

Getting Started with Capacity Planner

Capacity Planner 2.7

0.0

EN-000077-00



You can find the most up-to-date technical documentation on the VMware Web site at:

<http://www.vmware.com/support/>

The VMware Web site also provides the latest product updates.

If you have comments about this documentation, submit your feedback to:

docfeedback@vmware.com

© 2009 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at <http://www.vmware.com/go/patents>.

VMware, the VMware "boxes" logo and design, Virtual SMP, and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

VMware, Inc.
3401 Hillview Ave.
Palo Alto, CA 94304
www.vmware.com

Contents

- 1 Capacity Planner Overview 5**
 - Component Overview 5
 - System Requirements 7
 - Assessment Workflow 10
 - Data Manager Tasks 13
 - Dashboard Tasks 14

- 2 Application Analysis 17**
 - Software Profiles 18
 - Base Images 18
 - Sizing Virtual Machine Templates 18

- 3 What's Next 21**
 - Glossary 23
 - Index 25

Capacity Planner Overview

Capacity Planner is a capacity planning tool that provides integrated analysis, planning, and decision support functionality.

Using Capacity Planner, you can assess your capacity planning needs, consider virtualization, and run consolidation assessments. These assessments help to virtualize and consolidate data center infrastructure, redeploy strategic IT assets, and optimize workload capacity utilization.

Capacity planning includes the following tasks.

- Balancing workloads on existing computers
- Consolidating work so that fewer computers are needed to perform tasks
- Identifying outdated computers for retirement
- Planning for the purchase of new and more efficient hardware
- Introducing virtual machines to assume some of the work

This chapter includes the following topics:

- [“Component Overview,”](#) on page 5
- [“System Requirements,”](#) on page 7
- [“Assessment Workflow,”](#) on page 10
- [“Data Manager Tasks,”](#) on page 13
- [“Dashboard Tasks,”](#) on page 14

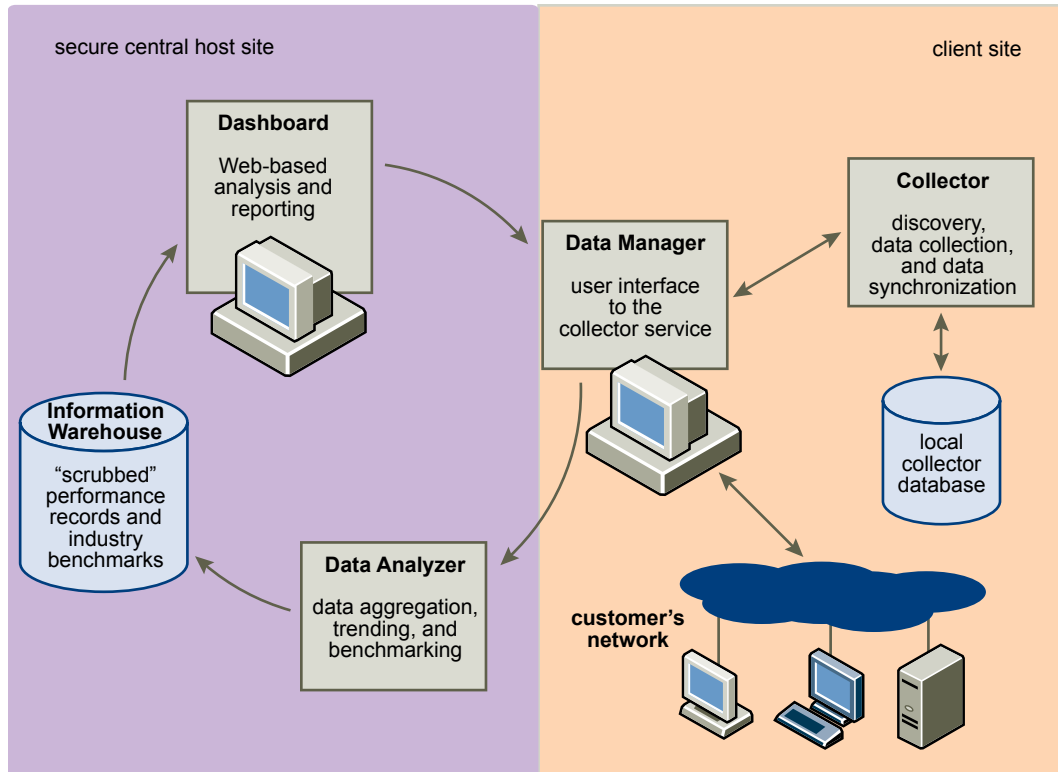
Component Overview

The full Capacity Planner architecture includes the Collector, the Data Analyzer, the Data Manager, and the Information Warehouse.

The Collector runs as a Windows service and performs its jobs in the background. Data that the Collector discovers and gathers is stored in a local database and uploads to a secure central host site.

Typically on an hourly basis, data is uploaded from the local collector database to the Capacity Planner Dashboard.

[Figure 1-1](#) illustrates how these components work together.

