Chapter 12

Monitoring VMware vSphere Performance

The monitoring of VMware vSphere should be a combination of proactive benchmarking and reactive alarm-based actions. vCenter Server provides both methods to help the administrator keep tabs on each of the virtual machines and hosts as well as the hierarchical objects in the inventory. Using both methods ensures that the administrator is not caught unaware of performance issues or lack of capacity.

vCenter Server provides some exciting new features for monitoring your virtual machines and hosts, such as expanded performance views and charts, and it greatly expands the number and types of alarms available by default. Together, these features make it much easier to manage and monitor VMware vSphere performance.

In this chapter, you will learn to:

- Use alarms for proactive monitoring
- Work with performance graphs
- Gather performance information using command-line tools
- Monitor CPU, memory, network, and disk usage by both ESX/ESXi hosts and virtual machines

Overview of Performance Monitoring

Monitoring performance is a key part of every vSphere administrator's job. Fortunately, vCenter Server provides a number of ways to get insight into the behavior of the vSphere environment and the virtual machines running within that environment.

The first tool vCenter Server provides is its alarms mechanism. Alarms can be attached to just about any object within vCenter Server and provide an ideal way to proactively alert the vSphere administrator about potential performance concerns or resource usage. I'll discuss alarms in greater detail later in this chapter in the section ''Using Alarms.''

Another tool that vCenter Server provides is the Resources pane on the Summary tab of both ESX/ESXi hosts and virtual machines. This Resources pane provides quick ''at-a-glance'' information on resource usage. This information can be useful as a quick barometer of performance, but for more detailed performance information you will have to search elsewhere—either elsewhere within vCenter Server, as I'll describe later in this chapter, or within the guest operating system itself. Because this tool provides only limited information, I won't discuss it further in this chapter.

Another tool that provides ''at-a-glance'' performance summary is the Virtual Machines tab, found on vCenter Server objects, datacenter objects, cluster objects, and ESX/ESXi hosts. Figure 12.1 shows the Virtual Machines tab of a cluster object. This tab provides an overview of

general performance and resource usage. This information includes CPU utilization, memory usage, and storage space utilized. As with the Resources pane, this information can be useful, but it is quite limited, so I won't discuss it any further in this chapter. However, keep in mind that a quick trip here might help you quickly isolate the one virtual machine that could be causing performance issues for the ESX/ESXi host on which it is running.

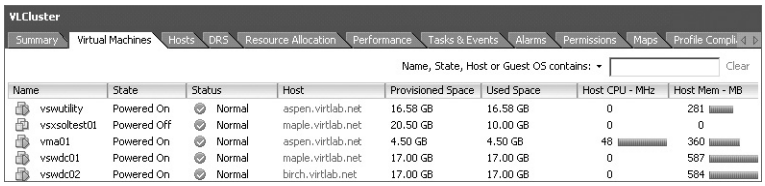

FIGURE 12.1
The Virtual Machines tab of a cluster object offers a quick look at virtual machine CPU and memory usage.

For ESX/ESXi clusters and resource pools, another tool you can use is the Resource Allocation tab. The Resource Allocation tab provides a picture of how CPU and memory resources are being used for the entire pool. This high-level method of looking at resource usage is useful for analyzing overall infrastructure utilization. This tab also provides an easy way of adjusting individual virtual machine or resource pool reservations, limits, and/or shares without editing each object independently.

vCenter Server also offers a very powerful, in-depth tool found on the Performance tab. The Performance tab provides a robust mechanism for creating graphs depicting the actual resource consumption over time for a given ESX/ESXi host or virtual machine. The graphs provide historical information and can be used for trend analysis. vCenter Server provides many objects and counters to analyze the performance of a single virtual machine or host for a selected interval. The Performance tab and the graphs are powerful tools for isolating performance considerations, and I discuss them in greater detail in the section ''Working with Performance Graphs.''

VMware also provides tools to run at the host level to help isolate and identify problems there. Because these tools require the presence of a Service Console, they work only with VMware ESX and not VMware ESXi. I'll take a look at these tools later in this chapter in the section ''Working with Command-Line Tools.''

Finally, I'll take the various tools that I've discussed and show how to use them to monitor the four major resources in a VMware vSphere environment: CPU, memory, network, and storage.

Let's get started with a discussion of alarms.

Using Alarms

In addition to the graphs and high-level information tabs, the administrator can create alarms for virtual machines, hosts, networks, and datastores based on predefined triggers provided with vCenter Server. Depending upon the object, these alarms can monitor resource consumption or the state of the object and alert the administrator when certain conditions have been met, such as high resource usage or even low resource usage. These alarms can then provide an action that informs the administrator of the condition by email or SNMP trap. An action can also automatically run a script or provide other means to correct the problem the virtual machine or host might be experiencing.

The creation of alarms to alert the administrator of a specific condition is not new in this version of vCenter Server. But the addition of new triggers, conditions, and actions gives the alarms more usefulness than in previous editions. As you can see in Figure 12.2, the alarms that come with vCenter Server are defined at the topmost object, the vCenter Server object. You'll also note that

there are far more predefined alarms in vCenter Server 4 than in previous versions of vCenter Server or VirtualCenter.

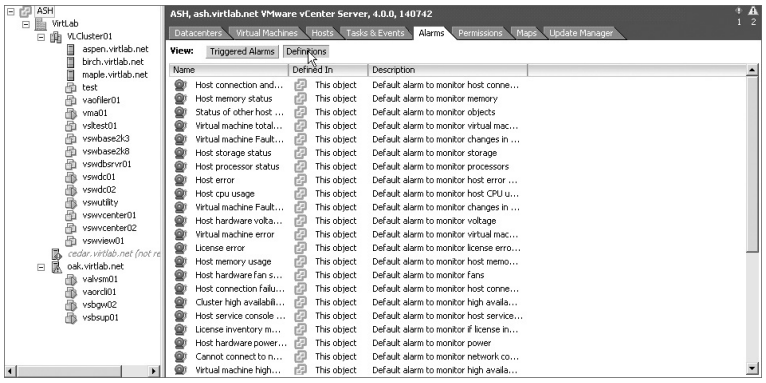

FIGURE 12.2
The default alarms for objects in vCenter Server are defined on the vCenter Server object itself.

These default alarms are usually generic in nature. Some of the predefined alarms include alarms to alert the administrator if any of the following happen:

- A host's storage status, CPU status, voltage, temperature, or power status changes
- A cluster experiences a VMware High Availability (HA) error
- A datastore runs low on free disk space
- A virtual machine's CPU usage, memory usage, disk latency, or even fault tolerance status changes

In addition to the small sampling of predefined alarms I've just described, there are many more, and VMware has enabled users to create alarms on just about any object within vCenter Server. This greatly increases the ability of vCenter Server to proactively alert administrators to changes within the virtual environment before a problem develops.

Because the default alarms are likely too generic for your administrative needs, creating your own alarms is often necessary. Before showing you how to create an alarm, though, I need to first discuss the concept of alarm scope. Once I've discussed alarm scope, I'll walk you through creating a few alarms. Then, in later sections of this chapter, I'll examine the use of those alarms along with other tools to monitor specific types of resource usage.

