Using Replication

Collected by
Ben Laiid

I. Introducing Replication

Replication in Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE) is based on Microsoft SQL Server 2000 merge replication. Merge replication is ideally suited to portable devices because it allows data to be updated autonomously on the portable device and the server. The data can later be merged when the device is connected to an instance of SQL Server.

In SQL Server CE, replication is message-based. SQL Server CE synchronizes with SQL Server by establishing an HTTP connection to the SQL Server Publisher through Microsoft Internet Information Services (IIS). By connecting through IIS, replication takes advantage of IIS authentication and authorization services. You can locate the Publisher behind a firewall or provide access to it using Microsoft Internet Security and Acceleration Server 2000. Replication can be performed over local area networks (LANs) and wide area networks (WANs). Using Microsoft ActiveSync®, you can communicate with SQL Server through the desktop computer's network connection, eliminating the requirement for the Windows CE-based device to have a network card or modem.

The replication communication protocol is designed for wireless transports. The protocol uses compression to reduce the amount of transmitted data, and supports the use of encryption to safeguard sensitive user data during transmission. Following a communication failure, transmission resumes from the last successfully transmitted message buffer.

SQL Server CE replication offers row and column filtering, which enable you to define and maintain subsets of data that can be published to different sites. Row filtering can be used to publish a subset of rows in a published table. For example, in a service organization, each service representative may only need to receive the customer records for the accounts they service. Column filtering can be used to publish a subset of columns in a published table. For example, a column filter can be used to eliminate large text or image columns that you may not want to publish to a Windows CE-based device.

Row and column filtering can be used on the same article within a publication. You can use filtering to specify the data that is published to each Windows CE-based device. For more information, see "Filtering Published Data" in SQL Server Books Online.

II. Planning for Replication

The following topics in this section provide information to consider before implementing replication in Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE). Typical SQL Server CE replication topologies are described. Information regarding what data can be replicated and limitations on published data is also provided.

II.1 Replication Topologies

You can configure replication in Microsoft® SQL Server™ 2000 Windows CE Edition (SQL Server CE) in a variety of ways. You should choose a configuration based on the following factors:

- Performance

  The performance of a system depends on the number of Microsoft Windows CE-based devices that must be supported, the frequency with which the devices must synchronize with the server, the connection bandwidth, the timeframe during which synchronization must occur, and the load that each device imposes on the server while synchronizing.

- Security

  The configuration you choose is influenced by the security requirements of the application you develop.

Basic Configuration Rules

The following basic configuration rules govern SQL Server CE replication configuration:

- The Subscriber is always an anonymous Subscriber.

  Individual Subscribers cannot be assigned a merge replication conflict resolution priority; the priority for Subscribers is always 0.

- SQL Server CE cannot be a Publisher or Distributor.

  Two or more Windows CE-based devices can share a publication only by subscribing to the same SQL Server 2000 publication.

- A Subscriber must always synchronize with the same Publisher.

  Synchronizing with alternate synchronization partners is not supported.

Topologies

You can choose from a variety of SQL Server CE replication topologies. The more common ways to configure SQL Server CE replication are:

- Single server
- Two servers
- Multiple Microsoft Internet Information Services (IIS) systems and SQL Server republishing

Single Server

In the simplest SQL Server CE replication topology, IIS, the SQL Server Publisher, and the SQL Server Distributor all reside on a single server. SQL Server CE Subscribers, running on Windows CE-based devices, synchronize by connecting to IIS on the Publisher. The Publisher can be located behind a firewall.

