

IBM Endpoint Manager for Software Use Analysis
Version 9.1 (includes update 9.0.1, 9.0.1.1 and 9.0.1.2)

Overview Guide

The IBM logo, consisting of the letters "IBM" in a bold, black, sans-serif font. The letters are composed of horizontal bars, with the "I" and "M" having a slight shadow effect.

IBM Endpoint Manager for Software Use Analysis
Version 9.1 (includes update 9.0.1, 9.0.1.1 and 9.0.1.2)

Overview Guide



Overview Guide

This edition applies to IBM Endpoint Manager for Software Use Analysis 9.0.1.2 (product number 5725-F57) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright IBM Corporation 2002, 2014.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Overview	1
What's new in this release	1
Features and functions	4
Limitations	5
Business benefits	7
Key concepts	8
Processor value unit licenses	8
Processor value unit table	10
Resource value unit licenses (RVU MAPC)	11
Resource value unit tier table	11
High-water mark	12
Products, components, and bundles	14
Main background application tasks	16

Raw data discovery and matched signatures	18
Contracts	18
Scenarios	18
Software and hardware discovery	19
License consumption verification and reconciliation	20
Continuous license compliance	20

Notices	23
Trademarks	24

Privacy policy considerations	25
--	-----------

Overview



Use this section to become familiar with IBM® Endpoint Manager for Software Use Analysis application update 9.0.1.2 infrastructure and key concepts that are necessary to understand how the application works.

What's new in this release

Software Use Analysis application update 9.0.1.2 provides new features and enhancements that facilitate your work with the application.

- **Update 9.0.1.2 Side-by-side migration form Tivoli Asset Discovery for Distributed 7.5 and 7.5.0.10**

You can migrate from Tivoli Asset Discovery for Distributed 7.5 to the current version of Software Use Analysis. It is a side-by-side migration. The Tivoli Asset Discovery for Distributed and Endpoint Manager agents must be installed on the same computer. The information that is migrated includes:

- Software bundlings
- Software instances that are excluded from pricing calculations
- Directories that are excluded from scanning
- Definitions of VM managers

Raw data and scan results are not migrated because the same data is discovered by Software Use Analysis application update 9.0.1.2 scans.

After the information is migrated, it is applied to the raw inventory that is discovered by the new scans. Your software inventory is managed by the new infrastructure. The migration is performed by the migration tool that provides the following options: export to CSV files, simulation, and migration.

- **Update 9.0.1.2 Scan Health widget**

New widget on the Infrastructure Administrator Dashboard provides the means to check the health of scans that are running in your infrastructure. You can check on which computers a particular type of scan failed, or was not initiated within the last 30 days as well as which computers did not upload scan results to the Endpoint Manager server. You can also learn which computers have an outdated version of the catalog and whether any shared disks in your infrastructure are not scanned.

- **Update 9.0.1.2 Information about software that is installed on shared disks on UNIX**

You can collect information about software that is installed on shared disks on UNIX. The disks are not scanned by default. To enable scanning of shared disks, select the option in the **Initiate Software Scan** fixlet.

- **Update 9.0.1.2 Optimized volume of scanned file data**

The Scanned File Data report provides information about files with particular extensions that were discovered on the computers in your infrastructure. The information is typically used for creating custom software signatures. You can optimize the volume of the monitored data by narrowing down the list of file

extensions that are collected to only those extensions that are typically used for creating signatures. It reduces the workload on the Software Use Analysis server and improves the application performance.

- **Update 9.0.1.2 Incremental import of usage data**
Data concerning software usage is only imported to Software Use Analysis when there are some changes. If usage data is not changed, it is not imported.
- **Update 9.0.1.2 Export of software signatures between instances of Software Use Analysis**
You can export custom signatures from one instance of Software Use Analysis and import them to a different instance of the application.
- **Update 9.0.1.2 Software Knowledge Base Toolkit can be installed at any point in time**
Previously, to avoid problems with reconciling catalog versions, you had to install Software Knowledge Base Toolkit before the first import of the software catalog to Software Use Analysis. Now, upload of a lower version of the catalog to Software Use Analysis is supported. Thus, you can install and start using Software Knowledge Base Toolkit at any point in time without the necessity of reconciling catalog versions. You can also switch back to using the built-in catalog management functionality.
- **Update 9.0.1.2 Improved import logs for serviceability purposes**
Detailed information about import steps and statistics for particular steps is written in the import logs.
- **Update 9.0.1.1 Support for IBM Endpoint Manager version 9.2**
Software Use Analysis can work on IBM Endpoint Manager version 9.2.
- **Update 9.0.1.1 Export of the entire software inventory to a CSV file**
When you create the audit snapshot of the **All IBM Metrics** report, the file `other_products.csv` is created as part of the snapshot. It contains information about software products that are not licensed under full or subcapacity licensing, and that do not generate any PVU or RVU MAPC consumption. It also lists details of the computers on which the software is installed.
Additionally, new columns are added to the following reports and panels:
 - IBM Software Classification
 - **Discovery Start**
 - **Discovery End**
 - **License Type** - the column is also visible when you reassign or share an instance.
 - All IBM Metrics
 - **Imported Part Numbers** - the column is also visible in the CSV report and the audit snapshot.
 - IBM PVU Subcapacity
 - **Imported Part Numbers** - the column is also visible in the CSV report and the audit snapshot.
- **Update 9.0.1.1 Support for Power Systems on Linux**
Endpoint Manager client version 9.2 can be installed on Red Hat Enterprise Linux 5, 6, and 7 as well as SUSE Linux 10 and 11 that run on IBM Power Systems. The endpoints are certified for PVU subcapacity reporting.
- **Update 9.0.1 Deployment Health widget**

New widget on the Software Asset Management Dashboard provides the means to check the status of endpoints that do not meet one or more of the deployment health conditions. You can check whether your endpoints are connecting to the IBM Endpoint Manager server or if the client version is outdated. You are also warned about the most important issues that might occur while endpoints are operating, such as problems with disk space or missing software prerequisites.

