preferred to be able to select one over another.

Settings apply to location: Notes Clients use 'Location' documents, these contain configurations that control how the Notes Client connects to servers, how replication is scheduled, what mail file is being used, etc. In this case the replication schedule that is being configured is for the listed location document. By checking "Apply changes to all locations" then all location documents on the Notes Client will be updated with the new replication options.

Note : All replication with a local database (a database not on a Server) is performed by the Notes Client. This cannot be initiated from the Domino Server side.

Overview of how replication works :

Replication will review the source and target database.

It then creates a list of which documents to check for changes (based on what has changed since the last complete replication history entry or if blank on all documents).

It will then use this list to compare documents to be replicated with the documents on the target server.

If two way replication is being implemented then this will be done individually for the Pull and the Push replication.

It then attempts to replicate the changes between the database - reviewing access levels at the server, database and document levels depending on the Certificates, ACLs and Reader fields.

Chapter 5 : Replication Configurations

In this chapter we will be looking at the different configuration options available while running replication. Not all of these are required for all replication situations, but it can be useful to understand each option and how it could be applied.

Replication options being covered here :

5.1 Replication Settings

To include :

- All tabs from the Replication Options dialog
- PIRC
- Selective replication
- Purge Interval

5.2 Replication and the ACL

5.3 Managed replicas through a policy

5.4 Local Replication settings through a policy

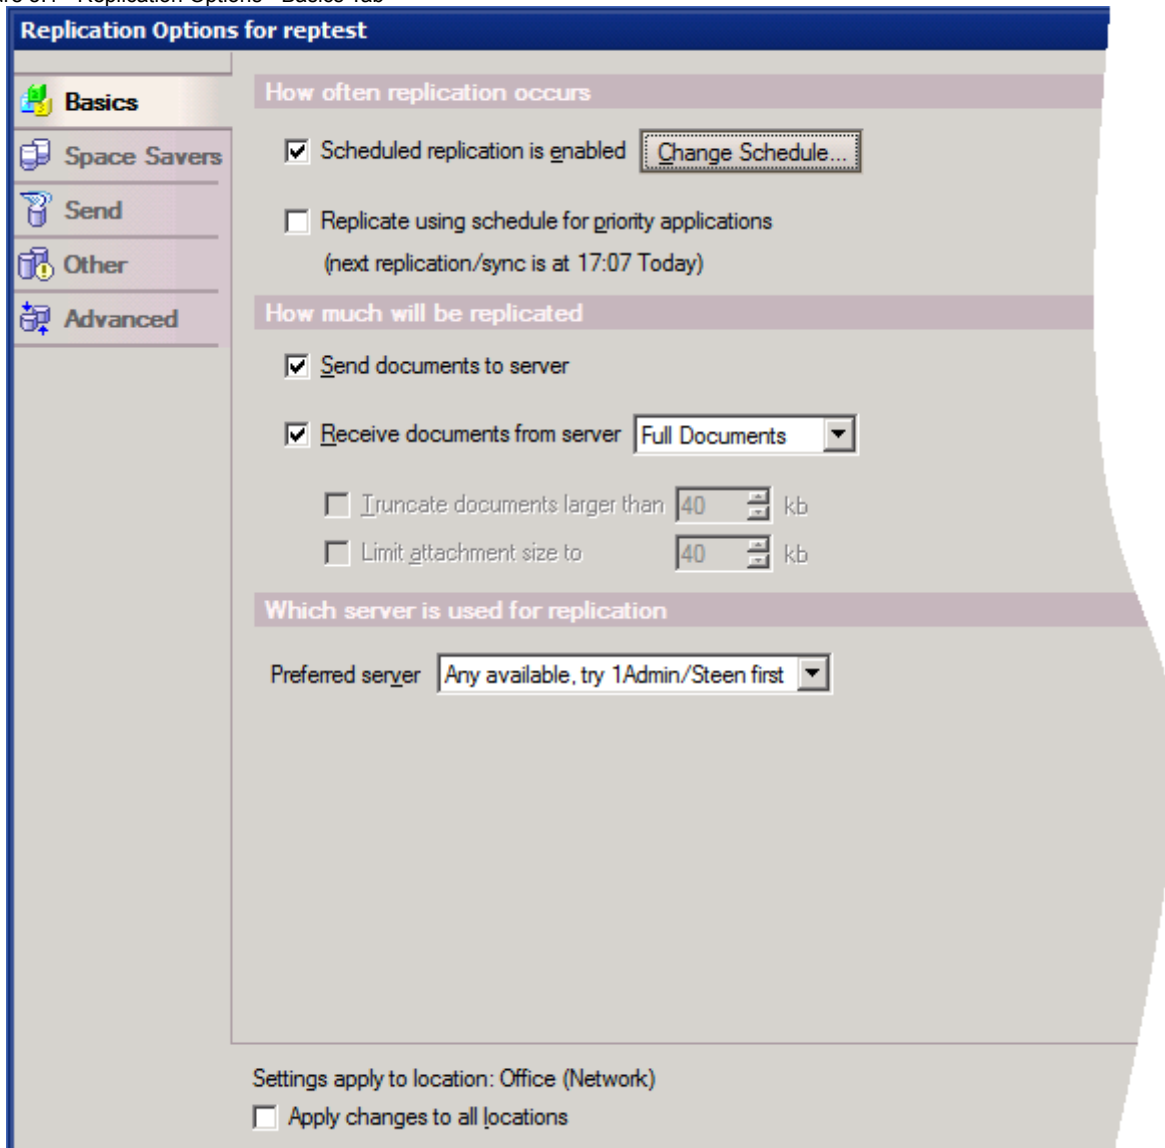
5.5 Replication settings on a Notes Client

5.1 Replication Options

The Replication options for a database can be accessed by selecting or opening the database with a Notes Client. Then clicking on File\Replication\Options for this application...

This will open the Replication Options dialog box :

Figure 5.1 - Replication Options - Basics Tab



This consists of 5 tabs covering different areas of replication.

Basics Tab

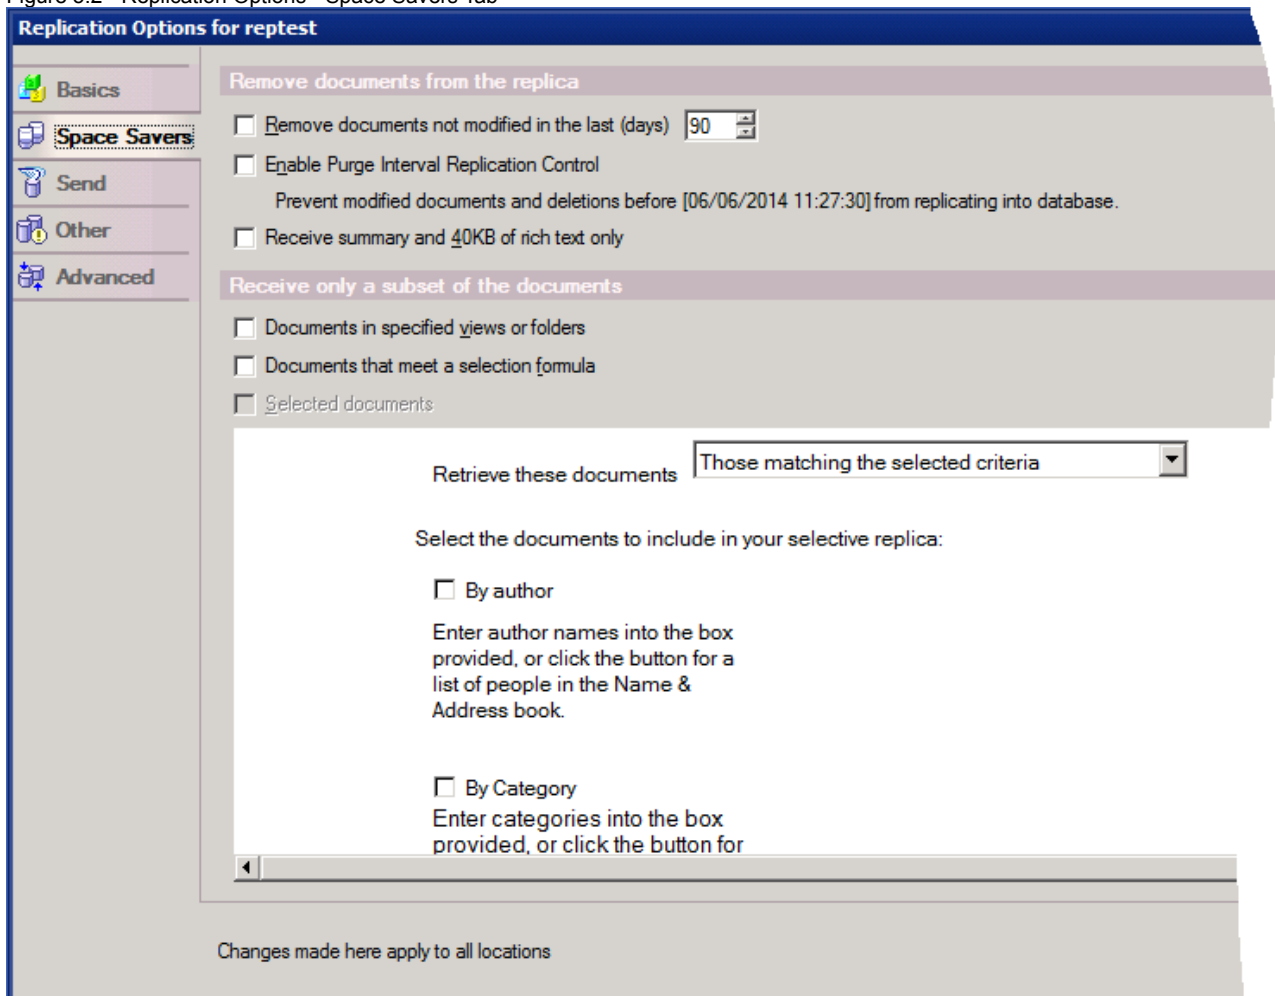
The Basics tab is the first tab displayed when the Replication Options dialog is opened (seen above). It is used to configure and schedule replication between a Client based replica and a server. If there is no Client based or local replica available then the options on this tab will be greyed out.

The options on this tab are covered in the Client to Server section of Chapter 4 – Running Replication.

Space Savers Tab

The Space Savers tab contains many of the useful options for replication and is also the source of a lot of confusion. We will be going through each options individually :

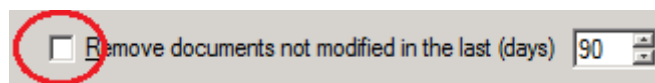
Figure 5.2 - Replication Options - Space Savers Tab



Remove documents not modified in the last (days)

This options controls 2 different areas depending on how it is configured.

When Unchecked



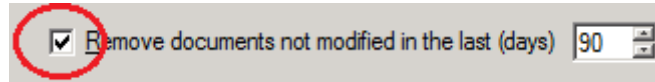
The days setting on the right had side will define when deletion stubs (there will be more on these later in the Chapter - Other areas of replication) are purged or removed from the database. Briefly, Deletion Stubs are a truncated document that is left in a database in place of the original document when it is deleted to indicate to the Replication task that the document should, in fact, be deleted from all other replicas.

If deletion stubs are not purged from a database they can build up in the database.

By default they will be purged when they are older than 90 days. Domino will check for old deletion stubs to be purged every one third of the Purge Interval setting. So, by default this will be every 30 days.

This will happen even if this option is unchecked.

When checked



If the option is checked then when Domino removes deletion stubs it will also remove documents that haven't changed within the specified number of days. These documents are purged, meaning no deletion stubs remain for the documents, so the documents aren't deleted in other replicas.

Note : If you select the check box on a non-replicated database, documents are lost and you can only recover them from a system backup.

Enable Purge Interval Replication Control

This option enables PIRC - a relatively new feature added to replication.

Simply put it can be configured to prevent the replication of old documents into a database. The most common cases occur when an old version of a NAB (Domino Directory or NAMES.NSF) is brought back online after several months or years either to recover or review some data. When this occurs any previously deleted person documents, group documents, server documents and other design elements can replicate back into the environment because the deletion stubs have already been purged from the current working replica. PIRC prevents these documents from replicating back into the environment. It can be enabled on any database and not just the NAMES.NSF.

How PIRC Works

PIRC prevents documents from replicating back in two ways :

The first is by modifying the 'since time' used by the replicator. If the server that is initiating the replication is release 8.5.3 or higher and therefore PIRC-aware, then the algorithm to determine which documents qualify for replication has been changed to use the later of the PIRC date or the Since time contained in the replication history. This prevents the older documents from being replicated. A DDM event is generated showing the number of documents skipped by PIRC allowing the administrator to monitor this option.

The second way in which PIRC prevents older documents from replicating in is by stopping the documents from being added to the database via NoteUpdate. This method occurs when a non-PIRC aware server is pushing older documents to the PIRC-aware server. In this case, NoteUpdate code will compare the SequenceTime of the document to the PIRC date. If the document is older, a DDM event will be logged and the document will be rejected for replication. Note that the source server is unaware that the document was rejected. Therefore its logs will show that all documents were successfully replicated. However, the PIRC server will log the server name and the documents that were rejected, thus allowing the administrator to track down old replicas of the database and remove them from the environment.

PIRC keys off the Cutoff date of a database. The cutoff date is the date when a purge last ran. Purge is the process which removes deletion stubs and, if selected, old documents from a database. Purge runs at one-third of the purge interval. Therefore PIRC will use this date to prevent documents older than this date from replicating into the database. If the cutoff date is not set in a database, which happens when replication history is cleared, the current date minus the purge interval is used.

In this way no documents older than the current documents or recently deleted documents can replicate into the database, however existing documents will remain on the database.

For example :

A user's mail file has document 10, 20, 30, 40, 50, 60, 70, 80, 90 and 100 days old along with some older documents that are needed on the database.

There is a backup of the mail file available with documents 80, 90, 100, 110, 120, 130, 140, 150, 160 and 170 days old.

The original database is running on a server with a Purge interval of 90 days and PIRC is NOT enabled.

A user accesses the backup through a Notes Client to review documents that are no longer available in the server based database.

The backup is accidentally replicated with the active mail file causing the older documents to replicate to the server based database. The server based replica now has old, previously deleted documents, in it along with all it's correct documents.

With PIRC enabled all documents that are older then the Purge Interval are excluded from replication - that is all documents over 90 days will not replicate to the server based replica and all documents older then the purge interval will remain on the server based replica.

There are three ways to enable PIRC on a database:

1) As above in the Replication options.

2) You can implement PIRC using compact. Compact '-PIRC On' and '-PIRC Off' will toggle the option. This can be used to enable or disable PIRC on multiple databases with one command.

For example :

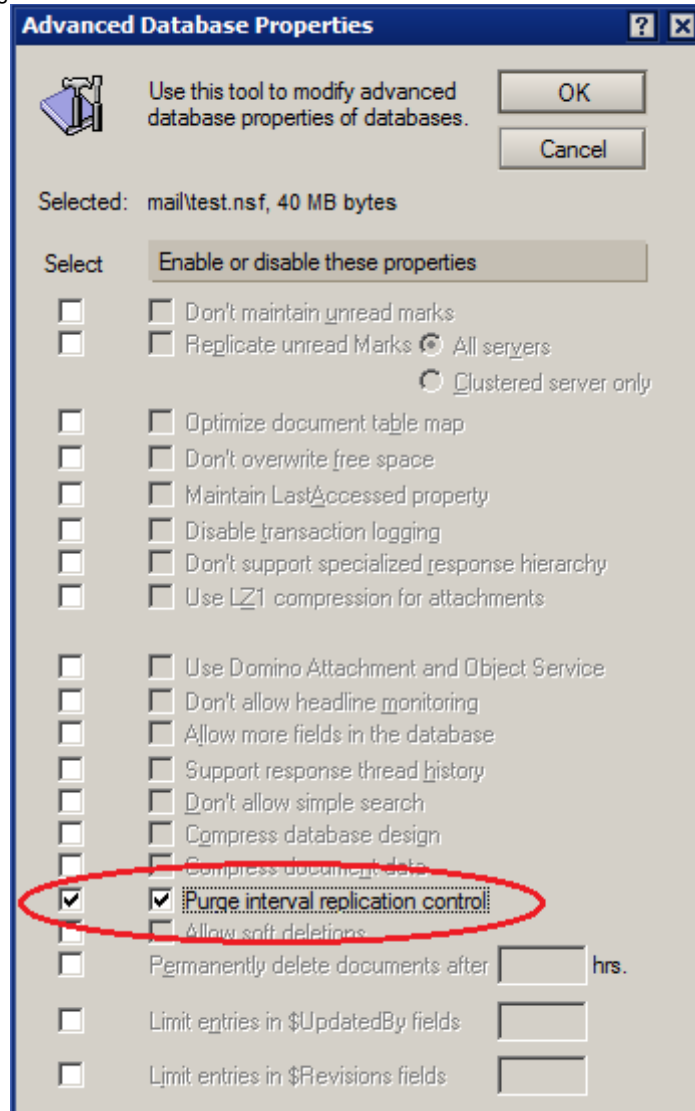
```
> lo compact mail\test.nsf -PIRC on
[0E10:0004-1050] 20/10/2014 15:30:00    Informational, PIRC has been
enabled for database mail\test.nsf.
```

```
[0E10:0002-1218] 20/10/2014 15:30:02    Database compactor process shutdown
> lo compact mail\test.nsf -PIRC off
[2DFC:0004-243C] 20/10/2014 15:30:31    Informational, PIRC has been
disabled for database mail test.nsf.
[2DFC:0002-24BC] 20/10/2014 15:30:33    Database compactor process shutdown
```

Note : This option is not case sensitive

3) You can set PIRC through the database options settings in the Administration Client. This also allows an administrator to set PIRC on multiple databases simultaneously.

Figure 5.3 - Enabling PIRC in the Administration Client



Considerations before enabling PIRC.

Before enabling PIRC you must make sure that all the documents you expect to replicate to the database have replicated. Therefore you must take care that the database has replicated with all the replicas in the environment prior to enabling it. This is especially important if some of the replicas are partial.

With PIRC enabled documents older than the Purge interval will not replicate into the database unless PIRC has been overridden.

Overriding PIRC.

A replication option '-NOPIRC' has been added to the replication task to temporarily override the PIRC setting without having to disable it. This is useful if there is ever the need to synchronize two replicas of a database but you do not want to expose the database to other replicas in the environment. This includes the server in the case of a PUSH. Also the NOPIRC switch overrides the replication history by setting the since time to zero. This way when the NOPIRC switch is used, a full replication is done between the two databases bringing them into sync without exposing either databases to other replicas in the environment.

for example :

```
rep testsrv1 test.nsf -nopirc
```

Design/Replace Considerations.