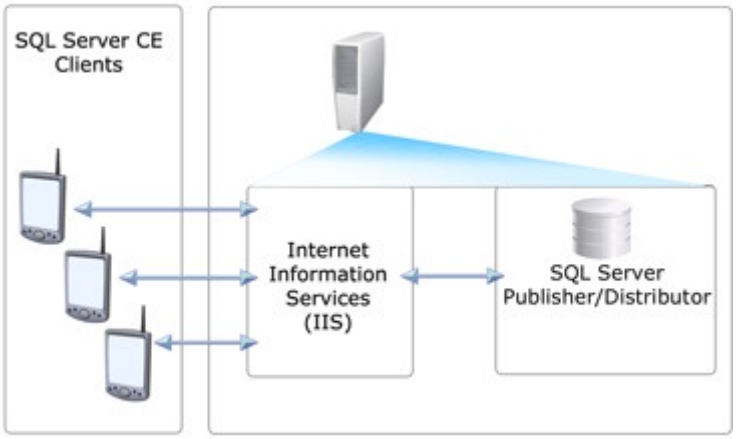

Two Servers

You can place IIS on one server and configure the SQL Server Publisher and Distributor on another server. The server running IIS can be isolated from the Internet by a firewall. SQL Server CE clients, running on Windows CE-based devices, synchronize by connecting to IIS. Another firewall can isolate the SQL Server Publisher and Distributor from the IIS system.

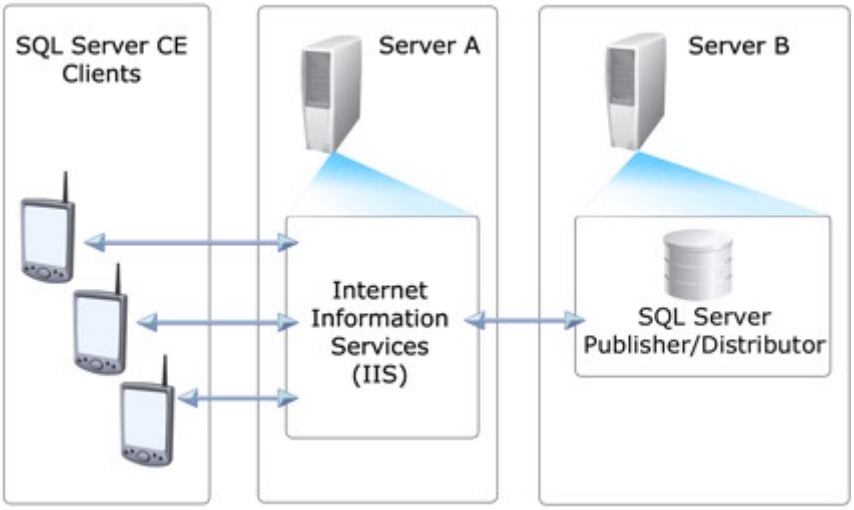

Multiple IIS Systems and SQL Server Republishing

If you must support very large numbers of SQL Server CE Subscribers that synchronize at the same time, you can partition the work across multiple computers running IIS.

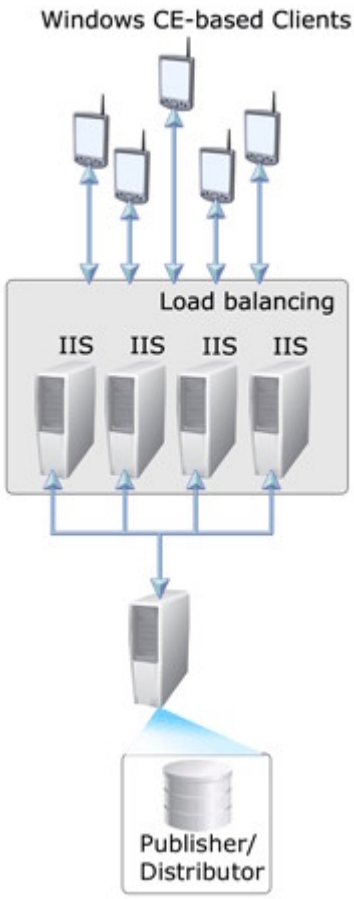

If further load balancing is required on the computer running SQL Server, you can create a republishing hierarchy on multiple computers. The top-level Publisher publishes data to Subscribers, which in turn republish the data, load balancing requests from the SQL Server CE clients. Load balancing is static in the sense that each SQL Server CE Subscriber is assigned to one of the republishers and always synchronizes with that republisher (SQL Server CE does not support alternate synchronization partners).