- **Update 9.0.1 Integration with SmartCloud Control Desk**

SmartCloud Control Desk is an integrated service management solution that helps you manage a comprehensive range of IT processes, services, and assets. You can use one of its components, Integration Composer, to import your hardware and software inventory together with PVU and RVU consumption data from Software Use Analysis to SmartCloud Control Desk where the data can be further analyzed and processed to create reports and license calculations.

- **Update 9.0.1 Managing the software scan from the user interface**

The software scan collects information about software applications that are installed on your endpoints. It is necessary that the scan runs on regular basis so that the data provided to Software Use Analysis is always up to date. The scan is initiated from the IBM Endpoint Manager console, however after you initiate it, you can modify its schedule and frequency directly from the Software Use Analysis user interface without the need to opening the console.

- **Update 9.0.1 Customized report views**

Reports in Software Use Analysis give you the means to specify report filters that limit the displayed data to particular entries that meet the specified conditions, and report options that allow you to view only the selected columns. Your settings can now be saved and displayed each time a report is being accessed. Such a customized report view can be displayed as a default view for all users or only the one who created it.

- **Update 9.0.1 REST API for exporting and importing customized report views**

Your customized report views can be exported from one Software Use Analysis instance and imported into another through REST API.

- **Update 9.0.1 REST API for retrieving computers, software, and license use**

New REST API functions allow you to quickly retrieve data about your computer systems, software instances installed on those systems, and license consumption that they generate. The retrieved data can be then passed to external applications.

- **Update 9.0.1 Automatic capacity configuration for Linux on z/VM[®]**

The capacity configuration for Linux on z/VM is required to collect complete data from your endpoints and to properly calculate the subcapacity consumption. The required values used to be manually entered in the Create Capacity Configuration for Linux on z/VM task in the IBM Endpoint Manager console. The configuration is now created automatically on System z10[®] E64 (type 2097) mainframes with z/VM 6.3 that supports the Store Hypervisor Information (STHYI) instruction. For those mainframes, the task in the console is not relevant and can be omitted.

- **Update 9.0.1 Setting the home page**

If you use a particular report or page frequently, you can set it as home page so that it is displayed each time you log in to Software Use Analysis.

- **Update 9.0.1 Support for the ISO/IEC 19970-2 standard**

ISO/IEC 19970-2 is an international standard for creating software identification tags. The tags are XML files that are used for discovering and identifying

software. They are delivered with software products and contain product-related information such as the publisher and version. They facilitate software discovery and the overall process of software asset management.

- **Update 9.0.1 CPU pooling on Linux for IBM System z®**

Starting from z/VM 6.3, CPU pooling is supported on Linux for IBM System z.

- **Native catalog export and import**

You can export the IBM software catalog together with your customized content to an XML file and import it to another instance of Software Use Analysis.

- **Catalog version**

The IBM software catalog is versioned according to its release date, so that you can easily check whether a newer version is available.

- **Fixlet® for diagnosing the Common Inventory Technology scanner**

The fixlet runs a diagnosing tool that gathers information about the Common Inventory Technology scanner. You can then send this information to the IBM support if any problems occur.

Features and functions

IBM Endpoint Manager for Software Use Analysis application update 9.0.1.2 provides useful features for managing virtualized environments. It discovers the software that is installed in your infrastructure, helps you to analyze the consumption data, and allows you to generate reports.

Software discovery and identification

Software Use Analysis scans your infrastructure to determine what software is installed on the monitored computers.

Signature discovery and creation

Software Use Analysis discovers signatures that allow products to be detected by inventory and usage scans. The application uses the software catalog to match signatures that are discovered with software products that are installed in your infrastructure. You can create new signatures for software items that are installed on your computers, but are not included in the catalog.

Software usage counting

Software Use Analysis provides you with information about how software that is installed in your enterprise is used.

Reports

Reports contain detailed information about the computers in your infrastructure and the software items that are installed on these computers. By viewing the All IBM Metrics report regularly, you can check whether actual installations reflect the software inventory information and ensure that the capacity values are assigned to products correctly.

Contracts

Contracts are used to manage information about licenses for the software products that are installed in your infrastructure and to track spending that is associated with those licenses. A contract contains information about the cost of acquiring and

maintaining a license, as well as its entitlement start and end dates.

IBM license compliance

Software Use Analysis provides the information that is needed to ensure compliance with processor value units (PVU), and resource value units (RVU MAPC) audit reports licensing terms.

Software Use Analysis calculates the maximum core capacity of the server that is available to the installed IBM software. The application also determines the number of processor value units (PVU) or processor-core entitlements that are required. If you deploy a supported virtualization technology, the tool provides processor core capacity information for the whole virtualized environment.

IBM license usage monitoring

You can generate PVU and RVU MAPC audit reports that show the product and processor value units. The audit reports can be viewed in the application or printed.

Limitations

Review the list of Software Use Analysis limitations.

- **Distributed Server Architecture (DSA) is not supported**

DSA involves setting up multiple IBM Endpoint Manager servers that can replicate data between themselves. If the main server is disabled, the backup one takes over its role ensuring continuous work.

This architecture is not supported because Software Use Analysis can be connected only to one of the replicated servers at a time. When you try to add another server, a message is displayed saying that such database is already connected.

- **Analysis properties exist in multiple sites**

The same analysis property can exist in multiple external sites in the Endpoint Manager console. In such a situation, the data that is gathered by the analysis might come from any of the sites, not necessarily the intended one. To avoid unexpected data, ensure that the analysis that you chose comes from the correct site. In the top navigation bar of Software Use Analysis, click **Management**, and then one of the properties: **Usage Properties**, **Package Properties**, or **UNIX Package Properties**. Click the name of the property, and in the Data Source Property list, check whether the site is correct. The name of the site is displayed below the name of the property. If the site is incorrect, change it.

- **Coexistence of Software Use Analysis version 2.x and application update 9.0.1.2 is not supported**

For more information, see **Coexistence considerations** in the *Installation* section.

- **Coexistence of Software Use Analysis and License Metric Tool is limited**

Due to major constraints, one Endpoint Manager server cannot be shared between Software Use Analysis and License Metric Tool. For more information, see **Coexistence considerations** in the *Installation* section.

