CHAPTER 1

Configuration Manager site hierarchy and distribution points

Microsoft System Center 2012 Configuration Manager helps empower people to use the devices and applications they need to be productive, while maintaining corporate compliance and control. It accomplishes this with a unified infrastructure that acts like a single pane of glass to manage physical, virtual, and mobile clients. It also provides tools and improvements that make it easier for IT administrators to do their jobs.

A good understanding of basic Configuration Manager concepts, processes, and practices is essential for being able to effectively troubleshoot problems when they arise. This chapter provides an overview of the Configuration Manager site hierarchy including determining when to use a central administration site, primary site, and secondary site; inter-site replication; distribution points; Active Directory requirements for sites; Forest Discovery and Publishing; boundaries and boundary groups; and cross-forest scenarios. The chapter also describes some best practices for installing the central administration site, primary sites, and secondary sites; performing unattended installations of sites; and using Prerequisite Checker. The chapter then concludes with some troubleshooting tips concerning database replication and the Configuration Manager console.

Configuration Manager site hierarchy

The Configuration Manager site hierarchy consists of the following site system roles:

- Central administration site
- Primary sites
- Secondary sites

There are other Configuration Manager roles, such as management point, distribution point, and so on, but this chapter mainly focuses on these three roles.

Central administration site

When setting up a Configuration Manager hierarchy, the central administration site is the first one you must install. The central administration site is always on the top of the hierarchy and cannot be joined or moved to an existing hierarchy. You can have only one central administration site per hierarchy.

The central administration site coordinates inter-site data replication across the hierarchy by using Configuration Manager database replication. The central administration site also allows the administration of hierarchy-wide configurations for client agents, discovery performance, and other operations.

The following are some considerations when deploying a central administration site:

- Used for all administration and reporting for the hierarchy
- Used for administration purposes only
- Can support up to 25 child primary sites simultaneously
- Participates in SQL database replication
- Does not process client data and does not support client assignment
- Does not support a management point so no client can report to this site
- Not all site system roles are available

Primary sites

Primary sites can be used to manage clients in well-connected networks. Primary sites cannot be tiered below other primary sites.

The following are some considerations when deploying primary sites:

- Can be a stand-alone site or a member of a hierarchy
- Only supports a central administration site as a parent site
- Use database replication to communicate directly to their central administration site
- Automatically configures database replication with its designated central administration site
- Attached directly to the central administration site
- Only supports secondary sites as child sites
- Can support up to 250 secondary sites
- Are responsible for processing all client data from their assigned clients
- Can have up to 100,000 clients attached to them

Secondary sites

Secondary sites control content distribution for clients in remote locations across links that have limited network bandwidth. The following are some considerations when deploying secondary sites:

- Are used to host site system roles to offload WAN link traffic.
- Can only be installed (pushed) from the Configuration Manager console.
- Can communicate with clients but never have clients assigned to them. A management point and distribution point are automatically deployed during the site installation.
- Can distribute content to other secondary sites (new in Configuration Manager 2012).
- Cannot report to another secondary site.

Determining when to use a central administration site

You should install a central administration site if you plan to install multiple primary sites. Use a central administration site to configure hierarchy-wide settings and to monitor all sites and objects in the hierarchy. This site type does not manage clients directly, but it does coordinate inter-site data replication, which includes the configuration of sites and clients throughout the hierarchy.

You can manage all clients in the hierarchy and perform site management tasks for any primary site when you use a Configuration Manager console that is connected to the central administration site. The central administration site is the only place where you can see site data from all sites. This data includes information such as inventory data and status messages.

You should configure discovery operations throughout the hierarchy from the central administration site by assigning discovery methods to run at individual sites. You can manage security throughout the hierarchy by assigning different security roles, security scopes, and collections to different administrative users. These configurations apply at each site in the hierarchy. You can configure addresses to communicate between sites in the hierarchy. This includes settings that manage the schedule and bandwidth for transferring file-based data between sites.