With PIRC enabled on a database special care is needed in order to push out design changes

because templates tend to contain design elements that do not change over a long period of time. Once these design elements are older than the Purge Interval PIRC would block them from replicating. To avoid this problem the Design Modification TimeDate was added to the database. This time/date is updated in the database whenever design refresh or replace is run, and it actually modifies one or more design elements. PIRC compares this time to the PIRC time and if it is newer PIRC will allow all non-data notes to replicate into the destination database. This will allow a new design or newly modified design elements to replicate out through the environment even if PIRC is enabled on the databases.

Receive summary and 40KB of rich text only

By selecting this setting Domino prevents large attachments from replicating and shortens the documents that the replica receives. The shortened documents contain only a document summary that includes basic information such as the author and subject, and the first 40K of rich text.

When users open a shortened document, they see (TRUNCATED) in the document title. To view the entire document, users open it and choose Actions -> Retrieve Entire Document.

Note :

Users can not categorize or edit shortened documents.

Shortened documents do not replicate unless the destination replica also has this option selected.

This is to prevent the complete original document being altered itself through replication.

Documents in specified views or folder

Once selected a list of the Views and folders from the database will be displayed. One or more of these can then be selected. Only the documents listed in these views and folders will be replicated into the database.

Documents that meet a selection formula

This option allows the use of a LotusScript SELECT command to select which documents are replicated. Variations of author, category, creation date, etc. can be used to limit the replication of documents that match the formula.

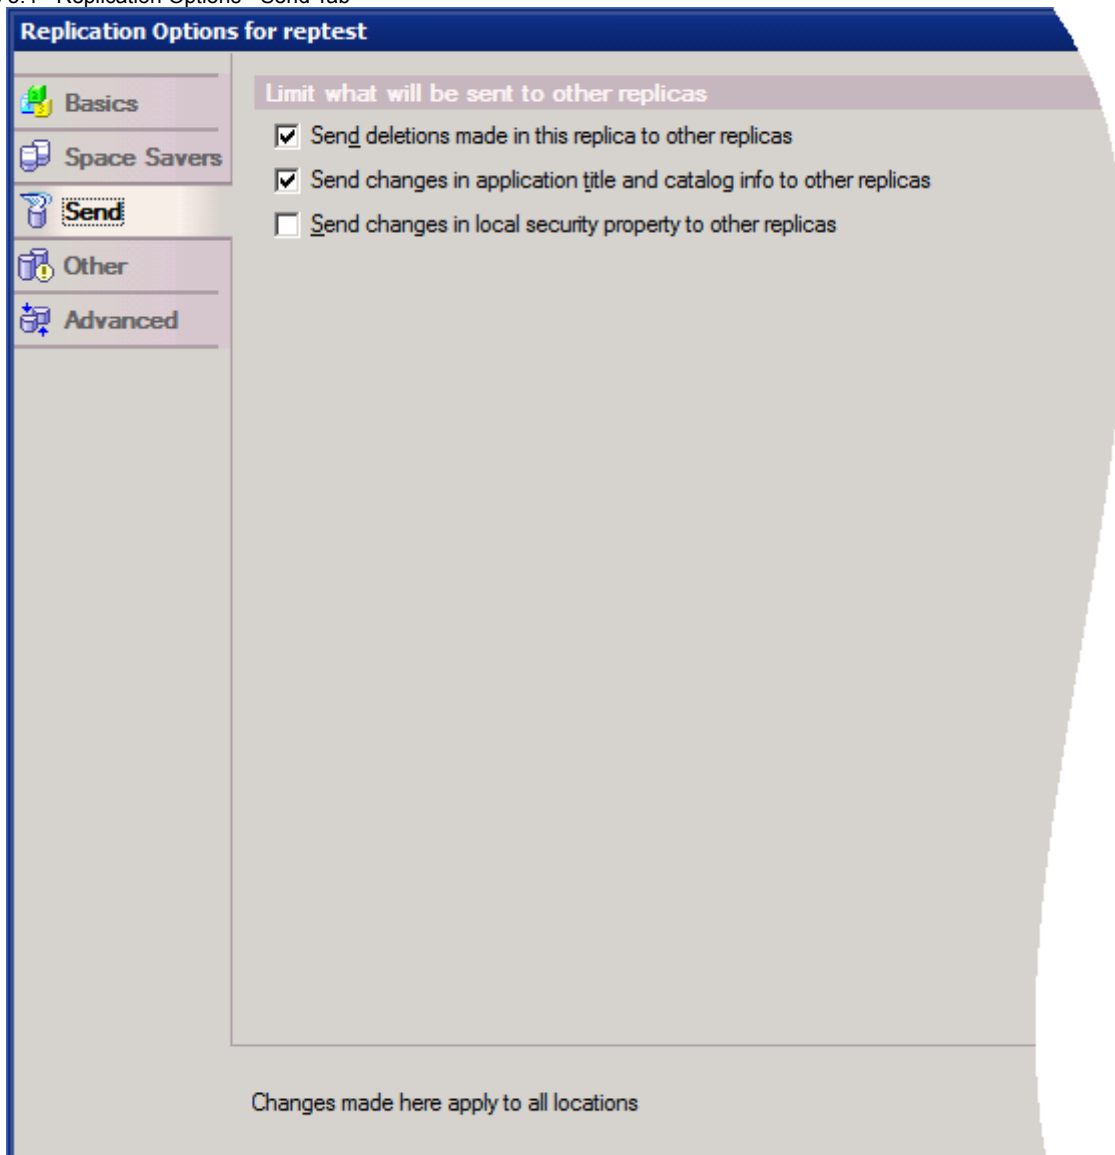
The default option is :

```
SELECT @All
```

This will allow the replication of all documents into the database.

Send Tab

Figure 5.4 - Replication Options - Send Tab



The send tab controls three options for what documents replicate out from this database.

Send deletions made in the replica to other replicas

This option controls the replication of deletion stubs out from this database. By default it is enabled allowing the replication of the deletion stubs. If disabled it prevents deletion stubs replicating out from this database.

Send changes in application title and catalog info to other replicas

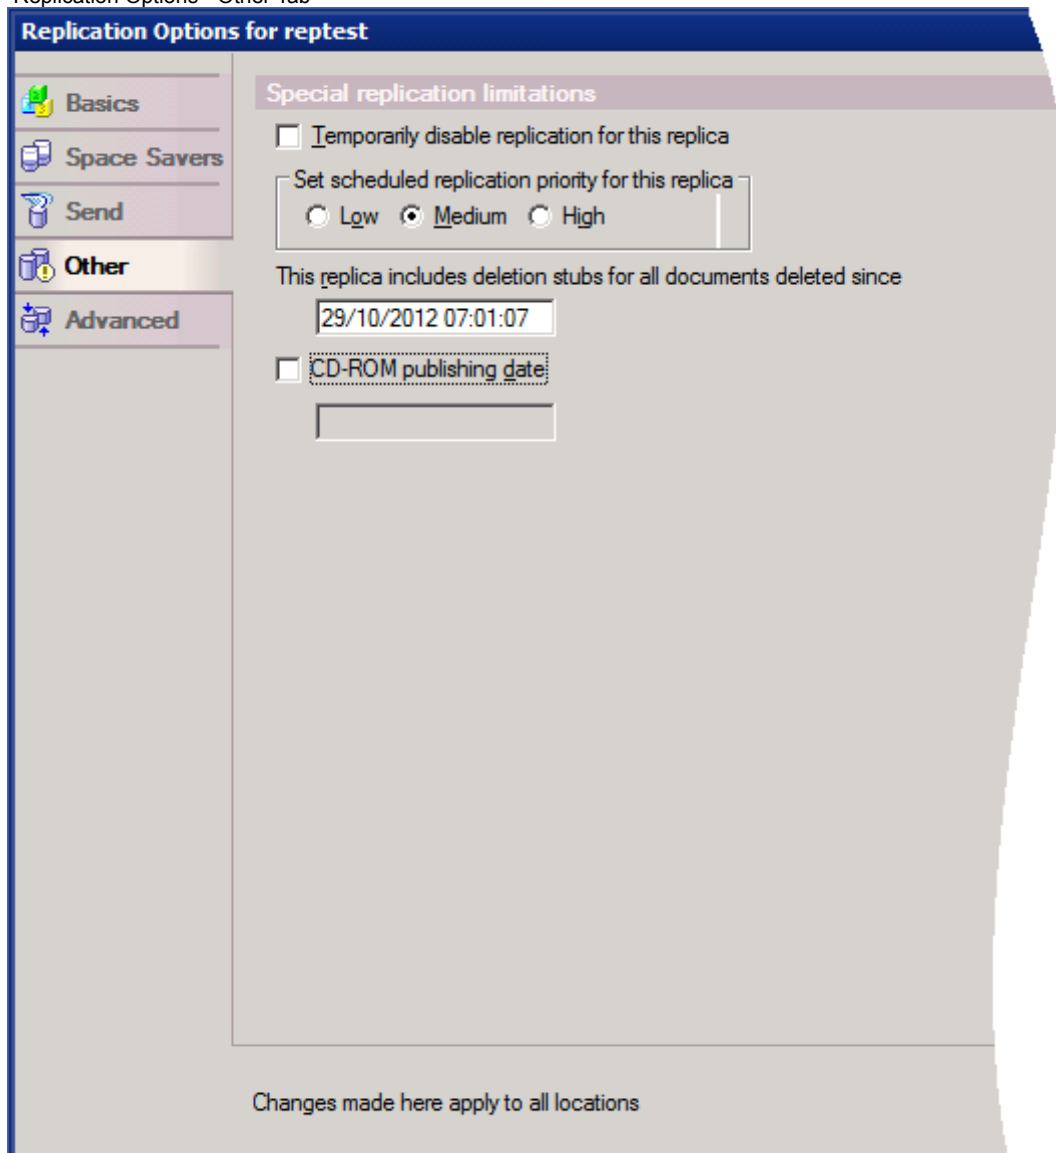
Selected by default this option enables the replication of the database title (not the filename) to other replicas and also catalog information about the database. If unchecked then this information will not replicate. As replication is dependent on the Replica ID and not the filename or database title if the database title differs between replicas then the replication will still continue.

Send changes in local security property to other replicas

Deselected by default, this setting, when deselected, prevents changes to the database Encryption property (set by choosing Encryption on the Basics tab of the Database Properties box). This is used mainly to prevent changes made to this property on a local replica from replicating to a server. For example, if this setting is selected and you disable the Encryption property on a local replica, the property remains selected on a server replica.

Other Tab

Figure 5.5 - Replication Options - Other Tab



This tab contains options that do fit in with the other tabs.

Temporarily disable replication for this replica

This option will disable all replication for this replica. Documents will not replicate to or from the database.

When enabled this message will be displayed if replication is attempted :

```
[2130:0002-2124] 04/09/2014 13:45:31 Replication is disabled for reptest.nsf
```

Set scheduled replication priority for this replica

As explained above (Chapter 2 : Elements of Replication - Section 2.5) with Connection Documents it is possible to configure separate schedules for replicas of different priorities. This is the setting that allows the configuration of that priority for a replica. By default it is set to Medium.

This replica includes deletion stubs for all documents deleted since

This field is automatically populated.

The cutoff date, which a replica uses to purge document deletion stubs deleted prior to the date and, optionally, unmodified documents

Which documents are scanned during the first replication after clearing the replication history

Note: By default, this date is periodically updated based on the number set for Remove documents not modified in the last x days.

CD-ROM publishing date

Unchecked by default. Some organizations distribute databases on CD-ROM rather than replicate them. To receive updates users replicate with a replica on the organization's server. The users specify the date the information was published on the CD-ROM so that the first replication with the organization's replica scans only documents created or modified since the publishing date. If users do not specify the date the initial replication unnecessarily scans the entire database which can be a slow process especially if it occurs over a slow connection.

This options works similarly to the replication history so that only documents that changed after the date are reviewed for replication.

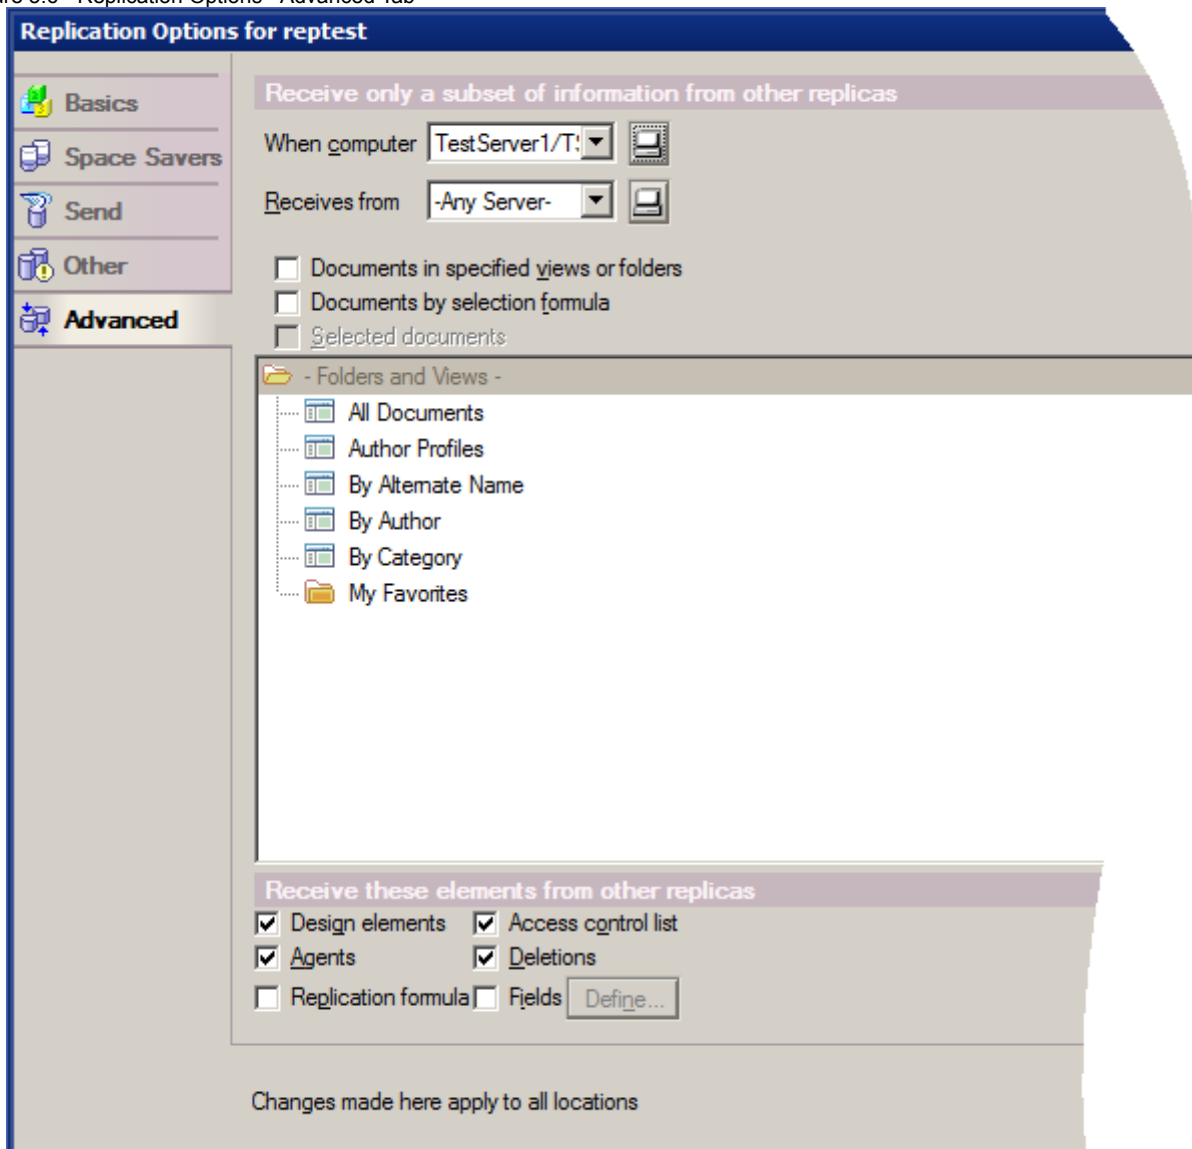
Advanced Tab

The Advanced Tab in the Replication options allows the administrator to control what documents get replicated INTO the database.

This differs from the Space Savers tab by the addition of extra options and by being able to select which servers which documents replicate from.

With the 'When computer' 'Receives from' options it is possible to replicate one set of documents from one replica and another set of documents from a different replica. This can be useful in controlling which documents are contained or available on different replicas.

Figure 5.6 - Replication Options - Advanced Tab



The first two options 'When computer' and 'Receives from' work together. 'When computer' configures the location of the replica to be affected by these settings and 'Receives from' controls what documents are replicated from that server.

Example :

With ServerA, ServerB, ServerC

When computer : ServerA

Receives from : ServerB

Configured not to replicate Design elements. So with this setup ServerA will not receive Design element changes from ServerB

On the same replica

When computer : ServerA

Receives from : ServerC

Configured to receive Design changes. With this configuration ServerA will receive Design element changes from ServerC.

Provided ServerB and ServerC are not independently replicating then ServerB and ServerC can all be replicas but with different Design Element, and ServerA will only receive changes from ServerC.

Documents in specified views or folder

Once selected a list of the Views and folders from the database will be displayed. One or more of these can then be selected. Only the documents listed in these views and folders will be replicated into the database.

Documents by selection formula

This option allows the use of a LotusScript SELECT command to select which documents are replicated. Variations of author, category, creation date, etc. can be used to limit the replication of documents that match the formula.

The default option is :

```
SELECT @All
```

This will allow the replication of all documents into the database.

The 'Fields & Functions...' and 'Formula Windows...' display the available options for this.

Design elements

Selected by default

Selects whether design elements from the defined replica are accepted. It is important to note that in this context Agents are considered separately from Design elements and if their replication is to be restricted they need to be selected separately.

Agents

Selected by default

Selects whether Agents from the design of the remote replica are replicated into this replica.

Replication formula

Not selected by default.

Selects whether the Replication formula from this tab are replicated from the other replicas of the database.

Access control list

Selected by default.

Controls if ACL settings are replicated into this database from the defined replica.

Deletions

Selected by default.

Controls if Deletion stubs are replicated into this database from the defined replica.