- **Contracts for IBM products are not supported**

Due to internal reasons, entries for some products are duplicated in the software catalog. Because of that, creating contracts for IBM products might not work correctly. Moreover, Software Use Analysis supports contracts for

installation-based licenses only. Such licenses are not applicable to the majority of IBM products. Thus, creation of contracts for IBM products is discouraged.

- **Modification of the catalog content for IBM products is limited**

IBM provides catalog content that allows for discovering practically all IBM products. Moreover, because of internal reasons, you might see duplicates of some IBM products. Thus, it is not possible to modify the correct entry. Modifications to IBM products are discouraged.

- **Signature definitions are not available on the Catalog Audit report for custom signatures that were deleted**

When you delete a custom signature from the software catalog, definition of this signature is not available on the Catalog Audit report. The definition section is empty.

- **Catalog search does not work during the import**

- **Automatic update of the software catalog from Software Knowledge Base Toolkit does not work if the SP800-131a standard is enabled in strict mode**

If you are using the SP800-131a cryptographic standard in strict mode, the software catalog cannot be automatically downloaded from Software Knowledge Base Toolkit. To update the catalog, you must manually download it from Software Knowledge Base Toolkit and upload it to Software Use Analysis.

- **SP800-131a compliance cannot be enabled with Microsoft SQL Server**

If the SP800-131a cryptographic standard is enabled in strict mode in Software Use Analysis and the enhanced security mode is configured on the Endpoint Manager and MS SQL servers, the connection between the Software Use Analysis server and the SQL Server cannot be established. This limitation exists because Microsoft SQL Server is not fully compliant with the SP800-131a standard.

- **The time range filter accepts incorrect values**

When you open the All IBM Metrics or IBM PVU Subcapacity report, and want to specify a time range for which the data is to be displayed, you are allowed to specify an alphabetic value instead of a numeric value. The value is accepted and no warning message is displayed. However, the time range that is specified in such a way does not narrow down the report in a correct way. The default value of one day is used instead.

- **Inventory data is not filtered by computer group**

If you open the Computer Groups report for a group that consists of computers with a particular operating system, the Inventory Data section is not properly filtered. For example, when you create a computer group for computers that run on Linux, but some computers in your infrastructure run on Windows, the Inventory Data section contains software items that are installed on both operating systems.

- **On some panels, users with access to selected computer groups can see data for all computers**

On the panels that are connected with reporting license usage and software classification, all users can see all computers in the environment regardless of the computer group that is assigned to them in their user profiles. This data is displayed in such a way because limiting the scope of computers would affect license usage information and would require recalculation of the values that are based on the current user access restrictions.

- **Duplicate agents exist in environments with Virtual Endpoint Manager**

When virtual machines on VMware are managed by Virtual Endpoint Manager, duplicate agents are visible in the Endpoint Manager console and Software Use

Analysis: an agent from the physical computer and an agent from the virtual guest. The actual agent is marked with an icon.

- **Language on the IBM Software Classification panel is set according to the browser settings**

Software Use Analysis can be configured to use either the web browser locale or the locale that is set in the application profile. The Software Classification panel does not support the customized locale setting that is specified in the user profile. It uses only the locale that is specified in the browser settings.

- **Incorrect texts on non-English locale**

Some texts are incorrectly displayed on the user interface or incorrectly formatted if the locale is set to other than English.

- **Export of the Hardware Inventory report to a PDF file takes a lot of time**

When you export the Hardware Inventory report to a PDF file, a lot of time is needed to generate the file. Because report generation is an asynchronous task, you can use Software Use Analysis user interface in the meantime.

- **Time is expressed in seconds in PDF reports**

In reports, some columns contain time values that are expressed in seconds while on the product user interface the values are presented in a more human-readable format.

- **Cannot add a datasource to Endpoint Manager on Linux in version 9.0 GA**

A data source cannot be added when configuring connections. The following message is displayed: Could not determine character set of the BES database. Are you sure the BES server is running an agent? You can work around this problem by manually setting the `_BESClient_DeploymentEncoding_IANAName` property to that of your deployment codepage.

- **After the port number that is used by the server is changed, the URL is not automatically redirected in Internet Explorer**

If you change the port that is used by the Software Use Analysis server when you have SSL enabled and use a self-signed certificate, the application URL is not automatically redirected in Internet Explorer. You can solve the problem by manually substituting the old port number with a new one in the application URL. The problem does not occur if you use a certificate that is signed by a certificate authority.

- **Changing the host name of the Software Use Analysis server is not supported**

Changing the host name of the Software Use Analysis server is not supported and it results in a failing data import. The host name also cannot be changed for IBM Endpoint Manager, since the host name is recorded into your license certificate during the installation. If you want to change the host name either for Software Use Analysis or for IBM Endpoint Manager, the applications must be reinstalled.

Business benefits

IBM Endpoint Manager for Software Use Analysis offers valuable benefits to your enterprise.

By using Software Use Analysis, you will:

Know what you have

Software Use Analysis identifies software, hardware, and license models in your infrastructure.

Manage what you have

By analyzing the consumption data, Software Use Analysis helps you to:

- Understand current license consumption
- Determine a root cause of license compliance discrepancies
- Prioritize and run follow-up actions

Optimize your license consumption

Software Use Analysis helps you to predict and plan future license consumption. Optimize your software consumption periodically to avoid not meeting your contractual licensing conditions.

Reduce your costs

Optimize your license costs by removing unused software. Plan your budget based on the Software Use Analysis reports. You can also avoid making license compliance payments by aligning software assets with legal licenses and ensuring that no unlicensed software is used.

Key concepts

There are several concepts that will help you to understand how the tool works and how to use it effectively.

Processor value unit licenses

A processor value unit (PVU) is a unit of measurement that is used for determining the licensing costs of IBM middleware products. It is based on the types of processors that are deployed on servers where the software is installed. For each product, the licensee must acquire the appropriate number of PVUs for the specific processor on which the software is deployed.

The number of required PVU entitlements is based on the processor technology and the number of processors that are available to the software product. For PVU-based licensing, IBM defines a processor as a processor core on a chip. For example, a dual-core processor chip has two processor cores. PVUs are assigned per core, not per processor. The number of PVUs that are assigned to a processor core is defined in the PVU per core table.