Consider installing a central administration site for any of the following reasons:

- When more than one primary site is present in a hierarchy
- When you need to scale-up the number of clients that can be managed
- When you need to off-load reporting and administration from your primary site
- When you need to monitor and report from all sites and objects in the hierarchy

NOTE **A central administration site can be installed only as a new installation.**

Determining when to use a primary site

Consider installing a primary site for any of the following reasons:

- When you need to manage clients directly.
- When you need to increase the number of clients you can manage. Each primary site can support up to 100,000 clients.
- When you need to reduce the possible result of failure of a single primary site.
- When you need to provide load-balancing support for clients across multiple servers.
- When you need to provide a local point of connectivity for administration.
- When you need to meet organizational management requirements. For example, you might install a primary site at a remote location to manage the transfer of deployment content across a low-bandwidth network.

Determining when to use a secondary site

Consider installing a secondary site for any of the following reasons:

- When a local administrative user is not required in a location
- When you need to manage the transfer of deployment content across low-bandwidth networks
- When you need to manage the transfer of client data across low-bandwidth networks
- When you need to establish tiered content routing for deep network topologies

Understanding site-to-site replication

Site-to-site replication is running behind the scenes when you create a collection, a package, or folders on a central administration site. The central administration site replicates that information using database replication to primary sites, and then the primary sites replicate their secondary sites.

The basic concepts and components involved in the process of replication of global and site data are as follows.

- **Database replication** Performs all non-content related site-to-site transfer of information such as inventory data, status messages, and Windows Server Update Services (WSUS) metadata. When you deploy a secondary site, Microsoft SQL Server Express is installed and used for replicating Configuration Manager data. In Configuration Manager, database replication is now used for replication between sites in all cases except for when data flows from a secondary site to a primary site; for that process the file-based replication employed by Configuration Manager 2007 is still used. In addition, file-based replication is used to initialize/re-initialize database replication by copying exported SQL data to another site server.

- **Global data** Global data is replicated to all primary sites but only a subset of it is replicated to secondary sites.
- **Site data** Site data is replicated between the primary site where clients are assigned and the central administration site.
- **Content** The content replicated to child sites includes software deployment packages and software update packages.

NOTE **If a data discovery record (DDR) is newly generated by Active Directory System Discovery in a child primary site, it is sent to the central administration site in the standard way by passing the DDR record up the hierarchy. After the DDR has been added to the database at the central administration site, any updates to the DDR information will use database replication to replicate.**

Global and site data

The differences between global and site data can be summarized as follows:

- Global data objects are replicated everywhere and consist of collections and their rules, packages, CIs, software updates, deployment data, and so on.
- Site or local data is replicated only between the primary site that created it and the central administration site. Site data, in general, is the data created by the system, for example data concerning collection membership (built by Collection Evaluator component) or hardware inventory (built by the client and processed by data loader).

Database replication

Configuration Manager database replication transfers data quickly and guarantees delivery by using SQL Service Broker (SSB) and SQL Change Tracking. However, this has nothing to do directly with SQL database replication technology from a Configuration Manager standpoint.

Database replication in Configuration Manager is more reliable than the file-based replication used in Configuration Manager 2007 which is based on file transfer using the Server Message Block (SMB) protocol. Configuration Manager 2007 environments sometimes used to experience site-to-site communication issues caused by anti-virus software when the customer environment did not have the proper filter configured in their anti-virus scanning software.

Database replication in Configuration Manager has the following characteristics:

- It can be one-way (for example, site data) or bi-directional (for example, global data).
- Site data is replicated from a primary site to a central administration site.
- Global data is replicated to all site servers. For example, a collection created on one primary site will show up in another primary site because the collection rule is global data.

File-based replication

As described previously, file-based replication is still used in certain circumstances by Configuration Manager. During the file-based replication process, files are transferred using the SMB protocol and TCP port 445. Filed-based replication works like this:

1. The sending component places a file into the Replmgr's outbound folder.
2. Replmgr creates a job file and places it into the Ready folder.
3. Scheduler picks up the job file and creates the sending request file, and generates the compressed files, which will be sent to the tosend folder.
4. Based on the job priority and other settings, the scheduler triggers the sender to write the files from the tosend folder to the despooler folder on the receiving site server.

