CA IT Client Management/ CA Unicenter Desktop & Server Management

ca
Transforming
IT Management

LEGAL NOTICE

This publication is based on current information and resource allocations as of its date of publication and is subject to change or withdrawal by CA at any time without notice. The information in this publication could include typographical errors or technical inaccuracies. CA may make modifications to any CA product, software program, method or procedure described in this publication at any time without notice.

Any reference in this publication to non-CA products and non-CA websites are provided for convenience only and shall not serve as CA's endorsement of such products or websites. Your use of such products, websites, and any information regarding such products or any materials provided with such products or at such websites shall be at your own risk.

Notwithstanding anything in this publication to the contrary, this publication shall not (i) constitute product documentation or specifications under any existing or future written license agreement or services agreement relating to any CA software product, or be subject to any warranty set forth in any such written agreement; (ii) serve to affect the rights and/or obligations of CA or its licensees under any existing or future written license agreement or services agreement relating to any CA software product; or (iii) serve to amend any product documentation or specifications for any CA software product. The development, release and timing of any features or functionality described in this publication remain at CA's sole discretion.

The information in this publication is based upon CA's experiences with the referenced software products in a variety of development and customer environments. Past performance of the software products in such development and customer environments is not indicative of the future performance of such software products in identical, similar or different environments. CA does not warrant that the software products will operate as specifically set forth in this publication. CA will support only the referenced products in accordance with (i) the documentation and specifications provided with the referenced product, and (ii) CA's then-current maintenance and support policy for the referenced product.

Certain information in this publication may outline CA's general product direction. All information in this publication is for your informational purposes only and may not be incorporated into any contract. CA assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA provides this document "AS IS" without warranty of any kind, including, without limitation, any implied warranties of merchantability, fitness for a particular purpose, or non-infringement. In no event will CA be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, lost investment, business interruption, goodwill or lost data, even if CA is expressly advised of the possibility of such damages.

COPYRIGHT LICENSE AND NOTICE:

This publication may contain sample application programming code and/or language which illustrate programming techniques on various operating systems. Notwithstanding anything to the contrary contained in this publication, such sample code does not constitute licensed products or software under any CA license or services agreement. You may copy, modify and use this sample code for the purposes of performing the installation methods and routines described in this document. These samples have not been tested. CA does not make, and you may not rely on, any promise, express or implied, of reliability, serviceability or function of the sample code.

Copyright © 2008 CA. All rights reserved. All trademarks, trade names, service marks and logos referenced herein belong to their respective companies.

TITLE AND PUBLICATION DATE:

CA IT Client Management /CA Unicenter Desktop & Server Management Green Book
Publication Date: March 28, 2008

ACKNOWLEDGEMENTS

THIRD PARTY ACKNOWLEDGEMENTS

CA PRODUCT REFERENCES

This document references the following CA products:

- CA CMDB

- CA Software Compliance Manager

- CA Support Bridge™

- CA Unicenter® Asset Intelligence

- CA Unicenter® Asset Management

- CA Unicenter® Asset Portfolio Management

- CA Unicenter® Desktop DNA®

- CA Unicenter® Network and Systems Management

- CA Unicenter® Patch Management

- CA Unicenter® Remote Control

- CA Unicenter® Service Desk

- CA Unicenter® Software Delivery

FEEDBACK

Please email us at greenbooks@ca.com (mailto:greenbooks@ca.com) to share your feedback on this CA Green Book. Please include the title of this Green Book in the subject of your email response. For technical assistance with a CA product, please contact CA Technical Support at http://ca.com/support (http://www.ca.com/support).

Contents

Chapter 20: Deployment of Microsoft Office 2007 279

Index 297

Chapter 1: Overview

IT departments are responsible for managing increasingly complex desktop environments during this time of unprecedented change. Multiple hardware platforms and disparate operating system versions, software license management, patch management, migration and refresh, and evolving security threats place an enormous management burden on IT. The result is an inconsistent desktop environment that is difficult to maintain and is unaligned with business goals.

This Green Book for CA IT Client Management (CA ITCM) focuses on the product solutions formerly known as CA Unicenter Desktop & Server Management (CA DSM) and IT Resource Management. This book has been created to address a number of topics that are relevant to managing the challenges involved in maintaining the compliance and availability of client and server computing using CA's solution. It provides an explanation of capabilities that you can deploy today to manage your client and server computing environment, and it provides implementation insight and best practice recommendations based upon our experience with numerous customer implementations around the world.

The book focuses on the three core components that make up CA IT Client Management:

- **CA Unicenter Asset Management** proactively manages desktops, laptops, servers, and PDA (personal digital assistant) assets. It provides full-featured asset tracking capabilities through automated discovery, hardware inventory, network inventory, software inventory, configuration management, software usage monitoring, license management, and extensive cross-platform reporting.

- **CA Unicenter Software Delivery** automates the deployment and migration of operating systems, application software, and patches across laptops, desktops, servers, and PDA systems in heterogeneous business environments. From distribution of software to management of system configuration and rollback across multiple platforms and locations, this scalable and secure management solution helps ensure the consistency and reliability of deployment and management of software.

- **CA Unicenter Remote Control** better enables IT administrators to reliably and securely access, control, view, manage, and modify remote systems. It offers exceptional features to facilitate Windows, Linux, and Macintosh administration, help desk support, virtual training, and telecommuting, enhancing the quality of services provided.

The CA IT Client Management/Desktop & Server Management Green Book includes guidance and recommended best practices for the following topics:

- Challenges in discovering and managing computers and the need for automation

- Securing your desktop and server management infrastructure against rogue agents or managers

- Collecting asset data from non-agented systems

- Best practices for deploying and configuring the components of the solution

- Ensuring the high availability of the CA Unicenter Desktop & Server Management (CA DSM) managers

- Architecting for a successful implementation

Who Should Read This Book?

This book provides the consultant, architect, or manager with some best practices for implementing and using the CA IT Client Management products to support your environment for desktops, servers, laptops, and so forth.

Implementation success is based on a combination of people, process, and technology, and this wide-ranging book provides process, technical, and architectural advice. Information on key topics includes Discovery and Management, Securing the Solution, and Architecture.

CA Unicenter Software Delivery, CA Unicenter Asset Management, and CA Unicenter Remote Control are the primary products addressed in this book.

Readers of the more technical areas in this book will benefit strongly from some prior familiarity with the primary products. Therefore, readers are encouraged to make use of the standard product documentation and to attend the relevant CA Education courses.

Although this book is not designed primarily for executives, some parts of it, especially some of the higher-level best practices, may be of interest to senior management.

Chapter 2: Architecture Planning and Implementation Best Practices

Architecture Sizing and Component Placement Considerations

Although the CA Unicenter Desktop & Server Management (CA DSM) solution architecture is designed to be highly scalable and flexible, designing the right architecture for your business needs requires more than just counting the number of desktops and servers to be managed. You need to understand:

■ **How your current environment is structured**. For example, what firewalls are in place? What network speeds are available between primary and smaller, remote offices? Will the MDB be shared with an existing CA Unicenter solution? Are there existing machines available to serve as scalability servers and domain managers or will new hardware need to be purchased?

■ **What software will be managed and how**. For example, how many software packages will be delivered and how often? How many assets will be managed and how frequently will asset queries be executed? Are there resource compliance mandates that must be considered? What service level agreements have been provided to end users and CA DSM administrators?

■ **What administrative resources are available and where**. For example, if a scalability server will be installed in one location to manage several remote offices will there be a resource available to maintain that server?

The *Unicenter Desktop and Server Management Solution Planning Guide* discusses these and other considerations. It provides sizing guidelines based on the results of extensive product performance and scalability testing. Topics include:

■ Local versus remote MDB

■ Ingres versus SQL MDB

■ Network latency

■ Clustering considerations

■ Multiple applications versus shared databases

■ Domain server capacity planning

■ Scalability server capacity planning

These guidelines and principles can help you develop an initial logical solution or architecture based on an understanding of business requirements, processes, policies, and network topology.

The *Unicenter Desktop and Server Management Solution Planning Guide* is available through the following CA Support Online link:

https://support.ca.com/phpdocs/0/common/impcd/r11/scalability/doc/DSMDoc/Unicenter%20DSM%20-%20Solution%20Planning.pdf

This document is part of the Implementation Best Practices pages which can be used in conjunction with this Green Book and the standard product documentation set to make the most of your product deployment.

To access the Implementation Best Practices pages directly, go to the following link:

CA Support Online: https://support.ca.com/phpdocs/0/common/impcd/r11/StartHere.htm

Chapter 3: Scenarios

Introduction

There are many situations that are common to many of our customers. In this chapter we will present some examples of common situations, the questions that are often associated with them, and recommended steps to address the related challenges.

Slow Bandwidth

A common scenario is a company that has branch offices or remote locations. Often times these locations are connected by slow links. These slow links can cause some special challenges when attempting to deliver software.

Some of the challenges are as follows:

■   The bandwidth of the link is low – under 512 kb

■   The bandwidth is needed for ongoing business operations throughout the day

■   The bandwidth is needed for backup operations

■   Additional bandwidth is needed for software delivery

The challenge when there is limited bandwidth is being able to manage the available bandwidth while ensuring that all needed operations can occur. In this scenario, the first step you must take is to limit the amount of data that gets sent over the link. One way to do this is to ensure that duplicate data is not sent over the link. The recommendation here is to install a scalability server at each remote location. Then the contents of software packages only need to be sent once.

In most cases, the scalability server should be a dedicated machine because it can get very busy. If the remote location only has a few machines (generally up to ten) the scalability server could be an existing machine. In most cases, the scalability server software should be installed on a server operating system, but it can be installed on a Windows XP operating system. Also make sure the machine has enough disk space to house all the software packages.

The scalability server also acts as a collection point for asset management, so some disk space will be needed for that as well. It is important to regulate the amount of data asset management generates. One way to do this is to regulate how often the asset management agent runs and what it collects. In most cases, a weekly collection is suggested. Scheduling the collection to occur on a day when there is not much other network traffic would be a good solution. See the Configuration Policy chapter later in this Green Book for more details regarding agent configuration.

Careful scheduling of bandwidth usage is also important. Things to consider are what else is using the network and when. For instance, it is best to stage software packages to the scalability servers when the network is not busy. The Data Transport Service (DTS) is used to stage packages to the scalability servers. DTS allows for associating a scalability server or a group of scalability servers with a calendar. The calendar allows for setting a time window of when it is OK to send data over the network. DTS will suspend and resume the data transfer based on the calendar. The calendar is associated with the machine in the 2D Map, and the calendar itself is configured in the event management component of common services.

In addition to configuring when DTS will transfer data, CA Unicenter Asset Management can also be configured to schedule when asset management data should be transferred over the link. The asset management engine job that collects data from the remote scalability server can be configured to occur only on certain days and/or at certain times. Additionally, the amount of data that is sent each time is also configurable (in the form of a maximum number of collected files)—but one must be careful with this setting. If the amount of data collected per session is too small and the time between collections is too far apart, the amount of data that needs to be transferred will continue to grow. This will cause the disk space requirement to grow and will compromise the reliability of the data being collected. To configure this schedule, look at the engine task section under the control panel. To configure the maximum number of collected files, right-click on the specific engine under Control Panel>Engines>All Engines, select 'Properties,' then select the 'Advanced' tab.

Limiting the amount of data and scheduling the time data is transferred is often not enough. Another factor to consider is how much bandwidth is used at any one time. DTS allows this to be configured as well. There are two values that help in this area: parcel size and throttle factor. These are configured in the DTS section of the 2D Map. See the Data Transport Service chapter later in this Green Book for more details.

CA Unicenter Remote Control can also use a large amount of bandwidth. Limiting the number of colors sent and not sending the background bitmap helps in this area.

To recap, it is important to control the amount of bandwidth needed, when the bandwidth is used, and how much bandwidth is used at any given time. Optimizing these parameters is key to ensuring that the CA IT Client Management product suite is transparent to users connected by a slow link.

Maintenance Window for Servers

Server uptime is very critical in most environments. But keeping servers current with patches is just as important. Many patches require certain services to be stopped or reset, and often times a reboot is also required. Given this, it is very important to be able to test that a patch will not cause any problems with the applications or with any other related components. Change management is a critical part of server maintenance. In addition to validating which patch should be applied, when it should be applied is just as important.

Many companies have a scheduled maintenance window, during which it is acceptable for services or servers to be brought down for maintenance. This time frame must be carefully

controlled so that it does not affect the business. At the same time, administrators do not want to have to remember what needs to be applied and when.

It is much easier to submit software patch jobs as they are approved but not have them execute until a certain time. It is best practice to dedicate a maintenance window for this process. The CA IT Client Management product suite (CA ITCM) allows for this function. A server or group of servers can be associated with a calendar. The calendar is then configured to match the maintenance window. In this way a job can be created, but it will not be launched until the calendar matches the approved time.

Configuring the calendar is a two-step process. The first step is to enter the event management section of Common Services and choose Calendaring. Here is where the calendar is created and configured. After the calendar is created, use the DSM Explorer to associate the machine(s) with the calendar.

In addition to using the calendar to schedule the upgrades, the patch approval process can also be automated. If CA Unicenter Patch Management is used, the process can easily be configured to automate the change management process. Out-of-the-box integration is available between CA Unicenter Patch Management and CA Unicenter Service Desk. See the Integrating with CA Unicenter Patch Management chapter in the *CA Unicenter Service Desk Integrations Green Book* for details. An approval process can be configured to manage the patch acceptance and approval process. Once the patch is approved, it is automatically submitted to CA Unicenter Software Delivery. CA Unicenter Software Delivery will deliver the package in the next maintenance window.

Scheduling is critical for wise server management. CA IT Client Management allows one to manage the upgrade process very efficiently.

Periodic Reporting of Changes to the Environment

A very important component of CA DSM is scheduled reporting. People in different roles require different types of reports. Accounting and auditing, for example, may need a report on software installed by manufacturer and by department. Other reports needed may be hardware by manufacturer or by other criteria. These types of reports are needed to ensure software license compliance or to be able to negotiate more favorable rates with the hardware vendor.

Another use of reports is to meet the needs of the desktop technicians. Some examples of these reports are machines with low disk space or machines that do not have enough memory to support the new application that is going to be rolled out.

Reporting can either be a report by item and by number, or a report by exception. Since an enterprise environment is very dynamic, often management only wants to know about devices that have changed or that have fallen out of spec. For example, management may not want a report showing all machines and how much memory they have. Instead, they need an alert only when a machine does not have a certain amount of memory available or when the hardware configuration has changed. Alerts can be raise to notify of these conditions using query-based and event-based policy. Those details reside in the policy section of the DSM Explorer.

Other times, actual reports are needed. This is handled with the DSM Reporter. In this case, it is suggested to schedule the reports to run on a periodic basis. These reports can sometimes take time to run, so it is best to schedule them to run on a weekly or monthly basis. Another recommendation is to publish the reports to a file or even to a web site. In order to optimize the reporting, it is best to use the same queries to collect the data whenever possible; using the same queries to generate multiple reports is the most efficient method.

By running the reports on a scheduled basis, the report will be current and available for anyone to access. Not doing this will cause delays in providing information to the decision makers. By running reports on a scheduled basis, the data is available in raw form. Then, by applying filters, the data can easily be filtered to supply just the information that is being requested without having to run the report again. Reporting is configured in the reporting module.

Another example of the need for current status is in the case of an upgrade. For example, upgrading Windows XP from SP1 to SP2 requires periodic reporting. One would need to generate a list of machines that are the initial targets. One option is to build a software policy that says install the SP2 package to any Windows XP machine that is not yet SP2. The advantage of this method is that any machine that has a software delivery agent and is running Windows XP will have the package installed. This method will do the upgrade automatically, but it is difficult to control when it will occur and how many machines will get the upgrade at one time.

Another option is to provide greater control over the number of machines to which the upgrade will be applied at one time. In this case, additional criteria will be employed to narrow the initial target list. It is best to build dynamic groups based on all Windows XP machines and some other criteria such as host name A* or B*. Another method may be by IP range.

It is important to be careful how many machines the upgrade is applied to. A major change such as this may generate support issues because of a new look and feel the upgrade may cause. Therefore, it is important to monitor how many support issues are being opened versus the number of machines already updated. The rollout should progress at a rate that can be handled by the service desk. After most of the machines are upgraded, then the software policy can be used to upgrade any new machines that may join the network.

Chapter 4: Remote Management Database

Architectural

One of the first decisions to make in architecting a CA IT Client Management solution is whether the Management Database (MDB) should be installed locally (on the same host as the application) or remotely (on another host accessible through the network). Technical arguments can be made for both designs but, based on performance testing and reduced cost, implementing a single CA application and MDB on the same host is considered the best practice—unless compelling business or technical justification exists for remote placement.

If the enterprise manager is to be integrated with other components of the CA solution then it is recommended that the MDB be implemented on its own host. It would also be our recommendation that the server hosting the Enterprise MDB be highly available (clustered).

To determine if deviation from the best practice is justified requires an understanding of all requirements (technical, operational, and business). With these in mind, you can then apply the following set of simple guidelines and principles to drive the decision-making process.

Network Latency

Network latency introduced when the Database Management Solution (DBMS) server is remote from the application server would negatively impact overall performance. Implementing an application server and remote database server also adds cost to the solution in the form of additional hardware and operating system licenses. In the course of reviewing the other guidelines and principles, factors that outweigh the negative impact of network latency may surface. However, if there is no compelling justification (political or technical) for implementing the MDB on a server remote from the application server, the recommendation is to install the MDB locally on the application server.

Dedicated Database Servers

Policy may dictate that all databases be hosted on dedicated servers for administrative and security reasons and that no applications may be installed on these servers. While installing the MDB remote from CA applications may not be optimum because of the network latency introduced, nevertheless, the reduction in administrative cost and security exposure may more than offset the performance impact. Compliance with such policy would be considered a compelling justification for implementing the CA application(s) with a remote MDB.

Multiple Applications/Shared Database

Additional considerations are required when multiple CA applications will be sharing the same MDB— either now or in future—particularly at the enterprise (top) tier of the architecture. For example, CA Unicenter Asset Portfolio Management (CA Unicenter APM) and CA Unicenter Desktop & Server Management (CA DSM) integration currently requires that the applications share the same MDB.

Installing multiple applications on separate hosts reduces potential administrative conflicts (such as access control and maintenance schedules). For example, if the MDB requires a reboot, any other application installed on that server will be brought down. On the other hand, separating the applications from the MDB (and each other) provides the ability to perform administrative tasks independently.

In general, when multiple applications will share the MDB in a medium to large implementation, the recommendation is to host the MDB on a separate server, remote from the application server(s).

Conclusion

When designing an architecture for a single CA application, it is clear that co-locating the application and MDB on the same server (or cluster) would provide the best performance/least cost best practice. Alternatively, installing the MDB remotely should only be considered after a careful review of all factors to determine if there are other compelling justifications that outweigh the performance and cost benefits. Bear in mind that, while the term 'remote' implies separate hosts, it should be understood that the hosts should be electronically close with reliable network connectivity.

In the end, the decision must be driven by the overall impact to the business.

Installation

There are three methods for installing a remote MDB from the CA IT Client Management (CA ITCM) media. If the MDB does not already exist, it should be installed prior to installing any CA ITCM application servers.

1. **CA ITCM without CCS (Common Services)** – When installing only CA ITCM, no special action is required. Step one is to use the Install MDB option from the CA ITCM media and complete the MDB installation on the database server.

2. **CA ITCM with CCS** – When installing both CA ITCM and CCS, you must include the WorldView (WV) manager component on the MDB database server. Even though you will be running the CCS application components from the application server, the WV component needs the WV Manager installed on the MDB server for communication purposes. There are two ways to accomplish this:

   ■ If available in your environment, you may use the CA Unicenter Network and Systems Management (CA Unicenter NSM) media to select and install only the WV Manager. This is the preferred method. By using this method, only that one

component will be installed. You would accomplish this by first installing the MDB using the CA ITCM media. After that installation has completed, insert the CA Unicenter NSM media and install only the WV component on the MDB server. Next, using the CA ITCM media select the Install Unicenter DSM choice from the installation menu and make certain that the list of components to install includes CCS. During the installation, point to the remote MDB. This will install full CCS including Enterprise Management, WV, and Continuous Discovery on the application server.

■ If you do not have access to CA Unicenter NSM media, you can use the CA ITCM media. You will have to install full CCS on the MDB database server. Very likely, however, you will not want to run most of those components on that server. Therefore, immediately after the installation of full CCS and before the services are started on the MDB server, you will want to disable certain services including:

> CA – Continuous Discovery Agent

> CA - Continuous Discovery Manager

> CA DIA 1.2 DNA

> CA DIA 1.2 Knowledge Base

> CA-Unicenter

> CA-Unicenter (NR Server)

> CA-Unicenter (Remote)

> CA-Unicenter (Transport)

3. **MDB Currently Exists** – If the MDB has been installed on the remote server by another CA application, you will need to run the Install MDB on that server from the CA ITCM media. The installation will recognize that an MDB currently exists and will install only the necessary updates. You must remember to shut down any other applications that are accessing the MDB prior to running the MDB update. The above two scenarios involving CCS remain the same after the MDB has been updated.

Trusted Relationship

From the *Unicenter Desktop & Server Management Implementation Guide*: Note that if the MDB is installed on a machine that is remote from the domain manager, a trusted connection must be established between the two machines. One way to accomplish this is to have them reside in the same Windows domain.

A workaround to the above scenario has been identified. You must create identical Windows user accounts on each system. The account must be a local administrator where the CA DSM Manager is to be installed, and *also* a SQL login with permissions to the MDB on the MDB server.

Chapter 5: Implementing for High Availability

Deciding to Cluster

The risk of an application failing is usually the driving factor in deciding whether or not to use a clustered-server approach to service resiliency. For instance, there may be a patch that is needed in the environment, but without the CA IT Client Management (CA ITCM) asset management capability, it would be difficult to assess the degree of risk to an enterprise or business. Without the CA ITCM software delivery functionality, remediation of the application or OS at risk becomes tedious, burdensome, dangerously slow, or virtually impossible. The amount of time needed to fully replace a server and working application depends on several factors which have a direct consequence in determining if you need a cluster. Unless there are ready reserves of equipment, getting replacement hardware can take quite a while—usually longer than a critical-need application would require. On top of that, the reloading of the operating system, installing the backup software, locating the needed backups, and retrieving information from those backups could create quite a lengthy delay in service.

Determining the need for a cluster boils down to a classic assessment of cost versus risk. Historically, the more readily available an application is, the more expensive the supporting recovery solution is. Effectively, you pay for speed of recovery. CA Unicenter Desktop & Server Management (CA DSM) 11.2 has better support for business continuity with the support of clustered Microsoft servers.

What Is Needed and Supported

CA DSM 11.2 supports two clustered 32-bit servers running Microsoft Windows and Microsoft SQL Server. Implementations running on the Ingres database are not supported in a clustered environment. At this time, 64-bit operating systems are not supported. Windows 2003 Enterprise or Windows 2003 Data Center are the operating systems that support clustering. Although there is a menu selection in Windows 2003 for Cluster Administration, you cannot create a cluster on that OS.

Generally, two servers running Windows 2003 and sharing disks are requirements for a Microsoft cluster. There is usually a *quorum* drive used in managing the control of the cluster, and a *shared* drive for data that needs to be visible to each node in the cluster. In addition to the OS, both servers must be members of an Active Directory Domain, since the account is used to create and manage services across both machines. Although a cluster can contain a domain controller, CA DSM cannot run in this environment. Local accounts are created during installation of CA DSM and domain controllers do not have a local security database, or SAM, so the installation will fail.

A total of six static IP addresses are required in total, five for the creation of the cluster and one for the Microsoft SQL cluster when it is installed. The two *public* Network Interface

Cards (NICs) (one on each server) and the two *heartbeat* NICs (one on each server) are checked as part of the cluster creation and will cause a failure if they are not set for static addresses. The last IP address is needed as the address for the cluster itself.

Cluster Considerations

There are a few things to consider in the deployment of CA DSM in a clustered environment. We will see later, as we step through the installation, that the two servers share one HostUUID, and only one server is actively running CAF (the Common Application Framework) at any one point. As a result, CA DSM cannot manage the two machines that are supporting the cluster. Any time the application switches from one node to the other, CA DSM will detect a change in the MAC address, although the name would still be the name of the cluster. Subtle changes in the configuration of the machines and perhaps the applications installed would begin to cause odd results and certainly call into question the accuracy of the data from the two servers.

The Boot Server is not supported in a clustered configuration as part of the manager installation. Usually, a Boot Server is deployed as part of a scalability server, and it is preferred that devices in the environment report to a scalability server instead of directly to a domain manager.

The Discovery Manager is also not installed in a clustered configuration, so continuous discovery and the automated agent deployment will not be running in this environment.

How to Begin

Build the Cluster

The first step is, quite simply, building the cluster. There is ample documentation from Microsoft on how to do this, so we will not go into great detail here. You can reference Microsoft's site for the following:

"Quick Start Guide for Server Clusters",

http://support.microsoft.com/?id=258750 (heartbeat configuration), and

http://support.microsoft.com/kb/817064/ (enabling network DTC access).

To set up the cluster, log in using a Domain Account to do the installation. This is a requirement because the installation will require an ID with rights on the other node. Even though a Local Account has the same password and privileges on both nodes, the installation will fail, citing credentials.

According to the Microsoft document for clustering, there are some changes to make to the heartbeat connection. Make sure that the heartbeat connection is static, and click the Advanced button in the Internet Protocol Properties dialog shown next:

On the DNS tab of the Advanced TCP/IP Settings dialog, make sure that the check boxes are cleared for "Append parent suffixes of the primary DNS suffix" and "Register this connection's addresses in DNS", as shown here:

On the WINS tab, select the option "Disable NetBIOS over TCP/IP":

Additionally, make sure the bind order for the network cards is correct, making the public connection the first in the list. This is illustrated in the next two screens:

To enable network DTC Access, launch "Component Services" from the Control Panel. Drill down to My Computer and right-click. Check the Network DTC Access box. Other options to select are Allow Inbound, Allow Outbound, and No Authentication Required.

In the Cluster Administrator, there are some selections to be made for networking. For the heartbeat connection, enable the connection for internal cluster communications. For the public connection, Microsoft suggests that "All communications (mixed network)" be selected, allowing for a duplicate path for internal cluster communications.

Once the cluster is built, one thing to check before moving forward is that the resources can "fail over" between nodes. In checking the shared disk resource (in the following examples, the S: drive), open Windows Explorer on the active node and create a file in the root directory.

Open the Cluster Administration application, right-click the resource, and select Move Group from the drop-down list.

Once the group has been transferred to the passive node, log in to the passive node. Open Windows Explorer on that server and check the existence of the test file. Create another test file from the passive node as well. Using the Cluster Administrator, move the group back to the primary node and make sure both files exist on that side as well.

Once the resources fail over *and* back, installing SQL Server is the next step.

Install SQL Server

The installation of SQL Server 2005 on a cluster is remarkably simple. The nice surprise is that installing from the active node will install on both nodes, including necessary services and cluster configuration. Again, there is plenty of Microsoft documentation available for this procedure so we will not repeat it here. Like the cluster before it, the SQL installation across the nodes requires a domain account. Not all the components of the SQL installation are cluster aware. In the following screenshot, SQL Server and Analysis Services are cluster aware, but the Reporting Services are NOT cluster aware.

Make sure that you specify the shared drive as the location for SQL files.

In the creation of the SQL Cluster, make sure that MSDTC is in a different resource group from the SQL Server itself. MSDTC is usually in the Cluster Group. Later in the CA DSM installation, CA Common Services (CCS) will automatically find the name of the SQL server as long as the resource group doesn't contain MSDTC. Otherwise, the CCS installation will identify MSDTC as the name of the SQL database.

Test the Failover

As we did with the base cluster installation, test the failover of the SQL Server. Query the database from both nodes, move the SQL resources to the other node, and rerun the query. Once the database successfully fails over AND back, then proceed with the installation of CA DSM.

Active Node Setup

The installation at this point looks the same as a normal installation with an important difference: the Management Database Server is the name of the SQL cluster, not the OS cluster.

When we click on Recovery, the Cluster Name is the name given the OS cluster, not the SQL cluster, and not the node on which we are installing. We will click Enable Recovery Support and designate this as the active node.

It is important that configuration data be stored on a shared drive. In this example, the database, DSM executables, and library are all on the J: drive. A specific architecture for a busy enterprise may indicate the need to put the database on its own shared drive, separate from the executables and library directories on a third shared drive. From the documentation: "For a cluster installation the installation path for the configuration data must point to the shared disk of the cluster. The Shared Components should be placed on a local disk." From the "Choose Destination Location" screen, you have to click on the "Advanced" button to make this distinction. You will see the effect of this choice during the CCS portion of the installation.

In a standalone, non-clustered, configured CA DSM installation, the CCS portion is silent. In this situation, the installation will then invoke a manual installation of CCS. CCS will install the High Availability Service as part of the installation. This is needed for cluster support.

The screen that comes up shows the default installation options for CCS. High Availability Services (HAS) will automatically configure the cluster resources during the CCS installation. Not all the CCS components are supported by HAS, but the default selections are correct and should remain as they are. The one mentioned earlier in this chapter that is not supported by HAS is the Continuous Discovery Manager.

The SQL server, when not in a cluster resource group that includes MSDTC, will be automatically identified by the CCS installation. In the following graphic, the Database Server field was filled in by the installer. Check both the Enterprise Management Database and the Worldview Repository logon information to make sure that the instance name is blank, meaning it is the default instance.

The CAF service is set to manual during the installation, which will require an administrator's intervention for startup. This is due to CA DSM being cluster tolerant as opposed to being cluster aware. There are no resource DLLs that can be used by the cluster to interrogate the status of the application.

Another change in the clustered configuration is in the registry settings for the HostUUID, with the addition of three keys: ClusterHostName, ClusterIPAddress, and ClusterModeEnabled. The comparative screens are shown in the graphics below with a standalone configuration shown first and a clustered configuration second:

After installation, you will be prompted for a reboot. Make sure the resources won't fail over to the secondary node and reboot.

Passive Node Setup

Setting up the passive node is pretty straightforward, but first, some housekeeping. First, stop CAF and, via the cluster administrators, the CCS services. The HAS (High Availability Service), installed during the CCS install, will control CAM and tell it which node is in control.

Second, move the groups that contain SQL and the shared disk to the passive node. Before you do this, set the Unicenter services to "Do not restart" or take them offline. If the services are set to restart, they will try to start upon failing over to the passive node. Since CCS has not been installed on the passive node, the 'move group' command will fail and the whole group will move back to the primary node. Don't forget to restore them to the original settings after the installation.

When proceeding with the installation, we make the same selection of Enable Recovery Support, but this time we choose Passive. The configuration data location should be entered into the field at the bottom. The installer will find DSMRecovery.ini and process the same settings that were selected during the active node installation.

The default for "Package Import" is for multi-language packages. If you selected "English Only" during the installation on the primary node, you will have to make that selection on the passive node.

Some Manual Activities

Shares

If you chose to implement shares as the way to access the library, the creation of the share resources for the cluster configuration will be a manual process. From the cluster administrator, create the resource of a 'file share' type, determine which cluster nodes can own it, and what the dependencies may be. In this instance, the SDMSILIB share/resource is dependent on the Shared Disk, the S: drive.

Permissions may be set to NULL, to an AD group, or to the user credentials defined to the agent. The agent credentials can be set in Configuration Policy and encrypted using SD_SETCNF.exe.

Web Console

On both nodes, there needs to be a change to the WACConfig.properties file. The two lines to change are the AMS_URL and the WEBSERVICE_URL variables; they should point to the DSM virtual server, as follows:

AMS_URL=http:// *DSM virtual servername* /AMS/login.do
WEBSERVICE_URL=http://*DSM virtual servername*/UDSM_R11_WebService/mod_gsoap.dll

Before making the change, stop the Tomcat service with CAF STOP TOMCAT. After making the changes, restart IIS.

Failing Over

If there is a system failure (meaning the primary server is down), the database, drives, and database will fail over to the secondary server by themselves. They are usually defined to behave that way in a cluster. The only thing left to do is:

■ Run ActivateManagerNode.bat.

If you are trying to fail the application to another server manually:

■ Stop CAF on the primary node.

■ Using the cluster administrator, move the drives and SQL Server to the Passive Node.

■ Run ActivateManagerNode.bat on the passive node, which will start CAF.

The ActivateManagerNode batch file basically makes a configuration change using CCNFCMDA, which makes registry changes by writing them to a file and executing the .reg file, and, if there are no errors, will start CAF.

If you start CAF before running ActivateManagerNode.bat, an error will pop up saying that you are trying to start CAF on the non-primary node in a cluster.

ActivateManagerNode.bat:

```
@echo off
set LASTERR=0

REM set activenode to an empty string
echo ccnfcmda -cmd SetParameterValue -ps /itrm/common/failover -pn activenode -v ""
ccnfcmda -cmd SetParameterValue -ps /itrm/common/failover -pn activenode -v ""
if ERRORLEVEL 1 GOTO error01

REM set failover state to 1
echo REGEDIT4 > "%TEMP%\DSMFailover.reg"
echo [HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Unicenter ITRM\failover] >>
"%TEMP%\DSMFailover.reg"
echo "state"=dword:00000001 >> "%TEMP%\DSMFailover.reg"
regedit /S "%TEMP%\DSMFailover.reg"
if ERRORLEVEL 1 GOTO error02

REM launch caf start
caf start
if ERRORLEVEL 1 GOTO error03

GOTO end
```

```
:error01
echo ccnfcmda.exe failing - batch file terminated
set LASTERR=1
GOTO end

:error02
echo regedit failing - batch file terminated
set LASTERR=1
GOTO end

:error03
echo caf start failing - batch file terminated
set LASTERR=1
GOTO end


:end
del "%TEMP%\DSMFailover.reg"
exit /B %LASTERR%
```

Chapter 6: Firewalls and Network Considerations

Introduction

In large distributed environments you will often find many network restrictions and limitations. These restrictions are typically seen in networks where internal network resources need access to the outside world, or where agents on the outside need access to resources in the internal network.

In this chapter, the outside network is referred to as the Demilitarized Zone or DMZ, and the internal network is referred to as the intranet.

Typically these restrictions are caused by two things that will be described in the following sections:

■ Port usage/availability

■ Network Address Translation (NAT) usage

Port Usage

One important thing to understand before discussing the architecture is how the different CA IT Client Management (CA ITCM) components communicate. In general, CA ITCM uses just two ports:

■ **CAM**: 4104 (UDP) or 4105 (TCP). CA Message Queuing (CAM) is used for the session messaging. You will have to choose one port or the other.

The UDP port 4104 is the default port. UDP (User Datagram Protocol) is a lighter weight protocol than TCP (Transmission Control Protocol). UDP is a minimal message-oriented connectionless protocol. With UDP, messages (packets) cross the network in independent units so it does not handle Network Address Translation (NAT) environments well. Although the total amount of UDP traffic found on a typical network may be only a small percentage, a number of key applications including the simple network management protocol (SNMP), the Domain Name System (DNS), and the Dynamic Host Configuration Protocol (DHCP), use UDP. It is also used by streaming media and Voice over IP (VoIP).

In contrast, TCP is a connection-oriented protocol. A connection can be made from client to server, and from then on any data can be sent along that connection. In NAT environments, the TCP port (4105) is more suitable.

■ **Port Multiplexer**: 4728 (TCP). This is the port where the entire data stream is transported.

These two port numbers can be changed, but in this book it is assumed that you will use the default ports.

The following table contains the most relevant ports used by CA Unicenter Desktop & Server Management (CA DSM) with the corresponding type of communication, and the description of their usage:

Communication	Port Number	Description
Database (Microsoft SQL)	1433	This is the default port for Microsoft SQL. You can choose to have another port, but then you need to set up both the database server and the manager.
Database (Ingres)	Ingres/Net 19016 & 19017, JDBC: 19023	
All CA ITCM internal communication	CAM/Multiplex (See above)	All internal communication.
WEB Console	80	When connecting from a workstation to a manager use the Web Admin Console (WAC) port 80. Port 80 is also used in connection with Web Services.
CA Unicenter Software Delivery Agent-Related (the standard CA ITCM ports also apply)		
Shares (NetBIOS)	nbname 137 nbdatagram 138 nbsession 139	Shares are used in CA Unicenter Software Delivery for two things: NOS deliveries and MSI packages. You can choose not to use shares. This is typically the case if the agent is in the DMZ, as you do not want to open Shares to the outside world. If you choose not to use shares you might lose some MSI features like Repair. Also, by using non-NOS connections (NOS-Less or DTS download) you get the benefit of network-tolerant download, but you also temporarily need three times the space of the package on the local drive (zipped version, unzipped version, and installed).
Echo/Ping	7	Used in deployment.
TFTP	69	Used by deployment and Operating System Installation Management (OSIM).
DHCP	67, 68	Used by OSIM.

Note: This is not a complete list and is only meant as an introduction to the ports used. For a full list of ports and their description and usage, see the *Unicenter Desktop & Server Management Implementation Guide*.

In general, infrastructure deployment and OSIM should not be used through slow lines. So the last few table entries shown above are normally not an issue when architecting enterprise type infrastructures. However, we have seen service providers that use firewalls

to separate client machines needed to deploy operating system and CA ITCM agents, which will introduce a challenge to that environment.

Now that we have established the port usage, it is important to understand who initiated the communication. The communication scenarios are illustrated in the following diagram:

* The Agent can be changed to pull this on a scheduled basis if, for example, you are running in a NAT environment. This is discussed in the next section.

Network Address Translation

NAT allows you to share network connections on other networks by translating the address from an internal address to an external address.

Please note that the following sections include a conceptual description of NAT. However there is much more to this great and widely used technology.

You can split NAT into three general categories:

■ **Static**: In a static environment, you define a translation table on both sides of the network. In that way two machines on opposite sides of the router can see (ping) each other by name and the addresses will be resolved. In an environment like this, most of the CA ITCM system will work seamlessly because the normal method of addressing a node in CA ITCM is by name.

■ **Dynamic**: In a dynamic NAT environment the outside network address is assigned when you try to access the network. This means that the connection only stays open and the address is only assigned during the session. The technology that ensures that you can communicate in this environment is the TCP socket—it stays open and acts as a tunnel during the connection. But once disconnected, no machines on the other network can contact you. In a CA ITCM environment this means that the initiation can only be done from one side of the firewall. This normally means that the agent can connect, but the DSM Manager/DSM Explorer can not reach the agent.

■ **Overload/PAT (Port Address Translation)**: This is very similar to the dynamic NAT. The only difference is that typically only one address is shared by multiple connections. This is what you experience with your home broadband connection, where one address is assigned to the house but multiple computers can connect to the Internet and use it simultaneously.

Typically, NAT routers are used for security and IP address re-use. In the former case, end systems are hidden behind a NAT router using either static or dynamic NAT; an end system's internal, configured IP address is never exposed to systems beyond the router. In the latter case, Port Address Translation (PAT) or NAT Overloading is used in response to the decreasing number of available registered IP addresses.

Note: PAT presents the biggest challenges for CA DSM, and is what you should be most aware of when architecting the CA DSM infrastructure. PAT essentially stops any direct connections from beyond the router from reaching systems connected to the local LAN.

CA DSM uses the following two proprietary communications technologies:

■ CA Message Queuing (CAM)

■ TCP Streams through Port Multiplexer (used by Data Transport Service (DTS), the NOS-less file transfer component, and the CA Unicenter Remote Control video stream)

Since CAM uses the message's source IP address to identify an end system, this may cause a problem in an overloaded NAT network where the source IP address will always be that of the router. When CAM receives a second connection from the same IP address it discards the first, potentially causing problems because response messages would not be able to get back to the system that made the request.

Fortunately, CAM will work when communicating out from a PAT configured network (agent initiated traffic), but only if something else from the same PAT configured network does not come along in the middle of a message exchange. The effect of this may be minimal in a normally quiescent network and even in a CA DSM network with moderate activity where application code may experience connection failures or timeouts—though the code should retry and recover. A very active CAM network, however, may be significantly affected.

As a rule of thumb: *All* manager components need bi-directional communications—meaning that you can ping from both components to the other using a short name or fully qualified DNS name. The agent only needs 'agent to scalability server' connectivity. This means that a Dynamic NAT should only be between the agent and the scalability server.

Configuring All Components for a NAT Environment

If you are running in a NAT'd environment, there are some configurations you will need to do on all the CA ITCM components in the network. The settings involved are mainly the CAM configurations.

If you need to configure CAM for routing or for any other action, you need to do it in the CAM.CFG file. It is important that you set it in the right place. You will find the CAM installation folder in the Environment variable **CAI_MSQ**. The default installation location is *C:\program files\CA\SharedComponents\CAM*. Once you have located the CAM installation, you need to find the cam.cfg file. The configuration file does not exist by default. To create it, use this command:

        bin\camsave config

This will save the current configuration to a file called save.cfg. Rename the save.cfg to CAM.CFG. After building the file, you can add the configurations from the following example in order to match the environment you have.

The example CAM config file:

```
# CAM config saved: Fri Dec 07 14:21:05 2007

*CONFIG
fixed_paths = no
close_time = 60
client_hold = 60
connect_retry = 60
udp_port = 4104
tcp_port = 4105
spx_port = 4905
cas_port = 3104
dg_log_files = 8
dg_log_size = 64
dg_log = *
au_log_files = 8
au_log_size = 256
au_log = *
tr_log_files = 8
tr_log_size = 1024
tr_log = *
trace = none

*PATHS

*ROUTING

*AUDIT
off
```

How to Configure TCP in CAM

By default, CAM is using UDP as the initial communication protocol. Since UDP is not a protocol that normally is routed, you will need to change the CAM layer on all components that are communicating through a NAT—regardless of what type of NAT—to be using TCP as the default communication protocol.

You can force TCP in two ways:

1. Disable the other protocols including UDP by setting the port number to 0 in the *CONFIG section. For example:

```
udp_port = 0
tcp_port = 4105
spx_port = 0
cas_port = 0
```

2. If you only need TCP for one or two nodes you can define a *PATHS section in the file, and specify the server and the protocol. For example:

```
*PATHS
SS1.ca.com protocol=tcp
192.168.1.123 protocol=tcp
```

How to Configure Routing in CAM

If you have a NAT environment that requires routing (see below) then you will need to define forward and routing information in the *ROUTING section. For example:

```
*ROUTING
forward localhost 150.100.1.100
agent1 192.168.1.129
agent2 192.168.1.130
```

Configuring a CA ITCM Agent for NAT Environments

If you choose to have the agent and the scalability server separated with a NAT'd network you will lose some functions:

■ **Job Check**: This feature allows you to send a manual job check from the DSM Explorer, or the automatic job check from that scalability server. It will not work, as the agent can not be contacted. To overcome this you will need to set up CA Unicenter Software Delivery to job check on a scheduled basis from the agent. See the Agent Best Practices chapter later in this book for more details.

■ **Common Configuration:** Normally the changes in the configuration are sent 'down' through the infrastructure. If the agent is in a NAT environment, the scalability server will not be able to reach the agent and the changes will not be applied. To change this behavior you can set the agent to poll the server for changes. This is done in the common configuration with the locally managed attribute:
ITRM/agent/cc/CsmPollInterval

To set this locally managed parameter to one hour, you can use the following command on the domain manager:

```
ccnfcmda -cmd SetParameterValue -ps"ITRM/agent/cc" -pnCsmPollInterval –v3600
```

Or you can run a CA DSM script like this:

```
ccnfSetParameterStr("ITRM/agent/cc/CsmPollInterval ","3600")
```

Note: You should not set this parameter if you are not using NAT as it does introduce unnecessary network traffic in a normal LAN environment. To disable the Poll parameter, set the interval to 0.

■ **Remote Control:** Based on the Connection initiated from the DSM Explorer to the Agent, so if you need remote control you need to have Static NAT, or use another solution like CA SupportBridge.

Architecting the Solution in a NAT Environment

There are many ways to architect a CA DSM solution in a distributed environment. The following sections describe some of the most common CA-recommended ways.

Agent and Scalability Server in the NAT'd Network

The most typical way of using NAT is to have a scalability server in the NAT'd network; this can be an internal network or the Internet (DMZ), together with all the agents. The benefit of this is that all connections from the agent are through the scalability server, with the exception of remote controlling an agent. This gives great flexibility in adding agents and connection methods.

Here is an example of that setup:

■ A CA DSM domain manager was connected to the example corporate network

■ A CA DSM scalability server was connected to the NAT router with a static NAT address appearing as 150.100.1.100 to connections through the WAN, and 192.168.2.10 locally

■ Two (2) DSM agents were connected to the DHCP-based, NAT'd (or PAT'd) LAN

This is illustrated in the following diagram:

```
                    ┌──────────────────┐
                    │   DSM Domain     │
  Example           │    Manager       │
  Corporate         │                  │
  Network           │ 150.100.100.100  │
                    └──────────────────┘
                             │
        ┌────────────────────┴───────────────────────┐
        │                    │
        │          ┌──────────────────────┐
        │          │    150.100.100.1     │
        │          │                      │
        │          │ ROUTER - CORPORATE   │
        │          │                      │
        │          │     150.100.1.6      │
        │          └──────────────────────┘
        │                    │
        │                    │                         WAN
        └ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─┼─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ─ ┘
                             │
                    ┌──────────────────────┐  DMZ /   ┌──────────────────┐
                    │    150.100.1.5       │  Static  │ Known externally │
                    │                      ├──────────│        as        │
                    │    ROUTER - NAT      │          │  150.100.1.100   │
                    │                      │          ├──────────────────┤
                    │     192.168.1.1      │          │  DSM Scalability │
                    └──────────────────────┘          │      Server      │
                             │                         │    (server1)     │
        ┌────────────────────┴─────────┐               │  192.168.2.10    │
   Local│                    │          │              └──────────────────┘
   Network                   │          │
        │          ┌──────────────┐  ┌──────────────┐
        │          │ 192.168.1.129│  │ 192.168.1.130│
        │          │   (agent1)   │  │   (agent2)   │
        │          │ CA DSM Agent │  │ CA DSM Agent │
        │          └──────────────┘  └──────────────┘
```

In this scenario the scalability server is essentially in a separate network. So both the WAN and NAT'd LAN can see it, but with different IP addresses. The DSM Agents are hidden from the WAN.

In this scenario the scalability server successfully registered with the domain manager, and agents successfully connected to the scalability server for registration and inventory upload. However, this information was not collected by the engine.

The challenge with this setup is that the CA DSM domain manager engine is unable to connect to the 'outside' scalability server and 'Unable to Open…' messages were written to the engine event log.

From the manager's perspective, the scalability server is known by its external address (150.100.1.100) and, therefore, CAM messages are directed to this address. The scalability server, however, actually knows itself by its NAT address (192.168.2.100). So attempts to

forward the messages sent from the manager to an address that is unreachable are not successful.

To solve this you will have to configure CAM on the scalability server to route messages through its NAT address through a simple cam.cfg rule. Add a ROUTING rule to the cam.cfg file that directs messages destined for 150.100.1.100 to localhost. For example:

```
*ROUTING

forward localhost 150.100.1.100
```

Agent as CAM Proxy

In this scenario all CAM communications are routed through one of the agents in CA DSM. This requires that agent to be 'up', and therefore requires the appropriate CAM ROUTING rules on the agents and scalability server. This scenario is illustrated in the following diagram:

Example Corporate Network

| DSM Domain Manager 150.100.100.100 | DSM Scalability Server (server1) 150.100.100.101 |

150.100.100.1

ROUTER - CORPORATE

150.100.1.6

WAN

150.100.1.5

ROUTER - NAT

192.168.1.1

DMZ / Static

Known externally as 150.100.1.100

DSM Agent CAM Proxy (agent1) 192.168.2.10

Local Network

192.168.1.130 (agent2) DSM Agent

Similar to the previous topology, the agent1 machine is placed in the segregated network. And, since it is known by a different IP address externally, a CAM ROUTING rule is required to ensure all traffic sent to its external address is processed locally.

To do this, add a ROUTING rule to the cam.cfg file that directs messages destined for 150.100.1.100 to localhost. For example:

```
*ROUTING
forward localhost 150.100.1.100
```

In addition, now that agent1 is acting as a CAM proxy, CAM traffic from the scalability server destined for agent2 must be forwarded to agent1. So a routing rule must be added to the scalability server. For example:

```
*ROUTING
forward 150.100.1.100 192.168.1.*
```

Disadvantages When Using NAT

Having the agents separated from the domain manager and from the DSM Explorer poses some other issues. One can argue that with the other settings you will not need them in a day-to-day operation. But here are some of the common functions, potential issues, and configuration options.

■ **Infrastructure Deployment**: Infrastructure Deployment fails to discover the end systems during the scan phase and, even if it could, file share and telnet access is not possible because the end systems are hidden from the manager.

■ **Ad-hoc Job Checks from DSM Explorer**: As the DSM Explorer does not have direct access to the agents, you will not be able to do an ad-hoc job check. You can set up a lookup the other way using the local host file and CMA (see below).

■ **Remote Control**: You will be able to take over (viewer – host) within the same network, but you can not take over across the NAT'd network. If you require remote control across NAT'd networks, you could use a product like CA SupportBridge.

Architecting CA DSM in a Firewall Environment

In this setup you will have the agent in the DMZ or in another segregated network, and the rest of the solution in the intranet. This is the typical setup for small offices and home workplaces. In this environment you will typically:

■   Open ports 4104 and 4728 for bi-directional traffic for all users

Agent and Scalability Server Outside the Firewall

You typically use this scenario when you want to shield the DMZ/external network even more, or when you have a larger number of agents in one remote site. The difference from the first scenario is that you can limit the open ports only for the two CA ITCM ports to a specific machine (the scalability server), leaving all the agents unable to get to the intranet.

The advantage of this solution is that you can limit the port opening to only the traffic between the scalability server and the domain manager.

The Whole Domain Behind a Firewall

This scenario is typically used by a service provider, when they place a management domain in the client's network and then connect this to a central enterprise for common management and reporting.

For this scenario you need to open up the firewall to the CA ITCM port and the database port(s), but ONLY for the specific manager nodes in question (scalability server, domain manager (engine) and, if remote, the database server) in both directions.

In this setup you will get full functionality within the domains. So things like MSI and file share will work, but you can still control software delivery and report on inventory from a central point.

Note: In this scenario it is important to set up the right security profiles to ensure that the users on the enterprise have access in the two domains but no access between the domains are needed. See the Security chapter later in this book for more details.

Chapter 7: Data Transport Service

Many of the features mentioned in this chapter require CA Unicenter Desktop & Server Management (CA DSM) 11.2 or higher, and require the installation and configuration of CA Common Services (CCS). Without CCS, only default settings are available.

