



Universal CMDB

Software Version: 10.31

Sizing Guide

Document Release Date: December 2016

Software Release Date: December 2016



Hewlett Packard
Enterprise

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise 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. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise 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.

Copyright Notice

© 2002 - 2016 Hewlett Packard Enterprise Development LP

Trademark Notices

Adobe™ is a trademark of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: <https://softwaresupport.hpe.com/>.

This site requires that you register for an HP Passport and to sign in. To register for an HP Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** on the HP Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support site at: <https://softwaresupport.hpe.com/>.

This website provides contact information and details about the products, services, and support that HPE Software offers.

HPE 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 valued support customer, you can benefit by using the support website to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HPE support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and to sign in. Many also require a support contract. To register for an HP Passport ID, click **Register** on the HPE Support site or click **Create an Account** on the HP Passport login page.

To find more information about access levels, go to: <https://softwaresupport.hp.com/web/softwaresupport/access-levels>.

HPSW Integrations Catalog accesses the new HPSW Integrations and Solutions Catalog website. This site enables you to explore HPE Product Solutions to meet your business needs, includes a full list of Integrations between HPE Products, as well as a listing of ITIL Processes. The URL for this website is <https://softwaresupport.hpe.com/km/KM01702731>.

Contents

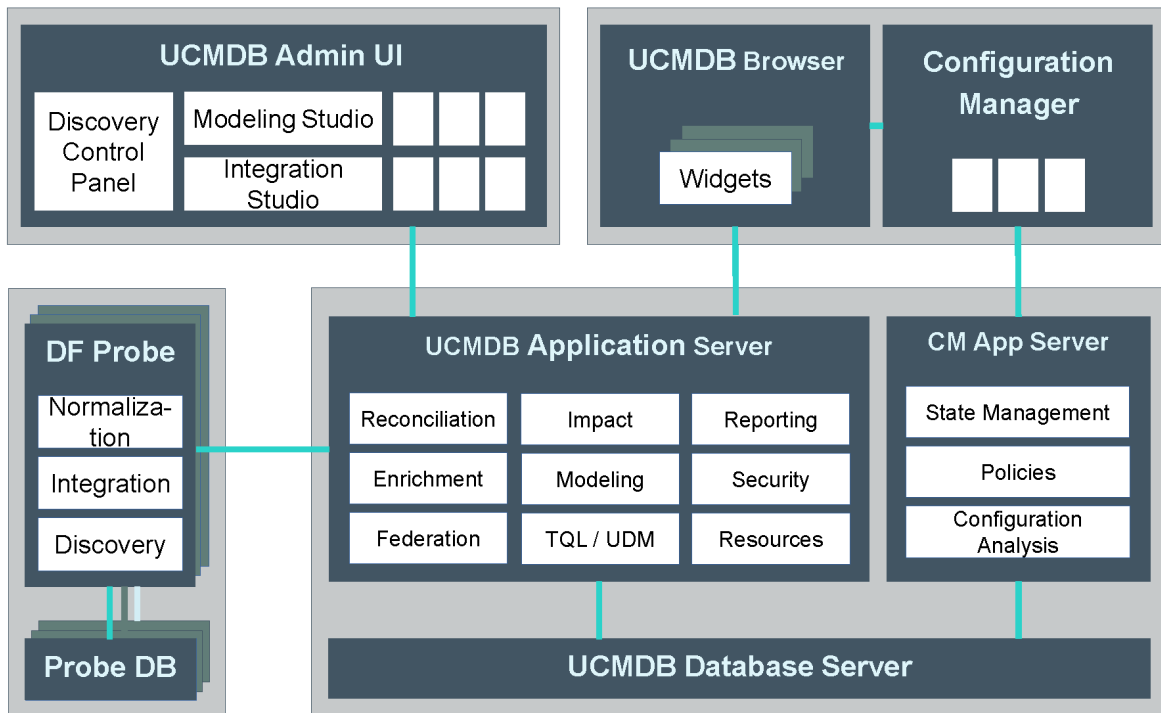
Chapter 1: CMS Sizing	4
Chapter 2: Right Sizing - High Performance HPE Configuration Management System	5
Chapter 3: Infrastructure Requirements	7
Chapter 4: Enterprise Grade Configuration	8
Chapter 5: Sizing Reference for Automated Service Modeling	10
Scenario 1. Service Discovery Only	10
Objectives	10
Test Environment	11
Configurations	12
Scenario 2. Service Discovery and Universal Discovery	13
Objectives	14
Test Environment	14
Discovery Schedule	15
Configurations	16
Chapter 6: Universal Discovery Capacity	18
Planning Your Discovery Schedule	19
Sizing XML Enricher	21
Metrics of Interest for 20K Scans Per Week - A Real Customer Environment	23
20,000 Scans Per Week	23
Metrics of Interest for 50,000 Scans Biweekly - Lab Environment	27
Test Summary	27
Test Results	29
Impact Factors for the Capacity - 50K	30
Metrics of Interest for 75K Scans Per Two Weeks	32
75,000 Scans Per Two Weeks	32
Enterprise Grade Deployment	33
Chapter 7: Other Recommendations	34
Send documentation feedback	35

Chapter 1: CMS Sizing

Welcome to the *HPE Universal CMDB Sizing Guide*. The purpose of this document is to help Universal CMDB (UCMDB) 10.22 (or later) and Universal Discovery (UD) 10.22 (or later) customers to right size your CMS infrastructure to support your CMDB strategy. The objective of the sizing is to estimate the system resource required to ensure the deployed system meets the performance objectives.