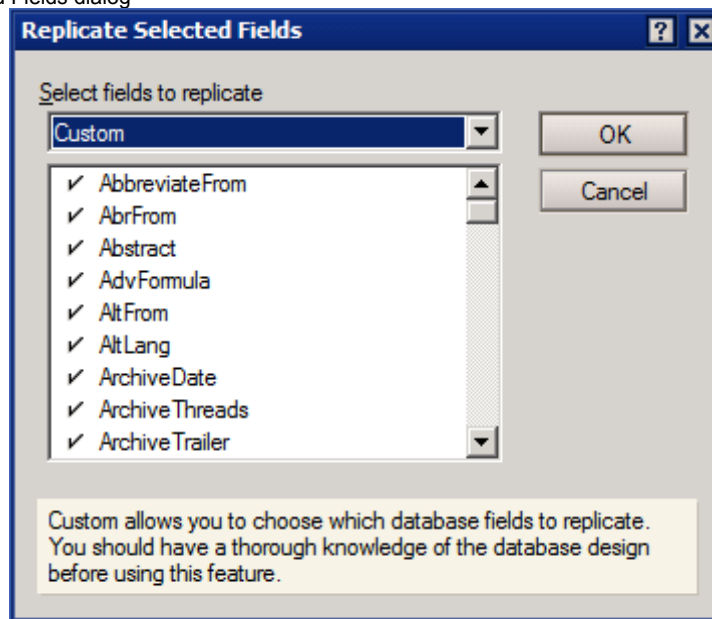
Fields

Not selected by default.

This option can be used to limit which fields from a document will replicate into this replica. Care must be taken with this option if used to ensure that all required field are replicated.

When selected the 'Define...' button is activated. Click this will open the Replicate Selected Fields dialog :

Figure 5.7 Replicate Selected Fields dialog



Selecting Customer as opposed to 'All Fields' in the 'Select fields to replicate' pull-down will display all the fields in the database. The fields not wanted to replicate can be unchecked.

Important note for the Selective Replication settings on the Advanced Tab

The selective replication options on the Advanced tab do not behave as some people expect them to behave. They work by reviewing all the documents in a database and removing any documents that should not be present based on these settings.

In this way documents that are already present on the replica will be removed when the replica replicates not only documents that have just been replicated.

Managers of these replicas should be aware of this before implementing this replication feature.

5.2 Replication and the ACL

In server to server replication it is the server ID that is used to authenticate the addition, editing and removal of documents.

So it is the SERVER ID that needs to be included in the ACL of the database being replicated into.

For example :

Server1 and Server2

With replication being started on Server1

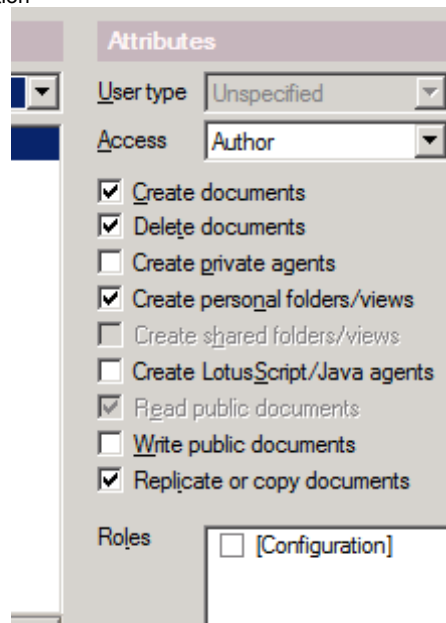
With a standard replication (Pull-Push) Server1 needs to have at least Reader access to Server2 to be able to read the data from it, it will need Editor access to its own database to create, update or remove documents.

When Pushing from Server1 to Server2 then Server1 will need sufficient access to Server2 to place the changes onto Server2.

In order to replicate all deletions, updates and new documents Editor access in the ACL is needed for the server. In order to replicate ACL or Design changes Manager access is needed in the ACL.

Replication can also be disabled in the ACL with the Attributes :

Figure 5.8 - ACL Attributes used for replication



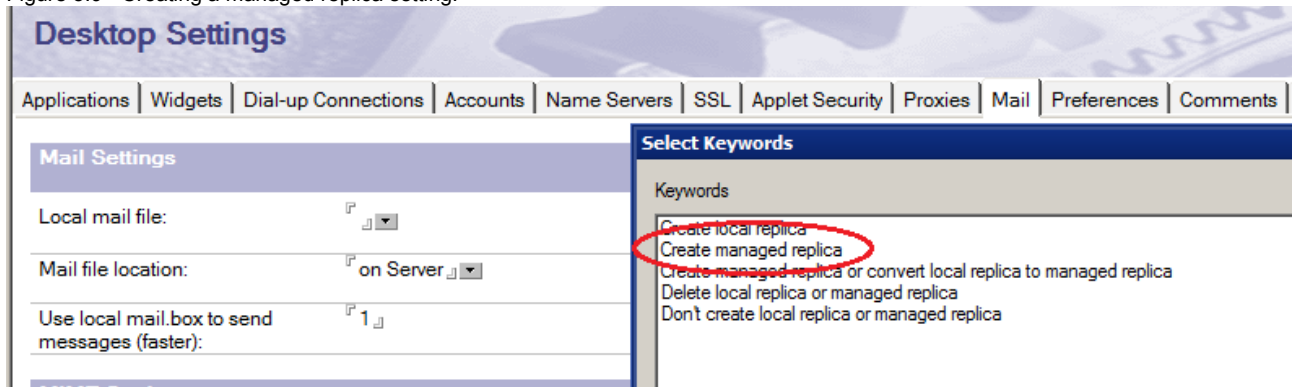
An ACL entry would be needed for Create documents, Delete documents for all changes to replicate. Similarly the Attribute "Replicate or copy documents" needs to be checked for replication to be allowed.

5.3 Managed replicas through a policy

An administrator can use the Managed Replicas Desktop Settings document to set a policy that, for users affected by the policy, marks any existing local mail replica as a managed replica and, if desired, creates a new managed replica if no local replica exists.

This can be configured in the Mail tab of the Desktop Settings by selecting "Local Mail file: Create managed replica".

Figure 5.9 - Creating a managed replica setting.



This enabled the "Managed Replica Settings" section :

Figure 5.10 - Managed replica settings.

Managed Replica Settings	How to apply this setting:	Inherit from parent policy:	Enforce in child policies:
Keep full document for most recent: <input type="text" value="30"/> days	<input type="checkbox"/> Don't set value	<input type="checkbox"/> Inherit	<input type="checkbox"/> Enforce
Selective replica: <input type="checkbox"/> Yes	<input type="checkbox"/> Don't set value	<input type="checkbox"/> Inherit	<input type="checkbox"/> Enforce
Amount of free space required before cache is created: <input type="text" value="0"/> Mb	<input type="checkbox"/> Don't set value	<input type="checkbox"/> Inherit	<input type="checkbox"/> Enforce
Amount of rich text to retrieve on truncated documents: <input type="text" value="0"/> Kb	<input type="checkbox"/> Don't set value	<input type="checkbox"/> Inherit	<input type="checkbox"/> Enforce
Amount of attachment data to retrieve on truncated documents: <input type="text" value="0"/> Kb	<input type="checkbox"/> Don't set value	<input type="checkbox"/> Inherit	<input type="checkbox"/> Enforce

A managed replica is a variation of a local replica of the user's mail file. The managed replica is used by default when the user is connected on a server-based location. It also has the capability of containing full documents for a section (based on time/date) of the mail file and truncated content for the rest (older data).

The managed replica extends the capabilities of a local mail replica and a local outgoing mailbox to those locations where the user remains connected to a Domino server. In this configuration, (connected location with managed replica) when the user opens the mail file, the client will detect that a managed replica for the file exists, open that managed replica (locally) and present the user with the view of the mail from the local file. Updates, sending of mail, and receipt of new mail will happen in the same way as when the user uses a standard local replica with a local outgoing mailbox. An administrator can configure the Notes client to perform this synchronization in as timely a manner as possible.

The typical user of a managed replica is connected to a Domino server for mail, but on a network or line that can have high latency issues. By moving all network I/O between local and server files to the background client replicator, the user's interface will never be blocked waiting for an I/O operation to complete. All the user's operations will occur on the local files.

In this way the Administrator is able to configure the replication schedule and availability for the individual users.

For more details, see Managed Replicas Explained on the Notes and Domino wiki.

http://www-10.lotus.com/ldd/dominowiki.nsf/dx/Managed_Replicas_

5.4 Local Replication settings through a policy

General replication configurations can be applied to all local replicas through policies. These will not just apply to mail databases but to all databases on the Notes Client.

These settings are available through the Desktop Settings, in the Preferences tab and then the Replication tab :

Figure 5.11 - Local Replication Settings.

Default settings for a local replica		How to apply this setting:
Amount of document that Notes should replicate:	<input type="text" value="0"/> ▾	Don't set value ▾
Limit document size of partial documents:	<input type="text" value="0"/> ▾	Don't set value ▾
Partial document size limitation:	<input type="text" value="0"/> KB	Don't set value ▾
Limit attachment size of partial documents:	<input type="text" value="0"/> ▾	Don't set value ▾
Partial document attachment size limitation:	<input type="text" value="0"/> KB	Don't set value ▾
Create a full-text index for faster searching:	<input type="text" value="0"/> ▾	Don't set value ▾
Encrypt replicas:	<input type="text" value="0"/> ▾	Don't set value ▾
When to create replicas:	<input type="text" value="0"/> ▾	Don't set value ▾
Default replication schedule		How to apply this setting:
Enable replication for all locations:	<input type="checkbox"/> Yes	Don't set value ▾
Normal priority replication:	<input type="checkbox"/> Replicate at schedule below for normal priority	Don't set value ▾
Replicate daily between:	<input type="text" value="0"/> ▾	Don't set value ▾
Days of week:	<input type="text" value="0"/> ▾	Don't set value ▾
High priority replication:	<input type="checkbox"/> Replicate at schedule below for high priority	Don't set value ▾
Replicate daily between:	<input type="text" value="0"/> ▾	Don't set value ▾
Days of week:	<input type="text" value="0"/> ▾	Don't set value ▾
Replicate when the user starts the client:	<input type="checkbox"/> Yes	Don't set value ▾
Replicate when Notes ends:	<input type="checkbox"/> Yes	Don't set value ▾

Default settings for a local replica

Amount of document that Notes should replicate:

- Options :
- Receive full documents
 - Receive partial documents
 - Receive summary only

This setting controls if full or partial documents are replicated. It acts the same as the "Receive documents from server" option in the Basics tab of the Replication Options dialog described in Chapter 4 – Running Replication.

Limit document size of partial documents:

- Options :
- Enable
 - Disable

This enables the limiting of the size of partial documents, if they are enabled.

Partial document size limitation:

- Options :
- Number can be inputted in KB

Limit attachment size of partial documents:

Options : Enable
 Disable

This enables the limitation of size of attachments to replicate.

Partial document attachment size limitation

Options : Number can be inputted in KB

Create a full-text index for faster searching:

Options : Enable
 Disable

If enabled a Full Text Index for the local replica will be created when the replica is created. This allows for faster and more accurate searches through the database.

Encrypt replicas:

Options : Do not encrypt
 Locally encrypt

Determines if locally created replicas will be encrypted by default on creation.

When to create replicas:

Options : Create immediately
 Create at next scheduled replication

Controls if a replica will be created immediately when requested by the user or at the next scheduled replication time.

Default replication schedule

Enable replication for all locations:

Options : Yes

Sets if replication will be enabled on all client locations.

Normal priority replication:

Options : Replicate at schedule below for normal priority

Enables the setting of the schedule for Normal priority replication on the Notes Client.

Replicate daily between

Options : User can input the times to replicate to the server. e.g 8:00am 10:00 AM - 6:00 PM

This setting can take individual times or a range of times.

Days of the week:

Options : Sun, Mon, Tue, Wed, Thu. Fri. Sat

Allows the days for this replication to be defined.

High priority replication:

Options : Replicate at schedule below for high priority

Enables the setting of the schedule for High priority replication on the Notes Client.

Replicate daily between

Options : User can input the times to replicate to the server. e.g 8:00am 10:00 AM - 6:00 PM

This setting can take individual times or a range of times.

Days of the week:

Options : Sun, Mon, Tue, Wed, Thu, Fri, Sat

Allows the days for this replication to be defined.

Replicate when the user starts the client:

Options : Yes

Configures if local replication is triggered when the Notes Client is launched.

Replication when Notes ends:

Options : Yes

Configures if local replication is triggered when the Notes Client is shut down.

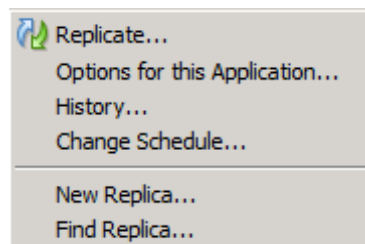
5.5 Replication settings on a Notes Client

There are several replication settings available through a Notes Client. As outlined in the Client to Server section of Chapter 4 the replication schedule for the current Location document can be accessed and modified through the Basics tab of the Replication Options that can be opened from the menu File/Replication/Options for this application...

File/Replication :

All the menu options available here are :

Figure 5.12 – Replication menu options



Replicate : Will trigger replication for the current database – with the option to run in the background or foreground.

Options for this Application... : Will option the Replication Options dialog as discussed in section 5.1 Replication Settings.

History : Will option the Replication History for the current database.

Change Schedule... : Will open the Replication/Sync Schedule for location dialog

New Replica... : Will open the New Replica dialog.

Find Replica... : Will search for replicas in selected server and add them to the WorkSpace.

The Replication and Sync page.

By clicking on the Replication and Sync icon on the left hand side tool bar of the Notes Client you will open the Replication and Sync page.

This icon is :



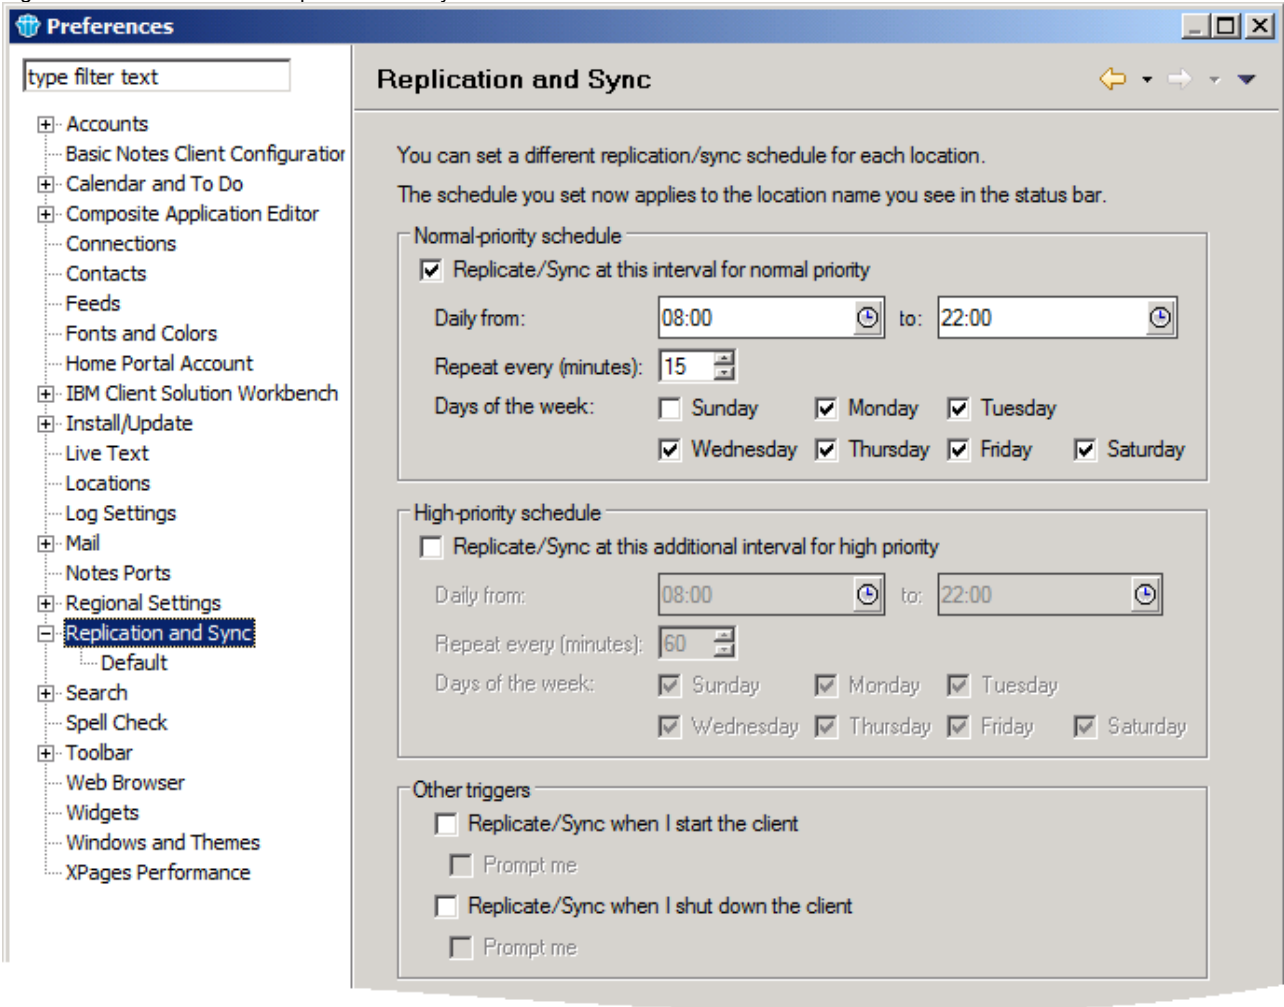
The Replication and Sync page looks like this :

Notes Client Preferences :

By selecting File/Preferences the Preferences dialog is opened. In the section on the left is a Replication and Sync line which when expanded opens to Replication and Sync/Default.