Processor value unit license types

Entitlements can be full capacity or subcapacity. Full capacity is based on the overall number of processor cores on the physical server on which a product is installed. Subcapacity licensing is based only on the PVUs that are available to the licensed software, not on the total PVU count.

The following table shows a simple full capacity license cost for a software product that is installed on a single server.

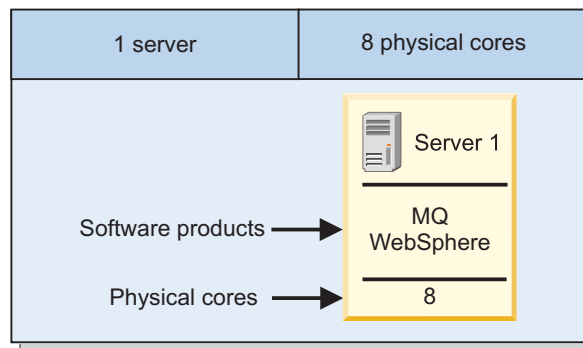
Description	Calculation
Physical processors	8
Cores per processor	2
PVUs per core	100
PVUs per server	$8 \times 2 \times 100 = 1600$
Software cost per PVU	\$50
License cost	$1600 \times \\$50 = \\$80,000$

Subcapacity licensing is based only on the PVUs that are available to the licensed software, not on the total PVU count for a server. Subcapacity licensing is only available in a virtualization technology environment. PVUs are transferable among systems by product within an enterprise.

Subcapacity licensing can significantly reduce licensing costs, especially with a move to a newer processor core technology. Subcapacity licensing requires that a number of requirements are met. For more information, see: Passport Advantage® Virtualization Capacity (Subcapacity) Licensing.

Example of full capacity: physical cores on one server

The license is counted as the highest number of PVUs on the computer where the product is installed. The licensee acquires licenses for all the physical processor cores.

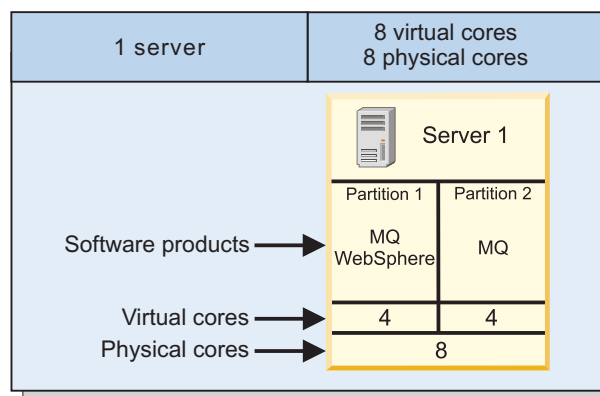


The licensee must acquire licenses for the full capacity in the server available to the software products.

Cores to license	Full capacity
MQ software	8
WebSphere® software	8

Example of subcapacity: virtual cores on one server

The license is counted as the highest number of PVUs that are available for the partition on which the product is installed. The licensee acquires licenses for the lower of virtualization capacity or full capacity of the server. Eligible virtualization technologies can be used to create partitions or virtual machines (VMs) which restrict the processor capacity available.

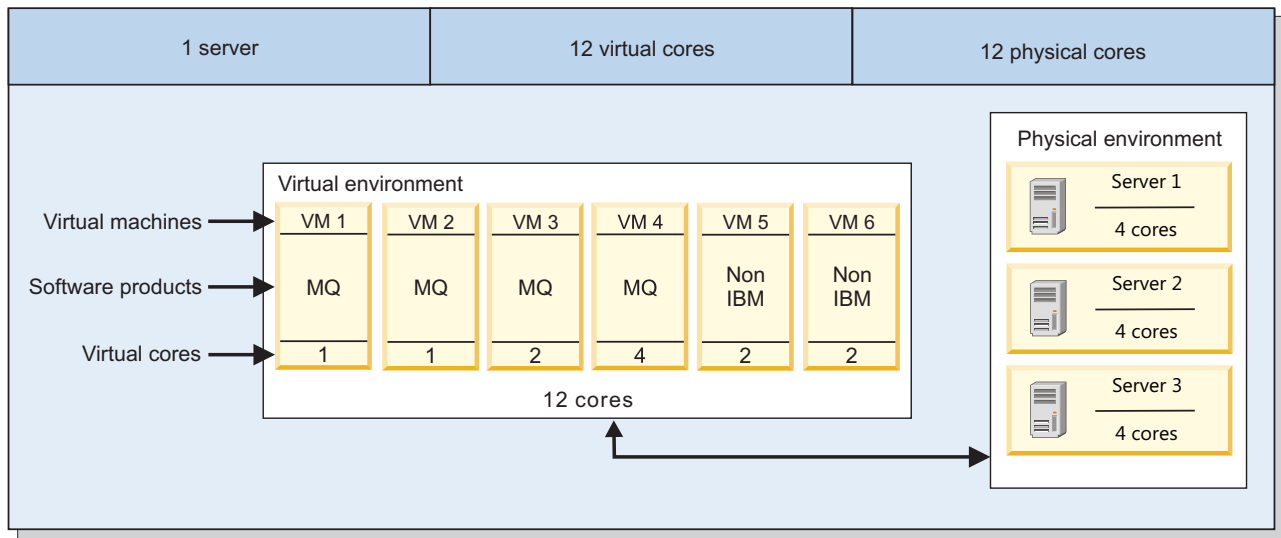


The licensee acquires licenses based on the lower subcapacity in the server.

Cores to license	Partition 1	Partition 2	Subcapacity	Full capacity
MQ software	4	4	8	8
WebSphere Software	4	-	4	8

Example of subcapacity: virtual cores on a server cluster

The license is counted as the highest number of PVUs that are available for the partition on which the product is installed. The licensee acquires licenses for the lower of virtualization capacity or full capacity of the group of servers. Eligible virtualization technologies can be used to create partitions or virtual machines (VMs) which restrict the processor capacity available to a subset of the server or group of servers.



The licensee acquires licenses for the lower subcapacity that is available in the server cluster.