Understanding Alarm Scopes

When creating alarms, one thing to keep in mind is the *scope* of the alarm. In Figure 12.2, you saw the default set of alarms that are available in vCenter Server. These alarms are defined at the vCenter Server object and thus have the greatest scope—they apply to all objects managed by that vCenter Server instance. It's also possible to create alarms at the datacenter level, the cluster level, the host level, or even the virtual machine level. This allows you, the vSphere administrator, to create specific alarms that are limited in scope and are intended to meet specific monitoring needs.

When you define an alarm on an object, that alarm applies to all objects beneath that object in the vCenter Server hierarchy. The default set of alarms that VMware provides with vCenter Server are defined at the vCenter Server object and therefore apply to all objects—datacenters, hosts, clusters, datastores, networks, and virtual machines—managed by that instance of vCenter

Server. If you were to create an alarm on a resource pool, then the alarm would apply only to virtual machines found in that resource pool. Similarly, if you were to create an alarm on a specific virtual machine, that alarm would apply only to that specific virtual machine.

As you'll see later in this chapter, alarms are also associated with specific types of objects. For example, some alarms apply only to virtual machines, while other alarms apply only to ESX/ESXi hosts. You'll want to use this filtering mechanism to your advantage when creating alarms. For example, if you needed to monitor a particular condition on all ESX/ESXi hosts, you could define a host alarm on the datacenter or vCenter Server object, and it would apply to all ESX/ESXi hosts but not to any virtual machines.

It's important that you keep these scoping effects in mind when defining alarms so that your new alarms work as expected. You don't want to inadvertently exclude some portion of your VMware vSphere environment by creating an alarm at the wrong point in your hierarchy or by creating the wrong type of alarm.

Now you're ready to look at actually creating alarms.

Creating Alarms

As you've already learned, there are many different types of alarms that administrators might want to create. These alarms could be alarms that monitor resource consumption—such as how much CPU time a virtual machine is consuming or how much RAM an ESX/ESXi host has allocated—or these alarms can monitor for specific events, such as whenever a specific distributed virtual port group is modified. In addition, you've already learned that alarms can be created on a variety of different objects within vCenter Server. Regardless of the type of alarm or the type of object to which that alarm is attached, the basic steps for creating an alarm are the same. In the following sections, I'll walk you through creating a couple different alarms so that you have the opportunity to see the options available to you.

Creating a Resource Consumption Alarm

First, let's create an alarm that monitors resource consumption. As I discussed in Chapter 7, vCenter Server supports virtual machine snapshots. These snapshots capture a virtual machine at a specific point in time, allowing you to roll back (or revert) to that point-in-time state later. However, snapshots require additional space on disk, and monitoring disk space usage by snapshots was a difficult task in earlier versions of VMware Infrastructure. In vSphere, vCenter Server offers the ability to create an alarm that monitors VM snapshot space.

Before you create a custom alarm, though, you should ask yourself a couple of questions. First, is there an existing alarm that already handles this task for you? Browsing the list of predefined alarms available in vCenter Server shows that although some storage-related alarms are present, there is no alarm that monitors snapshot disk usage. Second, if you're going to create a new alarm, where is the appropriate place within vCenter Server to create that alarm? This refers to the earlier discussion of scope: on what object should you create this alarm so that it is properly scoped and will alert you only under the desired conditions? In this particular case, you'd want to be alerted to any snapshot space usage that exceeds your desired threshold, so a higher-level object such as the datacenter object or even the vCenter Server object would be the best place to create the alarm.

Perform the following steps to create an alarm that monitors VM snapshot disk space usage for all VMs in a datacenter:

1. Launch the vSphere Client if it is not already running, and connect to a vCenter Server instance.

YOU MUST USE VCENTER SERVER FOR ALARMS

You can't create alarms by connecting directly to an ESX/ESXi host; vCenter Server provides the alarm functionality. You must connect to a vCenter Server instance in order to work with alarms.

2. Navigate to an inventory view, such as Hosts And Clusters or VMs And Templates. You can use the menu bar, the navigation bar, or the appropriate keyboard shortcut.
3. Right-click the datacenter object, and select Alarm ➢ Add Alarm.
4. On the General tab in the Alarm Settings dialog box, enter an alarm name and alarm description.
5. Select Virtual Machine from the Monitor drop-down list.
6. Be sure that the radio button marked Monitor For Specific Conditions Or State, For Example, CPU Usage, Power State is selected.
7. On the Triggers tab, click the Add button to add a new trigger.
8. Set Trigger Type to VM Snapshot Size (GB). For this alarm, you're interested in snapshot size only, but other triggers are available:
   - VM Memory Usage (%)
   - VM Network Usage (kbps)
   - VM State
   - VM Heartbeat
   - VM Snapshot Size (GB)
   - VM CPU Ready Time (ms)
9. Ensure that the Condition column is set to Is Above.
10. Set the value in the Warning column to 1.
11. Set the value in the Alert column to 2. Figure 12.3 shows the Triggers tab after changing the Warning and Alert values.
12. On the Reporting tab, leave both the Range value at 0 and the Frequency value at 0. This ensures that the alarm is triggered at the threshold values you've specified and instructs vCenter Server to alert every time the thresholds are exceeded.

CAUTION: COUNTER VALUES WILL VARY!

The Is Above condition is selected most often for identifying a virtual machine, host, or datastore that exceeds a certain threshold. The administrator decides what that threshold should be and what is considered abnormal behavior (or at least interesting enough behavior to be monitored). For the most part, monitoring across ESX/ESXi hosts and datastores will be consistent. For example, administrators will define a threshold that is worthy of being notified about—such as CPU, memory, or

network utilization—and configure an alarm across all hosts for monitoring that counter. Similarly, administrators may define a threshold for datastores, such as the amount of free space available, and configure an alarm across all datastores to monitor that metric.

However, when looking at virtual machine monitoring, it might be more difficult to come up with a single baseline that works for all virtual machines. Specifically, think about enterprise applications that must perform well for extended periods of time. For these types of scenarios, administrators will want custom alarms for earlier notifications of performance problems. This way, instead of reacting to a problem, administrators can proactively try to prevent problems from occurring.

For virtual machines with similar functions like domain controllers and DNS servers, it might be possible to establish baselines and thresholds covering all such infrastructure servers. In the end, the beauty of vCenter Server's alarms is in the flexibility to be as customized and as granular as each individual organization needs.