Understanding distribution points

A *distribution point* is a computer designed to deliver binary files/packages to Configuration Manager clients. Examples of such binary files can include applications, operating system deployment images, boot images, software updates, and so on.

In Configuration Manager, distribution points now use a new storage format called the content library. The content library replaces the SMSPKG shares as the default folder structure used to host content. The content library stores all content on the distribution point using single instance storage; this means each unique file is only stored once on the distribution point, regardless of how many times it is referenced by a package. In addition, the file is stored only once on the distribution point even if it is contained in multiple packages.

You should use a distribution point instead of a secondary site if:

- You are not concerned about network usage due to clients pulling policy or reporting status, inventory, or discovery to their primary site location. In this case you would use a distribution point instead of secondary site.
- You want to leverage Background Intelligent Transfer Service (BITS) access for clients, for example to use BranchCache, Operating System Deployment (OSD) multicasting, or Application Virtualization (APP-V) streaming.
- You find that client-side BITS does not provide enough bandwidth control for your WAN.

Active Directory requirements for sites

To install any Configuration Manager site, such as a central administration site, primary site, or secondary site, the server needs to be a member of an Active Directory domain. Though the Active Directory schema extension is optional, it is highly recommended that you extend the schema for Configuration Manager.

Active Directory schema extension

Extending the Active Directory schema is a forest-wide action and can only be done one time per forest. The following are some considerations for Active Directory schema extension in Configuration Manager environments:

- All Configuration Manager site systems must be members of an Active Directory domain.
- Configuration Manager Active Directory schema extensions provide many benefits for Configuration Manager sites, but they are not required for all Configuration Manager functions.
- If you have extended your Active Directory schema for Configuration Manager 2007, you do not have to update your schema for Configuration Manager.
- You can update the Active Directory schema before or after you install Configuration Manager.
- Schema updates do not interfere with an existing Configuration Manager 2007 site or clients.

Disjoint namespaces

A disjoint namespace happens when one or more domain member computers have a primary Domain Name Service (DNS) suffix that does not match the DNS name of the Active Directory domain of which the computers are members. For example, a member computer that uses a primary DNS suffix of corp.contoso.com in an Active Directory domain named na.corp.contoso.com is using a disjoint namespace.

The following are some considerations for disjoint namespaces in Configuration Manager environments:

- With the exception of out-of-band management, Configuration Manager supports installing site systems and clients in a domain that has a disjoint namespace.
- To allow a computer to access domain controllers that are disjoint, you must modify the msDS-AllowedDNSSuffixes Active Directory attribute on the domain object container. You must add both of the DNS suffixes to the attribute.
- To ensure that the DNS suffix search list contains all DNS namespaces that are deployed within the organization, you must configure the search list for each computer in a domain that is disjoint. Include in the list of namespaces the primary DNS suffix of the domain controller, the DNS domain name, and any additional namespaces for other servers with which Configuration Manager might interoperate. You can use Group Policy to configure the DNS suffix search list.

Single label domains

Single-label domain names are DNS names that do not contain a suffix such as .com, .corp, .net, or .org. For example contoso would be a single-label domain name while contoso.com, contoso.net, or contoso.local would not be single-label domain names. Configuration Manager does not support single-label domain names.

Extending the schema for Configuration Manager

You can extend the schema during setup, by using the Extadsch.exe command line tool, or by using the LDIFDE tool. The schema changes are stored in \SMSSETUP\BIN\x64\ConfigMgr_ad_schema.ldf

NOTE **The schema does not need to be extended again for Configuration Manager 2012 and later, if it has already been extended for Configuration Manager 2007.**

Forest Discovery and Publishing

In order to guarantee that clients are correctly assigned to Configuration Manager sites, and to guarantee that all software, software updates, and operating system images are available to Configuration Manager clients, it is necessary to make sure that the boundaries in Configuration Manager and Active Directory are correctly configured. Up-to-date boundary information results in efficient deployment of applications and software updates to managed client computers. Forest Discovery and Publishing helps clients not only discover the sites in the forest, but also publish existing sites that can manage clients across domains, thus eliminating the need to deploy additional sites.