Chapter 2: Right Sizing - High Performance HPE Configuration Management System

This chapter helps you to right size the UCMDB/UD architecture to scale the discovery needs for your enterprise. Right sizing enables enterprises to deliver value out of CMS design.



UCMDB/UD application performance depends on many factors. Amount of data (discoverable and non-discoverable) that will be stored in UCMDB/UD is just one aspect.

CMS Sizing



Small

< 2 million
Cls & Relationships



Standard

2 - 10 million
Cls & Relationships



Enterprise

10 - 150 million
Cls & Relationships

Chapter 3: Infrastructure Requirements

The performance of UCMDB Server, Data Flow Probe server, and UCMDB database server is critical for running CMS operations successfully.

UCMDB Server

Deployment	CPU	RAM	Disk Space
Small	4 cores	16 GB	60 GB
Standard	8 cores	32 GB	80 GB
Enterprise	24 cores	64 GB	100 GB

Data Flow Probe Server

Deployment	CPU	RAM	Disk Space
Small	4 cores	12 GB	100 GB
Standard	8 cores	16 GB	200 GB
Enterprise	16 cores	24 GB	300 GB

Right Sizing CMS Database

Choosing the right database based on the enterprise requirement will help you understand the max limits we have on the total number CIs you can discover with respect to different RDBMS options we have. Check the *HPE Universal CMDB Database Guide* to right size your CMS database.

Database	Maximum # CIs and Relationships
Oracle Database 12c Enterprise Edition x64 - 12.1.0.2.0	150 million CIs and relationships
MS SQL Server 2014 Enterprise Edition x64 - 12.0.2000.8	125 million CIs and relationships
PostgreSQL	12.5 million CIs and relationships

For an enterprise grade customer, HPE recommends to use physical database servers with 24 Cores and 64GB RAM.

Chapter 4: Enterprise Grade Configuration

Below configurations will help enterprise customers to scale inventory and agentless discovery jobs. Probe memory settings and thread configurations can be adjusted based on the discovery needs.

The following configurations are based on Union mode probes. The settings might change for Separate mode probes.

Property files on the probe	Setting Details	Comments
hp\UCMDB\DataFlowProbe\bin\ WrapperEnv.conf	set.GATEWAY_MIN_ MEM=2048 set.GATEWAY_MAX_ MEM=8192 set.MANAGER_MIN_ MEM=2048 set.MANAGER_MAX_ MEM=8192	Increase the amount of memory used by JVM, so that it can run more discovery jobs in parallel
hp\UCMDB\DataFlowProbe\conf\ postgresql.conf	shared_buffers = 1024MB	Increase the amount of memory that PostgreSQL can use for buffering the data in memory
hp\UCMDB\DataFlowProbe\conf\ DataFlowProbe.properties	appilog.agent.local.services. poolThreads=200 appilog.agent.local.services. defaultMaxJobThreads=40 appilog.agent.probe.max Connection=120	Increase the number of threads to be used by the probe. This increases the concurrency of operations and allows more discovery jobs to run simultaneously.

Value Reference for Global Settings

The following table describes maximum value from recent cases for some global settings. You can consider the values below as a reference if you want to increase the value of the following global settings.

Global Setting Name	Visibility	Description	Default Value	Max Value from Recent Cases	Reference
dal.link.condition.max.result.size	UI	The number of maximum allowed results of link condition	500,000	50,000,000	
dal.object.condition.max.result.size	UI	The number of maximum allowed results of objects condition	2,000,000	50,000,000	QCIM1H107996
quota.name.server.model.objects	UI	Max number of Objects in Server	20,000,000	40,000,000	QCIM1H109513
quota.name.customer.model.objects	UI	Max Number of Objects in the Customer Model	20,000,000	40,000,000	QCIM1H109513
reconciliation.connected.cis.fuse	JMX		20000		
tql.group.collectors.result.size	UI	TQL Group Collectors Result Size	200,000	300,000	QCIM1H111087
tql.instances.max.allowed	UI	Max instances allowed to be returned in instances request	50000		
tql.max.objects.visit.model.calc.task	UI	Max number of objects allowed visiting during a single search in the CMDB model	30,000,000	200,000,000	QCIM1H110392

Chapter 5: Sizing Reference for Automated Service Modeling

This chapter assists you to determine the required system resources and corresponding parameter settings of the UCMDB Server and the Data Flow Probe to run Service Discovery in different scenarios with satisfactory performance.