MQ software	VM 1	VM 2	VM 3	VM 4	VM 5	VM 6	Subcapacity	Full capacity
Cores to license	1	1	2	4	-	-		

Processor value unit table

A processor value unit table contains information about the conversion ratio that is used for available processor types.

The processor value unit table supports processor-based pricing models in which charges differ according to the type of processor on which the licensed product is installed and running. In the table, a number of units is assigned to each processor type for which this type of pricing model is available. It is necessary to ensure that the table is updated to reflect the ratios for all processors in use.

If a processor cannot be properly identified by an agent, it is shown in the processor table as OTHER. This case can occur because it is a new processor type and its discovery is not yet supported. In this situation, upload the newest processor value unit table.

Resource value unit licenses (RVU MAPC)

Resource value unit (RVU) licensing provides the means for pricing structures to be based on the number of activated processor cores used or managed by a product (MAPC - manage activated processor cores).

An activated processor core is a processor core that is managed or used by a product, regardless of whether the capacity of the processor core can be limited through virtualization technologies. Entitlements can be full capacity or subcapacity.

Agents retrieve information about the number of processors used to determine the level of license use in terms of resource value units.

Resource value unit license types

The following license types base the calculation of license use on the number of activated processor cores managed or used by products:

- **Full Capacity License**

The license is counted as the highest number of resource value units in the physical hardware environment that are managed by the product.

- **Subcapacity License**

The license is counted as the highest number of resource value units that are available for use by the product. The virtualization capacity license counting rules that define how many activated processor cores must be counted, is at: Virtualization Capacity License Counting Rules.

Resource value unit tier table

The resource value unit (RVU) tier table is used to apply a factor for different numbers of activated processor cores. The greater the number of activated processors, the fewer RVUs are required.

Table 1. The RVU tier table for activated processor cores

Tier	From quantity	To quantity	Factor
1	0	2500	1.00
2	2501	10000	0.80
3	10001	50000	0.60
4	50001	150000	0.40
5	150001	-	0.20

To license in the following core environment:

- Servers with 45,000 activated processor cores

Follow these steps:

1. Begin with the total number of activated processor cores to be purchased for a single RVU-based product.
2. Divide the quantity into tiers based on the RVU tier table. For example, if 15,000 activated processor cores are purchased, tier 1 has 2,500, tier 2 has 7,500, and tier 3 has 5,000.
3. Multiply the number for each tier by the applicable factor to determine the RVUs per tier.

4. Add the RVUs for each tier to determine the total RVUs required for the product.
5. Multiply the total RVUs by the product price.

The following table shows the 45,000 activated processor cores divided into quantity tiers.

Table 2. Activated processor cores divided into quantity tiers

Tier	Number of managed cores	RVUs
1	2500	2500
2	7500	6000
3	35000	21000
Total	45000	29500

The following calculation is used to determine the number of RVUs required to license the 45,000 activated processor cores in the server environment.

- The first tier is used to calculate the first 2,500 activated processor cores at the factor of 1 per core ($2,500 \times 1$).
- The second tier is used to calculate the activated processor cores from 2,501 to 10,000 at the factor of 0.8 per core ($7,500 \times 0.8 = 6,000$).
- The third tier is used for those activated processor cores from 10,001 to 50,000 at the factor of 0.6 per core ($35,000 \times 0.6 = 21,000$).
- Adding the RVUs together for each tier, 29,500 RVUs are required to license the 45,000 activated processor cores.

High-water mark


The high-water mark is the highest peak in a PVU or RVU value for a product over a reporting period. The license consumption for a product can fluctuate depending on the location and configuration of virtual machines in the physical environment. The high-water mark identifies a maximum consumption value that represents the number of PVUs or RVUs needed to license a product for a specified reporting period.

Software Use Analysis monitors your infrastructure and records the high-water mark in PVU and RVU consumption by each product for each virtual machine in your environment. You can compare your report totals to your inventory of license entitlements to confirm your compliance.

In these examples, the reporting period is four weeks. PVU is the license type. DB2 enterprise is the product that is licensed, and the number of PVUs consumed by DB2 are recorded for each week.

Example of high-water mark for a single server with one virtual machine

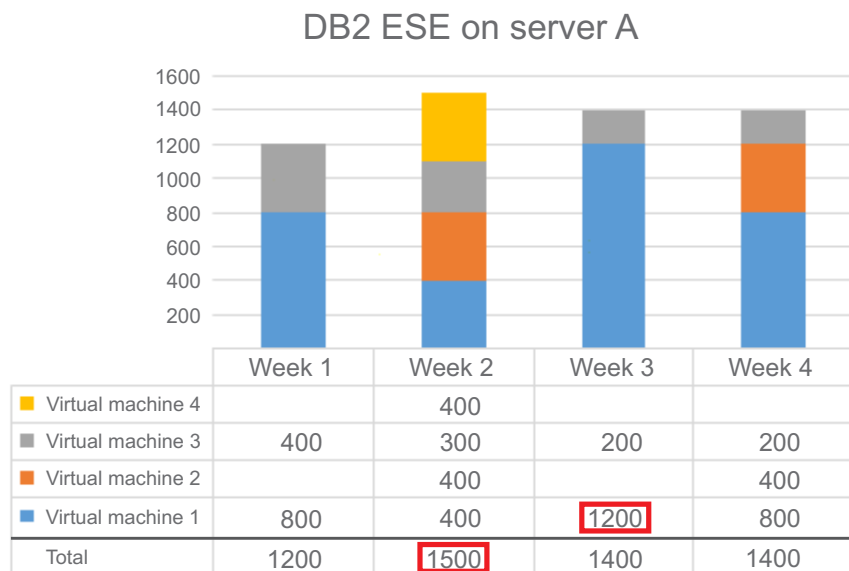
In this example, a single virtual machine has DB2 installed. The configuration of the virtual machine changes over time so that the number of processor cores available to DB2 changes. DB2 consumed 800 PVUs in week one, 400 in week two, 1200 in week three, and 800 in week four. The high-water mark for DB2 in this four-week period occurs in week three. 1200 is the number of PVUs assessed for this 4-week period.