Figure 1-1. Capacity Planner architecture

Data Manager

The Data Manager is the user interface for the Data Collector and is installed with the Collector.

The Data Manager provides an organized view of the collected information and is the administrative control for the Collector.

The Data Manager configures the Collector components to perform the following functions.

- Set up and maintain job schedules.
- Set user ID and password combinations for access to target systems.
- Manually execute jobs.
- View job progress, especially collection progress.
- Monitor execution messages written out by the Collector.

The Data Manager also performs the following functions.

- Manages the process by which collected data is sent to the Information Warehouse .
- Provides detailed and summary views and reports on all discovered objects, collected inventory information, and monitored performance data.
- Allows you to start and stop the Collector.

The Collector inventories the existing computers in an IT environment, including the type of computer, operating system, applications or services running on the computer, and so on. Then the Collectors gathers hundreds of performance statistics from each system, such as CPU busy time, CPU idle time, and memory utilization.

Dashboard

The Dashboard is a web-based application that allows you to interact with the discovered inventory and performance data collected from the customer site during the assessment.

You can also view the status of ongoing assessments in the My Assessments table of the Dashboard Home page. The Dashboard resides on a secure central host site along with the Data Analyzer and Information Warehouse.

Information Warehouse

The Information Warehouse is the repository for data that the Collector gathered from all deployed collectors.

The Capacity Planner Information Warehouse contains a growing set of industry reference data that you can use for comparative analysis and benchmarking. You can use this information to guide validated server consolidation and capacity optimization decisions for the enterprise.

The stored metrics that serve as industry benchmarks are used by the Data Analyzer to perform data analysis. Information in the Warehouse is anonymous. None of the information used in the research features can be traced back to a specific customer, which is how Capacity Planner ensures data privacy. Data collected from one organization cannot be viewed by another organization.

Data Analyzer

The Data Analyzer is the component that performs the evaluation required for capacity planning, such as comparing results collected for an organization to industry benchmarks and identifying trends.

After analysis, the analyzed data moves to a secure database called the Information Warehouse.

System Requirements

There are requirements for host and target systems.

Capacity Planner Host System Requirements

The host system can be a physical desktop, server, or virtual machine and must meet certain system requirements.

[Table 1-1](#) describes the requirements for a Capacity Planner host system.

Table 1-1. Host System Requirements

Category	Requirement	Notes
Operating system	Windows 2000 (Server or Professional) SP3, Windows XP Pro, or Windows Server 2003	The English language version is required. If the Collector is installed on Windows XP SP2, the Windows firewall must be turned off. If the collector is installed on Windows Server 2003, and collecting from Windows 2000 target systems, use the Windows 2003 R2 SP2 only.
WMI or remote registry	WMI (Windows Management Instrumentation) or remote registry must be installed	
CPU	Minimum 1.5GHz CPU	Runs on 32-bit and 64-bit, hosts in 32-bit mode
Memory	Minimum 1GB RAM	
Local Disk	2GB free disk space	

Table 1-1. Host System Requirements (Continued)

Category	Requirement	Notes
Network Connectivity	Connection to the Internet	Enables connection to the Capacity Planner Dashboard
Network Bandwidth	20,000 bytes per second during data collection	
Browser	Microsoft Internet Explorer 5.5 or higher	
Virus scanning	Some exceptions might be required for the duration of the assessment	

Authentication Requirements

To install and run the Data Manager service, the consultant must have an administrator-level user account for each host. This account must have the privilege of logging on as a service that is configured to start automatically.

Target System Requirements

You must make sure that the target systems meet the necessary criteria for collection.

Before collection starts, you must have the answers to the following questions.

NOTE You must know whether any Windows NT target systems are within the scope of the assessment. These systems require a different setup to enable data collection.

- What is the number of target systems to be included in the scope of the assessment?
- Does the customer want to identify target systems by using the discovery method or by importing a CSV file?
- How many of the systems are servers?
- How many of the systems are workstations?
- How many of the systems are Windows?
- Of the systems that are Windows, how many are NT systems?
- How many of the systems are Linux or UNIX?
- How many of the systems are targeted for inventory only?
- How many systems are targeted for a full analysis, including discovery, inventory, and collection of performance metrics?
- Are any of the target systems located behind a firewall in a perimeter network (DMZ) and, if so, how many?

All target domains and systems must be viewable in the network viewing area on the host. [Table 1-2](#) lists the requirements that must be met for each target system.

Table 1-2. Collection Requirements for Target Systems

For Microsoft Windows Target Systems	For Linux or UNIX Target Systems	For all Target Systems
Make sure that Windows Management Instrumentation (WMI) is enabled.	Make sure that port 22 is open.	
Enable Remote Registry.	Make sure that the sshd daemon is running.	Verify network bandwidth availability

Table 1-2. Collection Requirements for Target Systems (Continued)

For Microsoft Windows Target Systems	For Linux or UNIX Target Systems	For all Target Systems
Enable Performance Monitor (Perfmon).		Verify that the Collector can access the target systems with sufficient privileges.
Enable file and print services.		
On Windows NT target systems, install and enable WMI and Diskperf service.		