The Preferences for Replication and Sync are :

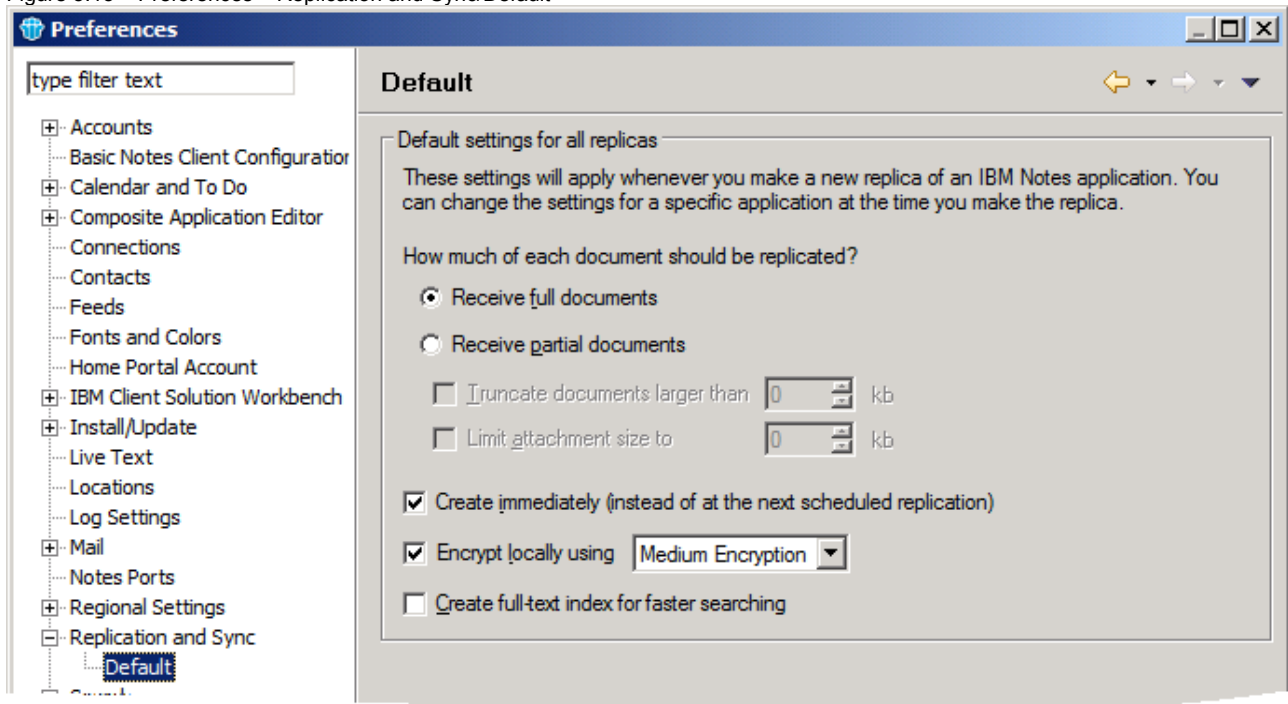
Figure 5.14 – Preferences – Replication and Sync



Here the replication schedule can be modified for the active Location.

The Preferences for Replication and Sync/Default are :

Figure 5.15 – Preferences – Replication and Sync/Default



As mentioned in the dialog itself these settings will be applied when a new replica of a Databases is created locally on the Notes Client.

Note : The Preferences configured through this dialog can also be configured through a Policy settings document by an administrator for multiple users. They are available in the Desktop settings through the Preferences tab and then the Replication tab.

Figure 5.15 – Desktop Settings, Policy document – Replication Preferences

Default settings for a local replica	Default replication schedule
Amount of document that Notes should replicate: <input type="text" value="0"/> KB	Enable replication for all locations: <input type="checkbox"/> Yes
Limit document size of partial documents: <input type="text" value="0"/> KB	Normal priority replication: <input type="checkbox"/> Replicate at schedule below for normal priority
Partial document size limitation: <input type="text" value="0"/> KB	Replicate daily between: <input type="text" value=""/>
Limit attachment size of partial documents: <input type="text" value="0"/> KB	Days of week: <input type="text" value=""/>
Partial document attachment size limitation: <input type="text" value="0"/> KB	High priority replication: <input type="checkbox"/> Replicate at schedule below for high priority
Create a full-text index for faster searching: <input type="text" value=""/>	Replicate daily between: <input type="text" value=""/>
Encrypt replicas: <input type="text" value=""/>	Days of week: <input type="text" value=""/>
When to create replicas: <input type="text" value=""/>	Replicate when the user starts the client: <input type="checkbox"/> Yes
	Replicate when Notes ends: <input type="checkbox"/> Yes

Chapter 6 : Other Areas of Replication

There are further areas and options that are affected or can affect replication, these are not involved directly with either the running or configuration of replication itself so are separate from those chapters.

Here we will be looking at :

- 6.1 Database Quotas
- 6.2 More on Deletion Stubs
- 6.3 Replication Conflicts
- 6.4 Networking
- 6.5 'Full replication' what can trigger it? Same as clearing the history
- 6.6 Database properties that DON'T replicate
- 6.7 Unread Marks
- 6.8 Soft deletions
- 6.9 DAOS and Replication
- 6.10 Replication and ODS
- 6.11 Replication and different server versions (32 and 64 bit)
- 6.12 Multiple replicas of the same database on a single server
- 6.13 Replicating folders (private and public)
- 6.14 CATALOG.NSF
- 6.15 System databases with related Replica IDs
- 6.16 nReplica command with the server shutdown
- 6.17 Triangulation
- 6.18 Choosing which replica to open on a Notes Client
- 6.19 Other types of replication

6.1 Database Quotas

Controlling the size of databases can be very important on a large server. The use of Quotas on a database allows an administrator to configure both a warning threshold and a cut-off for when a database reaches a defined size.

Information on the setting and use of database quotas can be found in the Domino Administration help document : Database size quotas

http://www-10.lotus.com/ldd/dominowiki.nsf/xpDocViewer.xsp?lookupName=Administering+IBM+Domino+9.0.1+Social+Edition#action=openDocument&res_title=Database_size_quotas&content=pdcontent

Where this becomes important for replication involves how replication handles Database quotas.

Replication, by default, will ignore database quotas.

Normally when a database quota is reached it will prevent the saving of new documents to the database. However, Quotas are in place to prevent the users from increasing the size of the database directly themselves and Replication is allowed to complete to ensure that all documents are present on the database.

In this way all documents will replicate into a database even when the quota has been reached but the user will not be able to save documents into the database.

This option can be disabled with the NOTES.INI parameter :

Repl_Obeys_Quotas=value

Description:

0 - Disables the Replicator from obeying quotas.

1 - Enables the Replicator to obey quotas.

Note : The REPL_OBEYS_QUOTAS setting will prevent the replicator task from replicating documents into a database that has grown past its quota, however the Cluster Replicator is a separate task from the Replicator and is also designed, by default, to replicate into databases past their quota. To disable this replication as well you should use the NOTES.INI parameter :

Clrepl_Obeys_Quotas=value

0 - Disables the Cluster Replicator from obeying quotas.

1 - Enables the Cluster Replicator to obey quotas.

Note 2 : Database quota settings DO NOT REPLICATE. Setting a quota on one replica will not affect another replica. This is to allow for different requirements on different systems. Whereas an administrator may want a limit on the size of a mail file on a server the user may want to keep all old mail documents on a local replica.

6.2 More on Deletion Stubs

As mentioned previously Deletion Stubs are truncated documents that are left in a database in place of the original documents when those documents are deleted to indicate to the Replication task that the document should, in fact, be deleted from all other replicas.

Deletion Stubs have the same UNID as the original document.

The deletion stub is the element that replicates when a document is deleted. When it replicates it replaces the target document, causing the target document to be deleted.

They do contain Time and Date information for when they were created.

If the target document has been edited and saved after the deletion stub was created then the deletion stub will not replicate and if two way replication is in place the deletion stub will be replaced with the document again.

Note : It is important that the Purge Interval, set in the Replication Options, is longer than the time between replicas. If it isn't then there is a possibility for deleted documents to return.

For example :

With a database on ServerA and ServerB

Purge interval is set to 5 days.

Replication is configured to replicate the database every 7 days.

Day 1 document is deleted from ServerA

Day 5 the deletion stub is purged from ServerA

Day 7 ServerA and ServerB replicate. The document is still present on ServerB and there is no matching document (with the same UNID) on ServerA so the replicator will replicate the document back to ServerA.

When the Replication frequency is higher than the Purge Interval the Deletion Stub will replicate to ServerB removing the document from ServerB before the Deletion Stub is purged from ServerA.

6.3 Replication Conflicts

In a scenario where there is a replica of a document on two or more servers it is possible that multiple replicas of the document could be individually edited and saved on different databases before they are replicated together. This leads to the documents both trying to replicate their changes to each other.

Normally the replicator task is not able to discern which of the updates should take precedence and so it will generate a Conflict document. These are referred to as "Replication or Save conflicts" as they can also be generated when two users edit a document on a single database at the same time.

One version becomes the "winner", meaning it becomes the main or parent document. The "loser" becomes a response document and is identified by a black diamond and the words "Replication or Save Conflict".

Figure 6.1 - Replication or Save Conflict



Deleting the "loser" but saving its information:

Open the document and see if it contains any information that should be added to the winner document. If

so, copy the information and paste it into the winner document and then delete the loser.

Deleting the "winner" and saving the "loser":

Open the loser document, switch to edit mode, and then save the document. This removes the document's Replication or Save Conflict status, removes the black diamond and promotes it to main document status.

When you edit the loser, Notes checks the winner to see if there are any response documents against it. While skipping over the conflict itself, Notes then updates the responses, changing their \$Ref fields to the UNID of the newly edited conflict document. Notes also checks the winner to see if it is a response to another document.

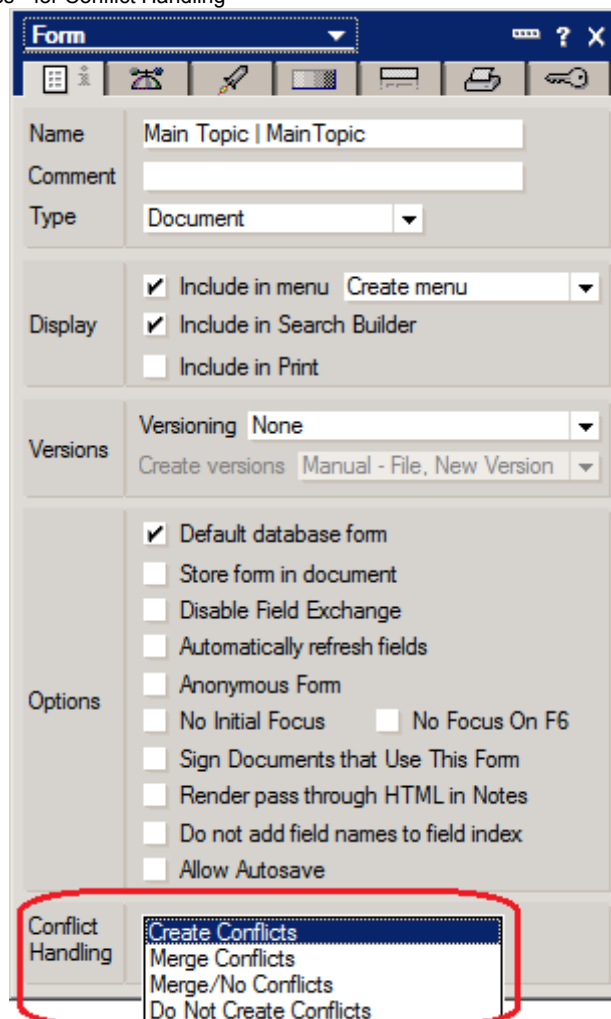
The loser's \$Ref field is set to match the \$Ref of the winner. This process results in replacing the winner with the loser in the response hierarchy. The winner can then be deleted from the database without orphaning any response documents.

A further options to reduce the number of Conflict documents generated is to enable "Merge Conflicts" in the Form Properties in the Domino Designer Client.

To access this option

- open the database or its template (and update the design of the database) in the Domino Designer client.
- select and open the Form that you want to enable this on.
- open the Form Properties - the 'Conflict Handling' is at the bottom of the first tab :

Figure 6.2 - Form Design Properties - for Conflict Handling



•**Create Conflicts** -- Creates conflicts so that a replication conflict appears as a response document to the main document. The main document is selected based on the time and date of the changes and an internal document sequence number.

- Merge Conflicts** -- If a replication conflict occurs, saves the edits to each field in a single document. However, if two users edit the same field in the same document, Notes® saves one document as a main document and the other document as a response document marked as a replication conflict. The main document is selected based on the criteria listed in the preceding bullet.
- Merge/No Conflicts** -- If replication occurs, saves the edits to each field in a single document. If two users edit the same field in the same document, Notes selects the field from the main document, based on time and date, and an internal document sequence number. No conflict document is generated, instead conflicting documents are merged into a single document at the field level.
- Do Not Create Conflicts** -- No merge occurs. Domino simply chooses one document over another. The other document is lost.

Note : In some configurations the View design can be set to not display response documents, in these cases conflict documents are not displayed in the View. This can lead to unexpected data loss in a documents and to a build up of old documents that are not being displayed normally in the database.

6.4 Networking

For replication to be possible between replicas it is necessary for there to be some form of connection. Normally this is some form of wired network, but it could also be dial up connections etc.

Replication between servers in the same Domain do not need connection documents. However, replication between servers from different Domains will require a Connection Document to link the Domains.

Similarly, from a Certification perspective the Servers need to be able to authenticate into each other.

If the server console command 'Trace' is unable to connect to the other server then replication will not be possible.

6.5 Full Replication what can trigger it? Same as clearing the history

As mentioned previously the replicator task will use the Replication History to confirm when the last replication occurred and use this information to limit the next replication to only review documents that have changed since the last replication.

Clearing the Replication History will cause the next replication to review all the documents in the databases to see if they match those in the replica. This can take longer to complete.

However, there are other actions that can also lead to this 'Full Replication' in the databases. Normally these are changes that could affect which documents could be replicated. If the change could cause documents that were previously configured to not replicate to now replicate then a full replication is triggered.

These actions include :

- Changes to the ACL.
- Changes to the Replication Options.

6.6 Database properties that DON'T replicate

Although there are several database properties that will replicate there are also several that don't. As mentioned Database Quotas do not replicate. Often this is to allow for different configurations between replicas.

This is a list of common database properties and whether they replicate or not :

BASICS TAB	
Archive settings	Yes
Replication settings	No
Encryption settings	No
Disable background agents for this database	No
Allow use of stored forms	Yes

Display images after loading	Yes
Web access: Use JavaScript when generating pages	Yes
Web access: Require SSL	Yes

DESIGN TAB	
Inherit design from template (template name)	Yes
Database is a template (template name)	Yes
List in Database Catalog	Yes
Show in 'Open Database' dialog	Yes
Do not mark modified documents as unread	No
Include in multidb indexing	Yes
Multilingual database	Yes
Default language,region, sort order	Yes

ADVANCED TAB	
Don't maintain unread marks	Yes
Replicate unread marks	Yes
Document table bitmap optimization	Yes
Don't allow simple search	No
Don't overwrite free space	No
Maintain LastAccessed property	Yes
Disable transaction logging	No
Allow soft deletions	Yes
Don't support specialized response hierarchy	Yes
Don't allow headline monitoring	No
Allow more fields in database	Yes
Limit entries in \$UpdatedBy fields	Yes
Limit entries in \$Revisions fields	Yes
\$Undelete Expire Time	Yes

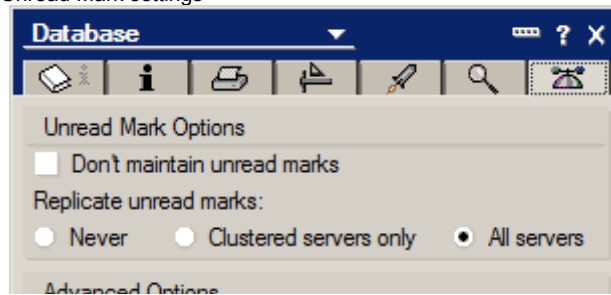
6.7 Unread Marks

Unread Marks are the method by which documents that have been added to a database, but not yet read by the user are marked.

Originally when users only had one mail file the replication of Unread marks was not required, however since users now often work from several locations with different computers and servers are often clustered the replication of Unread marks has become more important.

For individual databases the replication of Unread marks can be configured from the Database Properties. These are set in the Advanced tab :

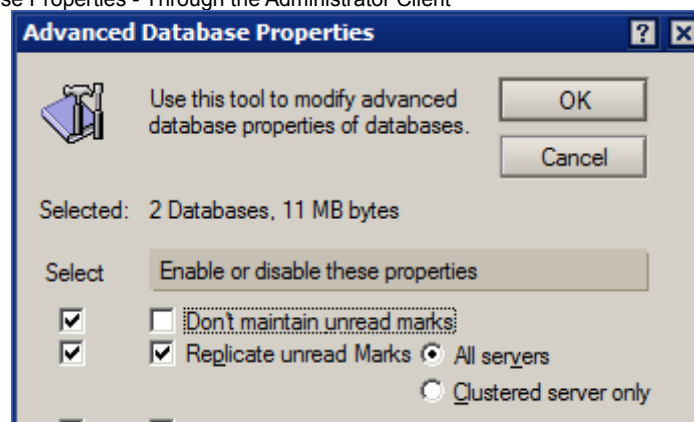
Figure 6.3 - Database properties - Unread Mark settings



Here you can configure if Unread Marks are to be maintained for the database and if the Unread marks will replicate.

These settings can be configured for multiple users through the Administrator Client. From the Files tab select the databases that need to be configured and click on Tools\Database\Advanced Properties...

Figure 6.4 - Advanced Database Properties - Through the Administrator Client



Here you can set if Unread marks are to be maintained and if they will replicate.

Note : Replicating unread marks in high-activity user databases other than mail databases is not recommended due to the performance impact.

More information on Unread Marks can be found in the Open Mic Webcast : Unread Marks

Title: Open Mic Webcast: Unread Marks - 4 June 2014 (Q&A, presentation, audio replay)

Doc #: 7042034

URL: <http://www.ibm.com/support/docview.wss?uid=swg27042034>

Note : As mentioned in Chapter 3 - Creating Replicas when creating a replica of a new database through AdminP the option "Exchange Unread marks on replication" must be enabled if the Unread marks should be replicated this is off by default and also requires a Notes.INI setting :

ADMINP_EXCHANGE_ALL_UNREAD_MARKS = 1

6.8 Soft deletions

If Soft Deletions are enabled on a database when a document is marked for deletion it is moved to the Trash folder and will remain there for a defined amount of time. After the Soft Delete Expire Time or if the Trash is emptied the document will be 'Hard' deleted converting it to a Deletion Stub that will replicate normally.

If a document is created on one replica, replicated to another replica and then Soft Deleted and replicated again - then the document will be moved to the Trash folder for both replicas. However, in the situation where the document is created on one replica and then soft deleted BEFORE it is replicated the document will be moved to the Trash folder on the original replica but will not be replicated to the second replica. When the Soft Deletion is then Hard Deleted the deletion stub will replicate normally.

6.9 DAOS and replication

DAOS stands for Domino® Attachment and Object Service.

This process allows attachments to be stored outside of the database with a link in the documents to the attachment file. In this way documents that use the same attachment can link to the same single file reducing the amount of disk space used.

However DAOS can also save in the amount of data needed to be transmitted during replication.

In a situation where both the source and target servers are DAOS enabled if the attachment file is already present on the target server then it is not replicated.