This includes:

Scenario 1. Service Discovery Only	10
Objectives	10
Test Environment	11
Configurations	12
Scenario 2. Service Discovery and Universal Discovery	13
Objectives	14
Test Environment	14
Discovery Schedule	15
Configurations	16

Scenario 1. Service Discovery Only

This chapter describes the scenario that you only run Service Discovery.

Objectives

The following table lists the expected time to complete all service discoveries based on the data we collected in performance tests.

Number of Services	First Run of Service Discovery	Scheduled Rerun of Service Discovery
<100	<4 hours	<4 hours
<500	<12 hours	<12 hours

Note: Scheduled Rerun of Service Discovery means rerun all service discoveries on a periodic basis, and the interval can be adjusted.

Test Environment

The performance tests are performed on an environment that meets the following preconditions:

- One service consists of less than 30 Running Software instances, which run on different hosts.
- One host can be discovered by the Host Discovery by Shell job within two minutes if everything works as expected.

Note:

- If the actual environment cannot meet these preconditions, the Service Discovery performance may not reach the above objectives. In this case, you can adjust several configuration settings to improve the Service Discovery performance.
- There is no CI and Relationship discovered by Universal Discovery.
- The numbers of CIs and Relationships reported by Service Discovery are shown in the table below.

CI Types	100 Services	500 Services
Node	6000	30000
IP Address	6000	30000
J2EE Domain	6000	30000
Running Software	12000	60000
J2EE Application	6000	30000
Service Connection Point	6000	30000
IP Service Endpoint	6000	30000
Consumer-Provider Link	12000	60000
Composition Link	18000	90000
Membership Link	6000	30000
Ownership Link	12000	60000
Usage Link	6000	30000

Configurations

This chapter lists the hardware and configuration settings that we used in performance tests. The real numbers may vary depending on the actual environment.

Infrastructure

The following table lists the configurations we used in performance tests for the scenario of Service Discovery only. They have been verified in testing environment and should be regarded as minimum conditions for Service Discovery only environment.

Role	CPU	RAM
UCMDB Server	8 Cores	32 GB
Data Flow Probe	8 Cores	16 GB
Database	8 Cores	32 GB

In an environment that has lower hardware configurations, the following results may occur:

- High CPU and memory usage
- Service Discovery cannot finish within the expected time

JVM Heap Size

The following table lists the JVM heap size configurations used in the performance tests.

It is recommended that you make sure the environment meets the following configurations before you run Service Discovery.

Role	-Xms	-Xmx
UCMDB Server	4096 MB	8192 MB
Data Flow Probe	2048 MB	4096 MB

You can adjust the JVM heap size in the following manner:

- For UCMDB server, set the **wrapper.java.maxmemory** parameter in the **<UCMDB_Server_Home>\bin\wrapper-platform.conf** file.

- For Data Flow Probe, set the **set.GATEWAY_MAX_MEM** and **set.MANAGER_MAX_MEM** parameters in the `<DataFlowProbe_Home>\bin\WrapperEnv.conf` file. For more information, see the *HPE Universal CMDB Data Flow Management Guide*.

Job Execution Threads

By default, the thread pool size on the Data Flow Probe is 80. Increasing this parameter can improve the Service Discovery performance, but will increase the workload on the Data Flow Probe. If you need to run a large number of service discoveries, you can increase this parameter accordingly.

The following table lists the thread distribution used in the performance tests, and we increased the thread pool size to 120.

Job Name	Number of Threads
Business Element Resolver	8
Host Discovery by Shell	64
Running Software Dependencies	16

You can adjust the thread pool size by setting the `appilog.agent.local.services.poolThreads` parameter in the **DataFlowProbe.properties** file. For more information, refer to the *HPE Universal CMDB Data Flow Management Guide*.

Number of Data Flow Probes

In the performance tests, we used only one Data Flow Probe. Tests show that a Data Flow Probe with the above configurations is capable of handling 500 service discoveries. If your Data Flow Probe cannot meet the above configurations, you can distribute the service discoveries to multiple Data Flow Probes.

If you want to employ more Data Flow Probes, keep in mind that more Data Flow Probes can speed up Service Discovery but will also increase the workload on the UCMDb Server. Therefore, consider using more Data Flow Probes if the UCMDb Server is not heavily loaded.

Scenario 2. Service Discovery and Universal Discovery

This chapter describes the scenario that you already have Universal Discovery and want to introduce Service Discovery.

Objectives

The following table lists the expected time to complete all service discoveries based on the data we collected in performance tests.

Number of Services	First Run of Service Discovery	Scheduled Rerun of Service Discovery
<100	<8 hours	<8 hours
<500	<24 hours	<24 hours

Note: Scheduled Rerun of Service Discovery means rerun all service discoveries on a periodic basis, and the interval can be adjusted.