Supported Target Systems

Capacity Planner 2.7 can collect data from the following target systems.

Windows Systems	Linux/UNIX systems
Windows NT 4.0 Server	HP-UX 10.xx (PA-RISC)
Windows NT 4.0 Professional Workstation	HP-UX 11 (PA-RISC)
Windows 2000 Server/Advanced Server/Datacenter	HP-UX 11.11 (PA-RISC)
Windows 2000 Server (64-bit Itanium)	HP-UX 11.22 (PA-RISC)
Windows 2000 Professional Workstation	HP-UX 11.23 (Itanium)
Windows XP Professional	Sun Solaris 7 (SPARC)
Windows XP Professional (EM64T, AMD 64)	Sun Solaris 8 (SPARC)
Windows 2003 Server	Sun Solaris 9 (SPARC)
Windows 2003 Server (64-bit Itanium)	Sun Solaris 9 (x86)
Windows 2003 Server (EM64T, AMD 64)	Sun Solaris 10 (SPARC)
Windows Vista Enterprise	Sun Solaris 10 (x86)
Windows Vista Business	SUSE Linux Enterprise Server 9
Windows Vista Ultimate	SUSE Linux 10
Windows 2008 Server	SUSE Linux 9
	SUSE Linux 8
	Red Hat Linux 8
	Red Hat Linux 9
	Red Hat Enterprise Linux (ES/AS/WS) 3
	Red Hat Enterprise Linux (ES/AS/WS) 4
	AIX 5.1
	AIX 5.2
	AIX 5.3

Target System Port Availability Requirements

The collector host must be able to connect to all target systems using specific protocols and ports.

[Table 1-3](#) describes port usage. All ports use TCP/UDP protocol.

Table 1-3. Port Availability Requirements for Target Systems

Port	Service	Description	MS Windows Services Using This Port
22	Secure Shell (SSH)	Used for secure logins, file transfers, and port forwarding	
135	loc-srv/epmap	Microsoft Data Communications Exchange (DCE) Locator, also known as End-point Mapper	DHCP Server, DNS Server, WINS Server
137	netbios-ns (NetBIOS names service)	Firewall administrators frequently see large numbers of incoming packets to this port because as users behind firewalls browse Windows-based Web sites, those servers frequently respond with NetBIOS look-ups.	DNS Server, WINS Server
138	netbios-dgm (NetBIOS datagram)	Port 138 is used primarily by the Server Message Block (SMB) browser service that obtains Network Neighborhood information.	
139	netbios-ssn (NetBIOS session)	Windows file and printer sharing	
445	DNS (Domain Name Service (DNS) Direct Hosting port)	In Windows 2000 and XP, redirector and server components support direct hosting for communicating with other computers running Windows 2000 or XP.	Active Directory

Target System Authentication Requirements

Account credentials are required on all target systems that are to be included in the assessment data collection.