Example of high-water mark for a single server with many virtual machines

In this example, a server has four virtual machines that have different PVU values that are discovered over time. To simplify this example, the virtual machines do not share any processor cores, and no capping is needed. The server level PVU calculation is the sum of all the virtual machine level PVU calculations. The chart and table show the number of PVUs consumed by DB2 for each virtual machine.

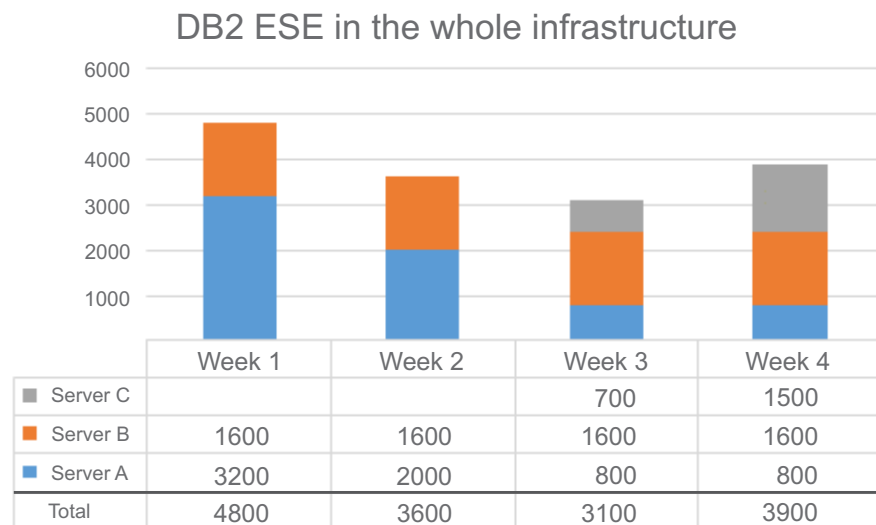
The single highest peak is 1200 PVUs in week three. However, DB2 is installed on more than one virtual machine on the server. To calculate the high-water mark, all of the values from all of the virtual machines must be counted, and not just one. The high-water mark for DB2 in this four-week period occurs in week two. 1500 is the number of PVUs assessed for this four-week period.



Example of high-water mark for the whole infrastructure

In this example, there are three servers with many virtual machines. The infrastructure PVU calculation is the sum of all of the server level calculations, and a server level calculation is the sum of the virtual machine level calculations on that server. The chart and table show the number of PVUs consumed by DB2 for each server.

The high-water mark for DB2 in this four-week period occurs in week one. 4800 is the number of PVUs assessed for this four-week period.



Products, components, and bundles

IBM Endpoint Manager for Software Use Analysis uses catalog based discovery.

The software catalog is a knowledge database that contains information about:

Software product (Offering)

A *software product* is a logical unit of software packaging and sharing that has a managed development and maintenance lifecycle and customer visible attributes. It can be a collection of components, software products whose licensing can be dependent on the licensing of the offering as a whole.

Component

A *component* is a unit of software that can be detected as installed or running on computer systems independently of other software items. It can be a part of a software product, and might be separately identified, but is not individually licensed.

Signature

A *signature* is a set of data that is unique to a given distributed software release. It is used to identify a distributed software release as installed or running on a computer system. There are two types of signatures:

- Inventory recognition signatures detect if a product or component release is installed on the system.
- Usage monitoring signatures detect if a product or component release is running on the system.

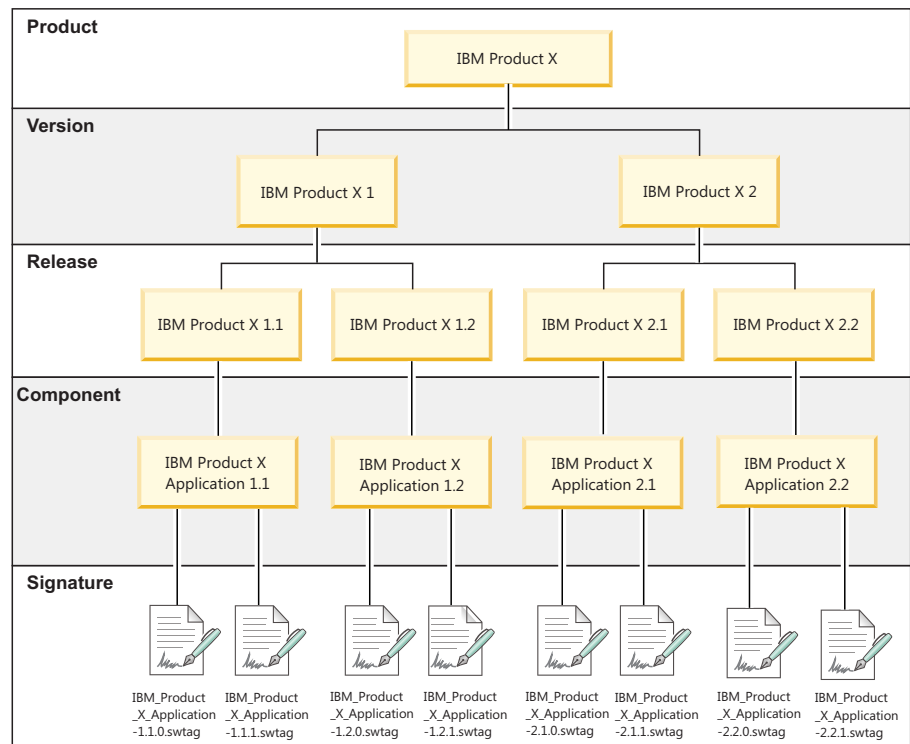
A signature can be a file, registry entry, or another identifier, that acts as a fingerprint for identifying software products. The catalog can be augmented with information about custom or proprietary software products in your IT environment.

Software catalog

A *software catalog* is a collection of software product definitions and data (signatures) that are used to detect software as installed or used in an IT environment.

An initial software catalog is included in the tool. Be sure to update it regularly, especially if you install new IBM software on your network. IBM delivers and maintains the IBM Software Catalog that includes IBM products and a set of non-IBM product definitions.