In a situation where both the source and target servers are DAOS enabled but the attachment file is not already on the target server only one copy of the attachment is replicated. In situations where multiple documents are using the same attachment this can reduce the amount of data replicated significantly.

In a situation where the source server is DAOS enabled but the target server is not then the documents will replicate normally and the attachments will be included in all replicated documents. Similarly if the source server is not DAOS enabled then all attachments will replicate normally.

It is important to Note that Notes Clients are not DAOS enabled, all replication to and from a Notes Client will include all attachments in the replication.

6.10 Replication and ODS

ODS stands for On-Disk Structure. It relates to the internal layout of a Domino Database. The ODS level used by a server can change from release to release.

Because of this and because Domino Domains often consist of multiple servers at different release levels it is very common for different replicas to be at different ODS levels.

This is not a problem and the documents between the different ODS levels will replicate normally.

Replicating a database of one ODS to a database of another ODS will not affect the ODS of the target database, only the data will be replicated.

6.11 Replication and different server versions (32 and 64 bit)

As mentioned above with ODS replication has no problems with working between servers and clients that are of different versions. This includes 32 and 64 bit versions.

6.12 Multiple replicas of the same database on a single server

It is possible to create or place multiple replicas of the same database on the same server. It is strongly recommended that this not be done. As replication uses the Replica ID to find the replica this can lead to situations where the wrong replica is updated, where updates go into one replica and out from another or even where changes are merged between the replicas on the same server. However none of these possibilities are predictable.

This situation can happen when administrators will create an Operating System copy of a database and change its filename assuming that this will differentiate it from the original database.

This can lead to unpredictable behaviour on the server especially if the duplicated replica is the NAMES.NSF or some other system database.

6.13 Replicating folders (private and shared)

Folders are containers used to store related documents or groupings of documents. Folders have the same design elements as views.

Primarily there are two types of folders - Shared and Private.

Shared folders will replicate normally.

Private folders are themselves of two types - desktop private and private.

Desktop private folder data is stored on the users Notes Client desktop - so is not available to the server for replication. However, if the user accesses a replica of the database the folder data will still be available for the replica through their desktop.

Private folder data is stored within the database. However the data uses a Reader field to prevent other users apart from the author or first user to access them, this prevents the server from replicating the data. If the user access the replica of the database the private folder information will not be available.

6.14 CATALOG.NSF

The CATALOG.NSF or Database Catalog contains a list of all databases on a server.

The information in the CATALOG.NSF includes (amongst other things) the filename, Replica ID and ACL information for each database. This information can be very useful in troubleshooting issues with replication and for checking to see if there are any databases with duplicate Replica IDs on the server.

6.15 System databases with related Replica IDs

Many of the system databases are required to replicate with each other throughout a Domino Domain, even if they are recreated. However, in most cases a newly created database will have a new Replica ID and will not replicate with the old copy.

Because of this some databases, when recreated, will have the same Replica ID again to allow them to replicate with their counterparts on other servers.

These Replica ID are based on the Replica ID of the Domain's NAMES.NSF.

for example :

```
NAMES.NSF 80257D2A:0070A553
ADMIN4.NSF 80257D2A:0370A553
CATALOG.NSF 80257D2A:0770A553
EVENTS4.nsf 80257D2A:0270A553
```

Note that it is only the first 2 digits of the second set of 8 digits that differ.

The following databases use this system allowing them to be recreated with the same Replica ID throughout the Domain :

```
ACTIVITY.NSF
ADMIN4.NSF
BILLING.NSF
CATALOG.NSF
DDM.NSF
EVENTS4.NSF
STATREP.NSF
VPUSERINFO.NSF
```

Every time these databases are recreated they have the same Replica ID allowing them to replicate throughout the Domain.

For example, if the Admin4.NSF database needs to be recreated if it is corrupt on one server. Once it has been re-made it will replicate all the process documents back into it from the replica on another server.

This process ensures that these system databases can replicate normally which is required for a lot of server functionality to run correctly.

If any of these databases is replaced with a Notes Copy of itself the Replica ID will be different and they will not replicate with each other and this could impede some server functionality.

If the NAMES.NSF is replaced with a Notes Copy of itself with a different Replica ID then if any of the other databases are then re-created they will be re-created based on the incorrect Replica ID of the new NAMES.NSF.

6.16 nReplica command with the server shutdown

There are a number of commands on a server that can be run while the server is actually shutdown. While these are normally restricted to database maintenance commands there are other options available as well - surprisingly including replication.

The options for Pull or Push only are only available through the 'mode' options :

For ONE TIME ONLY replication: You can specify optional parameters for replication, then the currently loaded replicator task will shut down.

Usage:

```
Load REPLICIA [[[server] database] mode]...
```

[server] Name of destination server to replicate with.

[database] Name of database or folder to replicate with destination server.

[mode]

- p Perform pull-only replication.
- s Perform push-only replication.
(Default is 2-way replication)

Example replication between TestServer1/TST and TestServer2/TST, from the TestServer1 machine :

```
C:\IBM\Domino>nreplica testserver2 reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:36 Database Replicator started
[21EC:0002-2C24] 09/09/2014 10:51:36 Replicator is set to Ignore Database
Quotas
[21EC:0002-2C24] 09/09/2014 10:51:37 Starting replication with server
TestServer2/TST
[21EC:0002-2C24] 09/09/2014 10:51:37 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Replicator added 3 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Replicator deleted 10 document(s) in
reptest.nsf from TestServer2/TST reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Pushing reptest.nsf to TestServer2/TST
reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Replicator added 7 document(s) to
TestServer2/TST reptest.nsf from reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Replicator deleted 12 document(s) in
TestServer2/TST reptest.nsf from reptest.nsf
[21EC:0002-2C24] 09/09/2014 10:51:37 Finished replication with server
TestServer2/TST
[21EC:0002-2C24] 09/09/2014 10:51:37 Database Replicator shutdown
```

6.17 Triangulation

In a situation where three or more servers are replicating there was a chance for inefficiency to be introduced.

With three servers :

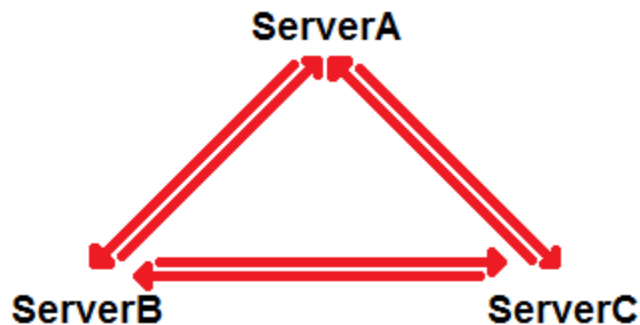
ServerA
ServerB
ServerC

When ServerA replicated to ServerB then ServerA replicated to ServerC technically ServerB and ServerC should match. However, previously, the replication history would still cause a check and replication of documents updated after the last replication between ServerA and ServerC.

This resulted in the use of extra network bandwidth when ServerA and ServerC replicated.

This led to the introduction of what is known as "Triangulation".


Figure 6.5 - Triangulation



So, when ServerA and ServerB and ServerA and ServerC have replicated fully it is considered that ServerB and ServerC have fully replicated as well.

This is only possible when ServerA and ServerB and ServerA and ServerC have completely replicated with each other. A way of marking when this happens was introduced into the Replication History of a database :

Figure 6.6 - Replication History showing the Asterix



After a complete replication an Asterix (*) is placed beside the Action.

Along with the asterix the Replication History itself is now replicated, this can lead to extra servers being seen in a Replication History for servers that the database has never replicated directly with.

Triangulation can be disabled with the following NOTES.INI settings :

```
NSF_REPLHIST_NO_TRI=1
REPL_NO_WS_TRI_HIST=1
REPL_NO_REMOTE_TRI_HIST=1
```

For Notes Client based replication :

```
NSF_REPLHIST_NO_TRI=1
    This will prevent existing triangulated entries from being read
REPL_NO_WS_TRI_HIST=1
    This will prevent new triangulated entries from being written
```

If there are restrictions on replication then the Asterix will not be added to the Replication History.

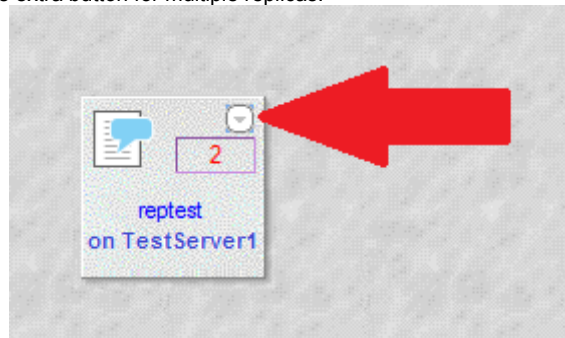
6.18 Choosing which replica to open on a Notes Client

If multiple replicas of a database have been opened through a Notes Client then there will be a link available to each individual replica.

Normally these links will be 'stacked' together into a single tile on the Notes Client Desktop.

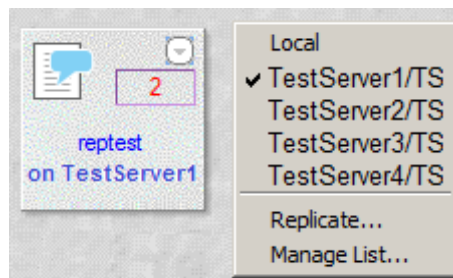
Then there are links to multiple replicas an extra button is displayed on the tile for the database :

Figure 6.7 - Database tile showing the extra button for multiple replicas.



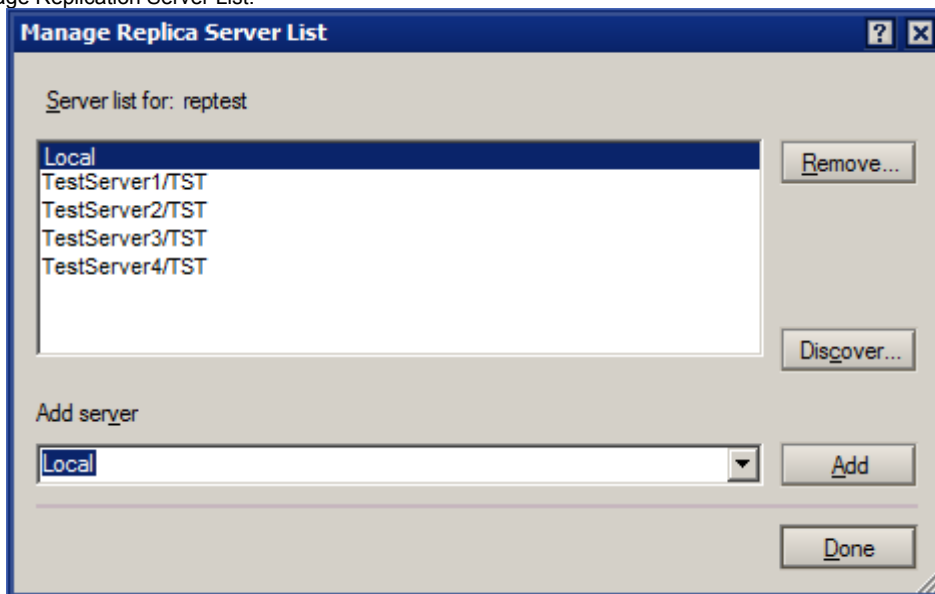
Clicking on this button will display a menu listing all the available replicas, the option to replicate with the selected replica from the local replica (if there is one) and the option "Manage list" :

Figure 6.8 - Manage Replication menu list.



Selecting the Manage List option opens the Manage Replica Server List :

Figure 6.8 - Manage Replication Server List.



Here it is possible to remove listed replicas so they are not displayed.

The Discover button will search for all known replicas listed in the Domain Catalog and add them to this list. When clicking on it a warning is displayed informing that this can take an extended period of time to complete.

As mentioned here normally the links to all the replicas are stacked together allowing for a smaller and easier way to view all databases on the Desktop. However, the option to have all the replicas stacked can be disabled by unchecking the Notes Client menu option View/Stack Replica Icons. If this is disabled then each replica for all databases will have an individual icon on the Notes Client desktop.

6.19 Other types of Replication

Cluster Replication

The most common 'other' type of replication used with a Domino Server is Cluster replication. This is a separate task to the Standard Replicator that we are going through here.

There are some fundamental differences between the Cluster Replicator and the standard replicator.

The Standard Replicator will compare databases, create a list of documents that have changed since the last replication (or list of all documents that don't match), compare the documents and replicate them as necessary. This happens either on a schedule or when manually triggered.

The Cluster Replicator logs changes made to a database and then replicates the changes. This happens automatically very frequently (every couple of seconds) to keep the documents matching. The Cluster replicator does not compare the documents to each other.

Cluster replication is what is known as 'event driven replication'.

Cluster replication can, in some cases, fail. This is most common when the number of changes introduced is too high for the cache and some are missed when the documents are cluster replicated.

To prevent this it is required that a Domino Cluster also maintain a standard replication schedule between all databases that are maintained in a cluster.

Archiving

Database Archiving is another 'type' of replication. This uses the Standard Replicator task to move documents from one database to another database (normally removing them from the original) based on a criteria that is used to select the documents.

This is one of the few times that the Standard Replicator task is able to function between two databases that do not have the same Replica ID.

Archiving does use the file name of the database to define which database to archive with.

Mail.BOX replication

When using a local replica of a mail file and sending a mail it is sent to a local MAIL.BOX and replicated from there to the MAIL.BOX on the server. This is another example of where the Standard Replicator is used but the source and target databases do not need to have matching Replica IDs.

Chapter 7 : Logging and Debug options

There are a number of issues regularly seen with replication and several of these scenarios, with solutions, will be reviewed in the Troubleshooting chapter below.

Before looking at these common replication issues section we are going to have a look at the common Logs and Debug options available that can be used to find and track down replication data.

What we're going to look at here :

7.1 Where and what to look for :

LOG.NSF
CONSOLE.LOG
DDM.NSF

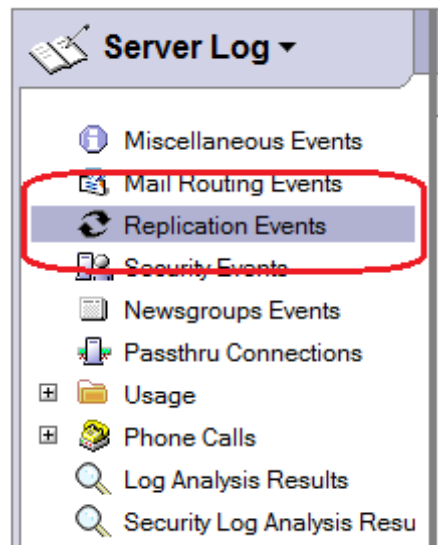
7.2 Extra logging and Debug parameters

7.1 Where and what to look for

LOG.NSF

The LOG.NSF is a very useful source for data relating to replication. Opening the LOG.NSF you can see that there is a list of Views on the left hand side including : Replication Events

Figure 7.1 - Views available in the LOG.NSF - highlighting Replication Events



Opening this View will display a list of all the servers that this Server (or Notes Client) has replicated with.

Each server section will expand out to different dates which, in turn, expand out to times for that day :

Figure 7.2 - Replication Events view in LOG.NSF

Search in View 'Replication Events'						
Search for						
With Server	Date	Starting Time	Ending Time	Minutes	Average	Initiated By
[-] TestServer1/TST				11.4	3.8	
	[-] 03/09/2014			0.3	0.3	
		16:57	16:57	0.3		TestAdmin/TST
	[-] 05/09/2014			11.1	1.2	
		11:19	11:19	0.0		TestAdmin/TST
		11:20	11:20	0.0		TestAdmin/TST
		11:23	11:23	0.1		TestAdmin/TST
		11:25	11:25	0.0		TestAdmin/TST
		11:25	11:25	0.0		TestAdmin/TST
		11:28	11:39	11.0		TestAdmin/TST
		11:31	11:31	0.0		TestAdmin/TST
		11:47	11:47	0.0		TestAdmin/TST
		16:55	16:55	0.0		TestAdmin/TST
	[-] 08/09/2014			0.0	0.0	
		10:33	10:33	0.0		TestAdmin/TST
[-] TestServer2/TST				77.1	77.1	
	[-] 05/09/2014			77.1	38.5	TestAdmin/TST

Opening each section will display information regarding the replication events that happened at that time :

Figure 7.3 - Replication Events view in LOG.NSF

Replication Events

Name: 1 TestServer1/TST
 Time: 2 12/08 11:41 - 12/08 11:41
 Remote Server: 3 TestServer2/TST
 Initiated By: 4 TestServer1/TST
 Elapsed Time: 5 0.3 minutes

Databases replicated:	Access:	Added:	Deleted:	Updated:	KB Rec'd:	KB Sent:
TestServer1/TST reptest.nsf 6 From: TestServer2/TST (PULL)Forms	7 Manager	0	0	4	10.75	0.80
TestServer2/TST reptest.nsf 6 From: TestServer1/TST (PUSH)Views	7 Manager	150	0	0	18.57	1,410.07

Here you can see :

- 1) The Source Server or Notes Client.
- 2) The time the replication was initiated.
- 3) The Server being replicated to - the target server.
- 4) Which Server initiated the replication.
- 5) How long the replication events took.
- 6) Listing of the database being replicated.
- 7) The ACL access of the actioning server to the other server.
This could be two different results as it would be TestServer1's access to TestServer2 and then TestServer1's access to itself as the Pull and Push replication are performed.
- 8, 9 & 10) The Added, Deleted and Updated number of documents.
- 11 & 12) The amount of data Received and Sent during the replication.

All of this can be very helpful if any issues with replication are found. The time of the replication can be used to cross-reference with other sources of data in case of network

outages etc.

The target server can be used to check if there were issues with that server.

The length of time for the replication can show if there was a network problem, matched to the amount of data sent and received - if the replication elapsed time is high but so is the amount of data sent and received this would be normal.

ACL access can explain why some changes have not replicated.

The document Added, Deleted and Updated counts can be useful to track replications.

Apart from the Replication Events View the Usage Views in the LOG.NSF can also be very useful in tracking changes to databases. The Usage\by Database View will show the replication and other events (Reads, Writes) that can be useful in tracking down changes to a database.

CONSOLE.LOG

The CONSOLE.LOG is a text log of all the data written to the Server Console. By default this will include all replication events between servers including the server being replicated with and the database being replicated (the level of replication reporting in use is defined with the NOTES.INI parameter LOG_REPLICATION - more on this in the debug section later on, but its default setting is '2')

By default the CONSOLE.LOG will be generated in the DATA\IBM_TECHNICAL_SUPPORT folder. It is re-created each time the server is restated with the old CONSOLE.LOG being renamed to

`CONSOLE_Server-Name_Year_Month_Day@Time.LOG`

Often errors occurring during replication will be reported here.