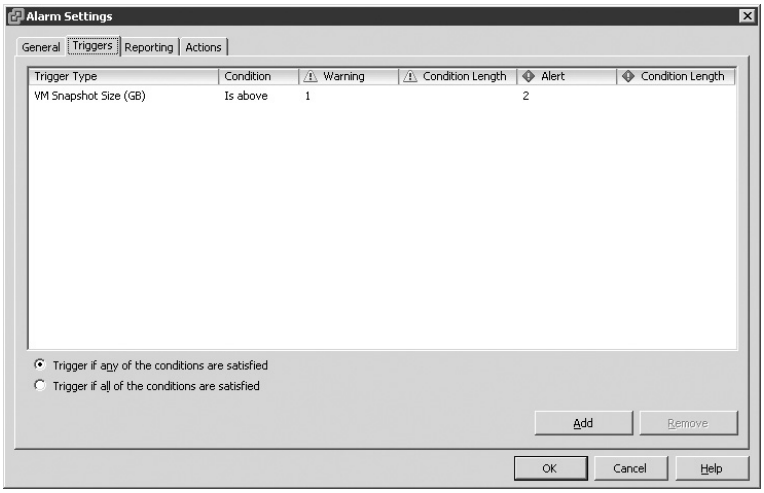

FIGURE 12.3
On the Triggers tab, define the conditions that cause the alarm to activate.

13. On the Actions tab, specify any additional actions that should be taken when the alarm is triggered. Some of the actions that can be taken include the following:

    - Send a notification email.
    - Send a notification trap via SNMP.
    - Change the power state on a VM.
    - Migrate a VM.
    - If you leave the Actions tab empty, then the alarm will alert administrators only within the vSphere Client. For now, leave the Actions tab empty.

Configuring vCenter Server for Email and SNMP Notifications

To have vCenter Server send an email for a triggered alarm, you must configure vCenter Server with an SMTP server. To configure the SMTP server, from the vSphere Client choose the Administration menu, and then select vCenter Server Settings. Click Mail in the list on the left, and then supply the SMTP server and the sender account. I recommend using a recognizable sender account so that when you receive an email, you know it came from the vCenter Server computer. You might use something like `vcenter-alerts@vmwarelab.net`.

Similarly, to have vCenter Server send an SNMP trap, you must configure the SNMP receivers in the same vCenter Server Settings dialog box under SNMP. You may specify from one to four management receivers to monitor for traps.

14. Click OK to create the alarm

The alarm is now created. To view the alarm you just created, select the datacenter object from the inventory tree on the left, and then click the Alarms tab on the right. Select Definitions instead of Triggered Alarms, and you'll see your new alarm listed, like in Figure 12.4.

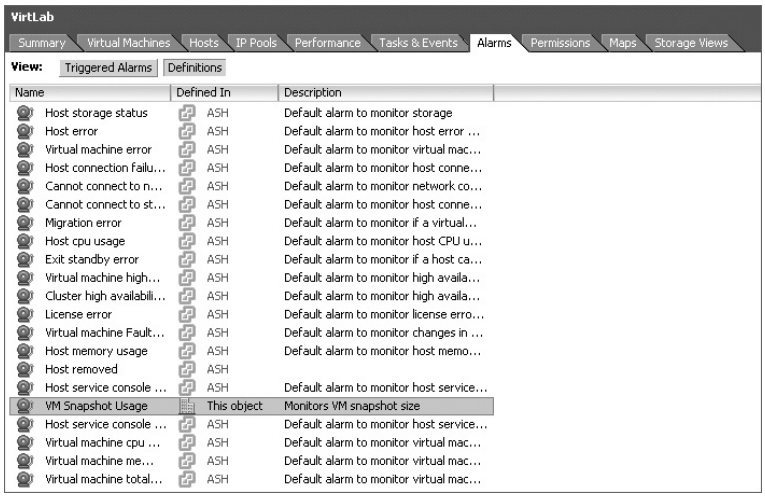

Figure 12.4
The Defined In column shows where an alarm was defined.

Using Range and Frequency with Alarms

Let's create another alarm. This time you'll create an alarm that takes advantage of the Range and Frequency parameters on the Reporting tab. With the VM snapshot alarm, these parameters didn't really make any sense; all you really needed was just to be alerted when the snapshot size exceeded a certain size. With other types of alarms, it may make sense to take advantage of these parameters.

The Range parameter specifies a tolerance percentage above or below the configured threshold. For example, the built-in alarm for virtual machine CPU usage specifies a warning threshold of 75 percent but specifies a range of 0. This means that the trigger will activate the alarm at exactly 75 percent. However, if the Range parameter were set to 5 percent, then the trigger would not activate the alarm until 80 percent (75 percent threshold + 5 percent tolerance range). This helps prevent alarm states from transitioning because of false changes in a condition by providing a range of tolerance.

The Frequency parameter controls the period of time during which a triggered alarm is not reported again. Using the built-in VM CPU usage alarm as our example, the Frequency parameter is set, by default, to five minutes. This means that a virtual machine whose CPU usage triggers the activation of the alarm won't get reported again—assuming the condition or state is still true—for five minutes.

With that information in mind, let's walk through another example of creating an alarm. This time you'll create an alarm that alerts based on VM network usage.

Perform the following steps to create an alarm that is triggered based on VM network usage:

1. Launch the vSphere Client if it is not already running, and connect to a vCenter Server instance.
2. Navigate to an inventory view, such as Hosts And Clusters or VMs And Templates.
3. Select the datacenter object from the inventory tree on the left.
4. Select the Alarms tab from the content pane on the right.
5. Select the Definitions button just below the tab bar to show alarm definitions instead of triggered alarms.
6. Right-click in a blank area of the content pane on the right, and select New Alarm.
7. Supply an alarm name and description.
8. Set the Monitor drop-down list to Virtual Machines.
9. Select the radio button marked Monitor For Specific Conditions Or State, For Example, CPU Usage, Power State.
10. On the Triggers tab, click Add to add a new trigger.
11. Set the Trigger Type column to VM Network Usage (kbps).
12. Set Condition to Is Above.
13. Set the value of the Warning column to 500, and leave the Condition Length setting at five minutes.
14. Set the value of the Alert column to 1000, and leave the Condition Length setting at five minutes.
15. On the Reporting tab, set Range to 10 percent, and set the Frequency parameter to five minutes.
16. Don't add anything on the Actions tab. Click OK to create the alarm.

Alarms on Other vCenter Server Objects

Although the two alarms you've created so far have been specific to virtual machines, the process is similar for other types of objects within vCenter Server.

Alarms can have more than just one trigger condition. The alarms you've created so far had only a single trigger condition. For an example of an alarm that has more than one trigger condition, look at the built-in alarm for monitoring host connection state. Figure 12.5 shows the two trigger conditions for this alarm. Note that the radio button marked Trigger If All Of The Conditions Are Satisfied is selected, ensuring that only powered-on hosts that are not responding will trigger the alarm.

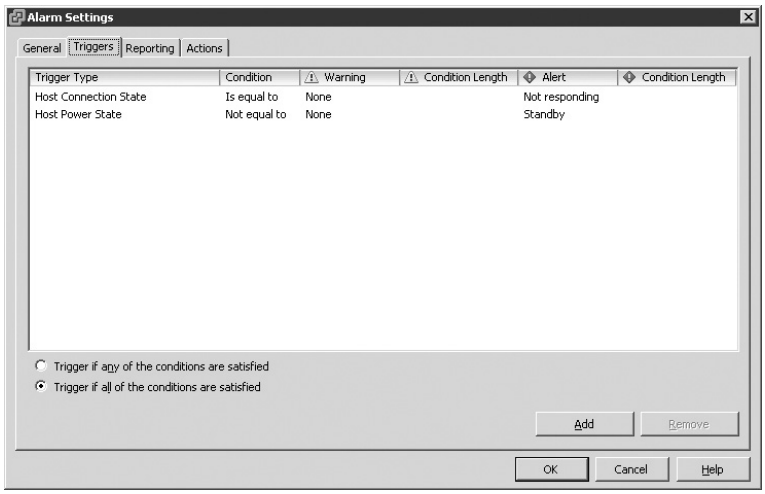

FIGURE 12.5
You can combine multiple triggers to create more complex alarms.

