

# HP Asset Manager

## Performance Test Results - Enterprise Discovery<sup>®</sup> and Asset Manager<sup>®</sup> Integration Using Connect-It<sup>®</sup>



Legal Notices .....	2
Introduction .....	3
Scenario architecture .....	3
Physical architecture .....	3
Servers' characteristics .....	4
Logical architecture .....	4
Information on data polled in Enterprise Discovery .....	6
Modifications performed to increase performance .....	6
Details of modifications .....	7
M1 – Asset Manager counters .....	7
M2 – Adding indexes in Asset Manager .....	7
M3 – Replace trigger on amExtensionCard .....	8
M4 – Adding indexes on counters tables .....	8
M5 – modification of Card_CardIDCompld index .....	8
M6 – adding an index on amWflInstance table .....	9
M7 – Query compilation optimization .....	9
Overview of test results .....	9
Details of tests results .....	10
Test results table .....	11
Thread count evolution performance study .....	12
Resource consumption on servers .....	13
ACSRV .....	13
CITSRV1 .....	14
CITSRV2 .....	15
EDSRV .....	16
For more information .....	18

## Legal Notices

© Copyright 1994-2008 Hewlett-Packard Development Company, L.P.

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.

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.

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 US trademark of Sun Microsystems, Inc.

Microsoft®, Windows®, Windows NT®, 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.

## Introduction

This document shows the performance gain that can be obtained by making specific modifications to the out-of-the-box Enterprise Discovery<sup>®</sup>-to-Asset Manager<sup>®</sup> Connect-It<sup>®</sup> integration scenario. It describes the various modifications that you make and the improved results that you may experience.

As this document shows, applying these modifications can decrease the scenario execution time to about 1/12<sup>th</sup> the execution time experienced using the out-of-the-box scenario.

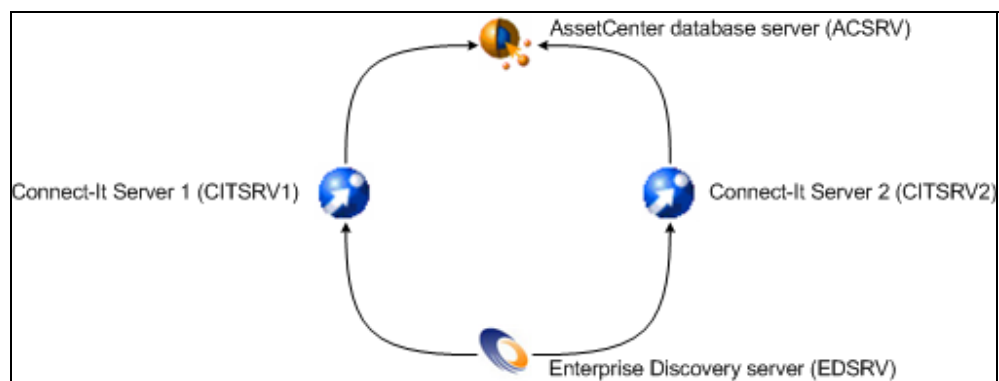
**Important:** This document is intended for customers who have large amounts of data and high performance expectations with Connect-It. Customers who have small amounts of data to integrate from Enterprise Discovery to Asset Manager will likely find that the benefits of modifying the out-of-the-box scenario, `edac.scn`, are not worth the required time and effort. In addition, customers who want to implement these modifications should already be familiar with both Connect-It and Asset Manager. It is also preferable that they have skills in SQL.

**Note:** The results of implementing these modifications in a production environment may vary based on a variety of factors, including hardware and software configurations; network hardware, software, and traffic volume; other applications running on client and server machines, including the database; and expertise in tuning Enterprise Discovery, Connect-It, and Asset Manager. All these factors are beyond the scope of this document.

**Note 2:** The testing was performed on a platform with AssetCenter 4.4.2. The results are still available for Asset Manager 5.0 and 5.1 as they were no significant change in the area concerned by the testing. In the results you still find reference to AssetCenter, former name for Asset Manager.

## Scenario architecture

### *Physical architecture*



Two servers hosting Connect-It will be used. Both extract data from Enterprise Discovery and insert it into Asset Manager. The scenarios on both servers are the same, except for the filter in the produced documents. (See the section Information on data polled in Enterprise Discovery).

## ***Servers' characteristics***

The following table shows specific information about the servers that were used in the benchmark test. All servers include Intel® Xeon® or Pentium® processors and run the Microsoft® Windows Server® 2003 operating system.

Server	RAM	CPU	OS	Software
ACSRV	4 Gb	4x Intel Xeon MP 3,16 Ghz	Windows Server 2003	AC 4.4
CITSRV1	2 Gb	4x Intel Xeon 3,6 Ghz	Windows Server 2003	CIT 3.6
CITSRV2	1,5 Gb	2x Intel Pentium III 662Mhz	Windows Server 2003	CIT 3.6
EDSRV	4 Gb	4x Intel Pentium III Xeon 550 Mhz	Windows Server 2003	ED 2.0.1

