HP Asset Manager

Software version: 9.30

Tuning

Document Release Date: 31 March 2011 Software Release Date: March 2011



Legal Notices

Copyright Notices

© Copyright 1994-2011 Hewlett-Packard Development Company, L.P.

Restricted Rights Legend

Confidential computer software.

Valid license from HP required for possession, use or copying.

Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Warranty

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services.

Nothing herein should be construed as constituting an additional warranty.

HP shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

For information about third-party and/or open source license agreements, or to view open source code, use one of the following ways:

- In the ThirdParty directory of the software installation CD-ROM
- In the directories in which the binary files of the third-party and/or open source tools are located after installation of the software.
- Through the component's url indicated in the Open Source and Third-Party Software License Agreements guide

Trademark Notices

- Adobe®, Adobe logo®, Acrobat® and Acrobat Logo® are trademarks of Adobe Systems Incorporated.
- Corel® and Corel logo® are trademarks or registered trademarks of Corel Corporation or Corel Corporation
 Limited
- Java is a registered trademark of Oracle and/or its affiliates.
- Microsoft®, Windows®, Windows® XP, Windows Mobile® and Windows Vista® are U.S. registered trademarks of Microsoft Corporation.
- Oracle® is a registered trademark of Oracle Corporation and/or its affiliates.
- UNIX® is a registered trademark of The Open Group.

Acknowledgements

This product includes software developed by the Apache Software Foundation (http://www.apache.org/[http://www.apache.org/]), which is Copyright © The Apache Software Foundation. All rights reserved.

This product includes software developed by The OpenLDAP Foundation, which is Copyright ©, The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. OpenLDAP® is a registered trademark of the OpenLDAP Foundation.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (http://www.openssl.org/), which is Copyright © The OpenSSL Project. All rights reserved.

This product includes software developed by the OpenSymphony Group (http://www.opensymphony.com/), which is Copyright © The OpenSymphony Group. All rights reserved.

This product includes code licensed from RSA Data Security.

This product includes software developed by the JDOM Project (http://www.jdom.org/), which is Copyright © Jason Hunter & Brett McLaughlin. All rights reserved.

Table of Contents

Introduction	5
What is the aim of this guide?	5 6
Chapter 1. Tuning a client	9
Network latency Tree view and table view Loading lists Type-ahead Access restrictions Properties of the objects Features List configuration	9 11 12 13 14 16 17
Chapter 2. Tuning the database	19
About hardware	19 21 21 22

Index	Inde	ex																										3
-------	------	----	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	--	---

Introduction

What is the aim of this guide?

This guide describes tuning strategies for Asset Manager. In particular, it deals with certain techniques to reduce bottlenecks caused by:

- the network,
- the database server,
- the Asset Manager client,
- the database engine

Who is this guide intended for?

This guide is mainly intended for Asset Manager administrators and database administrators. It contains information and procedures, which if misinterpreted, misapplied or implemented badly may provoke a range of problems from performance issues to data corruption.

Warning:

This document is intended to assist administrators in identifying and resolving performance-related issues within the software and database. However, it is not a comprehensive database-tuning guide, and although it is intended to contain as much information about tuning as possible, it is not guaranteed that the subjects mentioned herein are the only tuning options available. Also, HP cannot be held liable for any damages or data loss incurred as a result of improper use of the information provided herein. If you do not fully understand a procedure or how a procedure will impact your installation or information system, you are advised not to execute the procedure. As always, HP recommends that you perform a complete backup of your system and installation before making any major changes to it.

Reminder concerning Asset Manager architecture

Asset Manager relies on a two-tier architecture. The Asset Manager client connects directly to the database and falls into the category of what is usually called a "fat" client. A "fat" client resides completely on the client computer. The advantage of this kind of architecture is better performance on a fast LAN. The disadvantage is that the performance gets geometrically worse on a slow network. What this means is if your network latency is twice as bad as the recommended (10 milliseconds), your actual performance decrease will be approximately four times worse.

One of Asset Manager's key components is Asset Manager Automated Process Manager. Asset Manager Automated Process Manager does not exchange any data with the client. It connects to the database separately, and performs all of the background tasks that require regular triggering and/or monitoring. For example: an alarm that is set for a contract's end date and that issues an email message 30 days prior to that date. Asset Manager Automated Process Manager also processes workflows, calculates rents, and validates the authorization code in the database. Most of the tuning suggestions made in this document do not apply to Asset Manager Automated Process Manager.

Bottlenecks

There are basically four areas of Asset Manager where a performance bottleneck can occur: the network, the client, the database server, and the database engine. Network bottlenecks usually occur when you are utilizing Asset Manager on a slower network than it was designed to utilize.

Client bottlenecks are often actually database or network bottlenecks. Usually what happens is that the queries take a long time to return data (either because of a slow network or because it takes a long time for the database to respond). Database server bottlenecks occur when the hardware that the database server is running on is insufficient for the implementation of Asset Manager. Typically this is not a problem area because most implementations analyze their requirements in advance and make sure the hardware is good enough, but once in a while this can be a problem.

Database engine bottlenecks are the most common. What happens in this case is the database engine is not able to respond rapidly to the queries being sent to it. This can happen for a variety of different reasons; poor database tuning, complex queries, bad optimizer choices, or insufficient resources...



One thing to remember is just because you have a powerhouse of a server doesn't mean that the database is configured to take advantage of all of that power. Also, if you have several different applications sharing a single database server, that means the resources are going to be split among them, and that can affect performance.

Contents of this guide

Chapter Tuning a client

This chapter explains how to tune client computers.

Chapter Tuning the database

This chapter deals with tuning the database.

Conventions used in this guide

What follows is a list of the conventions that we use in this guide

Convention	Description
Java Script Code	Example of the code or command
Fixed-width characters	DOS command, function parameter or data
	format
	Portion of omitted code or command.
Note:	Informative note
Extra information	

Convention	Description
Important:	Important information for the user
Be careful	
Tip:	Tip to help you use the application
User tip	
Warning:	Extremely important information for the user
Exercise caution	
Object	Asset Manager interface object: menu, menu
	entry, tab or button.

The following conventions are also used:

- The steps that we ask you to follow in a specific order are shown as a numbered list. For example:
 - 1 First step
 - 2 Second step
 - 3 Third and last step
- All the figures and the tables are numbered according to the chapter in which they are found, and the order in which the appear in the chapter. For example, the title of the fourth table in chapter two will be prefixed by **Table 2-4**.

Tuning a client

This chapter explores the different performance-optimization strategies for Asset Manager. It discusses certain internal mechanisms of Asset Manager that can impact performance and provides troubleshooting tips.

Network latency

As we mentioned, the network has a direct impact on application performance. The recommended network latency for Asset Manager is 10 milliseconds or less.

Tree view and table view

Asset Manager enables you to view records in two distinct modes:

- A tree view of the records. This view shows you the hierarchy of the table, allowing you to unfold the different leaves and easily view the parent-child relationships being maintained in that table. Any table that has a **FullName** field can be viewed in tree view. Some examples are Employees/Departments (amEmplDept), Portfolio Items (amPortfolio), Locations (amLocation), etc.
- A list view of the records without any hierarchy information.