What Is DTS?

Data Transport Service (DTS) is a file transfer mechanism that ensures safe, reliable data transfers across a wide variety of transport media. DTS optimizes the available bandwidth by giving better control as to when and how to use the network.

DTS provides features such as:

- Parcel size configuration

- Bandwidth throttling

- Checkpoint restart

- Automatic route and protocol determination, such as broadcast, multicast, and fan-out

- Encryption

- Compression

- Transfer scheduling

- Multiple route utilization

- Polite Agent mode (real-time adjustment of bandwidth throttling based on platform performance)

DTS is a service included as part of CA DSM, and is used by CA Unicenter Software Delivery as well as other CA products. DTS has been included with CA Unicenter Software Delivery since version 1.7.

DTS can ALSO be used by CA Unicenter Software Delivery at the agent level. Care should be taken, however. DTS is not designed to be installed on all agents. DTS is designed to be used on agents that have connectivity issues such as dial-up or poor quality connections. Out of the box DTS is installed on all enterprise, domain, and scalability servers. On UNIX and Linux the DTS agent is installed to all agents but is not configured to run. It must be enabled if DTS is going to be used. For Windows platforms, a DTS agent plug-in is provided as a software delivery package that can be installed on any agent machine.

DTS Components

Note: For additional details, see the *Unicenter Desktop & Server Management Implementation Guide*.

DTS consists of four components, described in more detail in the following paragraphs:

- NOS

- TOS

- SOS

- Agent

NOS

The Network Object Server (NOS) responds to requests from the Transfer Object Server (TOS). The NOS is responsible for keeping track of variables about the network objects themselves. This data is accessed by the NOS through the WorldView component. The variables used by the NOS are actually configured in the 2D Map. Examples of these variables are:

- Parcel Size

- Throttle Factor

- Protocol

- Transfer Schedule (Calendar)

The NOS values are semi-permanent. They are kept indefinitely and are reused each time a transfer is made to the specific device(s). These configurations may be modified as the network changes or different transfer criteria are needed.

TOS

The Transfer Object Server handles requests to initiate and manage data transfers. In the case of CA DSM, software delivery communicates with the TOS to create and manage file transfers. The TOS will communicate with the DTS Agent to initiate a transfer and monitor the transfer status.

The TOS will maintain data about the transfer. The TOS keeps its data in a table in the MDB. The types of information maintained by the TOS include:

- Sending Machine Name (initiator)

- Receiving Machine Name (responder)

- File Name

- Number of Retries that should be tried in the event of a network issue

- Interval between retries (in seconds)

In the case of CA DSM, the host and file name is designated by CA DSM. Settings such as the number of retries and the retry interval are configured by the software delivery section of configuration policy. In this area, items such as compression and encryption are also configured.

The TOS also stores status information about the transfer. The agent will periodically report the current status of the delivery back to the TOS. CA Unicenter Software Delivery will periodically check in with the TOS to update the status of the transfer.

The data about a specific transfer is kept in the TOS until the TOS is directed to delete that information. If CA Unicenter Software Delivery creates the transfer, then it maintains the TOS record for that transfer. CA Unicenter Software Delivery will ask the TOS to delete the data about a transfer either when the transfer is successful or when the software delivery job timeout has expired.

The TOS is also responsible for transfer groupings. In other words, CA Unicenter Software Delivery sends a list of targets to the TOS. The TOS in turn collects information about those targets (from the NOS), or applies default settings if the NOS has no data about a specific target. After collecting the data, the TOS separates the transfer into specific transfer groups based on specific like properties. A transfer group must have identical parameters such as parcel size, throttle factor, protocol, etc. Once the TOS groups the transfers with like properties, it sends separate tasks to the agent based on these groups. The default maximum number of machines per transfer group is 30. This is a configurable parameter.

SOS

The Schedule Object Server (SOS) handles scheduling of DTS transfers. The SOS interfaces with CCS WorldView and CCS Calendaring. In the CCS Calendaring module you create calendars to map to production requirements, for example, only transfer files between 11 p.m. and 7 a.m. In WorldView one would associate a machine or group of machines with a calendar. The SOS evaluates these relationships and triggers suspension or resumption of the transfer as they meet the configured parameters.

DTA

The Data Transport Agent (DTA) handles the actual file transfer. DTS uses agent-to-agent communication. There is a DTS agent required on both ends of the transfer. In the case of CA DSM, the TOS communicates with the agent on the sending side (the domain manager) when the target is a scalability server. The sending machine (the initiator) spawns a slave process to manage the specific transfer. The slave communicates with the receiving machine (the responder) to begin negotiating the transfer. The responder also spawns a slave process. Once communication is established, the initiator begins sending the file. The number of slaves is a configurable parameter; the default is 10. This policy is set in Configuration Policy DTS>Agent>Concurrency.

The initiator reads the file to be sent and copies it to a working area on the disk. The first step in the process is running any type of filter that was requested. There is a mandatory read filter that is run and, also at this point, file level encryption or compression is run.

The initiator proceeds to break the file into parcels. At this point any parcel-level encryption or compression is also done. The initiator then sends the first parcel of data. The responder gets the data and verifies it for integrity. If the parcel is undamaged, it sends a response to the initiator and performs any parcel write filters, such as compression or encryption, and places the data into the staging area.

This process continues until all the parcels have been sent. The responder then performs any file-level filters and places the file into its permanent location.

If the responder does not send an acknowledgement to the initiator, that specific agent is put into failure mode. The initiator agent will then wait for the timeout interval to try to resend the parcel. It will continue to do this until it reaches its retry limit, at which point it will mark the job as failed.

The above paragraph refers to a point-to-point transfer (one initiator, one responder). However, in most cases DTS will send to more than one machine at once, such as a fan-out or multicast/broadcast transfer. In this case the initiator will send the parcel to multiple targets and wait for responses from all the targets. If one or more targets do not respond with an acknowledgement, they will be dropped out of the target list temporarily. The agent will continue to send parcels, in order, to all the machines that are still responding. When all the machines fall into retry status, the agent will then go back to the point of the first failure and add all the agents back into the list of targets. It then will start sending to all those machines from the point of the last successful transfer. The initiator will continue this cycle until all targets have received the file successfully or the retry limit has been reached.

The initiator then marks the status as complete. CA Unicenter Software Delivery gets that status and then orders the agent on the destination to proceed with the installation.

The agents on both sides keep a status file. So when an agent restarts after an interruption like a system crash or power outage, it knows where it left off. When the transfer is complete or the job is deleted from the TOS, the status and staging files are deleted. These files are also deleted by the agent if they have been inactive for seven days. This value is configurable.

DTS Configuration

DTS is primarily configured in two CA DSM locations, Configuration Policy in the DSM Explorer and the 2D Map.

In order to be able to configure features in the 2D Map, CCS integration must be enabled. This is a configuration setting is located in configuration policy at DTS\Network Object Server\Enable CCS Worldview Integration and ensure it is set to True.

The first item to consider is where all the components are installed. Following is a table of the out-of-box component locations:

Component	Locations
Enterprise Manager	MDB, WorldView, CCS Event Manager, Event Agent NOS, TOS, SOS, DTA
Domain Manager	MDB, WorldView, CCS Event Manager, Event Agent NOS, TOS, SOS, DTA
Scalability Server	CCS Event Agent, DTA
Agent Machine	None by default. the DTA can be installed as needed

The architecture of where these components need to reside in a specific installation is critical to the efficient administration of the product.

NOS

The NOS is configured to get its data from a specific WorldView repository. Out of the box, the NOS assumes this to be local. If the MDB is remote, then one must configure the NOS to point to the WorldView repository on the MDB machine. If the WorldView manager is not local, the WorldView client must be installed on the machine where the NOS is running. It is strongly suggested that the NOS be physically close to the MDB. It should not be located where access is across a WAN connection. This linkage is configured in the configuration policy under Data Transport Service\NOS. The property name is CCS **WorldView Repository Name**. Out of the box it is configured to Local Host. Change this value to reflect the host name or IP address where the WorldView repository is housed.

There are a few other NOS configuration settings to consider as well:

Security Mode: By default, this value is configured to *fail*. With the value set to fail you will be prompted to enter a valid username and password in the 2D Map GUI or when using the command line to access the NOS. You may want to consider setting this value to *quiet* when you are first configuring the system. Once you have most of your configurations set, then you should set it back to *fail*. By doing this it will make configuration setting in the 2D Map a little faster because you do not need to authenticate each time.

Self Discovery: Create DTS Objects in WorldView: By default, the value is set to *true.* This means the DTS Objects will automatically be created in WorldView when the agent is detected.

SelfDiscovery: Create machines: By default, this value is set to *false*. If CA Unicenter Network and Systems Management (CA Unicenter NSM) is installed, it is suggested you leave this value alone and let the policy designated by the CA Unicenter NSM administrator add the machines to WorldView. But if WorldView is only being used to administer DTS or a separate WorldView repository is being used to administer DTS, there will be benefits to changing this value. By setting the value to *true* the machine will automatically be

populated into the WorldView repository when the DTS agent is detected. This eliminates the need to discover the object by other means.

There is a disadvantage to allowing the system to add machines automatically. If an unauthorized machine is added to the network, it will automatically be discovered by the system and added to the 2D Map.

Update Machine's IP Addresses: By default, this value is set to *false*. If the value is set to *true,* anytime the IP address of a device changes, the NOS will update this value in WorldView. If the machines in the environment have addresses assigned by DHCP and tend to change, it is best to change this value to *true.* But if the hosts are assigned static IP addresses, it may be best to leave this value as *false.* If the IP address of a device does change, a typical change control process will include editing the object in WorldView. If WorldView has the wrong IP address, the transfer will fail. This will enable one to know the IP address did change and investigate why.

Another item to consider when contemplating how to configure the above settings is network traffic. If these values are set to true, more network traffic will be generated between the TOS, NOS, and WorldView Repository. Each time an agent reports in to the TOS, the TOS will communicate with the NOS and in turn the NOS will contact WorldView to check on the status of that agent—does it exist and what is its IP address. If these functions are disabled, that traffic will not occur.

TOS

Network Location of the NOS: This value is the host name or IP address of the machine where the NOS this TOS is supposed to use is located. Again this is an architecture question. By default the value is *local host*. There is a fair amount of network traffic that is generated between these two services, especially if the auto-create/update features are enabled. A TOS is associated with one NOS. One NOS can be associated with multiple TOSs. In most cases the NOS/TOS communication should be on the same subnet and should not cross a WAN link.

The downside of multiple NOSs is added administration of the WorldView repository. If there are multiple repositories, it is important to keep track of which machines are parts of which TOS and, therefore, ensure the proper machines are populated to the proper WorldView repository.

In an environment where there is a CA DSM Enterprise Server and a number of Domain Servers, each with their own NOS, TOS and SOS, the WorldView repository population should be as follows:

■ The Enterprise NOS should have entries for all the Domain Servers.

■ The Domain NOS will have entries for all its scalability servers and any agent machines that have DTS installed and enabled and associated with that Domain Server.

In the case of the one NOS to TOS relationship mentioned above, it is suggested you use the WorldView agent located on each individual CA DSM server to administer its WorldView repository. This minimizes the risk of populating machines to a repository in error.

Transfer Group Size: This is the maximum number of transfers an individual transfer group will manage. The default is 30. The importance of this value and its interaction with other settings will be discussed later in this document.

DTA

The settings in the DTA configuration to be concerned about are as follows:

Concurrency: This is the number of concurrent data transport processes allowed to occur at one time. The default is 10. Each time a transfer is requested, a slave process is spawned. The interaction of this setting and others will be discussed later in this chapter.

Concurrency to reserve for responders: This setting dictates how many slave processes are held for incoming transfers. This setting is critical on installations where domains are a member of an enterprise and when agents are configured to use DTS. The default is 1. Take the domain manager, for example, with the default setting of 1. This guarantees that one slave process on the domain is reserved for incoming transfers from the enterprise manager. On the scalability server, one is reserved for stage jobs from the domain manager. If the configuration does not have agents configured to use DTS, then the setting is not critical. Incoming transfers will use as many slave processes as are available. Outgoing transfers will use as many as they need, but will never take the ones reserved for incoming transfers.

Be careful if you adjust these settings upward. These processes do consume memory and processor resources, and increasing these limits can cause adverse effects on your system.

CA Unicenter Software Delivery Configuration

These settings are located under the Configuration Policy\Default Computer Policy\CA DSM\software delivery\file transfer tree node. The items of concern here are:

DTS: Session Retry Interval: This setting is used to specify how often DTS should retry after a connection failure due to a network problem. The default is 60 seconds.

DTS: Session Retry Limit: This setting specifies how many times DTS should retry to establish a connection to continue the transfer. The default is 10 and there is no maximum retry limit.

Configuration Considerations

There are a number of values that are configured in WorldView that affect how transfers are done. They are explained very well in the product documentation. This document will touch on just a few as they will affect the overall performance of DTS.

When configuring DTS parameters in WorldView, be consistent. Try to group like objects as much as possible. The more objects you have with the same parameters, the more efficient DTS will be. For example, create DT Containers as either static or dynamic to reflect the

properties of your network. Create a container for your 56 KB links, another for your 128 KB links, etc. Then populate those containers with the machines associated with those network links. Then configure the container accordingly.

As mentioned earlier, the TOS creates transfer groups. The maximum number of machines in a group is configurable (default 30), but that number will vary based on other factors. In order for a transfer to be part of a group it must have identical parameter values. For example, the file name must be the same, and the parcel size and throttle factor must be the same. The type of transfer must be the same, that is, point-to-point, fan-out, multicast, etc. This being said, it is critical to group like objects with the same values. In other words, develop standards and enforce them. Do not have five machines with a 56 KB parcel size and one with a 52 KB parcel size. That would require two transfer groups, one for the five machines and one for the single machine. The values are close enough that they all can be 56 or 52.

So why is this so important? If you remember, there is a limit to the number of slave processes (10 by default). If you have parameters that are not consistent, you can end up with one member per transfer group. You can only have ten transfer groups active at one time. So in this scenario you would only be able to transfer data to ten machines at a time. This is not very efficient.

Other values to keep in mind are the retry interval and retry count. Oftentimes customers will set these values very high to compensate for poor quality network segments. The problem here is that a transfer may be trying and trying for days to complete a transfer to the one machine with a downed network link. Remember the max concurrency setting again. If a slave process is open for days continuing to retry 900 times with 30 minutes in between, that job will not fail for 45 hours—tying up that slave. If you have one or two bad links per job, it means all 10 slave processes will be tied up for two days. It is best to keep the retry count lower and let the jobs that are going to completely finish in a reasonable time run. Then address the bad links on an ad hoc basis.

When working with slow links, limiting the parcel size is crucial. Ensure that the parcel size is set to a size such that a complete parcel can be sent in no more than 50 seconds. The initiator will be waiting for a response from the responder acknowledging receipt of the parcel. If it does not get that response, the initiator will assume the parcel did not get sent and resend it.

Lastly, one other item of note… If machines are populated in WorldView, make sure that both the initiator and responder are in the map and configured with values. The TOS will take the most restrictive value. For example, if the initiator has a parcel size of 256 KB and the responder has a parcel size of 56 KB, the data will be sent with a parcel size of 56 KB. However, if only one value is found in the NOS, then the TOS will use its default value of one half of one MB. Consider the scenario where you decide to configure the slow links in WorldView but do not do anything to the domain manager in WorldView. Since there is no configuration for the domain manager in WorldView, the transfer is done at the default setting and fails because the slow link cannot take all of that data at one time.

Chapter 8: Migrating to r11

Introduction

If you are running a previous version of CA Unicenter Desktop & Server Management (CA DSM), there are two routes to getting the agents to the new r11 infrastructure:

1. **The Migration Tasks**: These engine tasks allow you to move the data to the new infrastructure (or a shared infrastructure) so you can do common reporting from day one, and plan your migrations steps.

2. **The Agent Bridge**: This bridge allows an older agent to be managed directly from the r11 infrastructure. This is a good solution if you have a new infrastructure and do not need to inherit all of the data from the old environment. The Agent Bridge also allows you to use agents not yet ported to an r11 code set. This functionality is similar to the Windows Mobile functionality, which is described in the Mobile Devices chapter later in this book.

These two methods can be combined to achieve both common reporting and legacy management. These topics are addressed in the Agent Bridge section of this chapter.

The Migration Tasks

Detailed procedures and guidelines for upgrading to the r11 release from a previous version of CA Unicenter Software Delivery, CA Unicenter Asset Management, or CA Unicenter Remote Control can be found in the *Unicenter Desktop & Server Management Implementation Guide*. Following is an overview and additional tips to further assist you in upgrading.

Direct software upgrade to release 11 is not supported. Rather, the upgrade path uses a parallel or side-by-side approach. First, the new r11 components are deployed (domain managers, scalability servers, and other components) typically using a new, more optimal topology. Then data is migrated from the pre-r11 databases to the new r11 CA MDB. Following is a list of supported upgrade paths:

- CA Unicenter Software Delivery r11 can be migrated from Version r3.1 and Version r4.0

- CA Unicenter Asset Management r11 can be migrated from Version r3.2 and Version r4.0

- CA Unicenter Remote Control r11 can be migrated from Version r6.0

The migration to CA DSM r11 is about homogenization of three distributed solutions with different architectures and separate databases into one consolidated r11 architecture.

The change is a significant one in that three formerly disparate and independent technologies have now merged into a single technology employing a single management database structure. Additionally, as each product is multi-tiered, an upgrade approach would require extensive backward compatibility support between the tiers *during* the migration process that would not add any value post migration.

The end result is that a simple upgrade strategy will not suffice.

The *Unicenter Desktop & Server Management Implementation Guide* includes a migration chapter which describes the general migration approach, limitations, and challenges, as well as the specifics regarding ASM.CNF keys, MSI library migration, and much more. We strongly suggest that you read this chapter.

Migration to r11 is achieved by three distinct steps:

1. The installation of management infrastructure

2. The migration of data

3. The installation of new agent software

Data Migration

The data migration step is driven by engine tasks, with one task type per product (CA Unicenter Asset Management, CA Unicenter Software Delivery, Operating System Installation Management (OSIM), and CA Unicenter Remote Control). There may be multiple legacy managers and multiple tasks for each legacy manager—which could add up to many tasks.

Migration is a continuous process and all migrated computers are visible in the system group All Legacy Computers. The process ends when the r11 agents are registered with the MDB—at which point you should *stop* managing those assets through legacy managers.

CA Unicenter Software Delivery—Specific Migration Information

There is a slight difference between what can be migrated at the enterprise and domain levels in terms of:

- Software packages

- Software groups

- Software policies

- Computer and user groups

- Computers and users (domain only)

- Computer job history (domain only)

- Security

To use the Unicenter Software Delivery 4.0 SP1 SDOAPI to get to v4.0 data, follow these steps:

1. Install the legacy API:

    > On Windows, install the legacy API from DSM Installpath\SD\Legacy\usd40sp1\sdapi.w32 *IF and only IF* Unicenter Software Delivery 4.0 SP1 and CA DSM r11 are *not* co-hosted (If they are co-hosted, the API is already installed).

    > On Linux, the legacy API is installed with the CA DSM r11 manager and no additional actions need to be taken.

2. Migration of Software Packages:

    > Exports packages and added procedures and imports them to r11.

    > Is only done once.

    > Limited Group Membership update on each run – unlinking is not detected on legacy manager.

3. Migration of Software Groups:

> Is only done once.

> Limited Group Membership update on each run – unlinking is not detected on legacy manager.

4. Migration of Software Policies (aka Software Templates):

> Is only done once.

> Never sealed automatically.

> Sealed Software Policies are migrated on the domain only.

5. Migration of Computer and User Groups:

> Is only done once.

> Limited Group Membership update on each run – unlinking on the legacy manager is not detected.

> Limited query migration. Many queries cannot be migrated properly but they will appear in r11 as 'disabled' and can be redefined using the DSM Explorer post migration.

6. Migration of Computers and Users (Domain only):

> Is only done once.

> Limited Group Membership update on each run – unlinking on the legacy manager is not detected.

> Scoped migration can be performed by naming the Computer and User group on the legacy manager. Only members of that group will get migrated.

Note: Watch out for duplicate HOST UUIDs. HOST UUID is used as the primary identifier of computers in r11. Duplicates WILL cause problems. The following SQL can help you identify the duplicate UUIDs:

```
USE SDDATA
select cast(c_lanname as varchar) as name, cast(c_uuid as varchar) as
uuid
from cadb.T_computer
where c_uuid <> 0x0000000000000000000000
-- This is for Computers and Staging servers
-- For mobile and computer user use 7 and 2
and (c_type = 0 or c_type = 3)
and c_uuid in (select c_uuid from cadb.T_computer
where c_uuid <> 0x0000000000000000000000
-- This is for Computers and Staging servers
-- For mobile and computer user use 7 and 2
and (c_type = 0 or c_type = 3)
group by c_uuid having (count(*) > 1))
```

7. Migration of Computer Job History (Domain only):

> Migrated every time.

> Records replaced on consecutive runs.

> Records merged on last run.

8. Migration of SD Security:

> Only migrated once.

> Security Profiles are migrated regardless of security provider conflicts (winnt/unix). They can be remapped to other OS security groups post migration using the DSM Explorer.

> Object ownership is only migrated if security providers match (winnt/winnt, unix/unix, not winnt/unix).

The migration of computer and user data occurs in the following phases:

9. The Computer or User is created by the migration task and set in the 'locked by migration' state in the CA DSM r11 manager to prevent the CA DSM r11 manager from setting up jobs for it.

10. When the r11 Agent registers with the CA DSM r11 manager, the agent version is updated. The Computer or User *remains* in the 'locked by migration' state in the CA DSM r11 manager.

11. The next time the migration task is run, the installation history from the legacy manager is merged with the existing history in the CA DSM r11 manager. Then the Computer or User state is unlocked from migration.

At this point, management of the Agent using r11 manager can begin. You should STOP managing the agent using the legacy manager, as no further migration of job history will be performed!

With the INI files changed to a common configuration, you might have Software Delivery packages that will have to be modified or dead information in the ASM.CNF file like this:

```
DIM sdserver,asmenv as String
asmenv=EnvGetString("asmroot")
IF ReadIniEntry("PEERHOSTS","AGENTSRV",sdserver,asmenv+"\asm\conf\asm.cnf") = 0
THEN
  ReadIniEntry("PEERHOSTS","ASMEM_Local",sdserver,asmenv+"\asm\conf\asm.cnf")
ENDIF
```

An example of the migrated script is:

```
DIM sdserver as string
CcnfGetParameterStr("ITRM/agent/solutions/generic/serveraddress",sdserver)
```

Detailed information on how ASM.CNF is migrated into r11 comstore is documented in the *Unicenter Desktop & Server Management Implementation Guide*. It should be referenced for the migration of user data including User, Room, Phone, and Comment, and customized agent name.

Important CA Unicenter Software Delivery Migration Notes

■   Migration is invoked by the agent installer if selected by the user.

■   Migrated configuration settings are protected by CCNF from override by the CA DSM r11 default policy. Only custom policy will override the migrated settings.

■   On Managers, the library is copied by a migration task - it is NOT automated on scalability servers. You will need to run sd_sscmd import to copy or move packages from legacy staging servers.

■   MSILIB is not automatically migrated on domain managers and scalability servers. You will need to run sd_sscmd importmsi to copy or move packages from a legacy local or staging server MSILIB.

■   In r11, software packages are identified by DSM UUID – *not* by their MSI Product Code. This appears in MSILIB in the folder named with this DSM UUID.

■   The DSM UUID is maintained during export and Enterprise to Domain distribution.

■   It is important to use the same DSM UUID throughout the Enterprise and all Domains to improve roaming source list updates.

■   The new MSILIB share name is 'SDMSILIB' and a new MSI dictionary file is provided to improve roaming source list update. In addition, new MSI procedure macros replace old macros used to find admin installs in new locations. The old macros are updated during package import. For more details, consult the *Unicenter Desktop & Server Management Implementation Guide*.

CA Unicenter Software Delivery Data Migration Implementation

The process of migrating Unicenter Software Delivery-specific data is as follows:

■   The Engine loads the usdLegacy DLL/SLIB.

■   The usdLegacy DLL/SLIB calls the sdmgrmig executable using instructions it received from the Engine.

■   Sdmgrmig loads usd40sp1 DLL/SLIB to connect to the legacy manager.

■   Sdmgrmig loads ps DLL/SLIB to connect to the database.

■   Sdmgrmig loads coApi DLL/SLIB to connect to the r11 common manager.

■   Migration is done for each selected object class.

■   All tracing goes to TRC_MIGRATION_USD.

As with previous installations, all CA Unicenter Software Delivery events are also logged to the NT(/CCS) event console. For example:

Very limited feedback is written to the Engine task portal, but you can see if the task has completed and what the result of the run was. For example:

Task	Status	Last Executed	Actions
Default Directory Synchronization Job	OK	10/18/2005 9:25:05 AM	X → ↓
computer-name Collect	OK	10/18/2005 3:18:47 PM	X → ↓ ↑
MIG - USD - computer-name	OK	10/18/2005 9:41:11 AM	X → ↓ ↑
computer-name Collect	OK	10/18/2005 3:19:48 PM	X → ↓ ↑
computer-name Collect	OK	10/18/2005 3:20:49 PM	X → ↓ ↑
MIG - UAM - Chili	OK	10/18/2005 3:25:35 PM	X → ↓ ↑
MIG - URC - Chili	OK	10/18/2005 3:33:32 PM	X → ↓ ↑
MIG - USD - Chili	**Active**		X → ↓ ↑
MIG - OSIM - Chili	Waiting		X → ↑

Sdmgrmig can be executed from a command line as well. For information on parameters supported by this command, execute it without providing parameters.

The installer sdcnfmig executable is used for information regarding migration of the ASM.CNF file. This executable loads rdcnf DLL/SLIB to read local ASM.CNF and write to the r11 comstore. It can be executed on a command line as well. To view the list of supported parameters, execute the command without providing any parameters.

OSIM-Specific Migration Information

The following OSIM data is migrated:

■   OSIM OS images and OSIM Boot images. Default parameters are kept synchronized.

■   The OSIM Configuration. Current states of OSIM computers are migrated.

- OSIM Computer Groups.

The prerequisites include:

- Target computers must be migrated first (by running the CA Unicenter Software Delivery migration job).

- OS Images must be migrated manually (see the *Unicenter Desktop & Server Management Implementation Guide* for a description).

- Collect all information to access the Unicenter Software Delivery 4.0 database (including database type, host, database name, credentials).

- OSIM Migration imports data related to operating system installation from a Unicenter Software Delivery 4.0 Local Server into r11. This includes parameter values and Unicenter Software Delivery 4.0 OSIM groups.

There are two kinds of parameter sets:

- Default parameter values of operating system images. The parameters of each Unicenter Software Delivery 4.0 OS Image are mirrored to the r11 OS image with the same name.

- Current parameter values of target computers (taken from the last successful installation). The current parameters of each Unicenter Software Delivery 4.0 OSIM computer are mirrored to the r11 computer with the same name.

The log file used is TRC_Migration_CSM.log.

An r11 group is created for each Unicenter Software Delivery 4.0 OSIM group and an '-r4' suffix is added to the name of the new r11 groups. Memberships are updated on the next run.

OSIM Data Migration Implementation

OSIM data migration runs as an engine job and is implemented as a DLL (ccsmmig.dll) or a UNIX library (libccsmmig.so), depending on the platform.

See the Operating System Installation Management chapter later in this book for more details regarding OSIM.

CA Unicenter Asset Management—Specific Information

The following CA Unicenter Asset Management data is migrated:

- Computers – including those in the 'Legacy' state. Use the migration scope to identify which computers to migrate.

- Static Groups (also link members).

- Dynamic Groups - depending on which queries are migrated.

- Computer General Inventory.

- Computer Software Inventory - depending on which Software Definitions are migrated.

- Jobs. Note that these are not linked to groups or assets. Status will not be migrated.

- Templates.

- Configuration Files.

- Queries. Although most will be migrated automatically, those that are not migrated successfully will be flagged.

- Query-based policies - depending on which queries are migrated.

- Event Policies.

- Categories.

- Publishers.

- Application Definitions. Note that heuristic applications will *not* be migrated.

- Only Software Definitions with corresponding software in Unicenter Asset Management r4.0 will be migrated, unless the Software Definitions belong to a category labeled 'Migration.'

- Report templates.

- Scheduled Reports.

CA Unicenter Asset Management—Specific Data Migration Process

Following is the process by which CA Unicenter Asset Management-specific data is migrated:

- CA DSM connects to the legacy database.

- Important data is verified and loaded into cache.

- Each individual configured item is migrated.

- The legacy database is closed.

The list of legacy objects is maintained in the amlegacy_objects table. This table maps the Unicenter Asset Management r4.0 primary key and the r11 primary key for each object. Objects that are mapped will not be migrated again (with the exception of an update of inventory on Computers).

Objects will be mapped to existing r11 objects if the name matches (with the exception of Computers). Computers are mapped to existing computers if host_uuids are the same or if the host_name is the same. To identify duplicates you can use this SQL:

```
--
-- Find Duplicate AM Units
--
USE CAAMDB
select UNITID, NAME from UNIT
where TYPE=1 and DMUUID <> '' and DMUUID in
(select DMUUID from UNIT
        where TYPE=1
        group by DMUUID having (count(*) > 1))
```

A computer's general inventory and filescan-based software inventory will be migrated. Software Definitions for these, however, have to be migrated separately.

To verify migration status, check the TRC_MIGRATION_UAM_0.LOG which is found in the CA DSM logs folder. This file provides runtime logging.

CA Unicenter Remote Control—Specific Information

The following data is migrated for the CA Unicenter Remote Control component:

- Computer Groups.

- Global/Local Address Book Information.

- Computers. The state will be listed as 'migrated.' Use the migration scope to identify which computers to include.

- Configuration Policy Data.

All necessary data for each object type is copied from the CA Unicenter Remote Control r6 database to memory, and each object is inserted repeatedly. Some object types are updated if already present - depending on responsibility. Inserting duplicates does not interfere with migration. The following data is inserted:

Computer -> ca_discovered_hardware(machine), ca_agent(RC status 'Migrated'), ca_group_member(group membership)

ComputerGroup -> ca_group_def(r11 group), ca_agent, ca_group_member(group membership)

Address book data -> ca_group_def(r11 group), ca_agent, ca_group_member(group membership), urc_ab_group_member(indicating address book), urc_ab_computer(with connection addresses), urc_ab_permission(permissions for address book groups)

Important CA Unicenter Remote Control Migration Notes

■ You can delete objects in the DSM Explorer and migrate again to recreate them.

■ Note that configuration policies are not deleted instantly although you cannot see them in the GUI any longer.

■ CA Unicenter Remote Control r6 address book groups are transformed into standard r11 groups, with some extra address book properties. R6 configuration policies are added to the r11 configuration by CCNF client calls. The sequence of calls is generated from the r6 data in memory.

■ To view migration status, check the TRC_MIGRATION_URC_0.LOG file. You can also check the Engine Status. For example:

Task List

Task	Status	Last Executed	Actions
Default Directory Synchronization Job	OK	10/18/2005 9:25:05 AM	✕ → ↓
computer-name Collect	OK	10/18/2005 3:18:47 PM	✕ → ↓ ↑
MIG - USD - computer-name	OK	10/18/2005 9:41:11 AM	✕ → ↓ ↑
computer-name Collect	OK	10/18/2005 3:19:48 PM	✕ → ↓ ↑
computer-name Collect	OK	10/18/2005 3:20:49 PM	✕ → ↓ ↑
MIG - UAM - Chili	OK	10/18/2005 3:25:35 PM	✕ → ↓ ↑
MIG - URC - Chili	OK	10/18/2005 3:33:32 PM	✕ → ↓ ↑
MIG - USD - Chili	**Active**		✕ → ↓ ↑
MIG - OSIM - Chili	Waiting		✕ → ↑

CA Unicenter Remote Control Data Migration

Data migration for CA Unicenter Remote Control is managed through the rcLegacy.dll and the rcManagerR11Migration.exe. The Engine calls rcLegacy with the data provided by the Migration Wizard. rcLegacy builds a command from the data and executes it.

This command can also be run standalone. For information on the syntax, execute the following:

    rcManagerR11Migration.exe -?

The Agent Bridge

The Agent Bridge serves several purposes. In general you can use it either for migration or for managing the agents that have not yet been ported to a true r11 code base. This section will cover important points for using the Agent Bridge in connection with migration.

The Agent Bridge is basically an interface that sits on one or more scalability servers and acts as a converter between r11 and the following legacy agents:

■ CA Unicenter Software Delivery from Version 3.0 and up

■ CA Unicenter Asset Management from Version r3.2 and up

Note: For the r11.2 Agent Bridge, CA Unicenter Asset Management agents prior to version 4 need a patch for generating the HostUUID. This issue has been addressed in the r11.2 C1 version of the Agent Bridge.

If you choose to use the Agent Bridge, you just need to move the agent connection from the CA Unicenter Asset Management sector or the CA Unicenter Software Delivery staging server to the new r11 scalability server where the Agent Bridge is installed.

Following is an example of a Unicenter Software Delivery 4.0 script that can be used to move the agent for the 4.0 infrastructure to the r11 infrastructure:

```
DIM server AS STRING
DIM INIFile as STRING
DIM rhdl,Dummy as Integer

server = "newdsm.ca.com"

EXECUTE("SDACMD.EXE SetServerAddress "+ server + " -f")

'Read UAM Install Location from Registry
rhdl=RegOpenKey(HKEY_LOCAL_MACHINE,"SOFTWARE\ComputerAssociates\Unicenter Asset
Management\InstalledComponents")
RegQueryVariable(rhdl, "Agents", INIFile, Dummy )

INIFile=INIFile+"\ncc31com.ini"
WriteINIEntry("Sector Mapping","RPCServer",Server,INIFile)
```

If you choose to use the Agent Bridge for the migration you need to be aware of the fact that none of the data from the old infrastructure will be migrated. When you move to the new infrastructure, you will in effect have a new installation. You can, therefore, successfully use the migration jobs discussed above to move things like:

■ Group membership

■ Installation history

■ Security settings

Implementing the Agent Bridge

The Agent Bridge has been introduced as an add-on for the r11.2 code line. It is not part of r11.2 or r11.2 C1, but it will be included in a future version of CA IT Client Management/Desktop Server & Management. Until the Agent Bridge is included in the GA code, you can obtain a copy by contacting CA Support and requesting the TestFix for the Agent Bridge for r11.2 or r11.2 C1. You then need to install it according to the instructions in the TestFix.

You can install the Agent Bridge on one or more scalability servers, depending on how you plan to use it. For example, if it is solely for legacy support, it could be an advantage to have dedicated legacy scalability servers. But for migrations, you would want to have an Agent Bridge on all scalability servers to upgrade the agents seamlessly without moving the scalability server.

Once installed, you will need to enable it. This is done using CA DSM Common Configuration. Using the DSM Explorer, navigate to the appropriate configuration policy and then select 'Asset Management' under the 'Scalability Server':

You will see the following fields:

■ **Agent Bridge Interval** – The default value is 60 seconds. This parameter controls the frequency of running the amBridge.exe process.

■ **Enable Agent Bridge** – The default value is False. Agent Bridge AM server will be started automatically if set to True.

■ **Enable Legacy User Support** – The default value is False. The legacy user account will be registered if set to True. **Note**: This will not be the real r11 user, as it does not have all of the information on Domains and so on.

For Software Delivery:

Using the DSM Explorer, navigate to the appropriate configuration policy, expand 'software delivery' and choose 'Backwards compatibility support' under the scalability server:

You will see the following fields:

- **Backwards compatibility support: Enable** – The default value is False. Set it to True to switch on the SD Agent Bridge server.

- **Backwards compatibility support: Use the SDOUTPUT share** – The default value is False. This parameter enables/disables the support for the SDOUTPUT share. This parameter should be set to True if CA Unicenter Software Delivery agents of a version prior to 4.0 are connected and using NOS as the CA Unicenter Software Delivery library access method.

- **Enable uuid Generator:** This parameter controls the UUID generation behavior of the Agent Bridge. If enabled, the Agent Bridge will generate and maintain a host UUID for any agent that does not report one. The generated host UUID is derived from the host name of the computer. **Note**: This parameter was introduced in CA DSM r11.2 C1 and is not shown in the screen shot above.

Working with the Agent Bridge

Once you have pointed your legacy agent to the scalability server, you will see the legacy agent just like an r11 native agent. In certain cases you may note some missing information, but in general it is transparent.

To identify if the agent has been registered, right-click on the computer, select 'Properties,' and then the Agent tab:

With the Agent Bridge active, the data shows up as it does in r11. You can use the data and run reports as you do with native r11 agents.

Unsupported Features

The following functions are not supported with the Agent Bridge:

- **Online Metering**: You will not be able to set up an online metering environment, which includes 'deny' and 'warning' features. However, you can still use offline metering.

- **Limitations on Offline Metering**:

    > No version number will be monitored for the signature/application definition-based metering.

    > The metering is based on the binary file name.

- **No Signature Scan**: As the signature scan was introduced with r11, this feature is not supported for the legacy agents.

- **No File Scan**: In r11 the file scan is replaced by the signature scan. Therefore, the file scan is not supported for legacy agents.

    Note: File Management is supported. Even though the file scan is not supported, you can enable the File Manager function to collect meta data for all the files on the hard drives.

- **No Common Configuration**: The Legacy agents do not support the new r11 Common Configuration, so all configuration needs to be done using the old methods like SDCONF, ASM.CNF, NCCLIENT.INI, and so forth.

Upgrading the Legacy Agent Using the Agent Bridge and CA Unicenter Software Delivery

As you have full management over the legacy agents and have all the r11 agent's packages already registered with the CA Unicenter Software Delivery system, you can simply send the r11 agent package to all the agents using the r11 DSM Explorer. When upgrading, it is extremely important that you either disable or uninstall the old agents so that the system does not get confused with two agents connecting to the same system.

You can take two routes:

1. **Uninstall the agents when you install r11**: For the Windows agents, all the r11 packages have an MSI Property REMOVELEGACY. If you set it to '1', then the legacy agent is uninstalled as soon as the r11 agent is installed.

The easiest way to get this done is to open the existing CA Unicenter Desktop & Server Management packages and create an added procedure. Then change the embedded file to include the REMOVELEGACY parameter.

Note: You will need to do this on all of the packages that you are planning to use from the migration.

2. **Uninstall the agents after you install the r11 agent**: The system comes with a built-in CA Unicenter Software Delivery package that can remove the legacy packages after the installation.

This package allows you to specify whether to remove all or only some of the agents.

Chapter 9: Asset Registration and Reconciliation

During the continuous usage of CA Unicenter Desktop & Server Management (CA DSM) new machines will be registered, old machines will be decommissioned, and machines will be reinstalled with new names or with new hardware. It is important that CA DSM is aware of and able to handle this changing environment, and that it can identify if a machine is a new machine with an existing name or the same machine with a new name.

How CA DSM Recognizes an Asset as New or Existing

To handle this, CA DSM uses what is called a HostUUID. This is a UUID (Universally Unique Identifier) which is generated by the agent at first execution, and then is used as a key to determine if the machine is new or not.

For detecting the machine uniqueness, CA DSM uses two passes, one on the agent and one on the engine. The reason for the two passes is that the agent does not have direct database access. It could have been moved from one scalability server to another, or the information could have been deleted at the agent but still exist in the database.

Asset Registration Process

This Agent Registration Process (also referred to as the CAINF process) basically uses three sets of inventory attributes:

■ **System UUID:** This is a unique ID created by the vendor of the motherboard.

■ **MAC Address**: This will be a list of all the physical MAC addresses detected on the machine. The list is treated as a pool. A change will be noted only if *all* MAC addresses are changed.

■ **Disk Serial Number**: This is a list of the serial numbers on all hard drives in the machine. Like MAC addresses, they are all seen as a pool. So a change is noted only if *all* hard drives have changed.

These three attributes are compared to the previously detected attributes, and if at least two out of three have changed, a new HostUUID is generated. The following chart illustrates the possible combinations of changed attributes and the resulting action on the HostUUID:

System UUID	MAC Address	Disk Serial Number	New/Keep HostUUID
Same	Same	Same	Keep
Same	Same	Changed	Keep
Same	Changed	Same	Keep

System UUID	MAC Address	Disk Serial Number	New/Keep HostUUID
Same	Changed	Changed	New
Changed	Same	Same	Keep
Changed	Same	Changed	New
Changed	Changed	Same	New
Changed	Changed	Changed	New

The attribute values are stored together with the Host UUID in the registry:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\HostUUID

The r11 agent registration process also accommodates the case where the agent has been installed as part of an image, like Ghost or ImageCast. In previous versions of CA DSM, you had to delete the HostUUID manually. However, with the new r11 logic the change will automatically be detected and the HostUUID will be regenerated. It is, however, still recommended that end users remove HostUUIDs from images before cloning them.