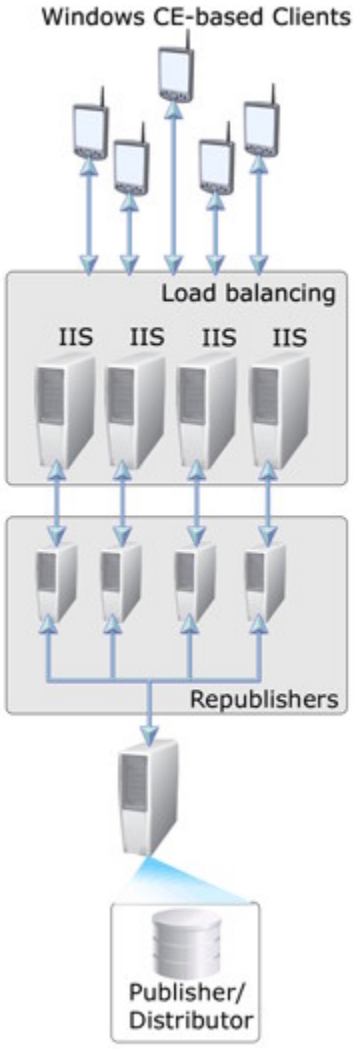

By using Microsoft ActiveSync®, a desktop system may be used to synchronize SQL Server CE to an instance of SQL Server. By using Microsoft Windows Powered Pocket PC 2002 and Microsoft ActiveSync 3.5, you can use the Pass Through functionality. If you are using ActiveSync 3.1, you must install SQL Server CE Relay.

Using Windows Clustering

SQL Server CE replication can be used on a Microsoft Windows 2000 cluster containing a farm of Web servers running IIS. However, the SQL Server CE Server Agent must maintain session state across client calls; therefore, you must use load-balancing hardware or software that is capable of maintaining a session state between the SQL Server CE Client Agent and the SQL Server CE Server Agent.

Load-balancing systems typically work as follows: Clients access the cluster using a common domain name with a single virtual IP address; the load-balancing hardware or software intercepts the incoming HTTP traffic and directs it to one of the servers in the cluster.

SQL Server CE can be used with a variety of load-balancing systems for example Windows 2000 Network Load Balancing (NLB).

Windows 2000 NLB is included as part of Microsoft Windows 2000 Advanced Server and Windows 2000 Datacenter Server. NLB can be used with SQL Server CE because NLB supports

client affinity, which is an option that ensures that the client is always directed to the same server running IIS within the cluster. Note, however, that if a server or network failure occurs during a client session, a new logon is required to reauthenticate the client and reestablish session state. Also, adding a new server running IIS to the cluster redirects some client traffic to the new server, which can affect existing sessions.

II.2 Supported Data Types and Data Type Mappings

The following table shows the data type mappings that are performed when replicating between Microsoft® SQL Server™ and Microsoft SQL Server 2000 Windows® CE Edition (SQL Server CE). The table lists mappings for each SQL Server data type and describes restrictions or special behaviors.

SQL Server data type	**SQL Server CE data type**
bigint (**int 8**)	**bigint** (**int 8**)
binary (n)	**binary** (n) or **image** If the length of the data is 510 bytes or less, SQL Server CE replication maps the SQL Server **binary** data to SQL Server CE **binary**; otherwise, it maps it to SQL Server CE **image**. If the length of the **image** data exceeds the length of the **binary** column, synchronization fails when the data is sent from the SQL Server CE Subscriber to the SQL Server Publisher.
bit	**bit**
character (synonym: **char**)	**national character** or **ntext** If the length of the data is 255 characters or less, SQL Server CE replication maps the SQL Server **character** data to SQL Server CE **national character** ; otherwise, it maps it to SQL Server CE **ntext**. If the length of the **ntext** data exceeds the length of the **character** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
character varying (synonyms: **char varying**, **varchar**)	**national character varying** or **ntext** If the length of the data is 255 characters or less, SQL Server CE replication maps the SQL Server **character varying** data to SQL Server CE **national character varying**; otherwise, it maps it to SQL Server CE **ntext** . If the length of the **ntext** data exceeds the length of the **character varying** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
datetime	**datetime**
decimal	See **numeric**.
double precision	**double precision**
float	**float**
image	**image**