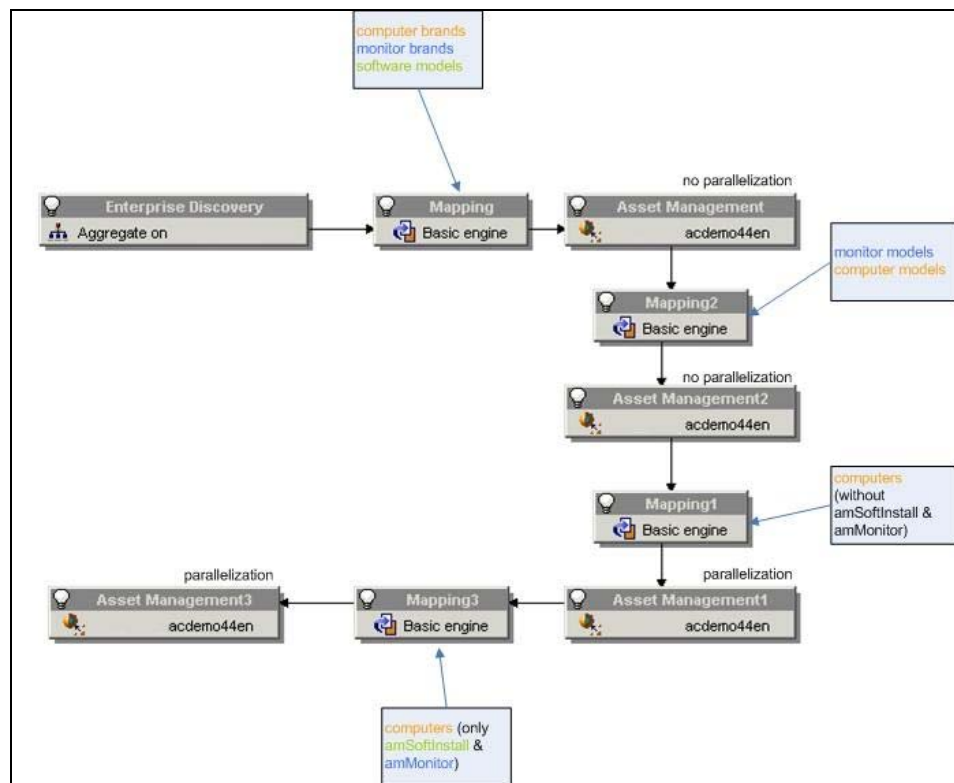
## ***Logical architecture***

The scenario extracts computer data, along with hardware and software information, from Enterprise Discovery and inserts it into Asset Manager. The parallelization capability of Connect-It, which allows multiple threads to perform operations, will be used to improve scenario execution time. Nevertheless, using Connect-It parallelization adds the risk of deadlocks, and the out-of-box scenario, which is intended to be run without parallelization, must be modified to be thread-safe.

Elements creation has been sequenced to allow parallelization to work correctly and to avoid deadlocks. Indeed, with the original scenario, enabling parallelization without modifications to the scenario will allow the creation of the same object in different threads. For instance, with insertions of computers having the same model, there can be a deadlock while creating the model or the associated brand.

The idea is to first create elements that will be used by other elements. For instance, in a first mapping, the computer brands will be created, then computer models (using the brands) will be created, and finally the computers themselves are created. Connectors performing insertions of shared elements (such as brands or models) must not be parallelized; only the connector performing insertions of unique elements (such as computers) can be parallelized.

Elements with no link can be treated in the same mapping box. For instance, computer brands and monitor brands can be created in the same mapping box. The diagram below shows the rebuilt scenario matching those requirements, and indicates where the different elements are created:



Here is a description of the different elements:

- **Enterprise Discovery:** The Enterprise Discovery connector produces documents for computers present in the Enterprise Discovery database. It produces documents only for computers, and not for other elements that are in the original out-of-the-box scenario `edac.scn`. (See the *Connect-It User's Guide* for more information.)
- **Mapping:** This mapping element is used for inserting computer brands, monitor brands, and software models. Insertion of computer brands and computer models, monitor brands, and monitor models has been done in two different mapping boxes, to prevent deadlocks from occurring in the Asset Manager database while creating the models and brands.
- **Asset Management:** The Asset Manager connector is used to insert computer brands, monitor brands, and computer models. This connector will never be multithreaded, to prevent deadlocks.
- **Mapping2:** This mapping box is used for inserting monitor models and computer models (using the brands created in the previous operations).
- **Asset Management2:** The Asset Manager connector is used to insert monitor and computer models. This connector will never be multithreaded, to prevent deadlocks.
- **Mapping1:** This mapping box is used to insert computers in Asset Manager. All are created except the `amSoftInstall` and `amMonitor` components. The `amSoftInstall` and `amMonitor` components are not mapped because there is a problem in multithreading with those components; they are inserted twice into Asset Manager (resulting in false information inserted into Asset Manager and in poor result performance). For instance, with one software installation, Connect-It tries to insert the same software installation into Asset Manager two times. The `amSoftInstall` and `amMonitor`

components must be inserted in a separated mapping (see Mapping3 below); this way, the right number of amSoftInstall and amMonitor elements will be inserted.

- **Asset Management1:** This Asset Manager connector is used to insert computers (without amSoftInstall and amMonitor). This connector can be multithreaded (using the Connect-It parallelization capability, an Asset Manager connector advanced property), without risk of deadlocks.
- **Mapping3:** This mapping box is used to complete the previous computers' insertions with amSoftInstall and amMonitor elements.
- **Asset Management3:** This Asset Manager connector is used to insert amSoftInstall and amMonitor elements. This connector can be multithreaded (using Connect-It parallelization capability, an Asset Manager connector advanced property), without risk of deadlocks.

**Note:** The optimized scenario named `edac-optim.scn` is attached to the white paper, so that the reader can refer to it and consult the modifications performed.

## Information on data polled in Enterprise Discovery

This table summarizes the data contained in the Enterprise Discovery database:

Element	Total	Average per computer
software	1 078 643	18.82
monitors	17 035	0.30
extensions cards	477 945	8.34
logical devices	522 810	9.12
physical devices	97 807	1.71
network cards	108 460	1.89
computers	57 327	

Considering the previous values, the number of portfolio items is 1,153,005 (including software, monitors, and computers).

So that the data polled on CITSRV1 and CITSRV2 is not the same, we perform an arbitrary filter that allows splitting Enterprise Discovery devices in two groups. The filter is applied on the field priority in the document that Enterprise Discovery produces. CITSRV1 polls computers whose priority is 1 (28,828 computers), and CITSRV2 polls computers whose priority is *not* 1 (28,499 computers).

