This education series consists of three modules that cover backup and recovery strategies for a Lotus® Domino® implementation on i5/OS®. This second module focuses on using replication and clustering for high availability.

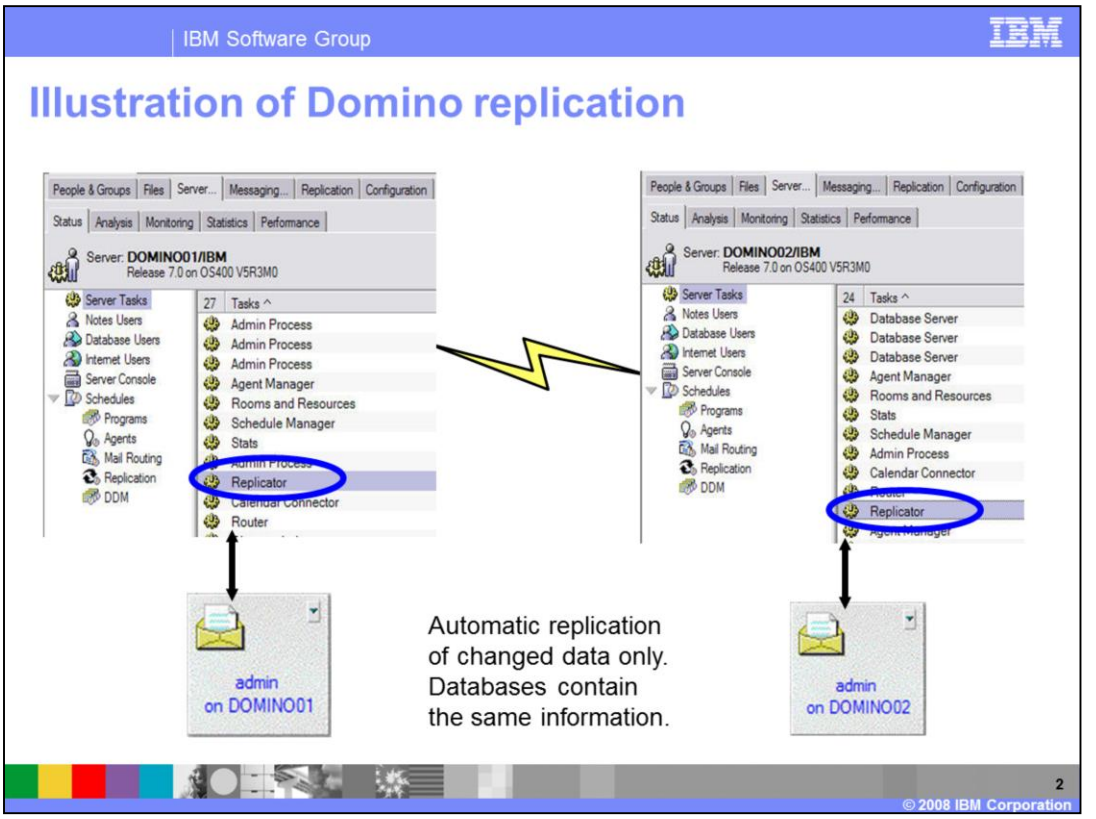

In this illustration, there are two Domino servers (Domino01 and Domino02) with the admin mail file on both servers. Domino replication sends the changes between the servers so that both databases contain the same information. This example is used in more detail throughout this presentation.

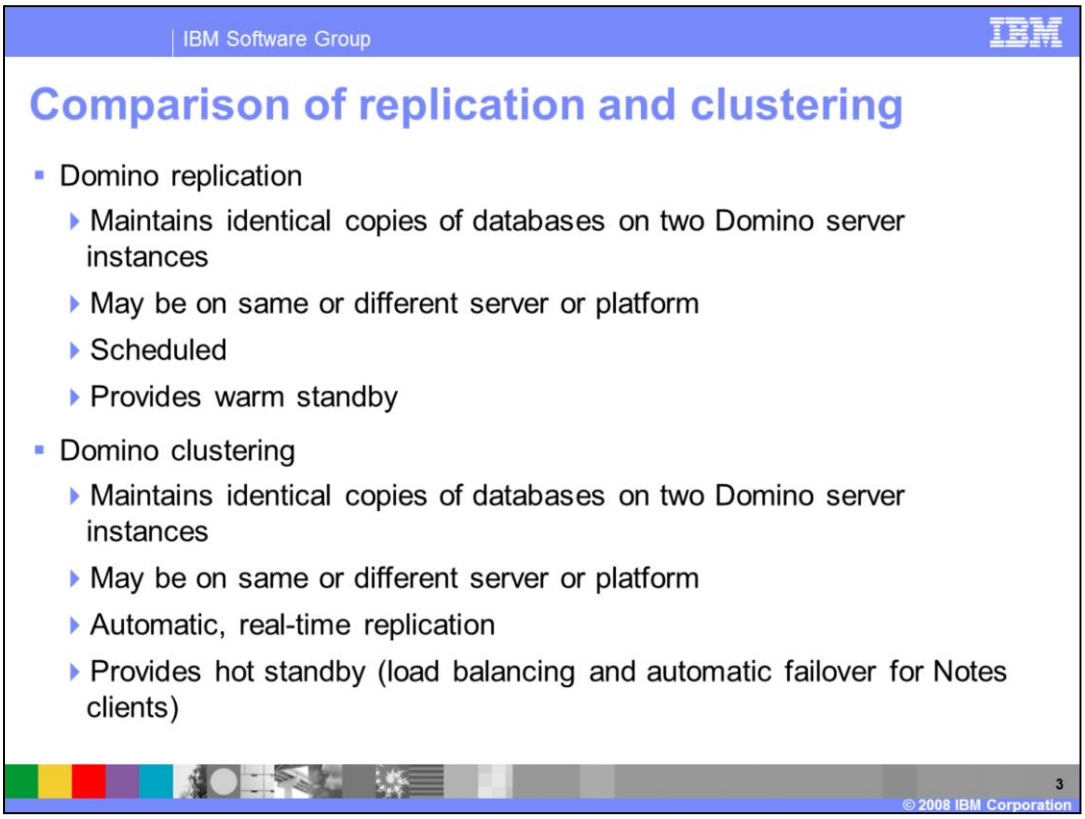

Domino replication maintains identical copies of a database on two Domino servers. These servers may or may not be on the same physical hardware. Replication can be scheduled to occur automatically. It can provide a warm standby for the server. Clustering provides real-time replication or a hot standby. It is typically used for load balancing and automatic failover for Lotus Notes clients.

What is Domino clustering?