The way the tree view list is built is very costly in terms of performance. If you are on a fast network with a table that is not very large, tree view will work fine. But the larger the table and the slower the network, the worse that performance is going to get.

Two tuning methods are immediately available to you:

- 1 Viewing the tables in the list view if the network is slow and/or the table is of considerable size.
- 2 Modifying the default display parameters of hierarchic lists so that they are initially collapsed. This speeds up the initial screen load. To do this, set the Initially collapsed option in Asset Manager to Yes. (Edit/ Options/ Lists/ Main lists and Edit/ Options/ Lists/ Other lists).

List-loading mechanism

When the **Initially collapsed** option is enabled, Asset Manager uses the following mechanism to display the lists (the following example concerns the opening of the **Departments and Employees** table in tree view):

1 Asset Manager sends a query to recover all the records at the top of the hierarchy (those for which the **sLvI** field is set to **0**) and whose **Full name** (**FullName**) field is greater than the first record in the list. The query sent to the database is the following:

```
SELECT E1.lIconId, E1.lEmplDeptId, E1.MrMrs, E1.Name, E1.FirstName, E1.
Title, E1.Phone, E1.Fax, E1.FullName FROM amEmplDept E1 WHERE E1.sLvl=0
AND E1.FullName >='/Admin,,ADMIN/' ORDER BY E1.FullName
```

2 When a branch is unfolded, a new query is sent to the database. This query recovers the records directly underneath and whose full name contains the name of the parent. For example (in this case, the Taltek branch of the tree is unfolded):

```
SELECT E1.lIconId, E1.lEmplDeptId, E1.MrMrs, E1.Name, E1.FirstName, E1.
Title, El.Phone, El.Fax, El.FullName FROM amEmplDept El WHERE El.FullNa
me LIKE '/Taltek/%' ESCAPE '\\' AND E1.sLvl=1 ORDER BY E1.FullName
```

Each time a branch is unfolded, a query of this type is sent by Asset Manager. If the record you are looking for is lower down in the tree, the amount of queries sent to the database is increased.

On the other hand, the same table opened in list view results in one single query as follows:

```
SELECT E1.lIconId, E1.lEmplDeptId, E1.MrMrs, E1.Name, E1.FirstName, E1.Tit
le, E1.Phone, E1.Fax, E1.FullName FROM amEmplDept E1 ORDER BY E1.lEmplDept
```

All the records returned by the query are used to populate the list, the size of which is limited as defined in the Asset Manager options.

Loading lists

Several options are available to define the number of records initially loaded in a list. By reducing or limiting the this number, you can improve response time when displaying a list. The options in question are located in the **Edit/Options/Lists/Main lists** and **Edit/Options/Lists/Other lists** sections. The available options are:

- **Do not load for longer than**: This option defines the maximum time (in milliseconds) during which Asset Manager tries to populate the list. When this time is overrun, the operation is suspended and the list if populated with those records recovered. By default, this value is set to 5000 milliseconds.
- **Do not load more than**: This option defines the maximum number of records loaded in a list. By default, 200 records are loaded.



This number also corresponds to the number of additional records that are loaded when clicking the + icon next to list.

The default values of these two options appropriate for the vast majority of cases. However, in specific circumstances when the list configuration is complex, these values may not suit you in terms of performance. If performance issues are encountered displaying lists, you may wish to:

- Reduce the value of the **Do not load more than** option to 30 or 50, for example.
- Change the value of the **ArrayFetchingSize** parameter in Asset Manager. This parameter specifies the number of records recovered by Asset Manager in each network packet. For performance reasons, this value must be the same as that defined in the **Do not load more than** plus one. For example, if the value of the **Do not load more than** option is 200, the **ArrayFetchingSize** parameter must be set to 201. This incrementation accounts for the NULL record, which is returned in the results of the query but not displayed by Asset Manager.

To modify the value of the **ArrayFetchingSize** parameter:

- 1 Load the amdb.ini file in a text editor,
- 2 In the section corresponding to the declaration of the database connection, find the **ArrayFetchingSize** parameter. The following is an example for the **AMDemo93en** database:

[AMDemo93en]
Engine=SQLAnywhere
Location=AMDemo93en

```
EngineLogin=itam
EnginePassword=7BC2423F1B1F6A42E2B3A27E396F52C3A9AD6828582AED88DAC2F
53E9A369BC03E6393911903124254200200
ReadOnly=0
CacheSize=5120000
ArrayFetchingSize=201
AmApiDll=C:\PROGRA~1\HP\ASSETM~1\ASSETM~2\bin\aamapi93.dll
```

3 Modify the value of the parameter and save your changes.



Changing this parameter may also prove useful when Asset Manager is used over a WAN (Wide Area Network).

Type-ahead

By default, Asset Manager uses type-ahead functionality to help users populate a link field. It works by waiting for a certain amount of time (which by default is after 500 milliseconds, or half a second) after a user stops typing in a link field, and then queries the database to find the first record that matches what the user has typed so far.

Most of the time this can be very handy; for instance, if you're typing in a person's name as a user of a portfolio item, it can help you populate the link after typing only a few letters of the person's last name. The bad side, however, is that once the type-ahead starts to execute the user will have to wait until the query returns before they can continue to type. In situations where there is a slow network or where the table the link connects to is large, this can be very expensive. In these cases it is recommended to disable type-ahead. To do this:

- Select the Edit/ Options/ Navigation/ Selection of linked records/ Type-ahead after menu,
- 2 Click the number located next to this option,
- 3 Change it to a very large number (for example 1,000,000) to prevent type-ahead from ever being triggered,
- 4 Click OK.

Access restrictions

An access restriction is part of an Asset Manager user profile. It corresponds to a record filter on a table. For example, you can make sure technicians only have access to the assets in their own department.

Because they are a designed to provide row-level security, access restrictions are commonly used, and rightly so. However, because they affect the query that gets sent to the database when a user opens a screen, a poorly written or poorly defined access restriction can cause all kinds of performance issues.

Access restrictions are always in the context of a specific table, and are broken down into two areas: Read and Write. Read access restrictions prevent users from seeing a particular set of records. By corollary, they also prevent users from modifying certain sets of records. Write access restrictions prevent users from modifying certain records. They do not affect their ability to read any data in the record, they simply receive an error message whenever they try to make a change to that record.

Access restriction mechanism

Read restrictions are applied via a WHERE clause added to the query sent to the database. For example, lets take an access restriction that a user may only view cost centers for which they are the supervisor. The query sent is as follows:

SELECT C1.lCostId, C1.Title, C1.AcctNo FROM amCostCenter C1, amEmplDept amEmplDept CurrentUser WHERE amEmplDept CurrentUser.lEmplDeptId = C1.lS upervId AND amEmplDept CurrentUser.lEmplDeptId = 25323 ORDER BY C1.lCos

This typical and simple case is not very costly in performance terms. But if the query is more complex, and in particular if it contains one or more sub-queries, the database engine will take more time in processing and sending the result to Asset Manager, which will result in a slower display. Once more, the size of the table concerned by the query plays an important role.