## Modifications performed to increase performance

The various modifications that were made to improve performance are named M1 through M7.

Tuning Asset Manager counters, adding indexes in the Asset Manager database, and using the Connect-It parallelization feature reduced scenario execution time by 11/12 (over 91%). The out-of-the-box `edac.scn` scenario takes 183 hours and 1 minute to complete the benchmark test, whereas the optimized scenario needs only 15 hours 3 minutes. (To review the all the optimizations, open the `edac-optim.scn` file, which is attached to this document) In the optimized scenario, 63 computers are inserted per minute into Asset Manager, which corresponds to 1,276 portfolio items inserted per minute.

This section describes the modifications in detail. The section *Details of modifications* summarizes the results obtained during the different tests performed.

During the scenario execution, ACSRV resources are not completely used: More than 70% of the CPU is available and only half of the memory is used. Nevertheless, all CPU resources on CITSRV1 and CITSRV2 are completely used. This means that the Connect-It servers are at their maximum, but not the Asset Manager server. So, it's still possible to have better performance results by adding a

third Connect-It server. The section *Resource consumption on servers* details the server profiles during testing.

## ***Details of modifications***

### **M1 – Asset Manager counters**

The first modification applied is the SQL Server procedure for managing counters in Asset Manager. Refer to the *Asset Manager Tuning* document, in the section *Eliminating locks and deadlocks*, to perform the necessary modifications. Here is a list of the counters that were modified in the benchmark:

amDashboarditem_Name	amInputEvent_EventNo	amOutputEvent_EventNo	amWfActivAlarm_Ref
amWfActivity_Ref	amWfEvent_Ref	amWfInstance_Ref	amWfOrgRole_Ref
amWfScheme_Ref	amWfTransition_Ref	amWfWorkItem_Ref	amCable_CableTag
amCableLink_BarCode	amColorCode_Code	amTraceHistory_BarCode	amCatalog_Code
InternalRef	amContract_Ref	amCntrRent_Code	amLoan_Code
amReturnEnv_Code	amDaDeplInstance_Name	amDaWfScheme_Ref	amBudget_code
amBudgetCategory_code	amBudget_BudgetNo	amCostCategory_Code	amCostCenter_Code
amExpenseLine_ItemNo	amFYDivision_Code	amPeriod_Code	amCbkStoredEvent_Event
amDecisionTree_BarCode	amDecisionTree_Code	amDowntimePeriod_Code	amEscalScheme_BarCode
amEscSchLevel_Code	amKnowlBase_Code	amProblemClass_BarCode	amSeverity_Code
amTicket_TicketNo	amComputer_Name	amComputer_Group	amConnection_CnxNo
amAsset_AssetTag	amBrand_BarCode	amCompany_Code	amModel_BarCode
amModel_ModelRef	amNature_Code	amNews_Name	amAssignment_Code
amMonitor_Serial	amEmplDept_BarCode	amEmplGroup_BarCode	amLocation_BarCode
amAdjustment_ItemNo	amFieldAdjust_ItemNo	amFieldAdjustModel_ItemNo	amProject_Code
amReqLine_ItemNo	amRequest_ReqNumber	amStock_Code	amTraining_BarCode
amWorkOrder_WONo	amInvoice_InvoiceNumber	amPOOrder_PONumber	amPOrdLine_ItemNo
amDeliv_Number	amResa_ItemNo	amItemReceived_ItemNo	amWfSyncPoint_Ref
amCategory_BarCode	amDeliv_DelivNumber	amExtensionCard_CardID	

### **M2 – Adding indexes in Asset Manager**

As the scenario reconciles the PhysicalAddress, TcplpAddress, and AssetTag fields, two of those fields must be indexed in the Asset Manager database to obtain better performance results. The AssetTag field is already indexed, so we need to index only PhysicalAddress and TcplpAddress. Create indexes on the Computers table with Asset Manager Application Designer (former name in AssetCenter 4.4.x and 5.0x: AssetCenter Database Administrator - AMDBA) having the following properties:

- Index on PhysicalAddress:
  - SQL Name: Cpu\_PhysicalAddress
  - Fields: PhysicalAddress
  - Type: Normal
- Index on TcplpAddress:
  - SQL Name: Cpu\_TcplpAddress

- o Fields: TcplpAddress
- o Type: Normal

### M3 – Replace trigger on amExtensionCard

The existing trigger amExtensionCard\_ExtNoDupNdx present in the Asset Manager database can lead to deadlocks. To avoid them, replace the existing trigger with the following script (which was written for SQL Server). You will need to perform a SQL query directly on the database; be careful while performing this operation:

```

DROP TRIGGER itam.amExtensionCard_ExtNoDupNdx
GO
CREATE TRIGGER itam.amExtensionCard_ExtNoDupNdx ON itam.amExtensionCard FOR INSERT,UPDATE AS
  IF UPDATE (CardID) OR UPDATE (lCompId)
BEGIN
  IF (
    NOT((SELECT MAX(CardID) FROM inserted) IS NULL AND (SELECT MAX(lCompId) FROM inserted) IS
    NULL))
  BEGIN
  IF(
    NOT EXISTS (SELECT * FROM deleted) OR EXISTS (SELECT i.lCardId FROM inserted i, deleted d
    WHERE i.lCardId = d.lCardId AND((ISNULL(d.CardID,'') <> ISNULL(i.CardID,''))OR(d.lCompId <>
    i.lCompId))))
  BEGIN
  IF(
    (SELECT COUNT(*) FROM inserted I, itam.amExtensionCard E (Index=Card_CardIDlCompId nolock)
    WHERE (I.lCompId = E.lCompId AND ISNULL(I.CardID,'')= ISNULL(E.CardID,''))) > (SELECT
    COUNT(*) FROM inserted) )
  BEGIN
  RAISERROR ('NoDupIndexTriggered : 'amExtensionCard.Card_CardIDlCompId'', 14, 3)
  IF 0 < (SELECT COUNT(*) FROM deleted)
  UPDATE amExtensionCard SET lCardId = D.lCardId, dtLastModif = D.dtLastModif, CardID =
  D.CardID, Name = D.Name, Revision = D.Revision, Type = D.Type, VendorID = D.VendorID,
  lCompId = D.lCompId, lVendorId = D.lVendorId FROM deleted D WHERE amExtensionCard.lCardId =
  D.lCardId AND amExtensionCard.lCardId IN (SELECT D.lCardId FROM deleted D)
  ELSE
  DELETE FROM amExtensionCard WHERE lCardId = (SELECT lCardId FROM inserted)
  END
  END
  END
  END
  END
  