The following diagram represents the catalog definition for the product IBM Product X. It shows how the various catalog definitions are related and how they provide a structure that supports identification of the software.



The catalog also contains information to support processor value unit (PVU), and resource value unit managed activated processor cores (RVU MAPC) license metrics. This information includes indicating which products use one of these metrics and how the license quantity is calculated.

Hierarchy

A *hierarchy* is a structure that represents software versions, releases, and variations.

Bundle

A *bundle* is set of products that are offered under a single entitlement or license with no dedicated components. In the catalog, a bundle is modeled as a software product with setup relationships to all of the software products that it consists of.

When a file systems is scanned, components are discovered by matching registry information against signatures that are defined in the catalog. A default product is selected for each component. You can change the default bundling.

Main background application tasks

The main goal of Software Use Analysis to generate PVU, and RVU audit reports based on the collected data. All calculations are done in accordance with PVU, and RVU license pricing rules that are described in the official subcapacity licensing documents.

Aggregation

Aggregation is the main calculation task in Software Use Analysis. The aggregation process is a scheduled background task that is run on a daily basis at a particular hour. By default it is performed when the server time is midnight. The task calculates the PVU, and RVU values based on the data that is collected from the agents during software and capacity scans.

Reaggregation

If the initial software bundles are correct after software discovery and rebundling is unnecessary, aggregation is the only calculation process that is required on the side of the Software Use Analysis server and the data that was once calculated is always correct. However, this is rarely the case. You must always modify some parts of the initial bundles that are proposed by Software Use Analysis. You must always confirm which bundles are correct for complex products. A complex product is a product that can be bundled with more than one software offering. After rebundling is complete, the PVU and RVU values that were already calculated must be refreshed. Reaggregation tasks were designed to recalculate or refresh PVU and RVU values that were already calculated. Manual actions that might trigger the data reaggregation include:

- Rebundling a software instance from one product to another
- Confirming the default bundle
- Including software instance in PVU, or RVU calculations
- Excluding software instance from PVU, or RVU calculations
- Sharing an instance

These five actions are very basic operations that the application users perform frequently to adjust the bundling data. In addition to manual actions, refresh of the calculated data can also be triggered by automated bundling.

Aggregation versus reaggregation

The aggregation process was designed to calculate the data from many agents for all products over a short time. In contrast, reaggregation process was designed to perform quick recalculation of PVU, and RVU values for a selected subset of products that were already aggregated. Aggregation of all products from all agents is much quicker (even hundreds of times) than reaggregation of the same amount of data. However, when you must recalculate the PVU and RVU values of only one product, reaggregation should be quicker than aggregation, which cannot recalculate the reporting value of only one product, but must reaggregate it for all discovered products simultaneously.

Inventory builder

Inventory builder is another background task that is executed periodically. During this time, the software inventory is built based on the data from the agent software scans. In other words, this task transforms a list of discovered software components to a list of the discovered software products. In most cases, the initial software bundling of detected components performed by the inventory builder has a very low level of confidence.

Automated bundling

Similarly to aggregation, automated bundling is a periodic background task. It is strictly related to the inventory builder task – when the execution of inventory builder ends, automated bundling starts.

When the automated bundling task runs, it determines the best bundle connections to all unconfirmed product instances. If the newly calculated bundles have a higher level of confidence than the current product bundles, automated bundling rebundles those product instances to the new product with the best match. In the vast majority of cases, subsequent automated bundling runs calculate the same level of confidence for most or even for all of the unconfirmed product instances. However, from time to time, especially in large environments with a large percentage of unconfirmed instances where it can actually happen frequently, the newly calculated confidence level can turn out to be higher than the old one. In this case, the software instance is rebundled. The most common reasons why automatic bundling is able to rebundle some old unconfirmed product instances are:

- Import of a new set of part numbers
- Import of a new software catalog
- After manual rebundling or manual confirmation of one product instance, other unconfirmed product instances can be better bundled by using partition or infrastructure collocation rules
- Detection of a new simple software component (a component that can be assigned to only one product) by agents might also change the calculations for other unconfirmed instances due to partition or infrastructure collocation automated bundling rules.

However, automated bundling does not replace the manual work that must be done to confirm or rebundle all unconfirmed product instances. The confidence level that is calculated by automated bundling is supposed to facilitate manual bundling by providing the best potential bundling options for all unconfirmed product instances.

Extract, Transform, Load (ETL)

In general, Extract, Transform, Load (ETL) is the process in database usage that combines three database functions that aim at transferring data from one database and placing it into another. The first stage, Extract, involves reading and extracting data from various source systems. The second one, Transform, converts the data from its original form into the form that meets the requirements of the target database. The last stage, Load, saves the new data into the target database, thus finishing the process of transferring the data.

In Software Use Analysis, the Extract stage involves extracting data from the IBM Endpoint Manager server. Such data includes information about the infrastructure,

installed agents and detected software. ETL also checks if the new software catalog is available, gathers information about the software scan and files that are present on the endpoints, and collects data from the VM managers.

The extracted data is then transformed to a single format that can be loaded to the Software Use Analysis database. This stage also involves matching scan data with the software catalog, calculating processor value units (PVU), processing the capacity scan, and converting information that is contained in the XML files.

After the data is extracted and transformed, it is loaded to the database and can be used by Software Use Analysis.

Raw data discovery and matched signatures

Raw data is the unprocessed data that can be used to create new signatures in the software catalog for unmatched software items.

Raw data includes a list of files that were found on the system and native installation registry entries. The data is obtained by monitoring process tables of the endpoints in your infrastructure. It can be used to create and diagnose signatures. Signatures are file names, registry entries, and other types of information that are unique to a particular software item and can be used to identify it.

Endpoint Manager clients, also called agents, are installed on every computer that is managed under Endpoint Manager. The clients collect information about the software that is installed on the endpoints in your infrastructure and send the raw data to the Endpoint Manager server. You import this data into the Software Use Analysis server by using the Extract, Transform, and Load (ETL) import process. The Software Use Analysis server matches discovery data with the software catalog and generates the inventory and usage reports. If you view the raw usage data report and find a software item that is not matched by any signature from your current software catalog, you can create a signature for that item that is based on the raw data.

