HP Universal CMDB

Software Version: 10.22

Sizing Guide

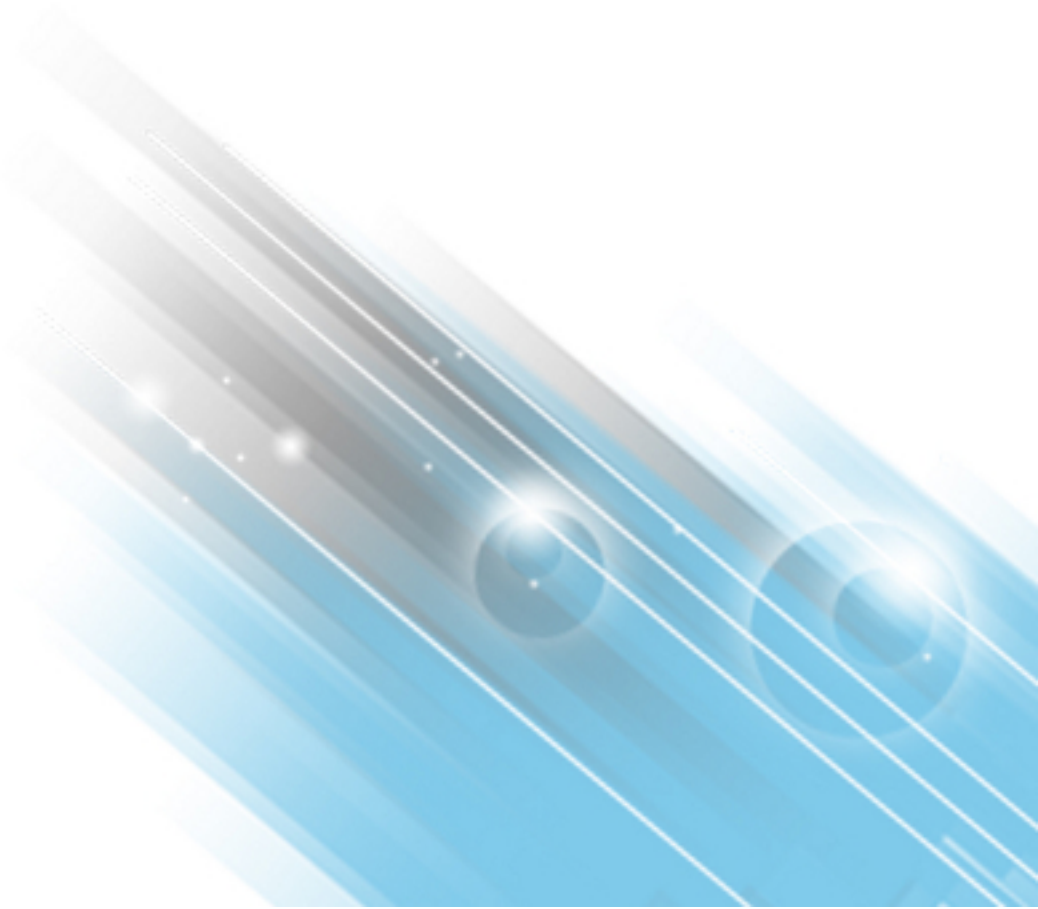

Document Release Date: December 2015
Software Release Date: December 2015

Legal Notices

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.

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.

Copyright Notice

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

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hp.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 HP 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 HP sales representative for details.

Support

Visit the HP Software Support site at: https://softwaresupport.hp.com.

This website 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 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 HP 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 HP 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.

HP Software Solutions Now accesses the HPSW Solution and Integration Portal website. This site enables you to explore HP Product Solutions to meet your business needs, includes a full list of Integrations between HP Products, as well as a listing of ITIL Processes. The URL for this website is http://h20230.www2.hp.com/sc/solutions/index.jsp.

Contents

CMS Sizing

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

This document includes:

Right Sizing - high performance HP Configuration Management System

This section 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.