integer (**int 4**) (synonym: **int**)	**integer** (**int 4**)
money	**money**
national character (synonyms: **national char**, **nchar**)	**national character** If the length of the data is 255 characters or less, SQL Server CE replication maps the SQL Server **national character** data to SQL Server CE **national character**; otherwise, it maps it to SQL Server CE **ntext**. If the length of the **ntext** data exceeds the length of the **national character** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
national character varying (synonyms: **national char varying**, **nvarchar**)	**national character varying** If the length of the data is 255 characters or less; SQL Server CE replication maps the SQL Server **national character varying** data to SQL Server CE **national character varying**; otherwise, it maps it to SQL Server CE **ntext**. If the length of the **ntext** data exceeds the length of the **national character varying** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
ntext	**ntext**
numeric (synonyms: **decimal**, **dec**)	**numeric**
real	**real**
smalldatetime	**datetime** If the precision of the **datetime** data exceeds the precision of the **smalldatetime** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
smallint (**int 2**)	**smallint** (**int 2**)
smallmoney	**money** If the precision of the **money** data exceeds the precision of the **smallmoney** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
sql_variant	**sql_variant** is mapped to **ntext**. If binary data exists in the **sql_variant** column, the binary data must be an even number of bytes or a conversion error occurs.
text	**ntext** If the length of the **text** data exceeds 1,073,741,823 characters, synchronization fails when the data is sent from the SQL Server publication to the SQL Server CE subscription.

timestamp	Not supported Data stored as **timestamp** data is vertically partitioned out of all SQL Server CE Subscriptions.
tinyint (**int 1**)	**tinyint** (**int 1**)
uniqueidentifier	**uniqueidentifier**
varbinary (n)	**varbinary** (n) or **image** If the length of the data is 510 bytes or less, SQL Server CE replication maps the SQL Server **varbinary** data to SQL Server CE **varbinary** ; otherwise, it maps it to SQL Server CE **image**. If the length of the **image** data exceeds the length of the **varbinary** column, synchronization fails when the data is sent from the SQL Server CE subscription to the SQL Server publication.
varchar	See **character varying**

Choose data types that are supported by both SQL Server and SQL Server CE, so it is not necessary for replication to perform data mapping. When this is not possible, the application should validate the values stored in SQL Server CE to ensure that replication can map these values between SQL Server and SQL Server CE.

You cannot publish a table having an index on **char**, **nchar**, **varchar**, or **nvarchar** columns with lengths greater than 255. Creation of the SQL Server CE subscription fails because these column types are mapped to **ntext**, and a primary key cannot be created on an **ntext** column.

II.3 Replication Limitations

Replication in Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE) has the following limitations:

- Dynamic filters

  When you use dynamic filters, SQL Server CE requires that you use the optimize synchronization option. You can set this option by using either the Create Publication wizard or stored procedure scripts:

- Using the Create Publication wizard

  Select **Yes, minimize the amount of data** to minimize the amount of data sent over the network when dynamic horizontal filters are defined. You cannot change this option for an existing publication.

- Using the sp_addmergepublication script

  Set the **@keep_partition_changes** option to **TRUE**. The **@keep_partition_changes** option specifies whether synchronization optimization occurs.

  For more information, see the topics "Optimizing Synchronization" and "sp_addmergepublication" in SQL Server Books Online.

Additionally, SQL Server CE cannot validate Subscriber information. When configuring a Publication with dynamic filters, you can request that the Subscriber information be validated; however, SQL Server CE ignores this request.

Note The **HostName** property specifies the dynamic filter value for the subscription. If you want to change the dynamic filter value for the Subscriber by specifying a new **HostName** property, you must call the **Reinitialize** method before you synchronize.

- Display of Subscriber names in SQL Server Enterprise Manager

  For SQL Server CE Subscribers listed in the SQL Server Enterprise Manager Replication Monitor, the Merge Agent subscription column is formatted as *subscribername:-agentid*. This guarantees uniqueness for each row in the Merge Agent if the Subscriber name supplied by the SQL Server CE application is not unique across all SQL Server CE subscriptions. The Subscriber name and the agent ID can be used to find the Subscriber ID (GUID) in the **msmerge_agents** system table in the distribution database.