Contracts

A contract contains information about the cost of acquiring and maintaining a license and its entitlement start and end dates. You use contracts to manage information about licenses for the software products that are installed in your infrastructure.

You can assign contracts to specific computer groups to indicate which computers are entitled to use particular software.

Contract usage reports show the number of software instances that you are entitled to, and the actual number of software instances that are used. You can reduce spending on unused licenses, and track computers that use unlicensed instances of a particular software product. You can also track the spending that is associated with software licenses and manage the costs of extending or downsizing licenses.

Scenarios

To understand the business goals that you can achieve with IBM Endpoint Manager for Software Use Analysis, familiarize yourself with the main user scenarios. Each scenario presents a real-life example of using the application and lists personas who are typically involved in achieving a business goal.

Software and hardware discovery

You can use Software Use Analysis to automatically discover hardware and software within your IT infrastructure.

Workflow



Participants who are involved in the scenario:

- Ian - Inventory Administrator

Ian wants to deploy a tool for license consumption monitoring and basic license compliance management. He also wants to create an inventory of all hardware in the enterprise. The following flow describes how he can accomplish his goals.

1. First, Ian obtains a baseline of computers that exist within the IT infrastructure of the enterprise. He can do this by referring to any existing inventory reports.
2. Then, he reviews a list of operating systems on which agents can be installed. He compares this list with operating systems on which the discovered computers are running and installs agents on endpoints with a supported operating system.
3. When the deployment is complete, Ian organizes computers into computer groups to avoid overloading the infrastructures.
4. Next, he also configures different types of scans so that they do not interrupt business workloads of the enterprise. To enable scanning of remote systems and applications, for example VM hypervisors, or cloud management software, he collects IP addresses and access credentials for these systems and applications.
5. Agents then perform an automatic capacity scan on the systems on which they are installed. They collect hardware information about processors that is required for pricing calculations, for example, processor vendors, brands, types, and number of cores.
6. Next, Ian finds out what software exists on the computers and meters its use. Agents automatically perform software inventory discovery. They scan discovered computers and use a software catalog to identify the software that is installed on these computers. They also collect information about the use of software products, including information about the level and duration of the use of a product.
7. Periodically, Ian verifies whether up-to-date artifacts are used for software and hardware discovery. He checks whether the latest approved software catalog or remote discovery plug-ins are used. He also reviews the operational status of the discovery. For example, he checks whether all infrastructure scans are successful. If any artifacts are outdated, or problems with discovery occur, Ian takes or requests a corrective action, such as endpoint troubleshooting, reconfiguration, redeployment; or performs ad hoc scans to gather the missing information.
8. When Ian verifies the operational status of discovery, he proceeds with verification of the quality of the discovered data. For products that are sold as a part of bigger offerings, he ensures that the discovered software is properly associated with its parent products. It is especially important for middleware software that can come with different bundles that influence software licensing. To support decision process, he reviews entitlement reports and existing inventory reports or retrieves details of particular software installations directly from system and application owners. If there are any discrepancies, he fixes the bundle assignment. For ISV and in-house developed tools, Ian periodically reviews endpoint operating system data to verify quality of software detection, and creates new software discovery signatures.

Results

Ian identified the scope of the deployment and deployed Software Use Analysis. He configured different types of scans, and verified discovery status. He also verified the quality of software inventory data. He possesses accurate and up-to-date data about hardware and software in the infrastructure of the enterprise. He can provide this information to Sam, the Software Asset Manager, who is responsible for identifying unlicensed and unauthorized software, license use optimization, and for software contracts negotiation.

License consumption verification and reconciliation

You can use Software Use Analysis to verify license consumption, reconcile it with software entitlements, and to identify and resolve any discrepancies.

Workflow



Participants who are involved in the scenario:

- Sam - Software Asset Manager

Sam wants to verify software license consumption. The following flow describes how he can accomplish his goal.

1. Sam reviews effective license entitlements. He analyzes the consumption data and compares it with license entitlements. He analyzes how license consumption was determined. For example, which software installations or which software and hardware metrics affect particular license consumption. He determines root causes of any discrepancies between license entitlement and consumption.
2. Sam retrieves detailed information from system and application owners for particular software installations to understand why each software product was installed. He takes corrective actions to properly classify software. His actions result in updated license consumption count. For software that is no longer used, Sam requests uninstallation or license transfer. For unlicensed software, he either requests uninstallation, or the purchase of missing licenses.

Results

When these two steps are performed, license compliance baseline is achieved. Sam can now focus on the continuous license compliance process.

Continuous license compliance

You can optimize your license spending and avoid non-compliance issues by periodically analyzing how changes to the IT infrastructure influence license consumption.

Workflow



Participants who are involved in the scenario:

- Sam - Software Asset Manager
- Ian - Inventory Administrator
- Robert - Resource Capacity Administrator

Sam wants to ensure that the most efficient licenses are purchased for products that are used in the enterprise. The following flow describes how he can accomplish his goal.

1. Ian analyzes changes in the IT infrastructure to determine discovery completeness. He checks whether there are any new systems provisioned or whether any systems are decommissioned. Having this information, Ian reviews changes to the discovered data and discovery statuses, and determines corrective actions to ensure full infrastructure coverage.
2. Periodically, Sam reviews changes to license consumption and how they were affected by changes in the infrastructure. If there are any discrepancies between license consumption and entitlement, Sam takes corrective actions.
3. Based on the license consumption trends and software usage view, Sam helps in budget planning and optimizes license cost by removing or transferring unused software. He also performs what-if analysis and compares available license models to identify most cost-effective schema of software licensing.
4. Sam periodically meets with Robert, the Resource Capacity Administrator, to analyze the impact of infrastructure on license cost and recommends infrastructure changes. He also analyzes planned infrastructure changes to assess their impact on license compliance, and to eliminate any potential non-compliance issues.

Results

By periodically reviewing changes in the IT infrastructure of the enterprise and how they affect license consumption, Sam can optimize license cost. He can request the removal or transfer of unused software, and identify the most effective licensing models for future purchases.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 79758 U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this information and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com)[®] are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.



Printed in USA