By default the size of the CONSOLE.LOG is fixed (the size depends on the Server Version) preventing the log from growing beyond a defined limit. When this limit is reached the logging will loop back to the start of the file and overwrite the first entries, this can lead to lost data. To allow the CONSOLE.LOG to grow beyond this size add :

`CONSOLE_LOG_ENABLED=1`
to the NOTES.INI and restart the server.

or enter the command :
Start consolelog

Note : using 'Start consolelog' will only affect the current running server, if the server is restarted and 'CONSOLE_LOG_ENABLED=1' is not set in the NOTES.INI then the old functionality will return.

In the case of replication on a Notes Client the same data can be collected using the "CONSOLE_LOG_ENABLED=1" parameter in the Notes Client's NOTES.INI file - once restarted the CONSOLE.LOG will be populated. It is also useful to know that the replication logging and debug options listed later on will also work on a Notes Client.

Running a standard replication will display the following information on the server console :

```
> rep testserver2 reptest.nsf
[5340:0002-3D68] 09/09/2014 11:20:13 Database Replicator started
[5340:0002-3D68] 09/09/2014 11:20:13 Replicator is set to Ignore Database
Quotas
[5340:0002-3D68] 09/09/2014 11:20:13 Starting replication with server
TestServer2/TST
[5340:0002-3D68] 09/09/2014 11:20:13 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[5340:0002-3D68] 09/09/2014 11:20:13 Replicator added 2 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[5340:0002-3D68] 09/09/2014 11:20:13 Replicator deleted 2 document(s) in
```

```
reptest.nsf from TestServer2/TST reptest.nsf
[5340:0002-3D68] 09/09/2014 11:20:13   Pushing reptest.nsf to TestServer2/TST
reptest.nsf
[5340:0002-3D68] 09/09/2014 11:20:14   Finished replication with server
TestServer2/TST
[5340:0002-3D68] 09/09/2014 11:20:14   Database Replicator shutdown
```

Here the additions and deletions from replication are listed.

Adding debug to the server will increase and alter the output data displayed on the server console. This can help in tracking down the source of issues with replication.

DDM.NSF

DDM is Domino Domain Monitoring, it uses a set of probes to monitor areas of the server. If an issue is found with one of the probe results a number of events can be triggered.

Replication has its own probe type.

The replication probe checks server replication. To configure the replication probe, specify the target servers to be probed.

Target servers are those servers that replicate with the probe server. For example, if the probe runs on ServerA, the target servers to monitor are those servers that replicate with ServerA. In this case, ServerB, ServerC, and ServerD.

Use the Replication/Replication Check probe to filter out replication of databases that do not require replication, and to closely monitor replication of those databases that require frequent replication. Critical system databases, such as NAMES.NSF, EVENTS4.NSF, and ADMIN4.NSF which are used across servers and that are required to be up-to-date, can be monitored more often than less-important databases. You can also define a lower tolerance for conditions that create events on system-critical databases than you would define for conditions that create events on non-critical databases.

Replication on selected servers and databases is monitored according to the schedule specified in the Probe document.

When you run the Replication/Errors probe replication occurs as it typically would, but events are reported to DDM. This probe may also report some conditions that are not errors, but that are not desired behaviour. For example, this probe may report Replication Note Errors when the replication is successful but some notes are not replicated. If the destination database does not have the source database in its ACL, and there have been changes to the source database, the notes in the source database which are new or updated are not replicated to the destination server because of the ACL. The replication probe will report this and include a document link to the notes that did not replicate. The Replication/Errors probe may also report replication conflicts in DDM and include a document link to the conflict document. This is not a true error condition for replication but it may not be desired behaviour; therefore, it is reported.

The Replication/Replication Check probe takes null replication into consideration. Null replication occurs when a database is up-to-date and no replication is required, but the replication is attempted. Null replication is not reported in the DDM database.

In hub-and-spoke topologies, replication is typically one-way only. Do not monitor replication in a direction in which replication is disabled.

7.2 Extra logging and Debug parameters

There are a number of options to increase the level of reporting from replication.

```
DEBUG_THREADID
LOG_REPLICATION
DEBUG_REPL & DEBUG_REPL_ALL
DEBUG_REPL_TIME
```

All of these options can be enabled on the Server console with the 'Set Config' command :

e.g

```
set config DEBUG_THREADID=1
```

However, if the server is restarted the parameter may not be active. For the parameter to remain active on a server it should be set in the NOTES.INI of the server.

The current setting for these can be confirmed on the server with the command :

```
show config <parameter>
```

e.g.

```
show config DEBUG_THREADID
```

or

```
show config DEBUG*
```

Note : These options will increase the level of logging from the server, in some cases these can SERIOUSLY affect the server performance, some care should be taken when enabling them to ensure that they are disabled once any tests or logging periods have completed.

DEBUG_THREADID

Although not specific to replication this is a very important debug option.

This displays each thread and task with a separate ID in the CONSOLE.LOG. With this enabled it is possible to follow each task through the console.log by the leading numbers on each line. This is very useful when there are multiple tasks being displayed on the console, even multiple replication events.

In this example the replication was run while an update of the View Indexes of the NAMES.NSF was in progress, each line in the CONSOLE.LOG is identified with the task ID number :

```
> rep testserver2 reptest
[2D08:0002-4E68] 09/09/2014 12:04:38   Informational, rebuilding view - user
specified REBUILD (reading C:\IBM\Domino\data\names.nsf view note
Title: '($Domains)
')
[5068:0002-184C] 09/09/2014 12:04:38   Informational, rebuilding view - user
specified REBUILD (reading C:\IBM\Domino\data\names.nsf view note
Title: '($Policies
ByGroup) ')
[46C4:0002-48FC] 09/09/2014 12:04:38   Database Replicator started
[46C4:0002-48FC] 09/09/2014 12:04:38   Replicator is set to Ignore Database
Quotas
[2D08:0002-4E68] 09/09/2014 12:04:38   Informational, rebuilding view - user
specified REBUILD (reading C:\IBM\Domino\data\names.nsf view note
Title: '($External
DomainNetworkAdresse')
[46C4:0002-48FC] 09/09/2014 12:04:38   Starting replication with server
TestServer2/TST
[5068:0002-184C] 09/09/2014 12:04:38   Informational, rebuilding view - user
specified REBUILD (reading C:\IBM\Domino\data\names.nsf view note
Title: '($Policies
Ext) ')
[46C4:0002-48FC] 09/09/2014 12:04:39   Finished replication with server
TestServer2/TST
[46C4:0002-48FC] 09/09/2014 12:04:39   Database Replicator shutdown
```

Here the replication thread is : [46C4:0002-48FC]

Since R8.0 DEBUG_THREADID is enabled by default on all server - however some administrators have disabled it by adding :

```
DEBUG_THREADID=0
```

into the NOTES.INI of the server.

Before starting any of the logging or debugging options `DEBUG_THREADID=1` should be enabled on the test server, either by adding "`DEBUG_THREADID=1`" to the `NOTES.INI` and restarting the server or by running the command "`set config DEBUG_THREADID=1`".

This option does not impact the performance of the server, but will slightly increase the size of the `CONSOLE.LOG`.

LOG_REPLICATION

This is a very useful and common replication logging option. By default (with no option configured) it runs at level '2' on a server. The options for this setting are :

`Log_Replication=<value>`

Description : Specifies whether the start and end of replication sessions are recorded in the Notes Log and displayed on the console:

`Log_Replication=0` (no replication logging)
`Log_Replication=1` (logs server replication events)
`Log_Replication=2` (adds logging of replication activity at the database level)
`Log_Replication=3` (adds logging of replication activity at the note level.)
`Log_Replication=4` (adds logging of replication activity at the field level)
`Log_Replication=5` (adds summary logging)

The above description of these options are what you will find in the Administrator Help document, however, in testing you will find that the replication output does not match the descriptions above :

- 0 & 1 Give the same output.
- 2 Adds some more data - and is the default logging level.
- 3 & 4 Are the same - source and target NoteID added to the output log but field level information is not included.
- 5 Some summary data on the replication is displayed.

From this, setting it to 2 can be useful to be able to see what databases are replicating. Setting it to 3 or 4 can be useful to get a list of all the documents (by NoteID) that are replicating.

Example output :

0 & 1

```
> set config log_replication=0
> rep testserver2 reptest.nsf
[2D60:0002-4E88] 09/09/2014 12:44:05 Database Replicator started
[2D60:0002-4E88] 09/09/2014 12:44:05 Replicator is set to Ignore Database
Quotas
[2D60:0002-4E88] 09/09/2014 12:44:05 Starting replication with server
TestServer2/TST
[2D60:0002-4E88] 09/09/2014 12:44:05 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[2D60:0002-4E88] 09/09/2014 12:44:05 Replicator added 3 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[2D60:0002-4E88] 09/09/2014 12:44:05 Pushing reptest.nsf to TestServer2/TST
reptest.nsf
[2D60:0002-4E88] 09/09/2014 12:44:05 Replicator added 3 document(s) to
TestServer2/TST reptest.nsf from reptest.nsf
[2D60:0002-4E88] 09/09/2014 12:44:05 Finished replication with server
TestServer2/TST
[2D60:0002-4E88] 09/09/2014 12:44:05 Database Replicator shutdown
```

Here you can see replication starting and stopping, the server being replicated with, the database being replicated and the number of documents replicated.

2

```
> set config log_replication=2
```

```

> rep testserver2 reptest.nsf
[3AB0:0002-2934] 09/09/2014 12:47:01 Database Replicator started
[3AB0:0002-2934] 09/09/2014 12:47:01 Replicator is set to Ignore Database
Quotas
[3AB0:0002-2934] 09/09/2014 12:47:01 Starting replication with server
TestServer2/TST
[3AB0:0002-2934] 09/09/2014 12:47:01 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[3AB0:0002-2934] 09/09/2014 12:47:01 Replicator added 3 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[3AB0:0002-2934] 09/09/2014 12:47:01 Replicator deleted 3 document(s) in
reptest.nsf from TestServer2/TST reptest.nsf
[3AB0:0002-2934] 09/09/2014 12:47:01 Replicator Modification Summary by
Type...
[3AB0:0002-2934] 09/09/2014 12:47:01 .....Data (3 Added, 3 Deleted)
[3AB0:0002-2934] 09/09/2014 12:47:01 Finished replication with server
TestServer2/TST
[3AB0:0002-2934] 09/09/2014 12:47:01 Database Replicator shutdown

```

Here the extra data added is :

```

Replicator Modification Summary by Type...
.....Data (3 Added, 3 Deleted)

```

3 & 4

```

> set config log_replication=4
> rep testserver2 reptest.nsf
[2E38:0002-3D24] 09/09/2014 12:51:53 Database Replicator started
[2E38:0002-3D24] 09/09/2014 12:51:53 Replicator is set to Ignore Database
Quotas
[2E38:0002-3D24] 09/09/2014 12:51:53 Starting replication with server
TestServer2/TST
[2E38:0002-3D24] 09/09/2014 12:51:53 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[2E38:0002-3D24] 09/09/2014 12:51:53 Deleted Data Note - DestID 2542, SrcID
2558
[2E38:0002-3D24] 09/09/2014 12:51:53 Deleted Data Note - DestID 2594, SrcID
2602
[2E38:0002-3D24] 09/09/2014 12:51:53 Deleted Data Note - DestID 2590, SrcID
2606
[2E38:0002-3D24] 09/09/2014 12:51:53 Added Data Note - DestID 2626, SrcID 2634
[2E38:0002-3D24] 09/09/2014 12:51:53 Added Data Note - DestID 2630, SrcID 2630
[2E38:0002-3D24] 09/09/2014 12:51:53 Added Data Note - DestID 2634, SrcID 2626
[2E38:0002-3D24] 09/09/2014 12:51:53 Replicator added 3 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[2E38:0002-3D24] 09/09/2014 12:51:53 Replicator deleted 3 document(s) in
reptest.nsf from TestServer2/TST reptest.nsf
[2E38:0002-3D24] 09/09/2014 12:51:53 Replicator Modification Summary by
Type...
[2E38:0002-3D24] 09/09/2014 12:51:53 .....Data (3 Added, 3 Deleted)
[2E38:0002-3D24] 09/09/2014 12:51:53 Pushing reptest.nsf to TestServer2/TST
reptest.nsf
[2E38:0002-3D24] 09/09/2014 12:51:53 Finished replication with server
TestServer2/TST
[2E38:0002-3D24] 09/09/2014 12:51:53 Database Replicator shutdown

```

Here you can see the NoteID of the documents added and deleted through replication :

```

[2E38:0002-3D24] 09/09/2014 12:51:53 Deleted Data Note - DestID 2590, SrcID
2606
[2E38:0002-3D24] 09/09/2014 12:51:53 Added Data Note - DestID 2626, SrcID 2634

```

5

```
> set config log_replication=5
> rep testserver2 reptest.nsf
[4A58:0002-2370] 09/09/2014 12:54:36 Database Replicator started
[4A58:0002-2370] 09/09/2014 12:54:36 Replicator is set to Ignore Database
Quotas
[4A58:0002-2370] 09/09/2014 12:54:36 Starting replication with server
TestServer2/TST
[4A58:0002-2370] 09/09/2014 12:54:36 Pulling reptest.nsf from TestServer2/TST
reptest.nsf
[4A58:0002-2370] 09/09/2014 12:54:36 Deleted Data Note - DestID 2614, SrcID
2622
[4A58:0002-2370] 09/09/2014 12:54:36 Deleted Data Note - DestID 2630, SrcID
2630
[4A58:0002-2370] 09/09/2014 12:54:36 Deleted Data Note - DestID 2634, SrcID
2626
[4A58:0002-2370] 09/09/2014 12:54:36 Added Data Note - DestID 2638, SrcID 2650
[4A58:0002-2370] 09/09/2014 12:54:36 Added Data Note - DestID 2642, SrcID 2646
[4A58:0002-2370] 09/09/2014 12:54:36 Added Data Note - DestID 2646, SrcID 2638
[4A58:0002-2370] 09/09/2014 12:54:36 Added Data Note - DestID 2650, SrcID 2642
[4A58:0002-2370] 09/09/2014 12:54:36 Replicator added 4 document(s) to
reptest.nsf from TestServer2/TST reptest.nsf
[4A58:0002-2370] 09/09/2014 12:54:36 Replicator deleted 3 document(s) in
reptest.nsf from TestServer2/TST reptest.nsf
[4A58:0002-2370] 09/09/2014 12:54:36 Replicator Modification Summary by
Type...
[4A58:0002-2370] 09/09/2014 12:54:36 .....Data (4 Added, 3 Deleted)
[4A58:0002-2370] 09/09/2014 12:54:36 0 Bytes Saved from Field Level
Replication
[4A58:0002-2370] 09/09/2014 12:54:36 0 Bytes Saved from Field Level
Replication
[4A58:0002-2370] 09/09/2014 12:54:36 Finished replication with server
TestServer2/TST
[4A58:0002-2370] 09/09/2014 12:54:36 Database Replicator shutdown
```

The extra information added by the level '5' option is :

```
[4A58:0002-2370] 09/09/2014 12:54:36 0 Bytes Saved from Field Level
Replication
[4A58:0002-2370] 09/09/2014 12:54:36 0 Bytes Saved from Field Level
Replication
```

In most situations this is unrequired.

Summary for LOG_REPLICATION :

From this, setting it to 2 can be useful to be able to see what databases are replicating. Setting it to 3 or 4 can be useful to get a list of all the documents (by NoteID) that are replicating.

DEBUG_REPL & DEBUG_REPL_ALL

When replication runs it creates a list of all documents changed since the last complete entry in the Replication History. It will then review this list comparing each document with its counterpart in the replica database. If a document is to be replicated it will replicate it.

Setting :

```
DEBUG_REPL=1
DEBUG_REPL_ALL=3
```

either in the NOTES.INI or using 'Set Config' will print out the list of documents to be replicated and their comparison to the CONSOLE.LOG.

This can be very useful in situations where one or more documents in a database are NOT replicating as you

can review :

If the document is actually being considered for replication.

If the document is being considered the reason why it was not replicating will be displayed.

All documents are identified with their UNID (Universal Identification number)

Note : This debug will only output for documents that are listed to be replicated, that is documents changed since the last entry in the Replication History. To get a complete listing of all documents and if they are replicating the Replication History must be cleared.

Note 2 : This debug option can be EXTREMELY verbose. It will list every document to be replicate, then each document again when it is being compared for replicated and then repeat this again for the reverse replication. In a situation where every document in a database is being checked on a large database this will SERIOUSLY IMPACT THE SERVER'S PERFORMANCE. It is strongly recommended that it is disabled after any testing is complete.

Example output :

In this example - 1 new documents is added to one replica, 1 existing document was edited and 1 was deleted.