Don't Modify Built-in Alarms

In Chapter 9 I discussed vCenter Server's roles, and I mentioned that you should create custom roles instead of modifying the built-in roles supplied with vCenter Server. That same recommendation applies here: instead of modifying one of the built-in alarms, disable the built-in alarm (using the Enable This Alarm check box at the bottom of the General tab), and create a custom alarm that meets your needs.

It might seem obvious, but it's important to note that you can have more than one alarm for an object.

As with any new alarm, testing its functionality is crucial to make sure you get the desired results. You might find that the thresholds you configured are not optimized for your environment

and either aren't activating the alarm when they should or are activating the alarm when they shouldn't. In these cases, edit the alarm to set the thresholds and conditions appropriately. Or, if the alarm is no longer needed, right-click the alarm, and choose Remove to delete the alarm.

You'll be able to edit or delete alarms only if two conditions are met. First, the user account with which you've connected to vCenter Server must have the appropriate permissions granted in order to edit or delete alarms. Second, you must be attempting to edit or delete the alarm from the object on which it was defined. Think back to my discussion on alarm scope, and this makes sense. You can't delete an alarm from the datacenter object when that alarm was defined on the vCenter Server object. You must go to the object where the alarm is defined in order to edit or delete the alarm.

Now that you've seen some examples of creating alarms—and keep in mind that creating alarms for other objects within vCenter Server follows the same basic steps—let's take a look at managing alarms.

Managing Alarms

Several times so far in this chapter I've directed you to the Alarms tab within the vSphere Client. Up until now, you've been working with the Definitions view of the Alarms tab, looking at defined alarms. There is, however, another view to the Alarms tab, and that's the Triggered Alarms view. Figure 12.6 shows the Triggered Alarms view, which is accessed using the Triggered Alarms button just below the tab bar.

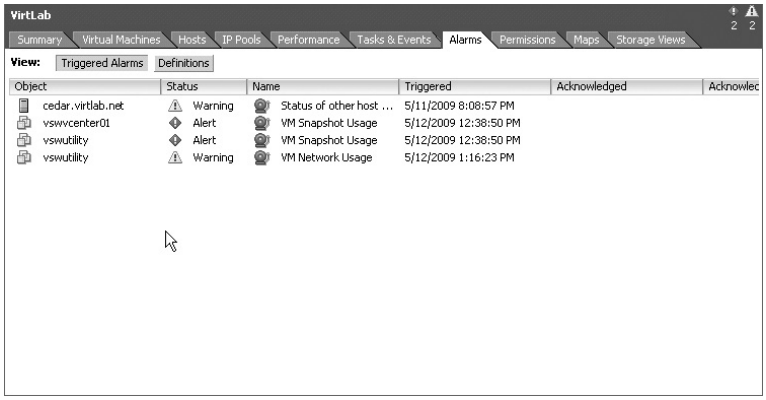

FIGURE 12.6
The Triggered Alarms view shows the alarms that vCenter Server has activated.

GETTING TO THE TRIGGERED ALARMS VIEW QUICKLY

The vSphere Client provides a handy shortcut to get to the Triggered Alarms view for a particular object quickly. When an object has at least one triggered alarm, small icons appear in the upper-right corner of the content pane for that object. You can see these icons in Figure 12.6. Clicking these icons takes you to the Triggered Alarms view for that object.

The Triggered Alarms view shows all the activated alarms for the selected object and all child objects. In Figure 12.6, the datacenter object was selected, so the Triggered Alarms view shows

all activated alarms for all the objects under the datacenter. In this instance, the Triggered Alarms view shows four alarms: one host alarm and three virtual machine alarms.

However, if only the virtual machine had been selected, the Triggered Alarms view on the Alarms tab for that virtual machine would show only the two activated alarms for that particular virtual machine. This makes it easy to isolate the specific alarms you need to address.

After you are in Triggered Alarms view for a particular object, a couple of actions are available to you for each of the activated alarms. For alarms that monitor resource consumption (that is, the alarm definition uses the Monitor For Specific Conditions Or State, For Example, CPU Usage, Power State setting selected under Alarm Type on the General tab), you have the option to acknowledge the alarm. To acknowledge the alarm, right-click the alarm, and select Acknowledge Alarm.

When an alarm is acknowledged, vCenter Server records the time the alarm was acknowledged and the user account that acknowledged the alarm. As long as the alarm condition persists, the alarm will remain in the Triggered Alarms view but is grayed out. When the alarm condition is resolved, the activated alarm disappears.

For an alarm that monitors events (this would be an alarm that has the Monitor For Specific Events Occurring On This Object, For Example, VM Powered On option selected under Alarm Type on the General tab), you can either acknowledge the alarm, as described previously, or reset the alarm status to green. Figure 12.7 illustrates this option.

FIGURE 12.7
For event-based alarms, you also have the option to reset the alarm status to green.

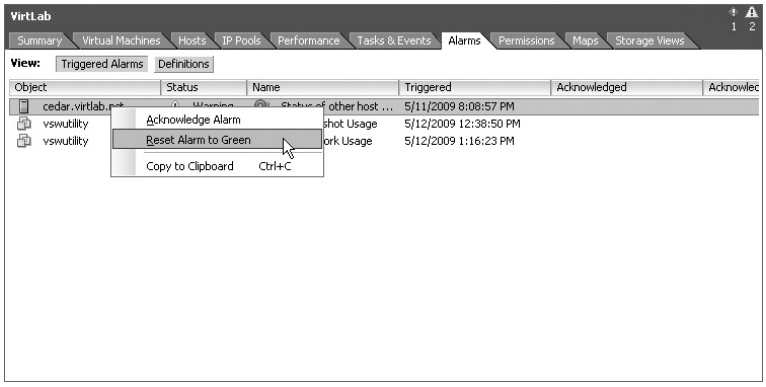

Resetting an alarm to green removes the activated alarm from the Triggered Alarms view, even if the underlying event that activated the alarm hasn't actually been resolved. This behavior makes sense if you think about it. Alarms that monitor events are merely responding to an event being logged by vCenter Server; whether the underlying condition has been resolved is unknown. So, resetting the alarm to green just tells vCenter Server to act as if the condition has been resolved. Of course, if the event occurs again, the alarm will be triggered again.

Now that you've looked at alarms for proactive performance monitoring, let's move on to using vCenter Server's performance graphs to view even more information about the behavior of virtual machines and ESX/ESXi hosts in your VMware vSphere environment.