```

### M4 – Adding indexes on counters tables

The counters tables we created above help prevent deadlocks in the Asset Manager database, but when running a scenario that inserts a lot of records, the stored procedure that purges records does not have the time to delete the items in the tables. Access time to the records increases over time. Indexes on the most used tables can help there. As counter tables are not exposed in AMDBA, a SQL query must be directly performed on the database. Be careful while performing this operation.

The following script (which was written for SQL Server) creates indexes on four tables:

```

create unique index idx_cnt_amWfInstance_Ref_luserspid_lValue
on itam.cnt_amWfInstance_Ref (luserspid, lValue)
create unique index idx_cnt_amExtensionCard_CardID_luserspid_lValue
on itam.cnt_amExtensionCard_CardID (luserspid, lValue)
create unique index idx_cnt_amComputer_Name_luserspid_lValue
on itam.cnt_amComputer_Name (luserspid, lValue)
create unique index idx cnt amAssiignment Code luserspid lValue
  
```

### M5 – modification of Card\_CardIDCompId index

The index (on amExtensionCard table in the Asset Manager database) is not optimal; its columns have to be inverted to be more effective. Because it is a system index, you directly perform a SQL query to invert the columns; be careful while performing this modification. Note that the following script was written for SQL Server:



```
CREATE INDEX [Card_CardID1CompId]
ON [itam].[amExtensionCard] ([lCompId], [CardID])
WITH
    DROP_EXISTING
ON [PRIMARY]
```

## M6 – adding an index on amWfInstance table

Even if this table is not directly called by the integration scenario, many “where” clauses are executed by Asset Manager during the scenario execution. The queries are time-consuming for the database because there is no index on this table. Create an index on Workflow instances table with AMDBA having the following properties:

- SQL Name: WfInst\_IDocRecordIdIWfSchId
- Fields: IDocRecordId, IWfSchId
- Type: Normal

## M7 – Query compilation optimization

Looking at MS SQL Server database activity with the Profiler tool, we see that queries were often recompiled. Also, some queries directly contained values (instead of using parameters), and so were systematically compiled. Parameters have been added to the `amdb.ini` file on CITSRV1 and CITSRV2 to avoid this behavior. The parameters are placed under the section corresponding to the connection used by the connector to reach the Asset Manager database:

```
[ACDemo44en]
Engine=MSSQL
...
StmtCache=500
ConstAsText=0
ConstAsTextForSelect=0
ConstAsTextForGet=0
```

For system ODBC connections, the file can be found under `C:\Windows\`; for user ODBC connections, the file can be found under `C:\Documents and Settings\`.

## Overview of test results

Tuning Asset Manager counters, adding indexes in the Asset Manager database, and using the Connect-It parallelization feature reduced scenario execution time by 11/12 (over 91%). The out-of-the-box `edac.scn` scenario needs 183 hours and 1 minute to complete, whereas the optimized scenario needs only 15 hours 3 minutes. In the optimized scenario, 63 computers are inserted per minute into Asset Manager, which corresponds to 1,276 portfolio items inserted per minute.

Refer to the section *Modifications performed to increase performance* to see the modifications in detail. The section *Details of tests results* summarizes the results obtained during the different tests made.

During the scenario execution, ACSRV resources are not completely used: More than 70% of the CPU is available and only half of the memory is used. Nevertheless, all CPU resources on CITSRV1 and CITSRV2 are completely used. This means that the Connect-It servers are at their maximum, but not the Asset Manager server. So, it is still possible to have better performance results by adding a third Connect-It server. The section *Resource consumption on servers* details the server profiles during testing.

## Details of tests results

The first tests (T1, T2, T3 and T4) were done without splitting the treatment of amComputer insertions from amSoftInstall / amMonitor insertions. This means that there were no “Asset Management3” and “Mapping3” elements in the scenario for those tests; all was done in “Mapping1” mapping box.

In the results table, the “Configuration” column displays the number of threads for:

- The “Asset Management1” component in T1, T2, T3 and T4 tests
- The “Asset Management1”/“Asset Management3” components in subsequent tests