- Connection timeout

  SQL Server CE connectivity solutions recover from communication failures by restarting from the last successfully transmitted block of data, as long as the connection recovers within a time-out period of approximately two minutes. This makes synchronization possible even if the underlying transport is not reliable.

- Snapshots

  SQL Server CE only supports the alternate snapshot location option, if the usual (default) snapshot location is not available.

  To specify an alternate snapshot location, you can set this option by using either the create publication stored procedure script or by using SQL Server Enterprise Manager to edit the properties of an existing publication:

- Using the sp_addmergepublication script

  Set the **@snapshot_in_defaultfolder** option to **FALSE** and **@alt_snapshot_folder**= location of the alternate folder for the snapshot.

- Editing Publication Properties of an existing publication

  On the **Snapshot Location** tab of the publications **Properties** dialog box, clear the **Generate snapshots in the normal snapshot folder** check box, select **Generate snapshots in the following location** check box, and specify the new location in **Folder**. You cannot change this option for an existing publication.

  Note SQL Server CE does not support FTP or compressed snapshot options for either the default or alternate locations.

- Ranged identity columns

  SQL Server CE identity columns must have a data type of **integer** (**int 4**). SQL Server CE identity columns cannot have a data type of **bigint**, **smallint**, **tinyint**, **decimal**, or **numeric**. If you subscribe to a publication having an identity column other than **integer** (**int 4**), the creation of that subscription fails on SQL Server CE.

  SQL Server CE allows you to modify the seed and increment values at the Subscriber by using the ALTER TABLE *table_name* ALTER COLUMN *column_name* IDENTITY (*seed,increment*) statement. This allows you to manage identity ranges manually, if you want. However, if your publication includes an identity column and the **AutoIdentityRange** property at the Publisher is TRUE, you should not modify the seed or increment values at the Subscriber. If you specify a seed that is above your allocated range identity, SQL Server CE returns an error when you attempt to insert a new record in the table. When you next synchronize, the Publisher corrects the problem by assigning your Subscriber a new identity range.

- Computed columns

  Not supported. Data stored as computed columns is vertically partitioned out of all SQL Server CE subscriptions.

- Information that is not propagated to the SQL Server CE Subscriber

  You can include the following items in a SQL Server publication, but they are not propagated to the SQL Server CE Subscriber:

  - CHECK constraints
  - DEFAULT definitions for columns
  - Extended properties
  - Stored procedures
  - Views
  - User-defined functions
  - Triggers

  Because SQL Server CE replication cannot propagate these items, you should implement equivalent logic in a SQL Server CE-based application. Doing so ensures that the SQL Server CE database remains consistent with the SQL Server database. For example, if the SQL Server database includes a CHECK constraint, the SQL Server CE-based application should implement the corresponding check in application code.

- Case sensitivity

  SQL Server CE databases are not case-sensitive; therefore objects and data are treated differently in a SQL Server CE database from the way in which they are treated in a SQL Server database. For example, two tables, **MYTABLE** and **mytable**, are treated as different objects in a case-sensitive SQL Server database. The same table names cause a naming conflict in a SQL Server CE database. Similarly, two values, **MYVALUE** and **myvalue**, cause a primary key violation in SQL Server CE, because they are treated as the same value, whereas they do not cause a violation in a case-sensitive SQL Server database.

- UNIQUE constraints

  SQL Server and SQL Server CE differ slightly in the way they implement UNIQUE constraints. SQL Server allows a single NULL value in a column with a UNIQUE constraint; SQL Server CE does not allow a NULL value in this case. If you subscribe to a publication that contains a NULL value in a column with a UNIQUE constraint, the creation of that subscription fails on SQL Server CE.

- Tables and columns

  For all replication tables, the number of available columns is 252 (255 – 3 system columns). SQL Server CE allocates three system columns for tracking.