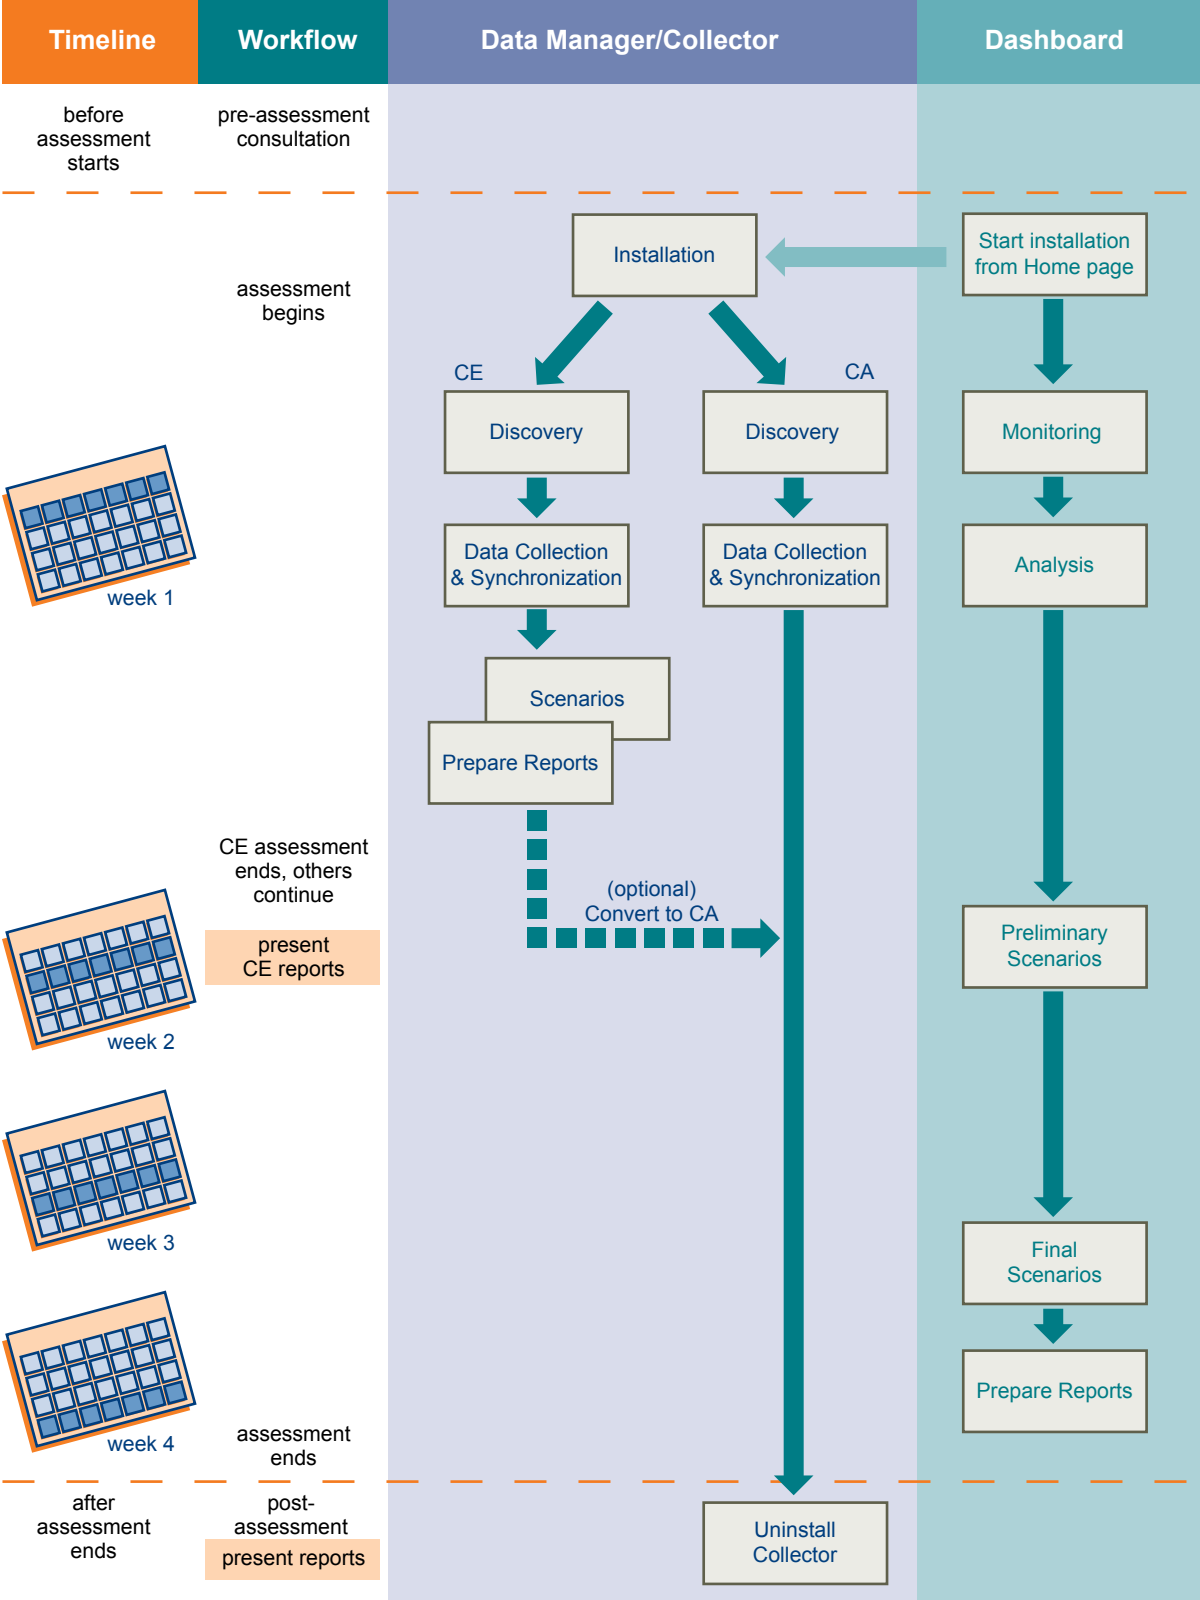
The onsite system administrator can provide a global connection account that has local administrator rights on all the target systems that are to be analyzed. If this cannot be arranged for security reasons, the onsite system administrator can set accounts (credentials) in the Data Manager. User account information is stored in the local Collector database, not in the Information Warehouse.

Assessment Workflow

The assessment timeline includes managing tasks from the Data Manager, Collector, and Dashboard.

The tasks shown in [Figure 1-2](#) illustrate the workflow during an engagement. Work performed by the Data Analyzer and Information Warehouse is not shown.

Figure 1-2. Engagement Timeline and Workflow



Types of Assessments

There are two types of assessments with different goals.