## Test results table

N°	Scenario / Modifications applied	Configuration		Time measurement		Avg Time	sec / computer	computer / sec	computer / min	nb rejected	% rejected	sec / portfolio item	portfolio item / sec	portfolio item / min
		CITSRV1	CITSRV2	CITSRV1	CITSRV2									
T1	edac		1 thread		183:01:00	183:01:00	11.49	0.09	5.4	0	0.00	0.57	1.75	105.00
T2	edacOptim - M1	1 thread	1 thread	102:55:00	44:43:00	73:49:00	4.64	0.22	13.2	0	0.00	0.23	4.34	260.40
T3	edacOptim - M1-M2	1 thread	1 thread	112:13:00	53:22:00	82:47:30	5.20	0.19	11.4	0	0.00	0.26	3.87	232.20
T4	edacOptim - M1-M2-M3	12 threads	8 threads	21:51:00	16:48:00	19:19:30	1.21	0.82	49.2	0	0.00	0.06	16.57	994.20
T5*	edacOptim - M1-M2-M3	20/10 threads	12/6 threads	19:52:00	16:01:00	17:56:30	1.13	0.89	53.4	0	0.00	0.06	17.85	1 071.00
T6	edacOptim - M1-M2-M3-M4	20/10 threads	12/6 threads	18:31:00	17:04:00	17:47:30	1.12	0.90	54	0	0.00	0.06	18.00	1 080.00
T7	edacOptim - M1-M2-M3-M4-M5	20/10 threads	12/6 threads	19:03:00	16:10:00	17:36:30	1.11	0.90	54	0	0.00	0.05	18.19	1 091.40
T8	edacOptim - M1-M2-M3-M4-M5	6/3 threads	6/3 threads	20:26:00	15:19:00	17:52:30	1.12	0.89	53.4	0	0.00	0.06	17.92	1 075.20
T9	edacOptim - M1-M2-M3-M4-M5-M6	6/3 threads	6/3 threads	15:03:00	15:28:00	15:15:30	0.96	1.04	62.4	0	0.00	0.05	20.99	1 259.40
T10	edacOptim - M1-M2-M3-M4-M5-M6-M7	8/3 threads	8/3 threads	12:39:00	17:28:00	15:03:30	0.95	1.06	63.6	0	0.00	0.05	21.27	1 276.20

The following formulas were used to calculate the values shown in the test results table:

- “Avg time” : If the treatment was done by CITSRV1 and CITSRV2, then “Avg time” equals (“Time passed on CITSRV1” + “Time passed on CITSRV2”) / 2.  
Otherwise it’s the time passed on the server used; the type is “hh:mm:ss”
- “sec/computer” = “Avg Time (seconds)” / “Number of computers scanned in Enterprise Discovery”
- “computer/sec” = “Number of computers scanned in Enterprise Discovery” / “Avg Time (seconds)”
- “computer/min” = “computer/sec” \* 60
- “sec/portfolio item” = “Avg Time (seconds)” / “Number of portfolio items to be inserted in Asset Manager”
- “portfolio item/sec” = “Number of portfolio items to be inserted in Asset Manager” / “Avg Time (seconds)”
- “portfolio item/min” = “portfolio item/sec” \* 60
- “%rejected” = “Number of computers documents on error” / “Number of computers scanned in Enterprise Discovery” \* 100

\* T5 is the first scenario having amSoftInstall and amMonitor insertions separated from amComputer insertions.

Note on threads number: As you can see, between T7 and T8 the scenario duration is quite similar, despite the important threads reduction. The threads count does not need to be too high to have an interesting performance gain; moreover, having too many threads causes high CPU consumption and latency in SQL Server treatments (because threads are waiting for other threads). Refer to the next paragraph to see the performance evolution in terms of thread count evolution.


## Thread count evolution performance study

To determine the best thread count to configure in the “Asset Management1” and “Asset Management3” connectors, we performed tests, varying the number of threads and checking the time required for the scenario to execute. During the tests, the number of documents to produce was limited to 2,000. The optimized scenario was used to perform the study. The results are in the above table and chart.

Configuration			Threads Total	Time needed (h:mm)		Avg (h:mm)	sec / compute	computer / sec	computer / min	
CITSRV1 threads	CITSRV2 threads	CITSRV1		CITSRV2						
0	0	2	2	2/2*	0:00	1:13	1:13	2.19	0.46	27.40
1	1	1	1	1/1	1:28	1:03	1:15	2.25	0.44	26.67
2	1	2	1	2/1	0:51	0:43	0:47	1.41	0.71	42.55
2	2	2	2	2/2	0:56	0:46	0:51	1.53	0.65	39.22
3	2	3	2	3/2	0:44	0:43	0:43	1.29	0.78	46.51
4	2	4	2	4/2	0:40	0:38	0:39	1.17	0.85	51.28
4	3	4	3	4/3	0:38	0:39	0:38	1.14	0.88	52.63
5	3	5	3	5/3	0:36	0:38	0:37	1.11	0.90	54.05
6	3	6	3	6/3	0:36	0:38	0:37	1.11	0.90	54.05
7	3	7	3	7/3	0:35	0:38	0:36	1.08	0.93	55.56
8	3	8	3	8/3	0:33	0:38	0:35	1.05	0.95	57.14
8	4	8	4	8/4	0:33	0:38	0:35	1.05	0.95	57.14
9	4	9	4	9/4	0:33	0:38	0:35	1.05	0.95	57.14
10	4	10	4	10/4	0:33	0:38	0:35	1.05	0.95	57.14
10	5	10	5	10/5	0:32	0:38	0:35	1.05	0.95	57.14
11	5	11	5	11/5	0:33	0:37	0:35	1.05	0.95	57.14
12	5	12	5	12/5	0:33	0:38	0:35	1.05	0.95	57.14
12	6	12	6	12/6	0:32	0:39	0:35	1.05	0.95	57.14

\* in this test, only CITSRV2 is used, the 2,000 computers are treated on this single server

In the configuration column table, the first value on a server is the number of threads in the “Asset Management1” connector; the second value is the number of threads in the “Asset Management3” connector.



The results show that beyond 8 threads allocated to the “Asset Management1” connector and 3 threads allocated to “Asset Management3”, there is no more performance gain. Therefore, this is the configuration that was applied to the final test (T10 in the results table in the section *Test results table*) for the connectors.