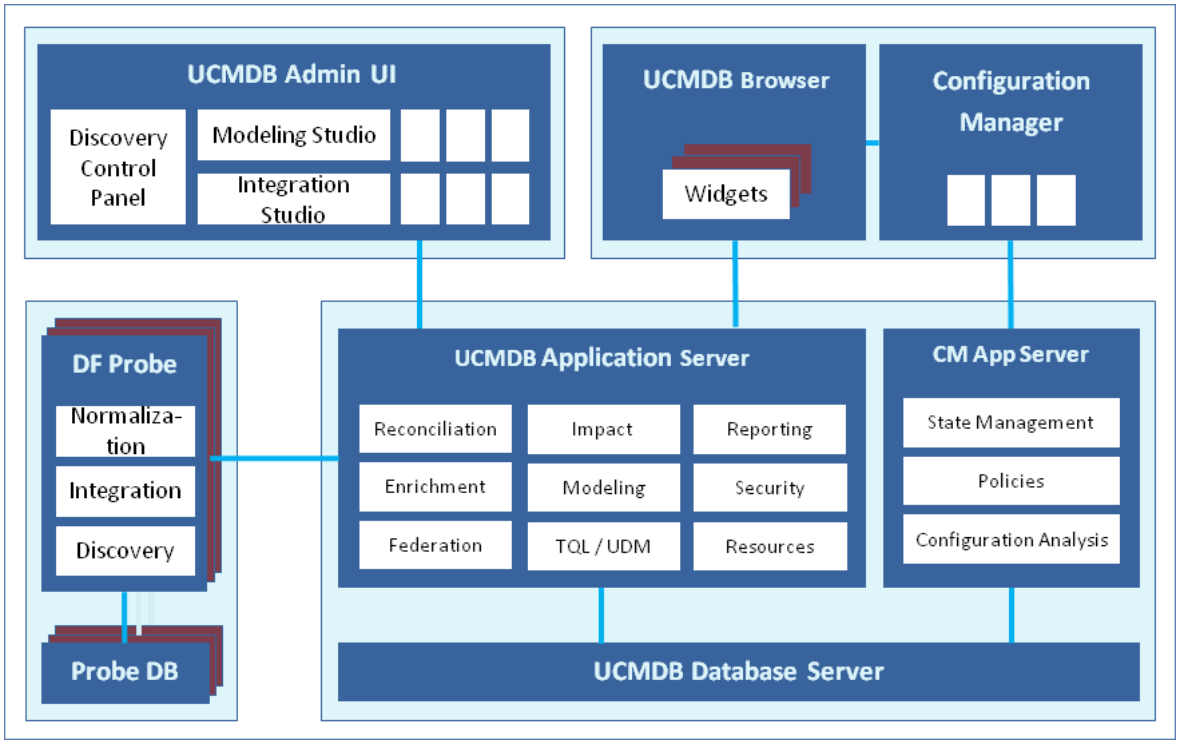

UCMDB/UD application performance depends on many factors. Amount of data (discoverable and non-discoverable) that will be stored in UCMDB/UD is just one aspect.

CMS Sizing

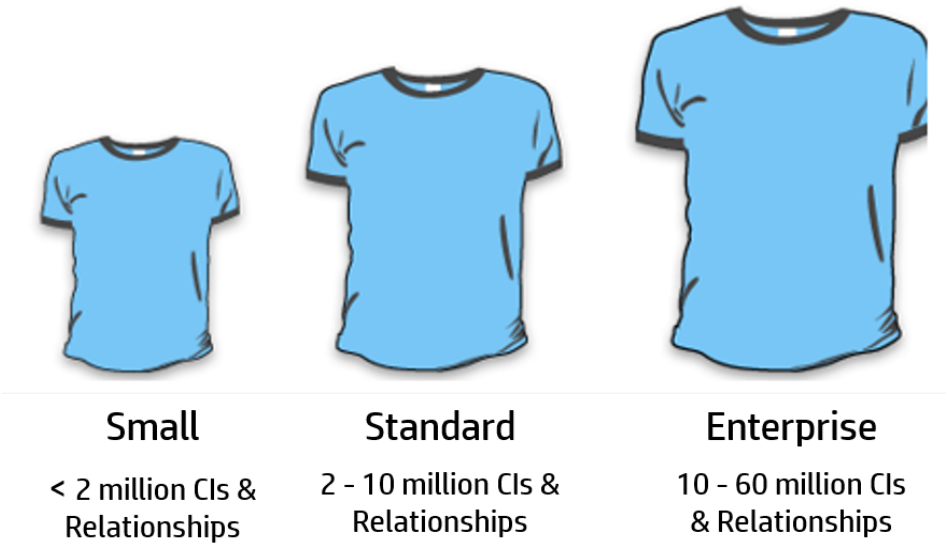
Small
< 2 million CIs & Relationships
Standard
2 - 10 million CIs & Relationships
Enterprise
10 - 60 million CIs & Relationships

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 *HP Universal CMDB Database Guide* to right size your CMS database.