Assessment Type	Description	Assessment Performed By	Results
Consolidation Estimate (CE)	Estimate of the potential savings from implementing an assessment. Only key inventory and performance metrics are collected over a period of several days.	Sales engineers, partners or consultants who have not attended Capacity Planner training as well as trained partners and consultants.	<ul style="list-style-type: none"> ■ Two consolidation and virtualization scenarios (one aggressive scenario, one conservative scenario) ■ Standard interim and final reports
Consolidation Analysis (CA)	A thorough study that collects inventory and performance metrics over a month or longer, taking into account monthly fluctuations. Provides a detailed analysis of potential savings available through consolidation and virtualization.	Only trained partners and consultants.	<ul style="list-style-type: none"> ■ Unlimited consolidation and virtualization scenarios ■ Standard interim and final reports as well as detailed analytical reports, such as forecasting, alerts, and anomaly detection

Starting an Assessment

An assessment begins with a consultation to gather and convey information.

Different types of information should be gathered before an assessment begins to inform the direction the assessment takes.

- Determine the organization's goals for performing the assessment and decide what type of assessment would meet those goals.
- Collect information about computers in the organization's environment, such as number, location, and type. Computers being assessed are referred to as target systems and can be Microsoft Windows, Linux, or UNIX computers (either servers or workstations).
- Scope the assessment to determine whether to evaluate all target systems in the organization or only a subset.
- Understand the organization's network and security infrastructure. Consultants performing an assessment need access to target systems to collect inventory and performance data.
- Discuss authentication, network connectivity, and security issues that might need to be resolved before data collection.
- Exchange contact information (telephone and email address) for all members of the assessment team.
- Decide on a host system on which to install the software to gather and upload the data.

Post-Assessment Tasks

Post-assessment tasks include presenting final reports and proposals, conducting any follow-up tasks that were previously agreed upon, and uninstalling the Data Manager from the host computer.

After the assessment has concluded, you can present the results of the data analysis along with a proposal for consolidation and virtualization.

You can uninstall the Data Manager from the host computer or deactivate the database ID. Either action stops data from being collected and uploaded to the Information Warehouse.

Data Manager Tasks

After the Data Manager has been installed, the next steps are identifying systems and data collection.

Identifying Systems

You can identify systems to collect data from either by discovery or by importing a .csv list.

An organization might want to evaluate all target systems in the organization or a subset. Target systems can be imported into the assessment by using one of two methods.

If connectivity issues are found, they can be corrected at this time by using Test Connection in the Data Manager.

You should specify a schedule for additional discovery jobs to take place during the assessment to find target systems that were not in service but are restored to service during the assessment period.

The Data Manager includes a navigation pane with a navigation tree that shows which target systems were imported and how the target systems might be grouped by domain, department, physical location, and so on. Review the target systems displayed in the navigation pane to verify that the systems shown match the scope of the assessment agreed upon in the pre-assessment consultation.

Discovering Systems

You can start the discovery process in the Data Manager.

This method has the advantage of helping organizations locate target systems that might be forgotten.

After Capacity Planner is installed and configured, you should confirm whether the target systems that an organization wants to evaluate have been discovered by the Collector. The discovery process validates that all target systems have been found.

Importing Systems

You can import a list of systems that you want to collect data from.

The target systems can be imported from a list provided in a comma-separated values (CSV) file. If this list is the exclusive source of systems to be assessed, then computers in the network that are not on the import list are not evaluated during the assessment. The list is typically agreed upon during the preassessment consultation.

Data Collection

Data collection focuses on two types of information: inventory and performance.

Inventory Data Collection

Inventory data collection finds hardware and software information for each target system, such as information about a system's CPU, memory, physical disks, network interface cards (NICs), services or daemons, applications, shares, and file systems. Inventory is attempted every day for systems that have not been inventoried.

Performance Data Collection

Performance data collection is a sampling of statistical information collected from each target system. The sampling is done using hundreds of performance counters, such as % Idle Time and % Processor Time. The counters provide average, mean, minimum, and maximum performance statistics. Performance collection occurs every hour.

Excluding Information from Collection

If you want to exclude some categories of information from inventory or some counters from performance collection, you can configure the Collector to do so. Excluded information is not evaluated by the Data Analyzer or uploaded to the Information Warehouse.

Typical Assessment

The typical assessment lasts four weeks to accommodate monthly fluctuations in activity, such as a month-end billing cycle. The exception is a CE assessment. Because a CE is only an estimate, collection is typically done for just a few days.

After installation and the discovery phase is complete, the Data Manager is left running on the host at the customer site and continues to collect inventory and performance data. You can monitor the activity remotely from the Capacity Planner Dashboard and make corrections as needed.

Capacity Planner collects data by using the following utilities that are provided with the target system's own operating system.

- Windows Management Instrumentation (WMI)
- Remote Registry
- Remote Procedure Call (RPC)
- Perfmon
- Secure Shell (SSH)

Dashboard Tasks

After setup is complete and data collection progresses, you can review the results, create consolidation scenarios to model potential changes, and create reports from the Dashboard.

You can perform a number of tasks in the Dashboard after the first week of data collection.

Monitoring

You can check the progress of your assessment in the My Assessments table.

You can check a summary of the assessment progress in the Dashboard My Assessments table. The Collection Dashboard provides the view of the inventory and performance success rate week by week. It displays the data collected in past twelve weeks. If any problems occur, such as the Collector being stopped, a notification is sent to you for action.