Working with Performance Graphs

Alarms are a great tool for alerting administrators of specific conditions or events, but alarms don't provide the detailed information that administrators sometimes need to have. This is

where vCenter Server's performance graphs come in. vCenter Server has many new and updated features for creating and analyzing graphs. Without these graphs, analyzing the performance of a virtual machine would be nearly impossible. Installing agents inside a virtual machine will not provide accurate details about the server's behavior or resource consumption. The reason for this is elementary: a virtual machine is configured only with virtual devices. Only the VMkernel knows the exact amount of resource consumption for any of those devices because it acts as the arbitrator between the virtual hardware and the physical hardware. In most virtual environments, the virtual machines' virtual devices can outnumber the actual physical hardware devices, necessitating the complex sharing and scheduling abilities in the VMkernel.

By clicking the Performance tab for a datacenter, cluster, host, or virtual machine, you can learn a wealth of information. Before you use these graphs to help analyze resource consumption, I need to help you get to know the performance graphs and legends. I'll start with covering the two different layouts available in performance graphs: the overview layout and the advanced layout. First up is the overview layout.

Overview Layout

The Overview layout is the default view when you access the Performance tab. Figure 12.8 shows you the Overview layout of the Performance tab for an ESX host. Note the horizontal and vertical scrollbars; there's a lot more information here than the vSphere Client can fit in a single screen.

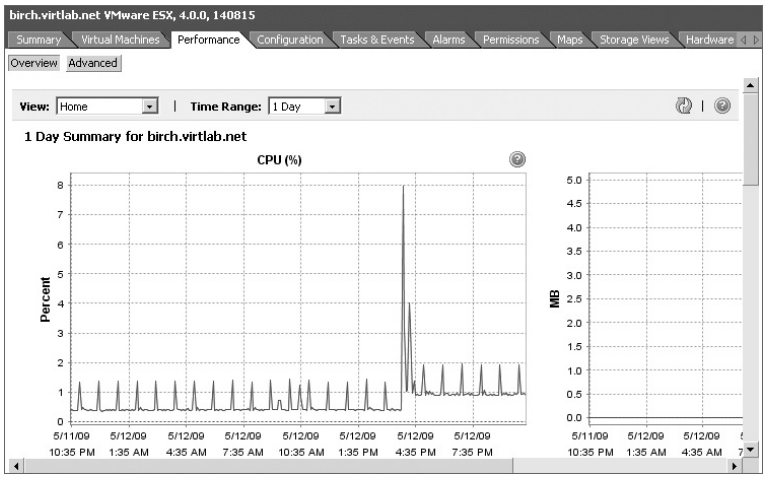

FIGURE 12.8
The Overview layout provides information on a range of performance counters.

At the top of the Overview layout are options to change the view or to change the date range. The contents of the View drop-down list change depending upon the object you select in the vSphere Client. Table 12.1 lists the different options available, depending upon what type of object you select in the vSphere Client.

Next to the View drop-down list is an option to change the date range for the data currently displayed in the various performance graphs. This allows you to set the time range to a day, a week, a month, or a custom value.

In the upper-right corner of the Overview layout, you'll see a button for refreshing the display and a button for getting help.

Below the gray title bar (where you'll find the View and Time Range drop-down lists, the Refresh button, and the Help button) are the actual performance graphs. The layout and the graphs that are included vary based on the object selected and the option chosen in the View drop-down list. I don't have the room here to list all of them, but a couple of them are shown in Figure 12.9 and Figure 12.10. I encourage you to explore a bit and find the layouts that work best for you.

TABLE 12.1: View Options on the Performance Tab

IF YOU ARE VIEWING THE PERFORMANCE TAB FOR THIS KIND OF OBJECT . . .	THE VIEW DROP-DOWN LIST CONTAINS THESE OPTIONS:
Datacenter	Clusters Storage
Cluster	Home Resource Pools & Virtual Machines Hosts
Resource pool	Home Resource Pools & Virtual Machines
Host	Home Virtual Machines
Virtual machine	Home Fault Tolerance Storage

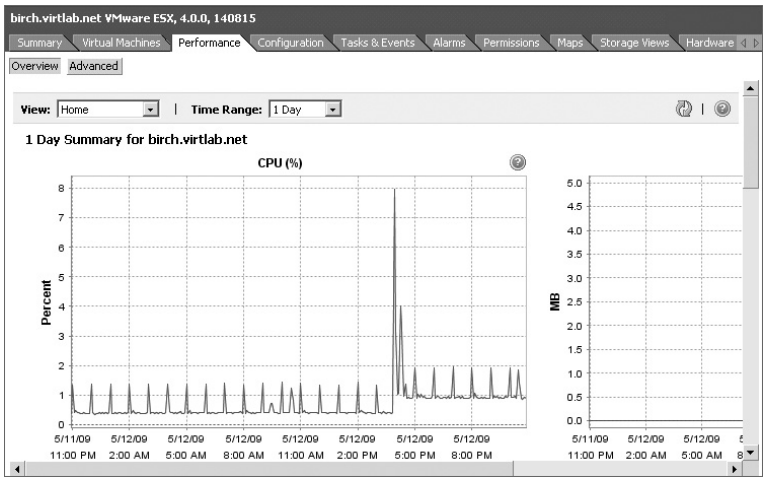

FIGURE 12.9
The Performance tab for an ESX/ESXi host in Overview layout includes eight charts, many of which are shown off-screen.

The Overview layout works well if you need a broad overview of the performance data for a datacenter, cluster, resource pool, host, or virtual machine. But what if you need more specific data in a more customizable format? The Advanced layout is the answer, as you'll see in the next section.

Advanced Layout

Figure 12.11 shows the Advanced layout of the Performance tab for a cluster of ESX/ESXi hosts. Here, in the Advanced layout, is where the real power of vCenter Server's performance graphs is made available to you.

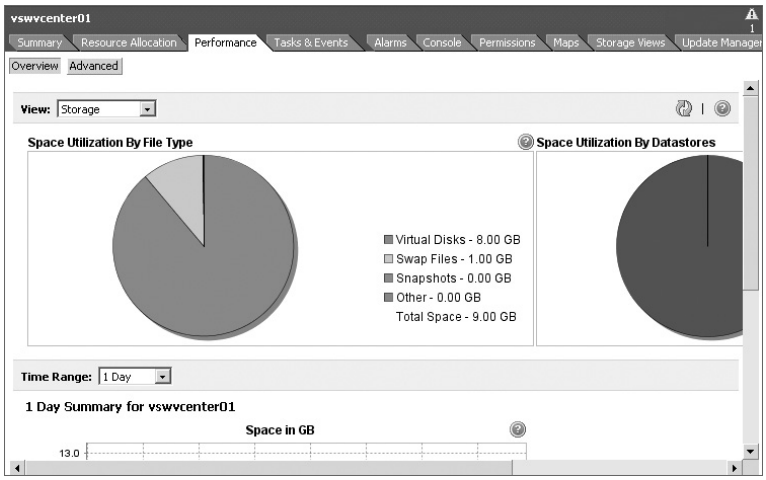

FIGURE 12.10
The Storage view of the Performance tab for a virtual machine in Overview layout displays a breakdown of storage utilization.