Note: If you want to disable the r11 agent logic, you can add this to the registry:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\HostUUID]
"LockHostUUID"="1"
```

When a new HostUUID is generated, an agent registration request is sent to the scalability server together with a full set of inventory.

Engine Registration

Once the data has been sent through to the scalability server, the engine will do a second pass on the host to see if the machine already exists in the MDB and if it needs creation, renaming, or updating.

The engine will do a database lookup based upon HostUUID and hostname, together with primary MAC address. Based on the change detected, the engine will handle the registration as follows:

■ **No Change:** Data is stored in the normal way.

■ **Hostname Changed:** The asset record is renamed and the data is collected as normal.

■ **HostUUID Changed:** The MAC address pool is checked to see if it is a new machine or a new HostUUID. If any of the MAC addresses in the list matches an existing one, the

asset record is updated with the new data. If none of the MAC addresses match, a new record is created.

Important! This will result in two records with the same name, but representing two different machines. This is new with CA DSM r11.

■ **Both Hostname and HostUUID Changed:** A new record is created.

Once CA DSM is ready to register the asset to the MDB, it uses the common registration process to ensure integrity between all the products sharing the same MDB. The next section will describe the common registration process.

Common Asset Registration for CA Products

The asset is central or prominent in many CA products. Assets can be inserted into the MDB by a variety of sources including discovery tools such as CA Unicenter Asset Management and CA Unicenter Network and Systems Management (CA Unicenter NSM), or ownership tools such as CA Unicenter Service Desk or CA Unicenter Asset Portfolio Management (CA Unicenter APM).

Even though an asset can be discovered by multiple products, the asset schema is designed to reconcile the fact that an asset coming from different sources is actually the same asset. This is accomplished through asset registration using the CORA (Common Object Registration API.)

The asset schema accomplishes the goal of providing a set of common asset tables for hardware and software assets. The schema allows for a cross-product view of assets and meets the following requirements:

■ Allow for the discovery of an asset by multiple sources (for example, CA Unicenter Asset Management and CA Unicenter NSM)

■ Allow for the discovery of multiple values of an identifying asset property (for example, multiple DNS names and/or MAC addresses)

■ Allow for multiple 'virtual' assets to be discovered and reconciled to one physical asset (for example, VMWare images or dual-boot scenarios)

Each of the products discussed here (CA Unicenter Asset Management, CA Unicenter APM, CA Unicenter Service Desk, and CA Unicenter NSM) leverages CORA to register an asset when it comes into view using the product's particular operations. For example, when CA Unicenter NSM discovery occurs, each discovered asset is registered. Similarly, when a CA Unicenter Service Desk or CA Unicenter APM user enters asset information using the data entry forms of those products, that asset is also registered.

Consider the scenario where we have run an initial first-level discovery with CA Unicenter NSM, or where a new asset is registered due to discovery by CA Unicenter NSM Continuous Discovery. That server or desktop asset is registered with the identifying properties known to CA Unicenter NSM, typically DNS name and MAC address. Subsequently, a CA Unicenter Asset Management scan is performed for that same server. CA Unicenter Asset

Management also registers the asset, also using the DNS name and MAC address, plus additional identifying properties. Because the DNS name/MAC address pair matches the previously registered asset, the information held by CA Unicenter NSM is now effectively joined to the information held and managed by CA Unicenter Asset Management.

If a user is creating a ticket in CA Unicenter Service Desk for that same server, when they enter in their information as prompted by the ticket creation web forms, CA Unicenter Service Desk registers the server, and it matches with the previously discovered information. Now, the additional information entered through CA Unicenter Service Desk is also available through the unified view of the asset in the MDB.

The following CORA section contains a detailed discussion of registration scenarios and how the matching and reconciliation occurs and is represented.

Common Object Registration API (CORA)

All Unicenter r11 products utilize the common MDB schema to store and manage their data. As the interface through which these assets are registered and as the only source for updating these tables, the CORA ensures that asset data flows consistently, thereby supporting the data and referential integrity of the MDB's master asset data model.

The master asset data model consists of the following three levels of asset references:

- ■ The asset source level, which consists of the **ca_asset_source** table, is used to track assets as they enter the system from different data sources, whether input manually or through discovery.

- ■ The logical asset level, which consists of the **ca_logical_asset** and **ca_logical_asset_property** tables, is used to store virtual assets. The logical asset level acts as a middle layer that exists between the data source and the physical level

to accommodate assets embedded in other assets such as VMWare sessions or dual-boot scenarios.

- Finally, the physical asset level, which consists of the **ca_asset** table, stores the identifiers that define the object as a distinct, physical asset.

After CORA is given a set of registration identifiers from the calling r11 application, it performs one of the following actions:

- **Return** the asset source reference if the registration identifiers match an existing asset, thus preventing duplicate assets from being registered.

- **Insert** a new physical, logical, logical property, or asset source record into the database depending on where the mismatch occurs. This step also prevents duplication of data by inserting records only at the appropriate levels. For instance, if there are no physical assets that can be identified by the registration identifiers, a new physical asset is created. However, if a physical asset can be identified by the registration identifiers, but not a logical asset, then a new logical asset is created and linked to the existing physical asset.

- **Update** an existing identifier(s) in the database with one of the registration identifiers. In this scenario, a single physical asset can be identified by the registration identifiers and one or more identifiers need to be updated.

- **Merge** two physical or logical assets together. In this scenario, CORA received information indicating that two or more physical assets are, in fact, the same asset. The existing physical assets are merged together to form one asset, and information for each asset is stored in ca_logical_asset_property table.

For r11.1, when a product registers an asset and CORA generates a UUID that matches an existing asset, CORA also automatically links (reconciles) Owned and Discovered information for that asset.

To determine which CORA version is being used by the product, execute the following command:

coraver

Asset Matching Logic

When an asset is registered, CORA generates the asset UUID (ca_asset) by applying black-box logic to the following six properties:

- Serial Number

- Asset Tag (appearing as Alt Asset ID)

- Host Name

- MAC Address

- DNS Name

■ Asset Label (Name)

CORA applies the following weighting system to these properties to determine if a match exists. Since certain properties are considered 'more important' than others, CORA will recognize a duplicate based on those values alone.

■ **Serial Number** is the most highly weighted field. Two assets with the same serial number will **always** be matched by CORA *unless* Asset Tag or Host Name is different.

■ **Alt Asset ID** is the second most highly weighted field. Serial Number and Alt Asset ID appear at the highest level of the asset registration schema in ca_asset. If Serial Number and Asset Tag match, CORA can create a new asset *only* if the Host Name is unique.

■ **Host Name** appears in the middle level (ca_logical_asset). If Serial Number and Alt Asset ID are blank, the Host Name takes precedence over DNS and MAC Address values. Although more than one DNS/MAC pair can be specified for the same Host Name, it will still be considered the same Asset.

■ **DNS Name** and **MAC Address** are weighted the same. CORA will recognize the same asset if DNS or MAC address match and will create a new asset when they do not.

■ Finally, although **Asset Label (Name)** is required to create an asset, you can have multiple assets with the same name as long all the other CORA fields are empty.

The following table shows how CORA determines uniqueness of an Asset. The intent here is not to show every single combination but to show enough of the behavior so that one could determine what would occur based on the properties they choose to include when registering an Asset in release r11.

Serial Number	Asset Tag	Host Name	DNS Name	MAC Address	Asset Label	Results
Unique	Unique	Unique	Unique	Unique	Unique, Duplicate or Null	New Asset
Unique	Null	Null	Null	Null	Unique, Duplicate or Null	New Asset
Null	Unique	Null	Null	Null	Unique, Duplicate or Null	New Asset
Null	Null	Unique	Null	Null	Unique, Duplicate or Null	New Asset
Null	Null	Null	Unique	Null	Unique, Duplicate or Null	New Asset
Null	Null	Null	Null	Unique	Unique, Duplicate or Null	New Asset
Null	Null	Null	Null	Null	Unique	New Asset
Null	Null	Null	Null	Null	Duplicate	Duplicate

Serial Number	Asset Tag	Host Name	DNS Name	MAC Address	Asset Label	Results
Null	Null	Null	Unique	Duplicate	Unique, Duplicate or Null	Duplicate
Null	Null	Null	Duplicate	Unique	Unique, Duplicate or Null	Duplicate
Null	Null	Null	Unique	Unique	Unique, Duplicate or Null	New Asset
Unique	Duplicate	Duplicate	Duplicate	Duplicate	Unique, Duplicate or Null	New Asset
Duplicate	Unique	Duplicate	Duplicate	Duplicate	Unique, Duplicate or Null	New Asset
Duplicate	Duplicate	Unique	Duplicate	Duplicate	Unique, Duplicate or Null	New Asset
Duplicate	Duplicate	Duplicate	Unique	Duplicate	Unique, Duplicate or Null	Duplicate
Duplicate	Duplicate	Duplicate	Duplicate	Unique	Unique, Duplicate or Null	Duplicate
Duplicate	Duplicate	Duplicate	Unique	Unique	Unique, Duplicate or Null	Duplicate
Duplicate	Duplicate	Duplicate	Duplicate	Duplicate	Unique, Duplicate or Null	Duplicate

Another behavior of CORA, which is represented below, shows other ways duplicates can occur.

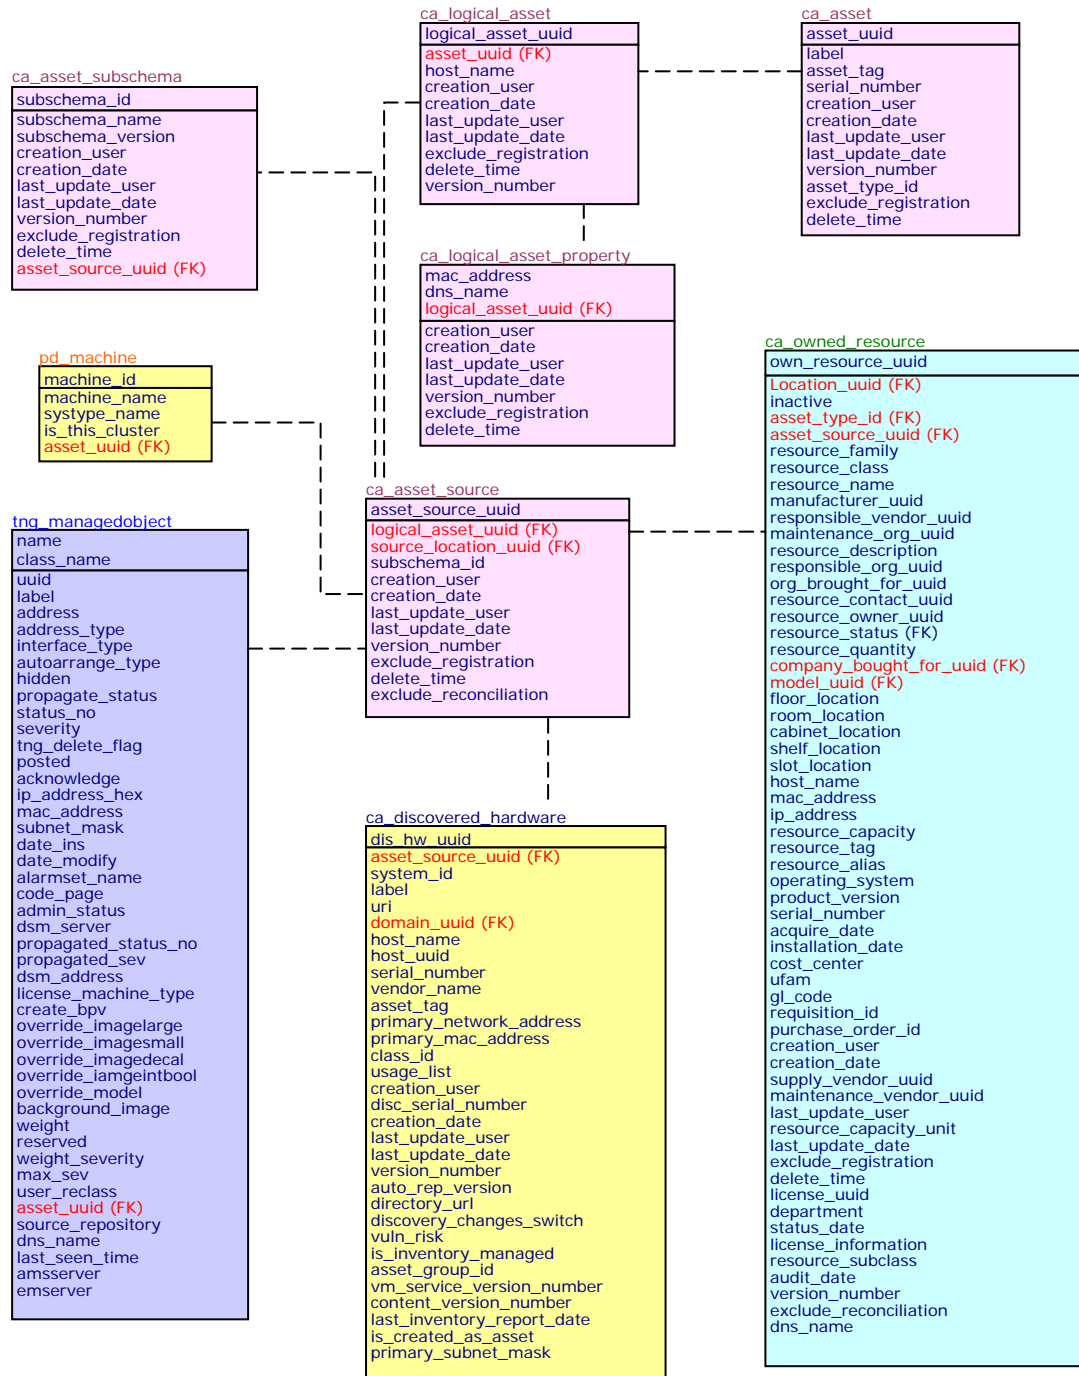| Serial Number	Asset Tag	Host Name	DNS Name	MAC Address	Asset Label	Results
Null	ABC	Null	Null	Null	Unique or Null	New Asset
123	ABC	Null	Null	Null	Unique, Duplicate or Null	Duplicate of ABC
789	XYZ	Null	Null	Null	Unique or Null	New Asset
Null	XYZ	Null	Null	Null	Unique, Duplicate or Null	Duplicate of XYZ

Discovered vs. Owned Assets

Attributes for each asset are divided into 'Discovered'and 'Owned' in order to facilitate reconciliation and verification capabilities. The primary tables used to identify data sources are:
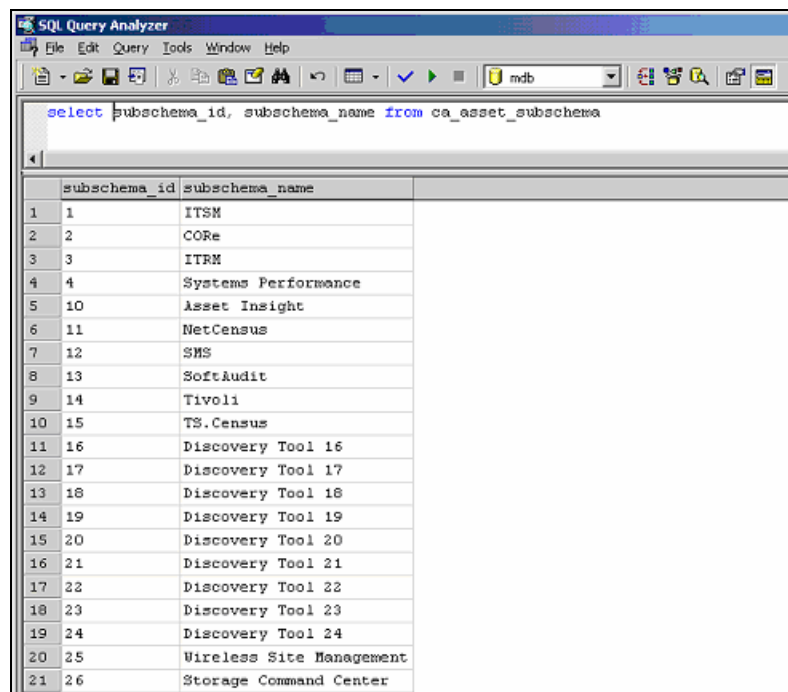
- CA_DISCOVERED_HARDWARE

- TNG_MANAGEDOBJECT

- PD_MACHINE

- CA_OWNED_RESOURCE

To understand how these tables relate to one another, consider the following graphic.

ca_logical_asset
- logical_asset_uuid
- asset_uuid (FK)
- host_name
- creation_user
- creation_date
- last_update_user
- last_update_date
- exclude_registration
- delete_time
- version_number

ca_asset
- asset_uuid
- label
- asset_tag
- serial_number
- creation_user
- creation_date
- last_update_user
- last_update_date
- version_number
- asset_type_id
- exclude_registration
- delete_time

ca_asset_subschema
- subschema_id
- subschema_name
- subschema_version
- creation_user
- creation_date
- last_update_user
- last_update_date
- version_number
- exclude_registration
- delete_time
- asset_source_uuid (FK)

ca_logical_asset_property
- mac_address
- dns_name
- logical_asset_uuid (FK)
- creation_user
- creation_date
- last_update_user
- last_update_date
- version_number
- exclude_registration
- delete_time

pd_machine
- machine_id
- machine_name
- systype_name
- is_this_cluster
- asset_uuid (FK)

ca_owned_resource
- own_resource_uuid
- Location_uuid (FK)
- inactive
- asset_type_id (FK)
- asset_source_uuid (FK)
- resource_family
- resource_class
- resource_name
- manufacturer_uuid
- responsible_vendor_uuid
- maintenance_org_uuid
- resource_description
- responsible_org_uuid
- org_brought_for_uuid
- resource_contact_uuid
- resource_owner_uuid
- resource_status (FK)
- resource_quantity
- company_bought_for_uuid (FK)
- model_uuid (FK)
- floor_location
- room_location
- cabinet_location
- shelf_location
- slot_location
- host_name
- mac_address
- ip_address
- resource_capacity
- resource_tag
- resource_alias
- operating_system
- product_version
- serial_number
- acquire_date
- installation_date
- cost_center
- ufam
- gl_code
- requisition_id
- purchase_order_id
- creation_user
- creation_date
- supply_vendor_uuid
- maintenance_vendor_uuid
- last_update_user
- resource_capacity_unit
- last_update_date
- exclude_registration
- delete_time
- license_uuid
- department
- status_date
- license_information
- resource_subclass
- audit_date
- version_number
- exclude_reconciliation
- dns_name

ca_asset_source
- asset_source_uuid
- logical_asset_uuid (FK)
- source_location_uuid (FK)
- subschema_id
- creation_user
- creation_date
- last_update_user
- last_update_date
- version_number
- exclude_registration
- delete_time
- exclude_reconciliation

tng_managedobject
- name
- class_name
- uuid
- label
- address
- address_type
- interface_type
- autoarrange_type
- hidden
- propagate_status
- status_no
- severity
- tng_delete_flag
- posted
- acknowledge
- ip_address_hex
- mac_address
- subnet_mask
- date_ins
- date_modify
- alarmset_name
- code_page
- admin_status
- dsm_server
- propagated_status_no
- propagated_sev
- dsm_address
- license_machine_type
- create_bpv
- override_imagelarge
- override_imagesmall
- override_imagedecal
- override_iamgeintbool
- override_model
- background_image
- weight
- reserved
- weight_severity
- max_sev
- user_reclass
- asset_uuid (FK)
- source_repository
- dns_name
- last_seen_time
- amsserver
- emserver

ca_discovered_hardware
- dis_hw_uuid
- asset_source_uuid (FK)
- system_id
- label
- uri
- domain_uuid (FK)
- host_name
- host_uuid
- serial_number
- vendor_name
- asset_tag
- primary_network_address
- primary_mac_address
- class_id
- usage_list
- creation_user
- disc_serial_number
- creation_date
- last_update_user
- last_update_date
- version_number
- auto_rep_version
- directory_url
- discovery_changes_switch
- vuln_risk
- is_inventory_managed
- asset_group_id
- vm_service_version_number
- content_version_number
- last_inventory_report_date
- is_created_as_asset
- primary_subnet_mask

The ca_asset_source table contains the subschema_id column which identifies the origin of the asset. The subschema_id values are maintained in ca_asset_subschema as shown with the following query:

The subschema_name column contains shortcut descriptions that refer to CA products or product components. The ones of immediate interest and their associated primary tables are:

■ 'ITSM' (IT Service Management) objects, which include 'Owned'sources such as CA Unicenter APM, CA Unicenter Service Desk, and CA CMDB (ca_owned_resource), have a subschema_id of '1.'

■ 'CORE'(CA Unicenter NSM WorldView Repository – Common Object Repository) or CA Unicenter NSM (tng_managedobject) objects have a subschema_id of '2.'

■ 'ITRM' (IT Resource Management) or CA Unicenter Asset Management objects (ca_discovered_hardware) have the subschema_id of '3.'

■ Systems Performance Management objects (pd_machine) have a subschema_id of '4.'

If an asset is registered in the MDB by different products, CORA only registers that asset once, and then links the information from the different data sources. As a result the ca_asset table has a single unique entry for each asset.

The following screenshots provide a walkthrough of the queries executed after a sample asset is registered by CA Unicenter APM, CA Unicenter Service Desk, CA Unicenter NSM, and CA Unicenter Asset Management. Note that the order in which the products register the asset is not relevant to the process.

First, from Machine name into ca_asset:


```sql
select * from ca_asset where label like '%Server1%'
```

	asset_uuid	label	asset_tag	serial_number
1	0xE7FC0779EE22424389C28A30FD23D607	Server1-topgun	NULL	56dd501

Then, from ca_asset into ca_logical_asset (using asset_uuid):


```sql
select * from ca_logical_asset where asset_uuid in (
select asset_uuid from ca_asset where label like '%Server1%')
```

	logical_asset_uuid	asset_uuid	host_name
1	0xECE60606A217C74D9E8D7C0D836901D8	0xE7FC0779EE22424389C28A30FD23D607	Server1-topgun

The ca_logical_asset_property shows the logical instances of the same asset. For instance, if the same asset is registered by CORA with a different DNS and/or MAC address but the same Host name, CORA recognizes it is the same asset and stores two logical instances in this table:


```sql
select * from ca_logical_asset_property where logical_asset_uuid in(
select logical_asset_uuid from ca_logical_asset where asset_uuid in (
select asset_uuid from ca_asset where label like '%Server1%'))
w
```

	dns_name	mac_address	logical_asset_uuid	creation_user	creation_date
1	Server1-TOPGUN	000C295ACD6A	0xECE60606A217C74D9E8D7C0D836901D8	NULL	1160141977
2	Server1-topgun.ca.com	000C295ACD6A	0xECE60606A217C74D9E8D7C0D836901D8	NULL	1160141840

Note: In this example the DNS name input by CA Unicenter APM (or CA Unicenter Service Desk, or CA CMDB) did not use the fully qualified name as it was discovered by CA Unicenter NSM. It is only an example.

Then, from ca_logical_asset to **ca_asset_source** (using logical_asset_uuid):


```sql
select * from ca_asset_source where logical_asset_uuid in(
select logical_asset_uuid from ca_logical_asset where asset_uuid in (
select asset_uuid from ca_asset where label like '%Server1%'))
```

	asset_source_uuid	logical_asset_uuid	source_location_uuid	subschema_id
1	0xD93E456B08FE604C9D8461E13E94DC55	0xECE60606A217C74D9E8D7C0D836901D8	NULL	1
2	0xED4DABA4834CDD4FBAB89D8BB31FDE01	0xECE60606A217C74D9E8D7C0D836901D8	NULL	2
3	0xEE19F52014604C4D8ED5C211C10CFCCA	0xECE60606A217C74D9E8D7C0D836901D8	NULL	3

Here you can see the different data sources from the subschema_id value.

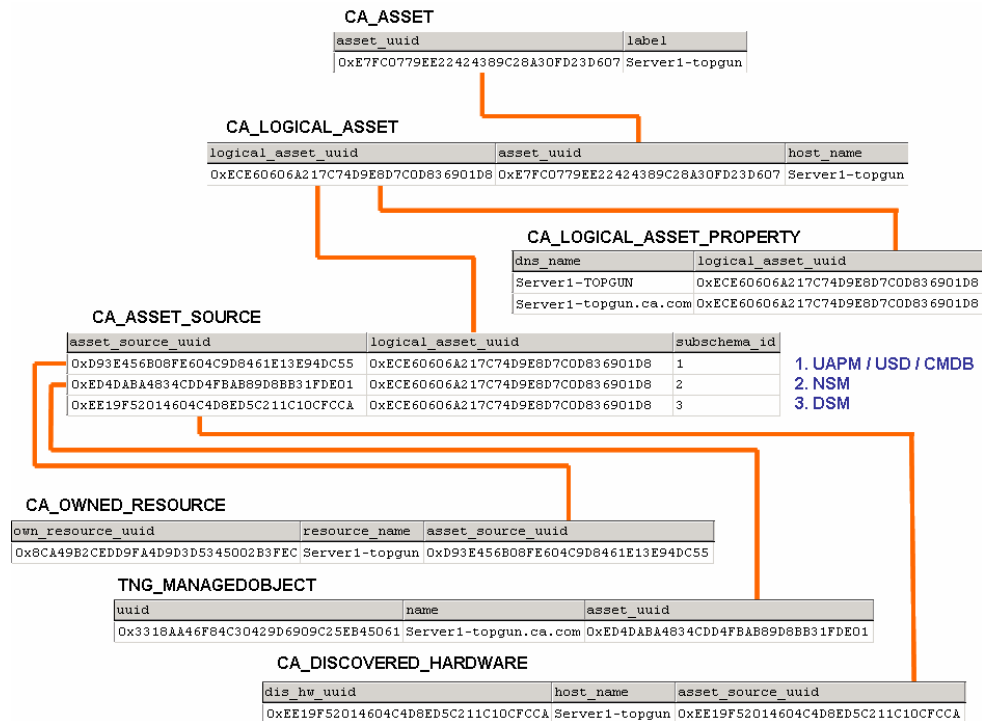Then, from ca_asset_source into CA Unicenter Asset Management ca_discovered_hardware (using asset_source_uuid):

Then, from ca_asset_source into CA Unicenter APM/CA Unicenter Service Desk **ca_owned_resource** (using asset_source_uuid):

Then, from ca_asset_source into CA Unicenter NSM **tng_manageobject** (using asset_source_uuid):

These links can be summarized in the following graphic:

CA_ASSET

asset_uuid	label
0xE7FC0779EE22424389C28A30FD23D607	Server1-topgun

CA_LOGICAL_ASSET

logical_asset_uuid	asset_uuid	host_name
0xECE60606A217C74D9E8D7C0D836901D8	0xE7FC0779EE22424389C28A30FD23D607	Server1-topgun

CA_LOGICAL_ASSET_PROPERTY

dns_name	logical_asset_uuid
Server1-TOPGUN	0xECE60606A217C74D9E8D7C0D836901D8
Server1-topgun.ca.com	0xECE60606A217C74D9E8D7C0D836901D8

CA_ASSET_SOURCE

asset_source_uuid	logical_asset_uuid	subschema_id
0xD93E456B08FE604C9D8461E13E94DC55	0xECE60606A217C74D9E8D7C0D836901D8	1
0xED4DABA4834CDD4FBAB89D8BB31FDE01	0xECE60606A217C74D9E8D7C0D836901D8	2
0xEE19F52014604C4D8ED5C211C10CFCCA	0xECE60606A217C74D9E8D7C0D836901D8	3

1. UAPM / USD / CMDB
2. NSM
3. DSM

CA_OWNED_RESOURCE

own_resource_uuid	resource_name	asset_source_uuid
0x8CA49B2CEDD9FA4D9D3D5345002B3FEC	Server1-topgun	0xD93E456B08FE604C9D8461E13E94DC55

TNG_MANAGEDOBJECT

uuid	name	asset_uuid
0x3318AA46F84C30429D6909C25EB45061	Server1-topgun.ca.com	0xED4DABA4834CDD4FBAB89D8BB31FDE01

CA_DISCOVERED_HARDWARE

dis_hw_uuid	host_name	asset_source_uuid
0xEE19F52014604C4D8ED5C211C10CFCCA	Server1-topgun	0xEE19F52014604C4D8ED5C211C10CFCCA

Notes about Asset Classes

Although 'class' is a mandatory field to for asset registration, it is not a field that is used by CORA. The concept of 'class,' in fact, is interpreted differently by different products. For example:

■ In CA Unicenter Service Desk, the concept of a 'family' is used to identify the highest level of definition for a CI. Each family can consist of one or more 'classes' to allow for a more granular categorization of CIs. Further, each family has an extension table that defines the attributes that are visible in the CI Detail page. When CA CMDB is implemented, it includes over 50 families and over 140 classes that are each stored in the MDB and shared between CA Unicenter Service Desk and CA Unicenter APM.

■ When the MDB used by CA CMDB is shared with CA Unicenter APM, those CMDB families are shared and are known to CA Unicenter APM as 'asset types' for 'models' and 'assets.' In other words, for CA Unicenter APM:

> CMDB Families = CA Unicenter APM Asset Types

> CMDB Classes = CA Unicenter APM Classes

In CA Unicenter APM, the Asset Type is a family_id field for ca_model_def and a resource_family for ca_owned_resource.

■ CA Unicenter Asset Management, on the other hand, does not use families and classes to register discovered assets. However, if CA Unicenter Service Desk is also installed and integrated with CA Unicenter Asset Management, CA Unicenter Service Desk

reflects the registration of a discovered asset as 'owned,' it uses the default family of 'Hardware' and class 'Discovered Hardware.'

■ CA CMDB content creates new families and classes. However, these classes are not the same classes that are used by CA Unicenter NSM to classify discovered objects. In fact, only a small number of CA Unicenter NSM classes match CA CMDB classes. However, procedures are provided in the *CA CMDB Administrator Guide* for mapping CA Unicenter Service Desk/CA CMDB classes to CA Unicenter NSM classes.

Note: As multiple CA Unicenter Service Desk/CA CMDB classes can be mapped to the same CA Unicenter NSM Class, the pdm_nsmimp cannot use the CA Unicenter NSM class to determine which Class to use when creating the asset. It does not know how to pick the correct class if multiple classes are mapped.

Useful Links and References

The full schema for the MDB is viewable through the Implementation Best Practices page (formerly the 'r11 Implementation CD') which is available on http://ca.com/support.

Common Asset Viewer

The Common Asset Viewer (CAV) (formerly known as AMS or the Asset Maintenance System) is a collection of browser based view-only screens available to various CA applications so that they can view details on any asset in the MDB. The CAV provides a common interface which can be used for viewing owned and discovered asset information.

In r11, the CAV is embedded with CA Unicenter Service Desk, CA CMDB, CA Unicenter APM, and CA Unicenter Asset Management applications. In CA Unicenter Service Desk and CA CMDB, when looking at a CI, CAV provides a common interface through which the consolidated asset details relating to the CI can be viewed. It also enables navigation from the asset data to other CA asset-related applications and allows users to see data that is stored about the CI in these other applications.

CAV contains three tabs: one for displaying owned asset information, one for displaying discovered asset information, and one for displaying network asset information. The asset

data contained in these tabs are typically read from and maintained by the following CA applications:

Common Asset Viewer Tab	Typical Source and Maintenance of Data
Owned Resources	CA Unicenter APM; CA Unicenter Service Desk; CA CMDB
Discovered Desktop/Server Resources	CA Desktop & Server Management (in particular, CA Unicenter Asset Management)
Discovered Network Resources	Unicenter Network and Server Management

Technical/Best Practice Information

■ Launching CAV

> CAV can be configured to display any combination of the three CAV tabs (Owned Asset, Discovered Asset, and Network Asset) by specifying the correct parameter when invoking its URL.

> CAV is accessed using a URL. Configuration parameters are available on the URL that allow users to configure how CAV appears (whether a pop-up window, what type of data is displayed, and so on).

> Using the URL, CAV also can be configured to provide links to CA Unicenter Service Desk, CA Unicenter APM, or CA Unicenter Asset Management applications. For example, you can begin in CA Unicenter Service Desk and launch into CAV to view details on an asset. This is accomplished by forwarding URL information from one application to the other.

■ Architecture

> CAV is a J2EE Struts/JSP application and connects to the MDB using JDBC.

> Each application that embeds CAV installs it into its own application directory. The CAV directory is *not* shared by different applications. However, as long as all instances of the CAV are pointing to the same MDB, they all display the same data.

> CAV uses a silent installer which is invoked by the parent application.

  Note: CAV is not able to support applications residing in multiple MDBs.

■ Implementation and Configuration

> CAV configuration information is stored in the AMS.properties file located in the CAV installation directory. This file can be edited manually but in most cases should be edited by using the AMSConfig.java program included with CAV. Instructions for using this program are in the top of the AMS.properties file itself.

> Error logging is controlled using the log4j.properties file located in the CAV installation directory. When experiencing problems with CAV, the first step to

take is to turn on enhanced logging by changing the 'rootcategory' property specified in this file to DEBUG.

> When experiencing problems with CAV, make sure to note the CAV version number. This can be found in the version.rel file located in the CAV installation directory.

Chapter 10: Roaming Between Scalability Servers

How to Roam

This chapter describes what happens when a computer roams between two scalability servers in the *same* CA Unicenter Desktop & Server Management (CA DSM) domain. When a computer registers with a scalability server in a different CA DSM domain we refer to this operation as a move. The move operation is not addressed here.

We will cover three areas of CA DSM functionality: Asset Management, Software Delivery, and Common.

There are several ways to initiate a roam. Three common methods are described below.

Method 1

The following command, executed on the local machine, will cause the DSM Agent running on the computer to register with a new CA DSM scalability server and hence roam to it:

```
caf setserveraddress newserveraddress
```

This command can, of course, be executed manually but can also be incorporated into a number of different automated roaming implementations. Known solutions include execution through login scripts and CA DSM DMScript. Of particular interest is a Field Developed Utility (FDU), available through CA Services, known as Location Awareness. The Location Awareness solution can direct a computer to an appropriate scalability server based upon an IP range.

Please note that FDUs are not supported by CA Support. To have this FDU supported, a support agreement must be created through the CA Services organization. Otherwise, the support of this utility and any questions or errors related to its use will be the responsibility of the customer organization. Please contact your CA Services representative for more information.

Method 2

The scalability server address can also be changed using CA DSM configuration policy. Known solutions exploit the possible association of CA DSM configuration policy with dynamic (query-based) groups. The scalability server name or address is identified and set within a configuration policy, and the configuration policy is then associated with a dynamic group. Now whenever a computer becomes a member of the group, the configuration policy settings are automatically distributed to the computer—hence changing the scalability server to which the agent reports. This method must be carefully evaluated because it has two disadvantages. The first is that the computer must have access to the old scalability

server (in order to receive the updated config policy), and the second is that the roam operation generally occurs later in the process.

Method 3

Some customers use the same Fully Qualified Domain Name (FQDN) for all scalability servers. DNSs (Domain Name Servers) are then used to associate different IP addresses with the same scalability server FQDN dependent on IP subnet.

Roaming Explained

There are three CA DSM tiers involved in roaming. Each tier consists of several plug-ins, so the information flow is complex. The tiers and their plug-ins include:

Agent:

Common Agent, Common Configuration Agent Plug-in, Asset Management Agent Plug-in, Software Delivery Agent Plug-in

Scalability Server:

Common Scalability Server, Common Configuration Scalability Server Plug-in, Asset Management Scalability Server Plug-in, Software Delivery Scalability Server Plug-in

Domain Manager

Common Manager, Engines, Common Configuration Manager Plug-in, Asset Management Manager Plug-in, Software Delivery Manager Plug-in

Common Information Flow

The Common Agent registers with the new scalability server. It compares the network address (which could be FQDN, hostname, or IP address) of the new scalability server with the network address of the previous scalability server. If it detects that a roam has occurred it will run additional plug-ins, specified in the 'afterroamplugin' configuration value of each solution (see itrm/agent/solutions/XXX/afterroamplugin and itrm/agent/solutions/XXX/afterroampluginparams in comstore).

The network address (which could be FQDN, hostname, or IP address) of the previous scalability server and the time of the roam is reported to the CA DSM domain manager. This is done so that the domain manager can ensure the most recent registration is used. This is important if the domain manager is processing registrations for the same computer from multiple scalability servers.

The registration triggers the Common Configuration Agent Plug-in to send a "hello" message to the Common Configuration Scalability Server plug-in (based on the scalability server address in comstore).

The Common Scalability Server plug-in receives the registration from the Common Agent and stores this in the inbox. The information will wait here until collected by the CA DSM Engine.

The Common Configuration Scalability Server plug-in forwards the hello message from the Common Configuration Agent Plug-in to the Common Configuration Manager.

The engine collects the computer's registration information from the scalability server. It will detect that the computer has roamed from the old scalability server to the new scalability server (in the MDB, ca_agent.server_uuid is different from the new scalability server's uuid). The engine will update the Common Agent information in the MDB, including pointing it to the new scalability server. During this process the engine submits a notification, NOID_COMMON_COMPUTER_UPDATED_ROAMING, using CA DSM common notification informing interested parties that the computer has roamed.

The hello message is received by the Common Configuration Manager and analyzed. Part of this message is the scalability server hostname to which the agent now reports. If the hostname is different from the scalability server ID currently stored in the common configuration tables (within the MDB), these will be updated. Any communication from manager to agent will from now on be routed through the new scalability server.

Asset Management Information Flow

When the Asset Management Agent Plug-in runs it will ask the Asset Management Scalability Server Plug-in running on the new scalability server if a configuration for itself (using the CA DSM HOSTUUID) is available. If it is not, it will wait for three minutes and then retry. If this fails, it will wait until its next scheduled run. Until the engine has processed Common Agent information on the new scalability server and pushed an asset management configuration file down to the Asset Management Scalability Server Plug-in, the Asset Management Agent Plug-in will not be able to deliver information to the new scalability server.

After the engine has collected information from the Common Scalability Server plug-in (see the Common Information Flow above) the engine will validate the Asset Management Scalability Server plug-in on the new scalability server. Here it will find that the agent's configuration file is missing and the engine will therefore push information about the agent, including configured collect tasks and asset management jobs, down to the Asset Management Scalability Server Plug-in. A configuration file for the agent is hereby created on the new scalability server. This will enable the Asset Management Agent Plug-in to report to the Asset Management Scalability Server Plug-in running on the new scalability server. It is important to note that the engine does not push down status information (collected on the previous scalability server) for Asset Management Modules/Jobs to the new scalability server.

When the Asset Management Agent Plug-in eventually retrieves its configuration, it will perform scheduled collect tasks and asset management jobs. Again, it is important to note that all configured asset management collect tasks and asset management jobs are considered waiting for the agent, because there is no status available on the new scalability server.

The hardware inventory delivered to the new scalability server as a result of the scheduled collect tasks will, however, be delta inventory. The agent will compare the newly collected inventory with the inventory previously collected and submitted to the old scalability server, and will only deliver the deltas.

The engine will collect the information delivered by the Asset Management Agent Plug-in from the new scalability server and store this in the database.

Software Delivery Information Flow

When the Software Delivery Agent Plug-in connects to the new scalability server, the Software Delivery Scalability Server Plug-in will see the agent as a new computer. The data retrieved from the agent will be cached on disk until the Software Delivery Manager notifies the scalability server that the agent is known.

The Installation Manager, a sub-component of the Software Delivery Manager, is responsible for handling all manager processing related to a roam. The processing starts when the Installation Manager receives a notification, NOID_COMMON_COMPUTER_UPDATED_ROAMING, from the engine. When receiving this notification, the Installation Manager will lock the computer with the roam lock (as seen in the DSM Explorer). After this, the Installation Manager will send a sync message of type VXML_SYNC_REASON_ROAM to the old Software Delivery scalability server. The old scalability server will then send a reply back. This reply is received into the Installation Manager's main message queue.

When the Installation Manager receives the reply, it will start the process of determining whether there are any active jobs on the old scalability server that need to be moved to the new scalability server. The first thing that happens in this phase is that the roam lock is removed. Note that if there are a lot of pending messages in the Installation Manager's queue, then the lock will remain on the computer until the reply from the old scalability server makes it to the front of the queue and is processed. Then the actual move of any jobs takes place. Finally, the Installation Manager will send a sync message to the new scalability server, including the HOSTUUID of the agent object—which is what the Software Delivery Scalability Server plug-in requires.

The Software Delivery Scalability Server Plug-in now receives a notification from the Software Delivery Manager that the computer is known (and unlocked). All the data cached for the agent is now sent up to the Software Delivery Manager (Installation Manager). A final response message is then sent to the Installation Manager.

Example Scenarios

From Office to Home Using VPN

The Scenario

A user's laptop is connected directly to the corporate network. The DSM Agent reports to a named scalability server. The IP address of the scalability server is resolved through a DNS server.

The user powers down the laptop or puts it into standby mode.

The user goes home.

The user powers up the laptop or brings it out of standby mode and connects from home through VPN to the corporate network.

In this case, we assume the Common Application Framework (CAF) and all agents have started before the VPN connection becomes active. This will typically be the case since either CAF will not have been stopped or because most VPN solutions require end users to authenticate.

The computers connected to the VPN are served by a different DNS server which has the *name* of the end point's scalability server registered, but the network address is different. In this way the DSM Agent is invisibly redirected to a different scalability server when connected through VPN.

How CA DSM Behaves

In this situation, when the VPN connection is established, the DSM Agent detects a change in IP address.

When an IP address change is detected, the DSM Agent registers with its scalability server, which in this case is a different scalability server from the one it previously registered with. However, because the name has not changed, the agent does not realize this fact.

The registration triggers the Common Configuration Agent Plug-in to send a Common Configuration hello message to the scalability server.

The scalability server receives the registration from the agent and stores this in an inbox. The information will wait here until collected by the engine.

The scalability server also forwards the hello message from the Common Configuration Agent Plug-in to the DSM Manager.

The engine collects the computer's registration information from the scalability server and detects that the computer has roamed from one scalability server to another. The engine updates the MDB appropriately and submits a general notification message informing

interested parties that the computer has roamed. It also pushes down an asset management configuration for the agent to the new scalability server.

The Manager receives the general notification from the engine and locks Software Delivery (SD) management of the computer. The manager then sends an SD sync message to the old scalability server.

When the manager receives the SD sync reply, it will start the process of determining whether there are any active SD jobs on the old scalability server that need to be moved to the new scalability server. The SD lock is removed, the actual move of any SD jobs takes place, and an SD sync message is sent to the new scalability server.

The Common Configuration hello message is received by the manager and analyzed. Part of this message is the scalability server hostname to which the agent now reports. If the hostname is different from the scalability server ID currently stored in the common configuration tables (within the MDB), these will be updated. Any communication from manager to agent will from now on be routed through the new scalability server.

The Asset Management Agent Plug-in runs on a regular basis, typically once a day. Whether it will run during the VPN connected time will obviously depend on how long the laptop is kept connected and when the last run took place.

If the Asset Management Agent Plug-in does indeed run while connected to the VPN (and hence a new scalability server), it will ask the scalability server if a configuration for itself is available. If it is not, it will wait for three minutes and then retry. If this fails, it will wait until its next scheduled run.

Unless the engine is under heavy load, the DSM Agent will have registered and an asset management configuration will have been made available on the new scalability server within the three minute retry interval.

When the Asset Management Agent Plug-in eventually retrieves its configuration, it will perform scheduled collect tasks and asset management jobs. Collected delta inventory and job status information will be uploaded to the scalability server.

The engine will collect the information delivered by the Asset Management Agent Plug-in from the new scalability server and store this in the MDB.

When the Software Delivery Agent Plug-in connects to the new scalability server, the scalability server may believe the agent to be a new computer. This will always happen after an agent registration—for instance on an IP address change. In this case, the data retrieved from the agent will be cached on disk until the manager notifies the scalability server that the agent is known. Once received, all the Software Delivery data cached for the agent is sent up to the manager.

Chapter 11: Data Replication between CA ITCM Domains and the Enterprise

Introduction

In a multiple CA IT Client Management (CA ITCM) domain architecture, it is desirable to have most of the data contained in a single Management Database (MDB). That way you can search, report on, administer, and perform a variety of functions from a centralized console. Because of the amount of data and for performance reasons, it is not desirable to replicate all of the data to a single MDB.

With the advent of CA DSM r11, CA has taken the guesswork and the often confusing elements out of the selection of items for replication and has developed a standardized methodology for the replication.

With the creation and linkage of the domain to an enterprise, an engine task is automatically created to handle the replication of data between domain and enterprise databases. Replication from the domain database to the enterprise database is carried out by a replication job that runs through the domain manager's engine process.

When replication begins, the engine determines which information needs to be pushed from the domain to the enterprise and also which information needs to be pulled from the enterprise down to the domain.

Typically, host-specific information, such as inventory attributes, is replicated upward while configuration information, such as asset groups, is replicated downward.

With each domain manager, a default engine is installed, called the System Engine. When the domain manager is linked to an enterprise manager, that engine is configured to perform the domain-to-enterprise replication tasks.

Database Objects Replicated

The following table from the CA ITCM documentation lists the database objects that are replicated from enterprise to domain (down) and from domain to enterprise (up).

Object	Replication Direction
Discovered Computers	up
Discovered Users	up
Discovered Computer Users (relations between Computer and Users)	up
External Assets Definitions	down
External Assets	up
Computers General Inventory	up
External Assets inventory	up
Query Definitions	down
Group Definitions	down
Group Membership	down
Custom Made Software Definitions	down
Custom Defined Manufacturer	down
Computers Software Inventory (found based upon signature scan)	up
Asset Management Jobs	down
Asset Management Job Status	up
Asset Management Modules	down
Asset Management Modules Status	up
Asset Management Configuration File Definitions	down
Asset Management Configuration Files **Note**: These files are only replicated to the enterprise manager if the request to collect this information has been defined on the enterprise manager. If the request has been defined on the domain manager, the data is not replicated upwards.	up
Asset Management Template Definitions	down
Asset Management Policy Definitions	down

Replication Functionality

The replication process can be configured during the installation procedure or after domain and/or enterprise level installations have been completed. If you have already installed an enterprise server, then during the installation of a domain you will be asked if you want to link the domain to the enterprise. That is all that is required.

Quite often, however, domains are created and an enterprise level is created after the fact. Even if it is planned, best practice dictates that you install at least the first domain and configure it before you link it to an enterprise. The descriptions below demonstrate how that process is completed.

To perform the linkage, go to the DSM Explorer on the enterprise server. As pictured below, go to Control Panel/Domains/All Domains and from the right side click New Domain.

The wizard will open. Type in the name or IP address of the domain you are trying to link and click next.

You should receive the next screen saying the operation was successful. If you do not, the first thing to check is if your DNS is set up properly and/or you are unable to reach the domain server through IP.

Click Finish and the process is started.

To confirm the linkage has been successful, you can go to the DSM Explorer and access the CA DSM domain you just linked. Then go to Control Panel\Engines\Engine Tasks and verify that a Replication Task has been created and assigned to the system engine. Having the Replication Task show up in the Domain console may take a couple of minutes depending on your situation.

To confirm the linkage has been successful you can go to the DSM Explorer, and then go to Control Panel\Engines\Engine Tasks and verify that a Replication Task has been created and assigned to the system engine. Having the Replication Task show up in the Domain console may take a couple of minutes depending on your situation

Please note that the actual replication is not instantaneous. The two MDBs are exchanging information using the CA ITCM Engines over whatever network setup you have in place, so this could take several hours to complete. However you should begin to see computers with inventory information showing up in the DSM Explorer enterprise console within a short time.

Once the replication is complete, presuming you have the necessary rights to the data, you can access your entire enterprise through the DSM Explorer enterprise console.

From this screen on the enterprise server (see the lower right hand corner) you can open the highlighted asset information directly from the domain server database. This is helpful if you are searching for information about an asset that may not have replicated. The DSM Explorer domain console will open automatically and display the information about that asset.

Unlink a Domain

Occasionally a replicated MDB may become corrupt, or for administrative reasons you may want to unlink a domain. One methodology for re-establishing clean data is to unlink the domain and then re-link after the unlinking is complete. Unlinking will remove all of the data from both the enterprise MDB and the domain MDB—so just as with replication, remember that unlinking is not instantaneous. Depending on the amount of data this may take several hours to complete.

When you want to unlink a domain you should shut down the replication engine on the domain manager. The Replication Task is assigned automatically to the System Engine. Therefore you need to shut down that engine. From the DSM Explorer domain console, go to Control Panel\Engines\All Engines\System Engine and right-click and select Stop Engine as shown below. Then go to the DSM Explorer enterprise console.

From the DSM Explorer enterprise console, go to Control Panel\Domains\All Domains*domainname*. Right-click and select Remove from Enterprise.

This completes the unlink process.

Summary

Replication from domain(s) to an enterprise MDB is very straightforward and can be redone at a later time if necessary. This functionality gives you a centralized administrative console, although you can still do administrative tasks at the domain level.

Chapter 12: CA Common Services

Overview

The underlying infrastructure for many CA solutions is CA Common Services™ (CCS). Composed of a comprehensive set of management, infrastructure, and visualization services, CCS is the 'foundation and glue' that automatically integrates a rich array of CA products into a collective enterprise solution.

As the core component of CA Unicenter Network and Systems Management, Common Services provides a single point of integration between CA solutions and user-written applications. Platform support, event management, communications, and visualization are just a few of the many core services available in CCS. An impressive host of functions continually improves the efficiency of your business while reducing resource utilization.

Understanding CCS

When installing from the CA IT Client Management (CA ITCM) media, CCS is installed as one package. This eliminates the complexity of component selection. See the CCS Installation section below for specific instructions.

The primary components in CCS are:

■ Enterprise Management

■ Enterprise Discovery

■ WorldView

■ 2D Map

■ Report Explorer

Enterprise Management

Enterprise managers are a collection of management components that employ a single, easy-to-use graphical user interface—the Event Console—to monitor and administer different events, including SNMP traps, application events, and system events.

Included in the Enterprise Management functionality is the ability to take action on any message delivered to the console. Through an easy-to-use interface, actions can be developed to go all the way from simply informing a technician of a changed condition, to

running complex internal and external utilities to automatically correct the changed condition.

As an example, you may want to ensure that only the standard desktop programs are allowed to run on all of the desktops in your environment. This may be for productivity reasons and it may also be for software compliance reasons. In conjunction with the CA ITCM Policy Violation capabilities, it is possible to send an event to the Unicenter Console after the CA ITCM Asset Management plug-in has run and detected a violation. That message can be used to send an alert in a wide variety of formats to inform someone of the violation. In a more advanced implementation, that message can be used to automatically trigger an un-installation job through the CA Unicenter Software Delivery component of CA ITCM.

Enterprise Discovery

The comprehensive Enterprise Discovery (IP Discovery) feature is used to easily identify your entire network. Depending on your requirements and the size of your enterprise, you can adjust Enterprise Discovery to run on a single network segment or over the entire enterprise. You can even schedule Enterprise Discovery to run at a certain time on a specific day.

As Enterprise Discovery finds computers, servers, routers, and so forth, it creates objects that represent these entities and stores the objects in the MDB. The Discovery Monitor keeps you apprised of the Discovery process as it searches your networks.

The discovery of Storage Area Network (SAN) devices is seamless with the rest of the managed devices.

WorldView

Whether your 'world' is several machines used for word processing in an office, or a global network of PCs, servers, routers, and sophisticated applications and business processes, you can explore it in the Real World Interface, WorldView. WorldView consists of two components, the 2D Map and the Association Browser. These components offer a complete and authentic view of your business.

2D Map

The 2D Map is a powerful user interface that acts as an infrastructure navigator and controller, letting you view any part of your business at any time, with the click of a button.

The 2D Map provides a graphical representation of your IT infrastructure and customized Business Process Views. It offers an overview of the structure of your network using icons to depict your resources. It is all right there on the screen— networks, sub-networks and segments, PCs, servers, routers—everything you need to see!

Report Explorer

Now that you have all this information, how can you use it to your greatest advantage? The Report Explorer lets you customize, view, and print reports so that the extensive amount of information available to you is presented in the most meaningful format for your particular requirements.

Additional Capabilities

Beyond the core components listed above, CCS contains a number of tools that allow for advanced customization to fit your organization's specific needs. They are:

- ObjectView

- Class Browser

- Dynamic Business Process Views

- Repository Import/Export

- SDK

ObjectView

The 2D Map provides a complete business overview of the entire organization. ObjectView lets you take a deeper look at each of the objects displayed on the maps. Each resource has Management Information Base (MIB) information—a collection of attributes of the device or the application—such as performance, status, and so on.

Accessed directly from the 2D Map or Object Browser, ObjectView displays minute details of the device or application, letting you know the exact status of that resource at any time of day or night. As is the case with other components, this is done from a single console.

Class Browser

The Class Browser displays all the properties of each class in an easy-to-understand and intuitive manner. The Class Wizard allows you to easily create or modify Class designations. The Class Specification lets you view or modify classes.

Dynamic Business Process Views

This facility lets you collect WorldView objects based on a query or set of queries. Based on your query, the system keeps the set of objects contained in the Dynamic Business Process View up-to-date based on WorldView notifications. If something changes in an object, the system checks whether the change fits the query. If so, it adds the object to the Dynamic Business Process View. If the change no longer fits the query, it is removed from the Dynamic Business Process View. Dynamic Business Process Views are similar to the billboard, letting you scope the query to objects at a level of your choosing in the topology.

To create Dynamic Business Process Views, first choose an object whose children you want to limit your view to in the 2D Map. Then, in Design mode, use the Tool Palette, drag the Dynamic BPV class (Classes, ManagedObject, BusinessView, DynamicBPV) to the 2D Map, and rename it to define the query.

To set up a query for this newly created Dynamic Business Process View, right-click on it and choose Viewers, Dynamic BPV Query View.

Repository Import/Export (TRIX)

TRIX, the repository import/export utility, is an invaluable tool for any size business. You can use TRIX to populate repositories with objects from other repositories. This means that you can populate a central corporate repository with objects from remote repositories. You always have a complete, current corporate repository, which ensures that you have the most reliable information about the status of your business.

Software Development Kit

The Software Development Kit (SDK) is not automatically installed as a component to CCS in version r11.x. When desired, however, the SDK can be acquired and will co-exist in this version of CA ITCM.

The SDK is a component that provides the APIs and utilities required to develop solutions that integrate with your CA product. It offers the opportunity to integrate custom applications with CA architecture, including WorldView, Agent Technology, and Enterprise Management. The SDK provides an extensive set of APIs that let you manipulate classes and objects in the MDB, create agents to instrument custom resources, and much more.

Tip: For more information about the Software Development Kit, see the *CA SDK Developer Guide*. Each CA solution has the runtime modules it requires embedded in the product.

Any third-party product that is enhanced with the SDK can communicate and integrate with any of the CA solutions containing the specific common services (such as Event Management) for which integration has been performed.

Two typical scenarios in which a vendor or third party can use CCS are:

■ Integrate an existing third-party product so that it can communicate with your CA solution.

■ OEM—Develop a third-party product in which the partner provides added value by embedding CA technology in the product. The embedded technology could be a complete CA solution (such as CA Unicenter Network and Systems Management) or a CCS runtime component (such as Event Management).

CCS Installation

CCS uses the same MDB as other r.11x products, however there are specific considerations as to where and how to install the components.

The first consideration is that the WorldView Manager must reside on the MDB server.

In a single domain architecture with the domain MDB on the same machine as the domain, you will simply include all the CCS components on that server. There is no need to install it separately, but you should use the Custom Install method and be sure that the box next to CCS is checked for inclusion in the installation.

In a multiple domain architecture with an enterprise server at the top tier, it is best practice to install only one instance of CCS at the enterprise level.

When installing with a remote MDB you must include the WorldView (WV) manager component on the MDB database server. Even though you will be running the CCS application components from the application server, the WV component needs the WV Manager to be installed on the MDB server for communication purposes. There are two ways to accomplish this:

■ If available in your environment, you may use CA Unicenter Network and Systems Management media to select and install only the WV Manager. This is the preferred method. By using this method, only that one component will be installed. You would accomplish this by first installing the MDB using the CA ITCM media. After that installation has completed, insert the CA Unicenter Network and Systems Management media and install only the WV component on the MDB server. Next, on the application server and using the CA ITCM media, select the Install CA Unicenter Desktop & Server Management (CA DSM) choice from the installation menu and make certain that the list of components to install includes CCS. During the installation, point to the remote MDB. This will install full CCS including Enterprise Management, WorldView, and Continuous Discovery on the application server.

■ If you do not have access to CA Unicenter Network and Systems Management media you can use the CA ITCM media. You will have to install full CCS on the MDB database server. Very likely, however, you will not want to run most of those components on that server. Therefore, immediately after the installation of full CCS and before the services are started on the MDB server, you will want to disable certain services including:

> CA – Continuous Discovery Agent

> CA – Continuous Discovery Manager

> CA DIA 1.2 DNA

> CA DIA 1.2 Knowledge Base

> CA-Unicenter

> CA-Unicenter (NR Server)

> CA-Unicenter (Remote)

> CA-Unicenter (Transport)

Using CCS Calendars with CA Unicenter Software Delivery

You can use CCS calendars in three ways when working with CA Unicenter Software Delivery.

DELIVERY CALENDAR

The Delivery Calendar is only applicable to DTS (Data Transport Service). It is the calendar you can define within WAC and also on the CA Unicenter Software Delivery job options tab. Calendars define only intervals and not points in time. The intervals defined within a calendar specify when it is acceptable to carry out an operation. So the delivery time you specify and the CCS calendar you define for a job work together. When the delivery time comes, CA Unicenter Software Delivery checks the calendar to see if it is OK to deliver the package. If not, it keeps checking at regular intervals until it is OK.

The use of this calendar needs to be fully understood because it is only used to control WHEN the DTS transfer of a package is ALLOWED to occur. There are three types of transfers:

- DTS transfer of the package from Manager to Scalability Server (Download method = NOS or NOS-less internal)

- DTS transfer of the package from Manager to Agent (DM = DTS)

- DTS transfer of the package from Scalability Server to Agent (DM = DTS)

When using this method in the first instance (Manger to Scalability Server) the Calendar will apply. If, however, the package is already staged on the scalability server and DTS is not selected, then this Calendar will not apply to any further transfers.

JOB EXECUTION CALENDAR

This calendar is the one that you attach to a computer or group within the DSM Explorer. This is applicable to all download methods and is defined for a computer, not a job. The scalability server uses this to determine when it's OK to activate the job after the activation time has expired. This requires a CCS EM agent/client to be installed on the scalability server since it's the scalability server that checks the calendar. This is not applicable when using DTS.

WorldView DTS Objects

You can associate a calendar with a DT computer or DT link within the Unicenter 2D Map. This is then applied to any DTS transfers and DTS will suspend transfers until an associated calendar says it's OK to begin the transfer.

Chapter 13: Discovery and Deployment

The Challenge of Bringing Servers, Desktops, and Laptop Computers under Management

There are various challenges in the management of computer devices, which is made increasingly more challenging by the proliferation of departmental and 'personal' firewalls. While these firewalls are sometimes deployed for very good reasons, they significantly hamper the ability to automatically classify and deploy the correct CA Unicenter Desktop & Server Management (CA DSM) Agent to that platform.

Regardless of the challenges, however, network discovery and classification are essential elements to properly managing your environment.

A clear example of this, which also has clear implications regarding security and compliance, is the discovery of unknown computers connected to your network. Given how easily these devices can be introduced to the network without any prior review or approval, the need for automated discovery becomes especially apparent.

Consequently tools must be used to monitor the network efficiently and bring these newly detected computers under management, and raise an alert when circumstances are detected that indicate this cannot be automatically completed. This reduces the risk of a non-standard machine disrupting the production environment and potentially disrupting business activities.

Infrastructure Discovery

Continuous Discovery uses manager/agent architecture to disperse collectors that can be deployed across the enterprise to reduce network traffic passing over Wide Area Networks (WANs). The discovery agents can be deployed to remote locations perhaps where scalability servers are located in order to improve the automatic detection rate for new undiscovered computers.

The install of a CA DSM domain manager with CA Common Services (CCS) will automatically install a Discovery Manager and Discovery Agent.

If you do not plan to use the CA DSM Auto Deployment capabilities then it is recommended that you do not run the CCS Discovery Services (CA-Continuous Discovery Agent and CA-Continuous Discovery Manager). To prevent these services from automatically starting they should be set to Manual or Disabled.

Discovery Strategy

The Continuous Discovery component employs a multi-tiered discovery architecture that can monitor an environment and automatically discover and classify new computers. The discovery of new computer entities uses three types of discovery mechanism:

- **Network traffic analysis** utilizes a passive discovery mechanism that listens for network-attached devices.

- **Network probing technology** is an active method that utilizes ping, Telnet, HTTP along with SNMP, Port Scan, and ARP (Address Resolution Protocol) cache discovery to identify devices.

- **DHCP listening technology** listens for DHCP (Dynamic Host Configuration Protocol) requests.

Once a device is located, the discovery infrastructure will consult information contained within its configuration rules and attempt further classification utilizing various techniques including SNMP, port scanning, and MAC address matching.

An example of a classification rule would be: Connect to port 80 (HTTP) and look for the string 'IIS.' If this is found, the device would then be recognized as a Windows device, due to the high probability that what has been detected on port 80 is an Internet Information Server (IIS) and is only available on Windows. The configuration rules supplied are prioritized, and classification automatically processes the list of rules to arrive at the best classification possible.

The methods and classification rules are supplied in an XML file that can be modified to incorporate new rules or the customization of the supplied rules to tailor them to your specific environment's needs.

A hierarchy of classes is used when classifying devices. This approach is used to address the common problem when a device is initially discovered and there is insufficient information available to ascertain the actual operating system version and other attributes.

Using a hierarchy of classes significantly improves the thoroughness of the classification. In the event that the device has no network profile and only the MAC address is available for classification, using this approach will still enable some measure of accurate classification as the system can identify the type of device based on MAC address. It can review the MAC values against industry-reserved MAC ranges, such as those for Dell_device, for example.

Parent Class	Child Class
Windows	WindowsNT, Windows9x, WindowsNT_Server, Windows2000, Windows2000_Server, WindowsXP, Windows_NetServerPocketPC
UNIX	AIX,Solaris, DG_UXLinux, HPUnix, NCRUnix, UnixWare, SCOUnix, Silicon, SiemenUX, FUJIUxp, Sequent_Server, OpenVMSICLUnix
Linux	RedHatLinux, SuSELinux, TurboLinux

Discovered devices, once classified, are added to the Management Database (MDB) through the Common Object Repository API (CORA). CORA, in turn, enables information made available from different CA products to be combined and reconciled to provide a single, well-defined object within the MDB that is referenced and populated by the discovering and consuming applications.
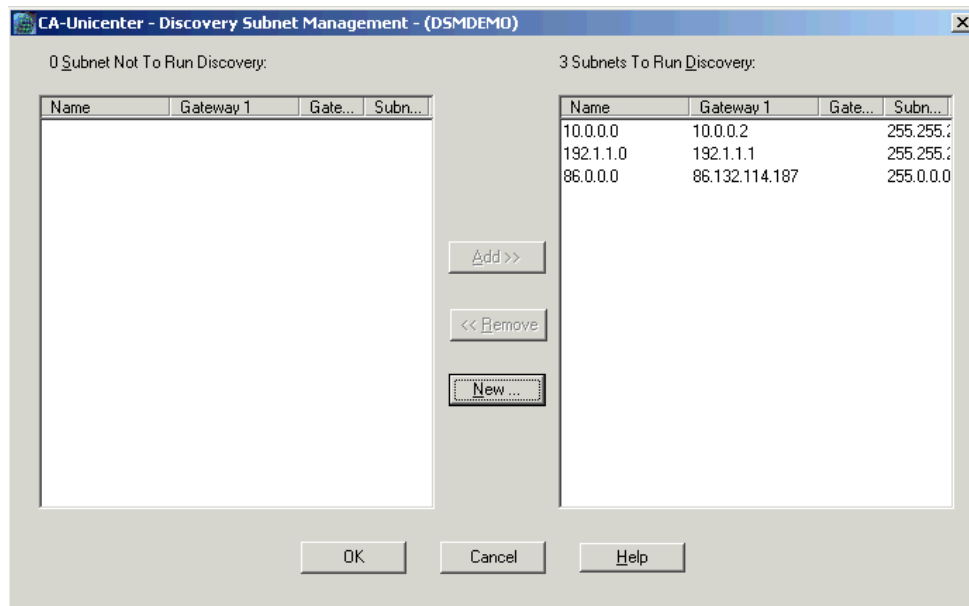
Discovery Configuration

Configuration of the discovery components is accomplished by updating the properties of the manager and agent objects within the CA Common Services WorldView 2D Map. This can be achieved by launching the 2D Map and navigating to the TCP/IP network object and drilling down through the subnet layer to the subnet segment containing the CA DSM domain manager. Drilling down from the domain manager you will find a Unispace object, and below the Unispace object will be two icons that represent the discovery manager and agent.

Select one of objects and right-click. Then select Open Details and a series of tabs will appear. These tabs are used to configure the relevant components. The example shown is of the Discovery Agent Runtime tab, which is used to define the subnets to be managed.

By default, the discovery agent will only look for computers on its own subnet. An alternate approach to the one above for adding subnets to be managed is to use the Classic Discovery Advanced subnet management as show in the screen shot below.

Classic Discovery is discussed in more detail later in this chapter.

A method of adding subnets to be managed is to use the discovery command line to discover a single system on each subnet that you wish to add. For example:

```
"dscvrbe –R DOMAIN_MANAGER_NAME -7 IPADDRESS_to_Discover -5 Yes –v 9"
```

Further information on Continuous Discovery can be found in Chapter 3 of the *CA Unicenter Network and Systems Management r11.x Administrator Guide.*

Note: The *CA Unicenter Network and Systems Management r11.x Administrator Guide* uses the MCC user interface when discussing configuration of the Continuous Discovery components. The MCC component is not installed by CA DSM, but all configuration steps can be done from the WorldView 2D Map.

Examples of Discovery Configurations

Different networks require different discovery mechanisms. There are many factors that can contribute to a network environment.

Examples of such factors contributing to the necessity of different configurations are shown in the following table:

Network Considerations	Limitations	Solution
DHCP servers	Location Configuration (different configuration options)	Configure routers to forward DHCP request to discovery agent
Switched networks	Switch locations Accessibility of network traffic	Connect agent to mirrored/spanned switch port
Firewalls	Blocked out protocols (ICMP, SNMP, UDP) Fixed ports with special permissions	Install discovery agent outside of firewall and configure a single port between the agent and manager through the firewall utilizing the CAM protocol
SNMP installations	Is it company policy to enable/disable SNMP for their computers Is SNMP traffic blocked	An extensible method of classification is carried out by the discovery agent allowing classification without the need for SNMP
Router configurations	Blocked traffic	Agents can be dispersed around the network to manage remote network segments

Different environments will have different configuration requirements. Accordingly, there is no single configuration approach that will work for every environment. By using the facilities inherent to Continuous Discovery, these components can be configured to address the particular requirements of even the most complicated environments.

The default settings included with the product were developed based on the experiences gained from supporting key client deployments for several years, and they represent the

combination of settings most commonly required. As shipped, these settings assume that they will be deployed in an environment like the following:

Switched networks with DHCP support with SNMP and Internet Control Message Protocol (ICMP) traffic enabled.

■   The manager has DHCP monitoring enabled.

■   The agent has the CTA (Common Traffic Analyzer) component enabled and will monitor the local network traffic by default.

■   The agent is configured to monitor the local subnet (the default). To expand the subnets that are monitored, you need to update the SubnetsToBeManaged property of the agent object with the subnets that are required for discovery. The discovery manager has an option called Workload Balancer (WLB), which by default is enabled. This functionality will ensure that when a device is discovered on a new subnet, Continuous Discovery will automatically manage the entire subnet. If this feature is not required, the Workload Balancer option should be disabled.

The default configuration has the following limitations:

■   Only the local subnet and devices under DHCP control will have full MAC address discovery. If there are machines on the network that are not under DHCP control (for example, lab machines with fixed IP addresses), MAC address discovery will depend on SNMP (which might not be enabled on every machine). Another issue would be that the MAC address classification rules would not work. This could possibly result in a device not being classified at all. Please see documentation about the Discovery Classification Engine.

■   In a large environment, the number of devices could exceed the maximum supported devices per agent. For better performance, the discovery load should be distributed over several agents.

There are possible limitations that could limit the devices Continuous Discovery can find. The following is a list of these conditions and how to address them through the Continuous Discovery configuration:

■   **Firewalls**

In an environment where traffic is blocked by a firewall, Continuous Discovery can still be used to monitor devices behind firewalls by placing an agent behind a firewall, with the CAM port enabled. (See a more detailed description of how to do this earlier in this section.)

■   **Fixed IP addresses**

In those environments where servers have been assigned fixed IP addresses and SNMP is not consistently deployed, there are three additional configuration options that can be used to achieve more comprehensive discovery results than would otherwise be possible in those environments:

> **ARP cache monitoring of routers**

The Continuous Discovery Agent can be configured to monitor the ARP cache of a router through SNMP (this is automatically enabled for the local gateway).

> **CTA enabled**

Run the agent of a router's network tab with CTA enabled and fully configured for all subnets. The advantage of this option is that only one agent needs to be deployed to cover all of the subnets that are covered by the router to get greater coverage.

> **Agent per subnet**

Deploy an agent into every subnet. This would require having access to at least one machine per subnet.

Installing Discovery Agents

In a perfect world you would have a discovery agent installed in each subnet in order to maximize the chances of quickly detecting new computers and also rogue network devices such as unauthorized Wireless Access Points. But in reality, this architecture for desktop management is impractical for most organizations.

To maximize the success rate of new computer discovery you need to have a good understanding of your company's network infrastructure. Discovery agents can be placed at strategic locations in the network where network traffic from the computer is most likely to be seen. Some examples are:

■ Web proxy server

■ Web server farm

■ Email server farm

■ Backbone switches or routers

For the CTA component of the discovery agent to function, it needs to see the network traffic of the devices. Today most networks are switched, meaning that the computer running the discovery agent would only see network traffic intended for itself. To overcome this problem, the network devices need to be configured to forward a copy of all traffic to the network port where the computer running the discovery agent is attached.

At the same time, specific network infrastructure components could be set to forward a copy of DHCP requests to the CA DSM domain manager, which is running the CCS Discovery Manager. This again increases the success rate of discovering new computers as they connect to the network.

Currently it is not possible to install just the discovery agent using the CA DSM installation media. If your architecture requires distributed discovery agents then you need to install these from the CA Unicenter Network and Systems Management (CA Unicenter NSM) media.

When deploying a remote discovery agent you must manually assign the agent to a discovery manager by editing the following registry key on the agent machine and restarting the agent:

HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Discovery\Shared Components\AgentEngine\AgentManager with the name of the domain manager machine.

The Discovery Agent Service must be restarted for the change to take effect.

Restricting the Type of Devices to be Discovered

The Continuous Discovery components are capable of discovering all types of devices, but in some cases you may want to restrict the discovery specific types. For example, perhaps you need to focus your implementation on managing just workstations, laptops, and servers and you are not be interested in printers or network infrastructure devices.

The discovery components include the flexibility to let you configure the operation to discard certain classes of objects and will only add the classes specified into the MDB.

Follow these steps to set up filters to exclude all classes from discovery other than workstation computers:

1. Choose Start, All Programs, Computer Associates, CA Common Services, WorldView. Then select the Object Browser. When the Object Starts, confirm the Repository name and enter.

2. In the left pane, expand TNGRoot, Reference, CaMtxReference, and click CaMtxClassFilterEntry.

   The CaMtxClassFilterEntry object opens in the right pane. This object contains a list of all the names of classes that are handled by the Discovery Engine.

3. To exclude a class or subclass, find the entry that corresponds to the class, highlight the entry, choose Object Delete from the main menu, and click OK on the confirmation window that appears.

   The class or subclass is deleted from the filter. Only objects in the classes and subclasses that appear in the filter are used during classification by discovery agents. Instances of classes that are excluded show up as unclassified TCP/IP devices.

4. Run the updateclassrules utility.

   Note: Be sure that the classdefinition.xml and classifyrule.xml files are writable before you run updateclassrules. The classdefinition.xml and classifyrule.xml files in the Discovery\Config folder are updated.

5. (Optional) If the discovery agent is not local, copy the classdefinition.xml and classifyrule.xml files to the corresponding folder on the discovery agent computer.

6. Restart the Discovery Manager and Discovery Agent.

Modify or Write Classification Rules

After discovery locates devices on your network, the classification engine then classifies these devices according to how you have configured the classification engine. The classification engine configuration files let you customize the discovery rules to your environment.

Classification means that a class and subclass are defined for each discovered object as it is added to the MDB. These classes are then referenced by the CA DSM domain manager when evaluating if it should try to automatically install an agent on the newly discovered computer.

You can modify the classification rules in the classify rule.xml file provided with CA Common Services in order to improve the classification success rate for servers, workstations, and laptops.

While you can write additional discovery classification rules to personalize the classification process, this is not recommended, as the CA DSM Auto Deployment functionality will only recognize the standard classes, such as Windows XP.

Discovery Classification Rules

To write discovery classification rules, follow these steps:

1. Modify existing rules or add new rules to the classifyrule.xml and methods.xml files. From a default installation this file can be found in the directory :\Program Files\CA\SC\CCS\Discovery\Config.

2. For *Continuous Discovery*, move classifyrule.xml and methods.xml to the '\config' folder on each discovery agent for which the rules apply. Then restart the agent.

3. For *Classic Discovery*, run ruletodbconverter.exe on the classifyrul.xml file in the discovery_install\config directory on the computer where the MDB resides. Then run Discovery.

Note: Each discovery agent can have a different set of rule files, but we advise against this.

For detailed steps on how to modify or write classification rules and for examples, see the How You Modify or Write Classification Rules chapter in the *CA Unicenter NSM Administrator Guide*.

Rediscovering Computers

The Continuous Discovery processes, and also when running, the Classic Discovery process, are looking for changes and updating/adding computers as needed. Sometimes, particularly in the test environment, it is necessary to take a previously discovered computer and discover it again. In order to do this we need to make it 'undiscovered' by removing it from the MDB. The Continuous Discovery processes use a caching system in order to improve efficiency across the network and when accessing the MDB. The steps needed to make a previously discovered computer undiscovered are as follows:

1. Stop the Continuous Discover Manager Service

2. Stop the Continuous Discover Agent Service

3. From the DSM Explorer, delete the computer from the All Computers list

4. From the 2D Map, delete the computer including any child objects

5. Start the Continuous Discover Manager Service

6. Start the Continuous Discover Agent Service

It takes a few minutes for the Discovery process to rebuild the cache and settle down, so we recommend waiting a few minutes before turning on the computer to be discovered.

Automating Discovery Events

When a previously undiscovered computer is detected on the network, there are a number of different views as to what should happen. Later in the chapter we will discuss the CA DSM view, which is to automatically attempt to bring the computer under management by installing the DSM Agent. If that action fails, open an incident with the service desk to have the computer located and removed from the network or brought under management by a manual agent installation.

A second view is that when a previously undiscovered computer is detected on the network, a change order should be opened requesting the computer be brought under management and the change request workflow should automatically install the DSM Agent.

To automate tasks, such as interaction with the service desk when new objects are discovered, we recommend the use of the CA Unicenter Event Management components, which can be used to automate the handling of discovery events sent by the discovery manager to CA Unicenter Event Management. The use of Event Management Message Record Actions (MRA) can then run any required actions.

Messages are sent to the Event Manager by the discovery components for the following events:

■ Address Change—A discovered device changes its IP address

■ Discovery Event—All discovery events enabled or disabled

■ Handshake Event—Communication events between the manager and agent

- New Device Events—New device discovered event

- New Subnet Events—New subnet discovered event

We recommend only enabling the New_Device_Events in order to reduce the number of events that need processing by the system.

About Classic Discovery

Classic Discovery is the discovery process that you can set up and run on demand to find and classify devices and then automatically register their existence in the MDB. Classic Discovery lets you decide which subnets you want to discover and when. You can start a Classic Discovery from the Discovery GUI or the command line (dscvrbe).

Note: Classic Discovery is typically needed only in those cases where you did not install the Discovery Agent and the Discovery Manager that continuously discover your network. If you are using the Continuous Discovery method, you do not need to run a Classic Discovery.

The IP Discovery process consists of the following main functions:

- **Ping**—IP Discovery identifies whether a network device exists and is able to communicate. The ICMP uses the ping utility to send requests to the designated computer at periodic intervals and waits for a response.

- **Simple Network Management Protocol (SNMP)**—After receiving a response and confirmation that a network device is valid, IP Discovery issues an SNMP request to the network device. This request asks for specific Management Information Base (MIB) information, which is used to classify and gather information about the network device.

Object descriptions and relationships based on the information in the device's SNMP Management Information Base (MIB) are then used by IP Discovery to create a managed object for this network device in the MDB. SNMP MIB agents typically are resident in network device firmware and are provided by each device's vendor.

Discovery also determines if a device provides Web-Based Enterprise Management (WBEM) data, and if so, creates a WBEM object in the device's Unispace.

Best Practices for Classic Discovery

We recommend that you run a combination of Classic and Continuous Discovery when you want to discover subnets. However, Classic and Continuous Discovery work differently depending on what options you select for both methods. Being aware of the different capabilities and mechanisms can help you make an informed decision about how to best utilize these components.

Combine Running Classic and Continuous Discovery

Classic Discovery and Continuous Discovery have access to different levels of the object (device) being discovered and may assign different names to the discovered device, accordingly:

■ Classic Discovery supports naming a device using its sysname (the MIB-II value for a device that supports SNMP), which is the default if no domain name server (DNS) name is available.

■ Sysnames, conversely, are not supported by Continuous Discovery. Continuous Discovery names devices based on their DNS names (except for routers that do not have valid DNS names for their IP interface cards).

■ To avoid the problem where the same device is discovered and registered using different names (when using a combination of Classic Discovery and Continuous Discovery), set the command dscvrbe -j option to IP to use the IP address if the DNS name cannot be found. Using IP addresses to name discovered devices ensures that objects are named using the same method and that no duplicates result. Set this option only if DNS is not enabled in your environment.

Note: When you run a full subnet discovery using Classic Discovery, stop the Continuous Discovery services.

■ Continuous Discovery discovers only subnets on which a managed agent is known to exist based on that agent being registered in the MDB. To use Continuous Discovery to monitor subnets where managed agents are not installed, run the Classic Discovery dscvrbe command to discover a router and all of the subnets it supports, or alternatively write a script using the dscvrbe -7 option to discover all of the gateways on the desired subnets.

Timeout Values

The values you specify for SNMP timeout and ping timeout greatly affect how successful and how long discovery takes to run. If you set higher timeout values, discovery takes

longer to run but has sufficient time to communicate with the devices and obtain the needed information. If you use lower timeout values, discovery runs faster, but devices may not be classified correctly, or even discovered at all. The length of the timeout is set by:

■ Using the command line by specifying the -W parameter on the dscvrbe command.

■ Using the Management Command Center Discovery or Advanced Discovery Wizard Timeouts page.

■ Using the Discovery Classic GUI Timeouts box on the Discovery page of the Discovery Setup dialog.

Subnet Filtering

Using a subnet filter on large networks with multiple large subnets is advantageous because you can limit your search to certain subnets within the network, which can mean a shorter discovery process.

Use a subnet filter to do the following tasks:

■ Limit the scope of discovery by confining it to a certain range of subnets and devices. For example, if you use the subnet filter 172.24.*.*, only the subnets from 172.24.1.0 to 172.224.255.0 are searched.

  If there is a subnet called 172.119.1.0, that subnet is not searched because it does not fall in the range specified by the subnet filter.

■ Enter a range of as many as ten filters. The filter statement uses a comma-separated format of a1.b1.c1.d1,a2.b2.c2.d2,...a10.b10.c10.d10.

  Only those subnets passing through filter1 (a1.b1.c1.d1) or filter2 (a2.b2.c2.d2) or filtern (up to 10) will be searched and created as TNG/IP_Subnet.

■ Use the default subnet filter of *.*.*.*, which does not limit the scope of the discovery process.

  After the selected subnets are searched by discovery, they are placed on a list in the right pane of the Discovery Subnet Management dialog.

Discovery Methods

You can use any of the following Classic Discovery methods to discover your network:

■ **ARP Cache**

  The ARP cache method starts at the gateway address (the address of the nearest router to the computer running discovery) for the current subnet and uses the ARP cache of that device to determine information about the devices. The ARP cache contains the IP-to-MAC (physical network) address mappings.

  Discovery retrieves the gateway address from the computer on which it is running and gets the IP list from the ARP cache on that router. It then discovers the subnets

nearest that router and for each subnet it discovers, queries its gateway, doing the same thing over and over again.

For each device found in the ARP cache, an SNMP request is initiated. If the device does not respond, it is assumed to be a non-SNMP device, just the IP address is retrieved, and the object is created as an Unclassified_TCP object.

■ **Ping Sweep**

A Ping Sweep pings all of the devices on the network based on the subnet mask, finds IP devices, and then retrieves SNMP information. If no SNMP information is retrieved, just the IP address is retrieved, and the object is created as an Unclassified_TCP device. This is the slowest but most thorough method.

■ **Fast ARP**

Similar to ARP cache, Fast ARP saves time by checking only the ARP cache of routers. Fast ARP is the best method for updating the MDB when you do not want to use the more intensive searches provided by Ping Sweep and ARP cache. This is the fastest way to discover your network.

■ **DNS Search**

The DNS Search limits the discovery of devices to those that are defined in the DNS. The IP address of each of these devices is combined with the defined subnet mask to determine whether or not to discover the device (In contrast, the Ping Sweep option tries to discover all active devices numerically, without regard to their definition in the DNS).

Each discovery method has advantages and disadvantages. The Ping Sweep method provides more comprehensive quantitative information—in the form of the number of devices—because each and every potential IP address on the network is pinged. Even devices not recognized by the router, which may not be discovered through the ARP cache method, can be discovered using Ping Sweep.

On the other hand, ARP cache provides the MAC and IP address information on all the devices that are found in the ARP cache of the router. Ping Sweep, however, generates additional network traffic and is thus more time consuming than ARP cache and Fast ARP. Sometimes, to discover every device in the network, a combination of Ping Sweep and ARP cache is required.

Number of Discovery Processes to Run Discovery

In the Discovery Setup\Service page you can specify the number of Instances (processes) discovery can use when doing a discovery. You can alter this number to be higher or lower with the following effects:

■ Having a low number of processes (1 or 2), the discovery process will consume less memory, take more time for discovery, and less network bandwidth will be used.

■ Having a high number of processes, the discovery process will consume more memory, the time taken for discovery will be shorter, and more network bandwidth will be used.

Note: The number of threads in the Discovery Setup\Discovery page also plays a part in how long discovery takes and network bandwidth used.

Preparation for Discovery

Before running discovery, use this checklist to ensure these prerequisites have been met:

■ The computer from which you are running discovery must be connected to the network and have a valid IP and gateway address. You can ping the gateway address to ensure TCP/IP connectivity.

■ You have the correct SNMP community names for all of your devices network infrastructure devices. The community name is case-sensitive and the default community name is 'public.'

■ The Host IP Address and the Gateway Address are displayed in the Discovery Setup dialog on the Discovery tab.

■ If you do not see these addresses, check your network setup in the Control Panel.

■ The subnet filter, subnet mask, and subnet count are set on the Discovery Setup dialog.

■ You selected the Enable the Trace Console checkbox if you want to monitor the progress of the discovery process.

Configuration of Classic Discovery

To run classic discovery (Discovery wizard)

1. On the Windows Start menu, select Programs\Computer Associates\CA Common Services\Discovery.

2. Click Auto Discovery.

   The Distributed Services window appears.

The Distributed Service container lists the database servers on which discovery may run, the name of the MDB on each server machine, and the current status and startup method for DISCOVERY. By default it should list two services—one for TCP and the other for XPX—with both having the Repository Name of your domain manager.

The Service's Host Selection buttons allow you to add database servers to the container list, as well as update the status of the list displayed. The Service Action buttons allow you to start and stop discovery and display the Discovery Setup dialog to set discovery service parameters, as well as the Add Server Machine dialog to add new remote servers to the existing list. Follow these steps:

1. Select CA-AutoDiscovery and click Setup.

    The Unicenter NSM Discovery Greetings Window appears.

2. Click Next.

    The Unicenter NSM window appears.

3. Select the correct MDB and click OK.

    The Unicenter NSM Repository Selection page appears and the repository you selected is displayed in the Repository Selection box.

4. Click Next.

    The Unicenter NSM Discovery Scope page appears.

5. Choose one of the following and click Next:

    > Discover the entire network.

    > Discover a subset of the network.

6. If you selected Discover, a subset of the network, the Unicenter NSM Subnet Management page appears.

    Add the Subnet Starting IP address, Gateway IP Address, and Community Names for the subnets you want to discover in the environment.

7. Once complete, click Next.

8. In the Unicenter NSM Discovery Methodology, select one of the following and click Next:

    >   Faster Discovery.

    >   Detailed Discovery.

9. In the Unicenter NSM Discovery DHCP Configuration page, enable the DHCP setting for discovery, if it is required, and specify the range in DHCP to run discovery on.

10. When these settings are configured as required, click Next.

11. In the Unicenter NSM Discovery Multiple Instances page, select Normal Discovery or Attempt Faster Discovery from the Multiple Discovery Processes.

    Note: Selecting Attempt Faster Discovery will increase the number of threads that will be used to run discovery and will increase the memory consumed by the discovery process.

12. Specify the SNMP/ICMP Timeout value.

13. When configured as required, click Finish.

    The discovery wizard closes.

14. In the Distributed Services window select CA-AutoDiscovery and click Start.

    The Unicenter Discovery Monitor comes up showing the status of the discovery, uptime for the discovery and the number of objects that have been discovered, processed, and added to the MDB.

    When discovery is complete, the status of the Unicenter Discovery Services window shows Stopped with a red icon and shows the number of objects in the Added column. You can close these windows now and view these objects in WorldView.

Running Advanced Classic Discovery

To run advanced classic discovery (discovery wizard)

1. From the Windows Start menu, go to Programs\Computer Associates\CA Common Services\Discovery\Auto Discovery.

    The Distributed Services window appears.

    The Distributed Service container lists the database servers on which discovery may run, the name of the MDB on each server machine, and the current status and startup method for discovery.

    The Service's Host Selection buttons let you add database servers to the container list, as well as update the status of the list displayed. The Service Action buttons let you start and stop discovery and display the Discovery Setup dialog to set discovery service parameters, and add the Add Server Machine dialog to add new remote servers to the existing list.

2. Select CA-AutoDiscovery and click Setup.

The Unicenter NSM Discovery Greetings Window appears.

3. Click Advanced.

The Discovery Setup window appears.

Use this dialog to configure when and how the discovery service runs, and to access the Discovery Subnet Management dialog to set which subnets are searched and the Enter New Community Name dialog to add a community name to the MDB. Review the history of the discovery processes that ran. The Discovery Setup dialog is divided into seven tabbed pages: Discovery, Repository, Included Classes, Services, DHCP, SAN Discovery, and Profiles.

4. In the Discovery page configure the options as required. The following fields are available:

Current TCP/IP Configuration box:

Host IP Address	IP address of the machine from which discovery is started. This is provided for information purposes only.
Gateway Address	Gateway address of the machine that discovery is started from. This is provided for information purposes only.

New Subnet Criteria box:

Subnet Filter	The Subnet Filter controls which subnets are searched, thus controlling the devices stored in the MDB. It is advantageous to use Subnet Filter on large networks with multiple large subnets because you can limit your search to certain subnets within the network.
Subnet Mask	The Subnet Mask is used when the MDB is empty in order to identify the subnet mask of your machine. A subnet mask is a 32-bit value that is used to extract the network ID and host ID from the IP address. This value affects the organization of devices displayed in the 2D Map.
Subnet Management button	Opens the Discovery Subnet Management dialog which allows you to select and deselect specific subnets to search, as well as add new subnets to the discovery process that have not yet been searched.
Subnet Count option button	When selected, specifies the number of undiscovered subnets to search using the text box to the right of this option button. If Subnet Count is set to 20, for example, discovery will run on the first 20 subnets it encounters, regardless of the subnet mask and subnet filter settings. This parameter can be set from 0 to 9999. **Note**: The Subnet Count does not include the subnets listed in the Subnets to Run Discovery container of the Subnet Management dialog. Discovery is run on the subnets listed in the Subnet Management dialog in addition to those undiscovered subnets specified by the number in the Subnet Count.
All option button	Causes discovery to search all subnets.

See the [Discovery Methods](#) section of this Green Book for details.

Threads: This field specifies how many parallel threads will generate SNMP queries on the network for discovery. This field works with the Run Multiple Instances field in the Service tab. See the [Number of Discovery Processes to Run Discovery](#) section of this Green Book for details.

Discovery Level Options box:

Rediscover Previous Subnets check box	Moves subnets from the Subnets Not To Discover list to the Subnets To Discover list of the Discovery Subnet Management dialog so that discovery will be run on those subnets, as well as new subnets (subnets on which discovery has not previously run). **Note**: A subnet will not be moved from the Subnets Not To Run Discovery list if it has been placed there manually.
Check Additional Ports	When selected and the device is non-SNMP, causes discovery to check the device's registry, FTP port, and Telnet port in order to match the device with an existing class stored in the MDB. These checks are run only if the SNMP check fails. When not selected, these additional checks (registry, FTP port, Telnet port) are not performed. If all checks fail, the device is classified as an Unclassified_TCP object in the MDB.
Discover SNMP Devices Only check box	When selected, discovery will create objects based on SNMP devices using more comprehensive information on these devices. Non-SNMP device information is limited to address only. When not selected, discovery treats all devices alike. The default setting is not selected.
Delete Old IP Interfaces check box	Controls whether discovery creates new router and interface objects in the MDB when the same router and its interface(s) are discovered during a subsequent run of Discovery. Initially, discovery creates router and interface objects when routers and interface(s) are discovered in the network. When a subsequent run of discovery finds the same router and its interface(s), the default is not to create the objects again and overwrite the existing router/interface objects in the MDB. When selected, discovery will overwrite the existing router and interface objects. Consequently, the new network topology can be displayed in the map.
Discover SAN Devices Only check box	Controls whether the discovery that runs on the subnets you specify in Subnet Management is limited to SAN objects. A typical IP discovery executes, but only SAN devices are added to the MDB. When the device discovery is complete, SAN links are determined and SAN Discovery will use the newly discovered SAN objects and those already existing in the MDB to determine the SAN configurations in the subnets.

Retries box:

ICMP Query Retry	Determines the number of pings, up to 5, sent to each device during discovery. This parameter is used for the Ping Sweep method only. ICMP Maximum Retry is typically set to 1 for small networks and 5 for very large networks. The higher setting for larger, more active networks allows for high network activity. The default setting is 2.
SNMP Query Retry	Determines the number of SNMP queries, up to 5, sent to each device during discovery. This parameter is used for the Ping Sweep method only. SNMP Maximum Retry is typically set to 1 for small networks and 5 for very large networks. The higher setting for larger, more active networks allows for high network activity. The default setting is 2.

Timeouts box:

ICMP Timeout	This is the ICMP timeout value when pinging a device and waiting for a response. The range of timeout values is 10 to 30000 milliseconds. The default timeout value is 2000 milliseconds.
SNMP Timeout	This is the SNMP timeout value when pinging an SNMP device and waiting for a response. The range of timeout values is 10 to 30000 milliseconds. The default timeout value is 2000 milliseconds.

Object Naming Options box:

Use Domain Name Server check box	Specifies whether a Domain Name Server (DNS) should be used. Use this option to avoid redundant names, which are not allowed, in the MDB. The default is selected.
Use IP Address Instead of sysName check box	Retrieves IP addresses and computer names (used for discovery) from each DNS.
Label Interfaces with DNS Name check box	Changes interface (router) labels to DNS name. For example: usgpko4F:999.999.999.99 is changed to jpnmachn: 999.999.999.99 where: usgpko4F is the router name associated with the IP address 999.999.999.99 is the IP address of the machine jpnmachn is the DNS name, which is determined by the IP address
Remove Suffix check box	This option allows you to truncate the name of objects based upon a list of suffixes. For example, if you want to remove the domain name suffix 'acme.com' from your device names, select Remove Suffix and then specify 'acme.com' in the entry field. During discovery, when devices are identified as device1.acme.com, the name saved to the MDB is device.1

5. After you have configured how you want the discovery process to run and what subnets you want to discover, click the Repository tab.

6. The Repository tab lets you set MDB criteria, and the SNMP Community Name.

    Add the SNMP community names used in your environment to this list by clicking New.

    You can also choose an update frequency for the MDB.

    > High option button: Allows updating of the MDB without any delay after discovering devices. This option depends on the machine that is contains the MDB.

    > Medium option button: Allows throttle time for updating MDB to be set to two seconds. This is helpful when multiple instances of discovery or multiple applications are accessing the MDB concurrently. This option depends on the machine that contains the MDB.

    > Low option button: Allows throttle time for updating the MDB to be set to four seconds or more. This can be done when multiple instances of discovery or multiple applications are accessing the MDB concurrently. This option depends on the machine that contains the MDB.

7. After the attributes in Repository tab are set, select the Included Classes tab.

    In the Included Classes page, include or exclude classes you want to discover. For example, unselect Workstation from the tree list if you do not want to discover workstations.

8. After the classes have been included and excluded as required, go to the Service tab.

    The Service page lets you specify a Startup Type and Execution Options.

9. Select from the following Startup Types:

    > Automatic

    The discovery service starts every time the system starts without using the Distributed Services dialog Start button. You can also start and stop discovery using the Start and Stop buttons on the Distributed Services dialog.

    > Manual

    Allows you to start and stop discovery using the Start and Stop buttons on the Distributed Services dialog.

    > Disable

    When Disable is selected, discovery will not run.

10. The following Execution Options are available:

    > Run Multiple Instances

When checked, allows more than one discovery service to access the MDB. You can then specify in the edit control text box the number of discoveries to run. Nine is the limit. If this number is left blank, discovery runs as if this option was not chosen.

> Show Trace Console check box

When selected, provides a DOS window that monitors discovery activity. The information is different than that of the message log. It is useful as a debugging tool.

> Message Log Level

When selected, sets the level of detail of error messages written to the log file. The higher the number entered in the text box to the right of this option button, the more detail. The log file can be found in the Install directory of WorldView under the subdirectory Log.

11. When finished configuring the Service tab, click OK.

The configuration is saved for discovery and the Discovery Setup page closes.

12. In the Distributed Services window select CA-AutoDiscovery and click Start.

The Unicenter Discovery Monitor shows the status of the discovery, uptime for the discovery, and the number of objects that have been discovered, processed, and added to the MDB.

If the Show Trace Console option was selected in the Service tab, a DOS console window appears when the discovery starts showing the trace of the discovery process.

When discovery is complete, the status of the Unicenter Discovery Services window shows Stopped with a red icon and displays the number of objects in the Added column. You can close these windows and view these objects in WorldView.

Running Classic Discovery from Command Line

You can run discovery from a command line using the command dscvrbe. The dscvrbe command line utility supports the following flags (defaults are within square brackets):

Standard Flags:

Flag	Description
-M [*.*.*.*]	Subnet filter. This filters the subnets defined in the discovery ipsubnet table.
-N subnet mask | [255.255.255.0]	Subnet mask.
-8 filename	Discover subnets defined in a file. This is a text file defining subnets, ranges, and exclusions. The subnet mask must be defined for entries.

Flag	Description
-S all | number | [0]	New subnet counter. Sets how many new subnets to discover.
-S 0	Will discover only the subnets defined in discovery ipsubnet.
-S all	Will try to discover everything it can, including new subnets added after discovery starts
-L [0] - 9	Level of Logging. The higher the logging level number, the more detail in the log file.
-V [Yes] | 1 – 9	Log level of console messages. Turns trace console on.
-5 Yes | [No]	Rediscover subnets previously discovered according to discovery ipsubnet table.
-U Yes | [No]	Port scan for unclassified devices (Telnet, FTP, HTTP, SMTP).
-21 0 – 65 [64]	Number of Threads for Discovery (higher thread level will result in speedier discovery).

Object labeling convention flags:

Flag	Description
-3 suffix	Suffix to be removed from device name. Example: -3 .ca.com means myserver.ca.com becomes myserver.
-4 Yes | [No]	Use DNS name for interface. DNS names for router and multi-home interfaces are not retrieved unless this option is on. If -4 Yes, name for interface would be dnsname.ip and label would be dnsname. If -4 No, Name and label for interface would default to parentname.ip.
-F [Yes] | No	Use DNS name or not. DNS name is now default name for all regular objects. Only the label will be affected by this switch. If –F No, the label will be either the sysName or IP according to –J option.
-J [sysName] | IP	If no DNS, use sysName or IP. If no DNS, this option will determine the name of the object. If –J sysName and no SNMP and no DNS, IP will be used. If –F No, this option also determines the label.

Router Flags:

Flag	Description
6 [Yes] | No | NoObj	Loopback for router or no objects.
	Use Loopback address for router address. If -6 No, the IP of the interface first accessed will be the router IP address.
	Example: challenger.ca.com
	Loopback = 172.24.0.4
	Interface = 172.24.4.1
	Interface = 141.202.243.249
	If dscvrbe is discovering the object by subnet 141.202.243.0, IP of router would be 141.202.243.249 without this option; with this option on, IP of router becomes loopback address, 172.24.0.4.
-9 Yes | [No]	Refresh routers and the interfaces.
	This option affects the update and creation of the router and its interfaces in the repository. If -9 Yes, the router will be updated and the interfaces will be dropped and recreated. If -9 No, the router will be updated, but the interfaces will not; new interfaces will be added.
-12 [Yes] | No	Perform crosschecks in Routing table. This option sets whether the Interfaces should be retrieved for a router.

Behavioral flags:

Flag	Description
-A Yes | [No]	Repeatedly discover the subnets.
-B Yes | [No]	Do not insert class B network.
-Q [Yes] | No	Stop discovery when finished.
-O Yes | [No]	Create SNMP device only.
-W millisecond | [800]	SNMP query timeout value.
-G [1] | 2 | 3 | 4 | 5 | 6 | 7 | 8	Number of SNMP query retries.
-E millisecond | [400]	ICMP Ping query timeout value.
-P 1 | 2 | [3] | 4 | 5	Number of ICMP pings per object.

Flag	Description
-T second | [0]	Database update throttle time.

DHCP environment flags:

Flag	Description
-K Yes | [No]	Delete old unclassified DHCP objects. If an existing object is unclassified, the old object will be deleted to make way for a classified object based on IP address.
-X Yes | [No]	Running in DHCP environment. If this is set, Fillsql will check the tng_dhcp_scope table to identify DHCP ranges.
-Z Yes | [No]	Update old classified DHCP object. If set, Fillsql will update the class of an existing DHCP object.

Miscellaneous flags:

Flag	Description
-1 string	Unique string identifying current run. This is the discovery id timestamp. It is used when starting the fillsql process.
-13 Filename	Preferred discovery log filename. Use this option to change the name of the log file (discover.log) to something unique.
-I unique number	Instance Number. This is the instance number used to control multiple instances of discovery.
-19 Profilename	Discovery Profiles. The Profile, created by tngsrvcs, is a configuration file for running discovery. This option is used by wvschdsv.exe for scheduled discovery.
-Y mm/dd/yyyy | [current day]	Date to start (also needs -H).
-H hh:mm | [current time]	Time to start (also needs -Y).

See the following examples of using the dscvrbe command:

■   Discover a single node

dscvrbe –R mdb_hostname -7 hostname_to_be_discovered –v 9

■   Discover multiple subnets using a file

dscvrbe –R mdb_hostname -8 discover.txt -v 9 -S ALL -D PINGSWEEP

Where the discover.txt should look similar to the following:

141.202.236.0,141.202.236.1,255.255.255.0

141.202.237.0,141.202.237.1,255.255.255.0

141.202.114.0,141.202.114.1,255.255.255.0

■   Discover without suffix

dscvrbe –R mdb_hostname -7 hostname_to_be_discovered -3 suffix_to_be_removed –
v 9

Infrastructure Deployment

Introduction

The Infrastructure Deployment functionality of the CA DSM domain manager facilitates the
initial deployment of CA DSM infrastructure components within a heterogeneous enterprise.
Infrastructure Deployment is also sometimes known as DMDeploy. The CA DSM
infrastructure components, such as agents and scalability servers, can be transferred and
installed on to a system that currently does not have the CA software installed. This
functionality is typically used for an initial roll out of the CA DSM infrastructure and
subsequent rollout of new agents and scalability servers.

The DSM Explorer provides access to the deployment functionality allowing the DSM
Administrator to both manually perform and monitor a deployment, or configure

deployment automation rules so that as Continuous Discovery detects new computers the DSM Agent is automatically deployed to them.

Both the creation of deployment jobs or deployment policies is done through a wizard. The wizard takes the administrator step by step through the actions needed to perform either process.

Deployment System Architecture

The Deployment system architecture consists of a number of components listed below:

Deployment Manager: (DMDeploy) The manager component of the deployment system residing in the domain manager system.

Deployment Management - Bootstrap Program: (DMBoot) A simple bootstrap installation program intended to install the DMPrimer on target systems.

Deployment Management - Primer: (DMPrimer) The client component of the deployment system. This component is distributed to the machines targeted for DSM Agent deployment. It is responsible for the transfer, install, logging, and reporting of agent installations as directed by authorized Deployment Managers.

Deployment Management - Wizard: Client GUI application of DMDeploy.

Deployment Management - Sweep: (DMSweep) The command line component of the deployment system. This component issues commands to the DMDeploy through the Deployment Management Application Program Interface (DMAPI) and receives scanning and monitoring data.

Deployment Management – Continuous Discovery for Deployment: The component of the deployment system that receives notification upon the discovery of a new asset. It evaluates this asset against deployment policy and instructs the deployment manager to build a deployment job.

While the scalability server is used to hold copies of the Agent Packages and Deployment Primer it has no running deployment processes.

Creating Deployment Policy

The CA DSM Continuous Discovery and Deployment component allows you to automatically deploy the DSM Agent to newly discovered assets. We have already discussed the need to configure the Continuous Discovery Agent/Manager in order to detect assets as they appear on the network. Now we will discuss the corresponding configuration of the CA DSM Continuous Discovery Deployment Policy.

The Continuous Discovery Deployment Policy wizard guides you through the eight steps for creating individual deployment policies. At a minimum, one policy will be required for each scalability server within your architecture. Additional policies may be required for the following:

■   Different platforms

■   Different subnets

■   Different security credentials

Step 4 of the wizard asks you to specify Target Criteria in the way of a TCP/IP range of addresses. This is different to the Continuous Discovery which is configured by subnet. For example, Continuous Discovery may be configured to monitor three subnets:

■   10.0.1.0

■   10.0.2.0

■   10.0.3.0

But this can be specified in the Deployment Policy as a single IP range.

Step 5 of the wizard asks you to specify the target platforms. You can select individual or multiple platforms to be targeted.

In the example above, the class of Windows is selected meaning ALL Windows versions.

Once the policy is complete, the Deployment Manager will automatically build a deployment job for only undiscovered computers that match this policy.

From the CA DSM perspective the asset can have one of three statuses:

■ **Undiscovered:** The computer has never been detected and no information written to the MDB.

■ **Unmanaged:** A Computer has been discovered and information written to the MDB but a DSM Agent is not installed.

■ **Managed**: DSM Agent is installed.

We will discuss deploying the DSM Agent to unmanaged computers in the next section.

Creating Deployment Jobs

Deployment jobs can also be created manually using either the Infrastructure Deployment Wizard or Command line.

The Deployment Wizard

The Continuous Discovery Deployment Policy wizard guides you through the eight steps for creating individual deployment policies. A deployment job can either be to stage a

deployment package at scalability servers or to deploy a deployment package to target computers. In this section, we will discuss some of the steps involved with the deployment of the DSM Agent to target computers.

Step 3 of the Deployment Wizard (Payload) asks you to select the package that will be used for this deployment Job. Only ONE package may be selected for each deployment job. If it is necessary to deploy multiple packages to the same computer(s), then multiple jobs must be created. The list of selectable packages is read from the deployment library of the domain manager (see Customizing the Agent Packages later in this chapter) unless the option to Transfer Packages From Scalability Server is selected. If this option is selected then you are asked to select the scalability server, and the list of available packages is read from that scalability server deployment library.

Step 4 of the Deployment Wizard (Target Criteria) asks you to select the method of discovering the computers that the agent package should be deployed too. Six different methods are available:

■   **Deploy to all computers in a specific domain** – The Deployment Manager issues a Microsoft network browse for the specified domain. Only computers that are a member of that domain and currently active on the network are listed.

■   **Deploy to specific computers –** You enter a host name or IP address. Multiple host names and IP address can be specified by using a comma to separate each entry.

■   **Deploy to computers within an IP address range –** The Deployment Manager will attempt to contact each IP addresses within the range specified. Optionally a Microsoft Windows Domain name can also be specified and then only computers found that are also a member of the specified domain are listed.

■   **Deploy to computers within a directory –** You are asked to enter the URI of the Directory, OU, or computer group. All computers specified by the URI will be listed. By selecting the Browse button you may browse the available directories, OU, or Computer Group(s) in order to make your selection—but only a single selection may be made per deployment Job.

■   **Deploy to computers specified in a Target Credentials File –** A list of host names and/or IP addresses with credentials for each target computer are listed in a Target Credentials File. How to create a Target Credentials File is covered in the next section (The Deployment Command Line).

■   **Deploy to computers specified by a query -** The Deployment Manager issues a CA DSM Query. Only computers that are returned from the results of the query are listed. For example, the query Unmanaged Windows Computers will list all computers found in the MDB with a status of Unmanaged.

Step 5 of the Deployment Wizard (Scan Targets) takes the list of computers from Step 4 and looks for them on the network. The scan process will show if the computer is active on the network and its current status. You can suspend the scan at any point and go back to previous pages, or you can move onto the next page in the wizard even though the scan is not complete. This is useful if the computers you require have already been discovered within the search criteria you specified.

Legend							
This table describes the meaning of the icons in the discovered computer table.							
	Description		Description		Description		Description
	Ready To Deploy		Requires Credentials		Already Deployed		Deployment Inprogress
	Newer Version Installed		Older Version Installed		No Response		Unknown

Step 6 of the Deployment Wizard (Target Selection) asks you to select from the discovered computers the ones to which you wish to deploy the agent. If required, you will be prompted to enter credentials for accessing the computer(s) that you wish to install the agent upon. An individual set of credentials can be specified for each computer, or one set can be specified for all those selected.

Step 7 of the Deployment Wizard (Agent Configuration) allows you to enter any optional installation parameters to the Agent and Deployment Primer install processes. See the Agent Packages section later in this chapter for further details.

When the Infrastructure Deployment wizard is completed a deployment job is created which manages the progress of the deployment to all the selected end systems. The status of this job can be viewed using the DSM Explorer, but please note that the job is NOT persistent; it is maintained in the memory space of the DMDeploy manager process. If the DMDeploy manager exits for any reason, is restarted, the machine is rebooted, and so forth, the job information and any unprocessed status messages from end system installations are lost. These installations, however, will continue if already started.

The Deployment Command Line

In a security-conscious era, dmsweep is sometimes perceived to have some security weaknesses. First of all, passwords entered on the command line are clearly displayed on the screen. Furthermore, they can be retrieved, on UNIX systems at least, by simply issuing the 'ps' command. To prevent this exposure of passwords on the screen, it is possible to put the password in a file for later use. But again, the password is often in clear text—not a very secure alternative.

This dmdeploy command line has two features which significantly reduce the risk of inadvertently exposing passwords to unauthorized parties. These are:

■   To render passwords entered on the command line invisible. To this end, the user will be able to state that they wish to enter a password but do not wish to have it displayed. They will subsequently be prompted for the password, which will not be echoed to the screen.

■   To allow passwords stored in a file to be held in an encrypted form.

Prompting for Passwords

There are three ways of specifying passwords on the command line:

■   **dmsweep deploy ... /tu < target user > /tp < password > ...**

In this case the password is displayed on the screen.

- **dmsweep deploy ... /tu < target user >**

  Here the password option (/tp) is not supplied. This causes a blank password to be used.

- **dmsweep deploy ... /tu < target user > /tp ...**

  Here in the third option /tp is given, *without* a password. As a result, the user will be prompted for the password. No characters will be echoed to the screen as the password is entered.

The Target Credentials File

The target credentials file can be used to hold a list of machines to deploy to, and the credentials needed to access those machines. The target credentials file is identified on the command line by the option **/targetcred** or **/tc**:

For example: *dmsweep deploy /tc < credentials file name > /pn < package number > /pparms < parameters >*

The following are examples of the different ways entries could be made in the file:

1. /ip machine1 /tu username1 /tp password1

2. /ip machine2 /tu username2 **/tp**

3. /ip machine3 **/tu /tp**

4. /ip machine4 /tu username4 **/ep < encrypted password >**

5. /ip machine5 /tu username5 **/ep**

6. /tu defaultUser1 /tp defaultPassword1

7. /ip machine6 /tu username6

8. /ip machine7,machine8,machine9 /tu /tp

9. /ip < IP address1 > /toip < IP address2 > /tu < username > /tp < password >

10. /domain < domain > /tu /tp

11. /ip machine10,machine11,machine12

Each entry in the file uses the same syntax used on the command line, for consistency, and here too it is possible to be prompted to supply the value for an option which is not in the file. If that option is a password, it will not be displayed on the screen when it is entered at the keyboard.

A fuller explanation of each line follows:

1. The first line in the file, for **machine1**, has a password in clear text. This option is still retained for those who wish to use it.

2. For **machine2**, the password option is present but no password is given. The user will be prompted for the password and it will not be displayed on the screen.

3. Dmsweep will prompt for both the username and the password associated with **machine3**, as both options have no values.

4. On the 4th line we see a new option, /ep. This is for encrypted passwords. Here a password is supplied but it is in encrypted form. The /ep option alerts dmsweep to this fact, and the password will be decrypted before being used.

5. **Machine5** also has an encrypted password associated with it but no value is given. As a result the user will be prompted to enter the *encrypted* password. This particular line can be of value in UNIX systems, for example, where the password is piped to the dmsweep command line.

   For example: cat $encryptedPasswordFile | dmsweep deploy /tu username /ep

6. On the sixth line we see a username and password without a machine. These become the *default* username and password, to be used for any subsequent entry in the file with no /tp option. They remain in force until another username and password without a machine are encountered in the target credentials file. The new values will then become the default username and password.

7. **Machine6** on the seventh line has no /tp option. The default password set in the preceding line will be used for this machine. No prompt will be issued.

The remaining entries show that it is possible to specify lists of machines, IP ranges or even domain names in the target credentials file. The last line has only a list of machine names, but no username or password. In this case, both the default username and default password will be used for each of those machines.

It is expected that each entry in the target credentials file occupies a single line.

In order to create the encrypted passwords for use in the credentials file, you must use the command CAF SAVECREDS. This command will output to a file called cafcreds.txt and from there you can copy and paste the information into the dmdeploy target credential file.

For example, *caf savecreds dummy user administrator password secret host* machine4 will result in the file cafcreds.txt containing a line 'machine4' *'dummy'* ' machine4*administrator'* *'4DDzqKuSJN5Ml6RsEY+o5Q'.'*

In the previous example, line 4 would now become:

   4. /ip machine4 /tu machine4\administrator /ep 4DDzqKuSJN5Ml6RsEY+o5Q

A full list of the parameters available for the Deployment Command line can be found in the CA Unicenter Desktop & Server Management installed documentation, *Command Line Reference Guide*.

Agent Packages

When the CA DSM domain manager was installed, packages were automatically set up for use by the deployment system and its local scalability server.

To reduce overall network traffic, deployment packages can be staged to the other scalability servers that make up the CA DSM infrastructure. This operation is basically a normal deployment except that the installation command is replaced by a copy into the staging area.

The deployment wizard or command line is the only method of staging agent packages to scalability servers, as they are held in the staging area as a single encrypted file ready for use by the deployment process.

On Windows, the staging area is a share called (DMDEPLOYSS$). On Linux/UNIX, the staging area is either accessed using SSH or an FTP server, depending on how your CA DSM architecture was specified during product installation.

When using the Deployment wizard and selecting the option Deployment from a Scalability Server it simply takes the agent package data from the staging area rather than the manager library area.

In r11.1, the primer installation image is transferred from the Manager, while in r11.2 it can be staged at the scalability server.

Even when a scalability server is used to stage and deploy software packages, control of the process is still handled by the Deployment Manager.

Removing Agent Packages

If, in your implementation, you only need to use a specific package (for example,the DSM Agent + all agent plugins 11.2.226.1085 windows_x86 ENU,DEU,FRA,JPN), then it makes sense to have only that one package available for selection. Having all the packages available for selection can be confusing and may lead to the wrong one being accidentally deployed.

CA. Unicenter® Desktop and Server Management

Deployment : Deployment Payload (Step 3 of 8)   [Back] [Next] [Finish] [Cancel]

DSM Explorer > DSMDEMO - Domain > Control Panel > Deployment > Infrastructure Deployment Wizard

| 2 Job Type | 3 **Payload** | 4 Target Criteria | 5 Scan |

☐ **Transfer packages from Scalability Server**   [Configure]
If you do not use a scalability server the packages will be transferred from the DSM Manager.

It is only possible to deploy a single package in a deployment job. Please select a package from those available below:

○ **CA Unicenter DSM Agent + Basic Inventory plugin (ENU) 11.2.226.4182 Linux_x86 ENU**
Deploy Agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Basic Inventory Plugin (English only Edition) 11.2.226.1085 Windows_x86 ENU**
Deploy agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Basic Inventory plugin 11.2.3.1658 Linux_x86 ENU,DEU,FRA,JPN**
Deploy agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Basic Inventory Plugin 11.2.226.1085 Windows_x86 ENU,DEU,FRA,JPN**
Deploy agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Basic Inventory Plugin (English only MBCS) 11.2.226.1085 Windows_x86 ENU**
Deploy agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Basic Inventory Plugin (MBCS) 11.2.226.1085 Windows_x86 ENU,DEU,FRA,JPN**
Deploy agent with Basic Inventory plugin

○ **CA Unicenter DSM Agent + Asset Management plugin (ENU) 11.2.226.4182 Linux_x86 ENU**
Deploy Agent with Asset Management plugin

○ **CA Unicenter DSM Agent + Asset Management Plugin (English only Edition) 11.2.226.1085 Windows_x86 ENU**
Deploy agent with Asset Management plugin

[Back] [Next] [Finish] [Cancel]

The packages are stored in the Packages folder of the CA DSM installation directory. They are stored in a folder structure of Agent Type, Version, Language, and Platform.

To have only the packages displayed that are relevant to your environment, all you need to do is remove the ones that are not required from Packages folder structure.

CA.