Analysis

After a week of collection, you can view the results of the data collection in the Dashboard.

Several data views in the Dashboard that provide this information.

Forecast Critical Processors	A list of target systems and the number of weeks until each target system reaches a critical usage point (greater than 50% of capacity).
System Processor Load View	A summary of how each server or workstation is used during standard business hours compared to use during the peak hour (the highest utilization in a 24-hour period).
Anomalies Summary	A summary of target systems with performance data that deviates from industry average values.
Vintage Systems	A view of target systems operating below a minimum performance threshold.
Base Image Manager	Allows you to work with templates based on software profiles with an approved operating system that are used to create a consolidation profile utilizing new virtual machines.

Modeling

Using data collected after three to four weeks, you can model consolidation and virtualization options.

As the assessment enters the third or fourth week of data collection, you can prepare preliminary consolidation and virtualization models called scenarios. By trying various scenarios, you can determine more efficient and cost-effective ways to structure the IT setup.

When preparing scenarios, consider several approaches. One scenario might suggest more aggressive recommendations and another a more moderate approach. A scenario might include recommendations for purchase of new hardware to increase efficiency or the introduction of virtualization technology on existing target systems to handle some of the workload.

Reporting

Preparing reports is a final step in a capacity planning engagement.

As the assessment period ends, you can run final scenarios and prepare reports for presentation of the assessment findings.

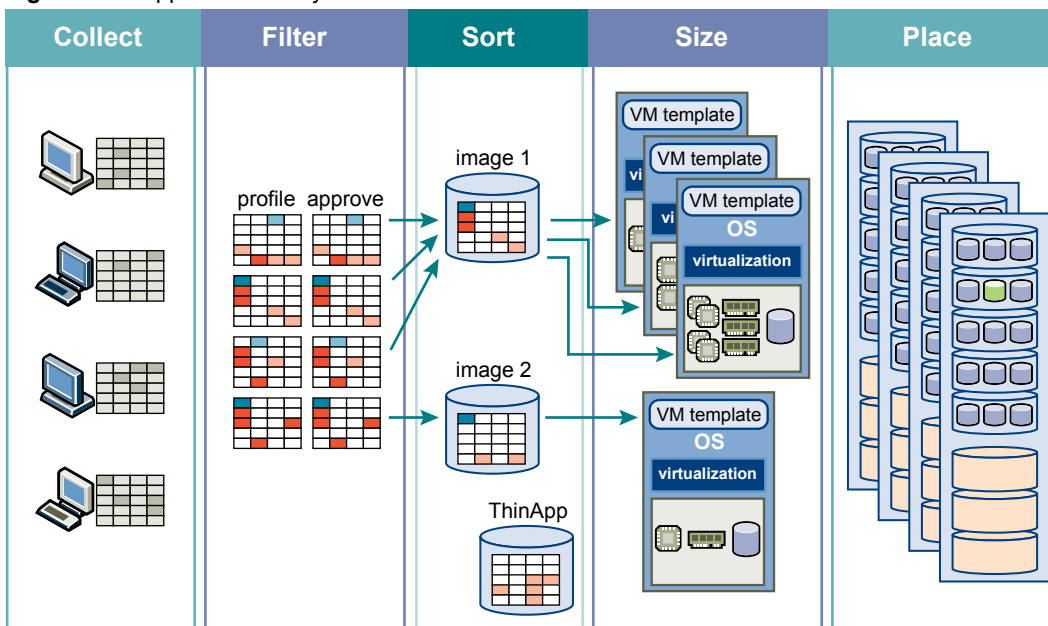
Application Analysis

2

Application analysis allows you to analyze applications and select applications for software profile management and base image management.

The application analysis workflow has several stages. Desktop applications are analyzed to determine what common applications are used by different users. From this analysis, software profiles are created. From these, base images are created. The steps in this process are collecting, filtering, sorting information into base images, sizing the virtual machines, and then placing them.

Figure 2-1. Application Analysis Process



This chapter includes the following topics:

- [“Software Profiles,”](#) on page 18
- [“Base Images,”](#) on page 18
- [“Sizing Virtual Machine Templates,”](#) on page 18

Software Profiles

Software profiles are used to analyze application usage and distill it so that base images can be created based on user patterns.

Software profiling allows Capacity Planner to standardize the use of applications and operating systems. Profiles can be linked to indicate that standardization should be done if an application appears. You can determine what applications or operating systems to use to create profiles.