Write restrictions work in a slightly different way. When a user tries to modify a record in a table with one or more write access restrictions, Asset Manager sends a separate SELECT query to the database specifying the access restriction in the WHERE clause. If the record that the user is trying to modify is not returned, a message informs the user that they do not have the required permissions and the changes are cancelled at the database level. If the record is returned, an UPDATE statement is sent to the database and the modifications are validated. The network traffic is potentially twice as high (a test query plus a modification query). In addition, as for read restrictions, the complexity of the queries and the size of the tables concerned also plays a role in reduced performance.



It is therefore crucial to carefully analyze how you require access restrictions to be applied and to implemented them in the most succinct and rational manner possible.

Properties of the objects

The properties of database objects (fields, links...) may have a significant impact on performance.



These properties can been seen via the Configure object shortcut menu or by using Asset Manager Application Designer.

The display properties

The Asset Manager database objects, fields and links in particular, have properties that determine how they are displayed. These properties are the following:

- **Mandatory**: The label of this field is shown in red and the field must be populated in order to be able to validate the record,
- **Read only**: The field is grayed out,
- **Irrelevant**: The field is not displayed.

These properties may be set to the following values:

- Yes
- No
- Script

In this last case, an interpreted Basic script determines the final value of the property.

A value of Yes or No for a property does not impact performance because it is not required to be evaluated. This is not the case for a script. When a script is loaded or updated, all the scripts associated with the record are executed. Simple scripts that work on local data are relatively transparent in terms of performance. On the other hand, scripts that query the database (for example, when using **AmDbGetLong()** type functions) will generate queries and have

a direct impact on performance. It is the same when the value of a field or a remote link is referenced in a script, for example:

[CostCenter.Supervisor.Location.City]

To resume:

- Analyze clearly the need to script a property,
- Bear in mind that all scripts that perform database queries may impact performance,
- Optimize your scripts for efficiency.

History

Asset Manager makes it possible to keep history of modifications made to database objects. This functionality is controlled by the **Keep history** property of fields and links. The history mechanism is as follows:

- 1 A user modifies a record. If any object concerned by this modification has a **Keep history** properties that evaluates to 1 (i.e. it is set to **Yes** or a script returns the value 1), the history agent is triggered.
- 2 The history agent creates a record in the History table (**amHistory**) storing both the old and new values of the object.

As for the **Mandatory**, **Read only** and **Irrelevant** properties previously mentioned, the **Keep history** property may cause performance issues if a complex script (querying the database) is associated with it. Another impact to performance is directly linked to the fact that an insert statement is sent to the database for each historized object. In ideal cases, if the number of historized objects is limited this will not cause any particular problems. This is not the case when dozens of historized objects are modified. Another factor in reduced performance is linked to the fact that the amount of time it takes for the query to complete is directly related to the size of the table. The more records there are in the table, the longer it takes. And typically if someone is keeping history on several fields, the amHistory table can get very, very large, in some cases upwards of 10 million rows.

We therefore recommend:

- For best performance, only keep history on those fields you absolutely have to. Do not historize fields systematically.
- Choose an amount of time after which you can delete history. You can write a workflow to delete history after 6 months, for example.

Validity check

Validity scripts are used by Asset Manager to check certain conditions of records before inserting them into the database. For example, this type of script is used to check that an end date in a contract does not come before the start date. If this occurs, the record is not inserted and an error is raised.

Scripts are subject to the same performance constraints as previously mentioned for the properties of objects.

Features

Features are a historical functionality of Asset Manager introduced prior to version 4.0.0. They provided the ability to define additional information for records when it was not possible for users to extend the database schema. Despite their power, features are often costly in performance terms. Taking the Locations table as an example, three tables are impacted when using the features, as illustrated below:

The Locations table (amLocation) does not actually hold any information about features. It is simply the target of the Feature Values (amFVLocation) table, which holds the values of the features associated with the Locations table. The Features (amFeature) table itself contains the physical data on the feature, such as its SQL name and Label, etc.

Each feature you assign to a record causes at least one extra join to be executed to pull in that data. The impact on performance may be significant, especially if the feature is displayed in the list.

Since Asset Manager 4.0.0, it is now possible to modify the database schema directly. Using features is less justifiable. However, you may still wish to do so in the following two cases:

- 1 When the piece of information is going to be on very few records (1% or less, but may vary depending on implementations), you may wish to use a feature rather than create a new field in the table. A feature only takes up space when it is populated (the only notable exception being when the **Force display** property is on). An additional field takes up space regardless. Under these circumstances, using a field may use up a large amount of storage space unnecessarily.
- 2 Modifying the database structure is relatively straightforward but can be risky. If something happens during the structure update that causes it to stop, your database will be in an inconsistent state and you will have to restore from a backup. For this reason HP recommends:
 - Backing up the database before modifying the structure,
 - Carefully planning changes and having them performed by personnel with the required competences.

During the planning stage, you may with to temporarily use feature before permanently creating a field to link to perform the same role.

List configuration

A majority of performance problems experienced when using the Asset Manager client are due to poor list configuration choices. When a user opens a screen, Asset Manager executes a query to pull in all of the list data. The query is dynamically built based on several settings (the columns in the list, sorts, filters, etc.), read access restrictions, table/tree view, etc.

The choice of columns in the list is particularly important as they are used in the SELECT clause of the query. Performance is optimal when querying one single table only. However, if you add a column that is a link to another table, that creates a new join in the query.



Note:

This is also true for features, for the reasons already mentioned.

Before adding the following items to a list configuration, we also recommend thinking through your needs:

- A calculated field: Basic-type calculated fields also slow things down because they must be evaluated once for every record in the list.
- The **self** description string: For certain tables, the description is complex (for example, the Portfolio Items table) and may require multiple outer joins.

Sorts

When a list is sorted on the value of a field or link, the database engine sorts the results that are sent to Asset Manager. This operation is not expensive when the object is indexed in the database. The following list resumes the sort strategies that can be applied to lists starting with the least expensive in performance terms:

- Sort on an indexed field in the table
- 2 Sort on an indexed field on a linked table one link distant (e.g. **User.BarCode** from the list of Portfolio Items)
- 3 Sort on an non-indexed field in the table
- 4 Sort on multiple indexed fields in the table
- 5 Sort on multiple non-indexed fields in the table
- 6 Sort on multiple fields (indexed or not) in a linked table one link distant
- Sort on the description string of a linked table
- 8 Sort on fields (indexed or not) in a linked table more than one link distant

Filters

Filters translate directly to the WHERE clause of the query that populates the list on a table. As such, you should consider how they are going to be treated by the database when you create them. Filters that involve linked tables are necessarily going to be slower than filters that only involve fields on the main table. Complex filters can, if improperly designed, make the list open extremely slowly. Some rules of thumb for writing filters:

- Whenever possible, try to use AND instead of OR
- Try to avoid using LIKE and NOT LIKE. If you must use them, try to only use a trailing wildcard and use an indexed field. For example, it is better to write:

```
FullName LIKE '/Taltek/%'
than:
FullName LIKE '%/Taltek/%'
```

Wildcards at the start of the string for comparison force a complex search through the whole database, which can prove very costly.