```
☐ 📁 Packages
    ⊞ 📁 Private
    ☐ 📁 Public
        ☐ 📁 CAUnicenterDSM
            ☐ 📁 AllAgents
                ☐ 📁 11.2
                    ☐ 📁 ENU,DEU,FRA,JPN
                        ⊞ 📁 Windows_x86
```

Every time the Deployment wizard starts it reads the lists of packages available from the folder structure and presents those for selection. In this example, we have removed all but the current version of the All Agents package, so that is the only selection presented by the wizard.

Adding Agent Packages

When a new Platform, Language, or Version is introduced, the new packages need to be imported onto the domain manager from the DVD. The dsmPush copy script is used to import packages for the specified products and platforms into the Infrastructure Deployment library or into the CA Unicenter Software Delivery library.

For the full description on how to use the dsmPush functions, please see the *CA Unicenter Desktop & Server Management Reference Guide* which is part of the CA DSM online documentation set.

Customizing the Agent Packages

When the Deployment Primer starts the installation process it passes a number of parameters to the installation program, some of which can be user defined. The user defined parameters are entered from the Deployment wizard during the process of defining the deployment job.

The predefined and user defined parameters are controlled by the settings contained in the file dmdeploy.dat. In the example of our All Agentspackage, the file is located in the directory:

```
:\Program
Files\CA\DSM\Packages\Public\CAUnicenterDSM\AllAgents\11.2\ENU,DEU,FRA,JPN\Windows_x
86
```

If we wanted to remove the risk of accidentally passing incorrect or invalid options to the installer, then we could do this by editing the dmdeploy.dat and deleting the **bolded text** in the example below:

```
[Deployment]
        ItemName = "Deployment"
        ItemVersion = 1.0
        OSType = 2
        Comment = "Deploy agent with all plugins."
        ProductDisplay = "CA Unicenter DSM"
        PackageDisplay = "Agent + all agent plugins"
        VersionDisplay = 11.2.226.1085
        SortOrderDisplay = 600
        WindowsParameters = " /qn AGENT_SERVER=$1$ ALLUSERS=1
ENSURE_CAF_STOPPED=1 $2$ EXITFILE=$E$ /l*v
```

```
       %TEMP%\ITRMAgentSetup.log "
       WindowsProcToRun = "DeployWrapper.exe"
       MSIProductCode = {501C99B9-1644-4FC2-833B-E675572F8929},{624FA386-3A39-4EBF-
9CB9-
       C2B484D78B29},{84288555-A79E-4ABD-BA53-219C4D2CA20B},{62ADA55C-1B98-431F-
8618-CDF3CE4CFEEC}
       REGProductCode = PRODUCT=[CA Unicenter DSM]PACKAGE=[Agent + all agent
plugins]VERSION=[11.2.226.1085]
       $1$ = Please enter the Scalability Server address to connect
to&^[^<>`~!/@\#}$%:;)(_^{*=|'+]{1,255}$
       &1&SCALABILITY_SERVER
       $2$ = Please enter any additional Windows install
options&.*&1&WINDOWS_OPTIONS
       $3$ = Due to the combined MSI package nature of this deployment package, MSI
feature list properties such as ADDLOCAL must not be specified. Otherwise, the
deployment of the package will fail. If specific features for a package need to be
listed, the recommendation is to use individual agent plug-in deployment packages
instead.
```

This results in the option to enter User Parameters being removed from the wizard for the deployment of this package.

In the same way, if we had a small environment and wished to hard code the name of the scalability server that the agent will report to, this can be done by changing the *1* on the line starting *WindowsParameters=* to the name of the server, and removing the line starting *1 = Please enter the Scalability Server address*.

Note: Changes made to the dmdeploy.dat file are not preserved by the upgrade process or by use of the dsmPush script. Any changes made will need to be reapplied.

Important: Although it is possible to specify additional MSI command line properties for the Windows versions of the deployment packages, special consideration must be taken for the Windows versions of the Agent + All Plug-ins and scalability server packages. Due to the combined MSI package nature of these packages, package-specific MSI feature list properties, such as ADDLOCAL, must not be specified; otherwise the deployment of the package will fail. If specific features for a package need to be listed, it is recommended to do this for the dedicated agent deployment package in question.

A full list of the parameters available to the DSM Agent and Scalability Server Installer can be found in the *Unicenter Desktop & Server Management Implementation Guide*.

Understanding the Deployment Process

The deployment process is the same regardless of whether the Wizard or Continuous Discovery provides the information necessary for the Deployment Manager to create a new deployment job.

All operating systems are not created equal. The Deployment Manager has different capabilities on different platforms. Two restrictions on the Linux Deployment Manager are that:

■   It cannot push out a primer to Windows shares since it cannot run the installation command.

■   It cannot enumerate Windows domains.

You will get an error if you try to do the latter. The Windows manager, however, is fully capable of SSH and Telnet/FTP deployment to UNIX variant targets.

The two diagrams below give an overview of the processes and data flows involved.

Automated Discovery and Deployment:

Manual Discovery with Automated Deployment:

We will now look in more detail at some of the steps of the Deployment Wizard, the Deployment Manager, and the configuration settings that affect their behavior. The diagram above shows the main steps:

1. Using Wizard or DMSweep, you select the payload to deploy and specify the list of target machines.

2. For each machine in the list, the Deployment Manager performs a number of steps, each conditional on the success of the previous.

   a. It will check to see if the machine is in DNS.

   b. It will then try to open TCP/IP socket 7 (ping the machine). The configuration setting Ping Check Skip can be used to completely disable the ping check.

   c. CAM pings to determine the presence of a DMPrimer.

   d. Lastly it checks for installed packages. A more detailed explanation of this step can be found later in this chapter.

3. The administrator selects machines for deployment and enters credentials if required.

4. For each machine in the list specified by the administrator, the Deployment Manager will now deploy the DMPrimer. The steps taken next depend on if the target operating system is Windows or Linux and whether the deployment is through a scalability server. If the deployment is from the domain manager, then the full DMPrimer is copied to the target, but if the deployment is from the scalability server then the small DMBoot is copied to the target and it in turn gets the larger DMPrimer from the

scalability server. The Deployment Manager will sequentially try three different methods to get the DMPrimer or DMBoot files to the target:

> Windows share ($admin)

> SSH/SFTP

> Telnet/FTP

If all three methods are unavailable, the message 'Failed to Telnet' is returned. This error is usually associated with the No Primer Transport status and is preceded by 'failed to map share' and 'failed to SSH' errors.

The DMPrimer installation is run, and once completed it notifies the Deployment Manager (using CAM) of successful installation so that package deployment can now be initiated. If the Deployment Manager does not receive the install success message, it results in the job status of No Primer Transport.

Further discussion of the Window Share and Linux/UNIX SSH methods can be found later in this chapter.

5. The Deployment Manager sends the payload data to the DMPrimer (using CAM messages). The data appears as a primer.packN file in the Primer installation folder ('N' is a small number). When the transfer is complete the DMPrimer unpacks the payload data and sends a 'package received' message to the Deployment Manager.

6. The Deployment Manager then sends the payload installation command in a CAM message to The DMPrimer which executes it.

7. Once the payload installation is complete, the DMPrimer sends the installation status to the   Deployment Manager and the job is status is shown according to the success or failure of the installation.

The job status showing as successful is not enough for the computer to appear in the DSM Explorer under All Computers. Once the installation is complete, CAF will automatically be started, the agent will then register to the scalability server, and when the engine next collects the inventory, the registration event will be processed. The new computer will then appear under All Computers in the DSM Explorer.

By default, the Deployment Manager will perform 10 concurrent deployment tasks; that is the agent or primer is being copied/installed on 10 target computers at the same time. If your network and server capacity is sufficient, then the number of concurrent deployment processes can be configured using the common configuration policy setting Deployment Thread Limit. If your agent deployments are through scalability servers, then this value can be higher than if the agent is deployed directly from the domain manager.

Deploying to Windows Over a Network Share

Deployment to Windows targets using DMDeploy is subject to a number of environmental obstacles. As with Linux/UNIX deployments, many things have to occur successfully before the deployment payload becomes active on a target system. Unfortunately, since no CA software is assumed to be present during deployment, it is not always possible to report accurate status and diagnostic messages during this operation.

Following is the sequence of events that occurs when Windows is deployed to a clean system using shares.

1. The Deployment Manager attempts to open a share which by default is admin$ on the target system. The share name used can be changed using the configuration parameterTarget Share Name.

2. If the share is opened successfully the primer installation package is copied to the target system. This consists of the following four files:

   > DMBoot.exe

   > dmkeydat.pmr

   > dmsetup.exe

   > msvcr71.dll

3. When the files have been successfully copied, the primer installation is launched by the MSI installer. After the primer installation has finished, you should see the following registry key:

   `HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\DMPrimer`

   The key's data will show where the primer files have been installed. The default location is:

   `C:\Program Files\CA\DSM\DMPrimer`

4. The primer should now have started and 'dm_primer.exe' should appear in the task manager. The primer log file, C:\Documents and Settings\<user name>\Local Settings\Temp\dmprimer.log, should also appear at this point.

   Some operating systems, such as Windows 9x, do not have a method for remote invocation. In these cases, it is necessary to configure the OS to install the primer on a significant operating system event, such as a reboot.

5. The transfer of the deployment payload (for example, CA Unicenter Asset Management agent) starts and you should see the 'transferring xx%' messages appear in the GUI or the CLI. This is quite quick in comparison to the primer upload.

6. Package installation starts and you should see an '.msi' process corresponding to the package being installed (for example, 'AgtAM.msi' for the Agent + Asset Management plug-in package) in the Task Manager. The 'Add or Remove Programs' utility, accessible from the Control Panel, will display the packages currently installed. After completion, the deployed package will be present, the exact entry obviously depending on which payload you installed.

The package installation status is sent back to the manager and displayed in the GUI/CLI.

Deploying Over SSH

Deployment to Linux or UNIX targets using DMDeploy over SSH (the default method) is quite complex and subject to a variety of environmental obstacles. Many things have to occur successfully before the deployment payload becomes active on a target system.

Unfortunately, since no CA software is assumed to be present during deployment, it is impossible to report accurate status and diagnostic messages during this operation. Therefore an understanding of the process is important.

Following is the sequence of events that occur when deploying to Linux using SSH, to a clean system:

1. The Manager connects to target system using ssh. You should see the connection attempt logged into the system log file /var/log/secure.

2. If the connection is successful, the primer installation package is pushed to the target system using ssh/sftp.

3. Primer image is uploaded to /tmp/dmprimer (You can do 'ls –ltr' in this directory to monitor the   growth of files as they are uploaded. This directory is not removed after the installation.

4. When the image stops growing, the primer installation (installdmp) is launched.

5. After the primer installation is finished, you should see two packages called 'ca-dsm-dmprimer' and 'ca-dsm-dmprimer-standalone' in the pif package list ('lsm –lOpif. Primer files are installed into the following directory:

   `/opt/CA/UnicenterDSM/dmprimer`

6. DMPrimer should then be started and you should see 'dm_primer start' appear in the process list. The DMPrimer log file (/tmp/dmprimer.log) should also appear.

7. The transfer of the deployment payload (for example, CA Unicenter Asset Management agent) starts. You should see the 'transferring xx%' messages appear in the GUI/CLI. This is quite quick in comparison to the primer upload.

8. Package installation starts. You will see the 'installdsm' process in the ps listing.

9. The command 'lsm –lOpif' will display the PIF packages currently installed on the Linux/UNIX. After completion, the package ca-dsm will be present. Other sub-components will be present depending on which payload you are installing.

Package installation status is sent back to the manager and displayed in the GUI/CLI.

Deploying Over Telnet/FTP

Deployment using FTP can seem back to front until you think about it. This method of deployment works as follows:

1. First the Deployment Manager connects to the target using Telnet.

2. Then the Deployment Manager issues FTP commands over Telnet on the target, pulling the primer installation image from manager to target—in other words, initiating the FTP 'get' request on the target. This is different for deployment using shares and ssh, which are pushes from the manager to target. This method only requires that a single FTP server be set up on the Manager, rather than having to have one on each target.

Detection of Deployed Packages

Detection of deployed packages differs based on the operating system involved. On Windows this is done through MSI product codes. The CA DSM payloads (packages) include MSI product codes within dmdeploy.dat. The primer uses this code and interrogates the MSI database to check whether a product has already been deployed. A payload may consist of multiple MSI products (for example, All Agents package).

On Linux/UNIX, registration is recorded in the following file:

`$CA_ITRM_BASEDIR/dmprimer/bin/dmdeploy.reg`

The installer registers/deregisters a product by calling dmprimer with special arguments. The Infrastructure Deployment Manager asks 'is product x version y installed?,' not 'what products have you got?'

The configuration parameter 'AlwaysDeploy' can be used to force a payload push/installation even if a newer version of the payload is detected on the target computer.

Upgrade of Primers

Primers are not automatically upgraded when deploying newer agents to a machine where the DMPrimer is already installed. In order for the DMPrimer to be upgraded, the configuration policy setting 'Always Deploy Primer' needs to be set on the Deployment Manager. This setting can also be used to force a reinstall of a primer that you suspect may not be functioning correctly.

Deployment of DMPrimer uses native OS data transport mechanisms or is done manually, but all subsequent network activity is performed using CAM (such as pushing payloads and returning payload installation status). This reliance on CAM isolates the main portion of the deployment process from networking issues and configuration.

Note: This process requires re-authentication (to push new encryption keys).

Deployment Security

Authentication between the Deployment Manger and Deployment Primer is done using public/private key exchange. The Deployment Manager generates a public + private key pair when it is installed.

If the primer is pushed out from the Deployment Manger, the public key goes with it, copied to the target share, or sent through SSH/FTP. At this time we know you are able to deploy to this machine because you have supplied a valid administrator user and password to connect to the share, or remotely log in to SSH or Telnet.

Once the public key is in place on a target, you do not need to supply credentials again. This is what the 'ready to deploy' state means. 'Credentials required' means that the public key is not in place or is invalid.

The primer is able to handle public keys from multiple managers. When authentication is successfully achieved from manager to target, the public key is renamed from the initial name dmkeydat.pmr to <manager name>.pmr.

After the major part of the implementation project is complete, we recommend that the deployment keys are recreated. Deleting the Deployment Manager's private key and restarting the deployment manager process will force the generation of a new key pair. When you next deploy you will need to re-authenticate with all existing primers, sending the new public key out using the normal mechanisms: shares, ssh, or FTP. This means that an administrator user and password will be requested again the next time you use dmdeploy.

Access to the Infrastructure Deployment functionality is controlled by CA DSM Class security.

This table summarizes what users can do for each permission bit they have set:

Permission Bit	Permitted Activity
-	None
V	View the version of DMDeploy
R	List packages, get rules
X	Scan, Deploy, Stage, Abort, Suspend, Resume, Enable rule
D	Delete machine from job, Delete rule
C	Create rule, Update rule

See the Security chapter in this Green Book for further discussion about using role-based security.

Chapter 14: Agent Best Practices

The CA Unicenter Desktop & Server Management (CA DSM) Agent is the one component of the solution that will be implemented on every single computer that you wish to manage (every computer in your company!). It will be installed on various platforms and operating systems from UNIX to Windows, and the requirements for its use will differ depending on the platform and its role of Server or Workstation.

This chapter, while applicable to both Servers and Workstations, will only consider some of the options available when the Agent is implemented in the Microsoft Windows environment.

This chapter also assumes that the following Agent components are deployed:

■   Unicenter Software Delivery

>   Data Transport Service

■   Unicenter Asset Management

■   Unicenter Remote Control

Unicenter CA DSM Architecture

The following graphic summarizes the CA DSM architecture:

DSM Architecture

The DSM Agent

The DSM Agent consists of a number of components, with the Common Application Framework (CAF) providing all of the common functions and environmental control to the functional plug-ins. The three main plug-ins provide the functionality of Software Delivery, Asset Management, and Remote Control.

What Happens at Startup?

When a machine running the DSM r11 Agent boots up, the DSM Service starts. This service in reality is CAF. Then CAF, in turn, will start up all configured plug-ins.

The following diagram provides an overview of the startup process.

Load config	Load configuration details from comstore
Start SM	Start session messaging plug-in
Start CCNF	Start Common Configuration Agent plug-in
Order plug-ins	Work out order of plug-in startup based on dependencies
Start plug-ins	Start all auto-start plug-ins in appropriate order
Start Scheduler	Start scheduler
Wait for msgs	Wait for messages and events

There are some exceptions to this which are started outside of CAF:

■ cfusrntf.exe is invoked transiently whenever a user logs in to a system. This is used to capture user accounts information.

■ sxplog32.exe is invoked persistently whenever a user logs in to a system. This is used to apply settings within a user context for packages created by the Software Delivery Packager SXP. It is only used when an SXP package is installed. This is started using the registry key:

    HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run\DsmSxplog

■ cf_SysTray.exe is invoked persistently whenever a user logs in to a system. This is used to provide a menu applet within the system tray area of the desktop. It is started using the registry key:

    HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Run\CAF_SystemTray

Some of the functional plug-ins remain in memory and others are started by CAF on request of the domain manager or the CAF Scheduler. For example, the Software Delivery plug-in (SDAGENT.EXT) is started with CAF, but once it has completed its check for any Software Jobs, it will exit.

Agent Registration

When a CA DSM Agent is installed on an end system, the first thing it does is attempt to register its existence with its domain manager through its scalability server. This registration process is common across all products and includes a limited amount of inventory information (known as 'Basic Inventory' or sometimes 'Basic Hardware Inventory'). By default, the agents reregister at a random time every 24 hours. After the first registration, only the inventory changes (the delta inventory) will be sent.

For some network types (for example, IDSN dial-up) it is best to increase the time interval in order to reduce call charges. This can be done from the central configuration by changing the settings of the '*Registration refresh scheduled job*' found under DSM\Common Components\CAF\Scheduler.

To manually force a re-registration of the Asset, you can use the right-click option from the systray icon or the command 'caf register all.'

The registration event is sent to the scalability server. The next time the Engine performs a 'Collect' task for that scalability server, the new Asset will be created or registration of existing Assets will be verified.

The time taken from the Agent installation to visibility of that Agent in the DSM Explorer, therefore, depends on the number of sectors and scheduling options defined for the Engine and its tasks. A number of domain manager processes are notified of the registration and may then take further action. For example, if this computer has moved to a new domain manager, then the Computer Mover process will start.

Interaction with the End User

When implementing the DSM Agent, a decision must be made regarding the level of interaction the Agent will have with the end user of that system. From our experience, the world is divided 50/50 with half the administrators saying that the Agent should be totally silent and have zero interaction with the end user, while the other half of the world believes that the user should be informed when actions are taking place on their system. In some cases, country or compliance-based laws may dictate the choice for you. A common example of this is notifying the end user when a Remote Control session is taking place.

Choosing the Right Language

While English is often the language of choice for the systems administrator, it may not be the first choice for the end user. The languages that are supported by the DSM Agent include the following:

- English (US)

- German

- French

- Japanese

- Korean

- Spanish (Traditional)

- Chinese (PRC)

Note: See the CA DSM Compatibility Matrix on http://www.ca.com/support for the latest list of languages supported by the DSM Agent.

The CA DSM Common components and individual functional components will use the chosen language, as shown in these examples:

During the manual installation of the DSM Agent you can select the language that the installation dialogs will use during the install process.

From the interactive installation, the CA DSM language (DSM_LANGUAGE) the Agent will use is automatically set to the language used to display the installation dialogs.

If you are deploying the 'All Agents' package using CA Unicenter Software Delivery or the Deployment Manager, then you can choose to deploy just the language pack that is required by that user or all language packs. The installation property DSM_LANGUAGE can be used to explicitly set the CA DSM language, but it has no effect if a suitable language package is not available.

If you deploy all the language packs but do not specify the DSM_LANGUAGE installation parameter, then the DSM Agent language is by default the system default locale, if such a language package is available.

If the language package for the CA DSM language is missing or removed, CA DSM falls back to English (ENU).

In all cases the native Operating System MUI (Multilingual User Interface) or specific language pack and Font must be installed. Otherwise, CA DSM falls back to English (ENU).

Changing the Language After Installation

The CA DSM language is set on the initial installation and cannot be changed by subsequent modifications or upgrades. The only way to change the language is by reconfiguration.

To reconfigure the DSM Agent to use a different language pack, run the **ccnfcmda** command on the agent host as follows:

```
ccnfcmda -cmd SetParameterValue -ps itrm/common/localization -pn language -v
lang
```

Replace lang with one of the following options:

lang Values	Language
CHS	Chinese (PRC)
DEU	German
ENU	English (US)
ESN	Spanish (Traditional)
FRA	French
JPN	Japanese
KOR	Korean

When the language is reconfigured, you must stop and restart CA DSM for the change to take effect.

When configuring the CA DSM language, make sure that the language package for the specified language has been installed because there is no check for availability. If no language package for the specified CA DSM language has been installed, CA DSM falls back to English (ENU).

The command to reconfigure the Agent language could be run on the target computer from either a CA Unicenter Software Delivery or CA Unicenter Asset Management job. The target systems can easily be identified by using a Computer Group built with a Query of the language setting returned by the Asset Inventory.

Common Components

Cfsystray

The **cfsystray** is the main launch point for setting the Common Agent Function (CAF) and the individual Agent functions. It can be displayed as a single icon from where right-click menu options are available for all components:

It can also be displayed as separate icons:

Out of the box, the default is to display the CA DSM systray icon. The configuration options allow for some choices. We can:

■ Display or hide the CA DSM Framework Service commands

■ Hide the system tray application

■ Display the CAF notification icon

These three options are controlled from the central configuration settings and can be found under DSM\Common Components\CAF\System Tray.

There are other configuration options available that are not visible out of the box. We can make these settings available for customization by updating the central configuration policy.

The first thing we need to do is make each system tray plug-in's trayvisible property managed. To make the plug-in property 'trayvisible' a managed property for CA Unicenter Software Delivery, CA Unicenter Remote Control, and CA Unicenter Asset Management in the default policy with a default visibility set to True, we need to perform a number of actions:

On the domain manager machine (and also enterprise manager, if applicable), create an XML file containing the following XML code:

```
<configuration>
  <allusers>
    <paramsection name="itrm">
        <paramsection name="common">
            <paramsection name="caf">
                <paramsection name="plugins">
                    <paramsection name="sdagent">
                        <parameter name="trayvisible" value="1">
                            <parameterinfo name="pi_ trayvisible">
                                <attribute name="type">bool</attribute>
                                <attribute name="desc">Load the SD system
tray plugin.</attribute>
                            </parameterinfo>
                        </parameter>
                    </paramsection>

                    <paramsection name="rchost">
                        <parameter name="trayvisible" value="1">
                            <parameterinfo name="pi_ trayvisible">
                                <attribute name="type">bool</attribute>
                                <attribute name="desc">Load the RC system
tray plugin.</attribute>
                            </parameterinfo>
                        </parameter>
                    </paramsection>

                    <paramsection name="amagent">
                        <parameter name="trayvisible" value="1">
                            <parameterinfo name="pi_ trayvisible">
                                <attribute name="type">bool</attribute>
                                <attribute name="desc">Load the AM system
tray plugin.</attribute>
                            </parameterinfo>
                        </parameter>
                    </paramsection>

                </paramsection>
            </paramsection>
        </paramsection>
    </paramsection>
  </allusers>
</configuration>
```

Ensure CA DSM is running and then execute the command:

```
"ccnfregdb -mlocalhost -f<name of XML file> -e"
```

The following screen shot shows the resulting change to the centrally managed configuration options:

Now that we have the extra settings available, we can use them to hide the main CA DSM System Tray and only display the CA Unicenter Remote Control system tray icon.

To hide the main system tray and display only the CA Unicenter Remote Control system tray

1. Create a new policy with the following properties set as shown:

    ..\DSM\common components\CAF\plugins\sdagent   trayvisible= False

    ..\DSM\common components\CAF\plugins\rchost   trayvisible= True

    ..\DSM\common components\CAF\plugins\amagent   trayvisible=False

    ..\DSM\common components\CAF\System Tray\System Tray : Hidden = False

    ..\DSM\common components\\CAF\System Tray\System Tray : Visible = False

2. Seal the policy and apply to the target computer group.

3. Recycle the CAF system tray on the end-user computer twice:

    Cfsystray stop

    Cfsystray show

    Cfsystray stop

    Cfsystray show

    Note: This could be done by writing a simple script deployed to the target computer group by either CA Unicenter Software Delivery or a CA Unicenter Asset Management job.

    The CA Unicenter Remote Control system tray icon should now be displayed.

Hiding the Agent from Add/Remove Programs

In most organizations the end user does not have the authority to install or uninstall software on their computer. However, for some users (for example, software developers) that need Admin privileges in order to perform their daily tasks, you can help reduce their

risk of inadvertently removing the DSM Agent by hiding its entries in the Add or Remove Programs list.

Hiding the Add or Remove Programs is done by marking them as System Components. This is done by adding 'SystemComponent'=dword:00000001 to each of the entries found under '[HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\Uninstall' listed in the following table:

Display Name	GUID
CA Unicenter DSM Agent + Basic Inventory Plugin	501C99B9-1644-4FC2-833B-E675572F8929
CA Unicenter DSM DMPrimer	5933CC13-52AB-4713-85DB-E72034B5697A
CA Unicenter DSM Agent + Asset Management Plugin	624FA386-3A39-4EBF-9CB9-C2B484D78B29
CA Unicenter DSM Agent + Software Delivery Plugin	62ADA55C-1B98-431F-8618-CDF3CE4CFEEC
CA Unicenter DSM Agent + Remote Control Plugin	84288555-A79E-4ABD-BA53-219C4D2CA20B
CA Unicenter DSM Agent + Data Transport Plugin	C0C44BF2-E5E0-4C02-B9D3-33C691F060EA
CA Unicenter Desktop & Server Management	C163EC47-55B6-4B06-9D03-2A720548BE86 InstallShield_C163EC47-55B6-4B06-9D03-2A720548BE86

Once the registry has been updated for all the DSM Agent components, they should no longer be displayed in the Add or Remove Programs window.

Customizing the Reboot/Logoff Process

When CA Unicenter Software Delivery or CA Unicenter Remote Control initiates a log off or a reboot of the target computer, the application framework (CAF) displays a dialog that shows a banner bitmap and informs the user of what is happening.

You can replace the default banner bitmap with one of your own by creating a bitmap image file (with file extension .bmp) 500x65 pixels in size. Store this file on disk and set the configuration policy, DSM\common\caf\general\dialogbitmap, to the pathname of the file.

When the dialog is displayed, it reads this file and displays the image.

The text displayed by the Common Application Framework (CAF) in the command line and in CAF dialog boxes can also be customized. Changing the text in this way will not cause any technical issues, but it could cause confusion if a support incident is raised based on the customized messages, because CA will not have any reference for these messages. Additionally, any changes made would need to be backed up before applying any future upgrade or posted fix, as the customizations would be lost if the resource file is updated in the upgrade or fix.

All of the text displayed by CAF is externalized in a resource file (for example, caf.enu) which by default is located in:

x:\program files\ca\Unicenter DSM\bin (up to r11.1)
x:\program files\ca\DSM\bin (r11.2)

It is necessary to select the correct resource file for the language installed.

The structure of this file is self explanatory, but the section likely to be of most interest is [cfCafDialog] which contains the strings used by the CAF Dialog.

As an example, if you wanted to change the title of the CAF Dialog displayed when a reboot is requested:

1. Locate the section [cfCafDialog]

2. Then modify the text for the parameter *IDS_REBOOTTITLE=*

No restart of CAF is necessary for the change to take effect.

It is not possible to change the resource file in the master image prior to installation. It is contained within a source cab file, so this is a post installation process. However, if you carry out an administrative installation of the agent plug-ins, then the cab files are fully

expanded. This would allow you to make these modifications prior to carrying out network installations of the agent plug-ins to target computers.

A number of configuration options are available that affect the dialog processing, such as 'Enable dialog' and 'Hide dialog on defer.' The options include:

Name	Value	Description
Caf Dialog: bitmap filename		The pathname of a bitmap file which is to be displayed on the reboot/logoff dialog. This must b
Caf Dialog: Enable dialog	True	Enable the reboot/logoff dialog.
Caf Dialog: Hide dialog on defer	False	Enable hiding the reboot/logoff dialog after clicking the defer button.
Caf Dialog: Reappearance time	60	This time (s) before reboot/logoff when the corresponding dialog, which has been hidden by po
Caf: Address change command		Script to run when caf detects a network address change
Caf: Enable remote CLI commands	2	Enable execution of remote CLI commands. 0=none, 1=only non-admin commands, 2=all comm
Caf: Reaper timeout	60	The time (s) after communicating with a worker process at which caf decides that the process h
Caf: Reboot command		Custom reboot command. When caf is asked to reboot the computer, caf can use this comman
Caf: Shutdown command		A command line to execute when caf shuts down. This is executed after all plugins except mes
Caf: Startup command		A command line to execute when caf starts up. This is executed after messaging and configura

When CAF is requested to reboot the computer, by default it uses the standard Microsoft API calls to perform this action. The command that is executed by CAF to reboot the computer can be customized if required. A common example of this would be if the target system was an ATM machine. Then we would change the 'Reboot Command' to be the application version that checks whether the ATM is in use before closing the application and rebooting.

Computer Naming

The name used to identify the Asset is the computer name as defined to the operating system. From r11, the computer name is not the unique identifier for the asset so it is possible have multiple computers with the same name listed.

Name	OS	Trust Level	Origin	AM Status	SD Status	RC Status	Scalability Server
xpclient1	Windows XP Professional	● ● ● ● ●	CA	Operational	Operational	Unknown	DSMDEMO.home.ca.com
xpclient2	Windows XP Professional	● ● ● ● ●	CA	Operational	Operational	Listening	DSMDEMO.home.ca.com

Having the same computer name for multiple assets is common in some environments. A good example of this would be a school where perhaps the Server is called 'server' and each computer is called 'PC1,' 'PC2,' and so on. This is fine, but when every school in the country has the same names for its assets, centralized administration becomes difficult because you cannot tell one 'PC1' from another 'PC1.' To overcome this problem, you can tell the Agent to use a name that you provide rather than the OS computer name.

The Name that the Agent registers with can be changed from the GUI or by command line. The command to change the Agent name is '*CAF setcomputername xxxxx*' where *xxxxx* is the name that you want. In the example of the school perhaps you would append the school number to each name, for example, '0123-PC1.'

CA Unicenter Software Delivery

The CA Unicenter Software Delivery Agent (SDAGENT) is started when CAF first starts. It starts again if CAF detects a change of IP address or when it is requested to do so by the scalability server. It is not a memory-resident task. Once it has completed its check for any Software Jobs, it will exit and only be started again on request. This process is often referred to as 'Job Check.'

The Job Check dialog can be either hidden or visible to the user, and as with the CAF dialogs, the logo and text strings can be customized.

All of the text displayed by 'Job Check' is externalized in a text file called Agent.txt which is located in:

x:\program files\ca\Unicenter DSM\sd\nls (up to r11.1)
x:\program files\ca\DSM\ sd\nls (r11.2)

The bitmap that is displayed is changed by replacing the file called agent.bmp. The file must be within a specific size range. The agent.txt file lists details of this.

Agent Privileges

The CAF Service is running as the Operating System User LocalSystem. When it starts the CA Unicenter Software Delivery Job Check process, then by default the same User is used. In most cases LocalSystem is the best user account for this process but there are some situations where a Domain Administrator account is required. An example is defining a new printer with drivers to a print server.

CAF can be configured to launch the Job Check process under a Domain Administrator account by using the command 'caf setcreds sdagent User XXXXXX Password YYYYYY' on the Agent system requiring the change.

The Domain User Account that is used must have the specific security privilege of 'replace a process level token' and the password should not be set to require regular changing. If it is necessary to change the password on a regular basis, then it is best to use two different Domain Administrator accounts and alternate between them at each password change. If a

CA Unicenter Software Delivery job is to be used to apply the configuration change, then this must be done before the password has expired!

Logon Shield

By default, the Logon Shield is not enabled by the installation of the CA Unicenter Software Delivery Agent. The Logon Shield can be used to prevent a user from logging on to the computer while a CA Unicenter Software Delivery job is active. The implementation adds the Logon Shield executable to the Operating System GINA (Graphical Identification and Authentication DLL) chain. So implementing the Logon Shield in your organization needs careful testing, particularly if multiple products that use GINA hooking are installed.

The Logon Shield is enabled by a CA Unicenter Software Delivery Configure procedure. A number of options to its use are available.

The procedure options are:

■ Disable - Disables the Logon Shield.

■ Disable and remove - Removes the Logon Shield.

■ Enable forced log off - Forces the user to log off every time a job is to be executed on the computer.

■ Enable wait until log off - Does not force the user to log off. Instead, the job remains pending until the user logs off. Once the user has logged off, the job is run.

■ Enable wait until log on - Keeps a pending job from executing until a user logs on to the computer on which the job is to be run.

■ Enable per job - Must be executed prior to specifying that you want to use the Job options flag. Prevents the user from being logged on while the job executes. Note, however, that jobs can be set up without using this Job options flag.

When the Logon Shield has been activated, a dialog informs you about the current state of the job. If you press CTRL+ALT+DEL, the dialog appears and continuously displays the progress of job execution. The dialog disappears automatically after the job has ended.

Setting up a job using one of the Enable Logon Shield configuration procedures requires a reboot of the computer in order to activate the Logon Shield. This is due to the fact that the Microsoft login gina.dll, msgina.dll, is to be replaced with sxpgina.dll. Conversely, setting up a job using the Disable Logon Shield configuration procedure for a computer also requires a reboot of the computer in order to deactivate the Logon Shield, because the disabling requires the sxpgina.dll to be replaced with the Microsoft msgina.dll.

The Disable and Remove Logon Shield procedure removes the special gina.dll that is installed to monitor and control the login process. The agent may reboot.

The Logon Shield functionality should be used with care. If the 'Force user to log off' policy is used, then the administrator should make sure that no jobs that require user interaction are scheduled for an agent. For example, if an installation is to be launched in the background as a dialog box and the user is logged off, there can be no user interaction. Because the job is running, the user cannot log on.

If the target computer is configured to follow the 'Enable forced logoff' policy and a file is open during Job Check, the related application provides a dialog to save this file. If 'Cancel' is selected in this dialog, the 'Forced logoff' procedure is terminated. The next logoff dialog appears when the next Job Check is scheduled. Meanwhile, at the DSM Explorer the state of a second job for the target computer is displayed with 'Job executing.' It is not set to 'Postponed,' since this subsequent job is waiting for execution.

If the procedure 'Logon Shield: Enable wait until logon' has been executed for a target computer, and a job is scheduled for the computer when it is logged off, the status for the job returns with 'Job execution postponed by user at the agent end.' The target computer is still not logged on when the next job is scheduled. After say 50 minutes, this job is still displayed with 'Job execution ordered,' and no 'Postponed' status is displayed for this job. Since Job status is returned only once - when the job is actually about to execute - the first job blocks the execution of the second job, until the computer is logged on.

The Logon Shield is ignored for Desktop User agents and Domain User agents (also known as User Profile agents). For example, running the procedure 'User Agent: Enable' and then trying to run the procedure 'Enable forced logoff' ends the latter procedure with 'Job Execution Error.'

User Parameters

Sometimes information is needed about the computer that is not automatically collected by the CA Unicenter Software Delivery Agent. An example could be the Department Number or Owner details for the Asset. This information can be manually entered and stored as 'User Properties' of the CA Unicenter Software Delivery Agent. The properties can be entered from the right-click, Software Delivery\Properties option of the systray icon.

The entered 'User Data' is returned to the domain manager as part of the Agent registration process. Once processed by the domain manager, it can be viewed by right-clicking an Asset in 'All Computers' and selecting properties. Please note that the field names on the domain manager are called 'User Data 1' to 'User Data 4.' It is these field names that you must use when building Queries and Reports.

The 'User Data' can also be entered by using the CA Unicenter Software Delivery Agent command line (sd_acmd.exe). For example:

```
"sd_acmd UserInfo USER=SBROWN, LOCATION=London, PHONE="+44, 1343 3454545"
    COMMENT="Finance Department""
```

If the sd_acmd is used from a script run from the Login script or run command to capture the data when the user logs onto the system, then follow it with 'caf register all' to have the data sent in advance of the daily CAF re-registration.

Running the CA Unicenter Software Delivery Agent

As previously mentioned, the CA Unicenter Software Delivery Agent is run when CAF starts, when CAF detects a change of IP address, and when CAF receives the command to start the Agent from the scalability server.

The Agent can be started manually from the systray icon or command line. The command to start the agent is 'CAF start sdagent' or 'sd_acmd jobcheck.' The sd_acmd option is best if being used from a script as it supports the '/wait' parameter. Without this, the Agent is started asynchronously.

In most situations the CA Unicenter Software Delivery Agent is loaded by request of the scalability server when it has a job for it to run. In some network configurations it is not possible for the scalability server to establish an outbound socket connector to the Agent. In these cases it is necessary for the CA Unicenter Software Delivery Agent to periodically check with the scalability server for jobs. This can be automated by using the CAF scheduler to start the CA Unicenter Software Delivery Agent.

The following screen shot shows an example of setting the Agent common properties to schedule a check of the scalability server for jobs every hour:

Defining the schedule on each individual Agent by hand is not a practical option! To add a new schedule entry to the central configuration management, you need to create an XML file containing the following XML code on the domain manager machine (and also on the enterprise manager, if applicable):

```
<configuration>
 <allusers>
  <paramsection name="itrm">
   <paramsection name="common">
    <paramsection name="caf">
     <paramsection name="scheduler">
      <paramsection name="sdagentschedule">
       <attribute name="dis_en">Run the Software Delivery Agent</attribute>
       <parameter name="commandline" value="start sdagent">
       <attribute name="dis_en">Caf Scheduler: command line</attribute>
       <parameterinfo name="pi_commandline">
       <attribute name="type">string</attribute>
       <attribute name="desc">The caf command which performs this
job</attribute>
```

```xml
        </parameterinfo>
        </parameter>
        <parameter name="enabled" value="0">
        <attribute name="dis_en">Caf Scheduler: Enabled</attribute>
        <parameterinfo name="pi_enabled">
        <attribute name="type">bool</attribute>
        <attribute name="desc">Set to true if this job is enabled</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="type" value="day">
        <attribute name="dis_en">Caf Scheduler: Type of Job</attribute>
        <parameterinfo name="pi_type">
        <attribute name="type">string</attribute>
        <attribute name="desc">Type of schedule interval. Possible values are:
day, hour and minute. You can also add a number of optional qualifiers. Add
&quot;random&quot; to run the job with a random time added to the specified
time, up to the value of randomminutes. Add random_hour to run at a random
hour during the day. Add random_minute to run at a random minute during the
hour. Add &quot;now&quot; to run the job within &quot;randomnowtime&quot;
seconds. Multiple values are separated by spaces.</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="excludedays" value="">
        <attribute name="dis_en">Caf Scheduler: Days to exclude</attribute>
        <parameterinfo name="pi_excludedays">
        <attribute name="type">string</attribute>
        <attribute name="desc">The list of days which are excluded from the
schedule. You can specify any combination of monday, tuesday, wednesday,
thursday, friday, saturday and sunday. Names are separated by
spaces.</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="excludehours" value="">
        <attribute name="dis_en">Caf Scheduler: Hours to exclude</attribute>
        <parameterinfo name="pi_excludehours">
        <attribute name="type">string</attribute>
        <attribute name="desc">The list of hours which are excluded from the
schedule. You can specify hours using the 24 hour clock. Hours are separated
by spaces.</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="hour" value="1">
        <attribute name="dis_en">Caf Scheduler: Hour</attribute>
        <parameterinfo name="pi_hour">
        <attribute name="type">int</attribute>
        <attribute name="desc">For daily schedules, this value is the hour at
which the job runs. For hourly and minute schedules it is not
used.</attribute>
        <attribute name="incl">1,2,3,4,5,6,7,8,9,0,</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="minute" value="0">
        <attribute name="dis_en">Caf Scheduler: Minute</attribute>
        <parameterinfo name="pi_minute">
        <attribute name="type">int</attribute>
        <attribute name="desc">For daily and hourly jobs, this value is the
```

```
minute past the hour at which the job runs. It is not used for minute
jobs.</attribute>
        <attribute name="incl">1,2,3,4,5,6,7,8,9,0</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="repeat" value="1">
        <attribute name="dis_en">Caf Scheduler: Repeat</attribute>
        <parameterinfo name="pi_repeat">
        <attribute name="type">int</attribute>
        <attribute name="desc">The time between repetitions of the job. This
depends on the type - e.g. for a daily job, this is the number of days
between jobs.</attribute>
        <attribute name="incl">1,2,3,4,5,6,7,8,9,0</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="randomnowtime" value="0">
         <attribute name="dis_en">Caf Scheduler: Random now time</attribute>
        <parameterinfo name="pi_randomnowtime">
         <attribute name="type">int</attribute>
        <attribute name="desc">If &quot;now&quot; is specified in the type
then this is the number of seconds within which the job runs.</attribute>
        <attribute name="incl">1,2,3,4,5,6,7,8,9,0</attribute>
        </parameterinfo>
        </parameter>
        <parameter name="randomminutes" value="10">
         <attribute name="dis_en">Caf Scheduler: Random minutes</attribute>
        <parameterinfo name="pi_randomminutes">
         <attribute name="type">int</attribute>
        <attribute name="desc">If &quot;random&quot; is specified in the job
type then a random number of minutes between 0 and this value is added to the
specified job time. This allows a job to run at &quot;fuzzy&quot; regular
intervals.</attribute>
        <attribute name="incl">1,2,3,4,5,6,7,8,9,0</attribute>
        </parameterinfo>
        </parameter>
       </paramsection>
      </paramsection>
     </paramsection>
    </paramsection>
   </paramsection>
 </allusers>
</configuration>
```

Ensure CA DSM is running and then execute the command:


```
"ccnfregdb -mlocalhost -f<name of XML file> -e"
```

The result should now be a new schedule entry that can be configured and then deployed to your Agents in the standard way.

Note: The CAF scheduler can only execute CAF commands and cannot be used as a general purpose scheduler.

CA Unicenter Remote Control

The CA Unicenter Remote Control Agent is also known as the CA Unicenter Remote Control Host (RCHOST). It is started by CAF and runs permanently on the computer. Unlike CA Unicenter Software Delivery and CA Unicenter Asset Management where the majority of tasks occur in background, the action of taking control of a remote computer happens in the foreground and therefore is in full view of the end user.

Interaction with the End User

The level of interaction with the end user during a CA Unicenter Remote Control session can vary depending on the reason for the session. If, for example, it is a technical issue that requires actions by the administrator in order to fix the problem, there could be no involvement needed by the end user. In contrast, an issue where an end user is asking for guidance on using a particular function would be very interactive.

The first decisions to make are regarding notifying the end user (or not) when a session is in progress, and in fact whether that session can even take place before the end user approves it. The configuration can be set such that only a click on 'OK' is required, or the user can be asked to enter their password. This is controlled for each Agent or Group of Agents through the central configuration management settings.

The choice to allow a secure session where the end user cannot see the actions performed by the Administrator controlling their workstation is also configured using a central configuration policy.

Enable secure control.	True	Enable Secure Control connections.

If this option is allowed, then the Administrator will be allowed to select Secure Control when establishing a remote control session.

When the Secure Session is established, a screen is displayed informing the end user and effectively hiding the display behind.

Some of the text displayed in this message, along with other text displayed to the user by the CA Unicenter Remote Control host, can be customized by editing the file called rcHostViewer.enu found in the directory ca\dsm\bin\. This is the US English file, so please select the correct language file for your installation.

So, for example, if we wanted to customize the text displayed on the secure connection window, we would edit the file as shown:

```
# ---------------------------------------------------------------------------
------
  [RCOS]
IDS_BHS_TITLE="Remote Control Session In Progress"
IDS_BHS_INFO="This computer is being controlled by an IT Administrator of
Forward Inc.\n\nWe have disabled the local display of the machine's desktop.
Your keyboard and mouse have also been disabled.\n\nIf you believe that this
Remote Control session should not be taking place then please contact IT
Security on Extension 123456. \n\nYour desktop will be released when the
Remote Control session ends."
IDS_BHS_COPYRIGHT="Copyright (c) 2006 CA. All Rights Reserved"
```

This results in the following secure connection screen:

Communication Flow

Unlike CA Unicenter Software Delivery and CA Unicenter Asset Management where 99% of the communication flow is between the scalability server and the Agent, for CA Unicenter Remote Control 99% of the communication flow is between the Host and Viewer (Agent and DSM Explorer). The communication flow for session establishment occurs between the domain manager, scalability server and Agent, but once validated the Viewer and Host communicate directly. This communication flow needs to be taken into account if any firewalls or a NAT (Network Address Translation) network is between the Viewer and Host. Please see the Firewalls and Network Considerations chapter earlier in this Green Book for further discussions on network configuration.

CA Unicenter Asset Management

The Asset Management Agent consists of three main components:

■   Asset Management Performance Agent (AMPMAGENT)

- Asset Management Software Usage Agent (AMSWSWMAGTW)

- Asset Management Agent (AMAGENT)

The first two components are started when CAF first starts and remain in memory as running processes. The third component is also started when CAF first starts, and then again if CAF detects a change of IP address, when it is requested to do so by the CAF Scheduler, or by a manual action. It does not remain in memory as a running process.

The Asset Management Agent is a wrapper process that can perform a number of actions:

- Software Inventory

- Hardware Inventory

- Custom Inventory modules

- Template Inventory

- Jobs

When the Asset Management Agent is started, it will connect to the scalability server in order to see which of the above modules need to run. If the scalability server is contactable, then it will perform the configured modules and send the results to the scalability server before exiting.

If the scalability server is not available, then it will exit and not perform any of the modules.

Scheduling

The Asset Management Agent is run from the CAF scheduler by default once every 24 hours. The schedule configuration is a centrally managed configuration set that can be

amended and applied to individual or groups of Assets. For example, you may want to run the Agent more frequently on Assets in a remote sales office compared to those on the DMZ.

As previously mentioned, the Asset Management Agent runs a number of different processes. Each process can be disabled or enabled, and can have specific scheduling rules. For example, by default the Hardware Inventory is configured to run every time, whereas the Software Inventory only runs once per day.

The CAF scheduler also invokes the Asset Management Agent every hour in HINT mode. Hint mode invokes a check to see whether any CA Unicenter Software Delivery jobs have been performed, and if so, the Software Inventory module is started. Assuming that the CA Unicenter Software Delivery job was installing or uninstalling software, and provided that a Software Signature entry was created for the changed application, then the detected software for this Asset will be updated.

Forcing a Rescan

Sometimes it may be necessary to manually force the Asset Management Agent to run and perform Hardware and Software Inventory. This can be done from the DSM Explorer by selecting a single or multiple Assets, right-clicking, and selecting 'Activate Job Check.' By default the Asset Management Agent will only report any changes in Hardware or Software since the last time it was run. By specifying the Collect option, a full re-collect of all Hardware or Software information is performed.

The CAF command line can be use to start the Asset Management Agent. If run without any arguments, it will connect to the scalability server and run any configured modules, and will obey any scheduling configuration in place for this Agent. The Asset Management Agent can be run with the following arguments:

■   /COLLECT - Do not diff, send complete information

■   /RESCAN_INVENTORY - Run any hardware inventory configurations linked to the unit regardless of scheduling

■   /RESCAN_SOFTWARE - Run any software inventory configurations linked regardless of scheduling

Examples

To force a re-scan and collect of all Hardware and Software inventory, you would run this command on the Agent:

```
CAF start AMAGENT args /RESCAN_INVENTORY /RESCAN_SOFTWARE /COLLECT
```

By the use of the HOST parameter, some CAF commands can also be run on remote machines. So to run the same command on a machine other than your LocalHost, you would enter:

```
CAF start AMAGENT args /RESCAN_INVENTORY /RESCAN_SOFTWARE /COLLECT HOST
NAME_OR_IP_OF_MACHINE
```

If the credentials you are logged on with do not have access to the remote computer, then different credentials can be passed to CAF using USER and PASSWORD. For example:

```
CAF start AMAGENT args /RESCAN_INVENTORY /RESCAN_SOFTWARE /COLLECT HOST
name_or_ip_of_machine USER admin_account PASSWORD secret
```

Preventing Jobs from Running

There are a number of business scenarios that could require you to only perform Asset Management Hardware and Software Inventory on a machine and prevent all other activities. This can be done for individual or groups of assets by using the following centrally managed configuration option:

Collecting Custom Inventory Information

The Asset Management Agent can execute a number of custom Inventory Modules, but by default these are not enabled. Some collect specific information, as their titles suggest, whereas the one called WBEM Inventory is open to configuration. You can either customize the WBEM module or use it as an example and create a new module.

The complete WBEM data source is a considerable size and if collected can be many megabytes, putting unnecessary overhead on the CA DSM system. It is strongly recommended that you restrict the information that is collected to just that which is necessary.

During the configuration process, the Launch button can be used to start the WBEM browser. This allows you to browse your local WBEM data and chose the settings that you want to collect. Only data found in your local machine is presented, but additional settings can be manually added in the configuration window.

Creating Your Own Inventory Module

Sometimes inventory data that is valuable to you exists in files, registry keys, or other locations on the machine that are not automatically read by the CA Unicenter Asset Management inventory modules. It is then necessary to create your own inventory module to read this data and write a MIF file for CA Unicenter Asset Management to process. This Inventory module could be a simple CA DSM script or a more complicated C++ or Visual Basic program.

A lot of additional inventory information on the Windows platform is made available as ActiveX components. ActiveX provides access to information from Active Directory, WMI, etc. and is very easy to access using Windows scripting. The only thing missing is a library to create MIF Files and a way to embed the Windows script into a recurring Asset Management job.

This section will discuss how you can use Visual Basic Scripting (.VBS) or Desktop Management Scripts to read information and build MIF files. If you prefer using Java

scripting (.JS), you can use the same steps as with VBS, just not for the MIF library as it is written in Basic.

Implementing .VBS

Once you have written your Visual Basic script, you can create an asset job that allows you to encapsulate the .vbs script into the job. The script will then be transported and executed on the agent through the CA DSM infrastructure.

To build an asset job

1. Find the group or individual asset that you want to schedule the job for.

2. Expand the tree, right-click the Jobs node, and select Create Asset Job.

3. From the Job wizard, select External Utility as the Task type.

4. Click Next and give the job a name and description.

5. Click Next and select cscript.exe as the Executable.

6. Enter a Working directory where the process is executed; if you leave it blank, it will point to the Unicenter DSM\bin directory.
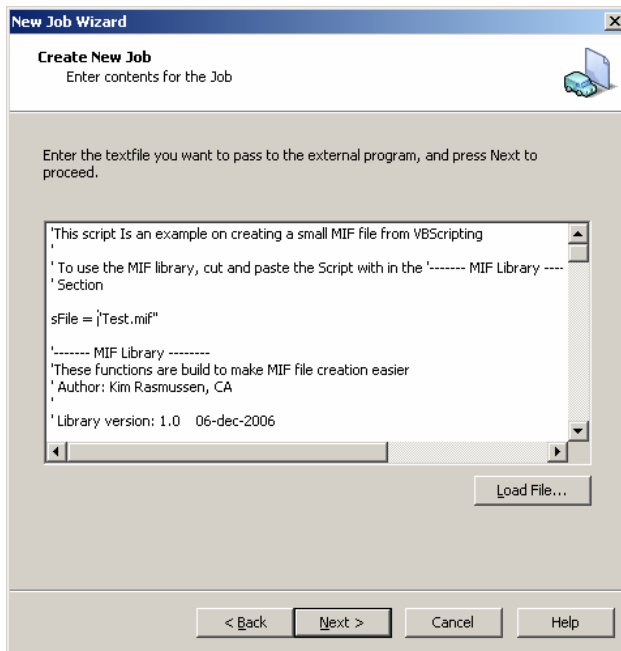
It is recommended that you use the default Agent directory, for example, C:\Program Files\CA\Unicenter DSM\Agent\units\00000001\uam, for the Working directory. In that way can you control the output folder from the CA DSM job configuration. This is to ensure the MIF is placed in the right folder to get collected.

7. Select the 'Create Text File' option and give the file a name.

   This will then create the file with the content from the next screen to that file name in the Working directory.

   Note: The file is deleted after the job has executed.

8. Click Next and insert the full .VBS script that you want to execute.

   See the example following this procedure.

9. When you click Next, you can schedule the module. If it is an inventory module, it is recommended that you run it regularly, at least once a day, to ensure up-to-date information.

Example

```
' To use the MIF library, cut and paste the Script with in the '------- MIF
Library --------' Section

sFile="test.mif"

'------- MIF Library --------
'These functions are built to make MIF file creation easier
' Author: Kim Rasmussen, CA
'
' Library version: 1.0    06-dec-2006
'
```

```
'---Library variables
Dim objMIFFileSystem, objMIFFile
Dim nGroupID
Dim nAttrID


'=======================================================================
'= CreateMIFFile function:
'=======================================================================
Function CreateMIFFile(strFile, strName, strDescription)
        Set objMIFFileSystem = CreateObject("Scripting.fileSystemObject")
        Set objMIFFile = objMIFFileSystem.CreateTextFile(strFile, TRUE)

        objMIFFile.WriteLine("Start Component")
        objMIFFile.WriteLine("  Name = """ & strName & """")
        objMIFFile.WriteLine("  Description = """ & strDescription & """")
        nGroupID = 0
        CreateMIFFile = 0
End Function


'=======================================================================
'= End MIFFile function:
'=======================================================================
Function EndMIFFile()
        objMIFFile.WriteLine("End Component")
        objMIFFile.Close
        EndMIFFile = 0
End Function


'=======================================================================
'= CreateMIFGroup function:
'=======================================================================
Function CreateMIFGroup(strGroupName, strDescription, strClass)
        nGroupID = nGroupID + 1
        objMIFFile.WriteLine("")
        objMIFFile.WriteLine("  Start Group")
        objMIFFile.WriteLIne("    Name = """ & strGroupName & """")
        objMIFFile.WriteLine("    ID = " & nGroupID)
        objMIFFile.WriteLine("    Class = """ & strClass & """")
        objMIFFile.WriteLine("    Description = """ & strDescription & """")
        nAttrID = 0
        CreateMIFGroup = 0
End Function


'=======================================================================
'= EndMIFGroup function:
'=======================================================================
Function EndMIFGroup()
        objMIFFile.WriteLine("  End Group")
        EndMIFGroup = 0
End Function


'=======================================================================
'= CreateMIFString function:
'=======================================================================
Function CreateMIFString(strName, strValue, strDescription)
        nAttrID = nAttrID + 1
```

```
        objMIFFile.WriteLine("")
        objMIFFile.WriteLine("    Start Attribute")
        objMIFFile.WriteLine("        Name = """ & strName & """")
        objMIFFile.WriteLine("        ID = " & nAttrID)
        objMIFFile.WriteLine("        Description = """ & strDescription & """")
        objMIFFile.WriteLine("        Type = String(64)")
        objMIFFile.WriteLine("        Value = """ & strValue & """")
        objMIFFile.WriteLine("    End Attribute")
        CreateMIFString=0
End Function


'========================================================================
'= CreateMIFInt function:
'========================================================================
Function CreateMIFInt(strName, intValue, strDescription)
        nAttrID = nAttrID + 1
        objMIFFile.WriteLine("")
        objMIFFile.WriteLine("    Start Attribute")
        objMIFFile.WriteLine("        Name = """ & strName & """")
        objMIFFile.WriteLine("        ID = " & nAttrID)
        objMIFFile.WriteLine("        Description = """ & strDescription & """")
        objMIFFile.WriteLine("        Type = Integer64")
        objMIFFile.WriteLine("        Value = " & intValue)
        objMIFFile.WriteLine("    End Attribute")
        CreateMIFInt=0
End Function
'------- MIF Library --------


'*********************************************--------
*********************************************
'********************************************* MAIN
*********************************************
'*********************************************--------
*********************************************

call CreateMIFFile(sFile, "Test MIF", "My first VB MIF File")
call CreateMIFGroup("MIF Group1","first group","")
call CreateMIFString("Var1","Value1","This is a String variable")
call EndMIFGroup()
call CreateMIFGroup("Group2","Secund group","")
call CreateMIFString("Var2","Value2","This is a String variable")
call CreateMIFInt("Var3",20002,"This is a Numeric variable")
call EndMIFGroup()
call EndMIFFile()
```

VBS and .MIF Files

To use VBS as an inventory module, the script needs to produce an output file in the commonly known file format .MIF. The sample script above contains commands that Build MIF: Component, Groups, String and Integer values.

If you need to use the MIF library in other scripts, you can *copy select* everything between the two Lines marked with: '------- MIF Library --------' and insert it at the top of your new script.

```
call CreateMIFFile(sFile, "Test MIF", "My first VB MIF File")
call CreateMIFGroup("MIF Group1","first group","")
call CreateMIFString("Var1","Value1","This is a String variable")
call EndMIFGroup()
call CreateMIFGroup("Group2","Second group","")
call CreateMIFString("Var2","Value2","This is a String variable")
call CreateMIFInt("Var3",20002,"This is a Numeric variable")
call EndMIFGroup()
call EndMIFFile()
```

In addition to the MIF library, this script is a small test script that contains an example in how to use the functions.

You can see from the code example that it is very similar to how the CA DSM Scripting MIF functions work.

Note: You need to do the EndGroup and EndFile yourself in the script.

The script will produce a MIF file that should look like this:

```
Start Component
Name = "Test MIF"
Description = "My first VB MIF File"

Start Group
Name = "MIF Group1"
ID = 1
Class = ""
Description = "first group"

Start Attribute
Name = "Var1"
ID = 1
Description = "This is a String variable"
Type = String(64)
 Value = "Value1"
End Attribute
End Group

Start Group
Name = "Group2"
ID = 2
Class = ""
Description = "Second group"

Start Attribute
Name = "Var2"
ID = 1
Description = "This is a String variable"
Type = String(64)
 Value = "Value2"
End Attribute

Start Attribute
```

```
Name = "Var3"
ID = 2
Description = "This is a Numeric variable"
Type = Integer64
Value = 20002
End Attribute
End Group
End Component
```

Once this MIF file has been processed by the Engine into the MDB, then it will be visible in the Additional folder under the asset's Inventory node.

Creating a .MIF File Using the CA DSM Script Language

In many cases complicated programming is not necessary. The additional inventory information can be collected using a CA DSM script to read the data and create the MIF file. One of the benefits of using the CA DSM scripting language is that it supports many different operating systems. As long as your script avoids any platform-specific OS calls, then it can be run on all operating systems that the DSM Agent supports. The CA DSM scripting language also offers advanced features regarding .MIF file manipulation.

The following example shows the CA DSM script for creating a MIF file for the additional inventory data:

```
'*******************************************************************************
'*
'*  This Script reads the Sound card on Windows NT and Windows 95 and create a MIF
file
'* Called ISND.MIF, in the Clientws
'*
'* Created by Kim Rasmussen    Computer associates.
 '* Version 1.0
'*
'*
'*******************************************************************************
***
```

```
DIM SoundCard AS STRING
DIM i AS INTEGER

'*****************************************************************************
***
'*
'* This function reds the Driver description in the Registry database
'*
'*****************************************************************************
***
FUNCTION get_soundcard(i as integer) AS STRING
  DIM driver AS STRING
  DIM dummy AS INTEGER
  DIM rhdl AS INTEGER

  IF (InStr(OsSystem,"NT")>0) THEN
    Rem read the Wave driver
    rhdl = RegOpenKey(HKEY_LOCAL_MACHINE,"SOFTWARE\Microsoft\Windows
NT\CurrentVersion\drivers32")
    IF rhdl <> 0 THEN
      RegQueryVariable(rhdl,"wave",driver,dummy)
      IF UCASE(driver)="MMDRV.DLL" THEN
        RegQueryVariable(rhdl,"wave1",driver,dummy)
      ENDIF

      RegCloseKey(rhdl)
    ENDIF

    REM Get the driver description
    IF driver <> "" THEN
      rhdl = RegOpenKey(HKEY_LOCAL_MACHINE,"SOFTWARE\Microsoft\Windows
NT\CurrentVersion\drivers.desc")
      IF rhdl <> 0 THEN
        RegQueryVariable(rhdl,driver,driver,dummy)
        RegCloseKey(rhdl)
      ENDIF
     ELSE
      driver = "N/A"
    END IF

  ELSEIF (InStr(OsSystem,"95")>0) THEN
    Rem Get the first Key in the wave group
    rhdl =
RegOpenKey(HKEY_LOCAL_MACHINE,"System\CurrentControlSet\control\MediaResources\wave"
)
    IF rhdl <> 0 THEN
      RegEnumKey(rhdl,0,driver)
      RegCloseKey(rhdl)

      Rem Read description for that driver
      rhdl =
RegOpenKey(HKEY_LOCAL_MACHINE,"System\CurrentControlSet\control\MediaResources\wave\
"+driver)
      IF rhdl <> 0 THEN
        RegQueryVariable(rhdl,"Description",driver,dummy)
        RegCloseKey(rhdl)
```

```
      ELSE
        driver = "N/A"
      ENDIF
    ENDIF
  ENDIF

  get_soundcard = driver
END FUNCTION    'Get_soundcard


'*********************************************************************************
*********************************** MAIN
*********************************************************************************
***********************************'
 IF (InStr(OsSystem,"NT")>0) OR (InStr(OsSystem,"95")>0) THEN
  SoundCard = Get_soundcard(i)

  CreateMifFile(WorkstationPath+"ISND.MIF","Multimedia","AMO Multimedia
information")
  CreateMifGroup(WorkstationPath+"ISND.MIF","Audio","Audio information","")
  CreateMifString(WorkstationPath+"ISND.MIF","Audio","Sound
Adapter",SoundCard,"Adapter information")
ENDIF
```

The .MIF commands and all the other script commands are described in the CA DSM Script
Editor online help (*Desktop Management Scripting Help*).

Chapter 15: Asset Collector

With CA Unicenter Desktop & Server Management (CA DSM) r11.2 C1, a new and useful technology called the Asset Collector was introduced.

Why the Asset Collector

With the CA DSM data area being used more and more by other CA products as a repository for discovery data, a more robust interface was needed to allow data from sources other than the CA DSM agents to be collected.

In older versions of CA Unicenter Asset Management, there was a MIF system that allowed you to register your own external assets. It was great and easy to use, but the data was stored differently from standard units, and you were only allowed to store hardware attributes.

The Asset Collector will take you to the next level and store both software and hardware data as exactly like the CA DSM agents do, for usage in products like CA Unicenter Patch Management and CA Software Compliance Manager. The only difference you will see is that the DSM Explorer displays the nodes slightly differently to indicate that you can not manage the units, and the value of the Origin field will not be 'CA':

Note that two new fields, Origin and Trust Level, were introduced in the r11.2 C1 release. They were added to help determine the data source and to gauge how accurate the data would be.

The Origin field values are also used at the DSM Explorer homepage to show the distribution of the managed computers in the database:

The Asset Collector works as an interface between a Standard XML file and the Standard Infrastructure. The Asset Collector is installed on all scalability servers. The Asset Collector will convert the XML file to the internal communications format and place it on the scalability server:

The Asset Collector Input

To feed the data through the Asset Collector into the CA IT Client Management (CA ITCM) system you will need to build an XML file and place it in a special folder on the scalability

server. The file must comply with standard XML specifications, and it is recommended that you keep it in a UTF-8 format if you are passing non-standard ASCII information.
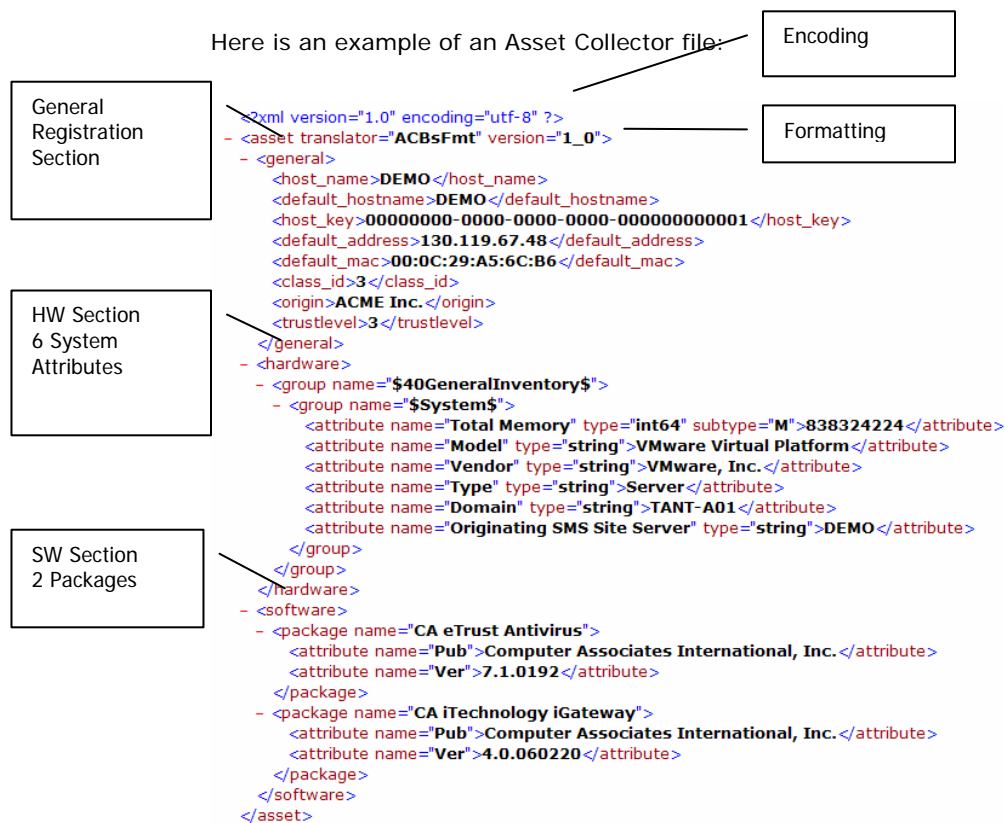
The XML File Format

The file typically consists of three XML Elements:

■ **General**: This is where all the registration data is located, including fields like Key, Origin, and Hostname. This section is mandatory and contains some mandatory fields (see below.)

■ **Hardware**: This is where you store the hardware attributes for the asset. This section is not mandatory.

■ **Software**: This is where you store the software inventory the software for the asset. This section is not mandatory.

Note: All tags must be in lowercase!

Here is an example of an Asset Collector file:

Encoding

Formatting

General Registration Section

HW Section 6 System Attributes

SW Section 2 Packages

```xml
<?xml version="1.0" encoding="utf-8" ?>
- <asset translator="ACBsFmt" version="1_0">
  - <general>
      <host_name>DEMO</host_name>
      <default_hostname>DEMO</default_hostname>
      <host_key>00000000-0000-0000-0000-000000000001</host_key>
      <default_address>130.119.67.48</default_address>
      <default_mac>00:0C:29:A5:6C:B6</default_mac>
      <class_id>3</class_id>
      <origin>ACME Inc.</origin>
      <trustlevel>3</trustlevel>
    </general>
  - <hardware>
    - <group name="$40GeneralInventory$">
      - <group name="$System$">
          <attribute name="Total Memory" type="int64" subtype="M">838324224</attribute>
          <attribute name="Model" type="string">VMware Virtual Platform</attribute>
          <attribute name="Vendor" type="string">VMware, Inc.</attribute>
          <attribute name="Type" type="string">Server</attribute>
          <attribute name="Domain" type="string">TANT-A01</attribute>
          <attribute name="Originating SMS Site Server" type="string">DEMO</attribute>
        </group>
      </group>
    </hardware>
  - <software>
    - <package name="CA eTrust Antivirus">
        <attribute name="Pub">Computer Associates International, Inc.</attribute>
        <attribute name="Ver">7.1.0192</attribute>
      </package>
    - <package name="CA iTechnology iGateway">
        <attribute name="Pub">Computer Associates International, Inc.</attribute>
        <attribute name="Ver">4.0.060220</attribute>
      </package>
    </software>
  </asset>
```

The <asset> and <general> tags are mandatory. Here is a list of the fields that can be used in the general section:

Asset Entries	Required	Description
vendor		String
serial_number		String
asset_tag		String
host_name	X	Not fully qualified (String)
host_key		Key will be generated if not supplied (string)
class_id		Platform name or code (string or number). See the *CA Unicenter Desktop & Server Management Implementation Guide* for a full description.
default_address	X	String
default_mac	X	String
default_hostname	X	String
default_subnet_mask		String
collect_time		Time in seconds (Number)
Trustlevel		Integer between 1 and 5
origin		String

The number of <general> attributes is constantly expanding. For an up-to-date list, see the *Unicenter Desktop & Server Management Implementation Guide* and the *Release Notes*.

Once you have built the collect file you need to give it a unique name, for example, the Hostname or the Internal Key. This is to ensure that the files do not get overwritten by some other process. The file extension has to be one of the following two types:

■ **.XIU** – Unsigned. The typical usage.

■ **.XIS** - Signed (See the section below.)

How to Collect the File

The Asset Collector uses three folders. They can be changed, but by default they are:

■ AssetCollectorBAK

■ AssetCollectorCollect

■ AssetCollectorOutput

The contents of these folders include the following information:

The next question would then be: How do we get the files to the Collect folder? There are many ways to do this, but here are three methods:

■ **Direct**: If your process is a converter from another data source or system, you can run the process on the scalability server. In that way, you will have access to the Collect folder. This option can be considered a safe solution as you are in control of all the application and data locally.

■ **FTP**: You could add the Collect folder as a FTP share. See the Microsoft Knowledge Base document, http://support.microsoft.com/kb/323384, for setup instructions. Once setup is complete, you can do an FTP put of the file to the server. For example:

```
ftp -s:script.txt ftp.server.com
```

The contents of script.txt could look like this, as an example:

```
Administrator
secret
bin
cd /AC
put DEMO.xiu
bye
```

■ **Share**: You could choose to share the Collect folder as a normal NetBios share. In this way, a remote node could map the drive or use Unified Naming Convention (UNC), for example, \\server\ac, to connect and deliver the file.

When using one of the network options, you will need to open up the AssetCollectorCollect folder for Write access. With that in mind, it is extremely important that you limit the access to only the users that need to deliver. You could also choose to make it a Write Only area, so that the external users cannot see or delete the files that are delivered to the collector.

Signed and Unsigned Files

If you have selected an option with an open network point to collect the data, you could choose to digitally sign the file and set the Asset Collector to only accept signed files. In this way, you can ensure that only the nodes that have the certificate installed will be collected.

If you want to sign your Collector file, you will need to install the CA DSM certificates on the machine (or if you already have the certificates there, you can use them.) Remember that the same certificate needs to be on both the sender and receiver. (See the X.509 V3 Certificates chapter later in this Green Book for more details about how to install a certificate.)

Example

This command is signing the DEMO.xiu file with the common CA ITCM Certificate (dsmcommon):

```
invsign sign DEMO.xiu dsmcommon
```

This will build a file called DEMO.xis that is ready for collection.

The signed file looks just like the unsigned file, except that the signed file has a key after the last <asset> Tag:

```
                </package>
        </software>
</asset>dsmcommonHª
0Ó¸] ;ÈôÏ±E ©Žyr]æœü Í÷3t³'>Ê´,Ð,  üðªŠê ]øà,  ± uÙr þ ,Ä¦ìù öhæ"Ôó j[£[J•Â'³C*Î˜ãvQ
•²z› -Grv
ÔÌ±¯ó4¬>!†á/ ZLÅ  a ©ë Öø¥_;ÎÂn"   €            ¥ˆ| -Ç
```

Configuring the Asset Collector

As with other CA DSM components, all settings for the Asset Collector are controlled through common configuration policy. Everything can be found under the Configuration\ Configuration Policy\Default Computer Policy\ DSM\Scalability Server\Asset Collector node:

Most of the settings are self-explanatory. Normally, you should not change any of these settings, with a few exceptions:

■ **Collection Folders**: If you for some reason you need to collect from multiple folders, you can specify multiple folders in a comma separated list. Alternatively, you could use **Recursive Folder Monitoring** and have multiple folders under the default 'AssetCollectorCollect' folder.

■ **Generate Deltas**: If you are on a fast network infrastructure, you could disable the deltas between the scalability server and the domain manager.

■ **Signed Files Only**: In either a secure environment or an open environment, you might want to only accept signed files to ensure that the sender is an approved source.

Usage Examples

With the new Asset Collector, writing your own inventory collectors just got much easier. The open interface provides a lot of different ways of reporting data. Following are some scenarios that could benefit from using the Asset Collector.

Writing Your Own Agent

The first thing that springs to mind is… I can write my own inventory agent. Typically, this would be for machines where CA does not provide an inventory agent, or for devices that do not have an OS. We have in the past seen the need for inventorying things such as vending machines, gas pumps, and cars, just to mention a few.

All you will need to do is to get the collector file created and then transported to the Collect folder (see methods above). Once this is passed on to the Asset Collector, your data will be treated just like any other inventory data.

Integration to an Existing Inventory System

Another important job for the Asset Collector is to help integrate information from other Inventory systems. Typically, you will write a small program to extract that data and build a collector file. One thing to remember in this scenario is that you need to create one file per each asset record in the system you are importing from.

Taking Inventory Offline

With the technology building on XML files, you would be able to generate offline inventory information on a USB or network disk and then copy this to the Asset Collector for importing and registering into your CA ITCM system. There are multiple initiatives in progress to create inventory based on this technology. See Product Announcements and Release Notes from CA in the future.

Why One Repository?

The huge advantage of using one repository for all of your inventory data is that inventory consumers then have only one place to look for data. CA products like CA Asset Intelligence, CA Unicenter Asset Portfolio Management, CA Unicenter Service Desk, and CA Software Compliance Manager would in one strike be able to use data from any source, as long as it had been collected by the Asset Collector.

So by using the Asset Collector, you will have one consolidated data source of all the assets you have discovered.

The Microsoft Systems Management Server (SMS) Connector

One of the first integrations created by CA that utilizes the Asset Collector is a connector to Microsoft Systems Management Server (SMS/SCCM). In general, the integration uses an Engine job that connects to the SMS site server on a scheduled basis. It exports all the asset data and builds Asset Collector files, one per each computer that is managed by SMS.

Setting up the SMS Connector

As mentioned before, you need to set up an Engine job to get the integration with SMS working. During the setup of the Engine job, using the New Task Wizard, you will first need to select the 'CA Asset Converter for Microsoft SMS' task type.

Once the Task Type has been selected and you have given the job a name, you will be asked for the SMS Service name, the SMS Site name (similar to the CA DSM domain), and the user credentials to get to the SMS information:

The Communications between the Engine and the SMS server is WMI, which is using the NetBIOS TCP port 135.

The next wizard page (Select information to collect) allows you the select whether you want to collect Hardware and/or Software:

Additionally, you can specify a scope using SMS asset groups. This allows you to limit the import to only a subset of the assets. Scopes are typically used if you need to divide the data into two CA ITCM Domains, or if you only need to manage some of the assets (typical for a service provider).

The last of the SMS Connector-specific wizard pages (Specify data delivery location) allows you to set up where the Asset Collector files will be placed:

Typically, the engine is running on a domain manager or a scalability server where the Asset Collector already is installed. The SMS Connector task will then read the Asset Collector configuration to find the right folder.

If you want the file in a different folder, or if the engine is running on a separate box, you can manually specify the folder. Remember that this will be relative to the Engine.

Note: If you are planning the have the Asset Collector remote from the Engine, then please note that the Engine by default runs as 'Local System.' In this case the Asset Collector share needs to be null session shared. For more information on Null Session Shares, see this Microsoft Knowledge Base Document: http://support.microsoft.com/kb/289655.

The Output

Once the SMS job has run, all the SMS assets are registered in CA DSM like any other assets.

Chapter 16: Security

Introduction

When we refer to the security aspects of CA Unicenter Desktop & Server Management (CA DSM), we can generally include four areas. They are:

■ Authentication

Authentication provides confidence that the requesting object is who it says it is. It is employed for login and machine-to-machine communications (application to application). Authentication is provided through operating system-specific mechanisms or X.509V3 certificate public/private key mechanisms.

■ Authorization (Permissions)

Authorization is about granting appropriate access to CA DSM data through class, object, and area permissions. Components use the verified identities to lookup rights and privileges to apply to the requested operation.

■ Encryption

Encryption is about keeping sensitive information secure, for example, encrypting remote control sessions to provide confidentiality and integrity.

■ Auditing

Auditing is about being able to audit operations or security violations that have or may occur in your CA DSM environment.

The standard product documentation for CA DSM contains a wealth of information about these areas, which we will not repeat here. We will focus on discussing some real world usage related to setting permissions and also X.509 certificate replacement and distribution. We will also describe some of the concepts.
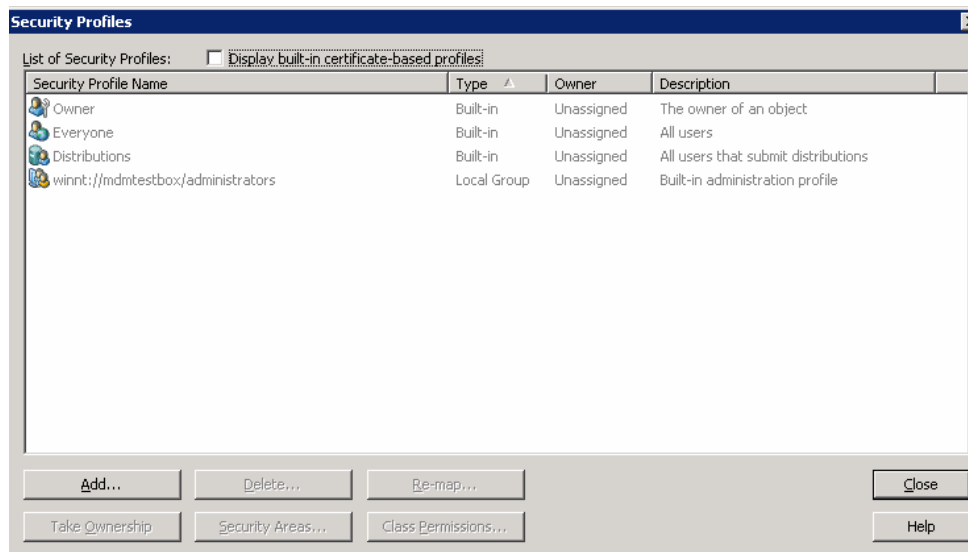
Permissions Model

The CA DSM permissions model is made up of the following components:

■ Security Profiles

Security profiles are roles, groups, or individual users mapped from a security provider (for example, an NT domain group). A security profile is an operating system user account or group in the domain manager (a local profile) or in its network domain (a domain profile). The security subsystem in CA DSM supports multiple security profiles. A security profile is either built-in (that is, created at installation time) or user-defined.

Security profiles ultimately allow the user to log in to the system. The authorized rights that user has once they are logged in are made up of the following types of permissions: Class, Object, and Area.

Note that a user can be a member of multiple profiles, through a group or individual user membership. If this is the case, then the cumulative permissions will be OR'd together to work out the resultant permissions. This means that the No Access permission does NOT override all other permissions. If a user has no access set through a group membership, but full control through another group membership, then the permissions will be cumulative, meaning the user will ultimately have full control. Also note that there is no replication of security profiles or permissions from an enterprise manager to a domain manager.

The above dialog box (retrieved from the DSM Explorer Security\Security Profiles menu) lists the default security profiles presented to you upon a fresh installation. You will see three profiles with a type of 'Built-in' and one with a type of 'Local Group'.

The 'Built-in' profiles are:

> Owner

  This profile is necessary for managing dynamically created objects so that every object will be assigned to an owner profile. Ownership may be changed using an application, such as a UI.

> Everyone

  This profile initially has the class permission NO_ACCESS for all classes.

> Distributions

This profile is essentially used when you have an enterprise. It has the class permission FULL_CONTROL for all objects.

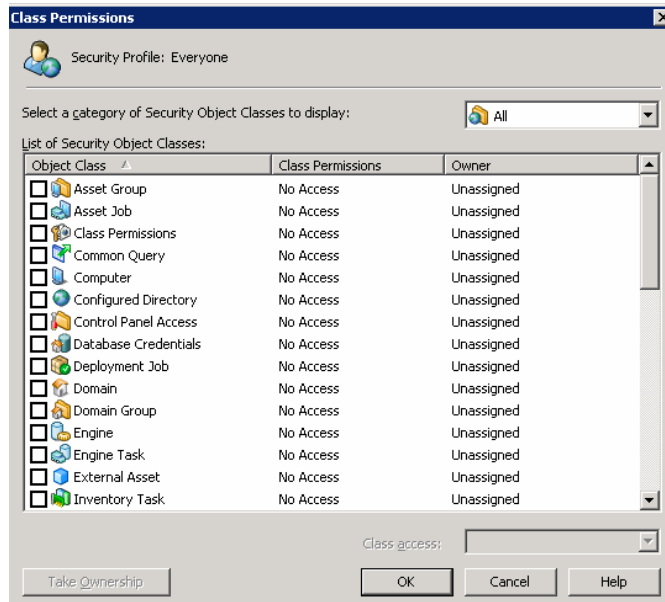And the Local Group profile is:

■ Local Administrators group

This profile initially has the class permission FULL_CONTROL for all objects.

Note: These default profiles *cannot* be deleted.
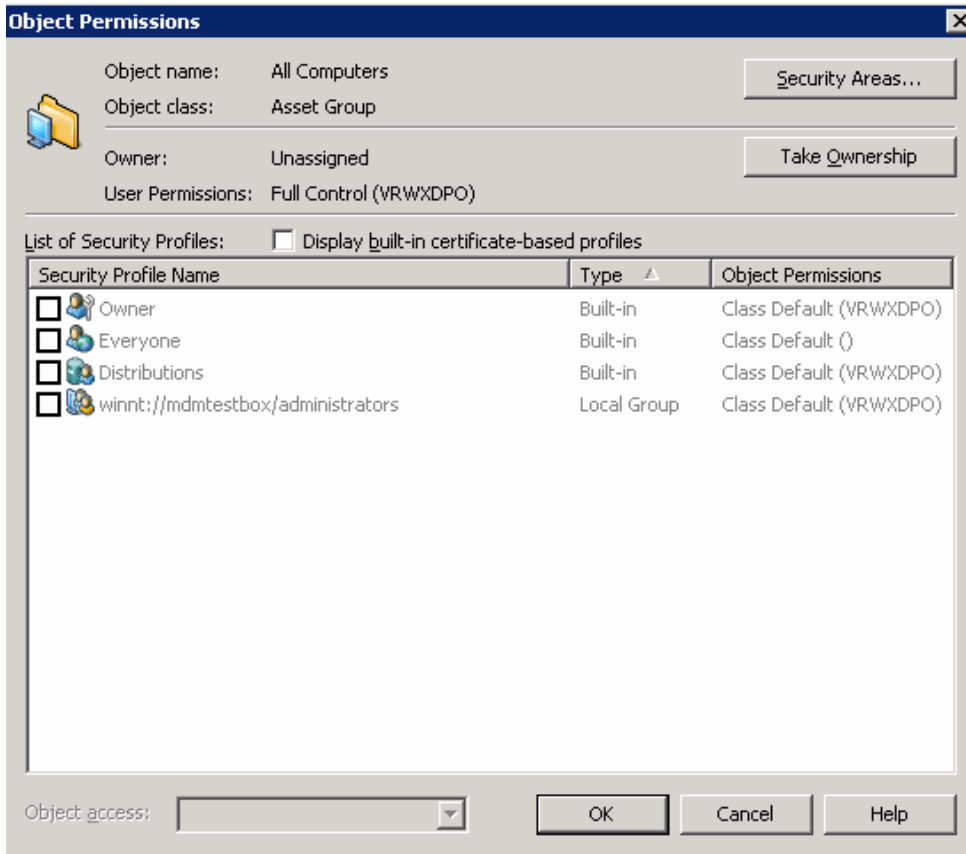
■ Class Permissions (Class-level security)

Class permissions are the default access rights that you can specify. This means that the class-level permissions that are in place will be used as the default rights for each object that is created as an instance of the Class. Think of this as the highest level of security (in terms of being less granular) that you can set. It can be sufficient to meet your needs in many situations.

Note: No Access is the default class permission for all new profiles that you create.

■ Object Permissions (Object Access Control Element)

Object permissions allow you to be more granular with your permissions. This means that you can, for example, give someone read access to all computers and groups through class permissions, but you may wish to make one or more specific objects (computers) not visible to them. This is what object-level permissions would enable you to achieve.

Let us discuss some of the things we can see in the above example:

■ Owner

An object is either owned by a user or is said to be unassigned. When a background process creates an object, that object is usually unassigned. If you create an object, your user ID is the owner of the object. As the owner of an object, you inherit all permissions associated with the Owner security profile. Owner is a special security profile, created at installation time and set to Full Access by default.

■ Take Ownership

A user gets the ownership of object.

- Object Permissions

    Permissions are based on the concept of an Access Control Entry (ACE), which is a bit-oriented integer. The following eight bits are currently used:

    V - View bit, allows you to show objects.

    C - Create bit, allows you to create objects.

    R - Read bit, allows you to read sub-objects of an object.

    W - Write bit, allows you to change an object.

    X - Execute bit, allows execution, depends on object type.

    D - Delete bit, allows you to delete objects.

    P - Permission bit, allows you to change the ACE itself.

    O - Ownership bit, allows you to take ownership of an object.

- Security Area Permissions

    A security area is a geographical, organizational, or topological division. A security area can be linked to one or more security profiles and one or more objects. A user can access an object if at least one security area linked to the object is also linked to at least one security profile of the user.

For important use cases and the descriptions of what area support is doing in the context of these use cases, see the section 'Security Area Support Use Cases' in the *Unicenter Desktop & Server Management Implementation Guide*.
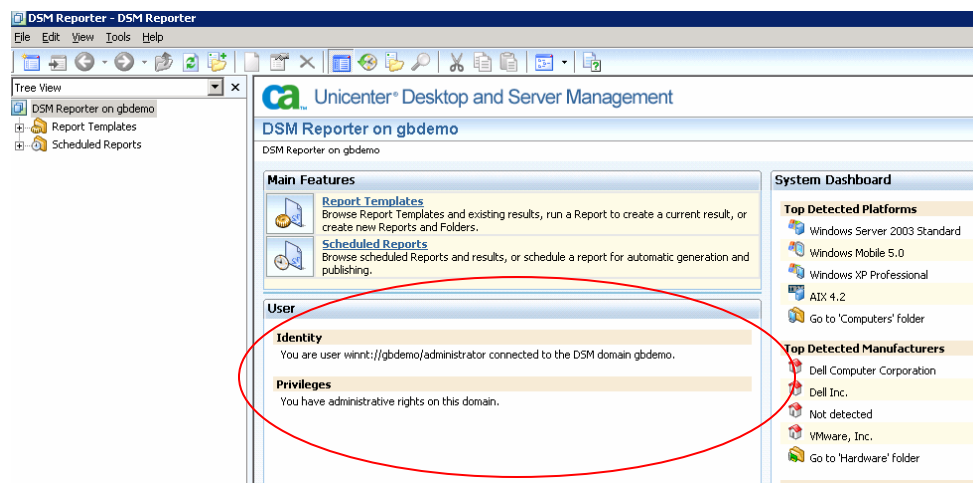

Applicable MDB Tables

Following are the Management Database (MDB) tables used by security with respect to class, object, and group level security:

■ Class-level permissions are maintained in the ca_class_ace table

■ Object-level permissions are maintained in the ca_object_ace table

■ Group-level permissions are maintained in the ca_group_ace table

Security Profiles and the DSM Reporter

Like the DSM Explorer, the DSM Reporter runs using the current Windows logged-on user profile (unified logon). Note that the user profile currently used, and the DSM Reporter privileges for this user, can be seen on the DSM Reporter front page, as illustrated in the User portlet below:

The DSM Reporter operates with two security classes: Report Templates and Report Scheduling. For these two classes, the permissions can be set individually for each user or

group of users. This description will focus on the Report Templates class, but the available permissions for Report Scheduling are used in the same way.

The minimum access level needed for reports is read access. If you do not have read access, the tree view with folders and templates will still be visible. But trying to access a report template will result in a message within the DSM Reporter informing you of your lack of security rights.

Here is a breakdown of the different Class Permissions and the resultant access:

Read

To access individual report templates, you need read access, as illustrated in the following security profile Class Permissions dialog box.

When having only Read access to report templates, vital functionality will be disabled in the DSM Reporter application. You will not be able to create reports or folders, and you will not be able to run reports.

Write

Write access is required to store new reports and folders in the database, to modify existing contents, and to remove (delete) contents from the database.

Create

The Create flag needs to be enabled to create new reports and folders (write access is also required, see above). Lack of this permission will result in New Folder and New Report options being unavailable (grayed out), as illustrated under Tasks below:

Execute

To execute reports and create new results, execute permissions are required. Lack of these permissions will result in the Run Report option being unavailable in menus and views, as illustrated in the Tasks portlet on the following screen shot. In addition, in order to schedule reports, engine permissions are required.

Recommended Settings

Typically, the advanced DSM Reporter user will need full access to report templates. This is required to be able to use the built-in reports, as well as to create new reports addressing a specific problem or to deliver information on request.

It is possible to create a limited user with only read and execute access, which will only allow that user to run existing reports.

By removing access to the report templates (No Access), or removing the Read security flag, the user will be prevented from doing any actions on report templates and will also be prevented from viewing the details of the templates. The main tree in the left-hand pane will still be visible.

Recommended Settings for Scheduling

If you wish a user to be able to schedule reports, Full Control should be provided to Report Scheduling. Scheduling also requires Write access to the engine security class, as scheduled reports need to be assigned to an engine. In turn, the engine will run the report at the specified time or at regular intervals.
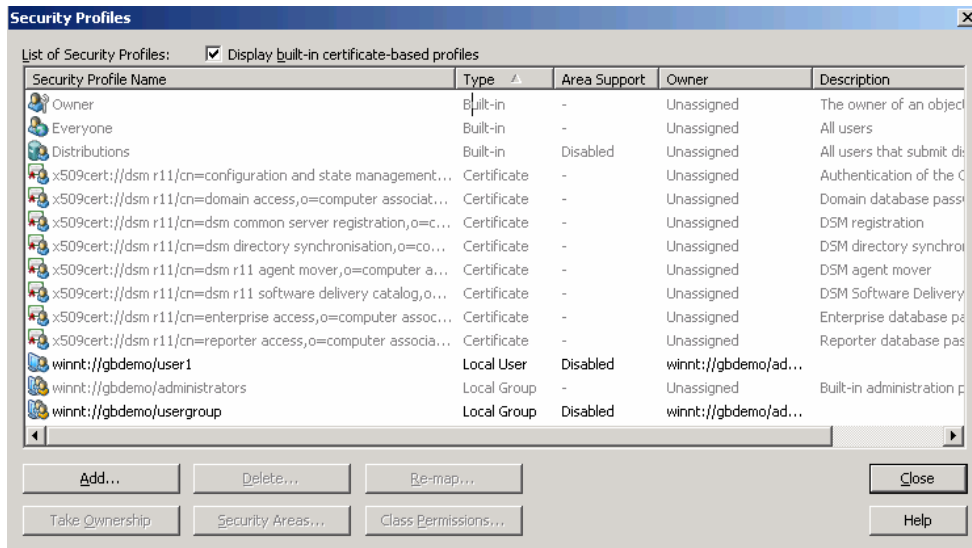
X.509 V3 Certificates

CA DSM uses X.509 certificates for authentication between its client processes and any service which requires authentication. Whenever a CA DSM client process connects to a CAF (Common Application Framework) plug-in that requires authentication, the client process must pass security credentials relevant to the target service's security requirements. Where the client process is running as an autonomous process, such as a Windows NT service or a UNIX daemon, the client process may authenticate using X.509 V3 certificates in the absence of any user credentials.

An X.509 certificate for CA DSM authentication comprises a set of attribute-value pairs packaged together with the public encryption key of an asymmetric key pair. The certificate is digitally signed and sealed by a root certificate. The certificate records the name of the subject to whom the certificate was issued, the issuing certificate authority name, and expiry information. The subject name is often referred to as the Distinguished Name (DN). The subject name is mapped to a Uniform Resource Identifier (URI) in the x509cert namespace, such as the following:

```
x509cert://dsm r11/CN=Basic Host Identity,O=Computer Associates,C=US
```

A CA DSM installation comes with a set of default standard certificates signed by a CA root certificate. The public root certificate is installed on every node within the enterprise. Using public key cryptography, clients will authenticate themselves to scalability servers upon request. A scalability server can then use the certified identity to perform subsequent authorization checks and commit audit records. The DSM Explorer enables certificate URIs to be seen as security profiles.

Security Profiles

List of Security Profiles: ☑ Display built-in certificate-based profiles

Security Profile Name	Type △	Area Support	Owner	Description
Owner	Built-in	-	Unassigned	The owner of an object
Everyone	Built-in	-	Unassigned	All users
Distributions	Built-in	Disabled	Unassigned	All users that submit di⋯
x509cert://dsm r11/cn=configuration and state management⋯	Certificate	-	Unassigned	Authentication of the C
x509cert://dsm r11/cn=domain access,o=computer associat⋯	Certificate	-	Unassigned	Domain database passⁿ
x509cert://dsm r11/cn=dsm common server registration,o=c⋯	Certificate	-	Unassigned	DSM registration
x509cert://dsm r11/cn=dsm directory synchronisation,o=co⋯	Certificate	-	Unassigned	DSM directory synchroⁿ
x509cert://dsm r11/cn=dsm r11 agent mover,o=computer a⋯	Certificate	-	Unassigned	DSM agent mover
x509cert://dsm r11/cn=dsm r11 software delivery catalog,o⋯	Certificate	-	Unassigned	DSM Software Delivery
x509cert://dsm r11/cn=enterprise access,o=computer assoc⋯	Certificate	-	Unassigned	Enterprise database pa
x509cert://dsm r11/cn=reporter access,o=computer associa⋯	Certificate	-	Unassigned	Reporter database pas
winnt://gbdemo/user1	Local User	Disabled	winnt://gbdemo/ad⋯	
winnt://gbdemo/administrators	Local Group	-	Unassigned	Built-in administration p
winnt://gbdemo/usergroup	Local Group	Disabled	winnt://gbdemo/ad⋯	

Add…   Delete…   Re-map…   Close

Take Ownership   Security Areas…   Class Permissions…   Help

There are three types of certificates in use:

■ Root Certificate

A root certificate is used to sign and validate all other certificates in the CA DSM authority. Without this we cannot validate that the other certificates are correct. The root can also be referred to as the trusted third party.

■ Basic Host Identity

The basic host identity is used by the low-level messaging components to enable a basic level of trust and setup of encrypted data channels, such as an Agent talking to a scalability server or DSM Explorer to domain manager. For example, before you have even logged into an Explorer this certificate has been checked. Every CA DSM node has a certificate that provides Basic Host Identity (BHI) installed by default.

■ Application Specific

Application certificates are used by various components of the CA DSM application to authenticate database access or server access. These certificates are used to control some of the application-level functions such as the software delivery catalog. Application certificates define which processes are allowed to receive the database credentials which run in the context of the machine, so we need a form of validation.

Other examples are common server registration of new assets, and the directory sync job. These processes access the MDB and use database credentials; without the right certificates, they do not get access. The DSM Reporter also utilizes an application certificate, as it has direct access to the MDB. When it starts, it asks the domain manager or enterprise manager for database credentials. This is allowed or not, depending on whether the certificate is valid or not. Another example is the CA Unicenter Software Delivery Agent mover process which is used when roaming between domains. Without the right certificate, the agent move process will not work.

If you have business requirements that require you to harden your CA DSM infrastructure, we recommend that you create and deploy your own root certificate, Basic Host Identity (BHI) certificate, and application-specific certificates.

For more information on certificates see the *Unicenter Desktop & Server Management Implementation Guide.*

Certificate Deployment

There are generally two scenarios regarding certificate deployment in your environment. Since the use of certificates is something that should be aligned with your company security and business practices, it is a recommended best practice to change the certificates *before* you actually deploy CA DSM. This means that you can update your installation images with the new certificates, and from then on your new infrastructure and agents will be deployed automatically utilizing the new certificates.

However, you may encounter a second scenario in which you have already deployed CA DSM and, due to a change in policy, you now need to change the certificates that are being used. We will discuss both these scenarios.

Note that the *Unicenter Desktop & Server Management Implementation Guide* contains details about certificate deployment and creation. This section complements that material.

Our first, and simplest, scenario is that the CA DSM application has not yet been installed. You can start by copying the contents of the CA DSM DVD media into a writeable directory, into which you can copy the new certificates.

Certificate Creation

To create new certificates you will utilize the cacertutil command line utility. It is necessary to install at least one CA DSM component (such as the DSM Explorer or DSM Agent) on a computer to gain access to this utility. The cacertutil utility can then be found in the \bin folder under the CA DSM installation directory.

Over the next few pages we will describe the process of creating new certificates and adding them to your installation media. We will do this with a descriptive approach and then reference the discussion in the form of a step 1, step 2 approach for easier reference.

Note: Any new certificates you create MUST have the same name as the original certificates, for example, itrm_dsm_r11_root. The reason for this is that the product installer is specifically looking for these certificate files during installation time.

The first step is to create a new set of certificates for our fictional company, Forward Inc. Start with a temporary directory on the system that has cacertutil installed. In our example we will be using c:\certs.

Replacing certificates involves copying the new files to the installation image. To make things simpler when creating your replacement certificates, you can consider creating a

batch file that will copy the files to the media master image. Some examples will be referenced in this document.

Create Certificates Batch File

```
cacertutil create -o:itrm_dsm_r11_root.p12 -od:itrm_dsm_r11_root.der -
op:rootpassword "-s:cn=newdsmroot,o=forwardinc,c=us" -d:3650
cacertutil create -o:basic_id.p12 -od:basic_id.der -op:password -oe "-
s:cn=bhic,o=forwardinc,c=us" -d:730 -i:itrm_dsm_r11_root.p12 -ip:rootpassword
cacertutil create -o:ccsm.p12 -od:ccsm.der -op:password -oe "-
s:cn=csm,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -ip:rootpassword
cacertutil create -o:itrm_dsm_r11_cmdir_eng.p12 -od:itrm_dsm_r11_cmdir_eng.der -
op:password -oe "-s:cn=dirsync,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12
-ip:rootpassword
cacertutil create -o:itrm_dsm_r11_agent_mover.p12 -od:itrm_dsm_r11_agent_mover.der
-op:password -oe "-s:cn=sdmover,o=forwardinc,c=us" -d:1095 -
i:itrm_dsm_r11_root.p12 -ip:rootpassword
cacertutil create -o:registration.p12 -od:registration.der -op:password -oe "-
s:cn=reg,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -ip:rootpassword
cacertutil create -o:babld.p12 -od:babld.der -op:password -oe "-
s:cn=babld,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -ip:rootpassword
cacertutil create -o:dsmpwchgent.p12 -od:dsmpwchgent.der -op:password -oe "-
s:cn=entaccess,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -
ip:rootpassword
cacertutil create -o:dsmpwchgdom.p12 -od:dsmpwchgdom.der -op:password -oe "-
s:cn=domaccess,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -
ip:rootpassword
cacertutil create -o:dsmpwchgrep.p12 -od:dsmpwchgrep.der -op:password -oe "-
s:cn=repaccess,o=forwardinc,c=us" -d:1095 -i:itrm_dsm_r11_root.p12 -
ip:rootpassword
cacertutil create -o:itrm_dsm_r11_sd_catalog.p12 -od:itrm_dsm_r11_sd_catalog.der -
op:password -oe "-s:cn=sdcat,o=forwardinc,c=us" -d:730 -i:itrm_dsm_r11_root.p12 -
ip:rootpassword
```

The following parameters are used in the above example:

-o	This specifies the output filename for the PKCS#12 packaged certificate.
-od	This specifies the output filename for the DER encoded certificate.
-op	This specifies a passphrase (password) used to encrypt the PKCS#12 certificate file.
-s	This specifies the subject of the certificate, such as the DN.
-d	This specifies the lifetime of the certificate in days. Here our root certificate is set for 3650 days/10 years. It is a good practice to make the BHI certificate valid for a lesser period of time than the Root Certificate, so in this example they have been set to 730 days/2 years. The application certificates have also been set to 1095 days/3 years. This is because these application certificates are normally more protected than the BHI certificates, as the application certificates are usually installed on infrastructure (except for the CA Unicenter Software Delivery catalog).
-oe	This generates a random encrypted version of the passphrase (password) used to decode the certificate and outputs it to the console. **Note**: All certificates are created with the -oe switch. This outputs an encrypted form of the password for subsequent use in the cfcert.ini file. The encrypted password output just gives a higher level of confidence that the certificates cannot be misappropriated, as the password is not in clear text.

The DN names are not specific and can be changed (CN=xxx,o=yyy etc.). You can choose whatever is most descriptive for your purposes here.

When you run the batch file to create new certificates, we recommend that you consider piping the output to a file so that you can keep the resultant outputted encrypted passwords. They will subsequently be used within your cacertutil import command line.
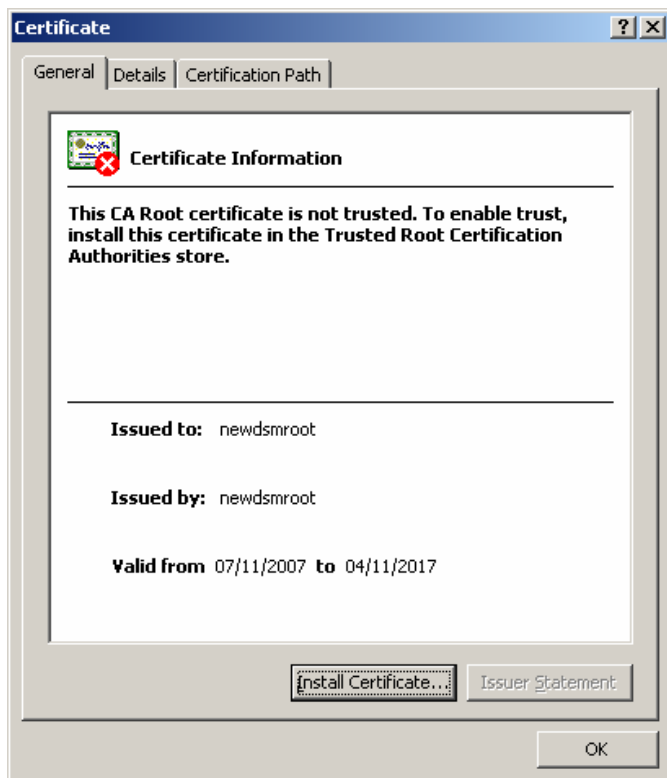
so: C:\certs\create_certs.bat > pwds.txt

Successful completion of the batch file will result in a new set of certificates, made up of DER files (Distinguished Encoding Rules) and p12 files (pkcs#12 public key cryptography standard) created in your temporary directory, along with the output file (pwds.txt):

The output file will contain multiple entries similar to those shown below. Remember that the encrypted passphrase will be re-used later.