Test Environment

The performance tests are performed on an environment that meets the following preconditions:

- One service consists of less than 30 Running Software instances, which run on different hosts.
- One host can be discovered by the Host Discovery by Shell job within two minutes if everything works as expected.
- Existing UCMDB data has less than 30,000 nodes, and less than 20 million CIs and relationships.
- Service Discovery and Universal Discovery never run in the same time window. For more information, see "[Discovery Schedule](#)" on the next page.

Note:

- This scenario is slower than Scenario 1, because the large amount of existing UCMDB data increases the workload of processing the results sent from the Data Flow Probes.
- If the actual environment cannot meet these preconditions, the Service Discovery performance may not reach the above objectives. In this case, you can adjust several configuration settings to improve the Service Discovery performance.
- The numbers of CIs and Relationships reported by Service Discovery are shown in the table below.

CI Types	100 Services	500 Services
Node	6000	30000
IP Address	6000	30000
J2EE Domain	6000	30000
Running Software	12000	60000
J2EE Application	6000	30000
Service Connection Point	6000	30000
IP Service Endpoint	6000	30000
Consumer-Provider Link	12000	60000
Composition Link	18000	90000
Membership Link	6000	30000
Ownership Link	12000	60000
Usage Link	6000	30000

Discovery Schedule

We strongly recommend that you schedule Service Discovery and Universal Discovery, so that Service Discovery, Universal Discovery, and Integration Service can run in different time slots.

Note: An overlapped execution of Service Discovery and Universal Discovery may,

- Slow down both types of discoveries.

Because the UCMDB Server uses the "first-come, first-served" mechanism, a later-arrived job must wait to be executed until the previous jobs are finished.

- Further increase the workload of the UCMDB Server and the Data Flow Probe.

For example, assume that you run Universal Discovery jobs once every day from Monday to Sunday before introducing Service Discovery. To avoid the overlapped execution of Service Discovery and Universal Discovery, consider the following policies:

- Define a job execution policy that disallows the execution of Universal Discovery jobs on weekends.
- Change the Service Discovery schedule so that Service Discovery is allowed to run on weekends.

For more information about job execution policies, refer to the *HPE Universal CMDB Data Flow Management Guide*.

Configurations

This chapter lists the hardware and configuration settings that we used in performance tests. The real numbers may vary depending on the actual environment.

Infrastructure

Due to the large number of CIs and relationships discovered by Universal Discovery, more resources are required in the performance tests.

Role	CPU	RAM
UCMDB Server	24 Cores	64 GB
Data Flow Probe	16 Cores	24 GB
Database	24 Cores	64 GB

In an environment that has lower hardware configurations, the following results may occur:

- High CPU and memory usage
- Service Discovery cannot finish within the expected time

JVM Heap Size

The following table lists the JVM heap size configurations used in the performance tests.

It is recommended that you make sure the environment meets the following configurations before you run Service Discovery.

Role	-Xms	-Xmx
UCMDB Server	8192 MB	16384 MB
Data Flow Probe	4096 MB	8192 MB

You can adjust the JVM heap size in the following manner:

- For UCMDB server, set the **wrapper.java.maxmemory** parameter in the **<UCMDB_Server_Home>\bin\wrapper-platform.conf** file.

- For Data Flow Probe, set the **set.GATEWAY_MAX_MEM** and **set.MANAGER_MAX_MEM** parameters in the `<DataFlowProbe_Home>\bin\WrapperEnv.conf` file. For more information, see the *HPE Universal CMDB Data Flow Management Guide*.

Job Execution Threads

If you isolate the executions of Service Discovery and Universal Discovery as recommended, you can consider increasing the thread pool size of the Data Flow Probe to 120 and distributing the threads to different Service Discovery jobs as listed below.

Job Name	Number of Threads
Business Element Resolver	8
Host Discovery by Shell	64
Running Software Dependencies	16

Number of Data Flow Probes

If there is no IP range overlap between Service Discovery and Universal Discovery, it is recommended to use a dedicated Data Flow Probe to run Service Discovery. Of course, you can use your existing Data Flow Probes, and Service Discovery will share Data Flow Probes with Universal Discovery.

If you want to employ more Data Flow Probes, keep in mind that more Data Flow Probes can speed up Service Discovery but will also increase the workload on the UCMDB Server. Therefore, consider using more Data Flow Probes if the UCMDB Server is not heavily loaded.

Chapter 6: Universal Discovery Capacity

This chapter includes:

Planning Your Discovery Schedule	19
Sizing XML Enricher	21
Metrics of Interest for 20K Scans Per Week - A Real Customer Environment	23
20,000 Scans Per Week	23
Metrics of Interest for 50,000 Scans Biweekly - Lab Environment	27
Test Summary	27
Test Results	29
Impact Factors for the Capacity - 50K	30
Metrics of Interest for 75K Scans Per Two Weeks	32
75,000 Scans Per Two Weeks	32
Enterprise Grade Deployment	33

Planning Your Discovery Schedule

Go to [Discovery Use Cases](#) and identify the list of discovery jobs you want to run to meet your operational needs.