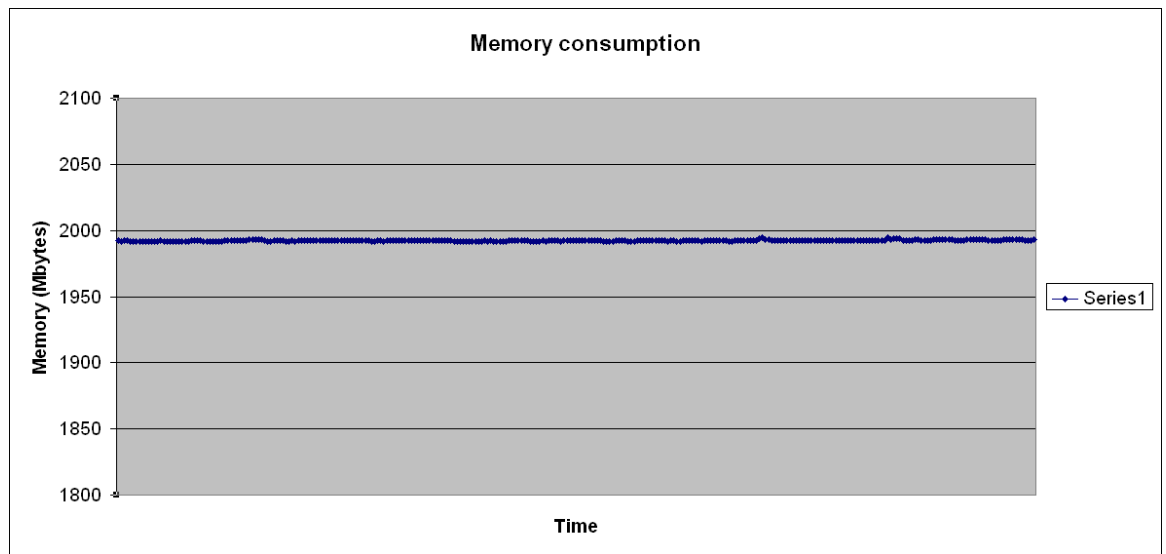
Don't be surprised that the computer/minute ratio is less important than in the results in the section *Test results table*. Indeed, there is a fixed time used to query the Enterprise Discovery database, and its impact on a test with 2,000 documents is more important than its impact on a test with 50,000 documents. So the computer/minute ratio will be more important as you treat more computers. What is important to remember here is that after a certain amount of threads is configured, there is no more performance gain. We suggest you perform the same tests in your environment to determine the best configuration, as it may vary according to your servers' capabilities and other factors.

### ***Resource consumption on servers***

Measures were performed on a one-hour time-frame. Memory and CPU evolutions stay the same after this period. The results concern the most optimized scenario (that used in T10).