Forest Discovery can discover IP subnets and sites in Active Directory and then add these as boundaries in Configuration Manager. Forest Discovery and Publishing can connect to all of your forests to build a complete map of your Configuration Manager environment. Forest Discovery and Publishing can also cross forest boundaries using specific credentials for each forest regardless of the trust type. The information obtained through Forest Discovery can be directly exported as either boundaries or boundary groups. Changes to discovered data are updated dynamically and aged out from the database when no longer present in Active Directory. The discovered data is also used when clients request a management point or distribution point to ensure they receive the best possible site for performance reasons. Credentials specified for each forest are used for both discovery and publishing and enable Configuration Manager sites to publish site information in both trusted and untrusted forests.

Publishing stores information, such as site system locations and capabilities, boundaries, and security information, required by client computers to establish trusted connections with site systems. It also stores information such as the client's trust relationship with the forest, and the management point's communication mode (HTTPS/HTTP) and the boundaries that are used to locate the most appropriate management point or distribution point to communicate with. This enables client computers to locate servers in a trusted forest to ensure user-targeted applications are successful.

IMPORTANT **As in Configuration Manager 2007, supernetting is not supported in Configuration Manager. However, when you run Active Directory Forest Discovery to discover your IP subnets it creates IP address ranges based on the subnet and mask defined in Active Directory.**

Boundaries and boundary groups

Boundaries represent network locations on the intranet where Configuration Manager clients are located. Boundary groups are logical groups of boundaries that provide clients access to resources. The sections below summarize some considerations concerning boundaries and boundary groups.

Boundaries

Each boundary represents a network location in Configuration Manager and is available from every site in your hierarchy. A boundary alone, however, does not enable you to manage clients at that network location. To manage a client, the boundary must also be a member of a boundary group. Boundaries can be any of the following:

- IP range
- IP subnet
- Active Directory site
- IPv6 prefix
- Boundary group for site assignment and/or content location

IMPORTANT **Overlapping site boundaries are supported for content location but are not supported for site assignment.**

Boundary groups

Boundary groups are used to manage your network locations. You must assign boundaries to boundary groups before you can use the boundary group. Boundary groups have the following functions:

- They enable clients to find a primary site for client assignment (automatic site assignment).
- They can provide clients with a list of available site systems that have content after you associate the distribution point and state migration point site system servers with the boundary group.
- To support site assignment, you must configure the boundary group to specify an assigned site for clients to use during automatic site assignment. To support content location, you must specify one or more site systems. You can only specify site systems with the distribution point or state migration point site system role. Both the site assignment and content location configurations are optional for boundary groups.
- When you plan for boundary groups, consider creating one set of boundary groups for content location and a second set of boundary groups for automatic site assignment. This separation can help you avoid overlapping boundaries for site assignment. When you have overlapping boundaries and use automatic site assignment, the site to which a client is assigned, might be too nondeterministic.

Cross-forest scenarios

Several cross-forest scenarios are possible when administering Configuration Manager environments:

- Simple client management in a different Active Directory forest. This scenario involves no object discovery, no added infrastructure, and manual client deployment.
- Managing clients using discovery and performing client push installations. This scenario involves cross-forest site publishing, cross-forest system discovery, and automated client installation. However, it does not add any additional infrastructure into the remote untrusted forest.
- Implementing child primary or secondary sites in a cross-forest environment. This requires a two-way forest trust.
- Installing site roles such as management point, software update point, and distribution point in a cross-forest environment.