Database

Database	MAXIMUM # CIS & RELATIONSHIPS
Oracle	60 Million CIs and Relationships
Microsoft SQL SERVER 2012	60 Million CIs and Relationships
Microsoft SQL SERVER - previous versions	40 Million CIs and Relationships
PostgreSQL	12.5 Million CIs and Relationships

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

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)

	Agent-based discovery	
Agentless Basic 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

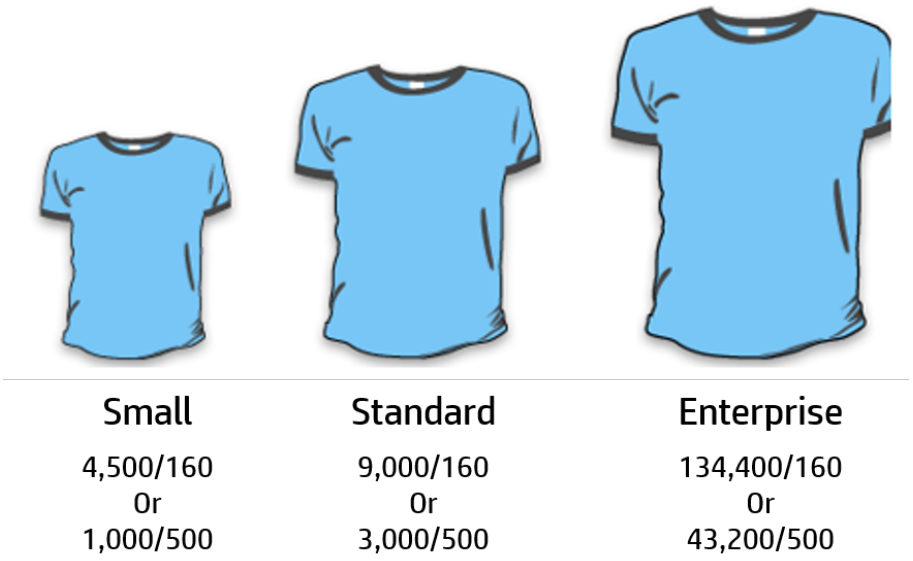

For example, in an enterprise deployment, if you are running 134,400 managed nodes, you can discover 160 node-related CIs for each managed node. If you are running only 43,200 managed nodes, you can discover 500 resource CIs for each managed node. This ratio will give you an idea on the scalability and 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.

HP 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.

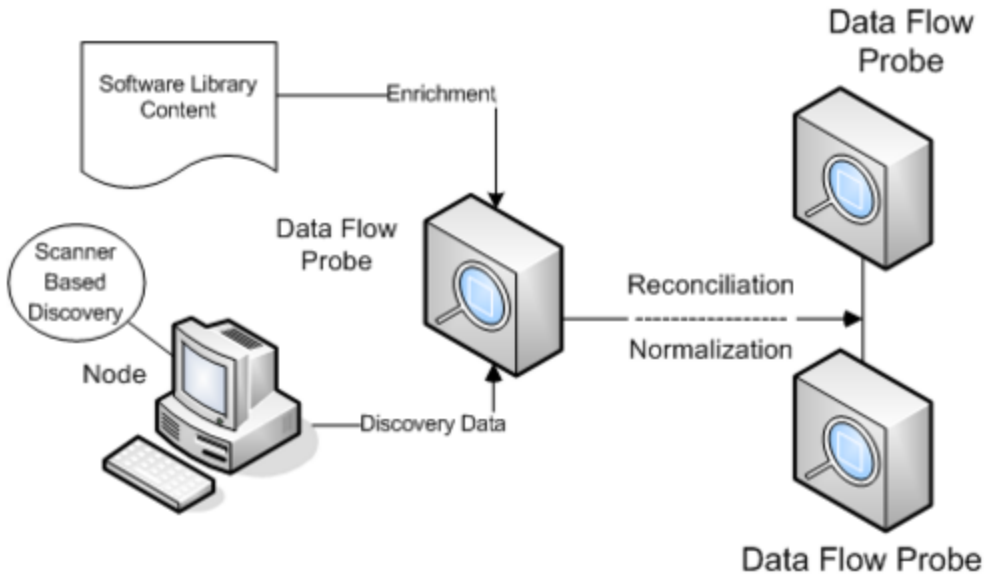

Metrics of Interest

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 two typical scenarios have 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. For more details, see 20,000 Scans per week.
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. For more details, see 75,000 scans per two weeks.

- **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.