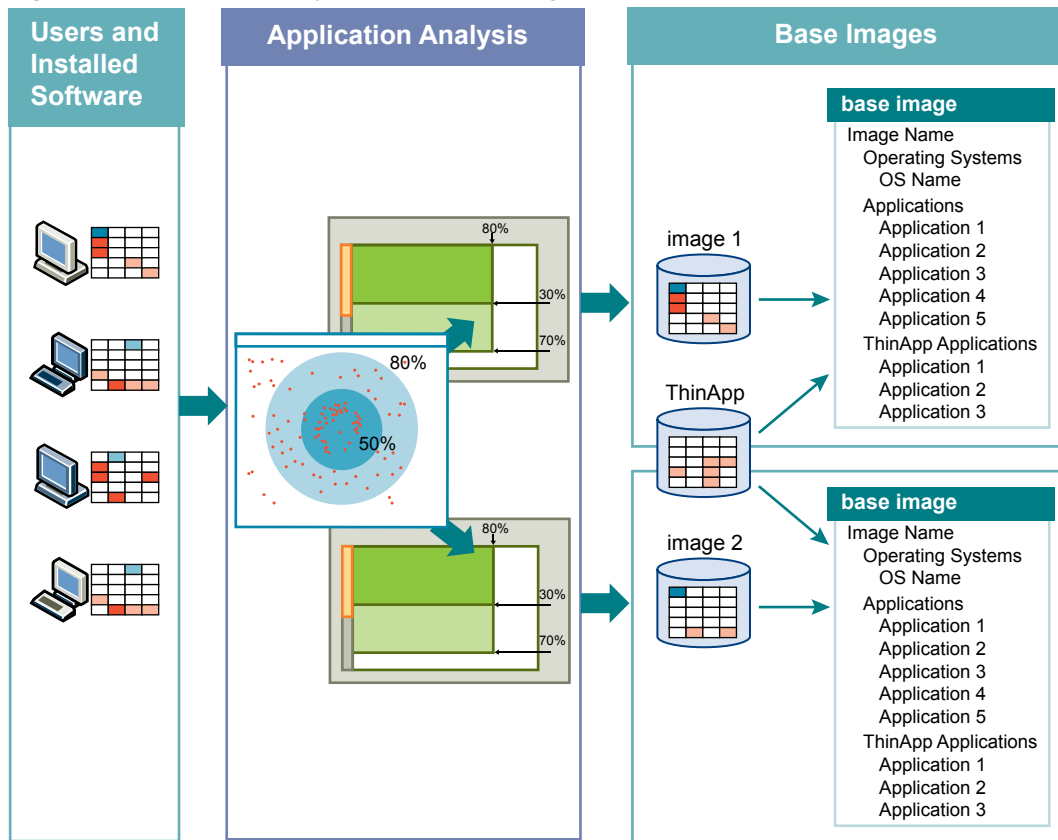
Base Images

Application inventory analysis creates base images and ThinApp images.

ThinApp images contain applications that are approved virtualization candidates. ThinApp images are used to virtualize application without virtualizing entire systems.

Base images contain software profiles, ThinApp applications, and an operating system.

Figure 2-2. Application Analysis Assists Base Image Creation

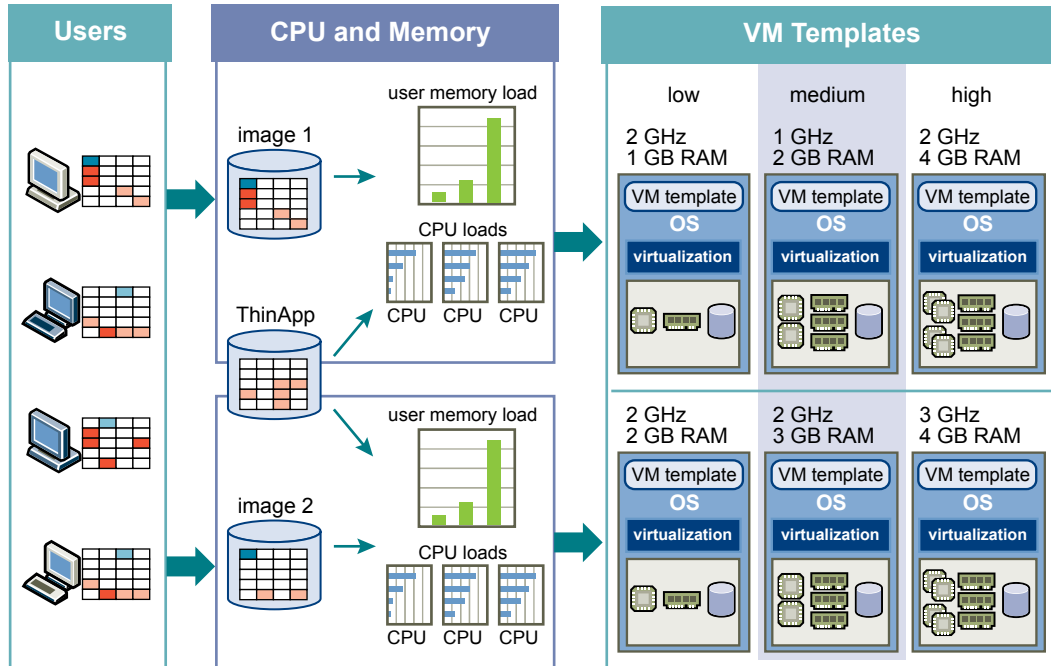


Sizing Virtual Machine Templates

Base images are used to create virtual machine templates, which can be sized.