It is better to use IN and NOT IN clauses with constants. For example:

```
Name IN ('Thomas',
                    'James')
```

Using these clauses with a sub-queries is very expensive.

2 Tuning the database

This chapter deals with tuning the database used by Asset Manager. It is not intended to replace specialized books on the subject but does describe several tuning strategies for Asset Manager. Whatever the case, it is highly recommended leaving this sort of operation to a qualified database administrator.

About hardware

The levels of database server hardware you need for your database is one of the hardest things to quantify. There are so many different factors involved, such as the number of users, the number of databases, the database engine, other applications running on the server, the quantity of data, and so on.

RAM

As a general rule of thumb, RAM is going to have the greatest impact on database performance. RAM is where the DBMS caches queries and data, so the more of it there is the more likely you are to get a cache hit when you execute a query. It's also where information for each user connection to the database is stored, so the more users you have the more RAM you need.

CPUs

The number and type of CPUs is also going to have an effect. If your DBMS is able to make use of multiple processors, then it is beneficial to have several, which allows the DMBS to do calculations in parallel. And of course, in general, the faster the CPU the better.

Hard drives

Hard disk I/O will also have a major effect on database performance. If you have a blazingly fast CPU with tons of RAM but your hard disk is slow, then any time the DBMS has to go to disk to retrieve data everything else has to wait until it's done. Ideally, you want the fastest hard disks you can get, typically SCSI drives, although IDE drives have come a long way in the past few years. You also may want to consider a RAID setup. RAID stands for Redundant Array of Independent Disks. It's essentially a way to combine multiple hard disks into one system for enhancing certain aspects of performance:

- If you want to set up a bare minimum RAID system, then you should consider RAID 1. RAID 1 is mirroring. That means that every disk in the array has another identical disk that duplicates it's data exactly. This allows a significant increase in Read I/O performance because both disks can be read in parallel (since they both contain the same information). As the majority of Asset Manager queries involve reading only, this will probably give you good performance. It also has the advantage of fault tolerance; if one disk crashes, causing data loss, you still have another disk with exactly the same data on it, so none of the data is actually lost. You just slip a new hard disk into the array and let it mirror the data again. Unfortunately there is no benefit to writing with this system. The other disadvantage is that you must have an even number of disks, so each one can be mirrored.
- If you have the money and the hardware, however, a RAID 10 system is better. This combines both mirroring and striping into one system. Striping is when the system writes data across several disks. This allows very fast write performance because each disk has a section of the data written to it simultaneously.

Network

As mentioned on several occasions, network performance is of extreme importance. The faster the network is, the greater the bandwidth and a greater number of concurrent users can be handled.

About queries

To better understand how query transactions impact performance, it is worthwhile detailing how a query is sent by Asset Manager:

1 Asset Manager builds the query and sends it to the database, but any parameters are not yet sent. The parameters are the hard-coded values of the query that are needed to specify what is returned. For example, in the following WHERE clause:

WHERE lEmplDeptId=12345

12345 is the parameter.

- 2 The database's query optimizer compiles an access plan for the data and returns it to Asset Manager. The value is irrelevant at this point. Using the above example, if the optimizer knows that it is going to get a value for **IEmplDeptId**, then it will be able to determine an appropriate access plan (such as using the index on **IEmplDeptId**).
- 3 Asset Manager binds the parameters to the query, then sends it back to the database.
- 4 The database executes the query based on the access plan and retrieves the data, then returns it to Asset Manager.



On some RDBMSs, steps 1 and 2 are not used. In that case, the query is just sent directly to the database with the parameters bound and without getting the access plan in advance. The impact on performance may be significant.

About database engines

This section contains a few suggestions that are valid for all database engines, although they will have differing levels of effectiveness depending on your DBMS.

- First off, when possible, keep indexes and data in separate tablespaces. Asset Manager uses a lot of indexes out of the box, and performance will typically be improved by having the indexes separated from the data. This is especially true when you can put the two tablespaces on separate disks, where the database can do simultaneous reads of both.
- Another recommendation is to calculate statistics for tables regularly. This can be done by either your database administrator as a regular database maintenance task, or by Asset Manager Automated Process Manager. HP

- recommends letting Asset Manager Automated Process Manager do this once a week. Table statistics are used to determine how often data is used. This lets the optimizer improve access to frequently used data.
- Rebuild your indexes regularly, as well, especially if you do a lot of inserts/deletes. After a while, the indexes become less efficient because of holes in their data and size. If you mostly do updates and there's not a lot of data being added or deleted, then you can probably do this once a quarter. Most likely, though, you should have the DBA perform this once a month a part of the regular database maintenance.
- If you're seeing poor performance overall, one thing you may want to look for is to see how much server swapping is going on. This happens when you have virtual memory specified and the operating system has to swap data out of RAM and onto the disk (in the pagefile) to make space for other data. or it needs to swap the data from the disk back into physical RAM in order to act on it. This can happen if there is insufficient physical RAM on the machine, or there is not enough physical RAM allocated to the RDBMS. If you see this happening (which you can see by looking at the Task Manager in Windows), then you either need to allocate more physical memory to the RDBMS (if possible) or add more physical memory to the server.

Eliminating locks and deadlocks

When several users are doing similar operations with quite large transactions involving counters, there is a significant probability that some concurrent transactions may end with deadlock on amCounter. This section deals with troubleshooting the deadlocks on several RDBMs.



Warning:

This information is intended for experienced database administrators.

Microsoft SQL Server

The proposed solution is to replace the amCounter table for most used counters and deport them to other tables. Each counter will have its own table to minimize locking.

Step 1 - Identify the deadlocking counters

Most likely, the counters with the highest numbers are the problem areas, but there might be others if some specific operations are leading to locks. Begin by identifying the counters and counter names. To do this:

- 1 Open MS SQL Server Query Analyzer
- 2 Run the following query against the Asset Manager database:

```
select * from amcounter
```

This query returns the current value of the counters stored in this table.



For the steps to follow, the following counters will be used:

amItemReceived ItemNo, amAsset AssetTag, amTicket TicketNo and amExpenseLine_ItemNo.

Step 2 - Backup

Backup the stored procedure **up GetCounterVal**.

Step 3 - Create tables for the counters

Run a CREATE TABLE script for each of the counters you are deporting using the following format:

```
CREATE TABLE [itam].[<NewTableName>] (
     [lValue] [int] IDENTITY (<IdentityStart>, 1) NOT NULL ,
     [luserspid] [int] NOT NULL
 ON [PRIMARY]
```

In this script:

- **NewTableName** is the name of the table to create, typically in the format cnt <Original counter name>.
- **IdentityStart** is the current value of the original counter table.