- A Domino cluster is a collection of two to six Domino servers which provides high availability and workload balancing.
- How does it work?
  - Two or more servers are selected to be in a cluster (up to six servers)
  - Each servers contain replica databases of the applications which are to be clustered
  - If user accesses a clustered server and it is not available, Notes opens a replica of the database on a different cluster server.
  - Domino continuously synchronizes databases

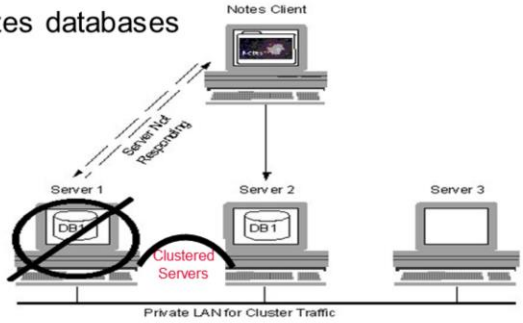

A Domino cluster is a collection of two to six Domino servers, which can provide high availability or load balancing for your users. When dealing with a cluster, each server will contain its own replica copy of the clustered database or application. The Domino servers will continuously synchronize the databases.

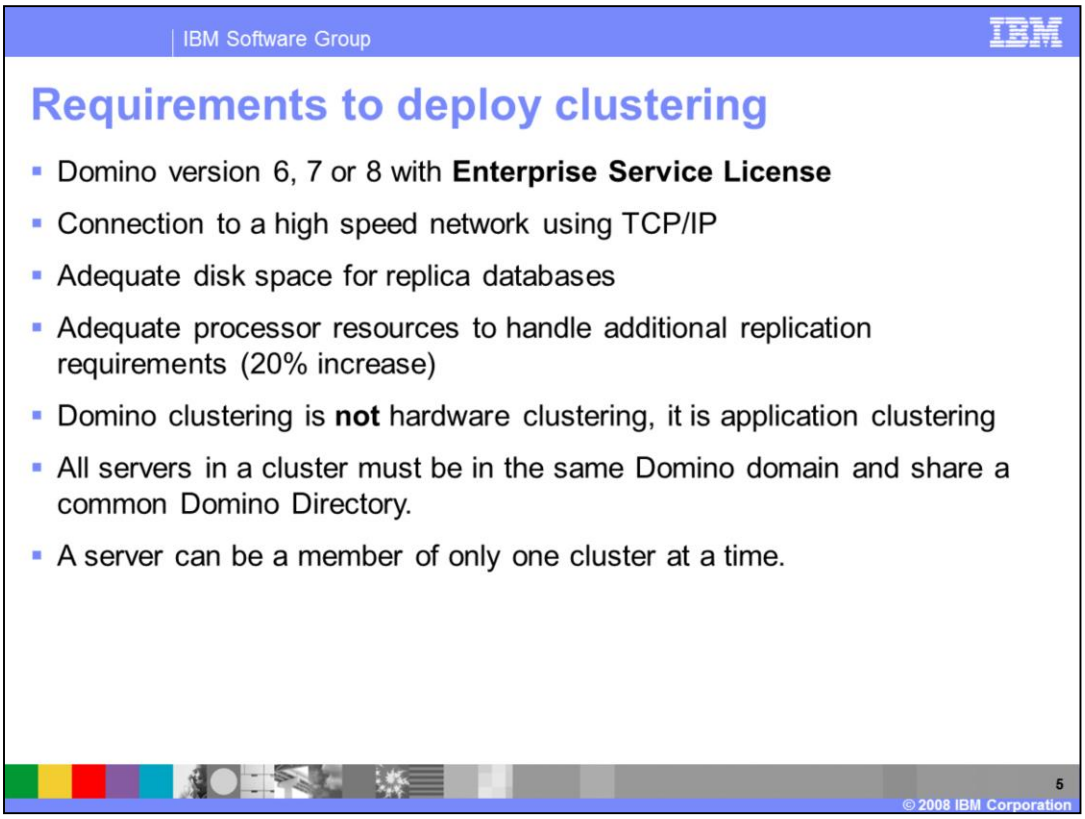

You must purchase a Domino Enterprise Service License to configure clustering in your environment. You will also need a high-speed network connection using TCPIP between the clustered servers. You must ensure that you have adequate disk and processor resources to handle the extra workload that clustering requires. Note that Domino clustering is application clustering and not hardware clustering. Also, note that the clustered servers must share the same Domino directory. Finally, a Domino server can only be a member of one cluster.