The ThreadID data has been removed and the font size reduced :

```
> sh config debug*
[2AF8:0008-4154] DEBUG_REPL=1
[2AF8:0008-4154] DEBUG_REPL_ALL=3
> rep TestServer2 reptest.nsf
ReplicateNotefile: EndTime=16/05/2015 02:54:18
ReplicateNotefile: ReplOpenDatabases->0x0=No error, stop=1
ReplicateNotefile: EndTime=16/05/2015 02:54:18
ReplicateNotes: NSFNoteReplicate()->OK, TotalNotes=3, NotesToCopy=2
ReplicateNotes: NSFNoteReplicate()->OK, TotalNotes=3, NotesToCopy=1
ReplicateNotes: NSFNoteReplicate()->OK, TotalNotes=3, NotesToCopy=0
ReplicateFolders: NSFStartFolderReplSource->0x3AE=Folders in database are up to date
Replicate Unreads: Src reptest.nsf Dst CN=TestServer2/O=TST!!reptest.nsf, SR 1, DR 1 SOC 1 DOC:1
Since=09/09/2014 13:43:17
09/09/2014 13:54:18 Database Replicator started
09/09/2014 13:54:18 Replicator is set to Ignore Database Quotas
09/09/2014 13:54:18 REPLICA: Replicator Debug Enabled
09/09/2014 13:54:18 REPLICA: Watching for ALL notes>
ReplUpdateCacheHistories: NSFDbWriteReplHistory(dest,CN=TestServer2/O=TST!!reptest.nsf)->This
replication was complete!
NSFDbWriteReplHistory> localWSInitiated: No; DestIsRemote = True; WS Flag: Off
09/09/2014 13:54:18 Starting replication with server TestServer2/TST
09/09/2014 13:54:18 REPLICA: Time on Local system: 09/09/2014 13:54:18
09/09/2014 13:54:18 REPLICA: Time on Remote system: 09/09/2014 13:54:06
09/09/2014 13:54:18 REPLICA: *** Starting replication of source database CN=TestServer2/O=TST!!
reptest.nsf to destination reptest.nsf
09/09/2014 13:54:18 REPLICA: *** Starting replication of source database reptest.nsf to
destination CN=TestServer2/O=TST!!reptest.nsf
09/09/2014 13:54:18 Pushing reptest.nsf to TestServer2/TST reptest.nsf
09/09/2014 13:54:18 REPLICA: Building SOURCE list (reptest.nsf) since 09/09/2014 13:43:17;
cutoff=29/10/2012 07:01:07; flags=0xC005011; SourceClass=0x77DF; AuxClass=0x0000
09/09/2014 13:54:18 REPLICA: ...Added note to Source list (UNID OF23218ACF:9852ACBD-
ON80257D4E:0046DF04;Note ID 0xA76;Class 0x0001)
09/09/2014 13:54:18 REPLICA: ...Added note to Source list (UNID OF526EBA69:552EAD01-
ON80257D4E:0045BBD9;Note ID 0xA5E;Class 0x8001)
09/09/2014 13:54:18 REPLICA: ...Added note to Source list (UNID OFDF79B525:6DB2ACC7-
ON80257D4E:0045BE4D;Note ID 0xA62;Class 0x0001)
09/09/2014 13:54:18 REPLICA: Source DB has newer note (OF23218ACF:9852ACBD-
ON80257D4E:0046DF04;SrcNote ID 0xA76;DestNote ID 0x0;Class 0x0001)
09/09/2014 13:54:18 REPLICA: ... (con't) SrcSeqNo = 1, DestSeqNo = 0; SrcSeqTime = 09/09/2014
13:54:06, DestSeqTime =
09/09/2014 13:54:18 REPLICA: Source DB has newer note (OF526EBA69:552EAD01-
ON80257D4E:0045BBD9;SrcNote ID 0xA5E;DestNote ID 0xA5E;Class 0x0001)
09/09/2014 13:54:18 REPLICA: ... (con't) SrcSeqNo = 3, DestSeqNo = 2; SrcSeqTime = 09/09/2014
13:53:51, DestSeqTime = 09/09/2014 13:42:28
09/09/2014 13:54:18 REPLICA: Source DB has newer note (OFDF79B525:6DB2ACC7-
ON80257D4E:0045BE4D;SrcNote ID 0xA62;DestNote ID 0xA66;Class 0x0001)
09/09/2014 13:54:18 REPLICA: ... (con't) SrcSeqNo = 3, DestSeqNo = 2; SrcSeqTime = 09/09/2014
13:53:58, DestSeqTime = 09/09/2014 13:42:24
09/09/2014 13:54:18 REPLICA: ...Skipping note in Destinationlist (UNID OF23218ACF:9852ACBD-
ON80257D4E:0046DF04 SrcNote ID 0xA76;DestNote ID0x0;Class 0x0001);
09/09/2014 13:54:18 REPLICA: ... (con't) Latest note is already in destination
```

```

09/09/2014 13:54:18 REPLICA: ...Skipping note in Destinationlist (UNID OF526EBA69:552EAD01-
ON80257D4E:0045BBD9 SrcNote ID 0xA5E;DestNote ID0xA5E;Class 0x0001);
09/09/2014 13:54:18 REPLICA: ... (con't) Latest note is already in destination
09/09/2014 13:54:18 REPLICA: ...Skipping note in Destinationlist (UNID OFDF79B525:6DB2ACC7-
ON80257D4E:0045BE4D SrcNote ID 0xA62;DestNote ID0xA66;Class 0x0001);
09/09/2014 13:54:18 REPLICA: ... (con't) Latest note is already in destination
09/09/2014 13:54:18 REPLICA: Folders in database TestServer2/TST reptest.nsf are up to date with
database reptest.nsf.
09/09/2014 13:54:18 Replicator added 1 document(s) to TestServer2/TST reptest.nsf from reptest.nsf
09/09/2014 13:54:18 Replicator updated 1 document(s) in TestServer2/TST reptest.nsf from
reptest.nsf
09/09/2014 13:54:18 Replicator deleted 1 document(s) in TestServer2/TST reptest.nsf from
reptest.nsf
09/09/2014 13:54:18 Replicator Modification Summary by Type...
09/09/2014 13:54:18 .....Data (1 Added, 1 Deleted, 1 Updated)
09/09/2014 13:54:18 Finished replication with server TestServer2/TST
09/09/2014 13:54:18 Database Replicator shutdown

```

The use of this output will be reviewed in the Troubleshooting Common Scenarios chapter.

DEBUG_REPL_TIME

In situations where replication is taking longer than expected enabling "DEBUG_REPL_TIME" can highlight where there may be an issue.

For each major replicator action, this parameter causes a message to be displayed on the Notes server console indicating the amount of time the action took. This information is NOT logged to LOG.NSF, because logging the message would adversely affect the timing of each action.

Example output :

```

> set config DEBUG_REPL_TIME=1
> rep TestServer2 reptest.nsf
REPLTIME: Opening source directory on CN=TestServer2/O=TST and building nameslist took 51 msec
REPLTIME: Building index of all (or specified) files locally took 2 msec ( 1 files )
REPLTIME: Searching the remote server for replicas of local files took 0 msec ( 1 Total Files
examined)
REPLTIME: Calling ReplicateNotefile for reptest.nsf from CN=TestServer2/O=TST!!reptest.nsf after 1
msec
REPLTIME: Initializing took 0 msec
REPLTIME: Reading replication history for destination, reptest.nsf, took 2 msec
REPLTIME: Reading replication history for source, CN=TestServer2/O=TST!!reptest.nsf, took 3 msec
REPLTIME: Source and Destinations Open and history compared after 0 msec
REPLTIME: Searching source database (Since = 09/09/2014 15:08:30)...
REPLTIME: Search of source database (Since = 09/09/2014 15:08:30) took 4 msec
REPLTIME: Searching destination database...
REPLTIME: Search of destination database took 2 msec
REPLTIME: Beginning traversal of index and actual replication of notes...
REPLTIME: Traversing index and actual replication of notes took 1 msec
REPLTIME: Replicating folders...
REPLTIME: Replication of folders took 1 msec
REPLTIME: Replicating unread marks...
REPLTIME: Replication of unread marks took 2 msec
REPLTIME: Writing source replication history took 1 msec
REPLTIME: Writing destination replication history took 1 msec
REPLTIME: Logging of info for this notefile took 0 msec
REPLTIME: Calling ReplicateNotefile for CN=TestServer2/O=TST!!reptest.nsf from reptest.nsf after 1
msec
REPLTIME: Initializing took 0 msec
REPLTIME: Opening and reading replication history for source, reptest.nsf, took 1 msec
REPLTIME: Opening and reading replication history for destination, CN=TestServer2/O=TST!!
reptest.nsf, took 3 msec
REPLTIME: Source and Destinations Open and history compared after 1 msec
REPLTIME: Reading Selective Replication formula for CN=TestServer2/O=TST!!reptest.nsf took 0 msec
REPLTIME: Searching source database (Since = 09/09/2014 13:54:18)...
REPLTIME: Search of source database (Since = 09/09/2014 13:54:18) took 1 msec
REPLTIME: Searching destination database...
REPLTIME: Search of destination database took 1 msec
REPLTIME: Beginning traversal of index and actual replication of notes...
REPLTIME: Traversing index and actual replication of notes took 5 msec
REPLTIME: Replicating folders...
REPLTIME: Replication of folders took 1 msec
REPLTIME: Replicating unread marks...

```

```
REPLTIME: Replication of unread marks took 19 msec
REPLTIME: Writing destination replication history took 1 msec
REPLTIME: Writing source replication history took 0 msec
REPLTIME: Logging of info for this notefile took 0 msec
REPLTIME: Updating collections took 0 msec
REPLTIME: Replication with Server CN=TestServer2/O=TST completed(1)
09/09/2014 14:14:28 Database Replicator started
09/09/2014 14:14:28 Replicator is set to Ignore Database Quotas
09/09/2014 14:14:28 REPLICA: Replicator Time Monitoring Enabled - Capture using notes.ini
parameter DEBUG_OUTFILE=debug.txt
09/09/2014 14:14:28 Starting replication with server TestServer2/TST
09/09/2014 14:14:29 Pulling reptest.nsf from TestServer2/TST reptest.nsf
09/09/2014 14:14:29 Pushing reptest.nsf to TestServer2/TST reptest.nsf
09/09/2014 14:14:29 Replicator added 1 document(s) to TestServer2/TST reptest.nsf from reptest.nsf
09/09/2014 14:14:29 Replicator deleted 2 document(s) in TestServer2/TST reptest.nsf from
reptest.nsf
09/09/2014 14:14:29 Replicator Modification Summary by Type...
09/09/2014 14:14:29 .....Data (1 Added, 2 Deleted)
09/09/2014 14:14:29 Finished replication with server TestServer2/TST
09/09/2014 14:14:29 Database Replicator shutdown>
```

Chapter 8 : Troubleshooting

Here we will be looking at several common problems that happen with replication and how they can be investigated and resolved.

- 8.1 Replication not happening for all databases
- 8.2 Replication not happening for one or a few databases
- 8.3 Replication not happening for one or a few documents
- 8.4 Scheduled replication not working
- 8.5 Database size differences between replicas
- 8.6 Document count matches but documents not displaying
- 8.7 Corruption
- 8.8 Performance
- 8.9 Deleted documents re-appearing
- 8.10 Database encryption
- 8.11 System databases not replicating
- 8.12 Documents being removed by replication

8.1 Replication not happening for all databases

In this situation the first step is to verify that replication is failing for all databases.

As replication is failing for all databases it is unlikely that the issue is in the databases themselves.

If replication can be run successfully for any database between the servers then it will confirm that, at least for some of the time, the Network and server authentication is working correctly.

If the replication is confirmed to be failing for all databases on a server then check the following :

- 1) Are both servers actually up and running normally? If not, restart the missing server or servers.
- 2) Is the Replicator task available on the server? If the command 'Show tasks' is entered is there an output for the replicator?

```
[2AF8:0008-4154] Replicator Idle
```

Is the task running or Idle?

If the task is not listed review the 'ServerTasks' parameter in the NOTES.INI, it should be added there if it is missing and the server restarted :

```
ServerTasks=Update,Replica,Router,AMgr,AdminP,CalConn,Sched,HTTP,LDAP,RnRMgr
```

- 3) If the source and target servers are in different Domino Domains :
 - Are the servers correctly cross-certified to be able to communicate?
 - Is there a connection document in place for communication between the source and target.

If either of these is missing or incorrectly configured they would need to be addressed. The Domains or the individual servers will need to be Cross-Certified.

If there is no connection document one will need to be created, if one exists but is incorrectly configured it will need to be updated with the correct data.

- 4) If the source and target are in the same Domain or if the servers are correctly Cross-certified with a Connection Document then it is possible that there is a network issue between the servers. This could be a problem with DNS settings or the Host file or any Firewall that may be in place between the servers. To test these :

- 4.1 Use the PING command and the DNS title of the server and the IP address of the server. This command runs from a command prompt.

If the IP address of the server connects but the title of the server does not then you should review your DNS settings, it is likely that the incorrect IP address is being used and preventing the servers from connecting.

The computers HOST file should also be reviewed to confirm it is correctly configured.

4.2 Use the server command 'trace [servername]' and 'trace [IP address]' to check the connection.

If this command fails to connect then the DNS settings and Host file should be reviewed again.

This command is run from the Server Console itself. If this fails then, as with PING, the DNS settings should be reviewed and the network connections between the servers.

A firewall could also be blocking required Domino network communications.

4.3 Use Telnet to PORT 1352 to check the connection using the 1352 port.

This is useful as in most connections Domino Servers use Port 1352, if this is being blocked then all communication including replication is likely to fail between the servers.

If any of these tests fail then there is a network configuration issue between the servers that needs to be addressed.

Any Firewall should be reviewed to confirm that Port 1352 is open for Domino Servers to interconnect.

If the Network itself is failing then this will need to be addressed to allow for replication.

8.2 Replication not happening for one or a few databases

Replication is seen to be working correctly on several other databases - confirming that the connection and network settings are correct.

However the replication of one or a few databases is still not running normally.

As replication is working normally for most databases it is likely that the cause of these issues is within the non-replicating databases.

1) During replication review any errors or warnings displayed in the CONSOLE.LOG or the LOG.NSF these can often point to where the issue is :

```
> rep TestServer2 reptest4.nsf> rep TestServer2 reptest4.nsf
[219C:0002-4D58] 09/09/2014 15:00:04 Database Replicator started
[219C:0002-4D58] 09/09/2014 15:00:04 Replicator is set to Ignore Database
Quotas
[219C:0002-4D58] 09/09/2014 15:00:04 Starting replication with server
TestServer2/TST
[219C:0002-4D58] 09/09/2014 15:00:04 Unable to replicate reptest4.nsf: File
does not exist
[219C:0002-4D58] 09/09/2014 15:00:04 Unable to replicate with server
TestServer2/TST: No databases to replicate on that server
[219C:0002-4D58] 09/09/2014 15:00:04 Database Replicator shutdown
```

In this example there was no matching replica for the database on the target server. To resolve this a new replica would be needed on the target server.

Other errors will point to where the problem is.

2) Verify that the Replication IDs of the source and target databases are the same. This can be seen on the Database properties dialog box on the Info tab (second from the left). This ID number must be the same between replicas for replication to occur. The Title and file name of the database do not need to be the same.

A regular cause of this is the use of the File\Database\New Copy option in Notes. This will create a new copy of the database but this will have a new replica ID and will need to be replaced with a proper replica of the database.

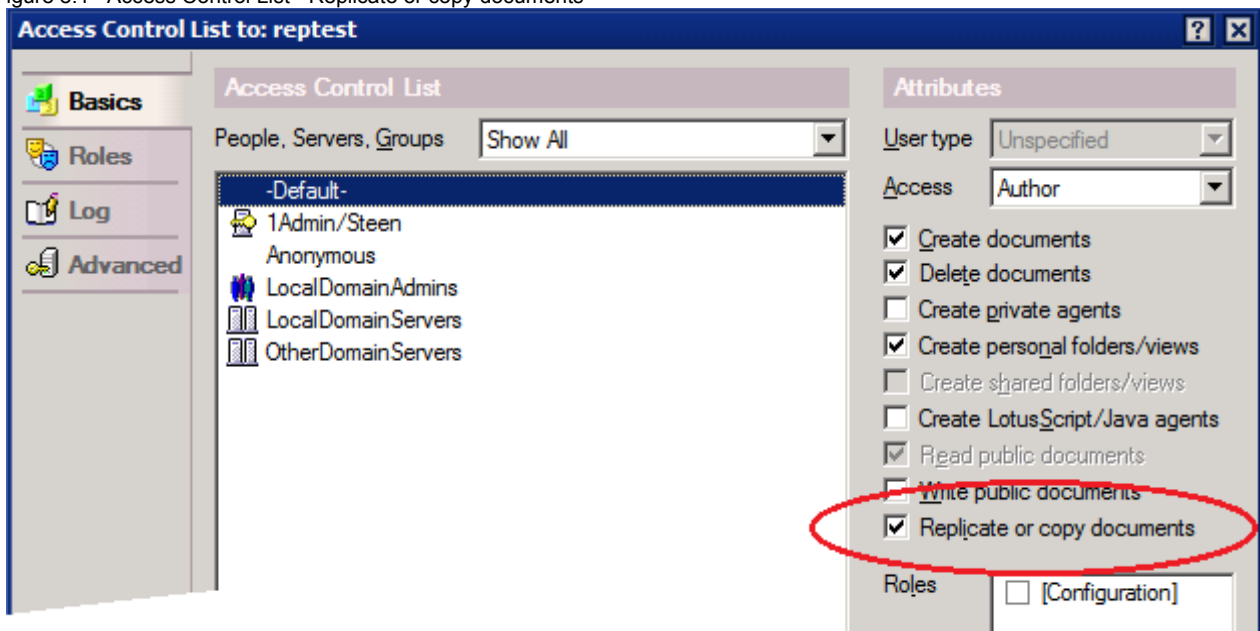
3) Check the ACL of the source and target databases.

Check the access level of each Server in the ACL of the database. For replication to work a server needs a minimum of Reader access for the source database and Editor for the Target, this will replicate standard documents but only in one direction and will not replicate design elements. Also verify if there are any roles being used on the database as it is possible that the server will need to be included in them..

For total replication of all settings, including the ACL and the Design documents then Manager access for the source and target servers will be needed.

Also, There is an access setting 'Replicate or copy document' this will need to be checked for the servers replicating.

Figure 8.1 - Access Control List - Replicate or copy documents



NOTE : If a server is included in a group with full access to a database BUT also listed itself in the ACL then the settings under its own name will be used. For Example, if a server is included in the group LocalDomainServers as manager but is listed in the ACL under its own name as Reader only then it will only have Reader access to the database.

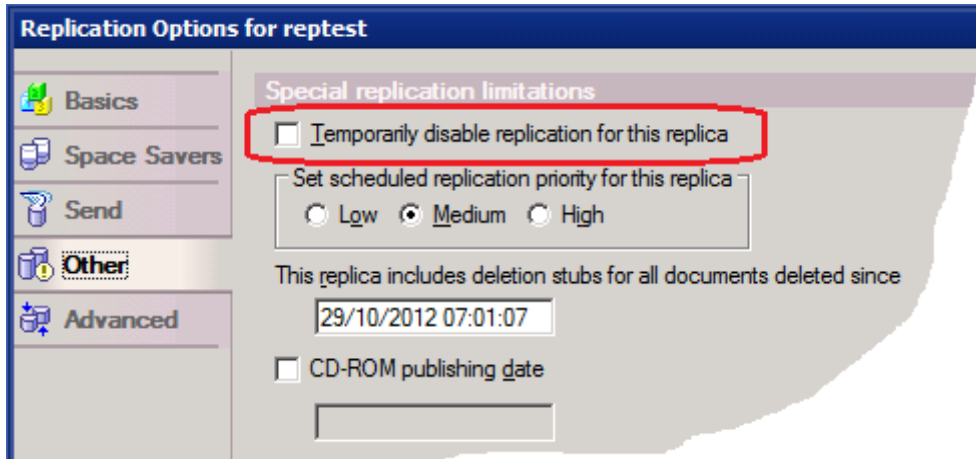
4) Verify that replication is not being prevented in the Replication Options of the database.

As explained in the Replication Configurations chapter - the Other tab of the Replication Options has a setting :

Temporarily disable replication for this replica

If checked this option will disable all replication for this replica. Documents will not replicate to or from the database.