A typical SQL query that can be used to generate these CREATION STATEMENTS on the database would be:

```
select 'CREATE TABLE [itam].[cnt '+Identifier+'] ([lValue] [int] IDENTITY
('+ltrim(rtrim(str(lValue)))+',1) NOT NULL ,[luserspid] [int] NOT NULL)
N [PRIMARY]' as sqlquery from amCounter where lCounterid<>0
```



In this example the owner of the tables will be specified as itam, be sure it matches your database and modify the script if necessary.

Using the examples of counters defined above, the statements sent will be the following:

```
CREATE TABLE [itam].[cnt amAsset AssetTag] (
     [lValue] [int] IDENTITY (17697, 1) NOT NULL ,
     [luserspid] [int] NOT NULL
 ON [PRIMARY]
```

```
GO
CREATE TABLE [itam].[cnt amExpenseLine ItemNo] (
     [lValue] [int] IDENTITY (10735, 1) NOT NULL ,
     [luserspid] [int] NOT NULL
) ON [PRIMARY]
GΟ
CREATE TABLE [itam].[cnt amItemReceived ItemNo] (
     [lValue] [int] IDENTITY (4, 1) NOT NULL,
     [luserspid] [int] NOT NULL
) ON [PRIMARY]
GO
CREATE TABLE [itam].[cnt amticket ticketno] (
    [lValue] [int] IDENTITY (140000, 1) NOT NULL ,
    [luserspid] [int] NOT NULL
) ON [PRIMARY]
GO
```

Step 4 - Create a stored procedure

The next step is to create the stored procedure to handle counters. We suggest that you use the following example, replace the counters name in the stored procedure and add/remove as many procedure blocks as needed. Also, make sure the counter name appears in the NOT IN () statement used to distinguish counters using their own table from those still using **amCounter**.

```
sp dropmessage 50895
sp addmessage 50895, 16,'up GetCounterVal increment is %d which is not sup
ported when modified for non-locking optimization. Please edit the AMDB.IN
I file on your computer and add ImportCounterCache=1 in the current databa
se connexion description'
CREATE PROCEDURE up GetCounterVal @CounterName varchar(64), @CounterIncrem
ent int AS
BEGIN
 DECLARE @CounterIdent int
 'Change the NOT IN() list in this next line to reflect the names of the
new counters in your database.
 IF @CounterName NOT IN ('amItemReceived ItemNo', 'amAsset AssetTag', 'am
Ticket TicketNo', 'amExpenseLine ItemNo')
 BEGIN
   DECLARE amCounterLock CURSOR FOR SELECT lCounterId FROM amCounter WHER
E Identifier = @CounterName FOR UPDATE
   SELECT @CounterIdent = MAX(lCounterId) FROM amCounter WHERE Identifier
= @CounterName
   IF @CounterIdent IS NULL
     INSERT INTO amCounter (lCounterId, Identifier, Description, lValue,
terIncrement, GetDate() FROM amCounter
   ELSE
   BEGIN
```

```
OPEN amCounterLock
        UPDATE amCounter SET lValue = lValue + @CounterIncrement, dtLastMo
dif = GetDate() WHERE lCounterId = @CounterIdent
   SELECT lValue FROM amCounter WHERE Identifier = @CounterName
   IF @CounterIdent IS NOT NULL
     CLOSE amCounterLock
   /* END IF*/
   DEALLOCATE amCounterLock
  END
 ELSE
 BEGIN
 IF @CounterIncrement <> 1
   RAISERROR (50895, 17, @CounterIncrement)
  /* END IF */
-- See generation sql below
 --Add an IF block (next 5 lines) for EACH new counter table. Insert the
original counter name in all the bold faced areas.
 IF @CounterName = 'amTicket TicketNo'
   BEGIN
      insert into cnt_amTicket_TicketNo (luserspid) values(@@SPID)
      SELECT max(lValue) FROM cnt amTicket TicketNo WITH (NOLOCK) where lu
serspid=@@SPID
   END
 IF @CounterName = 'amAsset AssetTag'
   BEGIN
      insert into cnt amAsset AssetTag (luserspid) values (@@SPID)
      SELECT max(1Value) FROM cnt_amAsset_AssetTag WITH (NOLOCK) where lus
erspid=@@SPID
 IF @CounterName = 'amExpenseLine ItemNo'
   BEGIN
      insert into cnt amExpenseLine ItemNo (luserspid) values(@@SPID)
     SELECT max(lValue) FROM cnt amExpenseLine ItemNo WITH (NOLOCK) where
luserspid=@@SPID
 IF @CounterName = 'amItemReceived ItemNo'
      insert into cnt amItemReceived ItemNo (luserspid) values (@@SPID)
      SELECT max(lValue) FROM cnt amItemReceived ItemNo WITH (NOLOCK) wher
e luserspid=@@SPID
   END
-- end of generation sql below
 END
END
GΟ
```



The content of the NOT IN() statement and the IF block statements can be respectively generated with the following queries:

```
select ''''+Identifier+''', ' as sqlquery from amCounter where lCounterid<
>0
and

select ' IF @CounterName ='''+Identifier+ '''
BEGIN
insert into cnt_'+Identifier+' (luserspid) values(@@SPID)
SELECT max(lValue) FROM cnt_'+Identifier+' WITH (NOLOCK) where luserspid=
@@SPID
END' as sqlquery from amCounter where lCounterid<>0
```

Step 5 - Check the stored procedure

Run the following script:

```
exec up_getcounterval 'amAsset_AssetTag',1

SELECT COUNT(*) from cnt_amAsset_AssetTag
exec up_getcounterval 'amItemReceived_ItemNo',1

SELECT COUNT(*) from cnt_amItemReceived_ItemNo
exec up_getcounterval 'amTicket_TicketNo',1

SELECT COUNT(*) from cnt_amticket_ticketno
exec up_getcounterval 'amExpenseLine_ItemNo',1

SELECT COUNT(*) from cnt_amExpenseLine_ItemNo
```

Each and every new counter table should theoretically have one record.

Step 6 - Modify the stored procedure up_GetId

Locking on ID generation may also occur sometimes. To minimize locking, modify the stored procedure **up_GetId** as follows:

```
DROP procedure up_GetId

GO

CREATE procedure up_GetId as
insert into LastId (Value) values(@@SPID)
return 1+(SELECT max(IdSeed) FROM lastid WITH (NOLOCK) where value=@@SPID)
)*30

/*

create procedure up_GetId as if (select count(*) from LastId WITH (READU
NCOMMITTED) )>20 delete from LastId insert into LastId (Value) values(@@S
PID) return 1+@@IDENTITY*30

*/
GO
```

Step 7 - Create a clean-up stored procedure

The purpose of this step is to build a stored procedure that will wipe out each counter table periodically when no inserts are pending.