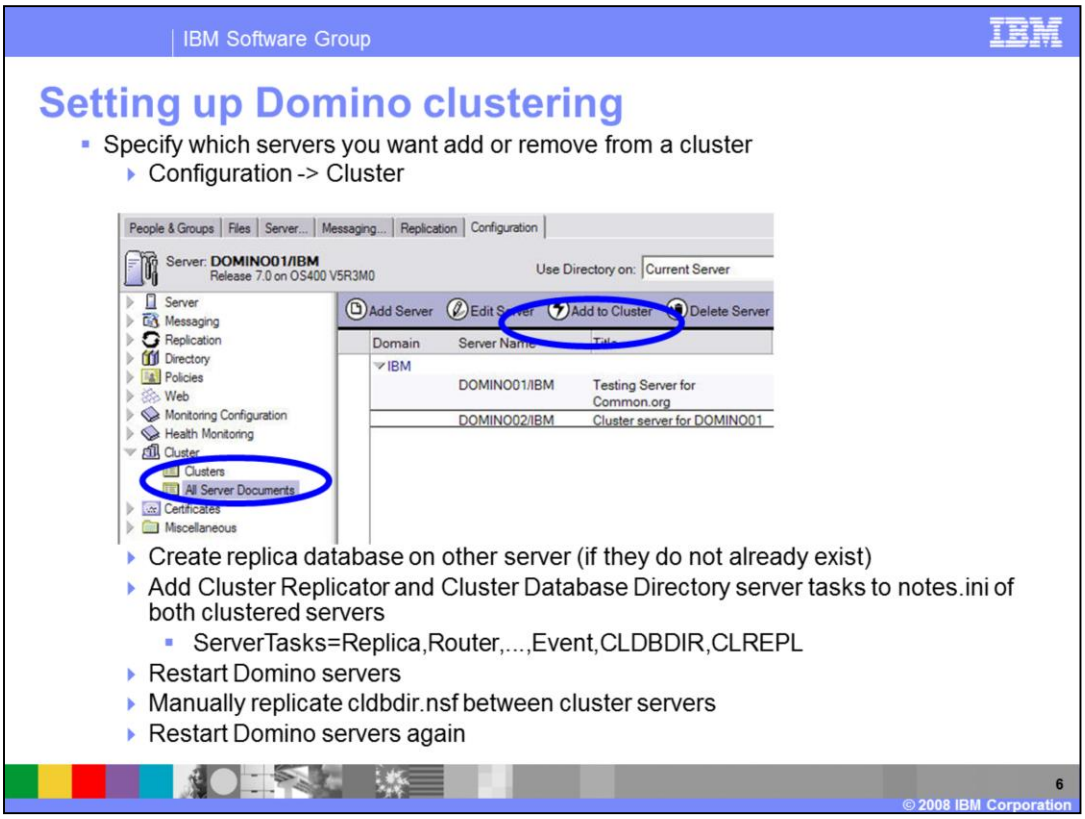

You can configure a new cluster from the Configuration tab of the Domino Administrator client. From there, you can select to add a server to a cluster. This process adds the cluster database directory (cldbdir) task and the cluster replicator (clrepl) tasks to the server. The process will also create a cldbdir.nsf database, which you must replicate to the other servers in the cluster. Note that only databases where replica copies exist on both servers will have failover capabilities in the new cluster. To ease the cluster setup, it is a good idea to create replica databases on the other server before creating the cluster.

Setting up Domino clustering (continued)

- Open the Cluster Database Directory (cldbdir.nsf) on the clustered servers
  - Configuration > Cluster > Cluster Directory(6)
  - Manage your replica copies and choose which databases to cluster

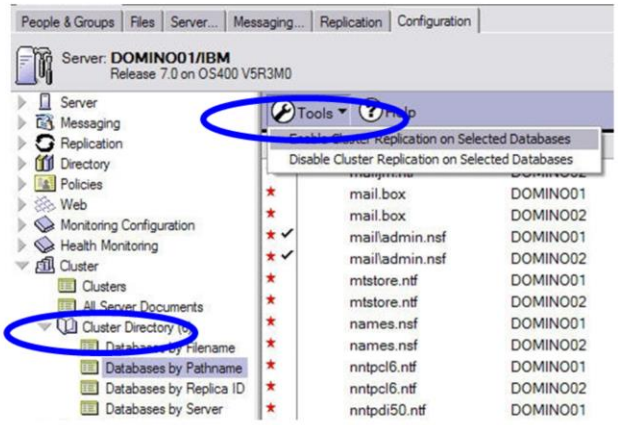

After you configure your cluster, you will then specify which databases to cluster. Perform this step by opening and editing the cldbdir.nsf database. Automatically clustered are databases that have a replica on the other server - meaning they contain the same replica ID on both servers.

What really happens with clustering

- Replication on steroids
  - When new mail is received it replicates it real-time to the replica copy in the Domino cluster
  - Replication synchronizes database replica copies within the cluster
  - Server > Status > Server Console > Live

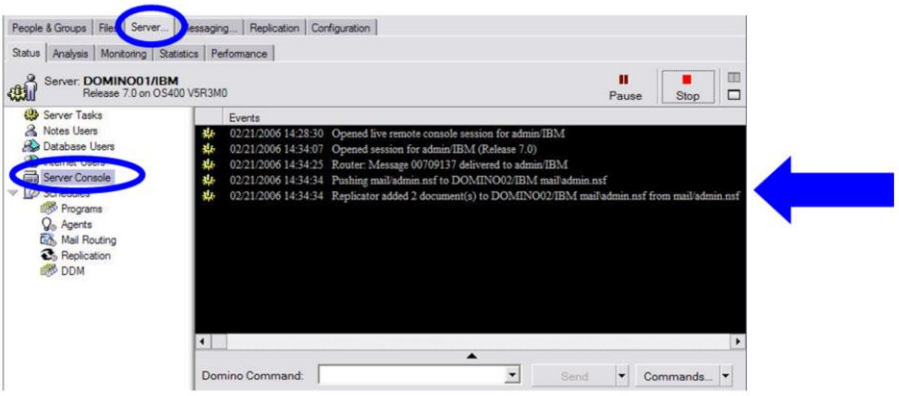

In a cluster, the server sends database changes to the other server.

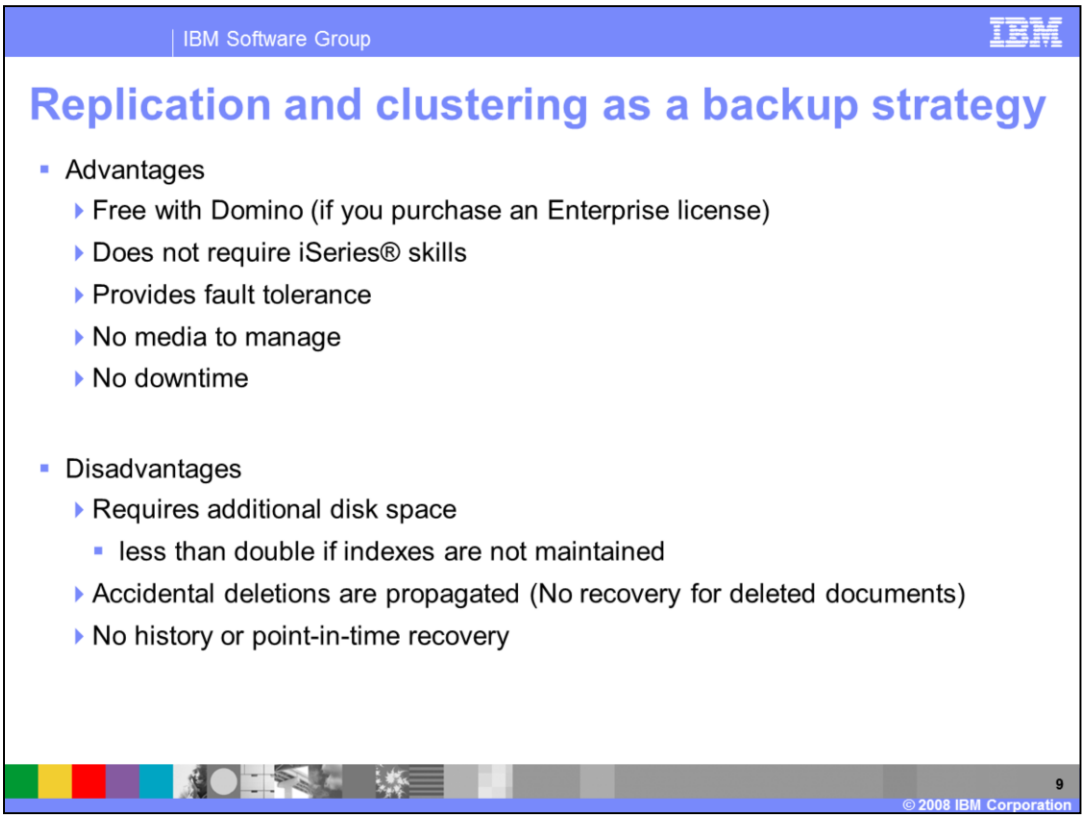

You should consider the advantages and disadvantages of using replication or clustering as a backup strategy. If you already purchased an enterprise service license for Domino, you do not need to purchase any additional software. You will also always have a hot backup available. The disadvantage is that you need double the disk space if you want to cluster all files. Also, deletions propagate. Thus, you should not use replication and clustering as the exclusive backup for your environment.

