

Advantage Replication FAQ

What is replication?

Replication provides the capability to automatically distribute changes from tables in one database to tables in another database without any work required on the part of client applications. Replication allows for a consistent view of a database to be maintained across multiple Advantage Database Servers.

Is replication functionality built into the Advantage Database Server?

Yes. Advantage Replication is available with Advantage Database Server v8.0 and newer as separately licensed functionality.

What server operating systems are supported with Advantage Replication?

Advantage Replication is supported on Microsoft Windows NT/2000/2003 and Linux. Replication with NetWare versions of Advantage Database Server is not supported.

How does replication compare to synchronization?

Advantage replication is not a synchronization process. Replication provides the ability to distribute table changes from one database to another. It does not verify that the tables are synchronized.

Does replication guarantee that source and target tables are always in sync?

No. Advantage replication simply distributes updates from the source database to the target database(s). If updates made to the target are not replicated to the source, the tables will be out of sync. Advantage replication does not perform any synchronization actions.

What is synchronous vs. asynchronous replication? Which does Advantage support?

With synchronous, or real-time, replication the databases involved in the replication process do not commit a transaction until all databases at distributed sites that are being replicated have processed that transaction. Synchronous replication is a solution that is narrowly suited for extreme, high availability networks involving very few distributed replication sites. This is due to the network latency introduced into each transaction as well as the high cost of the bandwidth required to attain an acceptable level of performance in this environment. Asynchronous replication imposes much less overhead on client applications. When an update is made by a client, the update information is placed on a queue to be processed independently of the client update. Asynchronous replication results in little performance degradation, nor does it require special hardware and high-bandwidth connections in order to be functional. Advantage Replication uses an asynchronous model.

Will replication cause performance degradation?

Some extra processing is required, but there should be very little performance degradation due to replication. When an update for a client application occurs, the update is stored in the replication queue. That is the primary cost of replication that may affect a user application. The bulk of the replication work is performed asynchronously by other threads and, thus, reduces the direct impact on existing applications.

Does Advantage Replication require Data Dictionaries?

Yes, Advantage Replication does require that all tables are bound within an Advantage Data Dictionary.

Does Advantage Replication support DBF tables?

Yes, Advantage Replication supports DBF and ADT tables. Both table types must be bound within an Advantage Data Dictionary.

How is replication different from the online backup functionality?

While there are some conceptual similarities between replication and backup functionality, they are implemented with very different solutions in mind. The backup functionality is intended for making backups of your data for subsequent restore operations. Replication is used to keep two or more live databases logically consistent with each other.

- Replication sends updates to another usable live database. Backup sends the data to a backup location that is not intended to be accessed directly by an Advantage application.
- Replication requires the target location to be running Advantage Database Server. Backup functionality does not.
- Replication operates on an update-by-update basis. If the target database is online, the update is sent to the target immediately after the original update occurs. Backup functionality works on a request basis. Updated records are tracked by the backup functionality and then all processed when requested by a user application or by the backup scheduling utility.

Can I replicate non-Advantage databases?

Advantage Replication is only available for data managed by Advantage Database Server. However, Extended Systems does offer replication and synchronization functionality of non-Advantage databases via our OneBridge Mobile Data Suite products, which is available separately. Contact your Advantage distributor or Advantage Sales representative for more information on replicating to non-Advantage databases.

What Advantage products are required for Advantage Replication?

Advantage Replication supports server-to-server replication between two or more Advantage Database Servers. Therefore, all machines within the replication environment must be running the Advantage Database Server v8.0 or greater. Licenses for Advantage Database Server and Advantage Replication are purchased separately.

Will Advantage Replication work with versions of Advantage Database Server older than 8.0?

No. Advantage Database Server v8.0 or greater must be used in order to use Advantage Replication.

Does Advantage Replication provide support for replicating data to laptops?

Yes. Advantage Replication can replicate data from a server to any Windows desktop or laptop that is running the Advantage Database Server.

How is security handled?

Advantage Replication uses the Advantage Client Engine to communicate with target servers. This allows for the same security that standard client applications enjoy. For example, the communications with the target server can be encrypted and authenticated. In addition, the replication queue can be encrypted at the source database to prevent viewing of the pending replication updates. Advantage will replicate only to an Advantage Data Dictionary; it will not replicate to a non-authenticated free connection.

Can changes be made to the databases being replicated while a replication session is occurring?

Yes. Advantage Replication does not open tables exclusively nor otherwise keep the tables involved with the replication process from being updated by another application. Individual records will be locked as they are being updated, but they will be available for update once those records have been written.

Can the frequency of replication between servers be scheduled?

No. Replication updates are processed asynchronously and on a continuous basis.

How is conflict resolution handled?

Conflict resolution can be handled with "CONFLICT" triggers on target tables. The actions taken by the CONFLICT triggers are up to the developer.

Can I replicate tables containing Auto-Updating fields?

Yes. However there are some special considerations when replicating tables with AutoIncrement, RowVersion, and ModTime fields. When replicating a record with an AutoIncrement field, the AutoIncrement value is copied to the target table. If a record already exists at the target with that same AutoIncrement value, either a replication conflict will occur or the target table will have two records with the same AutoIncrement value. A unique index on the AutoIncrement field can prevent duplicate values from being entered into the target table. RowVersion and ModTime field values are ignored by the replication processing; the values for those fields are always generated at the target.

What changes to my application are necessary?

If you are using an Advantage Data Dictionary, then no changes should be necessary. In general, replication requires no changes be made to the database, and thus replication should have no effect upon the application. The only exceptions to this would be if changes were required to a table's structure to allow all replicated tables to have the ability to uniquely identify each row.

Can I replicate tables that have Referential Integrity constraints defined on them?

Yes. When replication is performed, referential integrity rules are enforced to maintain primary and foreign key relationships. However, cascade operations are not performed at the target; they are only performed at the source. If an RI cascade rule exists, then you should replicate all tables involved in the cascade operation.

What does Advantage replication cost?

Advantage Replication pricing is based upon the user count of the Advantage Database Server you will be adding replication functionality to. Contact your local Advantage distributor or Advantage Sales representative to get Advantage Replication pricing information.

Can I use the Advantage Local Server product to perform replication of the server database?

No. Due to functionality requirements, the Advantage Database Server is required for Advantage Replication.

Can I replicate both tables and views?

Yes. The server can replicate directly to tables or to SQL views.

Is one-way and two-way replication available?

Yes. Advantage Replication supports one-way replication of updates from one server to another server only, or two-way replication of all updates on both servers to each other.

How do I incorporate transaction processing into the replication process?

Advantage Replication respects transactions in your current databases. If the updates occur in a transaction at the source, the replication updates will occur in an equivalent transaction at the target.

Can I limit what data is replicated?

Yes. Filters can be specified so that only a subset of records from a given table is synchronized.

Do the tables being replicated require primary keys?

No. You can choose which columns to use to uniquely identify the record in the target database.

Must all fields in a source table be replicated to the target table?

Yes. In the first version, there is no vertical filtering available.

Can fields from multiple tables on the server be synchronized to fields in a single table on the client?

No. Advantage Replication operates at the table level.