1 Get the object ID for each counter table. Run the following query:

```
select id, name from sysobjects where name like 'cnt%'
```

2 You should get a result set like this:

3 Generate the stored procedure for wiping out specific counter tables:

```
To generate all the main code block below for all deletes, run the sql
query
select 'delete from [cnt '+Identifier+'] WITH (TABLOCKX)' as sqlquery f
rom amCounter where lCounterid<>0
DROP PROCEDURE AM Cleanspeccounters
CREATE PROCEDURE AM Cleanspeccounters AS
BEGIN
     declare @locknb int
     set nocount on
     --while 1=1
     'Insert a code block (Next 1 line) for EACH new counter. Insert th
e appropriate ID or counter name in the bold faced areas.
     delete from cnt amAsset AssetTag with (TABLOCKX)
     delete from cnt amExpenseLine ItemNo with (TABLOCKX)
     delete from cnt amItemReceived ItemNo with (TABLOCKX)
     delete from cnt amticket ticketno with (TABLOCKX)
--Next line is to be commented when you define a new SQL Agent Task to
kick off the procedure each time on the server (it is in fact recommend
--waitfor delay '000:05:00'
        'here the stored procedure is fired each 5 minutes
        CONTINUE
      end
end
```

- 4 Define a new SQL Agent Task to kick off the procedure each hour or so on the server.
- 5 Verify that it works. After the programmed delay all specific counter tables should be empty.

Step 8 - Edit the amdb.ini file

- 1 Open the amdb.ini file of Asset Manager
- 2 Find the section called [Connexion] and add this line:

```
ImportCounterCache=1
```

3 Repeat this procedure on each client computer.



This step is absolutely mandatory.

Oracle

The proposed solution is to modify the **UP_GETCOUNTERVALUE** stored procedure.

This procedure will improve performance in wizards or external applications calling Asset Manager via APIs. Specifically, it will speed handling of large transactions involving the insertion of records/objects which need default values taken of from counters. On large transactions with several inserts needing default values from counters, counters may lock or deadlock (especially if several transactions are played at the same time, all needing a new value for the same counter). Current lock contention on counters can be monitored through a SQL query on the **v\$session** view (when connected as an administrator) with the following statement:

```
Select username, machine, sid, serial# from v$session where lockwait != Nu l1;
```

If this returns a large number of locks, then you may benefit from the workaround listed below. The purpose of the workaround is to replace counters whenever possible by sequences (which are non-locking mechanisms).



This solution eliminates your ability to query counter values through Asset Manager (the Tools/Administration/Counters menu item will display a screen with out-of-date values).

Step 1 - Preparation

- 1 Before you begin, make sure no one is using the database (or at least is doing inserts using counters). This is very important, otherwise the sequences you generate will not have the correct numbers.
- 2 Before altering the database, make sure you have a backup and tested copy of the database.

Step 2 - Replace the counters

- 1 Connect to SQLPlus as the Asset Manager database owner
- 2 Identify the counters to be replaced by sequences with the following query:

```
Select * from amcounter;
```

3 Move the counters into sequences:

```
SET CHARWIDTH 200;
SET PAGESIZE 9999;
SET HEADING OFF;
SET ECHO OFF;
Spool seqgen.sql

SELECT 'CREATE SEQUENCE ' || IDENTIFIER || ' INCREMENT BY 1 NOMAXVALUE
MINVALUE '|| (LVALUE + 1) || ' NOCYCLE CACHE 20 ORDER;' AS TEXT FROM AM
COUNTER where identifier IS NOT NULL;

Spool off
```

This yields a result similar to the following:

```
CREATE SEQUENCE amEmplDept_BarCode INCREMENT BY 1 NOMAXVALUE MINVALUE 6 0155 NOCYCLE CACHE 20 ORDER; CREATE SEQUENCE amLocation_BarCode INCREMENT BY 1 NOMAXVALUE MINVALUE 31 NOCYCLE CACHE 20 ORDER; CREATE SEQUENCE amCategory_BarCode INCREMENT BY 1 NOMAXVALUE MINVALUE 101 NOCYCLE CACHE 20 ORDER; CREATE SEQUENCE amProduct_CatalogRef INCREMENT BY 1 NOMAX VALUE MINVALUE 34076 NOCYCLE CACHE 20 ORDER; ...

CREATE SEQUENCE amOutputEvent_EventNo INCREMENT BY 1 NOMAXVALUE MINVALUE 11 NOCYCLE CACHE 20 ORDER;

49 rows selected.
```

4 The seggen.sql file can then be then edited to keep only the CREATE SEQUENCE statements.

Step 3 - Update the stored procedure

Replace the stored procedure **UP_GETCOUNTERVAL** to take advantage of sequences where possible.

1 Execute a basic statement to generate most of the sequence calls into the stored procedure. For example:

```
spool UP_GETCOUNTERVAL_gen.SQL
SET CHARWIDTH 200;
SELECT 'ELSIF CounterName = '||CHR(39)|| IDENTIFIER ||CHR(39)||' THEN '
```

```
|
|| ' SELECT '||IDENTIFIER||'.NEXTVAL INTO CounterValue FROM DUAL;'
FROM AMCOUNTER;
spool off;
```

This will generate a statement like the following:

```
ELSIF CounterName = '' THEN SELECT .NEXTVAL INTO CounterValue FROM DUA
L;

ELSIF CounterName = 'amEmplDept_BarCode' THEN SELECT amEmplDept_BarCod
e.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amLocation_BarCode' THEN SELECT amLocation_BarCod
e.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amCategory_BarCode' THEN SELECT amCategory_BarCod
e.NEXTVAL INTO CounterValue FROM DUAL;
...
```

- 2 Make the following modifications to the generated statement:
 - 1 Change the generated statement to properly handle IF Blocks.
 - 2 Remove the line with that deals with the empty counter:

```
ELSIF CounterName = '' THEN SELECT .NEXTVAL INTO CounterValue FROM DUAL;
```

3 Add the declare and end part of the UP_GETCOUNTERVAL stored procedure.

Your final statement should look like the following example:

```
CREATE OR REPLACE PROCEDURE UP GETCOUNTERVAL (CounterName IN VARCHAR
2, CounterIncrement IN NUMBER, CounterValue OUT NUMBER) AS
CounterIdent NUMBER;
BEGIN
IF CounterIncrement <> 1 THEN
RAISE APPLICATION ERROR (-20004, 'Error in counter increment with th
e Up getcounterval stored procedure altered to minimize locking, You
must use a counter increment of 1, please add the importcountercache
=1 in the amdb.ini file, in the parameters of your current database
connection');
Code.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amLocation BarCode' THEN                                SELECT amLocation Bar
Code.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amCategory BarCode' THEN                                SELECT amCategory Bar
Code.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amProduct CataloqRef' THEN                              SELECT amProduct Ca
talogRef.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amProduct BarCode' THEN                                  SELECT amProduct BarCo
de.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amProdCompo ItemNo' THEN                                SELECT amProdCompo It
emNo.NEXTVAL INTO CounterValue FROM DUAL;
.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amConnection CnxNo' THEN                                 SELECT amConnection C
nxNo.NEXTVAL INTO CounterValue FROM DUAL;
TVAL INTO CounterValue FROM DUAL;
```