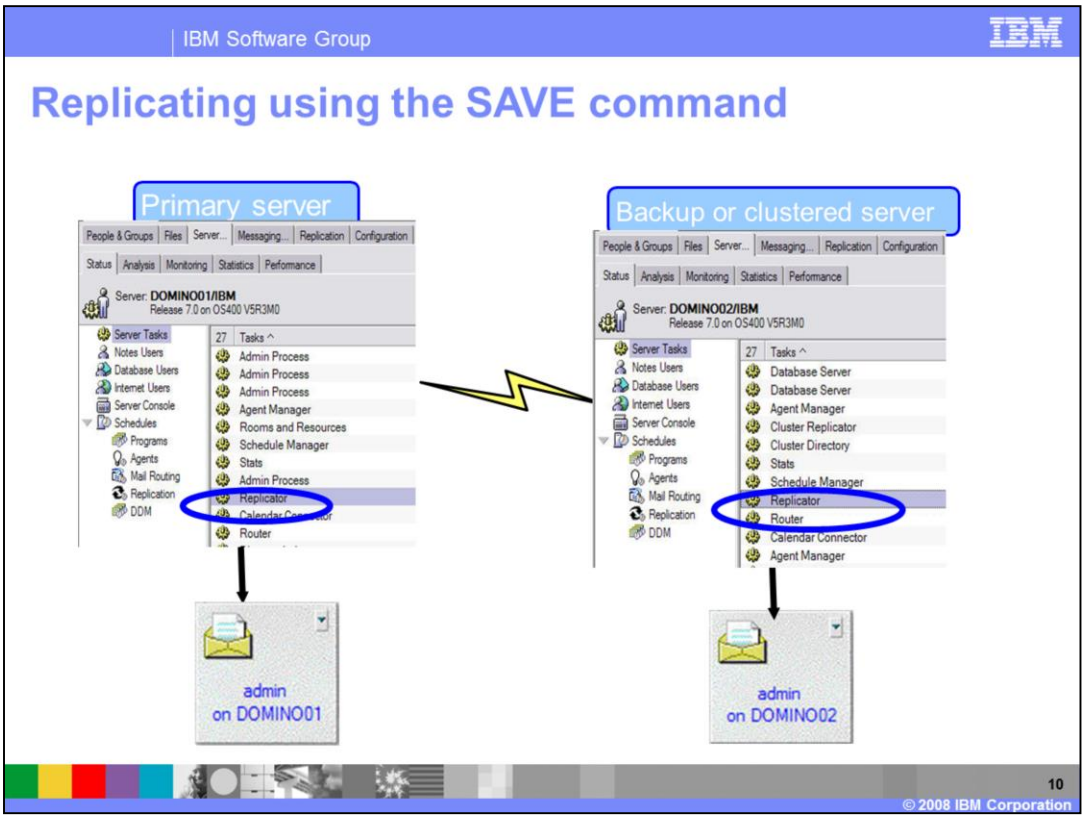

Using replication or clustering with the i5/OS Command Language (CL) SAVE commands is an attractive option for some. In this way, because there is always a warm backup available, you will not need to take the primary server down for backups.

.

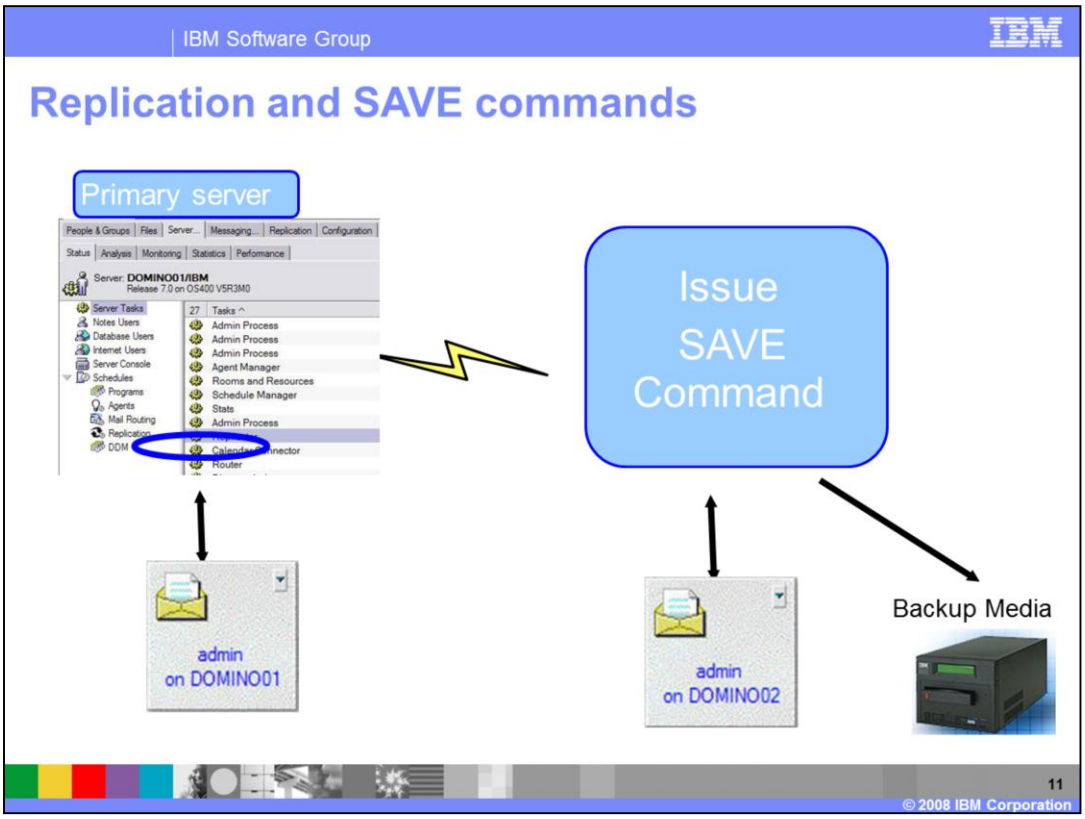

When using replication with the SAVE commands, you will first replicate all of the databases, and then end the secondary server to perform the backup.