- NOT FOR REPLICATION constraints

  SQL Server CE does not support the NOT FOR REPLICATION option. Do not create constraints using this option. If constraints in a database have the NOT FOR REPLICATION option, remove the constraint and then re-create it. If the NOT FOR REPLICATION option is specified, the constraint is still created on the SQL Server CE Subscriber, but it does not include the NOT FOR REPLICATION syntax.

II.4 Schema Changes on Publication Databases

You can add and drop columns in a table published in a Microsoft® SQL Server™ publication without dropping and re-creating the publication. If you add or drop a column in a published SQL Server table, that schema change is published to the subscription database in Microsoft SQL Server 2000 Windows® CE Edition (SQL Server CE) when the publication is next synchronized. You must make schema changes to a published table through the replication publication properties dialog box in SQL Server Enterprise Manager or through replication stored procedures. You cannot make schema changes to published tables in any other way. For more information, see "Schema Changes on Publication Databases" in SQL Server Books Online.

If a column is dropped from the SQL Server publication database, any changes made to data in that corresponding column in the SQL Server CE subscription database (that have not already been propagated to the Publisher) are discarded during synchronization.

You can add articles to an existing publication. When an article is added, the SQL Server CE Subscriber must be reinitialized.

Important The SQL Server CE Subscriber must also be reinitialized if the column that is added to an existing published SQL Server table is an identity column.

II.5 Replication Conflict Detection and Resolution

When Publishers and Subscribers are reconnected and synchronization occurs, the Merge Agent running on the Publisher detects conflicts and then determines which data is accepted and propagated to other sites based on a resolver specified when the publication was created. Each Subscriber merges with the Publisher, but conflicts typically occur between updates made at different Subscribers.

Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE) supports row-level tracking only. The SQL Server CE Database Engine invokes tracking when a row is inserted, updated, or deleted. When conflicts are detected at the row level, changes made to corresponding rows are considered a conflict, regardless of whether the changes are made to the same column. For example, a change is made to the address column of a row at the Publisher, and another change is made to the phone number column of the corresponding row in the same table at the Subscriber. With row-level tracking, a conflict is detected because changes were made to both rows.

Note Even if the Publisher is configured with column-level tracking, SQL Server CE masks the row as if every column has changed. Row-level tracking involves less tracking overhead, preserving valuable storage space on the device.

Because SQL Server CE implements row-level tracking only, column-level tracking is ignored, even if it is configured on the Publisher. Therefore, a conflict occurs even if the columns affected in the corresponding rows are not the same.

After a conflict is detected, the Merge Agent launches the conflict resolver selected for the article. In addition to the default resolver, merge replication allows you to use a variety of custom resolvers to deal with conflict situations. The accepted changes are chosen according to the rules of the conflict resolver. SQL Server CE Subscriber conflicts are always detected, resolved, and logged at the Publisher. For more information about conflict detection and resolution, see "Merge Replication Conflict Detection and Resolution" in SQL Server Books Online.

III. Implementing Replication

The topics in this section describe the configuration requirements for using replication in Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE). The topics contain procedures for performing configuration tasks; some procedures are optional, depending on your requirements.

To implement SQL Server CE replication, you must configure SQL Server. Configuring SQL Server involves the following steps:

III.1 Creating the Publication

You create publications for Microsoft® SQL Server™ 2000 Windows® CE Edition (SQL Server CE) by using standard SQL Server replication tools and techniques. The easiest way to create a publication is by using the Create Publication Wizard. For more information, see "Replication Wizards" in SQL Server Books Online.

When you create a publication for SQL Server CE Subscribers by using the Create Publication Wizard, it is important to select the following options:

- On the Select Publication Type page, select **Merge publication**.
- On the Specify Subscriber Types page, select the **Devices running SQL Server CE** check box.

By selecting **Devices running Windows CE**, the following options are automatically configured for the publication:

- Support for anonymous subscriptions
- Character mode snapshot

If you use dynamic filters, SQL Server CE requires that you use the optimize synchronization option. Set this option by using either the Create Publication Wizard or the stored procedure **sp_addmergepublication**.

If this is the first publication created for this server, the Create Publication Wizard prompts you to configure the snapshot folder location. The snapshot folder holds the snapshot files that contain the schema and data for published tables. The SQL Server CE Replication Provider must be able to read these snapshot files so it can download the files to the Windows CE-based device.