A virtual machine template contains an operating system, software, and virtual machine sizing parameters. Each virtual machine is sized, based on usage, into three templates: low, medium, and high.

Figure 2-3. Creating Virtual Machine Templates Workflow



What's Next

Now that you have become familiar with the components and basic tasks, you can do several things.

Table 3-1. What To Do Next

Task	Where to find it
Create your account in the Dashboard	http://optimize.vmware.com
Start a new assessment	Dashboard > Home > Start new assessment
Set up the company to be assessed	Dashboard > Administration > Company Setup
Download and install the Data Manager	<i>Capacity Planner 2.7 Installation Guide</i>
Add accounts for target systems	Data Manager > Objects > <system>
Connect to the target systems and test	Data Manager > Home > Test Collection
Register the Collector	Data Manager > Home > Register Collector
Register a database ID	Dashboard > Assessment > Register Database IDs
Synchronize the data upload from the collectors to the Dashboard and test the synchronization	Data Manager > Home > Synchronize Data
Create hardware templates	Dashboard > Assessment > Assessment Tools > Hardware Templates
Create software profile templates	Dashboard > Assessment > Assessment Tools > Software Profile Templates
Create a new user	Dashboard > Administration > Users
Set up notifications	Dashboard > Administration > Notifications
Run reports	Dashboard > Reports

Glossary

Anomalies Summary

A summary of target systems with performance data that deviates from industry average values.

anomaly

In Capacity Planner, an indicator that server performance is significantly different from the industry performance averages of like servers. Capacity Planner notes an anomaly when it detects performance that is more than three standard deviations from the industry average.

authorization role

A set of privileges grouped for convenient identification under names such as Administrator.

Base Image

A template based on application analysis with an approved operating system that is used to create a consolidation profile using new virtual machines.

Capacity Planner

A capacity planning tool that provides integrated analysis, planning, and decision support.

Consolidation Assessment

A month-long analysis of performance and inventory data collected from a company.

Data Manager

The user interface that manages the Collector.

discovery

The process of locating servers and workstations within domains.

EULA (end user license agreement)

The software license that details any restrictions placed on users.

Forecast Critical Processors

A list of target systems and the number of weeks until each target system reaches a critical usage point, which is greater than 50% of capacity.

group

A set of users assigned a common set of privileges.

Information Warehouse

A repository for collected data and a set of industry reference standards.

NIC (network interface card)

An expansion board that provides a dedicated connection between a computer and a network. Also called a network adapter.

Perfmon

A tool that enables user-level applications to collect and access performance statistics. Some form of performance monitoring is available on all Windows, Linux, and UNIX platforms, although the specific information collected and made available varies.

performance counter

Information collected about an entity (such as a host or virtual machine). Counter information includes the group to which the counter belongs (for example, memory), counter ID, counter name, key, statistic type, rollup type, and the kind of entity for which the counter is reported. CPU usage is an example of a performance counter.

read-only user

A role in which the user is allowed to view the inventory but not allowed to perform any tasks.

vintage server

A server that does not meet minimum CPU speed requirements, as defined in the Dashboard consolidation scenarios by the Information Warehouse Administrator, company Administrator, or company Power User.

WMI (Windows Management Instrumentation)

An API in Microsoft Windows that enables administrators or management software to monitor, configure, and control system and network resources.

Index

A

analysis 15
anomalies 15
application analysis 17
applications 18
assessment 15
assessment workflow 10
assessments 12
authentication 10

B

base image 18
base image manager 15
base images 17, 18
basic tasks 21
benchmarks 7

C

capacity planning 5
collection dashboard 14
collector 5, 6
Collector 7
collectors 7
configure tasks 6
connectivity 13
consolidation 15
consolidation analysis 12
consolidation estimate 12
consolidation scenarios 14
CPU 7
critical processors 15
CSV 8, 13

D

Dashboard 5, 7, 14
Data Analyzer 5, 7
Data Manager 5, 6, 13
Data Manager tasks 13
discovery 13
DNS 9

F

forecast critical processors 15

G

global connection account 10

I

images 18
importing systems 13
Information Warehouse 5, 7
information, exclude 13
inventory data 13
IT environment 6

L

Linux 13

M

MDAC 7
memory 7
metrics, stored 7
modeling 15
monitoring 14
My Assessments table 14

N

navigation pane 13
navigation tree 13
NetBIOS 9
network connectivity 7

O

operating system 7

P

Perfmon 13
PerfMon 8
performance data 13
ports 9
post-assessment 12
pre-assessment consultation 13
printer sharing 9
processor load 15

R

remote registry 7
Remote Registry 13
reports 14, 15

repository 7
RPC 13

S

scenarios 15
secure shell session 13
sizing 18
software profiles 18
SSH 9, 13
standardization 18
system requirements 7

T

TCP/UDP protocol 9
template,virtual machine 18

ThinApp images 18
trends 7
typical assessment 13

U

UNIX 13

V

vintage systems 15
virtual machine 18
virtual machine template 18
virtualization models 15

W

WMI 7, 8, 13