Cross-forest tips

The following are some tips for cross-forest scenarios:

- Inner-site communication (site-to-site communication) can use both file-based replication (SMB Port 445) and database replication (TCP/IP port 4022 by default) so configure your perimeter network firewalls accordingly.
- Site system roles (management point, distribution point, and so on) with the exception of the out-of-band service point and the application catalog web service point can be deployed in an untrusted forest.
- The Server Locator Point (SLP) functionality is now performed by a management point.
- Each Configuration Manager site can only host two software update points, one for intranet located clients and one for internet located clients. This needs to be considered when designing a multiforest (non-trusted) Configuration Manager site.
- You can add the forest you need on the Configuration Manager console through the Active Directory Forest Discovery method.
- You can use Publish to publish information to the client's Active Directory forest.
- To install and configure a child site (primary or secondary), the child site server must be located in the same forest as the parent site or reside in a forest that contains a two-way trust with the forest of the parent (central administration or primary) site.

Client approval

After client installation, the client remains in an unapproved state if you are using the default setting Automatically Approve Computers In Trusted Domains. You will therefore need to approve such clients after their installation.

Using Prerequisite Checker

The Prerequisite Checker (prereqchk.exe) is a stand-alone application that verifies server readiness for a site server or specific site system roles. Before site installation, Setup runs the Prerequisite Checker (see Figure 1-1). You can manually run the Prerequisite Checker on potential site servers or site systems to verify server readiness. This allows you to remediate any issues that you find before you run Setup.

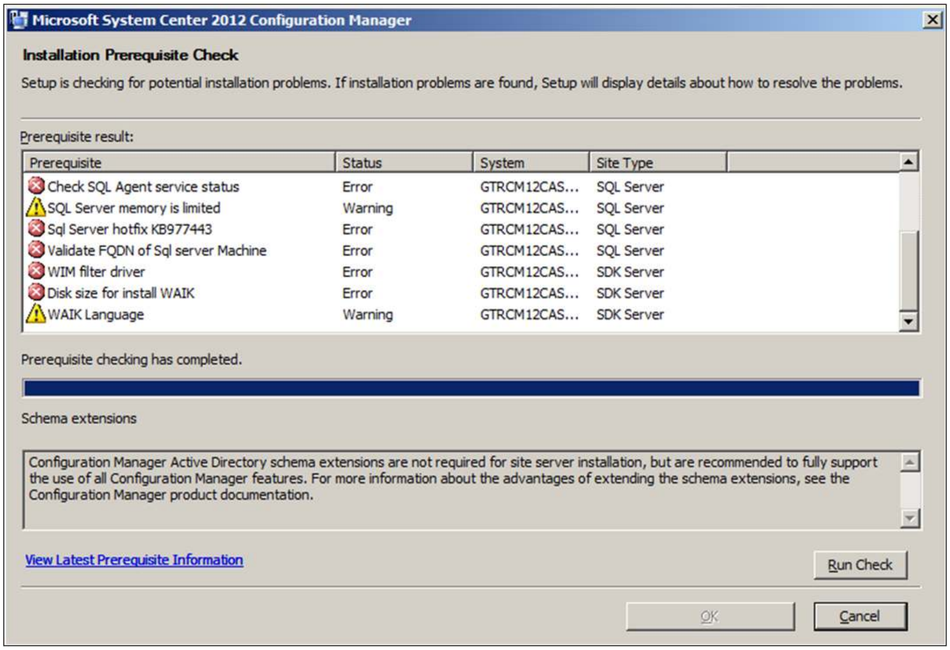

Microsoft System Center 2012 Configuration Manager

Installation Prerequisite Check

Setup is checking for potential installation problems. If installation problems are found, Setup will display details about how to resolve the problems.

Prerequisite result:

Prerequisite	Status	System	Site Type
Check SQL Agent service status	Error	GTRCM12CAS...	SQL Server
SQL Server memory is limited	Warning	GTRCM12CAS...	SQL Server
Sql Server hotfix KB977443	Error	GTRCM12CAS...	SQL Server
Validate FQDN of Sql server Machine	Error	GTRCM12CAS...	SQL Server
WIM filter driver	Error	GTRCM12CAS...	SDK Server
Disk size for install WAIK	Error	GTRCM12CAS...	SDK Server
WAIK Language	Warning	GTRCM12CAS...	SDK Server

Prerequisite checking has completed.