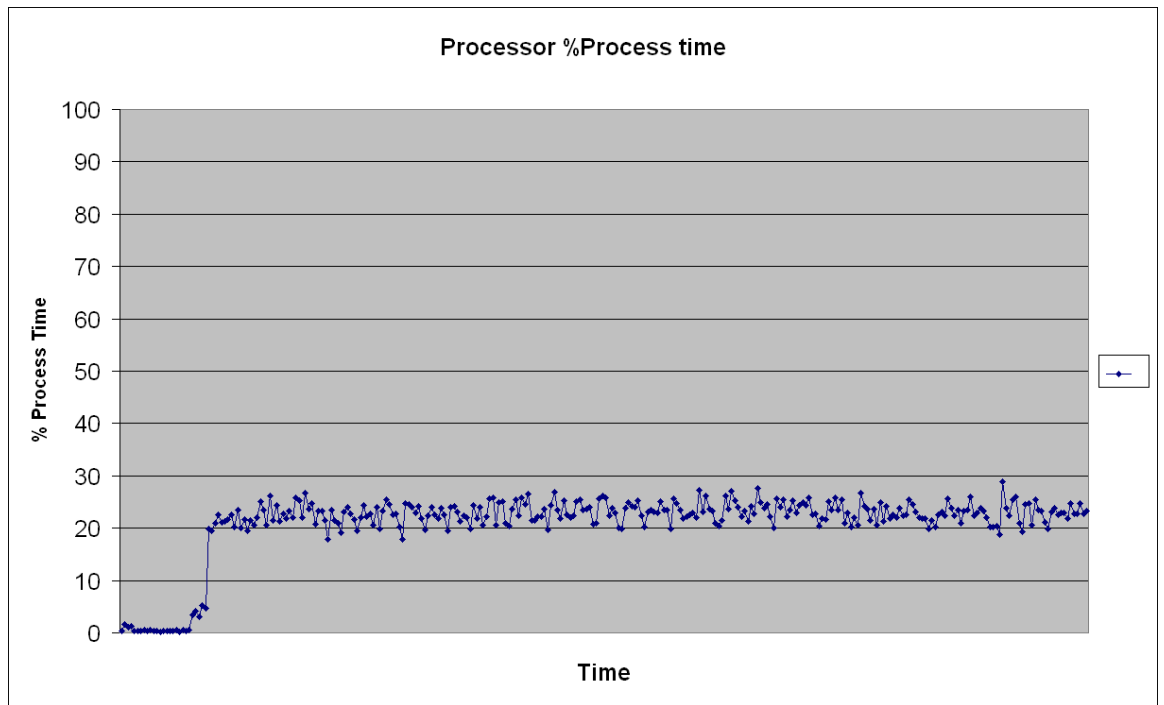
#### **ACSRV**

##### ***Memory consumption***




**Note:** ACSRV has 4 GB memory.

### ***CPU consumption***



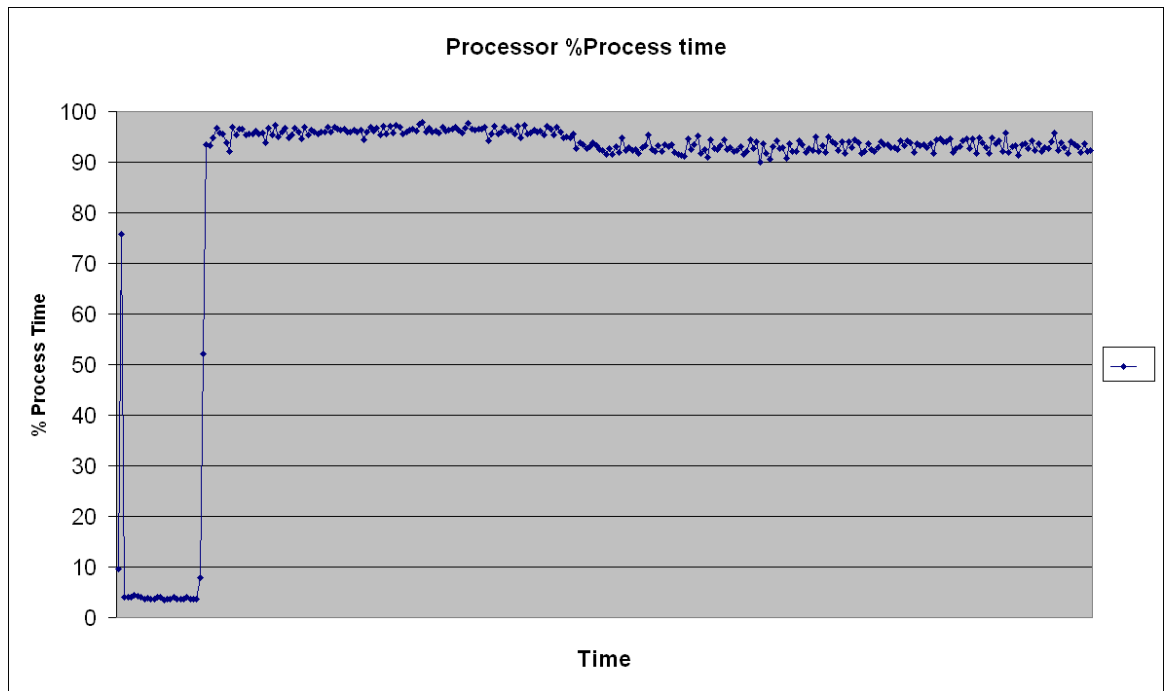
### **CITSRV1**

#### ***Memory consumption***



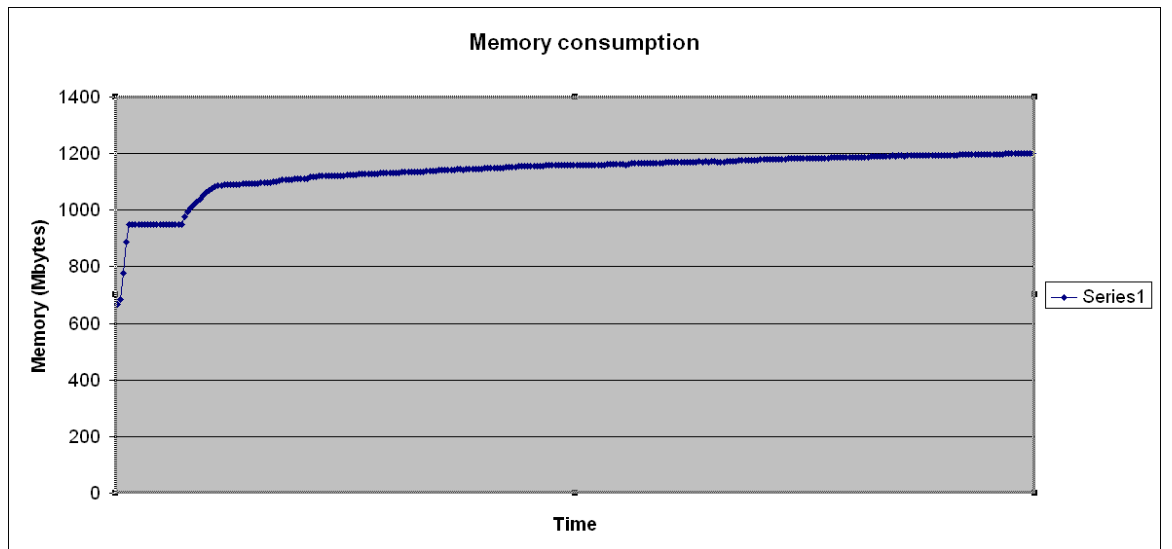
**Note:** CITSRV1 has 2 GB memory.