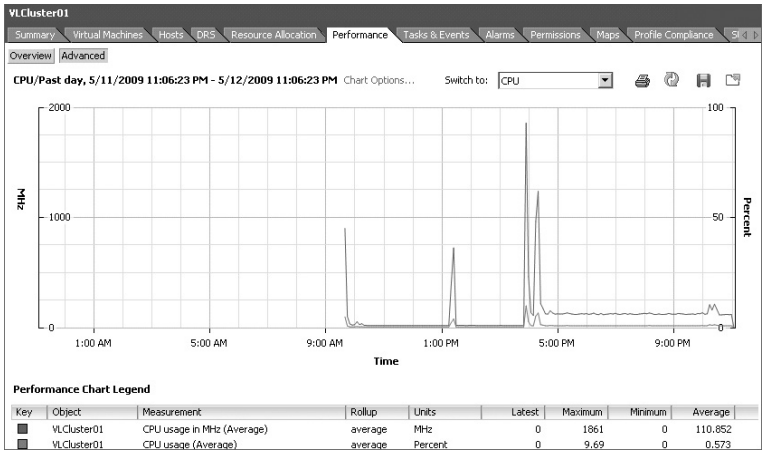

FIGURE 12.11
The Advanced layout of the Performance tab provides much more extensive controls for viewing performance data.

Starting from the top left, you'll see the name of the object being monitored. Just below that is the type of the chart and the time range. The Chart Options link provides access to customize settings for the chart. To the right, you'll find a drop-down list to quickly switch graph settings, followed by buttons to print the chart, refresh the chart, save the chart, or view the chart as a pop-up chart. The Print button allows you to print the chart; the Save button allows you to export the chart as a JPEG graphic. I'll discuss this functionality in the section ''Saving Performance Graphs.'' The Refresh button refreshes the data. The pop-up button opens the chart in a new window. This allows you to navigate elsewhere in the vSphere Client while still keeping a performance graph open in a separate window. Pop-up charts also make it easy to compare one ESX/ESXi host or virtual machine with another host or virtual machine. On each side of the graph are units of measure. In Figure 12.13, the counters selected are measured in percentages and megahertz. Depending on the counters chosen, there may be only one unit of measurement, but no more than two. Next, on the horizontal axis, is the time interval. Below that, the Performance Chart Legend provides color-coded keys to help the user find a specific object or item of interest. This area also breaks down the graph into the object being measured; the measurement being used; the units

of measure; and the Latest, Maximum, Minimum, and Average measurements recorded for that object.

Hovering the mouse pointer over the graph at a particular recorded interval of interest displays the data points at that specific moment in time.

Another nice feature of the graphs is the ability to emphasize a specific object so that it is easier to pick out this object from the other objects. By clicking the specific key at the bottom, the key and its color representing a specific object will be emphasized, while the other keys and their respective colors become lighter and less visible. For simple charts such as the one shown previously in Figure 12.11, this might not be very helpful. For busier charts with many performance counters, this feature is very useful.

Now that you have a feel for the Advanced layout, take a closer look at the Chart Options link. This link exposes vCenter Server's functionality in creating highly customized performance graphs. Figure 12.12 shows the Customize Performance Chart dialog box. This dialog box is the central place where you will come to customize vCenter Server's performance graphs. From here, you select the counters to view, the time ranges, and the kind of graph (line graph or stacked graph) to display.

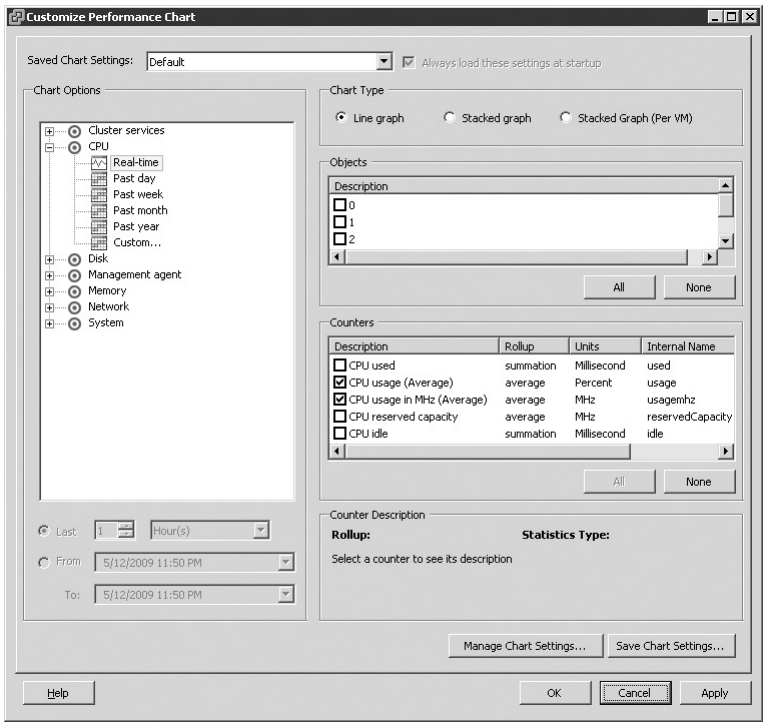

FIGURE 12.12
The Customize Performance Chart dialog box offers tremendous flexibility to create exactly the performance graph you need.

Because there is so much information available in the Customize Performance Chart dialog box, I've grouped the various options and types of information into the sections that follow.

CHOOSING A RESOURCE TYPE

On the left side of the Customize Performance Chart dialog box, you can choose which resource (Cluster Services, CPU, Disk, Management agent, Memory, Network, or System) to monitor or

analyze. The actual selections available in this area change depending upon the type of object that you have selected in vCenter Server. That is, the options available when viewing the Performance tab for an ESX/ESXi host are different from the options available when viewing the Performance tab of a virtual machine, a cluster, or a datacenter.

Within each of these resources, different objects and counters are available. Be aware that other factors affect what objects and counters are available to view; for example, in some cases the real-time interval shows more objects and counters than other intervals. The next few sections list the various counters that are available for the different resource types in the Customize Performance Chart dialog box.

If a particular counter is new to you, click it to highlight the counter. At the bottom of the dialog box, in a section called Counter Description, you'll see a description of the counter. This can help you determine which counters are most applicable in any given situation.

Setting a Custom Interval

Within each of the resource types, you have a choice of intervals to view. Some objects offer a Real-Time option; this option shows what is happening with that resource right now. The others are self-explanatory. The Custom option allows you to specify exactly what you'd like to see on the performance graph. For example, you could specify that you'd like to see performance data for the last eight hours. Having all of these interval options allows you to choose exactly the right interval necessary to view the precise data you're seeking.

Viewing CPU Performance Information

If you select the CPU resource type in the Chart Options section of the Customize Performance Chart dialog box, you can choose what specific objects and counters you'd like to see in the performance graph. Note that the CPU resource type is not available when viewing the Performance tab of a datacenter object. It is available for clusters, ESX/ESXi hosts, resource pools, and individual virtual machines.