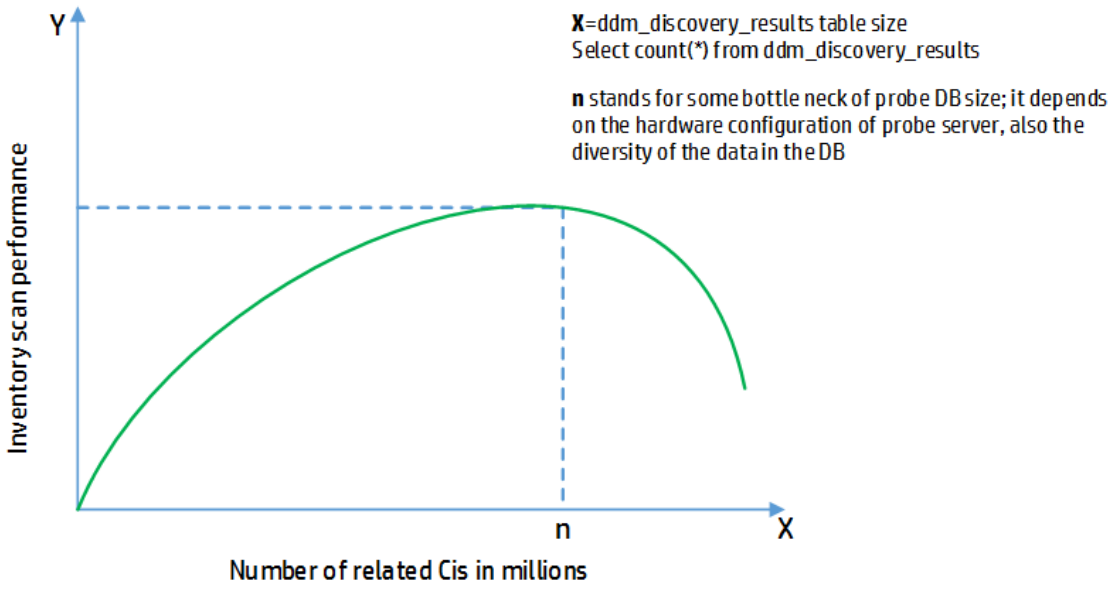

The numbers below are benchmark results from lab tests. Test results are based on enterprise grade UCMDB and Data Flow Probe servers. Time in seconds shows the amount of time it took to run the discovery jobs in the HP Lab environment per discovery trigger (that is, per node or IP).

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 Grade Deployment
- Enterprise Grade Configuration

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/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.

Impact Level indicates how seriously the capacity will be impacted. More stars means higher level of impact. Stars are the weight of each factor.

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

Impact Factors	Impact level (Low 1-5 High)	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	★★★★	15,000 client devices: every 7 Days	

Impact Factors	Impact level (Low 1-5 High)	Value used in testing	Comments
DHCP			
The number of management zones used	★★★★★	0	Refer to the management zone configuration.
Connect failure ratio	★★★★		The connection may fail 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	★★★★	**For fixed IPs:** ◦ Range IPs by ICMP ◦ Host Connection by Shell ◦ Inventory Discovery by Scanner	

Impact Factors	Impact level (Low 1-5 High)	Value used in testing	Comments
		For Dynamic IPs: ◦ 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	

Impact Factors	Impact level (Low 1-5 High)	Value used in testing	Comments
		ddm_map_ objectid: 3 million	

- **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

  - Enterprise Grade Deployment
  - XML Enricher Enterprise Mode

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\xmlenricher\ **WrapperEnricher.conf**	wrapper.java.maxmemory=8192	**For large probes:** Change the XML Enricher configuration to use "Enterprise" deployment
hp\UCMDB\DataFlowProbe\bin\ **WrapperGateway.conf**	Modify the following lines: • JRE 7 and older: wrapper.java.additional.19=-XX:PermSize=512m wrapper.java.additional.20=-XX:MaxPermSize=768m • JRE 8 and later: wrapper.java.additional.21=-XX:MaxMetaspaceSize=768m	Increase other heap settings, so that there is enough special memory available for JVM. These options are needed for JRE 7 only.
hp\UCMDB\DataFlowProbe\conf\ **enricher.properties**	max_enricher_thread_number=4	**For large probes:** 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

, continued

Property files on the probe	Setting Details	Comments
		memory
hp\UCMDB\DataFlowProbe\conf\ **DataFlowProbe.properties**	appilog.agent.local.services.poolThreads=200 appilog.agent.local.services.defaultMaxJobThreads=40 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/UD UI Changes (in the admin UI – Infrastructure Settings)

Global Setting Name	Default Value	Value for a large deployment
Max number of Objects in Server	20,000,000	30,000,000
Max Number of Objects in the Customer Model	20,000,000	30,000,000
TQL Group Collectors Result Size	200,000	400,000

JMX Settings

Global Setting Name	Default Value	Value for a large deployment
tql.max.objects.visit.model.calc.task	30,000,000	70,000,000
tql.group.collectors.result.size	200,000	600,000
tql.instances.max.allowed		600,000
reconciliation.connected.cis.fuse		40,000

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 • Intel Dual Core XEON Processor 2.4 GHZ or later • AMD OPTERON Dual Core Processor 2.4 GHZ or later	32	32	48	300	• Windows 2008 R2 64-bit • Red Hat Enterprise LINUX Server Release 6
Data Flow Probe	Enterprise		24	n/a	36	300	• Windows 2008 R2 64-bit
Database	Enterprise		64	64	96	500	• Oracle • Microsoft SQL

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 virtul 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 10.22)

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!