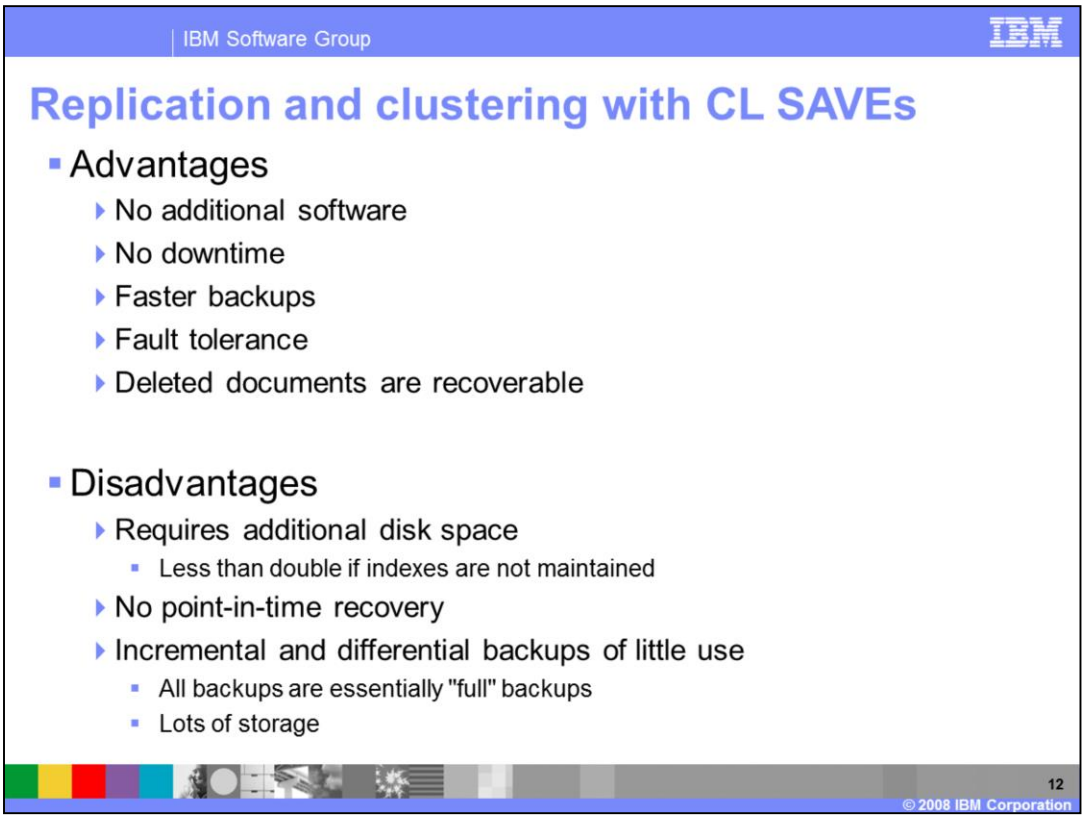

Here are some immediate benefits of using replication and clustering with i5/OS CL SAVE commands. It requires no additional software, no downtime of the primary server and deleted documents are recoverable. The disadvantages include the fact that it requires additional disk space and there is no option for incremental saves or point in time recovery.

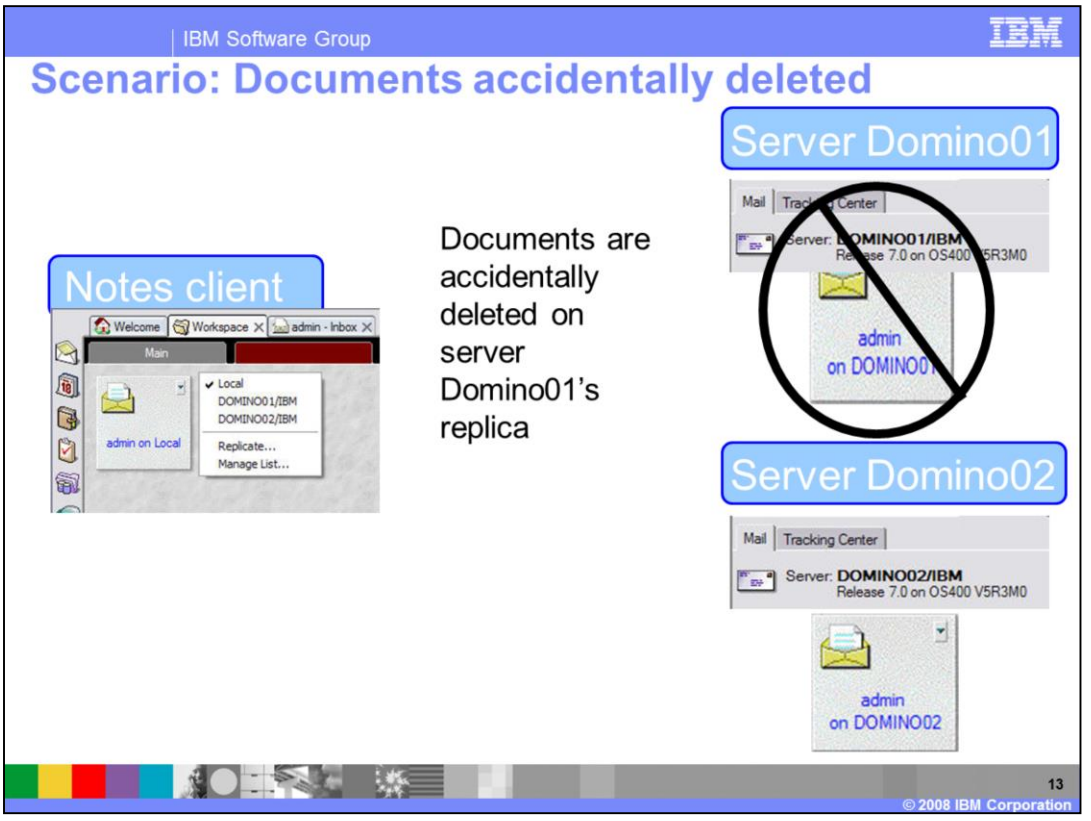

When you have replica databases in your environment, you must take special care when recovering a database. Here is an example to help you understand what can happen. Assume that the admin database has replica copies Domino01, Domino02 and their Lotus Notes client. The user accidentally deletes some documents from the copy on Domino01.