Table 12.2 lists the objects and counters available for CPU performance information. Because CPU performance counters are not available at the datacenter object, the DC column is shaded. Not all these counters are available with all display intervals.

Quite a bit of CPU performance information is available. In the section ''Monitoring CPU Usage,'' I'll discuss how to use these CPU performance objects and counters to monitor CPU usage.

Viewing Memory Performance Information

If you select the memory resource type in the Chart Options section of the Customize Performance Chart dialog box, different objects and counters are available for display in the performance graph. The memory resource type is not available when viewing the Performance tab of a datacenter object. It is available for clusters, ESX/ESXi hosts, resource pools, and individual virtual machines.

In Table 12.3 you'll find the objects and counters available for memory performance information, depending upon the inventory object and display interval selected. As in Table 12.2, the DC column is shaded because memory counters are not available at the datacenter object. Not all these objects are available with all display intervals.

Later, in the section ''Monitoring Memory Usage,'' you'll get the opportunity to use these different objects and counters to monitor how ESX/ESXi and virtual machines are using memory.

TABLE 12.2: Available CPU Performance Counters

COUNTER	DC	CL	ESX	RP	VM
CPU used		✓	✓	✓	✓
CPU usage (Average)			✓		✓
CPU usage in MHz (Average)		✓	✓	✓	✓
CPU reserved capacity			✓		
CPU idle			✓		✓
CPU ready					✓
CPU system					✓
CPU wait					✓
Cluster total				✓	
CPU entitlement				✓	✓

VIEWING DISK PERFORMANCE INFORMATION

Disk performance is another key area that vSphere administrators need to monitor. Table 12.4 shows you the performance counters that are available for disk performance. Note that these counters aren't supported for datacenters, clusters, and resource pools, but they are supported for ESX/ESXi hosts and virtual machines. I've shaded the DC, CL, and RP columns in Table 12.4 because these counters are not available for datacenter, cluster, or resource pool objects. Not all counters are visible in all display intervals.

You'll use these counters in the section ''Monitoring Disk Usage'' later in this chapter.

VIEWING NETWORK PERFORMANCE INFORMATION

To monitor network performance, the vCenter Server performance graphs cover a wide collection of performance counters. Network performance counters are available only for ESX/ESXi hosts and virtual machines; they are not available for datacenter objects, clusters, or resource pools.

Table 12.5 shows the different network performance counters that are available. The DC, CL, and RP columns are shaded because network performance counters are not available for datacenter, cluster, and resource pool objects.

You'll use these network performance counters in the ''Monitoring Network Usage'' section later in this chapter.

VIEWING SYSTEM PERFORMANCE INFORMATION

ESX/ESXi hosts and virtual machines also offer some performance counters in the System resource type. Datacenters, clusters, and resource pools do not support any system performance counters.

TABLE 12.3: Available Memory Performance Counters

COUNTER	DC	CL	ESX	RP	VM
Memory usage (Average)		✓	✓		✓
Memory overhead (Average)		✓	✓	✓	✓
Memory consumed (Average)		✓	✓	✓	✓
Memory total		✓			
Memory shared common (Average)			✓		
Memory granted (Average)			✓		✓
Memory balloon (Average)			✓		✓
Memory shared (Average)			✓		✓
Memory swap in (Average)			✓		✓
Memory active (Average)			✓		✓
Memory zero (Average)			✓		✓
Memory heap (Average)			✓		
Swap out rate			✓		✓
Memory state			✓		
Memory unreserved (Average)			✓		
Memory reserved capacity			✓		
Memory used by VMkernel			✓		
Swap in rate			✓		✓
Memory swap out (Average)			✓		✓
Available heap memory			✓		
Memory swap used (Average)			✓		
Memory entitlement				✓	
Memory balloon target (Average)					✓
Memory swap target (Average)					✓
Memory swapped (Average)					✓

More information on the system performance counters is available in Table 12.6. Because system performance counters are not available for datacenter, cluster, and resource pool objects, these columns are shaded in Table 12.6.

TABLE 12.4: Available Disk Performance Counters

COUNTER	DC	CL	ESX	RP	VM
Kernel disk command latency			✓		
Disk read rate			✓		✓
Physical device command latency			✓		
Queue write latency			✓		
Disk commands issued			✓		✓
Physical device read latency			✓		
Disk write requests			✓		✓
Kernel disk read latency			✓		
Disk write latency			✓		
Stop disk command			✓		✓
Disk write rate			✓		✓
Queue command latency			✓		
Disk bus resets			✓		✓
Disk command latency			✓		
Disk read latency			✓		
Disk read requests			✓		✓
Queue read latency			✓		
Kernel disk write latency			✓		
Physical device write latency			✓		
Disk usage (Average)			✓		✓
Highest disk latency			✓		

TABLE 12.5: Available Network Performance Counters

COUNTER	DC	CL	ESX	RP	VM
Network data receive rate			✓		✓
Network packets received			✓		✓
droppedRx			✓		
Network usage (Average)			✓		✓
Network packets transmitted			✓		✓
droppedTx			✓		
Network data transmit rate			✓		✓

The majority of these counters are valid only for ESX/ESXi hosts, and they all center around how resources are allocated or how the ESX/ESXi host itself is consuming CPU resources or memory. As such, I won't be discussing them in any greater detail later in this chapter. I've included them here for the sake of completeness.

VIEWING OTHER PERFORMANCE COUNTERS

These are the other available performance counter types:

- ESX/ESXi hosts also offer a resource type (found in the Customize Performance Chart dialog box in the Chart Options section) marked as Management Agent. This resource type has only two performance counters associated with it: Memory used (Average) and Memory swap used (Average). These counters monitor how much memory the vCenter Server agent is using on the ESX/ESXi host.
- ESX/ESXi hosts participating in a cluster also have a resource type of Cluster Services, with two performance counters: CPU fairness and Memory fairness. Both of these counters show the distribution of resources within a cluster.
- The datacenter object contains a resource type marked as Virtual Machine Operations. This resource type contains performance counters that simply monitor the number of times a particular VM operation has occurred. This includes VM power-on events, VM power-off events, VM resets, VMotion operations, and Storage VMotion operations.

I've included this brief description of these counters for the sake of completeness, but I won't be discussing them any further.

MANAGING CHART SETTINGS

There's one more area of the Customize Performance Chart dialog box that I'll discuss, and that's the Manage Chart Settings and Save Chart Settings buttons in the lower-right corner.