ELSIF CounterName = 'amWorkOrder_WONo' THEN SELECT amWorkOrder_WONo .NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amDecisionTree_Code' THEN SELECT amDecisionTre e_Code.NEXTVAL INTO CounterValue FROM DUAL

ELSIF CounterName = 'amReqLine_ItemNo' THEN SELECT amReqLine_ItemNo
.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amDeliv_DelivNumber' THEN SELECT amDeliv_DelivNumber.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amItemReceived_ItemNo' THEN SELECT amItemReceived_ItemNo.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amItemReturned_ItemNo' THEN SELECT amItemRetur
ned_ItemNo.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amInvoice_InvoiceNumber' THEN SELECT amInvoice
_InvoiceNumber.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amFieldAdjust_ItemNo' THEN SELECT amFieldAdjust_ItemNo.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amFieldAdjustModel_ItemNo' THEN SELECT amField AdjustModel ItemNo.NEXTVAL INTO CounterValue FROM DUAL;

ELSIF CounterName = 'amLoan_Code' THEN SELECT amLoan_Code.NEXTVAL I NTO CounterValue FROM DUAL;

```
ELSIF CounterName = 'amExpenseLine ItemNo' THEN SELECT amExpenseLin
e ItemNo.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amWfScheme Ref' THEN SELECT amWfScheme Ref.NEX
TVAL INTO CounterValue FROM DUAL;
EXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amWfActivity    Ref' THEN                                SELECT amWfActivity Ref
.NEXTVAL INTO CounterValue FROM DUAL;
Ref.NEXTVAL INTO CounterValue FROM DUAL;
AL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amWfTransition_Ref' THEN                                 SELECT amWfTransition
Ref.NEXTVAL INTO CounterValue FROM DUAL;
.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amWfSyncPoint Ref' THEN                                 SELECT amWfSyncPoint R
ef.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amWfInstance Ref' THEN SELECT amWfInstance Ref
.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amInputEvent EventNo' THEN                               SELECT amInputEvent
EventNo.NEXTVAL INTO CounterValue FROM DUAL;
ELSIF CounterName = 'amOutputEvent EventNo' THEN                             SELECT amOutputEve
nt EventNo.NEXTVAL INTO CounterValue FROM DUAL;
ELSE
SELECT MAX(lCounterId) INTO CounterIdent FROM amCounter WHERE Identi
fier = CounterName;
IF CounterIdent IS NULL THEN
SELECT MAX(lCounterId)+1 INTO CounterIdent FROM amCounter;
INSERT INTO amCounter (lCounterId, Identifier, Description, lValue,
dtLastModif) VALUES (CounterIdent, CounterName, CounterName, Counter
Increment, SYSDATE);
ELSE
= SYSDATE WHERE lCounterId = CounterIdent;
SELECT lValue INTO CounterValue FROM amCounter WHERE Identifier = Co
unterName;
END IF;
END;
```

- 3 Make a backup of the original **UP_GETCOUNTERVAL** stored procedure.
- 4 Update the stored procedure using the script previously generated.

Step 4 - Edit the amdb.ini file

- 1 Open the amdb.ini file of Asset Manager
- 2 Find the section called [**Connexion**] and add this line:

```
ImportCounterCache=1
```

3 Repeat this procedure on each client computer.



This step is absolutely mandatory.

DB2

The proposed solution is to modify the **UP_GETCOUNTERVALUE** stored procedure.

This procedure will improve performance in wizards or external applications calling Asset Manager via APIs. Specifically, it will speed up handling of large transactions involving the insertion of records/objects that need default values from counters.

If this returns a large number of locks, then you may benefit from the workaround listed below. The purpose of the workaround is to replace counters whenever possible by sequences (which are non-locking mechanisms).



This solution eliminates your ability to query counter values through Asset Manager (the Tools/Administration/Counters menu item will display a screen with out-of-date values).

Step 1 - Preparation

- 1 Before you begin, make sure no one is using the database (or at least is doing inserts using counters). This is very important, otherwise the sequences you generate will not have the correct numbers.
- 2 Before altering the database, make sure you have a backup and tested copy of the database.

Step 2 - Replace the counters

- 1 Connect to CLP as the Asset Manager database owner
- 2 Identify the counters to be replaced by sequences with the following query:

```
Select * from amcounter;
```

3 Move the counters into sequences:

```
SELECT 'CREATE SEQUENCE ' | IDENTIFIER | | ' INCREMENT BY 1 NOMAXVALUE
START WITH '|| CHAR(LVALUE + 1) || ' NOCYCLE CACHE 20 ORDER' AS TEXT FR
OM AMCOUNTER where identifier IS NOT NULL;
```

This will generate a statement like the following:

CREATE SEQUENCE amInputEvent EventNo INCREMENT BY 1 NOMAXVALUE START WI TH 1001 NOCYCLE CACHE 20 ORDER

```
CREATE SEQUENCE amOutputEvent EventNo INCREMENT BY 1 NOMAXVALUE START W
ITH 1001 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfActivAlarm Ref INCREMENT BY 1 NOMAXVALUE START WITH
1001 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfActivity Ref INCREMENT BY 1 NOMAXVALUE START WITH 1
005 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfEvent Ref INCREMENT BY 1 NOMAXVALUE START WITH 1011
NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfInstance Ref INCREMENT BY 1 NOMAXVALUE START WITH 1
029 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfOrgRole Ref INCREMENT BY 1 NOMAXVALUE START WITH 10
01 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfScheme_Ref INCREMENT BY 1 NOMAXVALUE START WITH 101
O NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfTransition Ref INCREMENT BY 1 NOMAXVALUE START WITH
1009 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amWfWorkItem Ref INCREMENT BY 1 NOMAXVALUE START WITH 1
029 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amCable CableTag INCREMENT BY 1 NOMAXVALUE START WITH 1
001 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amCableLink BarCode INCREMENT BY 1 NOMAXVALUE START WIT
H 1001 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amColorCode Code INCREMENT BY 1 NOMAXVALUE START WITH 1
001 NOCYCLE CACHE 20 ORDER
CREATE SEQUENCE amTraceHistory BarCode INCREMENT BY 1 NOMAXVALUE START
WITH 1001 NOCYCLE CACHE 20 ORDER
. . . . . .
73 rows selected.
```

4 The seggen.sql file can then be then edited to keep only the CREATE SEQUENCE statements.

Step 3 - Update the stored procedure

Replace the stored procedure **UP_GETCOUNTERVAL** to take advantage of sequences where possible.