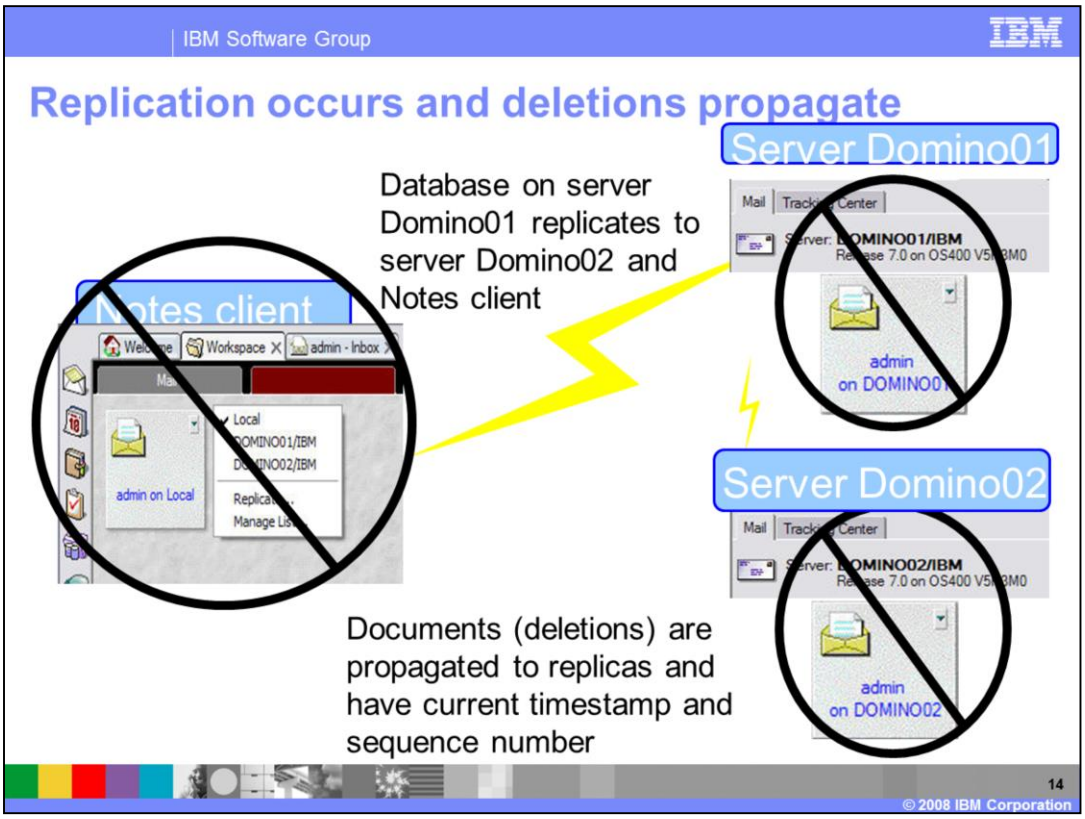

Since the administrator configured replication to occur automatically, the deletions replicate to the other database copies. Now the documents have been deleted from all three copies of the database.

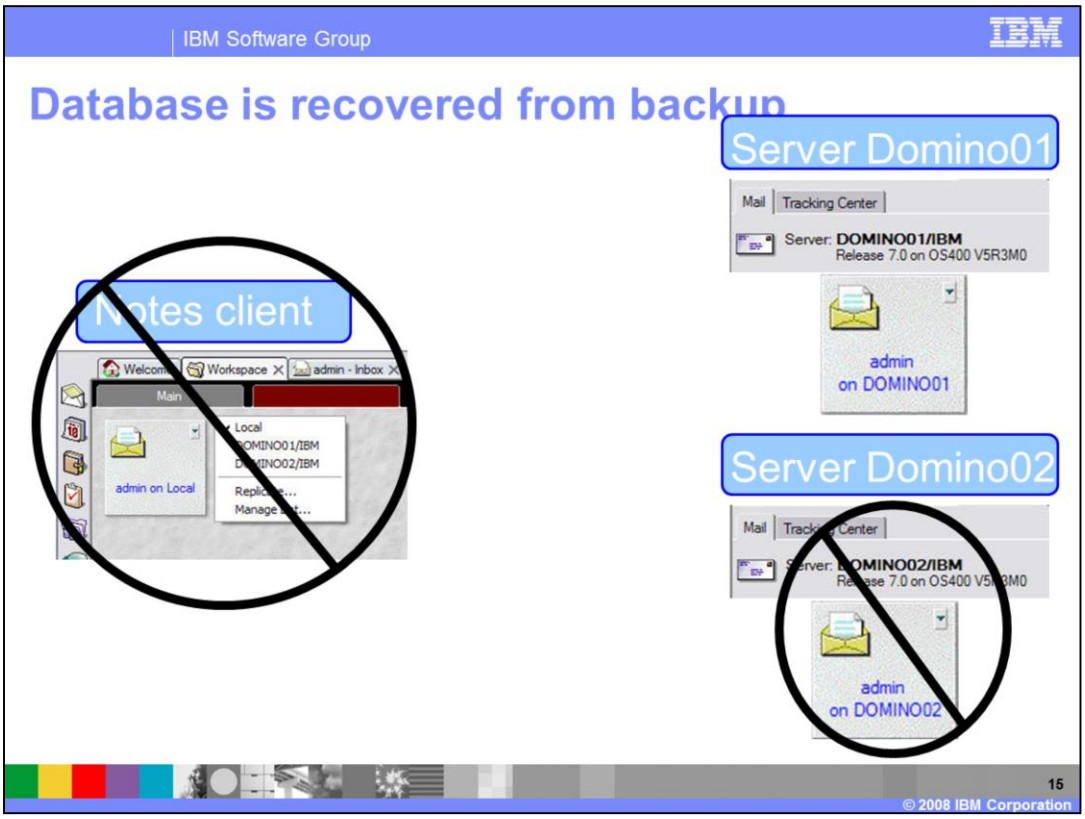

The administrator has been notified that the documents have been deleted and the database has been restored from backup to server Domino01.