```
C:\certs>cacertutil create -o:basic_id.p12 -od:basic_id.der -op:password -oe "-
s:cn=bhic,o=forwardinc,c=us" -d:730 -i:itrm_dsm_r11_root.p12 -ip:rootpassword
Encrypted pass phrase:
enc:AhRau5Xy0ukms5NfbrHOpnbXfaZytytk
...Operation OK.
…
```

Once you have created your own specific certificates, you can also perform a quick test to verify they have been created correctly (dates, DN etc). Simply double clicking one of the DER files such as itrm_dsm_r11_root.der will result in the following dialog:

Double clicking one of the P12 files gives you an easy way to ensure the passphrase you have utilized is valid:

Click through a few steps in the wizard, and then enter your unencrypted passphrase here. If it is accepted, you can then CANCEL this dialog.

Your next step is to edit the CFCERT.INI file. Ultimately this will allow the CA DSM installer to be able to install your new certificates. Note that CFCERT.INI exists on the installation media in multiple locations. You need to copy only one instance of the file to your temporary certificate location (c:\certs, for example) so that you can perform the necessary

edits. A section of the default cfcert.ini file contents are listed below. The full contents are described in the *Unicenter Desktop & Server Management Implementation Guide*.

```
[CAF]
files=itrm_dsm_r11_root.der,basic_id.p12
[Configuration]
files=ccsm.p12
[Files]
itrm_dsm_r11_root.der=cacertutil import -i:itrm_dsm_r11_root.der -it:x509v3
basic_id.p12=cacertutil import -i:basic_id.p12 -
ip:enc:uAa8VNL4DKZlUUtFk5INPnr2RCLGb4h0 -h -t:dsmcommon
ccsm.p12=cacertutil import -i:ccsm.p12 -t:csm -
ip:enc:IWhun2x3ys7y1FM8Byk2LMs56Rr8KmXQ
[Tags]
dsmcommon=x509cert://DSM r11/CN=Generic Host Identity,O=Computer Associates,C=US
csm=x509cert://dsm r11/CN=Configuration and State Management,O=Computer
Associates,C=US

etc
```

Each section of the cfcert.ini file declares the certificates that are required to be installed by the associated installer. The installer reads the 'files=' entry from its associated section in cfcert.ini and installs each certificate listed in turn by using the command located in the [Files] section of the cfcert.ini file.

For example, the common application framework (CAF) installer finds that it needs to install the certificates itrm_r11_dsm_root.der and basic_id.p12. In the [Files] section, the CAF installer finds the cacertutil commands associated to these certificates in the first two lines, and executes these commands.

The [Tags] section allows you to create new certificates which do not use the standard certificate URIs. When installing a CA DSM manager node, the installation components will read this section and set up security profiles for the named URIs. By convention, the file names listed in the 'files=' entry in each section of cfcert.ini are the same as the names of the underlying certificate file. This allows for easier maintenance of the cfcert.ini initialization file.

Since we are going to replace the default certificates with our new certificates, we need to change the following:

1. [Files] section to reflect the new passphrase (password)

2. [Tags] section to reference our new DN URI

The [Files] section will resemble this after editing is complete:

```
[Files]
itrm_dsm_r11_root.der=cacertutil import -i:itrm_dsm_r11_root.der -it:x509v3
basic_id.p12=cacertutil import -i:basic_id.p12 -
ip:enc:AhRau5Xy0ukms5NfbrHOpnbXfaZytytk -h -t:dsmcommon
ccsm.p12=cacertutil import -i:ccsm.p12 -t:csm -
ip:enc:E4H1xo9u2z4U9O775lZQVnZRQuDmJ4XF
itrm_dsm_r11_cmdir_eng.p12=cacertutil import -i:itrm_dsm_r11_cmdir_eng.p12 -
ip:enc:ewKioCf4EEnQbvr0NJBA16qAduWn3r27 -t:dsm_cmdir_eng
itrm_dsm_r11_sd_catalog.p12=cacertutil import -i:itrm_dsm_r11_sd_catalog.p12 -
ip:enc:SHkCxBQyt7KHPCwmU08KwM43fE3MOq9y -t:dsmsdcat
```

```
itrm_dsm_r11_agent_mover.p12=cacertutil import -i:itrm_dsm_r11_agent_mover.p12 -
ip:enc:HnkJ2gLBGpl8D3rOz02UJQE11wibLhpj -t:dsmagtmv
registration.p12=cacertutil import -i:registration.p12 -
ip:enc:U2z8skmxRzdjaAVQDMWwYTHv3LJbtFvo -t:dsm_csvr_reg
babld.p12=cacertutil import -i:babld.p12 -ip:enc:0H9c4YD4yb1xPR8Q3SSpr7EBzLAY9j84
-t:babld_server
dsmpwchgent.p12=cacertutil import -i:dsmpwchgent.p12 -
ip:enc:fBYRa9OPtZDMaVTi1FDCm5uw97Dnj7rk -t:ent_access
dsmpwchgdom.p12=cacertutil import -i:dsmpwchgdom.p12 -
ip:enc:qhDmdQAeEPaUiVuVmuFaHrlAH5IlVNGX -t:dom_access
dsmpwchgrep.p12=cacertutil import -i:dsmpwchgrep.p12 -
ip:enc:QH0qnofjHDO24PmWwBAfQTvu81gOXZzm -t:rep_access
```

What we have changed here is the enc:<encryptedpassphrase> section. We have copied
the resultant encrypted passphrase from the output file we created earlier (pwds.txt).

Next, we need to edit the [Tags] section as follows:

```
[Tags]
dsmcommon=x509cert://DSM r11/cn=bhic,o=forwardinc,c=us
csm=x509cert://dsm r11/cn=csm,o=forwardinc,c=us
dsm_cmdir_eng=x509cert://dsm r11/cn=dirsync,o=forwardinc,c=us
dsmsdcat=x509cert://dsm r11/cn=sdcat,o=forwardinc,c=us
dsmagtmv=x509cert://dsm r11/cn=sdmover,o=forwardinc,c=us
dsm_csvr_reg=x509cert://dsm r11/cn=reg,o=forwardinc,c=us
babld_server=x509cert://dsm r11/cn=babld,o=forwardinc,c=us
ent_access=x509cert://dsm r11/cn=entaccess,o=forwardinc,c=us
dom_access=x509cert://dsm r11/cn=domaccess,o=forwardinc,c=us
rep_access=x509cert://dsm r11/cn=repaccess,o=forwardinc,c=us
```

The change we have made here is to edit the original URI line (x509cert://dsm
r11/cn=domain_access,o=Computer Associates,c=us) so that it now reads as
cn=domaccess,**o=forwardinc**, and so forth, as this was the DN used when we first created
the new certificates.

The last step is to overwrite all instances of the certificate files (.DER and .P12) and
cfcert.ini files that exist in the installation media. This need only be done once, but is best
achieved through the use of a batch file. Search for all instances of DER, P12, and
CFCERT.INI files on the installation media, and overwrite with your new copies.

After you have successfully replaced the certificates and cfcert.ini files within the
installation image, installation can start as usual.

Review of Steps for Updating Certificates Before the Product Has Been Installed

1. Copy the CA DSM media to a writeable directory.

2. Create a temporary directory on the system that has DSM Explorer installed (c:\certs
   for example).

3. Create new certificates (using cacertutil through a batch file if preferred).

4. Copy an instance of CFCERT.INI file to the temporary directory.

5. Edit the [Files] section in CFCERT.INI with a new encrypted passphrase and a new URI in the [Tags] section (TAG section optional).

6. Copy the new certificates (*.DER, *.P12) to the installation media.

7. Copy the new CFCERT.INI to the installation media.
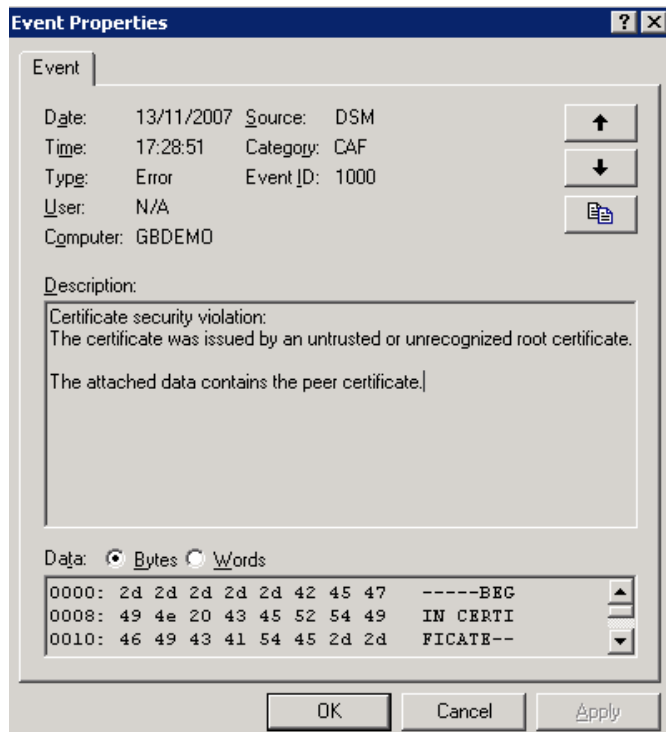
8. Install the product as normal.

Upon completion of the installation, the new certificates will be in use and imported into comstore. Any new deployments will utilize the new certificates. You will also see that the new security profiles will be created using the new DN from the [Tags] section of the CFCERT.INI (note the reference to Forwardinc in the URI) as follows:

In the following screen shot we can see what happens when an agent with the wrong certificates (or the standard out-of-box certificates) tries to connect to this secured CA DSM infrastructure. Note the 'encryption component failed' message on the client:

And the violation message in the Manager event log (note as of 11.2 C1 the IP Address of the violating node will be displayed in the event message):

Manager Already Deployed with Standard Certificates

The next scenario is encountered when you already have a Manager installed that is utilizing the out-of-box certificates. The steps to import and utilize the new certificates are listed below:

1. Utilize your updated CA DSM master image, by simply running Setup.exe from the media root. Select the Install Unicenter DSM option, which will then allow you to select the Repair option:

2. When the repair completes, this will result with the new security profiles being created (note that the original profiles will always continue to be listed in this scenario):

The existing original certificates will still be registered into comstore at this time. You can confirm this by running cacertutil list; this will show that only the original certificates are actually installed so far.

Replacing Certificates When Manager Is Already Installed but No Infrastructure Is Deployed

The next steps depend on whether you have at this stage already deployed an infrastructure (agents, for example) that will still be utilizing the original certificates. If you have only deployed a Manager and no other infrastructure, then follow the next steps. Otherwise, skip to the next paragraph (Replacing Certificates when Manager is Already Installed AND Infrastructure is Deployed).

3. Run cacertutil import to import your new certificates (use the commands in the cfcert.ini [Files] section). This will import the new certificates into comstore. In our example this needs to be run from your c:\certs directory that contains the certificates you originally created earlier in this chapter. The commands used in our example are below for reference although note the encrypted password below is unique to our example:

```
cacertutil import -i:itrm_dsm_r11_root.der -it:x509v3
cacertutil import -i:basic_id.p12 -
ip:enc:AhRau5Xy0ukms5NfbrHOpnbXfaZytytk -h -t:dsmcommon
cacertutil import -i:ccsm.p12 -t:csm -
ip:enc:E4H1xo9u2z4U9O775lZQVnZRQuDmJ4XF
cacertutil import -i:itrm_dsm_r11_cmdir_eng.p12 -
ip:enc:ewKioCf4EEnQbvr0NJBA16qAduWn3r27 -t:dsm_cmdir_eng
cacertutil import -i:itrm_dsm_r11_sd_catalog.p12 -
ip:enc:SHkCxBQyt7KHPCwmU08KwM43fE3MOq9y -t:dsmsdcat
cacertutil import -i:itrm_dsm_r11_agent_mover.p12 -
ip:enc:HnkJ2gLBGpl8D3rOz02UJQE11wibLhpj -t:dsmagtmv
cacertutil import -i:registration.p12 -
ip:enc:U2z8skmxRzdjaAVQDMWwYTHv3LJbtFvo -t:dsm_csvr_reg
cacertutil import -i:babld.p12 -
ip:enc:0H9c4YD4yb1xPR8Q3SSpr7EBzLAY9j84 -t:babld_server
cacertutil import -i:dsmpwchgent.p12 -
ip:enc:fBYRa9OPtZDMaVTi1FDCm5uw97Dnj7rk -t:ent_access
cacertutil import -i:dsmpwchgdom.p12 -
ip:enc:qhDmdQAeEPaUiVuVmuFaHrlAH5IlVNGX -t:dom_access
cacertutil import -i:dsmpwchgrep.p12 -
ip:enc:QH0qnofjHDO24PmWwBAfQTvu81gOXZzm -t:rep_access
```

4. The output from running the import command should be similar to that shown below for each new certificate:

```
C:\certs>cacertutil import -i:dsmpwchgrep.p12 -ip:enc:QH0qnofjHDO24PmWwBAfQTvu81

gOXZzm -t:rep_access

Importing...

Import of certificate allowed - replacing default with:

        'CN=repaccess,O=forwardinc,C=us'.
```

Running cacertutil list will also confirm that the new certificates are now imported into comstore. The original certificates will no longer be listed. This is because you can have only one certificate per tag name; the originals have been overwritten.

5. As the Root certificate does not utilize a specific tag it still exists in comstore at this stage, so you still have two root certificates installed. You now need to remove the original root with the cacertutil remove command:

```
cacertutil remove "-s:CN=DSM Root,O=Computer Associates,C=US"
```

6. Once this is done, no out-of-box agents will be able to communicate with your Management infrastructure.

7. The next step is to ensure that all newly deployed infrastructure nodes will utilize the correct certificates. To do this you need to use the dsmPush tool to update the software packages in the CA Unicenter Software Delivery Library and within the Deployment Manager. First, you will need to delete the CA DSM Software Packages from within the DSM Explorer. If this is the first time you have done this, you will receive a message box advising you that removing software from the Library will erase all installation records for that library item. You will need to confirm this action:

8. You do NOT need to delete the Deployment Packages that are used by the Deployment Manager function. For your information these are stored on a standard domain manager installation in C:\Program Files\CA\DSM\Packages. These will be updated automatically by the dsmPush tool.

9. Start the import of the CA Unicenter Software Delivery and Deployment Packages from your updated media by using the dsmPush tool. For example: *dmscript dsmpush copy* run from the root of your updated master image will start the process. See the *CA DSM Online HTML Documentation* (installed as part of a custom installation) for more information about the dsmPush tool.

   Note: This may take some time to complete. To confirm the update of the library and deployment manager has completed you can monitor the dmscript process within Task Manager. When this process completes, it indicates that the update is over.

   dsmPush also logs its action into your temporary directory. Once this process is complete, restart your manager. Then you will be ready to continue your deployment as normal, utilizing your new certificates.

Replacing Certificates when Manager Is Already Installed AND Infrastructure Is Deployed

This is similar to the process above, but the important difference is that you have to deploy your new certificates to your already-installed infrastructure (agents and so forth). This needs to be done in a specific order; otherwise, you risk disabling your agent infrastructure from being able to communicate with your manager. You need to start by deploying the new Root certificate *everywhere*. So the certificate deployment process could be achieved with a software delivery package.

1. Follow steps 1 through 4 as in the previous section, but *do NOT follow step 5*. Do not delete your original root certificate.

2. After completing steps 1–4 on your Manager, you need to import the new Root Certificate on any already-deployed infrastructure nodes. The actual command line for this is:

```
cacertutil import -i:itrm_dsm_r11_root.der -it:x509v3
```

You can use the above command within a CA Unicenter Software Delivery package. Cacertutil list on an agent would then return:

```
C:\certs>cacertutil list

Tag      : <empty>

Subject : CN=newdsmroot,O=forwardinc,C=us

Tag      : dsmsdcat

Subject : CN=DSM r11 Software Delivery Catalog,O=Computer Associates,C=US

Tag      : <empty>

Subject : CN=DSM Root,O=Computer Associates,C=US

Tag      : dsmcommon

Subject : CN=Generic Host Identity,O=Computer Associates,C=US
```

3. As this agent and the manager now have common trusted root certificates, they can connect and establish mutual trust. The next step would be to import the rest of the new relevant certificates *after* the root. For example, on an agent this would be the BHI and Catalog. Again, a software delivery package could achieve this for you.

4. Once you have deployed the new certificates to the rest of your infrastructure, you should then plan to remove the original root certificate from this infrastructure. Until this is removed, the agents are vulnerable to a rogue out-of-the-box manager. It is important to understand that this scenario can present challenges as some of the messages used in CA DSM are of a Persistent Store and Forward nature. So there may be messages in the infrastructure that may end up with delivery issues due to being issued prior to the change of certificates.

   Note: It is recommended that you should not delete the original root certificate for at least *14* days.

Chapter 17: Configuration Policy

The CA Unicenter Desktop & Server Management (CA DSM) products are configured centrally and locally through a shared component typically referred to as 'common configuration' or 'CCNF.' This component acts like an 'enhanced Windows registry,' in that it manages the runtime configuration of practically all of the CA DSM subcomponents and infrastructure features. This functionality allows your desired configuration settings and related business rules to be centrally managed, deployed, and enforced across your CA DSM environment.

Overview

Configuration policies can be viewed and manipulated within the DSM Explorer under \Control Panel\Configuration\Configuration Policy:

You will see what we call the 'Default Computer Policy,' which contains all known manageable parameters. The Default Computer Policy is sent to all computers automatically without any manual interaction as soon as the systems are registered into your CA DSM environment, meaning that every system has the same basic set of parameters.

You can change parameters by creating a new Configuration Policy, which is automatically derived from the original Default Policy. Any new Policy is referred to as a *custom policy*. Before you decide to create a custom policy, be aware that it is possible to edit the Default Computer Policy—this makes sense if the change you want to make is one that needs to be applied to *all* systems in your environment. We recommend that you change the Default Computer Policy if this is the case, and all computers will get the same new value. If it is a specific change for only a certain part of your population, then you should create a Custom Policy that contains your specific changes.

Note: In the event that you install a CA-supplied patch, new parameters and attributes may be added to the Default Policy — but your existing values will not be overwritten.

In real world scenarios, different groups or locations may require configuration parameters that will differ from other groups or locations. A common approach is to create a Custom Policy that can be assigned to multiple computers or groups in order to simplify administration of the configuration parameters. Note that from the administrator's point of view, configuration policies are created and maintained independently of any specific computer or group.

When you start to deploy Custom Policies to different groups, you may find that there could be some overlap in the parameters defined within your Custom Policies. In this case, the same parameter could be defined differently in more than one Policy and you could have a machine that is a member of more than one group. Since only a unique parameter value can be set on a computer, the following rules are used to determine how to proceed when configuration policies overlap:

- Policies assigned to a group are inherited by the children of the group. A child can be a group or computer.

- With a hierarchy, policies assigned on the child level override the ones on the parent level. In other words, all parameters defined on the parent level are also defined for the child. However, when a child policy overlaps with a parent policy, the child policy 'wins.'

- In all other cases where overlapping policies present a conflict, the user will be notified that there is a conflict to resolve. For example, imagine a scenario whereby you create two new policies that you wish to apply to two different groups that contain a number of agents where some of these agents are actually members of both of these groups. You have decided to make the Software Delivery Reboot Prompt Timeout value to be different for each group (remember—your target machines may exist in both groups).

  Application of the first policy to the first group would be fine. Policy application of the second policy to the second group would result in the Configuration Job failing to apply for the specific target(s) machine(s) that are a member of both groups. You will see a failed configuration job in the DSM Explorer for affected targets:

You will also see a reason for the policy application failure under the specific target's configuration policy node:

Here we can see a Policy Error with the following text: Policy 'Conflict Policy 2' assigned to 'Group2' conflicts with policy 'Conflict Policy 1' assigned to 'Group1': parameter '/itrm/usd/agent/RebootPromptTO.'

Important! This example is why we recommend that if you wish to assign a Configuration Policy to the All Computers group, then you should consider editing the Default Policy rather than applying a Custom Policy. Otherwise it is likely you will end up with Policy conflicts.

Activating and Distributing Configuration Policy

When the time comes for a computer configuration job to be activated, the manager collects the parameters that need to be sent down to the computer and does this through a Configuration Job. Configuration Job states can be viewed in the DSM Explorer under \Control Panel\Configuration\Configuration Jobs. Note that when a job is sent successfully to a target system, the result will not be stored under this tree node. If you see old Configuration Jobs listed here, that will indicate a problem with the pushing of the job. Typically, the problem would be that the target machine is not online or there is a policy conflict as described earlier. In terms of the target not being available, there is a timeout associated with a Configuration Job. Note that the value is controlled by the managed parameter msgtimeout in /itrm/manager/cc.

In R11.1 the default timeout is 7 days. In R11.2 the default is 60 days.

Note: If the system has been upgraded from r11.1 to 11.2, the existing default (7 days) is kept. This default parameter can be extended beyond 7 days.

Be aware that changes to the Default Policy will result in the changes being propagated immediately in your environment. It is important to understand the impact of the deployment of the update as custom policies are always derived from the Default Policy. In effect this means that the systems utilizing a custom policy will still receive a Default Policy update (automatically), even if they are running a custom policy.

Custom policy distribution time can be specified to occur at a specific time and you can view the combined values for a target. The Schedule Policies dialog has a Customize and Preview button to show the resulting configuration parameters:

Reported Configuration

The agent reports configuration parameter settings to the manager. On the manager, these settings are stored in the database where they are referred to by the GUI as the 'reported configuration':

A full report of all parameters is returned after the very first configuration job has been applied. Subsequent reports only contain deltas.

About Configuration Parameters

Configuration parameters can be:

- Centrally managed

  Configuration parameters are set up through the DSM Explorer and stored in the MDB. They are then evaluated and transmitted down to end systems through the CCNF and CSM (Configuration and State Management) sub-systems.

- Locally managed

  Here, although the MDB contains entries for the configuration parameters, the values are set and stored locally.

- Locally unmanaged

  This state can be set and reset only locally through the ccnfAgentApi. In other words, locally unmanaged parameters can be set to 'centrally managed' by the local end system. This essentially puts the manageability of these parameters under end-system control.

  These 'managed' parameters are collected together hierarchically under a configuration policy.

  Tip: When viewing managed parameters within the DSM Explorer, by default you will see their display names. To view their 'real' internal names, right-click on the list view column header and select the display 'internal names' option.

- Unmanaged

  Configuration parameters exist only on the end systems. These parameters typically contain values that are specific to and only useful to the processes that execute on the managed computer.

- Special Note on migrated parameters

  When, for example, a Unicenter Software Delivery (SD) v4 Agent is migrated to CA DSM r11, the migration procedure maps the Unicenter Software Delivery v4 Agent parameters to r11 parameters and writes the resulting values to the local comstore. This ensures that the r11 SD Agent continues using the previous defined configuration. Also note that migrated parameters are not modified by the r11 default configuration policy, which is pushed down automatically when the r11 Agent registers, as migrated parameters can only be changed by custom policies. Once a migrated parameter has been changed by a custom policy, it loses its 'migration' status and becomes an 'ordinary' parameter.

Enterprise and Domain Policies

On the Enterprise, the following rules apply to policies:

■ Enterprise policies are replicated to Domains.

■ Enterprise group configuration jobs are replicated to Domains.

■ Enterprise policies can only be applied to groups.

On the Domain, the following rules apply to policies:

■ Enterprise default policy replaces Domain default policy.

■ Policies can be applied to group assets or to individual assets.

■ Reported configuration is only available at the Domain level.

■ Replicated policies cannot be modified at the Domain level. This also applies to the Default policy.

Property Storage

Property storage (also known as 'persistence') is maintained in different locations—in the MDB and on the end system itself.

In the MDB

Centrally managed properties and their values, as well as locally managed property definitions (without values), are stored in the following MDB tables

■ csm_property

■ csm_object

■ csm_class

■ csm_link

The configuration manager processes the configuration policies applied to a specific computer as well as policies applied to any groups that the computer belongs to. It will eventually determine which properties should be set to what value and then will push these property values down to the end system.

On the End System

The configuration settings for a CA DSM system (manager/server or agent) end up on each node in an encrypted XML file store, made from comstore.enc and userstore.enc.

comstore.enc is stored in <DSM install path>\Agent\CCNF while userstore.enc is stored in <Documents and Settings>\<DSM install path>\Agent\CCNF.

These files contain the settings an agent will be using. As an example, if you wanted to know what property values are actually in force on a particular machine, this is a good place to look. Since the files are encrypted to avoid direct manipulation, you cannot open the file within a text editor to view the values but you can view the file in a variety of ways. You can utilize the CCNFCMDA command line interface—for example, to view a specific value, such as the Reboot Timeout Prompt value on an agent. In this scenario you could type the following (remembering we are using the 'internal' name to reference the values as mentioned previously):

```
ccnfcmda -cmd GetParameterValue -ps /itrm/usd/agent/ -pn RebootPromptTO
```

In addition the file can be decrypted and written to a standard XML format file using the following command:

```
ccnfcmda –cmd GetConfigStore –fi c:\MyComStore.xml
```

Examples of CA DSM Configuration Settings

The following content example shows some nested parameter sections (for example, 'itrm/usd/agent'), as well as a parameter 'rebootpromptto' and its value '1900':

```
- <paramsection name="usd">
  + <paramsection name="filetransfer" orgname="FileTransfer">
  + <paramsection name="filecompression" orgname="FileCompression">
  + <paramsection name="manager" orgname="Manager">
  - <paramsection name="agent" orgname="Agent">
      <parameter name="swddeltamode2" orgname="SWDDeltaMode2" entity="Manager" value="10" />
      <parameter name="swddeltamode" orgname="SWDDeltaMode" entity="Manager" value="100" />
      <parameter name="guimode" entity="Manager" value="2" />
  + <parameter name="ws" entity="Client" value="cscript.exe">
  + <parameter name="wsf" entity="Client" value="cscript.exe">
  + <parameter name="vbs" entity="Client" value="cscript.exe">
  + <parameter name="js" entity="Client" value="cscript.exe">
  + <parameter name="dms" entity="Client" value="dmscript.exe">
  + <parameter name="ips" entity="Client" value="dmscript.exe">
      <parameter name="allowapplicationlogoffreboot" orgname="AllowApplicationLogoffReboot" entity="Manager" value="0" />
      <parameter name="ignorerebootbeforelogoff" orgname="IgnoreRebootBeforeLogoff" entity="Manager" value="0" />
      <parameter name="skipoutputfiles" orgname="SkipOutputFiles" entity="Manager" value="0" />
      <parameter name="disablequote" orgname="DisableQuote" entity="Manager" value="0" />
      <parameter name="promptinbatch" orgname="PromptInBatch" entity="Manager" value="0" />
      <parameter name="disablejcewindowclose" orgname="DisableJCEWindowClose" entity="Manager" value="0" />
      <parameter name="hintsoftwarescanner" orgname="HintSoftwareScanner" entity="Manager" value="OFF" />
  + <parameter name="calculateziphash" orgname="CalculateZipHash" entity="Client" value="0">
      <parameter name="msiuserloggedon" orgname="MSIUserLoggedOn" entity="Manager" value="1" />
      <parameter name="msisourceupdate" orgname="MSISourceUpdate" entity="Manager" value="1" />
      <parameter name="hidejobcheckicons" orgname="HideJobCheckIcons" entity="Manager" value="0" />
      <parameter name="noslessswitchallowed" orgname="NOSLessSwitchAllowed" entity="Manager" value="1" />
      <parameter name="rebootpromptto" orgname="RebootPromptTO" entity="Manager" value="1900" />
      <parameter name="rebootpromptrt" orgname="RebootPromptRT" entity="Manager" value="10" />
      <parameter name="forcedreboot" orgname="ForcedReboot" entity="Manager" value="0" />
```
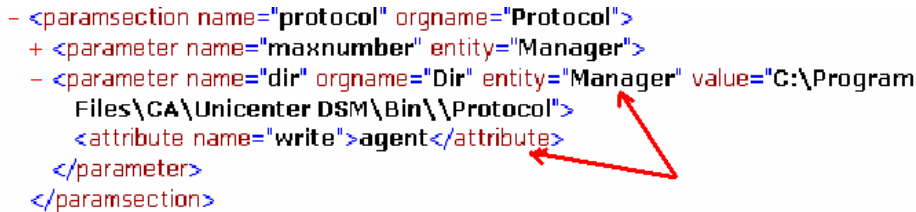
A managed parameter is indicated by an entity value of 'Manager.' For example:
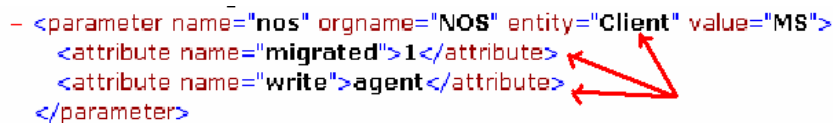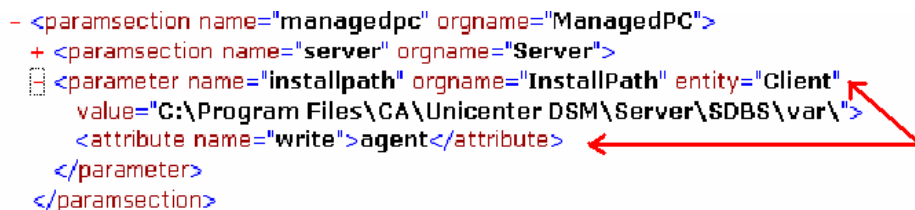
```xml
- <paramsection name="general">
  + <parameter name="ipchangescript" value="" entity="Manager">
    <parameter name="enableremotecli" value="2" entity="Manager" />
```

When the attribute 'write' is set to 'agent,' this means the local end system may update the parameter value. For example:

```xml
- <paramsection name="protocol" orgname="Protocol">
  + <parameter name="maxnumber" entity="Manager">
  - <parameter name="dir" orgname="Dir" entity="Manager" value="C:\Program
      Files\CA\Unicenter DSM\Bin\\Protocol">
    <attribute name="write">agent</attribute>
  </parameter>
</paramsection>
```

This is an example of a migrated parameter:

```xml
- <parameter name="nos" orgname="NOS" entity="Client" value="MS">
    <attribute name="migrated">1</attribute>
    <attribute name="write">agent</attribute>
</parameter>
```

This is an example of locally unmanaged parameters:

```xml
- <paramsection name="host">
    <attribute name="managed">local</attribute>
  + <parameter name="installpath" value="n/a">
```

This is an example of an unmanaged parameter:

```xml
- <paramsection name="managedpc" orgname="ManagedPC">
  + <paramsection name="server" orgname="Server">
  + <parameter name="installpath" orgname="InstallPath" entity="Client"
      value="C:\Program Files\CA\Unicenter DSM\Server\SDBS\var\">
    <attribute name="write">agent</attribute>
  </parameter>
</paramsection>
```

Extending the Configuration Policy

The Default Configuration Policy has been built to be extensible by our Development group so that we can add additional parameters to Configuration Policy to support new and future functionality. Please note, policy changes are only supported when they are provided by CA in the form of a product update.

Resetting Configuration Policy Parameters

When making changes to Configuration Policy Parameters within the DSM Explorer you have the option to discard changes that you have made, but this will only take effect for changes made while you have unsealed the policy you have just been working on. In the event that you have made changes to the Default Policy but you wish to reset the parameter values back to the out-of-the-box values, there is a command you can utilize for this:

```
ccnfregdb -mlocalhost -d"c:\program files\ca\dsm\Manager\CCNf" -o
```

This will read the XML files located in the Manager\CCNF directory and rewrite the Default Policy. Again, changing the Default Policy will result in the policy being redistributed to all agents.

Chapter 18: Mobile Devices

Introduction

CA Unicenter Desktop & Server Management (CA DSM) provides support for Microsoft Windows CE-based devices through the use of a native agent, and also supports other mobile and docked devices. In this chapter we will focus specifically on Microsoft Windows CE and the native device agent. First, we shall clarify the terminology around native and docked agents. See the points below:

■ Native refers to the fact that there is an actual CA Unicenter Software Delivery or CA Unicenter Asset Management agent physically installed on the device itself. This agent communicates over TCP, which can be WiFi, Bluetooth, GPRS (used by GSM mobile phones), or a LAN connection.

■ Docked refers to devices that are set up to synchronize with a desktop/laptop—for example, a CE device that synchronizes using the Microsoft ActiveSync program. The desktop/laptop will need to have a DSM Agent installed, which in turn will act as a proxy and allow you to communicate with the docked device from the perspective of CA DSM. For example, software can be targeted to the device and it will be installed when the device is next docked and synchronized with its host machine.

See the *Unicenter Desktop & Server Management Implementation Guide* chapter, 'How to Enable a Docking Device on Windows' for instructions on enabling the docking device support, as it is not enabled by default in r11. Once enabled, the device will be visible in the DSM Explorer.

Whereas r11 supports devices through the docked device and synchronization method, there is no actual r11 version of the Native Agent at this time. To take advantage of the r4 Native Agent, you can utilize the Agent Compatibility Bridge that is available for r11 that enables legacy CA Unicenter Software Delivery and Asset Management r4 Native Agents to communicate with and be managed by an r11 infrastructure. The Agent Compatibility Bridge is available on request from Technical Support. The first release of the Agent Compatibility Bridge supported Windows Mobile 5 devices and was compatible with CA DSM r11.2a.

With that in mind, when we talk about the Native Agent on the following pages we are referring to the CA DSM r4 version of the agent, in the context of being managed by an r11 infrastructure.

Deployment Options

Since mobile devices are mobile by their very nature—capable of moving from place to place—targeting them for deployment can present challenges. The original standard Unicenter Software Delivery r4 procedures, which involved sweeping the network and deploying using that method, are presented with challenges in the mobile device world. To elaborate, when running Unicenter Software Delivery r4, if the Agent SD Primer executable

is running on a CE Device, then it is possible to deploy to the device over the network using the r4 Deployment Wizard or SDSWEEP command line. However, the reality is that this is not a viable option considering the limited time these devices spend on the network and the inherent bandwidth challenges.

From a Unicenter Asset Management r4 perspective, since there is no automated sweep type mechanism like there is in Unicenter Software Delivery r4, the CA Asset Management Agent must be installed using either CA Unicenter Software Delivery or Microsoft ActiveSync. Generally the same rules apply here as they do in the desktop world. Once CA Unicenter Software Delivery is installed then any other relevant applications, including CA Asset Management, can be installed using that mechanism.

A scenario that is perhaps more manageable and more common is that the Unicenter Software Delivery r4 Agent is included in the actual build of the CE Device. In that case, the agent is active by the time a user receives the device. This can be achieved fairly simply, as the CA Unicenter Software Delivery r4 Agent is distributed in the CAB file format. This Agent CAB file can be copied to the device and the installation can be run as the device is first built, either automatically through your build scripts or manually depending on your requirements.

Device Naming

Before you begin any deployment, it is important to understand how the r4 Agents deal with the device naming concept in the CE world. Since many devices are provided new, named as Pocket_PC, the Unicenter Software Delivery r4 Native Agent will create a new name for each device upon the agent installation. This ensures that the device name is unique to CA Unicenter Software Delivery. The unique device identifier is an encoded string made up from calls made to the Windows CE kernel. Note that this name change is only from the perspective of the CA Unicenter Software Delivery Agent. The device name is not physically changed on the device.

You may wish to control the name that is given to the device. To do that you need to create a registry key on the device that has the name pre-set. An exported key is shown here:

    [HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\ID]

    "DeviceID"="<new_name_here>"

Following this will result in the installation of the Unicenter Software Delivery r4 Agent using the name it finds in the registry, and as such gives you control over the naming of the device.

Agent Configuration

The Unicenter Software Delivery r4 Native Agent can be configured using the registry on the device. Following are some examples of useful configurable options for the agent.

The following key will change the temporary download location used for a NOS-less installation to be a directory called 'temp'.

[HKEY_LOCAL_MACHINE\SOFTWARE\ComputerAssociates\Software Delivery\ASMCNF\Public]

  "InstTempPath" = "/temp"

Note: the quotes around the location are not needed when entered into the Device Registry; they are a result of a Registry export.

To enable CheckPoint Restart capabilities, this Registry Key needs to be set on the device, as follows (**Note**: Your agent needs to be running the correct genlevel to activate this functionality):

[HKEY_LOCAL_MACHINE \SOFTWARE\ComputerAssociates\Software Delivery\ASMCNF\Public]

CheckPointRestartActive = 1

To hide the Progress Window that is seen when the files are being downloaded to a device before package execution, you can configure the Registry Key as follows:

[HKEY_LOCAL_MACHINE \SOFTWARE\ComputerAssociates\Software Delivery\ASMCNF\ASMINST]

DisplayProgressWindow = 0

Note: This is not the same as hiding the Job Check window which can be achieved through the standard mechanism in CA Unicenter Software Delivery.

Another point to be aware of is related to the agent job check interval. On a CE Device the agent job checks upon device wake up, so when a device wakes up it will by default send a job check. This can happen many times during the day depending on how the device is used.

The job check feature can also be used to help determine a degree of agent health. When an agent job checks with its r11 server, you could consider utilizing a query based upon certain values. For example, the Last Hardware Scan value (shown in the following screen shot) can be queried and desired action taken as appropriate.

The screenshot below is from a CA DSM r11 domain manager that has the Agent Compatibility Bridge component installed. This view shows you some of the r4 Agent Procedures that are available, such as 'Retrieve all trace log files from an agent,' which can assist with certain troubleshooting tasks.

Chapter 19: Operating System Installation Management

Introduction

The *Inside OS Installation Management Guide* provided with r11.2 C1 of CA Unicenter Desktop & Server Management (CA DSM) contains a great deal of information related to OSIM. It is the first point of reference when seeking details on how to implement this feature.

The following sections add to or expand upon the information contained in the *Inside OS Installation Management Guidee*.

Migration of Legacy Boot and OSIM Images

OSIM Images

The *Unicenter Desktop & Server Management Implementation Guide* contains full details on how to migrate the OSIM images from a legacy version to r11.

Boot Images

It is best practice to create new OSIM boot images for use with r11. If you do not have access to a Win98 SE system or an NT Server CD, then please follow the instructions in the section Creating DOS Boot Images without Access to a Win98 or NT CD below.

Boot File Creation

WinPE

The *Inside OS Installation Management Guide* provided with CA Unicenter Desktop & Server Management r11.2 C1 contains detailed instructions on the creation of the WinPE boot images.

The versions of WinPE supported by r11.2 C1 are:

■   WinPE 2005

    WinPE 2005 can be downloaded from the Microsoft Software Assurance Windows Benefits. A sign-in by an authorized license user is required:

Please search for the WinPE 2005 Microsoft Part Number: T94-00005.

Note: CA Unicenter Desktop & Server Management only supports the use of Windows Server 2003 SP1 as the source media with this version of WinPE for 32-bit systems and Windows XP Professional x64 for 64-bit systems.

■   WinPE 2.0

WinPE 2.0 can be downloaded as part of the Windows Automated Installation Kit (AIK) from the Microsoft download sites.

Adding Drivers

WinPE 2.0

When creating the WinPE 2.0 image, the copype script will prompt to provide any additional required drivers as part of the build process.

WinPE 2005

If the copype script is used to create the WinPE 2005 image, it will prompt for additional drivers.

If the manual method is used to create the WinPE image, the additional drivers need to be copied into the WinPE image before creating the disk image. See the *Inside OS Installation Management Guide* for specific details on where to place the drivers.

From CA DSM r11.2 C1, there is an additional folder provided in the WinPE templates on the Image Prepare System (IPS).

If additional plug and play drivers are needed within the WinPE Boot Image they can be added to this folder in the template. They will then be incorporated into any WinPE Boot Image created from this template.

Note: There is a separate folder for each of the WinPE versions.

Partitioning a Computer's Hard Drive During an OSIM Build When Using WinPE as the Boot Image

Prior to the support of WinPE, the hard drive of the target PC could be partitioned using the cafdisk utility. This was controlled by the <OS imagename>.par file.

When using a WinPE-based boot image, it is no longer possible to use cafdisk so another solution is required.

OSIM uses the diskpart.exe from Microsoft to partition the hard drive in a WinPE environment. Diskpart.exe is controlled from our script file osimdisk.txt located on the Image Prepare System (IPS) located by default in:

   c:\program files\DSM\CA\osimips\os-templates\updates\winpe\ca-osim\osimdisk.txt

This file is copied from the IPS to the WinPE image as part of the procedure for creating the WinPE image for use with OSIM. See the *Inside OS Installation Management Guide* provided with CA Unicenter Desktop & Server Management r11.2 C1 for detailed instructions on creating the WinPE image.

The default configuration supplied with OSIM is to create a single partition that uses the entire disk capacity of Disk 0.

Osimdisk.txt can be modified prior to the creation of the WinPE.iso to configure alternate partitioning of the target PCs. A full description of the diskpart commands and scripts is contained within the WinPE Help files from Microsoft.

An example where the C: drive is partitioned with 10GB and the D: drive is partitioned with the remainder of the space in Disk 0 is shown below:

```
REM ++++ input for diskpart. ++++
REM ++++

REM ++++ partition the disk 0
select disk 0
clean
create partition primary size=10240
assign letter=c
active
create partition extended
create partition logical
assign letter=d
exit
```

Note: The lack of a size = variable creates a second partition that uses the entire remaining hard disk.

Important: If the hard disk is smaller than the first partition defined, this will cause the script to fail—therefore causing the WinPE phase of the OSIM build to fail.

If the target system already had an operating system installed in the past, then the CD drive will already be assigned to a drive letter. If the previous build only had one partition, then the CD drive will most often have been assigned to D:.

This would cause the above sample to fail as the second partition was defined to D:.

In this case, the CD drive letter needs to be changed prior to the partitioning of the hard drive. We will change it to E: in the following example:

```
REM ++++ input for diskpart. ++++
REM ++++

REM ++++ set the CD-drive letter to e:
select volume 0
remove
assign letter=e

REM ++++ partition the disk 0
select disk 0
clean
create partition primary size=10240
assign letter=c
active
create partition extended
create partition logical
assign letter=d
exit
```

DOS

Creating DOS Boot Images without Access to a Win98 or NT CD

If you need to create a DOS-based boot image and there is no access to a Windows 98 SE system or Windows NT Server CD, the following process will guide you through creating a valid boot image for use with either FTP or share mode. (**Note**: The MSCLIENT will work with the NIC installed in the target computer). This section expands on the details provided in the *Inside OS Installation Management Guide.*

Follow these steps:

1. Download wbootess.exe from the Internet and save it to your desktop or other folder. A simple search will identify the location; CA does not endorse any particular website for this.

2. Insert a blank floppy disk.

3. Run wbootess.exe. You will see the following screen:

   a. Accept the prompt to format.

   b. Wbootess.exe will run and create a Windows 98 SE boot disk.

4. On completion, the floppy will now have the following files (shown in the next screen shot):

Note: Some of these are hidden system files.

A:

Name ▲	Size	Type	Date Modified	Attributes
ATTRIB.EXE	15 KB	Application	23/04/1999 22:22	A
AUTOEXEC.BAT	1 KB	Windows Batch File	02/06/2001 21:10	HSA
COMMAND.COM	92 KB	Application	23/04/1999 22:22	HSA
config.sys	1 KB	System file	02/06/2001 20:14	A
DELTREE.EXE	19 KB	Application	23/04/1999 22:22	A
DOSKEY.COM	16 KB	Application	23/04/1999 22:22	A
EDIT.COM	69 KB	Application	23/04/1999 22:22	A
EDIT.HLP	11 KB	Help File	23/04/1999 22:22	A
EMM386.EXE	123 KB	Application	23/04/1999 22:22	A
EXTRACT.EXE	92 KB	Application	23/04/1999 22:22	A
FDISK.EXE	63 KB	Application	23/04/1999 22:22	A
FORMAT.COM	49 KB	Application	23/04/1999 22:22	A
HIMEM.SYS	33 KB	System file	23/04/1999 22:22	A
IO.SYS	218 KB	System file	23/04/1999 22:22	RHSA
MEM.EXE	32 KB	Application	23/04/1999 22:22	A
MORE.COM	11 KB	Application	23/04/1999 22:22	A
mouse.exe	5 KB	Application	01/01/2001 01:08	A
MOVE.EXE	27 KB	Application	23/04/1999 22:22	A
MSCDEX.EXE	25 KB	Application	23/04/1999 22:22	A
MSDOS.SYS	1 KB	System file	23/04/1999 22:22	RHSA
OAKCDROM.SYS	41 KB	System file	23/04/1999 22:22	A
REBOOT.COM	1 KB	Application	07/02/1988 06:00	A
SCANDISK.EXE	141 KB	Application	23/04/1999 22:22	A
SCANREG.EXE	162 KB	Application	23/04/1999 22:22	A
SMARTDRV.EXE	45 KB	Application	23/04/1999 22:22	A
SYS.COM	19 KB	Application	23/04/1999 22:22	A
XCOPY32.MOD	41 KB	MOD File	23/04/1999 22:22	A
XCOPY.EXE	4 KB	Application	23/04/1999 22:22	A

5. Remove all unnecessary files from the floppy. The required files are shown below. It is important that no other files are on the floppy.

A:

Name ▲	Size	Type	Date Modified	Attributes
COMMAND.COM	92 KB	Application	23/04/1999 22:22	HSA
EMM386.EXE	123 KB	Application	23/04/1999 22:22	A
FORMAT.COM	49 KB	Application	23/04/1999 22:22	A
HIMEM.SYS	33 KB	System file	23/04/1999 22:22	A
IO.SYS	218 KB	System file	23/04/1999 22:22	RHSA
MSDOS.SYS	1 KB	System file	23/04/1999 22:22	RHSA
SMARTDRV.EXE	45 KB	Application	23/04/1999 22:22	A
SYS.COM	19 KB	Application	23/04/1999 22:22	A

This floppy is now ready for the CreateBTImages process.

6. Create a folder MSCLIENT on the Image Prepare System (IPS), and download the MSCLIENT files from the Microsoft FTP server:

ftp://ftp.microsoft.com/bussys/Clients/MSCLIENT

Place it in the MSCLIENT folder on the IPS.

7. Extract both self-extracting files into this folder. There will be three prompts for overwriting; answer Yes to all of these prompts.

8. Create the boot image with the following command:

    createbtimages -l <path to msclient>\msclient

    For example, for the path in screenshot above:

    createbtimages -l "C:\Program Files\CA\DSM\osimips\MSCLIENT"

    When this procedure is completed, the Image Prepare System creates two boot images named osinstal.2 and osinstal.3, which are stored at the specified location. Use the command createbtimages -x to view the location at which the images are located.

9. These images are now ready for registering into the domain manager using the Registerbtimages command. See the *Inside OS Installation Management Guide* for details on this procedure.

Networking

Routing When the Target PC is in a Remote Subnet

If the CA Unicenter Desktop & Server Management boot server and target PC are in separate subnets, the router/switch must be configured to forward the UDP (User Datagram Protocol) PXE (Preboot Execution Environment) broadcast requests to the CA Unicenter Desktop & Server Management boot server. If this is not correctly configured, the target PC will not be able to contact the boot server to request the boot files during the PXE boot.

For example, if the router is a Cisco router then an IP helper needs to be defined:

    IP helper-address <address of boot server>

WinPE (4011)

In addition to a standard IP helper, when using a WinPE-based boot image, UDP packets over port 4011 must be routed to our boot server as well:

IP forward-protocol UDP 4011

After the initial PXE boot process, the WinPE image is loaded. Once it has loaded, the OSIM tools included in the WinPE need to locate the boot server and obtain the parameters to control the OS Installation.

The sdmpcimg.exe command is used to achieve this. The first thing that sdmpcimg.exe must do is discover the boot server and obtain the parameter file. It is the discovery that requires the use of port 4011.

If this is not configured then, even if the target PC can contact the boot server during the PXE boot process, the WinPE image cannot discover the boot server and will therefore fail to build the operating system.

Tracing the Network Traffic

In the event of a failure to contact the boot server during the PXE or WinPE processes, it is necessary to configure a network trace tool such as Wireshark or Microsoft Network Monitor to monitor the network traffic on both sides of the router to identify where the failure is.

To trace this successfully, a good understanding of what occurs during a PXE boot is required.

In the next example, the following IP addresses are used:

- DHCP (Dynamic Host Configuration Protocol) Server   = 192.168.100.1

- Boot server   = 192.168.100.10

- Router   = 192.168.101.2

The PXE boot executes as follows:

1. On power up, the target PXE client sends a DHCP discovery broadcast which contains PXE-specific options.

2. The boot server responds with an offer of the boot server's IP address and PXE options.

3. The DHCP server responds with an offer of an IP address.

4. PXE requests the IP address from the DHCP server.

5. The DHCP server acknowledges the requests and grants the IP address.

6. The PXE client requests a boot file (**Note**: It now has a valid IP Address).

7. The DHCP server acknowledges the request but does not know the location of the boot file.

8. The boot server acknowledges the request and details the location and name of the boot file to use.

9. The PXE client makes a TFTP (Trivial File Transfer Protocol) request for the boot file.

10. The boot server acknowledges the request and the boot file is transferred to the client through TFTP.

In the case of a remote DHCP or boot server there will be multiple Discover, Offer, ACK, and Request packets, as there will be one from each router (relay).

Because the traffic is going through the router, all the requests come from the router's IP address. Responses from the DHCP server and the boot server are sent to the router, which then forwards them on to the broadcast address on the remote subnet.

It is also possible to identify which IP Helpers are defined, as the Discovery packets on the routed trace are sent to specific IP addresses as defined by the IP Helpers. On the local trace, the discovery packets are sent to the Broadcast address on the subnet.

Understanding what you should expect to see in the network trace can help with determining what the problem is.

For example, if the Discovery Broadcast is not seen on the subnet of the boot server, then there is no 'IP Helper' or equivalent to forward the PXE Discover Broadcast from the target's subnet to the boot server.

However, if the Broadcast is seen on the boot server's subnet but there is no offer from the boot server, then the problem is not with the network and the server itself needs to be checked.

It is possible that there are no 'IP Helpers' configured at all. In this case, as well as no OFFER from the boot server, there would also be no OFFER from the DHCP server and the target PC would not receive an IP address.

Serial ATA

Most new desktops and laptops now ship with SATA (Serial Advanced Technology Attachment) hard disk drives instead of IDE (Integrated Drive Electronics) drives. This commonly causes issues with the OSIM process. The primary cause of these problems is that a DOS environment has problems addressing a SATA hard disk.

Problems that are often seen include failure to partition and copy errors while Windows Setup copies files to the disk. There are several things that can be attempted to resolve this issue:

WinPE

This is the preferred resolution. It provides a Windows environment which can therefore support the SATA drives.

IDE Emulation

On some systems it is possible to configure in the BIOS the hard disk to run in an IDE Emulation Mode which enables the DOS environment to work with the SATA drive.

Vendor Restrictions

Some system manufacturers now specifically state that their systems do not support a DOS environment. In this case, the only option is to use a WinPE-based boot file.

Adding Drivers to an Operating System (OS) Image

Plug and Play Drivers

If the target PC has hardware for which the operating system (OS) image does not contain suitable drivers, then it is possible to provide additional drivers to the OS image. This is particularly the case with Windows operating systems.

The steps to add additional plug and play drivers to a Windows OS image are as follows (**Note**: There is no limitation in path length with CA DSM r11.1):

1. Open the OEM folder in the OS Image located by default in:

    ...\DSM\Server\SDBS\var\managedpc\images\<imagename>\<imagename>\i386

2. Create a $1 folder in the i386 folder.

3. Create a folder to contain the driver files and folders. For example, pnpdrvrs:

    ...i386\$1\pnpdrvrs

4. Add the required drivers to this folder. Place each driver in a separate folder. For example:

    ...i386\$1\pnpdrvrs\audio

    ...i386\$1\pnpdrvrs\network

    It is possible to have one image used to image multiple models of PC. In this case, it might be beneficial arrange the folders so that they are identifiable per model. For example:

    ...i386\$1\pnpdrivers\model1\audio

    ...i386\$1\pnpdrivers\model1\network

    ...i386\$1\pnpdrivers\model2\audio

    ...i386\$1\pnpdrivers\model2\network

5. Edit the <imagename>.inf file located by default in:

    ...\DSM\Server\SDBS\var\managedpc\camenu

6. Add the OEMPnPDriversPath entry to the [Unattended] section of this file. For example:

    OEMPnPDriversPath = 'pnpdrvrs\audio;pnpdrvrs\network'

Setup will create the pnpdrvrs folder on the root of the system drive during the text mode setup and will then scan this location for additional plug and play drivers during the setup of the OS.

From CA DSM r11.2 C1 this process is simplified. The templates on the IPS now contain a driver folder within the oeminst folder.

Plug and play drivers can now be added to the image template by adding them to this folder prior to creating the OS Image.

It is also possible to add them to the image between the creation and registration.

```
☐ 📁 SDBS
        📁 inst
    ☐ 📁 var
        ☐ 📁 managedpc
            ⊞ 📁 agents
            ⊞ 📁 camenu
            ☐ 📁 images
                ⊞ 📁 dosboot
                ⊞ 📁 vista64b
                ⊞ 📁 vistax64
                ☐ 📁 WXPPSP2
                    ☐ 📁 WXPPSP2
                        ☐ 📁 i386
                            ☐ 📁 $oem$
                                ☐ 📁 c
                                    ☐ 📁 oeminst
                                            📁 driver
```

There is no need to modify the <imagename>.inf for this as it has already been configured to look in this folder for added drivers.

It is still possible to add each driver in a separate folder as the setup routine will traverse the driver folder and subfolders.

Mass Storage

Some systems require additional mass storage device drivers before the device will be useable. This could be to support a RAID (Redundant Arrays of Independent Disks) system, SCSI (Small Computer System Interface) device, or a SATA drive.

These drivers need to be added to the plug and play drivers as above, but also need to be added to the text mode portion of the setup so that the devices can be addressed during this phase. To do so, follow the instructions below:

1. Open the OEM folder in the OS Image located by default in:

   ...\DSM\Server\SDBS\var\managedpc\images\<imagename>\<imagename>\i386

2. Create a folder 'Textmode' in the i386 folder:

   ...\i386\OEM\Textmode

3. Copy the mass storage drivers to this folder. The driver files are normally:

   <driver>.sys, <driver>.dll, <driver>.inf, <driver>.cat, Txtsetup.oem.

4. Edit the <imagename>.inf file located by default in:

   ...\DSM\Server\SDBS\var\managedpc\camenu

5. Insert a new section containing the file names from the Textmode folder:

    [OEMBootFiles]

    <driver>.sys

    <driver>.dll

    <driver>.inf

    <driver>.cat

    Txtsetup.oem

6. Insert a new section with the name of the device:

    [MassStorageDrivers]

    "Name of the device"="OEM"

Note: Information about the device name can be obtained from the Txtsetup.oem file, which is provided by the hardware vendor.

Server Download Method

The default server download method for a boot server is TFTP unless specified otherwise at installation. All r11 OSIM images are prepared to work with TFTP access with the following exceptions:

■   Ghost and Powerquest images will only work with share access.

■   OSIM LINUX images always use NFS (Network File System) shares. The NFS share must contain the LINUX CDs/DVD directory structure.

Changing the Download Method

The boot server setup installs a switch tool, sdbsswitch, which can be used to change the boot server access method later.

The sdbsswitch tool creates or removes the OSIM shares and adjusts the OS images in the image store according to the access method:

■   sdbsswitch -t switches from share to TFTP access.

■   sdbsswitch -s switches from TFTP to share access.

■   sdbsswitch -l shows the current configuration of the boot server.

Share

■   Installing an OS on a target is faster from a share because the setup installs directly from the share.

- OSIM shares are read-only for a special OSIM user. If your enterprise policy does not allow shares, the TFTP download method is required.

TFTP

- OSIM has implemented a special extended TFTP protocol controlled by the OSIM boot server. In the case of TFTP, all needed data will be downloaded to the target before the installation starts.

Use of Imaging Tools

Inclusion of Agent in Image

When deploying an OS using an imaging tool, it is necessary to decide whether or not to include the DSM Agent as part of the image.

The recommendation is to not include the agent. The CA Unicenter Software Delivery Agent will be installed as part of the OS installation, and this then allows CA Unicenter Desktop & Server Management to deploy the other plugins as required.

If it is decided to include the agent within the image, some actions are required on the template machine if the CA Unicenter Remote Control agent is part of the image.

CA DSM r11.1, CA DSM r11.0

Before generating the image, perform the following actions on the template machine:

1. Start the agent and register it with a manager to ensure proper verification of the management server information.

2. Run the following commands:

   a. caf stop

   b. ccnfcmda -cmd DeleteParameter -ps itrm/rc/host/managed -pn convertedhostuid

   c. ccnfcmda -cmd DeleteParameter -ps itrm/cfencrypt -pn LOCALID

   d. ccnfcmda -cmd DeleteParamSection -ps itrm/rc/security/providers/winnt/users

3. Delete the following registry key:

   HKLM\SOFTWARE\ComputerAssociates\HostUUID

4. Generate the image to be deployed BEFORE starting the DSM Agent again.

5. The deleted values are machine-specific and will be properly regenerated on the template machine, as well as on the target machines at agent start.

CA DSM r11.2

Before generating the image, perform the following actions on the template machine:

1. Start the agent and register it with a manager to ensure proper verification of the management server information.

2. Run the following commands:

    a. caf stop

    b. ccnfcmda -cmd DeleteParameter -ps itrm/rc/host/managed -pn convertedhostuid

    c. ccnfcmda -cmd DeleteParameter -ps itrm/cfencrypt -pn LOCALID

    d. ccnfcmda -cmd DeleteParamSection -ps itrm/rc/security/providers/common/users

3. Delete the following registry key:

    HKLM\SOFTWARE\ComputerAssociates\HostUUID

4. Generate the image to be deployed BEFORE starting the DSM Agent again.

5. The deleted values convertedhostuid, LOCALID, and HostUUID are machine-specific and will be properly regenerated on the template machine, as well as on the target machines at agent start.

The default users for the security providers will not be recreated automatically by the host. Therefore, if the security mode of the host is not set to centralized security, the following command to re-create the default users for the host must be executed on the template machine, as well as on each target machine:

    rcUtilCmd.exe CreateDefaultUsers

Note: rcUtilCmd.exe can be found in the CA Unicenter Desktop & Server Management installation's directory.

Ghost/PowerQuest

The *Inside OS Installation Management Guide* supplied with CA Unicenter Desktop & Server Management r11.2 C1 contains a detailed section on the creation of OS images based on the Ghost and PowerQuest Imaging Tools.

Drivers

Adding plug and play drivers to a Ghost image is very similar to the procedure for a standard unattended installation.

In this case, the drivers would be placed in the target folders on the template system prior to running the sysprep utility. For example:

C:\pnpdrvrs\audio

C:\pnpdrvrs\network

The same modification is made to the <imagename>.inf file to define the OEMPnPDriversPath.

From CA DSM r11.2 C1 this process is simplified. The templates on the IPS now contain a driver folder within the oem folder.

Plug and play drivers can now be added to the image template by adding them to this folder prior to creating the OS Image. It is also possible to add them to the image between the creation and registration.

The drivers will be copied by OSIM to the C: drive before the Windows Mini Setup is launched. There is no need to modify the <imagename>.inf for this as it has already been configured to look in this folder for added drivers.

When the image is deployed, the Windows Mini Setup will check these folders for the additional plug and play drivers.

Ghost32

With the inclusion of WinPE support it is now possible to use 32-bit Ghost.

This removes the requirement to use FAT16\32 partitions, as the WinPE environment can handle the NTFS (Windows NT File System) partition.

Default Passwords

Canonprv

The canonprv user account is used to access the OS image when the server is running in share mode. In order to preserve security this account does not have a default password. Instead, the OSIM Manager automatically changes this password on a daily basis unless otherwise configured in a policy.

Change Process

The password change interval and password complexity rules can be defined within a configuration policy. This is fully documented in the *CA Unicenter Desktop and Server Management Advanced Topics Guide*, available on the following link:

CA Support Online:
https://support.ca.com/phpdocs/0/common/impcd/r11/troubleshooting/doc/DSM_Adv_topics_r11.pdf.

Local Administrator

When a target PC is built, the local administrator password is defined as a parameter for that configuration. If the default value has not been changed, the local Administrator password for the target PC after completion of the build will be 'default.'

Agent Plugin Installation

Unless the InstallAgent parameter is set to 'No' the DSM Agent + Software Delivery Plugin will be installed on the target PC as part of the operating system installation.

Status Change to Current

Agent install=

This parameter controls when the status change for a specific configuration will change from Installing to Current.

Agent install = No

With this option, the status will change to Current the first time the target boots from hard disk after the boot images have completed their processes.

Agent install = Yes

With this option, the status will remain Installing until the agent plug-in has been installed. When the new agent registers with its scalability server this information is passed to the OSIM system and the status is then changed to current.

BMSConfigID

The BMSConfgID is the ID of the OS installation job in the database. The manager sends the BMSConfgID as a job parameter to the target.

After the installation of the common agent the agent (CSM plug-in) reads the BMSConfigID from c:\osimconf.ini and sends it with the report to the manager.

The manager compares the reported BMSConfigID with the job ID in the database. If it fits, the installation job becomes current. Otherwise we assume the installation has not finished correctly and it is set to Failed with the following error:

[OSM010018] A started OS installation has been aborted on the target.

C:\osimconf.ini is created during the execution of the custom.cmd file contained in the OSImage. If the Custom.cmd is customized, care must be taken to not change any of the entries relating to BMSConfigID or osimconf.ini. If these are modified, it can lead to all future OS installations failing because the agent will not be able to report the BMSConfigID to the Manager.

Packages

The packages used for the agent installation are located separately from those in the CA Unicenter Desktop & Server Management software library and those used by Infrastructure Deployment.

The store for these packages is:

    \CA\DSM\Server\SDBS\var\managedpc\agents

This store is automatically populated based on the content of the software library.

If the server has been upgraded to a new version through a patched master image or new CD, the packages registered in the library may need to be upgraded in order to ensure the new Agent package is in the Agent Store.

Use the dsmPush utility to update the packages in the software library. The dsmPush utility is fully documented in the *Command Line Reference Guide* contained in the CA Unicenter Desktop & Server Management documentation.

Remote Scalability Server

When the boot server is on a remote scalability server, the agent store is still populated based on the content of the software library.

Therefore, in order that the agent can be installed as part of an operating system installation, the DSM Agent + Software Delivery Plugin package must be staged to the software library of the remote scalability server.

Move An OS-Image Store to Another Disk on Scalability Server

If the Boot Server is not set to TFTP mode use sdbsswitch –t to do so

Caf stop

Copy the entire directory c:\program files\CA\DSM\server\SDBS\var\managedpc to another directory, for example, d:\newstore\managedpc

Modify the comstore with the new path

Ccnfcmda –cmd setParameterValue –ps /itrm/scalability_server/osim/ManagedPC –pn InstallPath –v d:\newstore\

Caf start

In order to create the OSIM shares correctly on the new location call Sdbsswitch –s

Boot Server on Microsoft Secondary/Backup Domain Controller

The r11.1 ,r11.2 ,C1 Boot Server can not be installed on Microsoft secondary Domain Controller.

Such Domain Controller does not provide a local user management, and therefore the Boot Server setup fails when creating the special OSIM user. In this case, the Boot Server is not completely installed.

Upgrading an Existing Operating System Using the Package Created by OSIM

The OS image created by OSIM is used to build or rebuild computers, and can also be used to upgrade them from their existing OS.

The registerosimage command not only makes the OS image available for the OSIM system, it also creates a software library item.

This has the following three procedures:

1. Add to boot server

2. Remove from boot server

3. Upgrade

Procedures 1 and 2 are used by the OSIM system. Procedure 3 can be used to upgrade your existing computer's OS.

To upgrade, simply deliver this procedure as you would any other software package. In order for the upgrade to be successful, there is one thing that has to be considered: how do you provide the correct product key to the new OS for the upgrade?

The upgrade procedure uses a standard Microsoft Windows-based unattended installation response file, just like the OSIM installation.

The response file for the upgrade update.inf is separate from the response file used by the OSIM system. It is created by the createosimage procedure in the i386 folder of the osimage. The simplest way to ensure the product key is configured correctly is to use the -k switch with createosimage. When this is used, it primes the default.ini of the OS image with the Product key. It also makes the following entry to update.inf:

    ProductID=$ProdKey$

When you then run the registerosimage command the product key is read from the default.ini file and the $ProdKey$ variable is replaced by the actual Product key from default.ini.

The upgrade procedure will now successfully upgrade the OS of the target computer.

If you do not use the -k switch then update.inf will have the entry:

ProductID=000-0000000

If this is left, then the upgrade will start but it will stick at the prompt for the product key. You therefore need to edit this prior to running the registerosimage command so that it contains the correct value.

This can either be the actual product key or the $ProdKey$ parameter. If you use this parameter then you will need to define the default value for ProductID in default.ini. See the OSIM documentation for details on editing default.ini.

Once the update.inf file is modified, you can run registerosimage and the upgrade will work.

If you have already run registerosimage and the product key is not correctly defined in update.inf, then you either need to register a new image after having made the modification above or modify the update.inf in the actual library image. You would modify the library image in the normal way you would modify any software library image.

Note: The upgrade procedure is subject to the same restrictions as a manual OS upgrade. The upgrade must be supported by Microsoft and you must have verified that all installed applications also support the upgrade of the OS.

For details of the supported upgrade path for your operating system, see the following articles from Microsoft.

Windows 2000 Upgrade Paths

Windows XP supported upgrade paths

Windows Server 2003 Supported Upgrade Paths

Windows Vista: Upgrade Paths from Previous Versions

Chapter 20: Deployment of Microsoft Office 2007

Introduction

The installation procedure of Microsoft Office 2007 has been substantially redesigned in this the latest version of the Microsoft Office solution. Some of the installer features that differ between previous releases and Microsoft Office 2007 are listed in the following table:

Previous Version	Microsoft 2007 Office Release
Windows Installer (Msiexec.exe)	Setup program (Setup.exe)
Administrative installation point	Local installation source
One MSI file per product	Multiple MSI files per product
Core English version plus MUI Pack	Language-neutral architecture
Setup.ini file	Config.xml file
Setup command line	Config.xml file
Custom Installation Wizard	Microsoft Office Customization Tool
Custom Maintenance Wizard	Microsoft Office Customization Tool
Microsoft Office Profile Wizard	Group Policy system policies

The changes made by Microsoft in the Microsoft Office 2007 installation process mean that the MSI (Windows Installer, formerly Microsoft Installer) integration within CA Unicenter Desktop & Server Management (CA DSM) is no longer the best method to use to deploy Microsoft Office 2007. CA DSM can be used to deploy Microsoft Office 2007 using the MSI approach that is also used when deploying with Microsoft Active Directory Group Policy, but due to the limitations this imposes we cannot recommend this approach.

The purpose of this document is to provide guidance on how Microsoft Office 2007 can be deployed by CA DSM using the existing generic capabilities of the product.

Prerequisites to Install Microsoft Office 2007 Using CA DSM

In order to use CA DSM to deploy Microsoft Office 2007, a number of prerequisites are assumed. They include the following:

■   An understanding of CA DSM 11.x

■   An understanding of the Microsoft Office Customization Tool (OCT)

■   The correct media for the edition of Microsoft Office 2007 you wish to deploy

Creating the MSP Files Using OCT

The Office Customization Tool (OCT) is used to create the Setup customization file that will be used to configure the installation process of Microsoft Office 2007. The high level steps necessary for the creation of the customization file are provided below. See the Microsoft Office website for details on the OCT layout:

1. Copy the entire contents of the office media to a temporary folder.

2. Start the OCT using the command 'setup.exe /admin' from the location where the office media is kept.

3. The OCT will appear.

Click the OK button to create a new setup.

4. The Welcome Screen will be displayed.

5. Select 'Install location and organization name' in the left hand pane and fill in the installation path and organization name in the right hand pane.

6. Select 'Licensing and user interface' in the left hand pane and enter the valid product key. Then check the box 'I accept the terms in the License Agreement' in the right hand pane.

About the Display Levels

There are three levels of display that are available to the administrator to select, but only None is recommended for an unattended installation:

■ **None**:

It installs Microsoft Office without any user interface.

■ **Full - Default**:

The installation proceeds with fewer interruptions and your customizations are set by default for all users.

■ **Basic**:

The Basic display level shows the user a welcome screen, a simple progress bar, and error messages, but does not require any input from the user during the installation. If you choose this display setting, you can select the following options on the same screen of the tool:

> Select the **Suppress Modal** check box to hide error messages and other dialog boxes that might interrupt the installation.

> Select the **No Cancel** check box to prevent users from cancelling the installation process before it completes.

> Select the **Completion Notice** check box to add a message at the end of the installation that lets users know that Microsoft Office 2007 has been installed.

7. Select 'Set features installation states' in the left hand pane, and select the features to be installed in the right hand pane.

8. Select and make any further changes as necessary for your environment and use of Microsoft Office 2007.

9. Select the File\Save As dialog in the menu bar of OCT to save the file with customized settings.

10. Using the OCT, define the customizations required for the deployment of Microsoft Office 2007 in your environment. You can create as many customizations as you need for your different requirements. Once all required customizations have been defined, use the File\Save As menu item in the OCT to save the MSP file into a subfolder called Install under the folder created in step 1.

Creating the CA Unicenter Software Delivery Package

Multiple methods exist for the creation and delivery of software packages. In this section, we will show an example of using a generic package from the CA Unicenter Software Delivery Library. For this example, it does not matter if the CA Unicenter Software Delivery Agent is configured to use shares (NOS), File transfer (NOS-LESS), or DTS (NOS-LESS Data Transport Service).

The Microsoft Office 2007 installation process creates a cache of the installation files that it can refer back to at any time after the installation has completed. The cache is often referred to as the MSOCACHE and is created automatically by the installation. The use of the cache means that access to the original source files is often only required if the local MSOCACHE becomes corrupted (for some reason!).

Follow these steps to create the CA Unicenter Software Delivery package:

1. From the DSM Explorer, Software Package Library, select 'New SW package.'

2. The name and version for this package must use the correct values as supplied by Microsoft. This is because the setup program is still using the MSI installer, so using the same name helps with MSI software detection. Enter the correct details and click OK.

3. This will create a new package that is Open (unsealed).

Add the Source Files to the Package

1. Double click this package and select Source.

2. Select New Volume, enter a name, and browse to the temp folder which contains the office source files. Click Choose. Do not check Source is on Manager. Click OK.

3. Wait while the files are copied into the package source.

4. When completed, the source will be displayed.

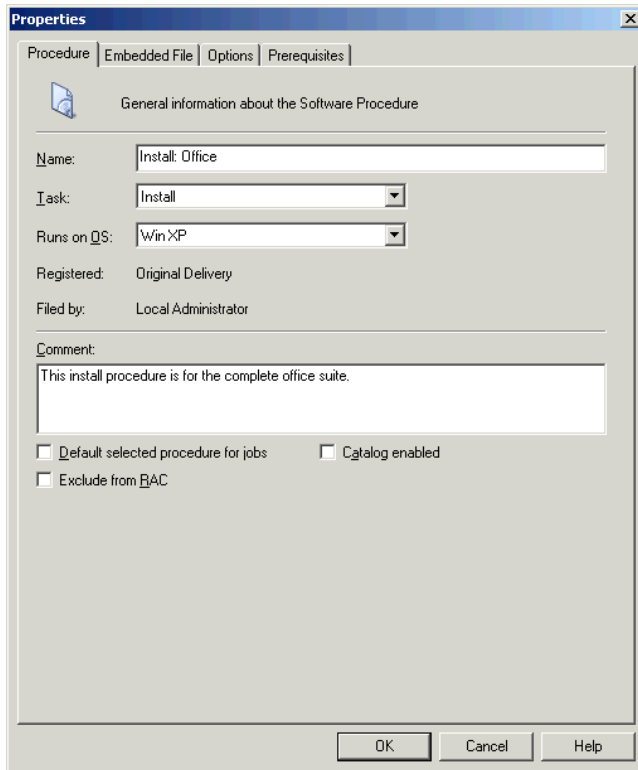## Create the Installation Procedures

1. Expand the source until the setup.exe can be seen:

2. Right-click setup.exe and select Copy.

3. Right-click on Procedures and select Past\Embedded file to Register Procedure.

4. Define the Name of the procedure and the Runs on OS value if required.

5. Select the Embedded File tab, and leave the pre-filled fields as-is. If you are using a single Item Procedure and if the MSP files are in the Update directory of the package, then the /adminfile is not necessary. In this case, the setup.exe will automatically process ALL msp files found in this directory. Otherwise, for multiple item procedures for different customizations, the additional command line parameters required for the setup.exe to use the specified msp file created by the OCT must be specified.

6. Because of the nature of this new installation method, it is no longer possible for CA DSM to use its inbuilt MSI support, which knows about the non-zero MSI return codes that indicate success. As it is possible for the installation of Microsoft Office to require a reboot, MSI may return a 1641 return code. To prevent the software job from being seen as failed when this code is returned, a further macro needs to be added to the parameters entered in step 5: $#ec:1641 $#ec:0. This instructs the agent to ignore this return code and allows the job to be marked successful.

7. Click OK. The procedure is now created.

8. Repeat these steps to create individual Item Procedures for each of the msp files created with the OCT.

Create the Uninstall Procedures

The command to uninstall Microsoft Office is *setup.exe /uninstall <product id>*. However, this command does not run silently and requires the end user to confirm that they wish Microsoft Office to be uninstalled. To avoid this, a config.xml file can be used in a similar fashion to the msp for the installation.

1. The first step is to identify the correct product id required for the uninstall command.

   This is found in the setup.xml file from the folder xxx.WW in the Microsoft Office source image, where XXX is the Microsoft Office product you are installing (Enterprise).

2. Open this file and find the Setup Id= value:

```xml
<?xml version="1.0" encoding="utf-8" ?>
<!--
  SIG=m2hhh+sBXmY5rmWoXsCXTO3rXKGukaTW3shasPaWxH1P6s6rpY9PU9KbCOfDUNn+TnlmG020Xw8+KzV2HX7iIAAoMSSVLvFOyoYNVEq
- <Setup Id="Enterprise" Type="Product" ProductCode="{90120000-0030-0000-0000-0000000FF1CE}">
    <PIDTemplate Value="89388&lt;````=````=````=````=`````&gt;@@@@@" />
    <Option Id="AlwaysInstalled" DefaultState="Local" DisallowAbsent="yes" DisallowAdvertise="yes" Hidden="yes" />
    <Option Id="Gimme_OnDemandData" DefaultState="Local" DisallowAbsent="yes" DisallowAdvertise="yes" Hidden="yes" />
```

Note: In this example the product ID is Enterprise (this is the Enterprise version of Microsoft Office 2007).

Create the config.xml

1. This config.xml file will contain code similar to the following:

```xml
<Configuration Product="Enterprise">

<!-- <Display Level="full" CompletionNotice="yes" SuppressModal="no"
AcceptEula="no" /> -->

<!-- <Logging Type="standard" Path="%temp%" Template="Microsoft Office
Professional Setup(*).txt" /> -->

<!-- <PIDKEY Value="BCDFGHJKMPQRTVWXY2346789B" /> -->

<!-- <USERNAME Value="Customer" /> -->

<!-- <COMPANYNAME Value="MyCompany" /> -->

<!-- <INSTALLLOCATION Value="%programfiles%\Microsoft Office" /> -->

<!-- <LIS CACHEACTION="CacheOnly" /> -->

<!-- <SOURCELIST Value="\\server1\share\Office12;\\server2\share\Office12" /> -->

<!-- <DistributionPoint Location="\\server\share\Office12" /> -->

<!-- <OptionState Id="OptionID" State="absent" Children="force" /> -->

<!-- <Setting Id="Reboot" Value="IfNeeded" /> -->

<!-- <Command Path="msiexec.exe" Args="/i \\server\share\my.msi" QuietArg="/q"
ChainPosition="after" Execute="install" /> -->

</Configuration>
```
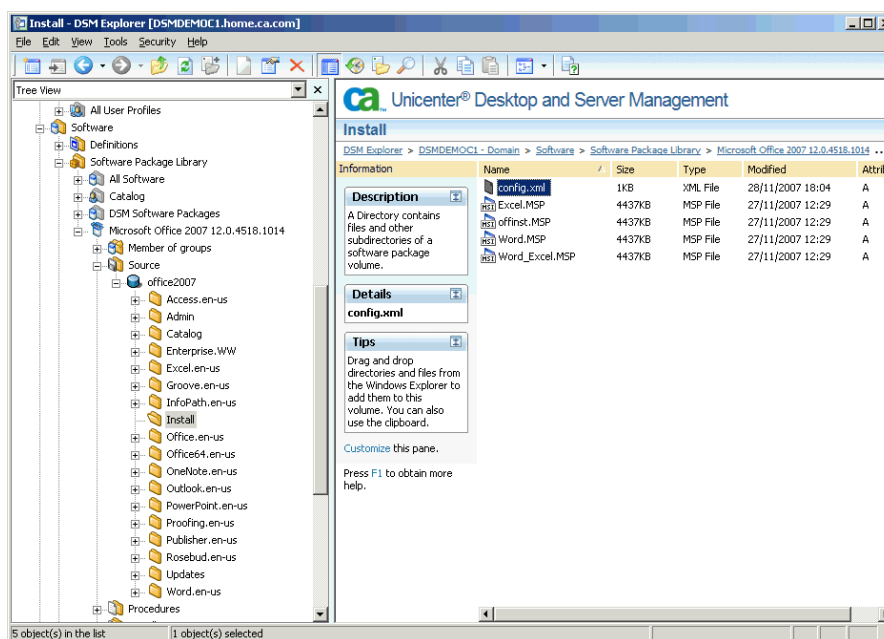
2. To make the Uninstall completely silent, modify the second line to read:

```xml
<!-- <Display Level="None" CompletionNotice="no" SuppressModal="yes"
AcceptEula="no" /> -->
```

3. Finally, uncomment this line so that the end result is:

```xml
<Configuration Product="Enterprise">

<Display Level="None" CompletionNotice="no" SuppressModal="yes" AcceptEula="no"
/>

</Configuration>
```

4. Save this file.

5. Right-click the file and select Copy.

6. From the DSM Explorer, open the Source of the Microsoft Office 2007 package and select the installation directory. Then right-click and paste the config.xml file into the package.
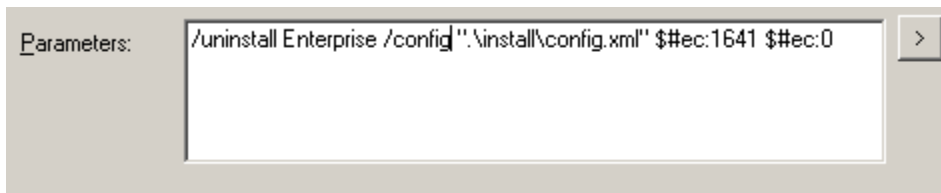

Create the Uninstall Procedure

1. Copy and paste the setup.exe from source to procedures as per steps 2 and 3 of the Create the Install Procedures paragraph above.

2. Name the procedure and specify the task Uninstall, and define the Runs on OS, if required.

3. In the parameters dialog of the Embedded Files tab, enter the uninstall parameter and the /configfile parameter in the same way as used when creating the installation procedures.

Be sure to direct the /configfile to the config.xml file created in the previous section. It is also necessary to add the $#ec:1641 $#ec:0 as before to allow for any reboot requested by the uninstall.

| Parameters: | /uninstall Enterprise /config ".\install\config.xml" $#ec:1641 $#ec:0 | > |

Create the Repair Procedure

1. Copy and paste the setup.exe from source to procedures as per steps 2 and 3 of the Create the Install Procedures paragraph above.

2. Name the Procedure and specify the task Configure, and define the Runs on OS, if required.

3. In the parameters dialog of the Embedded Files tab, enter the repair parameter and the /repair parameter in the same way as used when creating the installation procedures.

   Be sure to direct the /repair to the config.xml file created in the previous section. It is also necessary to add the $#ec:1641 $#ec:0 as before to allow for any reboot requested by the repair.

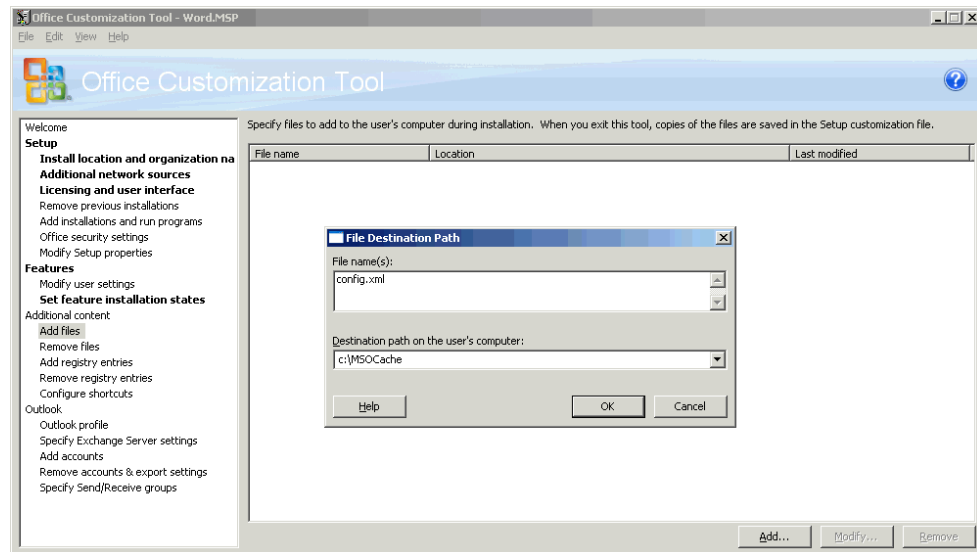| Parameters: | /repair Enterprise /config ".\install\config.xml" $#ec:1641 $#ec:0 | > |

Creating External Uninstall and Repair Procedures

The Item Procedures we have created so far are of the type Internal. Internal item procedures use files that are contained in the package, so the complete package must be visible to the agent in order for the procedure to run. In the case of an agent configured to use NOS-Less access to the CA Unicenter Software Delivery library, the full Microsoft Office package is always delivered to the target computer. If this is an Uninstall or Repair operation, then this is only necessary if the MSOCache is no longer available on the target computer.

An External item procedure uses files that are already on the target computer, so no files from the CA Unicenter Software Delivery Library are made available to the target computer with this type of procedure.

Not all of the files needed for an unattended uninstall or repair of Microsoft Office 2007 are automatically copied to the target system. The file config.xml needs to be added to the office installation using the OCT.
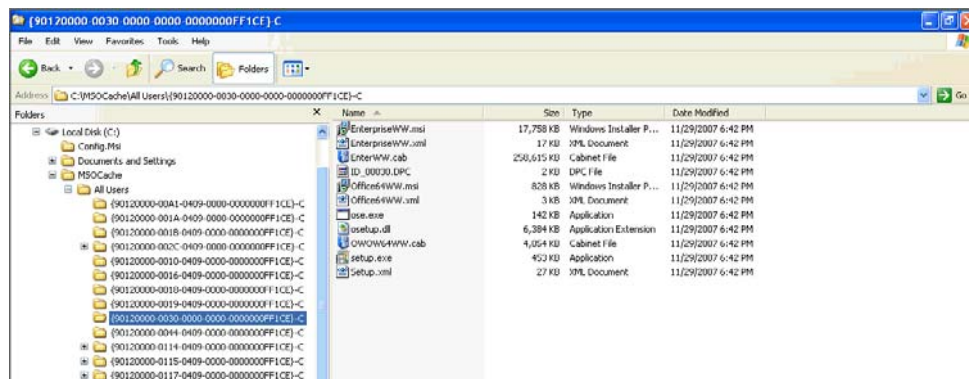
From the OCT select the option Additional content/Add files. The following example assumes that the MSOCache is located in its default location, c:\MSOCache.
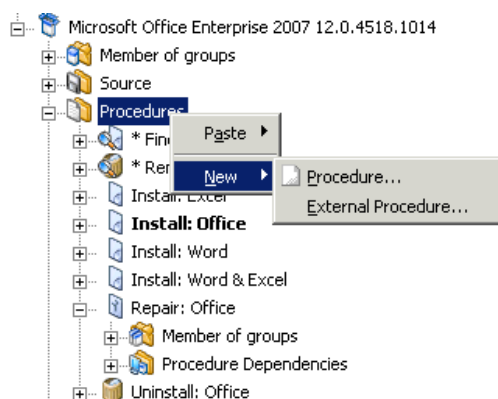

The creation of the External Item Procedure is very similar to that of an Internal Item Procedure.

The key information needed for item procedure creation is the location of the Setup.exe and config.xml files.

The setup.exe is located in the Microsoft Office products cache directory. The directory name may differ depending on the edition of Microsoft Office 2007 being deployed.
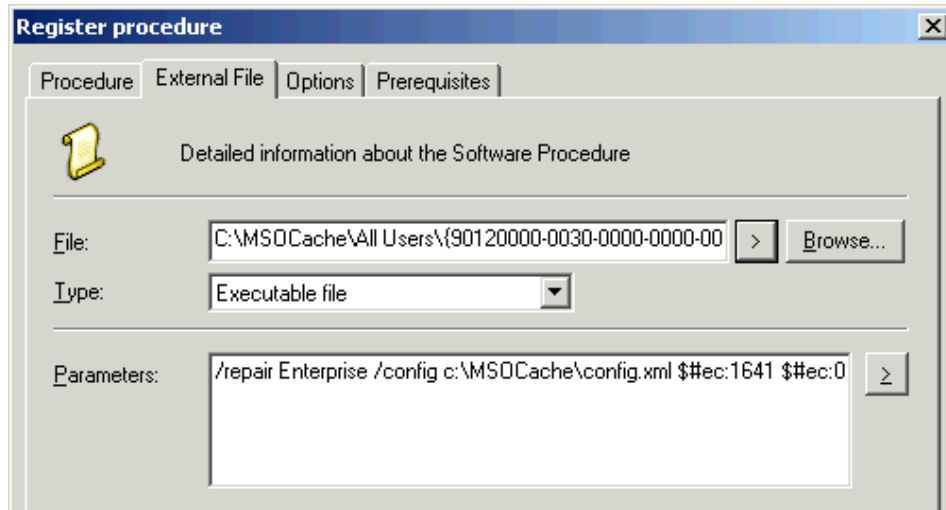
From the DSM Explorer we need to select the option to create an External Procedure.
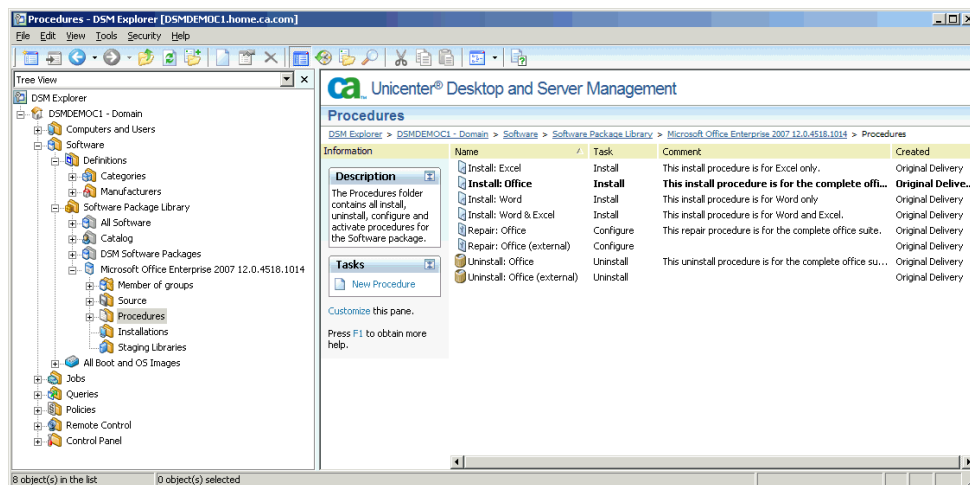

As done previously, name your procedure and chose the procedure Task. Then click on the External File tab.

In the parameters dialog of the External Files tab, enter the full patch and file name for setup.exe. Also enter the repair parameter and the /repair parameter in the same way as used when creating the installation procedures.

Be sure to direct the /repair to the config.xml file located in the c:\MSOCache directory. It is also necessary to add the $#ec:1641 $#ec:0 as before to allow for any reboot requested by the repair.

You can continue to add as many Item Procedures as necessary to manage the different customization options needed for the different users of Microsoft Office 2007 within your organization.

Using Customized Installs for Microsoft Office 2007with a config.xml File

Microsoft also provides a method of installing Office 2007 with customized settings held in XML rather than MSP files. According to Microsoft Office Customization Guide, the XML configurations do not make available all of the settings that the OCT has when producing MSP files. Care must be taken when editing any XML files to ensure that the file structure remains intact.

A similar procedure can be followed for the installation of Microsoft Office 2007 using MSP files through CA DSM. Slight modifications are needed to use the correct command switches to the setup program.

In the parameters sub-pane of the Embedded Files tab, enter the /config parameter and the config file name in the same way as used when creating the installation procedures. Be sure

to direct the /configfile to the config.xml file that you have created and placed in the installation directory of the Microsoft Office 2007 package. It is also necessary to add the '$#ec:1641 $#ec:0' as before to allow for any reboot requested by the installation.

Parameters: `/config ".\install\config.xml" $#ec:1641 $#ec:0`

Deploy the Package to Target Agents

To deploy the Microsoft Office package to target agents, follow the standard procedure for deploying software by selecting the required procedure from those created. The package and procedures can also be catalog-enabled to allow users to perform self service Microsoft Office installations, if required.
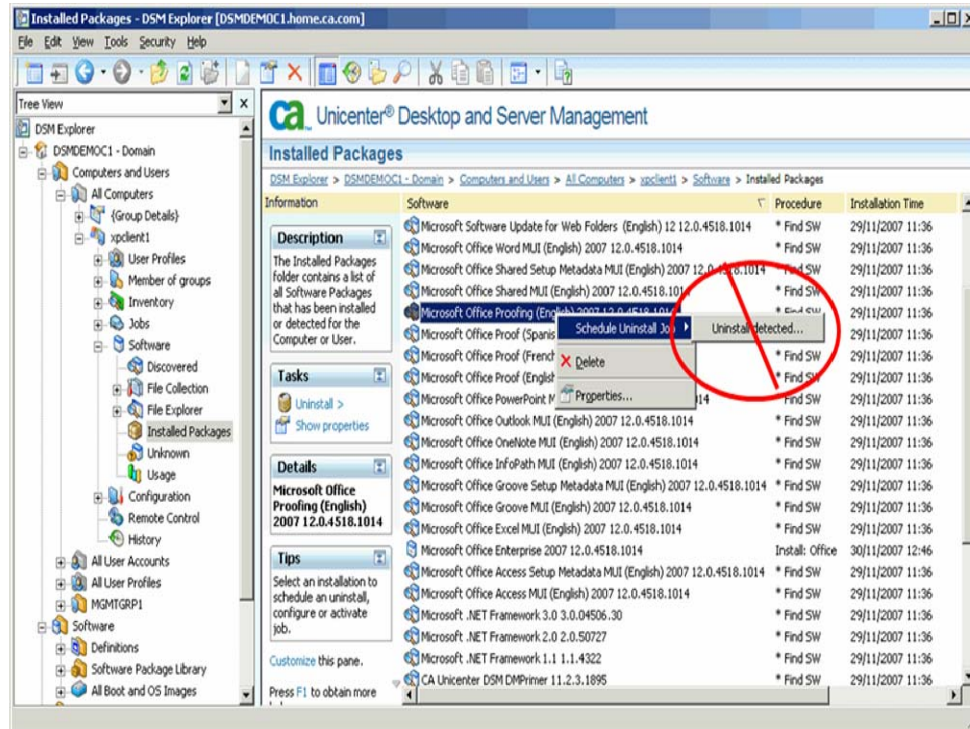
Uninstalling from Target Agents

The uninstall procedure can be deployed in the standard way either by selecting from the software library or by right-clicking the installation record and selecting the Uninstall option. The uninstall can also be catalog-enabled in the standard fashion if required.

Detecting MSI Applications

The Microsoft Office 2007 installation utilizes MSI. This enables the CA Unicenter Software Delivery Scan MSI function to detect and show the listed products for the target computer, even if the application was not installed by CA Unicenter Software Delivery. The Microsoft Office 2007 solution consists of many individual products which are individually registered to MSI. This function is useful for updating an agent's installed products as this table drives the Software Policy automation functions.

CA Unicenter Software Delivery automatically adds an MSI Uninstall procedure for the detected applications—but for Microsoft Office 2007 this must NOT be used. If you need to remove a product from an installed Microsoft Office suite, or the complete suite, then a custom MSP file can be added to the package and delivered in the normal way.

Index

2

2D Map
  discovered objects • 134

A

agent privileges • 185
agent registration • 173
Asset Collector
  configuring • 216
  default folders • 214
  overview • 211
  signed and unsigned files • 216
  SMS Connector • 218
  usage examples • 217
  XML files • 212
asset registration and reconciliation
  asset classes • 94
  Asset Registration Process • 83
  common asset registration • 85
  Common Object Registration API
    (CORA) • 86
  discovered assets • 89
  engine registration • 84
  HostUUID • 84
  overview • 83
  owned assets • 89

C

cfsystray • 177
classic discovery
  advanced • 139
  best practices • 134
  command line • 144
Common Asset Viewer (CAV) • 95
Common Object Registration API (CORA)
  • 86
configuration policies
  custom policy • 245
  Default Computer Policy • 245
  DTS • 60
  hierarchical rules • 246
  overview • 245
  Software Delivery • 63

D

Data Transport Agent (DTA)
  configuration • 63
  overview • 59
Data Transport Service (DTS)

  components • 58
  configuration • 60
  overview • 57
  Software Delivery configuration • 63
discovery
  automating discovery events • 132
  best practices • 134
  by command line • 144
  classic discovery wizard • 137
  Common Traffic Analyzer • 127
  configuration • 124
  configuration examples • 126
  continuous • 127
  firewalls • 127
  fixed IP addresses • 127
  IP discovery • 133
  methods • 135, 141
  network devices • 121
  overview • 122
  preparation • 137
  processes • 136
  restricting discovered devices • 130
  subnet filtering • 135

E

engine registration • 84

I

installation language • 174, 176
integration
  overview • 20
inventory modules • 200, 201

L

logoff process • 181
Logon Shield • 186

N

Network Object Server (NOS)
  configuration • 61
  overview • 58
New Task Wizard • 219

P

policy rules • 246

R

reboot process • 181

S

Schedule Object Server (SOS)
   overview • 59
SMS Connector
   engine job • 219
   output • 221
   overview • 218
   setting up • 219
subnet filtering • 135
subnet mask • 140

T

Transfer Object Server (TOS)
   configuration • 62
   overview • 58