1 Execute a basic statement to generate most of the sequence calls into the stored procedure. For example:

```
THEN Values NEXTVAL FOR '||IDENTIFIER||' INTO CounterValue ' FROM AMC
OUNTER
```

This will generate a statement like the following:

```
ELSEIF CounterName = 'amInputEvent EventNo' THEN Values NEXTVAL FOR amI
nputEvent EventNo INTO CounterValue;
ELSEIF CounterName = 'amOutputEvent EventNo' THEN                            Values NEXTVAL FOR am
OutputEvent EventNo INTO CounterValue;
ELSEIF CounterName = 'amWfActivAlarm Ref' THEN Values NEXTVAL FOR amWfA
ctivAlarm Ref INTO CounterValue;
ELSEIF CounterName = 'amWfActivity Ref' THEN Values NEXTVAL FOR amWfAct
ivity Ref INTO CounterValue;
ELSEIF CounterName = 'amWfEvent Ref' THEN Values NEXTVAL FOR amWfEvent
Ref INTO CounterValue;
```

```
ELSEIF CounterName = 'amWfInstance Ref' THEN Values NEXTVAL FOR amWfIns
tance Ref INTO CounterValue;
ELSEIF CounterName = 'amWfOrgRole Ref' THEN Values NEXTVAL FOR amWfOrgR
ole Ref INTO CounterValue;
ELSEIF CounterName = 'amWfScheme Ref' THEN Values NEXTVAL FOR amWfSchem
e Ref INTO CounterValue;
ELSEIF CounterName = 'amWfTransition Ref' THEN Values NEXTVAL FOR amWfT
ransition Ref INTO CounterValue;
ELSEIF CounterName = 'amWfWorkItem_Ref' THEN Values NEXTVAL FOR amWfWor
kItem Ref INTO CounterValue;
ELSEIF CounterName = 'amCable CableTag' THEN Values NEXTVAL FOR amCable
CableTag INTO CounterValue;
ELSEIF CounterName = 'amCableLink_BarCode' THEN Values NEXTVAL FOR amCa
bleLink BarCode INTO CounterValue;
ELSEIF CounterName = 'amColorCode Code' THEN Values NEXTVAL FOR amColor
Code Code INTO CounterValue;
ELSEIF CounterName = 'amTraceHistory BarCode' THEN Values NEXTVAL FOR a
mTraceHistory BarCode INTO CounterValue;
Code INTO CounterValue;
ELSEIF CounterName = 'InternalRef' THEN Values NEXTVAL FOR InternalRef
INTO CounterValue;
ELSEIF CounterName = 'amContract Ref' THEN Values NEXTVAL FOR amContrac
t Ref INTO CounterValue;
```

2 Make the following modifications to the generated statement:

- 1 Change the generated statement to properly handle IF Blocks.
- 2 Remove the line with that deals with the empty counter:
- 3 Add the declare and end part of the UP GETCOUNTERVAL stored procedure.

```
Your final statement should look like the following example:
CREATE PROCEDURE UP GETCOUNTERVAL (IN CounterName VARCHAR(64), IN Co
unterIncrement INTEGER, OUT CounterValue INTEGER) LANGUAGE SQL BEGIN
DECLARE Counterident INTEGER; DECLARE ERR MSG VARCHAR(255);
IF CounterIncrement <> 1 THEN
SELECT 'Up getcounterval stored procedure altered to minimize lockin
g, You must use a counter increment of 1, please add the importcount
ercache=1 in the amdb.ini file, in the parameters of your current da
tabase connection' INTO ERR MSG FROM SYSIBM.SYSDUMMY1;
SIGNAL SQLSTATE '75002' SET MESSAGE TEXT = ERR MSG;
ELSEIF CounterName = 'amInputEvent EventNo' THEN Values NEXTVAL FOR
amInputEvent EventNo INTO CounterValue;
ELSEIF CounterName = 'amOutputEvent_EventNo' THEN Values NEXTVAL FOR
amOutputEvent EventNo INTO CounterValue;
ELSEIF CounterName = 'amWfActivAlarm_Ref' THEN Values NEXTVAL FOR am
WfActivAlarm Ref INTO CounterValue;
ELSEIF CounterName = 'amWfActivity Ref' THEN Values NEXTVAL FOR amWf
Activity Ref INTO CounterValue;
ELSEIF CounterName = 'amWfEvent Ref' THEN Values NEXTVAL FOR amWfEve
nt Ref INTO CounterValue;
ELSEIF CounterName = 'amWfInstance Ref' THEN Values NEXTVAL FOR amWf
Instance Ref INTO CounterValue;
```

```
ELSEIF CounterName = 'amWfOrqRole Ref' THEN Values NEXTVAL FOR amWfO
rgRole Ref INTO CounterValue;
ELSEIF CounterName = 'amWfScheme Ref' THEN Values NEXTVAL FOR amWfSc
heme_Ref INTO CounterValue;
ELSEIF CounterName = 'amWfTransition Ref' THEN Values NEXTVAL FOR am
WfTransition Ref INTO CounterValue;
ELSEIF CounterName = 'amWfWorkItem_Ref' THEN Values NEXTVAL FOR amWf
WorkItem Ref INTO CounterValue;
ELSEIF CounterName = 'amCable_CableTag' THEN Values NEXTVAL FOR amCa
ble CableTag INTO CounterValue;
ELSEIF CounterName = 'amCableLink BarCode' THEN Values NEXTVAL FOR a
mCableLink BarCode INTO CounterValue;
ELSEIF CounterName = 'amColorCode_Code' THEN Values NEXTVAL FOR amCo
lorCode Code INTO CounterValue;
ELSEIF CounterName = 'amTraceHistory BarCode' THEN Values NEXTVAL FO
R amTraceHistory BarCode INTO CounterValue;
ELSEIF CounterName = 'amCatalog Code' THEN Values NEXTVAL FOR amCata
log Code INTO CounterValue;
ELSEIF CounterName = 'InternalRef' THEN Values NEXTVAL FOR InternalR
ef INTO CounterValue;
ELSEIF CounterName = 'amContract Ref' THEN Values NEXTVAL FOR amCont
ract Ref INTO CounterValue;
dentifier = CounterName;
IF CounterIdent IS NULL THEN SELECT MAX(|CounterId|)+1 INTO CounterId
ent FROM amCounter:
INSERT INTO amCounter (lCounterId, Identifier, Description, lValue,
Increment, CURRENT TIMESTAMP);
ELSE UPDATE amCounter SET lValue = lValue + CounterIncrement, dtLast
Modif = CURRENT TIMESTAMP WHERE lCounterId = CounterIdent;
END IF;
SELECT lValue INTO CounterValue FROM amCounter WHERE Identifier = Co
unterName;
END IF;
END
GΟ
```

- 3 Make a backup of the original **UP_GETCOUNTERVAL** stored procedure.
- 4 Update the stored procedure using the script previously generated.

Step 4 - Edit the amdb.ini file

- 1 Open the amdb.ini file of Asset Manager
- 2 Find the section called [**Connexion**] and add this line:

ImportCounterCache=1

3 Repeat this procedure on each client computer.



This step is absolutely mandatory.