Schema extensions

Configuration Manager Active Directory schema extensions are not required for site server installation, but are recommended to fully support the use of all Configuration Manager features. For more information about the advantages of extending the schema extensions, see the Configuration Manager product documentation.

FIGURE 1-1 Error and warning messages are displayed in Prerequisite Checker.

When you run Prerequisite Checker without the command-line options, the local computer is scanned for an existing site server and only the checks that are applicable to the site are run. If no existing sites are detected, then all prerequisite rules are run.

You can run Prerequisite Checker from a command prompt and specify the specific command-line options to perform only checks associated with the site server or site systems specified in the command-line. When you specify another server to check, you must have Administrator rights on the server for Prerequisite Checker to complete the checks.

The following are some tips on using Prerequisite Checker:

- You must have Administrator rights on the server.
- Prerequisite Checker creates a log file ConfigMgrPrereq.log in the root of the C: drive.
- The following files are needed to run Prerequisite tool independently.
  - prereqchk.exe
  - prereqcore.dll
  - basesql.dll
  - basesvr.dll
  - baseutil.dll

  These files are located under <ConfigMgrSourceFiles>\SMSSETUP\BIN\X64\ folder.

Best practices for installing a central administration site or primary site

This section summarizes some best practices when installing a central administration site or a primary site.

Security rights

Before starting the central administration installation, verify that the administrative user who runs Setup has the following security rights:

- Local Administrator rights on the central administration site server computer
- Local Administrator rights on each computer that hosts one of the following:
  - The site database
  - An instance of the SMS Provider for the site
  - A management point for the site
  - A distribution point for the site
- Sysadmin (SA) rights on the instance of SQL Server that hosts the site database

Site naming

Be sure to plan your site codes and site names carefully before you deploy your Configuration Manager hierarchy. Configuration Manager site naming should adhere to the following guidelines:

- Site codes and site names are used to identify and manage the sites in a Configuration Manager hierarchy. In the Configuration Manager console, the site code and site name are displayed in the <*site code*> - <*site name*> format.
- Every site code that you use in your Configuration Manager hierarchy must be unique. If the Active Directory schema is extended for Configuration Manager, and sites are publishing data, the site codes used within an Active Directory forest must be unique even if they are being used in a different Configuration Manager hierarchy or if they have been used in previous Configuration Manager installations.
- During Configuration Manager Setup, you are prompted for a site code and site name for the central administration site, and each primary and secondary site installation. The site code must uniquely identify each Configuration Manager site in the hierarchy. Because the site code is used in folder names, never use Microsoft Windows reserved names for the site code, such as AUX, CON, NUL, or PRN.
- To enter the site code for a site during Configuration Manager Setup, you must enter three alphanumeric characters. Only the letters A through Z, numbers 0 through 9, or combinations of the two are allowed when specifying site codes. The sequence of letters or numbers has no effect on the communication between sites. For example, it is not necessary to name a primary site ABC and a secondary site DEF.

- The site name is a friendly name identifier for the site. Use only the standard characters A through Z, a through z, 0 through 9, and the hyphen (-) in site names.

IMPORTANT **Changing the site code or site name after installation is not supported.**

Evaluation media

If you install Configuration Manager as an evaluation edition, after 180 days the Configuration Manager console becomes read-only until you activate the product with a product key from the Site Maintenance page in Setup.

Best practices for installing a secondary site

This section summarizes some best practices when installing secondary sites.