Let us assume two basic use cases:

- Agentless discovery
- Agent-based discovery (two options, with or without call home setting)

Agentless Basic Discovery	Agent-based discovery	
	Option 1: For client devices with call home setting	Option 2: For datacenter devices without call home setting
ICMP Ping	Call home processing	ICMP Ping
Host Connection by Shell	Inventory discovery by scanner	Host Connection by Shell
Host Resources by Shell		Inventory discovery by scanner
Host Applications by Shell		


When you think about these basic use cases, a couple of questions will arise:

- How many probes are required to complete your discovery?
- How frequently can you run these jobs? What is the valid business use case on the rediscovery interval?
- Is there an upper limit on the number of trigger CIs a probe can handle?

Planning your discovery capacity based on your use cases will give you good handle on the number of related CIs per discovered node. When planning capacity, among other issues, you should consider the ratio of managed nodes in your CMDB to node-related CIs. Node-related CIs include all CIs of types that are subclasses of Application Resource, Node Element, or Running Software.

The following diagram gives you an idea on the number of node-related CIs you can discover for each managed node based on the size and use cases. This number depends on the size of your deployment and the number of managed nodes the more managed nodes you maintain in the CMDB, the fewer node-related CIs you can discover for each managed node.

Managed Nodes to Node-Related CIs Ratio



For example, in an enterprise deployment, if you are running 336,000 managed nodes, you can discover 160 node-related CIs for each managed node. If you are running only 108,000 managed nodes, you can discover 500 resource CIs for each managed node. This ratio will give you an idea on the scalability and deployment strategy.

Sizing XML Enricher

The XML Enricher is used in UD Inventory scan jobs to process scan files. If your CMS environment leverages UD Inventory jobs, then you need to plan for extra memory requirement on the probes for the XML Enricher. Starting from Universal Discovery version 10.10, the XML Enricher is re-written in Java and uses 64-bit JRE.

The table below summarizes the memory and thread settings required to run XML Enricher process:

Probe Deployment Mode	XML Enricher Thread Count	XML Enricher JVM Memory (MB)	XML Enricher Mode
Small	1	3584	Database
Standard	2	5120	Memory
Enterprise	4	8192	Memory

Small deployment of the probe. In this mode the XML Enricher is configured to use the DB mode to save memory.

DB mode for the XML Enricher is an option which can be leveraged for small size deployment where it stores the largest portion of the SAI files and the file data in the temporary database instead of keeping it in memory. It saves quite a bit of memory, but the processing speed for incoming scan files is degraded and disk I/O on the server running the probe will show some spikes.


Standard deployment of the probe. The number of XML Enricher threads plays a vital role in the probe sizing for the Standard probe deployment.

Enterprise deployment of the probe. The Enterprise mode uses more threads and more memory for processing, but has the best performance (provided that the probe's hardware is sized appropriately). If one needs to scale even more (if there is a build-up of scan files in the **Incoming** directory of the XML Enricher), one can increase the amount of threads further.

HPE CMS product R&D team has tested up to 8 XML Enricher threads working in parallel. But in this case large amounts of memory (up to 16 GB) needs to be allocated for its JVM, so it does not run out of the probe memory when processing very large scan files.

Note: The XML Enricher runs as a separate process from the probe and has its own JVM. Even if the XML Enricher runs out of memory, it does not affect the probe process that runs separately. Increasing the XML Enricher thread count helps to improve scan processing performance.

Higher numbers of threads that are allocated to the XML Enricher service results in higher amounts of CPU and memory that the XML Enricher process consumes. As a result, the processing of scan files runs faster and faster. However, too many threads may slow down the processing if there is not enough CPU or memory resources available. In this case, disk performance becomes a constraining factor.



Metrics of Interest for 20K Scans Per Week - A Real Customer Environment

Note: This scenario was verified on UCMDB version 10.22.