**CPU consumption**



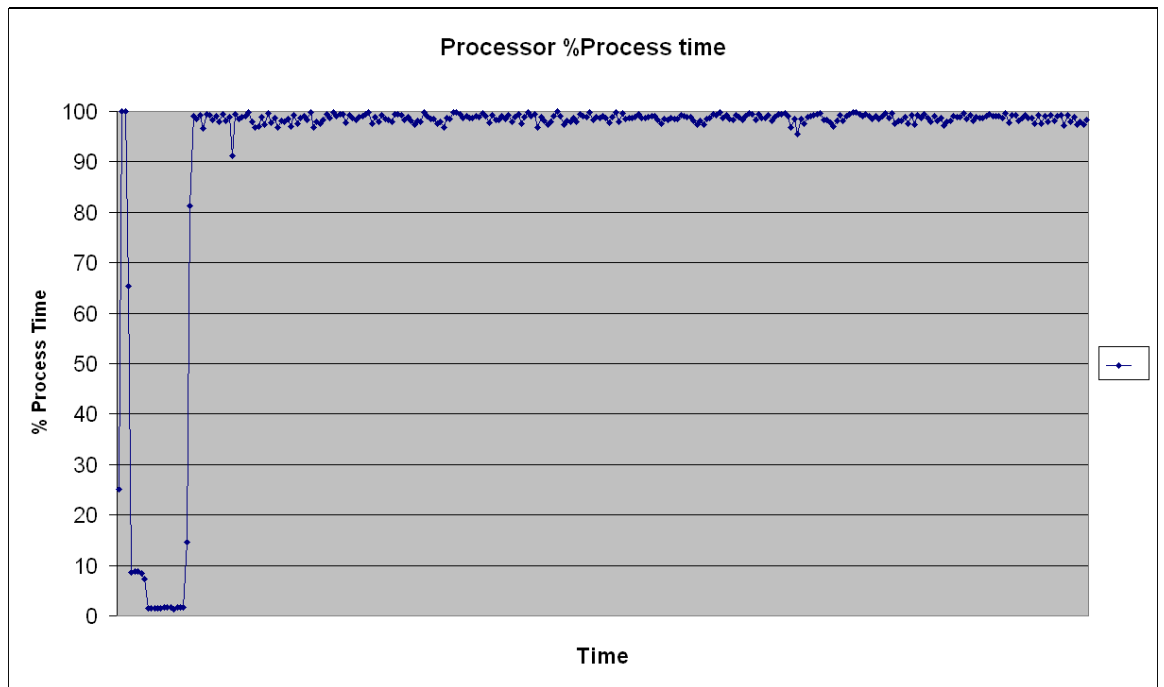
**CITSRV2**

**Memory consumption**



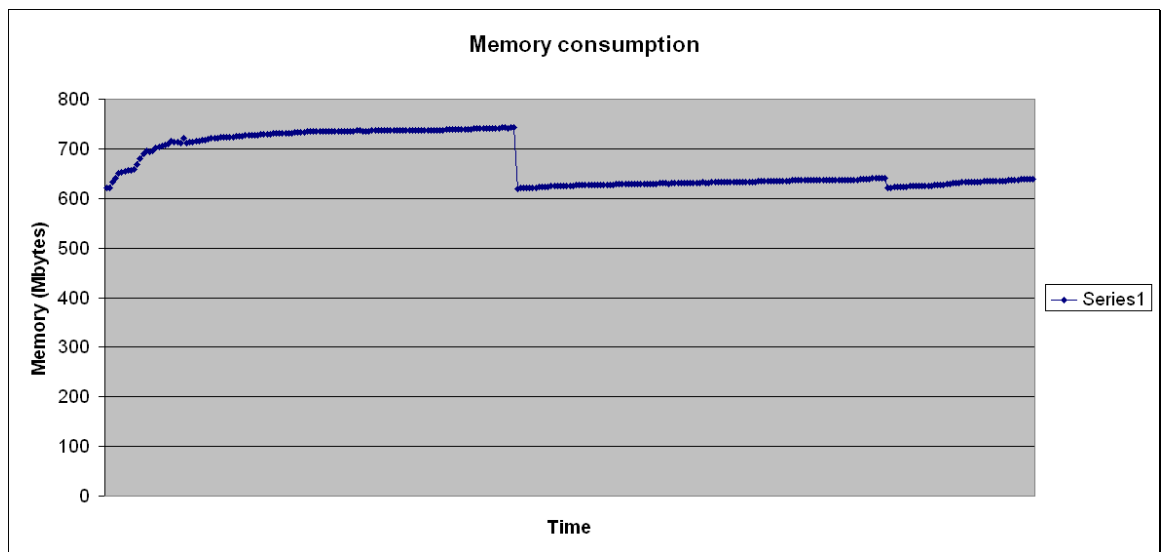
**Note:** CITSRV2 has 1.5 GB memory.

### CPU consumption



### EDSRV


#### Memory consumption



**Note:** EDSRV has 4 GB memory.



## CPU consumption



Memory and CPU consumption is not the same in the first 5 minutes of the scenarios' execution. During this period, ACSR, CITSRV1 and CITSRV2 do not consume a lot of CPU and memory, whereas EDSRV consumes a lot. It is indeed the period of document production, where Connect-It queries the source database; where resources are consumed on the source database to perform the queries (such as on EDSRV), and other servers wait for the result.

During the treatments, ACSR resources are not completely used: more than 70% of the CPU is available and only half of the memory is used. Nevertheless, CITSRV1 and CITSRV2 CPU are completely used. This means that the Connect-It servers are at their maximum, but not the Asset Manager server. So it is still possible to achieve better performance results by adding a third Connect-It server.

## For more information

Please visit the HP Software support Web site at:

<http://www.hp.com/managementsoftware/support>

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valuable support customer, you can benefit by being able to:

- Search for knowledge documents of interest
- Submit and track progress on support cases
- Submit enhancement requests online
- Download software patches
- Manage a support contract
- Look up HP support contacts
- Review information about available services
- Enter discussions with other software customers
- Research and register for software training

**Note:** Most of the support areas require that you register as an HP Passport user and sign in. Many also require an active support contract.

To find more information about support access levels, go to the following URL:

[http://www.hp.com/managementsoftware/access\\_level](http://www.hp.com/managementsoftware/access_level)

To register for an HP Passport ID, go to the following URL:

<http://www.managementsoftware.hp.com/passport-registration.html>

### ***Limited responsibility clause***

Asset Manager is integrated with several third-party applications. Examples: Database engines, Web servers, single sign-on software, load-balancing and clustering hardware and software solutions, reporting software such as Crystal Reports, etc.

Support for these applications is limited to their interface with Asset Manager. Support does not cover installation problems, setup and customization problems nor malfunctioning of the third-party application.

White papers contain examples of implementations that may work in your environment with or without customization. There is no guarantee that this will be the case. It could also be that some of the solutions covered by white papers appear as standard features in a future release of the software. When this is the case, there is no guarantee that you will be able to upgrade the solution you implemented based on the white paper to the equivalent standard feature.

© 2008 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. 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.

Itanium is a trademark or registered trademark of Intel Corporation in the U.S. and other countries and is used under license.

Using AM 5.x and ED to start a license management process .doc