Note The snapshot folder location defaults to an administrative share (for example, C$); you should use an explicit share instead of the default because the administrative share can be accessed only by an administrator account.

III.2. Securing the Publication

Securing the Microsoft® SQL Server™ publication involves the following steps:

a. Configuring Database Access

When you create a publication, Microsoft® SQL Server™ creates a publication access list (PAL) for the publication. The PAL is similar to an access control list (ACL). To add a user to the publication access list, you must first grant database access based on how you configured Microsoft Internet Information Services (IIS) and SQL Server authentication.

To configure the Database Access option

1. On the computer running SQL Server, start SQL Server Enterprise Manager: On the **Start** menu, point to **Programs**, **Microsoft SQL Server**, and then click **Enterprise Manager**.
2. In the console tree, expand **Microsoft SQL Servers**, **SQL Server Group**, the computer containing the publication, and **Security**; and then click **Logins**.
3. In the details pane, right-click the user ID under which the Merge Agent executes, and then click **Properties**. The user ID of the Merge Agent depends on how you configured IIS and SQL Server authentication.

IIS authentication mode	**SQL Server authentication mode**	**Must grant database access to**
Anonymous access	Windows authentication	*computername***IUSR_***computername** or the configured IIS anonymous user account.
Basic authentication	Windows authentication	The IIS client's user or group.
Integrated Windows authentication	Windows authentication	The IIS client's user or group. (This configuration is supported only if SQL Server and IIS are running on the same computer.)
Anonymous, Basic,	SQL Server	The user specified in the

or Integrated Windows authentication	authentication	**DistributorLogin** or **PublisherLogin** property of the Replication object.

4. Click the **Database Access** tab, and then select the **Permit** check box for the database.

b. Configuring the Publication Access List

When you create a publication, Microsoft® SQL Server™ creates a publication access list (PAL) for the publication. The PAL is similar to an access control list (ACL). It contains the list of logins that have been granted access to the publication. During synchronization, SQL Server checks the PAL for the Microsoft Internet Information Services (IIS) user, in the case of Windows authentication, or the SQL Server login, in the case of SQL Server authentication.

By default, the PAL contains the members of the **sysadmin** fixed server role and the login of the user who created the publication. Expand or restrict access to a publication by adding or deleting logins from the PAL. You must grant access to the publication based on how you configured IIS and SQL Server authentication.

To grant or revoke access to a publication

1. On the computer running SQL Server, start SQL Server Enterprise Manager: On the **Start** menu, point to **Programs**, **Microsoft SQL Server**, and then click **Enterprise Manager**.
2. In the console tree, expand **Microsoft SQL Servers**, **SQL Server Group**, the computer containing your publication, **Databases**, and the database you are publishing, and then click **Publications**.
3. Right-click the publication name, click **Properties**, and then click the **Publication Access List** tab.
4. Use the **Add** or **Remove** buttons to add or delete logins from the PAL. You must add the user to SQL Server before you can add the user to the PAL. For more information, see "Adding a Windows NT User or Group" in SQL Server Books Online.

IIS authentication mode	**SQL Server authentication mode**	**The PAL must grant access to**
Anonymous access	Windows authentication	*computername***IUSR_***computername** or the configured IIS anonymous user account.
Basic authentication	Windows authentication	The IIS client's user or group.
Integrated Windows authentication	Windows authentication	The IIS client's user or group. (This configuration is only supported if SQL Server and IIS are running on the same computer.)
Anonymous, Basic, or Integrated Windows authentication	SQL Server authentication	The users specified in the **DistributorLogin** and **PublisherLogin** properties of the Replication object.

5. **Note** If the Publisher and Distributor are on different computers, the new logins must exist in the PALs for both the Publisher and the Distributor. If the IIS user, in the case of Windows authentication, or the SQL Server login, in the case of SQL Server authentication, is not contained in the PAL, an Access Denied message is displayed at the Subscriber.

c. Configuring the Check Permissions Option

Use the **Check Permissions** option to provide an enhanced level of security to your publication. The **Check Permissions** option ensures that the Merge Agent has the authority to upload data changes to a Publisher. When you specify this option, Microsoft® SQL Server™ verifies that the Merge Agent login has the permissions to perform INSERT, UPDATE, and DELETE statements on the publication database.