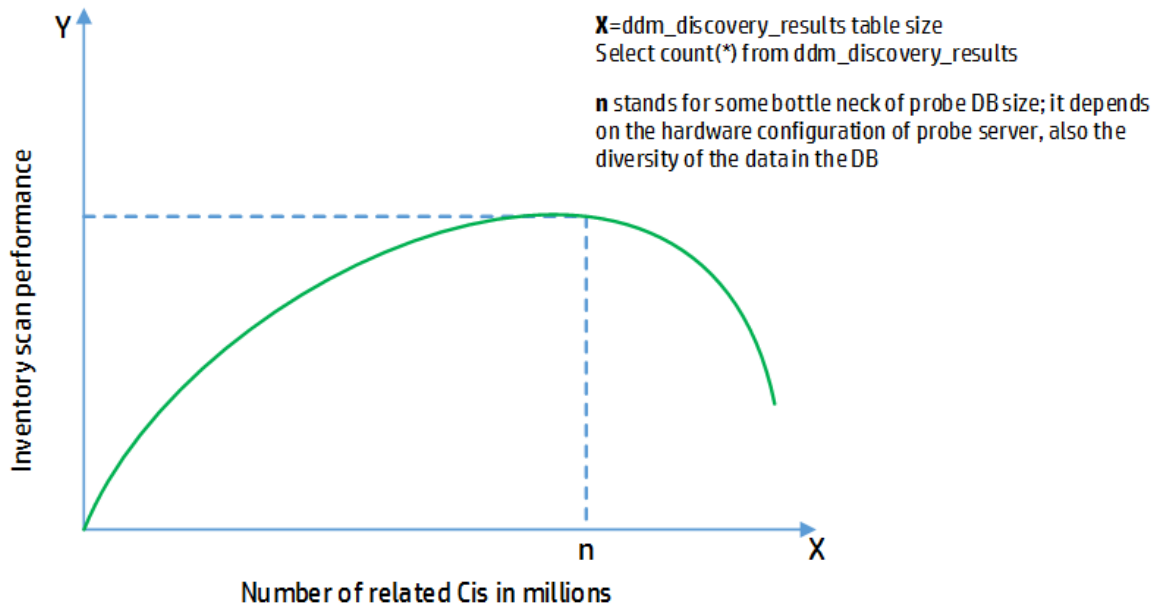
UCMDB/UD can be deployed in many ways and you can configure discovery in various ways. For example, you have the option to choose **Zone-Based Discovery** to run inventory scans and **Discovery Modules/Jobs** to configure other discovery jobs. The metrics below provides a guidance on the sizing of a probe server that the CMS product team has fully tested in customer environments.

The following typical scenario has been tested for your reference.

Scenario	Description
20,000 Scans per week	This scenario is based-on some impact factors. For example, average number of CIs per scan, DB size, network latency, device online timeslot, and so on.

20,000 Scans Per Week

An enterprise grade probe server is capable of handling at least 20,000 scans per week on the condition that the average number of related CIs for each scan is 250. This means that the total count of related CIs on the probe DB (**ddm_discovery_results**) will reach 5 million CIs. When the total count of related CIs on the probe DB continues to grow while more scans are executed, a probe server performance decrease is observed. This depends on the diversity of the data in the DB. As shown in the following diagram, there is a capacity threshold on the probe DB, beyond which the capacity decreases a lot.



The numbers below are benchmark results from lab tests. Test results are based on enterprise grade UCMDB and Data Flow Probe servers. Time in seconds shows the amount of time it took to run the discovery jobs in the HPE Lab environment per discovery trigger (that is, per node or IP).

Discovery Job	Low Network Latency (sec)	High Network Latency (sec)
Host Connection by Shell	2.28	4.72
Host Resources	365.33	474.929

IOPS Numbers from sampling analysis for Inventory scan on virtual environment client devices

The VMware vRealize tool was used to measure IOPS peak usage. 2 to 3 scan files are received per min to the incoming folder. Number of XML Enricher is set to 2 threads.

Scanner Configurations	# Threads	Peak IOPS
Key hardware data and no file data	8	900
Key Hardware and File Data (Key installation Directories)	4	2300
Key Hardware and all file data	4	3400

Test Bed – UCMDB/UD deployment and configuration

- Enterprise edition in ["Infrastructure Requirements "](#) on page 7

Note: There are 3 types of deployment, and the test bed we used is Enterprise edition.

- ["Enterprise Grade Configuration"](#) on page 8

Note:

- VMware virtualization were used to spin up probes machines and dedicated VCPUs were assigned.
- High end performance storage
- Inventory scans were tested with NTCMD, SSH, and UDA protocols.

Impact factors on the discovery capacity per probe per week

The data in the table below are collected during the testing for the probe capacity of 20,000/250 scan per week.

Impact Factors show what will impact the probe capacity.

The **Value** column shows the value used for the impact factor during our testing.

Impact Factors	Value used in testing	Comments
Average number of related CIs per scan	250	The type of operating system has impact on the number of installed software CIs. For example, enabling BaseUnixOs.sai for the XML Enricher to process Unix scan files will introduce a lot of installed software.
Frequency of reassign IP for client device by DHCP	15,000 client devices: every 7 Days	
The number of management zones used	0	Refer to the management zone configuration.
Connect failure ratio		The connection may fail