Security rights

To create a secondary site, verify the user that runs Setup has the following security rights:

- Local Administrator rights on the secondary site computer
- Local Administrator rights on the remote site database server for the primary site (if it is remote)
- Infrastructure Administrator or Full Administrator security role on the parent primary site

Other considerations

Some other considerations when installing secondary sites include:

- You need to specify a site code for the secondary site because it will be participating of the Configuration Manager hierarchy.
- When you choose the Use The Source Files At The Following Network Location or Use The Source Files At The Following Location On The Secondary Site Computer options, the location must contain the folder named Redist as a subfolder with the prerequisite redistributables, language packs, and the latest product updates for Setup.
- Use Setup Downloader to download the required files to the named folder Redist before you install the secondary site. Secondary site installation will fail if the files are not available in the Redist subfolder.
- Secondary site will use SQL Server Express or an existing SQL Server instance for the site database, and then configure the associated settings. Install and configure a local copy of SQL Express on the secondary site computer.

Unattended installation of a central administration site or primary site

By default, an unattended script named ConfigMgrAutoSave.ini is saved in the %TEMP% folder. Figure 1-2 shows an example of an unattended script.

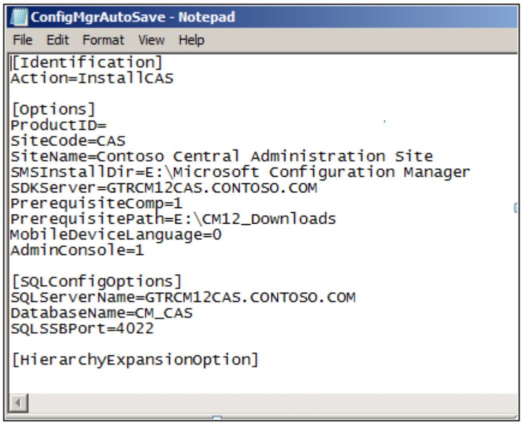

```
[Identification]
Action=InstallCAS

[Options]
ProductID=
SiteCode=CAS
SiteName=Contoso Central Administration Site
SMSInstallDir=E:\Microsoft Configuration Manager
SDKServer=GTRCM12CAS.CONTOSO.COM
PrerequisiteComp=1
PrerequisitePath=E:\CM12_Downloads
MobileDeviceLanguage=0
AdminConsole=1

[SQLConfigOptions]
SQLServerName=GTRCM12CAS.CONTOSO.COM
DatabaseName=CM_CAS
SQLSSBPort=4022

[HierarchyExpansionOption]
```

FIGURE 1-2 An example of a sample unattended script.

Troubleshooting database replication and console issues

This section shares some experience gained from the field concerning troubleshooting database replication and console issues with Configuration Manager.

Troubleshooting database replication

As explained earlier in this chapter, Configuration Manager site-to-site communication uses the SSB feature to replicate data between the site databases instead of the file-based replication used in previous versions of Configuration Manager. With SQL replication, the performance and reliability is improved. However, as a result of this change, it might be a little more difficult for some Configuration Manager administrators to troubleshoot replication issues when they occur. This section provides some tips on how to troubleshoot SQL replication in Configuration Manager by using the following step-by-step approach:

1. Using Replication Link Analyzer
2. Examining the log files
3. Performing SQL queries
4. Reinitiating replication

Step 1: Using Replication Link Analyzer

Your first step for troubleshooting replication should be to use the Replication Link Analyzer. In the Configuration Manager console, start by viewing the current status of the replication links. When you have problems, the first place you should check is the Replication Link Analyzer.

1. Click Monitoring | Overview | Site Hierarchy.

   If the link icon is green, everything is fine. If it is not green, continue with this procedure to use the Replication Link Analyzer to troubleshoot.
2. Click Monitoring | Overview | Database Replication.
3. Select the link.
4. In the lower portion of the window, you can see the detailed status of this link. This information includes whether the replication is active, and the status of the global data replication link and the site data replication link.
5. Right-click the link name to open the Replication Link Analyzer Wizard.
6. Follow the wizard to remediate if necessary, and then review the result files:
   - ReplicationLinkAnalysis.log
   - ReplicationLinkAnalysis.xml

The Replication Link Analyzer works by examining both sites and checking whether:

- The SMS service is running
- The SMS Replication Configuration Monitor component is running
- The ports required for SQL replication are enabled
- The SQL version is supported
- The network is available between the two sites
- There is enough space for the SQL database
- The SSB service configuration exists
- The SSB service certificate exists
- There are any known errors in SQL log files
- There are any replication queues disabled
- Time is in sync
- The transmission of data is stuck
- A key conflict exists

The Replication Link Analyzer can find and fix most but not all database replication problems. If Replication Link Analyzer has not helped you resolve your problem, you should proceed with step 2.

