Considerations for a Multi-Library NOTEbookS Environment

NOTEbookS has the ability to share information among libraries within a company. This is accomplished technically through the use of Notes' selective replication.

Notes Selective Replication – A Brief Tutorial

In every Notes database and document there is a unique alphanumeric identifier. For databases this is called the replica ID. Databases having common replica Ids are termed "replicas". Only one replica of a database is allowed on a server.

Similarly, every document within a database has a unique identifier, called the "UNID". Replication in Notes occurs when a program (typically run periodically as a server task) goes through databases having common replica Ids. Each document is compared to other documents in the other replica having the same UNID. The ones with the newest "last modified date" overwrite those with older "last modified date". If documents are deleted the associated documents in the other replica(s) are also deleted. And any new documents are added to the replica database(s). Typically this replication synchronizes document changes in both replica databases. Documents that have been modified on both servers are at risk of overwriting data from the other. Notes saves both versions of the documents. One becomes a "replication/save conflict". Replication/save conflict documents are written to both replicas.

Notes has the ability to filter which documents are synchronized in this fashion. This is done through "Selective Replication". A formula is put into the database doing the synchronizing, and this formula filters only certain records for synchronization.

How This Works in NOTEbookS

Before it was stated "Only one replica of a database is allowed on a server." This means that each library in NOTEbookS must reside on a separate server in order to have replication take place.

In NOTEbookS we have one field in each record called "Replicate". It is set to "Yes" or "No". To ensure against any document which is missing the field, then, our selective formula is "Replicate not = 'No'" (in other words, replicate anything that has "Yes" or where the field is missing).

What Can Be Synchronized in NOTEbookS

A decision needs to be made on which databases are to be synchronized. The databases in question are the OnLine Catalog, Tables, Serials, Acquisitions and Research.

Each database in NOTEbookS has one (and only one) Configuration Document. These documents carry library specific information. As an example they carry, for the Online Catalog, the last accession number used and the accession suffix for that library. They are not to be replicated.

Default values for the replicate field are stored within each database's configuration document.

Online Catalog

The NOTEbookS Online Catalog holds bibliographic data. There is at least one record for each title at each location (typically multiple copies are recorded together in one document, but multiple locations are always represented in a separate documents). In this manner when catalog records are replicated, documents from one replica (one library) are merged with documents from other replicas (other libraries), forming a union catalog. So the first decision that needs to be made are which records are to be merged into a union catalog. To do this set the Replicate field at the bottom of the documents to "Yes". Records that are not to be included in the union catalog are set to "No".

Tables File

Each library has a group of librarians who modify the library's records. To provide autonomy we keep these librarians' names in an "Author Names" field called DocEditors in each record. When the replication takes place and documents are merged between replicas, the documents from library A will have Library A's library staff names in the DocEditors field, and the documents from Library B will have Library B's library staff names in their DocEditor fields. In a multi-library configuration the LibraryStaff group members are given Author access in the database Access Control List (ACL). This ensures that each library can modify their own records, but not those of the other library. Note that in a single library configuration this is not an issue and the LibraryStaff group is typically set up with Editor access (versus Author access) in the database ACLs.

The LibraryStaff names are extracted from each library's NOTEbookS Tables file (and the LibraryStaff table documents are not replicated between libraries…'Replicate' is set to "No"). So if there is to be one "master librarian", his/her name should be in both libraries' tables file, but other library staff names should only be stored in their own library's tables file.

There are a number of other tables entries which are not to be replicated from tables file to tables file. These include, as an example, the name and address of the particular library. Vendor and publisher information may, depending on the company, be different for each library (and therefore would not be replicated). For some companies this might not be the case, and these records should be replicated. In general it is recommended that the tables data dictionary be printed out and a decision made on a field by field basis on what is to be replicated, and what is not.

Similarly, subject keywords are kept in the tables file. If these are to replicated between libraries (as is commonly done), then these documents need to replicate.

A full list of fields that are not to be replicated (or that should be examined and a decision made) is:

AcquisitionsDBName (optional)
BudgetCategory (optional)
CatalogFileName
CircOverdueMessages
CityA
CompanyA
CompanyName
CopyrightMessage (Optional)
CustomReqMessage
DBSInstalled
DeliveryInstructionsA
Format (Optional)

HomeLocation
LibraryMgr
LibraryName
LibraryStaff
Location (optional)
LocationSC
OtherDB
OverdueDays (optional)
PaymentMethod (optional)
PublicationPeriod (optional)
PurDeptA
RecallMessages (optional)
ResearchStaff
ResSources (optional)
SerialsAcquisitionsDBName (optional)
SerialsDBName (optional)
SerialsDistributionDBName (optional)
SerialsDistributionFileName
SerialsFileName
SerialsRID
SerialsUsage
StateA
StreetA
SubscriptionTerm
Usage
WebCatText
WebGlobalGreeting
WebResText
WebSerText
ZipA

Serials

The Serials database has 3 types of records (plus a configuration document). They are: Master Holding (Subscription) records, Issue records, and article abstract records. There is at least one master holding record for each subscription/location.

A decision needs to be made on whether master holdings records, issue records, and article abstract records are to be replicated or not.

Research

The Research database has three types of records, research requests (in various stages of completion) knowledgebase documents (completed research), and experts records. A decision needs to be made on which of these records will be shared between libraries, and those that will not be shared.

Acquisitions

Some libraries have centralized acquisitions, others keep this function decentralized (each library does their own purchasing). If centralized purchasing is desired then the Acquisitions database should be non-selectively replicated. Otherwise it should not be replicated at all.

Deletion Stubs

When a document is deleted in Notes and there are other replica copies of the document (the same UNID) in other replica databases, the replica documents are also deleted.

There is a situation that must be fully understood to ensure documents are not inadvertently deleted. This situation occurs when the Replicate field is changed. If a document is set to Replicate = “Yes” and it is replicated around, then the field is set to “No”, the selective replication filter disallows comparisons between the two databases for that document. The document will then be deleted in the “parent” database.

To make sure this doesn't happen, the safest thing to do is to:

- change the field to "No"
- do a cut and paste of the document within the database (this will assign a new UNID to the document)
- let replication occur

This needs to be thoroughly tested, especially in a hub and spoke environment where selective replication may or may not be in place at every replication point.

Therefore, a decision should be made before replication commences on which records are going to be shared and which are not to be shared. Also, once a record has been replicated it needs to be carefully changed according to the steps above so that it will not replicate.