Impact Factors	Value used in testing	Comments
		due to network issue or during moving. This depends on your environment.
Network latency between probe & discovery devices	In the same LAN: ~320 ms	Usually this is the ping time from probe to discovery nodes.
Number of IPs	0.5 million	
Devices online timeslots	Data Center: 24 hours/day Client: 8 hours/day	
Change Ratio for the discovered CIs	Less than 11%	The higher change ratio of discovered CIs, the longer time it takes to push the changed CIs to the UCMDB server.
Discovery jobs	<p>For fixed IPs:</p> <ul style="list-style-type: none"> • Range IPs by ICMP • Host Connection by Shell • Inventory Discovery by Scanner <p>For Dynamic IPs:</p> <ul style="list-style-type: none"> • Call home processing • Inventory Discovery by Scanner 	
# Max worker number per probe	200	
# Threads of Inventory Discovery by Scanner	40	
Schedule of Inventory Discovery by Scanner	Weekly: 24x7, no blackout policy	
Number of total CIs in UCMDB	Data Center: 18 million (CIs & relationships) Client: 25.2 million (CIs & relationships)	Pay attention to installed software CI. Usually it plays a great portion.
The size of probe DB tables	ddm_discovery_results: 3.5 million ddm_map_objectid: 3 million	

Metrics of Interest for 50,000 Scans Biweekly - Lab Environment

Test Summary

Conclusion

This testing is running against 50,000 devices (per device per scan) in a simulated environment. The time spent on scanning the remote devices are not counted in; the network latency does not exist since the simulated environment is running on the lab in the same VLAN, which means there is no scanner connect failure. The auto-deletion seldom happen because there is no device change in the simulated environment.

The whole testing finished in one week. Considering this result, it is confident that in a biweekly schedule, 50,000 devices can be scanned in a real customer environment, thus Data Flow Probe capacity can be increased to 50,000 devices, with each device contributing 360 CIs and relationships, and the total data volume will be 17.7 million CIs and relationships in UCMDB server.

Background

This testing is designed to simulate a scenario of 50,000 devices by running 3 jobs (**Range IPs by ICMP**, **Host Connection by Shell**, and **Inventory Discovery by Scanner**), in order to measure how long it takes to finish discovering these devices.

The testing environment is populated first with these 50,000 scan files, with all probe data cleared. This is to make sure there is no redundant and touching which will speed up the whole data-in, the operations on server side are mainly update and merge.

Test Environment

Role in Group	CPU	RAM	OS	Hard Drive	Port	SID
Database Server - Oracle	32 cores 2.40Ghz	128GB	Red Hat Enterprise Linux 7.0 64-bit	1.4 TB	1521	sid=pcoe
UCMDB Server 10.30	16 cores 2.40Ghz	128GB	Windows Server 2008 R2 64-bit	1.3 TB		
Data Flow Probe 10.30 simulated environment	32 cores 2.40Ghz	128GB	Windows Server 2008 R2 64-bit	1.3 TB		

Probe Configuration

Property files on the probe	Setting Details	Comments
hp\UCMDB\DataFlowProbe\bin\xmlenricher\ WrapperEnricher.conf	wrapper.java.maxmemory=8192 (the first 4 days) wrapper.java.maxmemory=16384 (the last 2.5 days)	Change the XML Enricher configuration to use "Enterprise" deployment
hp\UCMDB\DataFlowProbe\conf\enricher.properties	max_enricher_thread_number=4 (the first 4 days) max_enricher_thread_number=8 (the last 2.5 days)	Change the XML Enricher configuration to use "Enterprise" deployment
hp\UCMDB\DataFlowProbe\bin\ WrapperEnv.conf	set.GATEWAY_MIN_MEM=2048 set.GATEWAY_MAX_MEM=8192 set.MANAGER_MIN_MEM=2048 set.MANAGER_MAX_MEM=8192	Increase the amount of memory used by JVM, so that it can run more discovery jobs in parallel
hp\UCMDB\DataFlowProbe\conf\postgresql.conf	shared_buffers = 1024MB	Increase the amount of memory that PostgreSQL can use for buffering the data in memory
hp\UCMDB\DataFlowProbe\conf\ DataFlowProbe.properties	appilog.agent.local.services.poolThreads=110 appilog.agent.local.services.defaultMaxJobThreads=60 appilog.agent.probe.maxConnection=120	Increase the number of threads to be used by the probe. This increases the concurrency of operations and allows more discovery jobs to run simultaneously.

UCMDB Server Configuration

Impacting Factors	Category	Default Value	Value for a large deployment	Comment
Max number of Objects in Server	Global Setting Name	20,000,000	30,000,000	These settings are increased to prevent errors that occur once certain fuse types are reaching their default limits
Maximum number of elements	Global Setting	100,000	4,000,000	

UCMDB Server Configuration, continued

Impacting Factors	Category	Default Value	Value for a large deployment	Comment
(nodes) in view result	Name			

Job configuration and schedule

Job Name	Schedule	# of Threads
Range IPs by ICMP	daily	1
Host Connection by Shell	daily	60
Inventory Discovery by Scanner	Bi-weekly	40

Test Results

XML Enricher speed:

With `wrapper.java.maxmemory=8192` and `max_enricher_thread_number=4`, it processed about ~6700 scan files per day.