Figure 8.2 - Replication Options - Other tab, Temporarily disable replication for this replica



When enabled this message will be displayed if replication is attempted :

```
[2130:0002-2124] 04/09/2014 13:45:31 Replication is disabled for reptest.nsf
```

5) Check to see if it is possible to create a new complete replica from the source and target databases. If one fails then it may be corrupt and should have maintenance run on it or it could be replaced with a replica from the working server.

8.3 Replication not happening for one or a few documents

In this situation some documents are not replicating between replicas or the document counts are different.

1) The first thing to check is the actual document count of both the source and the target servers. The best way to do this is with the 'show database [database path and file name]' command.

This will output the number of documents, deletions, design elements and the View Index size of the database. The document count data here should be compared between the replicas :

```
> sh database reptest.nsf
[2AF8:0008-4154] reptest

[2AF8:0008-4154] Document Type                Live        Deleted
[2AF8:0008-4154] Documents                  3           93
[2AF8:0008-4154] Info                      1           0
[2AF8:0008-4154] Form                      1,197       1
[2AF8:0008-4154] View                      21          154
[2AF8:0008-4154] Icon                      1           0
[2AF8:0008-4154] Design                   1           0
[2AF8:0008-4154] ACL                      1           0
[2AF8:0008-4154] Help Index                   0           0
[2AF8:0008-4154] Help                      1           0
[2AF8:0008-4154] Macro                     29          1
[2AF8:0008-4154] Shared Field                 10          0
[2AF8:0008-4154] Repl. Formula                 0           0
[2AF8:0008-4154] Private Design                0           0
[2AF8:0008-4154] Unknown                      0           1

[2AF8:0008-4154] View sizes                Bytes
[2AF8:0008-4154] By Author                    0
[2AF8:0008-4154] ($Profiles)                42,192
[2AF8:0008-4154] (Default View)                 0
[2AF8:0008-4154] All Documents                100,704
[2AF8:0008-4154] By Category                    0
```

[2AF8:0008-4154]	My Favorites	0
[2AF8:0008-4154]	Author Profiles	0
[2AF8:0008-4154]	(LookupAuthorProfiles)	42,192
[2AF8:0008-4154]	By Alternate Name	0
[2AF8:0008-4154]	(LookupInterestProfiles)	0
[2AF8:0008-4154]	(ProfileThreads)	0
[2AF8:0008-4154]	(FlatCategory)	0
[2AF8:0008-4154]	(\$xpCategoryCloud)	0
[2AF8:0008-4154]	(\$xpAuthorCloud)	0
[2AF8:0008-4154]	(\$xpConfigProfile)	0
[2AF8:0008-4154]	(\$xpAuthorProfiles)	0
[2AF8:0008-4154]	(\$xpAuthorPosts)	0
[2AF8:0008-4154]	(\$xpAllDocuments)	0
[2AF8:0008-4154]	(\$xpMostRecent)	0
[2AF8:0008-4154]	(\$xpCategorized)	0
[2AF8:0008-4154]	(\$xpByAuthor)	0

Using this you can confirm that there is are documents not replicating between the replicas.

2) Clear the replication history and run replication again, in most cases this will cause the outstanding documents to replicate.

As shown above Replication will review the Replication History and only replicate documents that have been edited since the last complete replication.

Clearing the Replication History will result in the Replicator reviewing every document in the database for replication. This usually will cause any mismatches to be replicated.

3) Check to see if the Documents that are not replicating have Reader Fields in use. If they do then the source and target servers should be added to the Reader fields, either in a group or under their own names.

If the servers are not included in the Reader Fields then the servers will not have access to the documents and will not be able to replicate them.

Note : If using the debug options DEBUG_REPL=1 and DEBUG_REPL_ALL=3 (as in section 6 below) and the missing documents can be identified, then it is possible to review the reason for the replication failure. In the case of Reader fields this will be listed as :

Does not match Selective Replication formula

4) Review the Replication Options. Some documents may be being prevented from replication as there is a selective replication configuration in place. These would be configured in the Space Savers and Advances tab of the Replication Options.

5) If the documents that are not replicating can be identified verify that they can be opened from the database they are on. If they are corrupt and cannot be opened then replication will not work and the corruption needs to be addressed. More on this in the Troubleshooting Corruption section.

6) If some documents are still not replicating you can enable the following debug options on the server where the replication is initiated :

```
DEBUG_REPL=1
DEBUG_REPL_ALL=3
```

Both these are needed to enable this debug.

The replication history should be cleared and then a manual replication initiated between the servers.

Note : This debug is EXTREMELY verbose and will affect the performance of a server.

As the output from the debug is so large it is a good idea to identify which documents are not replicating. Then the references to these can be searched through the debug output. Each document may have several entries so it is important to review the entire debug output and not just the first matches.

If the non-replicating documents are not known it may be necessary to create a complete list of all documents from each database and compare them for differences. Then extract the UNID for the missing documents and review the debug output for entries referring to these documents.

More information on this debug can be found in Chapter 7 7. Logging and Debug options in the section "DEBUG_REPL & DEBUG_REPL_ALL"

7) If the number of documents not replicating is low or even one then another useful quick solution is to 'Dummy Edit' the document. This resets the modified by field and should replicate the document in the next replication.

To do this :

Open the document.

Edit the document.

Enter a single space on some field in the document.

Remove the added space.

Save the document.

8.4 Scheduled replication not working

If replication is configured to run on a schedule but is failing.

1) Check to see if a manually initiated replication will complete. If it completes then continue the investigation as outlined here. If it does not replicate continue depending on the output from the replication and investigate accordingly - depending on if one or all databases are not replicating.

2) For server based replication enter the server console command :

show schedule

```
> show schedule
[05EC:0008-4C80] Scheduled                               Type           Next
schedule

[05EC:0008-4C80] TestServer2/TST                       Replication
09/09/2014 16:30:00
[05EC:0008-4C80] TestServer2/TST                       Mail Routing
09/09/2014 16:30:00
```

This should output the next scheduled replication (and other) events on the server. If this replication is not listed then review the Connection Documents, ensure that they are correctly configured for your environment, that the database is included in the databases to be replicated and that they are Enabled.

3) If the scheduled replication is still failing, shutdown the server. Navigate to the DATA folder and rename the file "jobsched.njf" to "jobsched.bak".

Then restart the server. The file will be recreated.

This file contains the schedule information from the Connection Documents and the Program documents from the server, occasionally it may need to be re-created due to corruption.

8.5 Database size differences between replicas

This is normal between two replicas even when replication is otherwise working correctly with the same number of documents on all replicas.

Areas that can cause this are :

White space	Replicas can have different amounts of white space unused within them. Running a normal compact on them will not affect the file size. Running a Compact -B (this is case sensitive) will reduce the size of the database. If this is run on all replicas it can bring their sizes closer together.
View Indexes	The document display in each database is stored in View Indexes. Although View design (the criteria for what will be displayed) will replicate between replicas the View Indexes themselves do not replicate and these can be significantly large. Different replicas may have different needs for the Views indexes, where some View Indexes may not be used on one server but will be used on another. The server console command 'show database DatabaseFilename' will display all the View indexes and their sizes within the database so you can compare them to show any size differences.
Deletion Stubs	Deletion stubs are generated so that when a document is deleted the deletion of the document can replicate and the document is deleted on the replica database. These can be purged from the database and are done so after a set amount of time. This can be different from one replica to another, also if a deletion stub purge interval is for 3 days but the stub is replicated after one day the stub may remain on the replica for three days after the replicated. Although replication stubs are very small they can add up quickly if there has been a lot of deletions and can cause size differences between replicas. Again the server console command 'show database DatabaseFilename' will display the number of documents and deletions within a database so can be used to compare the differences.
Soft Deletions	<p>If soft deletions are enabled on a database when a document is deleted it will be moved to the Trash folder where it will be Hard deleted after a configurable amount of time (usually 48 hours). This is to allow users to recover accidentally deleted documents.</p> <p>If the document is added to the database and soft deleted to the Trash folder before replication to the replica then the document itself will not replicate from the Trash folder. If the document is replicated before being soft deleted then the move to the Trash folder is replicated.</p> <p>Once the document is Hard deleted it is converted to a normal Deletion Stub which does replicate normally.</p>
DAOS	If one replica is DAOS enabled then the attachments will be stored outside of the database, if the other replica is not DAOS enabled then the attachments will be stored in the database causing it to be bigger. Even if both are DAOS enabled then the size can be different as the minimum attachment size for DAOS to run on the attachment could be different between the servers.

8.6 Document count matches but documents not displaying

All documents appear to have replicated normally, however they are not displaying on the target database.

This can happen if there is an issue with the View Indexes on the target database or if they just have not updated.

The best option in this situation is to run the command :

```
load update -r [database path and file name ]
```

This will rebuild all the View Indexes in the database and display all documents that match the View design selection criteria.

If there is an 'All Documents' view available within the database check to see if the documents are visible there.

The View designs themselves may not be configured to display the documents in question. If the documents never displayed correctly then the View designs should be reviewed to see that the selection formula in use should select the missing documents.

8.7 Corruption

Any file can become corrupt - this is not just limited to Domino Databases but can happen any file on a computer.

Unfortunately as Domino databases are high usage they can become corrupt from a number of different sources.

This is usually caused by influences external to the Domino Server, including -
Crashes
Hardware failure
External software accessing the databases directly

All of these factors, and others, can lead to database corruption.

Corruption can lead to total database failure, or just minor issues in a database and can often be resolved using the 'load updll', 'load compact' and 'load fixup' commands.

From a replication perspective it can manifest in many ways. If the database is completely corrupt then replication will not be possible.

If just some of the documents are corrupt then they will not replicate.

A simple test for corruption affecting replication is to create a new replica of a database (either locally or on a different server) and comparing the document counts. If they are different, and there are no other causes for documents not to replicate, it is likely to be because of corruption.

A useful point with Replication is that Corruption itself does not replicate. In this way, if you can replicate a database the new resulting replica will be less likely to be corrupt. This is due to all the internal tables and properties being rebuilt and that the View Indexes are not replicated.

Another feature of the Compact command is the '-replica' switch. This will create a new replica of a database and replace the original with the replica. This can be a useful fix for low level corruption, however it is recommended to have a backup of the original database before running this.

The compact replica option can be initiated with :

```
load compact -replica [database path and file name]
```

8.8 Performance

If replication appears to be running slowly you can try the following :

1) Check to see if the issue is specific to one database. If it is check the LOG.NSF to see if a larger than normal amount of data is being replicated.

Run maintenance (fixup, compact and updll) on the database.

Enable DEBUG_REPL_TIME and review the times for the replication and see if there is anything taking longer than expected.

2) If all databases are replicating slowly.

Review the amount of data being replicated - this can be seen in the LOG.NSF

Review the servers performance - it may be overloaded.

Review the hardware (including the network) - this may be failing or overloaded.

Additionally the replication process is multi-processed, which means more than one replicator can run at the same time.

To enable this :

1. Use the NOTES.INI setting :

Replicators=n

where n represents the number of replicators that can be run concurrently.

or

2. Add additional "Replica" tasks to the Server Tasks= line of the NOTES.INI

ServerTasks=replica,replica,replica,update,stats

There is no limit to the number of replicator processes that can be run concurrently. The more replicator processes allowed, the greater potential impact on server performance.

8.9 Deleted documents re-appearing

This can occur if the replication interval is longer than the configured purge interval of the deletion stubs.

What can happen is this :

- ◆ The document is marked as deleted leaving a deletion stub of the document.
- ◆ The purge interval then removes the deletion stub.
- ◆ The replicator then sees that there is a document on one database that is not on the other and replicates it back.

To resolve this the replication interval should be set below that of the deletion stub purge interval so that the deletion stub replicates before it is purged.

This also often occurs when an old backup of a database is restored, replication will access the database and replicate old documents that were deleted from the original database. It is recommended to isolate restored replicas from other production replicas to prevent this, unless the older documents are needed.

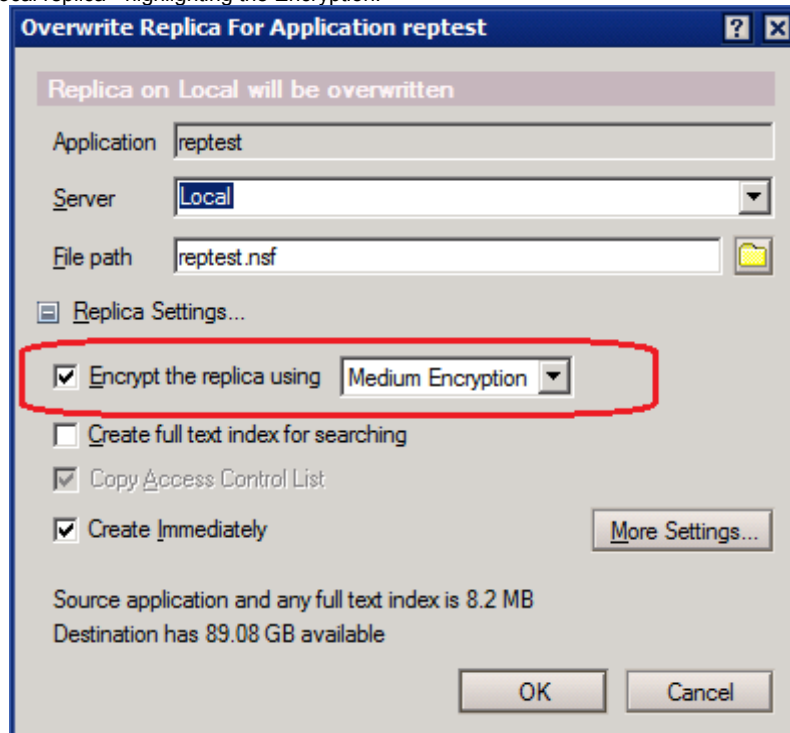
The new replication setting PIRC was introduced to help prevent this scenario. More information on this setting can be found in the Replication Configuration chapter.

8.10 Database encryption

It is possible to encrypt a database using a Notes ID. This then results in only that ID being able to access the database.

This can be an issue for replication as when a new local replica of a database is created the option to encrypt the new replica is enabled by default.

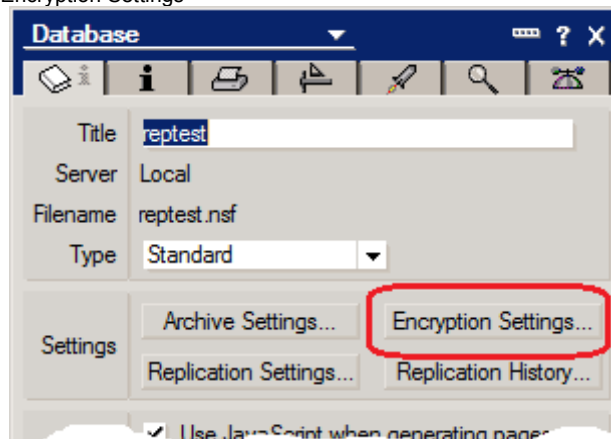
Figure 8.3 - Create new local replica - highlighting the Encryption.



Once the new replica is created it will only be accessible by the ID that created it. This can be a problem if the database is then placed on a server. The server itself will not be able to access the database.

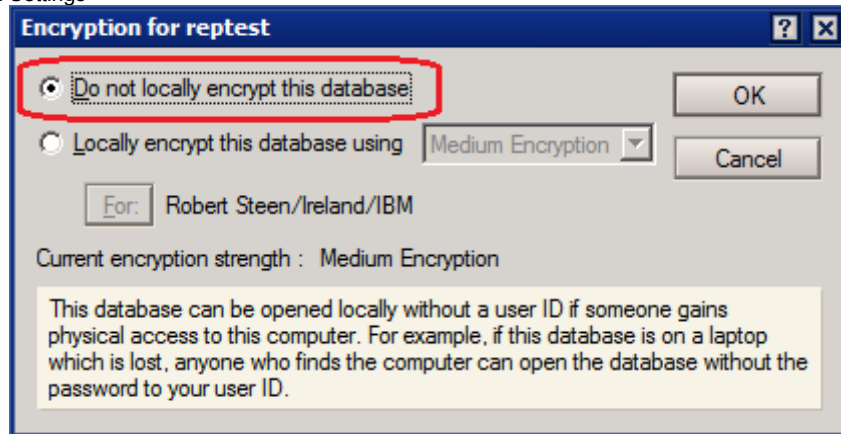
- To remove the encryption :
- Using the ID that encrypted the database :
- Open the database properties.
- On the first tab click on 'Encryption Settings' :

Figure 8.4 - Database properties - Encryption Settings



This will then open the "Encryption for" dialog :

Figure 8.5 - Encryption Settings



Click on "Do not locally encrypt this database" and then OK.

Note : The database will need to be compacted for this to be implemented.

Another factor with encryption worth mentioning is that it is possible to encrypt a database with a selected ID file not just the ID the user is working with at that time. In this way a database can be encrypted using the Server's ID file.

Although this will not affect access to the database when used through the server, that is replication and users will not be affected when opening the database through the server, anyone trying to access the database directly, or if the database is moved to another server using the operating system will not be able to access the database.

This can be used as an extra security layer on the databases - but needs to be administered carefully as it could potentially cause access issues if the wrong ID is used, or if the ID of the server is ever changed.

8.11 System databases not replicating

In some cases the system databases are found to be not replicating with each other.

As mentioned in "6.15 System databases with related replica IDs" the Replica ID of the system databases that need to replicate is based on the Replica ID of the NAMES.NSF in the Domain.

If the NAMES.NSF has been replaced with a Notes Copy of itself then it will have a new Replica ID and will not replicate with the other servers.

Along with the NAMES.NSF the other system databases, if they are replaced with Notes Copies of themselves, will now have a new Replica ID and will need to be resolved. If the NAMES.NSF is correct then in some cases these databases can be removed and recreated by re-starting the server, otherwise they will need to be replaced with replicas with the correct ID.

The incorrect NAMES.NSF will need to be replaced with a replica from another server and the other system databases recreated to correct their Replica IDs.

If the NAMES.NSF has been replaced and the Replica ID is incorrect for the Domain and then any of the other system databases are removed and allowed to be re-generated then the newly re-generated system databases will have a new Replica ID based off the incorrect Replica ID of the copy of the NAMES.NSF. The NAMES.NSF issue should be resolved and then the system database either replaced with a correct replica or allowed to re-generate.

8.12 Documents being removed by replication

In some cases it is seen that documents that are present in a database are removed after they replicate, even when the documents are present, not deleted, on the replica.

This can happen for two, related, reasons :

1) If selective replication is enabled on the database and the documents do not match the selection then they will be removed from the database - even if they are already present on the database and not replicating into the database.

2) If there is a reader field in use on the documents that does not include the source server. The server does not have access to read the documents and the replicator will remove them as not matching the selection formula - even when the selection formula is set to replicate all documents.