Step 2: Examining the log files

If you are still having difficulties after using the Replication Link Analyzer, your next step should be to check the following two log files for all involved sites:

- rcmctrl.log
- replmgr.log

During the troubleshooting process, you might not get extra details with default logging. You need to turn on verbose logging using the following registry key:

> *HKEY_LOCAL_MACHINE\Software\Microsoft\SMS\Component\SMS_REPLICATION_CONFIGURATION_MONITOR\Verbose logging*

- Set the Value 0 for Errors and key messages (the default value)
- Set the Value 1 for Errors, key messages, warnings and more general information
- Set the Value 2, which is Verbose, to see everything

Step 3: Performing SQL queries

If you are still unable to find the root cause of the issue, you need to run SQL queries using Microsoft SQL Server Management Studio on the central administration site or primary site to get more information. Specifically, you should:

1. Run the spDiagDRS script. The resulting output contains useful information about the general status of the database replication, the current replication link status, and the last sync time for each replication group.
2. Examine the vLogs view. These logs show more detailed information about the process. For example, when the database replication checks for changes, when it receives the BCP (bulk copy data) from the publisher, when it ProcessSyncDataXml, and when a specific table is updated.
3. Check the SSB log found at:

   C:\Program Files\Microsoft SQL Server\MSSQL10.MSSQLSERVER\MSSQL\ErrorLog

For more information on using profiler and SSB-related tables to troubleshoot service broker problems, see *http://www.sqlteam.com/article/how-to-troubleshoot-service-broker-problems*.

Step 4: Reinitiating replication

Reinitiating replication by sending a subscription invalid message should be the last step you try because it causes all the data to be re-replicated between the sites, which will generate a lot of network traffic.

To reinitiate the global data, run the following SQL command:

```
EXEC spDrsSendSubscriptionInvalid 'SiteCode', 'SiteCode', 'Configuration Data'
```

Troubleshooting the Configuration Manager console

Sometimes when you open the Configuration Manager console you will see the warning message shown in Figure 1-3.

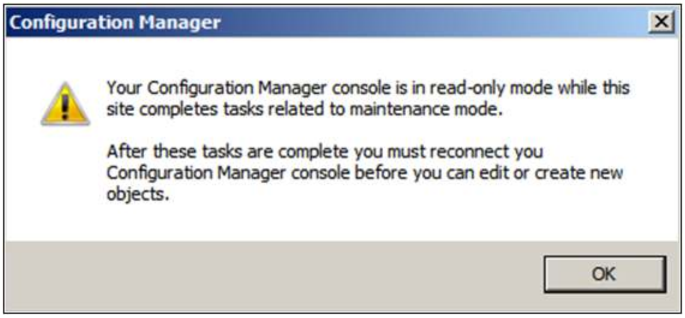

Your Configuration Manager console is in read-only mode while this site completes tasks related to maintenance mode.

After these tasks are complete you must reconnect you Configuration Manager console before you can edit or create new objects.

FIGURE 1-3 A warning indicates when Configuration Manager sessions are read-only.

This warning message tells you that the Configuration Manager console will be opened in Read-Only mode, which means you won't be able to make any changes to your Configuration Manager console. This happens under the following circumstances:

- The primary site did not complete site installation yet.
- The primary site has inter-site replication problems.
- The primary site is running a site restoration.
- The primary site is initializing global data.

In some cases, you will need to wait until the replication site restoration or the site server installation is completed; after that you must close and reconnect the Configuration Manager console to establish a normal session.