With `wrapper.java.maxmemory=16384` and `max_enricher_thread_number=8`, it processed about ~10000 scan files per day.

The full discovery cycle:

7 days for 43372 devices.

Data-in CI Rate:

~96 CIs/second for Host Connection by Shell

~300 CIs/second for Inventory Discovery by Scanner

Conclusion:

The test confirms that 50K devices per probe is possible.

However, the speed of processing the scan files in your environment may vary depending on the CPU type, disk speed, and crucially scan file size.

Impact Factors for the Capacity - 50K

The data in the table below are collected during the testing for the probe capacity of 50000/360 scan per week, ignoring the time scanner connects and scanner runs in the target node to generate the scan file.

Impact Factors show what will impact the probe capacity. The ones listed here are very important ones.

The **Value** column shows the value used for the impact factor during our testing.

Impact Factors	Value used in testing	Comments
Number of devices in the environment	50,000	
Average number of related CIs per scan	380	Operation system types impact the number of installed software CI a lot. For example, enabling BaseUnixOs.zsai for the XML Enricher to process UNIX scan files will introduce a lot of installed software CIs.
Frequency of new IP address assignment by DHCP	50,000 data center devices: Fixed IPs	
The number of management zones used	0	Refer to the management zone configuration.
Connect failure ratio	0	
Network latency between probe & discovery devices	In the same LAN: < 5ms	Usually this is the ping time from the Data Flow Probe to the nodes being discovered
Number of IPs	50,000	
Devices online timeslots	Data Center: 24 hours/day	
Change Ratio for the discovered CIs	0	The higher change ratio of discovered CIs, the longer time it takes to push the changed CIs to

Impact Factors	Value used in testing	Comments
		the UCMDB server
Discovery jobs	For fixed IPs: <ul style="list-style-type: none"> • Range IPs by ICMP • Host Connection by Shell • Inventory Discovery by Scanner 	
Total number of discovery threads per probe	110	
# Threads of Host Connection by Shell	60	
# Threads of Inventory Discovery by Scanner	40	
Schedule of Inventory Discovery by Scanner	Bi-Weekly: 24x7, no blackout policy	
Number of total CIs in UCMDB	17.7 million (CIs & relationships)	InstalledSoftware and UserSoftwareUtilization CIs (if software utilization is enabled) usually contribute the largest share of CIs/relationships.
The size of probe DB tables	ddm_discovery_results: 8 million ddm_discovery_touch_results: 8 million ddm_map_objectid: 8 million	When the size of DB tables reach 8 million, the operation of getting Data Flow Probe Status from UI may encounter timeout issue sometimes. From the testing, the timeout ratio is 50%. When increasing to 10 million, getting Data Flow Probe Status from UI operation always encounters timeout issue.

Metrics of Interest for 75K Scans Per Two Weeks

The following typical scenario has been tested for your reference.

Scenario	Description
75,000 scans per two weeks	This scenario is based-on the collaboration of Scanner Scheduler, Store and Forward and XML Enricher at enterprise grade configuration.

75,000 Scans Per Two Weeks

An enterprise grade probe can support up to 75,000 scans per two weeks. To achieve this, the combination of the following configurations should be performed.

- Scanner Scheduler
- Store and Forward scenarios: One enterprise probe server and two store and forward servers (specifications with enterprise probe grade server with no other applications running)
- XML Enricher is at enterprise grade configuration
- Only Inventory discovery by Manual Scanner Deployment running on the probe

Test Bed – UCMDB/UD deployment and configuration

- ["Infrastructure Requirements " on page 7](#)
- XML Enricher Enterprise Mode

Enterprise Grade Deployment

Role	Deployment	CPU	Memory (GB)	Linux Swap (GB)	Windows Virtual Memory (GB)	Free Disk Space (GB)	OS / 3rd-Party SW
UCMDB	Enterprise	24 cores <ul style="list-style-type: none"> Intel Dual Core XEON Processor 2.4 GHZ or later AMD OPTERON Dual Core Processor 2.4 GHZ or later 	32	32	48	300	<ul style="list-style-type: none"> Windows 2008 R2 64-bit Red Hat Enterprise LINUX Server Release 6
Data Flow Probe	Enterprise	24 cores Processor 2.4 GHZ or later	24	n/a	36	300	<ul style="list-style-type: none"> Windows 2008 R2 64-bit
Database	Enterprise		64	64	96	500	<ul style="list-style-type: none"> Oracle Microsoft SQL

Chapter 7: Other Recommendations

Virtual Machine Workload Management

Effective virtual machine workload management practices, such as setting metrics, can help you achieve the most efficient workloads and avoid the mistake of over-allocating resources to a virtual machine.

Send documentation feedback

If you have comments about this document, you can [contact the documentation team](#) by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Sizing Guide (Universal CMDB 10.31)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to cms-doc@hpe.com.

We appreciate your feedback!