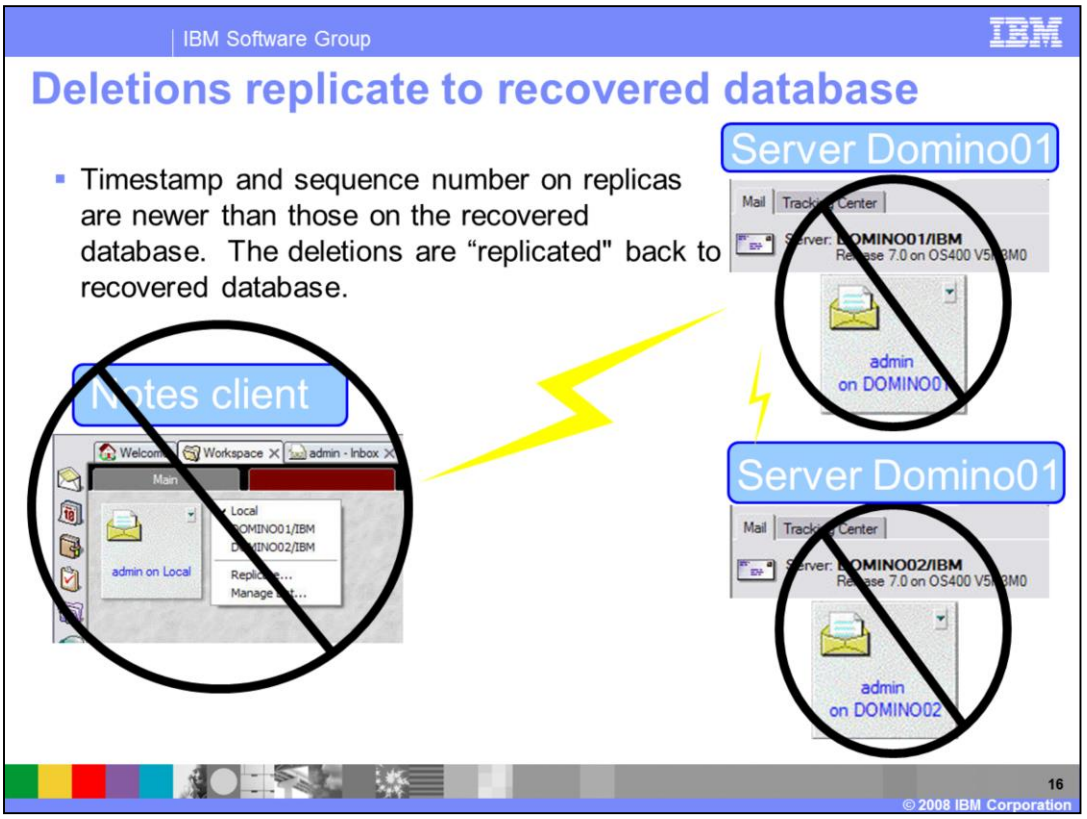

Because the timestamp and sequence number of the deletion is newer than the documents in the recovered database, the documents are again deleted from the database on Domino01.

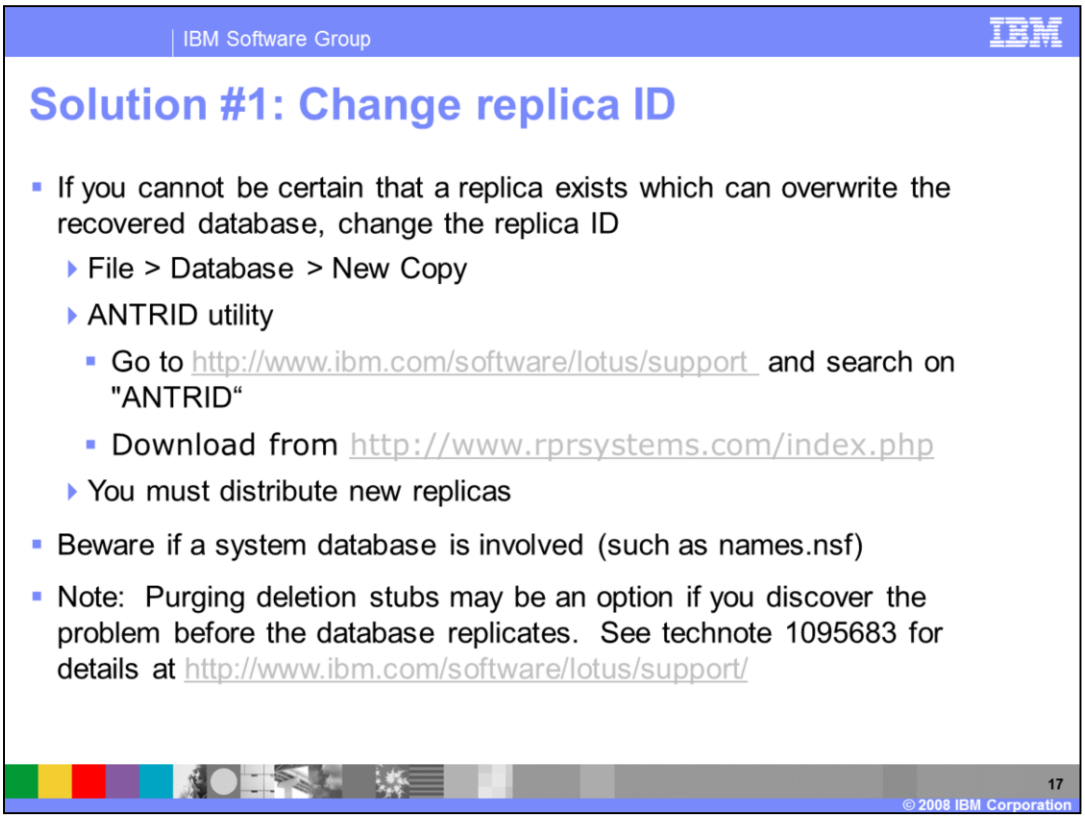

One option to resolve this problem is to change the replica ID of the recovered database. You can download a utility to update the replica ID. However, this will break replication because you must create new replicas on Domino02 and the Lotus Notes client. Thus, be aware if a system database like names.nsf is involved. Names.nsf must replicate between all servers in the Domino domain. If you change the replica ID of names.nsf or another system database, you must place a new replica copy of names.nsf on all other Domino servers in the domain. One thing to note is that if you discover the deletions before replication occurs, it is possible to purge the deletion stubs from the affected database. In this situation, you can follow the instructions in technote 1095683.