Using **Check Permissions** is optional. When you specify **Check Permissions**, both **Check Permissions** validation and PAL verification are performed.

If you set the **Check Permissions** option after the initial snapshot has been generated, a new snapshot must be generated and reapplied at the Subscriber for permissions to be validated when data changes are merged.

To configure the Check Permissions option

1. On the computer running SQL Server, start SQL Server Enterprise Manager: On the **Start** menu, point to **Programs**, **Microsoft SQL Server**, and then click **Enterprise Manager**.
2. In the console tree, expand **Microsoft SQL Servers**, **SQL Server Group**, the computer containing your publication, **Databases**, and the database you are publishing; and then click **Publications**.
3. Right-click your publication name, click **Properties**, and then click the **Articles** tab.
4. Click the properties button (**...**) for the published article you want to configure. This displays the **Table Article Properties** dialog box.
5. Click the **Merging Changes** tab.
6. Under **Check permissions**, you can select the **INSERT**, **UPDATE**, and **DELETE** check boxes. When these boxes are selected, SQL Server verifies that the Merge Agent is running under a user identity that has permission to insert, update, or delete data on the Publisher.

If you specify **Check permissions**, you must grant database access based on how you configured Microsoft Internet Information Services (IIS) and SQL Server authentication.

III.3. Configuring the Snapshot Folder

The snapshot folder holds the snapshot files containing the schema and data for published tables. The SQL Server CE Replication Provider must be able to read these snapshot files to create the subscription on the Microsoft® Windows® CE-based device. You control access to the snapshot folder by assigning file system access permissions to it.

When you assign Read permissions to the snapshot folder and share, remember that the identity under which the SQL Server CE Replication Provider runs depends upon how Microsoft Internet Information Services (IIS) authentication is configured:

- When IIS is configured to use Anonymous access, the SQL Server CE Replication Provider runs under the identity of the default IIS anonymous user account (the Internet Guest Account: **IUSR_***computername*). If you configure another Windows user account as the IIS anonymous user account, the SQL Server CE Replication Provider runs under the identity of that account.
- When IIS is configured to use Basic authentication, the SQL Server CE Replication Provider runs under the identity of the Windows user account corresponding to the login and password supplied by the Subscriber for the **InternetLogin** and **InternetPassword** properties of the Replication object.
- When IIS is configured to use Integrated Windows authentication, the SQL Server CE Replication Provider runs under the identity of the Windows user account corresponding to the login and password supplied by the Subscriber for the **InternetLogin** and **InternetPassword** properties of the Replication object. Unlike Basic authentication, the login and password are never transmitted over the network with Integrated Windows authentication.

Whether IIS and the SQL Server Distributor are located on the same computer or different computers, it is strongly recommended that you configure the snapshot folder as described in Configuring an Explicit Snapshot Folder.

III.4. Checking the Status of the Snapshot Agent

You can use the Replication Monitor in SQL Server Enterprise Manager to check the status of the Snapshot Agent and ensure that the snapshot has been created.

To check the status of the Snapshot Agent

1. On the computer running Microsoft® SQL Server™, start SQL Server Enterprise Manager: On the **Start** menu, point to **Programs**, **Microsoft SQL Server**, and then click **Enterprise Manager**.
2. In the console tree, expand **Microsoft SQL Servers**, **SQL Server Group**, the computer containing the Distributor for the publication, and **Replication Monitor**; and then click **Agents**.
3. In the left pane, click **Snapshot Agents**. The status of the Snapshot Agent is displayed in the right pane.
4. To display more information about an entry listed in the right pane, double-click the entry.

If the Snapshot Agent for the publication has not started, in the right pane, right-click the agent entry and select **Start Agent**.

Reference:

http://msdn.microsoft.com/library/default.asp?url=/library/en-us/sqlce/htm/_lce_connectivity.asp