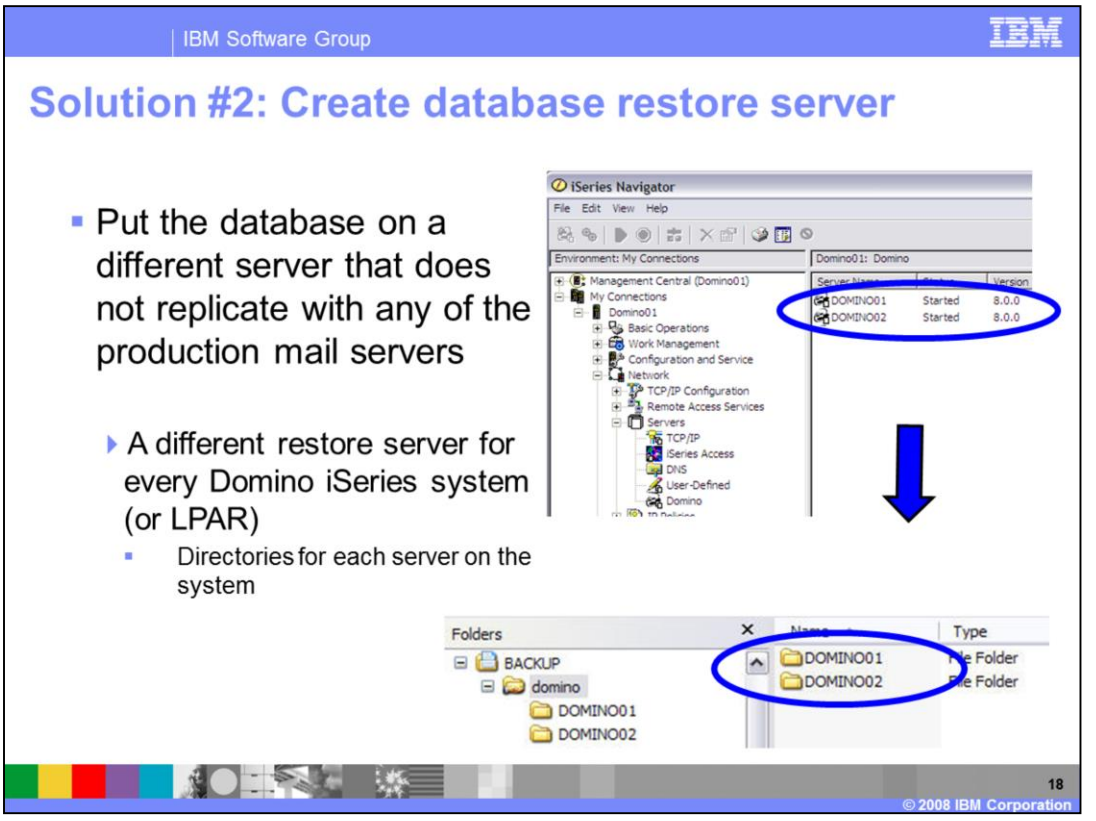

Another option to resolve this problem is to have a database restore server. The purpose of a restore server is to have a place to restore Domino databases where replication is not scheduled to occur. When you use a database restore server, you must open the database on the restore server and then paste the missing documents into replica copies of the database. You can then let the regularly scheduled replication put the documents back on the other server and Lotus Notes client. In this example, a backup Domino server called “Backup” exists. On this server, a subdirectory for each server is used as a temporary holding place for restored databases.

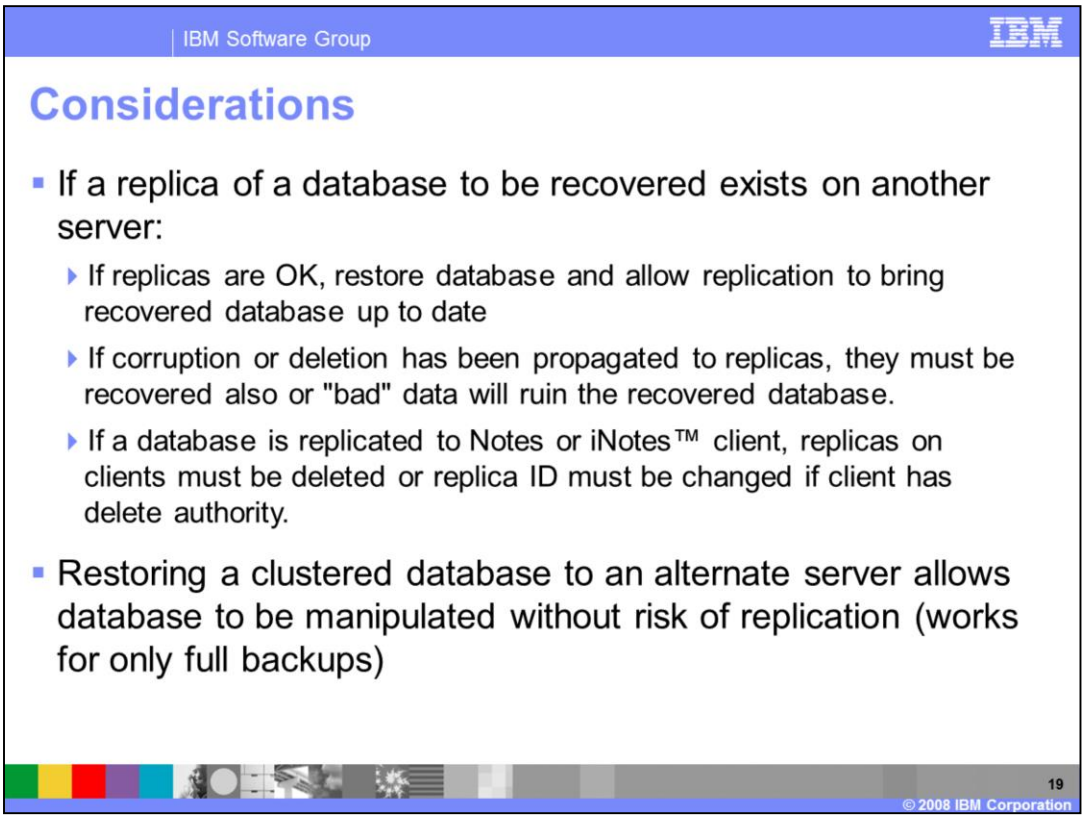

In summary, care must be made when a replicated database is recovered from backup. If you have removed the database from the server where the documents were incorrectly modified or deleted, but the change had not replicated to the remote server you can restore the database to the server where the problem existed and allow Notes replication to run with no further action. However, if the deletions or problem with the database has already replicated to other servers you need to restore the database on both the original server and the cluster or replica server. If you do not do this, the recovered database will be damaged. This same thought process exists for clients that have a replica databases. If the client replica does not contain the problem, no action is needed when the server copy is restored to the server. If the client replica is also damaged or missing documents, a new client replica must be created to protect the server copy from getting the newest changes back into the database. Finally, if you are going to restore the database to a database restore server or another server that does not replicate, you can open the restored database, copy and paste documents from it without having to consider how it will affect your other replica databases.