TABLE 12.6: Available System Performance Counters

COUNTER	DC	CL	ESX	RP	VM
Resource CPU usage (Average)			✓		
Resource memory allocation maximum (in KB)			✓		
Resource CPU running (1 min. average)			✓		
Resource memory overhead			✓		
Resource memory mapped			✓		
Resource memory shared			✓		
Resource memory swapped			✓		
Resource memory zero			✓		
Resource memory share saved			✓		
Resource memory touched			✓		
Resource allocation minimum (in KB)			✓		
Resource CPU maximum limited (1 min.)			✓		
Resource CPU allocation (in MHz)			✓		
Resource CPU active (5 min. average)			✓		
Resource CPU allocation maximum (in MHz)			✓		
Resource CPU running (5 min. average)			✓		
Resource CPU active (1 min. average)			✓		
Resource CPU maximum limited (5 min.)			✓		
Resource CPU allocation shares			✓		
Resource memory allocation shares			✓		
Uptime			✓		✓
cosDiskUsage			✓		
Heartbeat					✓

After you've gone through and selected the resource type, display interval, objects, and performance counters that you'd like to see in the performance graph, you can save that collection of chart settings using the Save Chart Settings button. vCenter Server prompts you to enter a name for the saved chart settings. After a chart setting is saved, you can easily access it again from the drop-down list at the top of the performance graph advanced layout. Figure 12.13 shows the Switch To drop-down list, where two custom chart settings—VM Activity and Cluster Resources—are shown. By selecting either of these from the Switch To drop-down list, you can quickly switch to those settings. This allows you to define the performance charts that you need to see and then quickly switch between them.

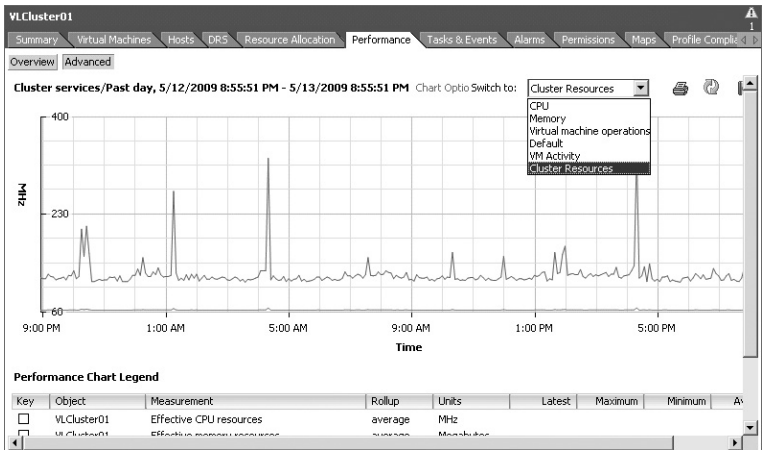

FIGURE 12.13
You can access saved chart settings from the Switch To drop-down list.

The Manage Chart Settings button allows you to delete chart settings you've saved but no longer need.

In addition to offering you the option of saving the chart settings, vCenter Server also allows you to save the graph.

SAVING PERFORMANCE GRAPHS

When I first introduced you to the Advanced layout view of the Performance tab, I briefly mentioned the Save button. This button, found in the upper-right corner of the Advanced layout, allows you to save the results of the performance graph to an external file for long-term archiving, analysis, or reporting.

When you click the Save button, a standard Windows Save dialog box appears. You have the option of choosing where to save the resulting file as well as the option of saving the chart either as a graphic file or as a Microsoft Excel spreadsheet. If you are going to perform any additional analysis, the option to save the chart data as an Excel spreadsheet is quite useful. The graphics options are useful when you need to put the performance data into a report of some sort.

There's a lot of information exposed via vCenter Server's performance graphs. I'll revisit the performance graphs again in the sections on monitoring specific types of resources later in this chapter, but first I'll introduce you to a few command-line tools you might also find useful in gathering performance information.

Working with Command-Line Tools

In addition to alarms and performance graphs, VMware also provides a couple command-line utilities to help with monitoring performance and resource usage. Unless stated otherwise, these tools work only with VMware ESX and not VMware ESXi, because they rely upon the presence of the Linux-based Service Console present only with VMware ESX.

Using esxtop

You can also monitor virtual machine performance using a command-line tool named `esxtop`. A great reason to use `esxtop` is the immediate feedback it gives you after you adjust a virtual machine. Using `esxtop`, you can monitor all four major resource types (CPU, disk, memory, and network) on a particular ESX host. Figure 12.14 shows some sample output from `esxtop`.

FIGURE 12.14
`esxtop` shows real-time information on CPU, disk, memory, and network utilization.

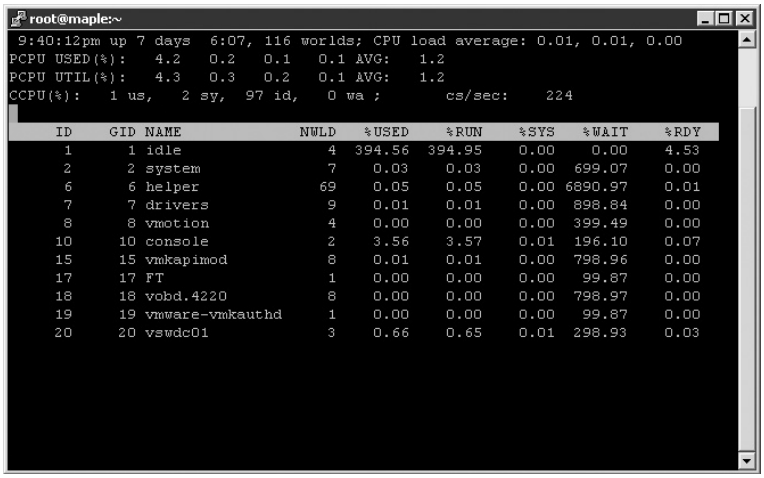

esxtop Is Only for VMware ESX

Because `esxtop` runs in the Linux-based Service Console, it works only on VMware ESX and not VMware ESXi. VMware supplies a separate tool called `resxtop` that supports VMware ESXi. I discuss that tool later in this section.

Upon launch, `esxtop` defaults to showing CPU utilization, as illustrated in Figure 12.14. At the top of the screen are summary statistics; below that are statistics for specific virtual machines and VMkernel processes. To show only virtual machines, press V. Be aware that `esxtop`, like many Linux commands you'll find in the ESX Service Console, is case sensitive, so you'll need to be sure to use an uppercase V in order to toggle the display of VMs only.

Two CPU counters of interest to view with `esxtop` are the CPU Used (%USED) and Ready Time (%RDY). You can also see these counters in the virtual machine graphs, but with `esxtop` they are calculated as percentages. The %WAIT counter is also helpful in determining whether

you have overallocated CPU resources to the VM. This might be the case if, for example, you've allocated two vCPUs to a virtual machine that really needs a single vCPU only. While in CPU mode, you can also press lowercase e to expand a virtual machine's CPU statistics so that you can see the different components that are using CPU time on behalf of a virtual machine. This is quite useful in determining what components of a virtual machine may be taking up CPU capacity.

If you switch away to another resource, press C (uppercase or lowercase) to come back to the CPU counters display. At any time when you are finished with `esxtop`, you can simply press q (lowercase only) to exit the utility and return to the Service Console shell prompt.

`esxtop` Shows Single Hosts Only

Remember, `esxtop` shows only a single ESX host. In an environment where VMotion, VMware Distributed Resource Scheduler (DRS), and VMware High Availability (HA) have been deployed, virtual machines may move around often. Making reservation or share changes while the virtual machine is currently on one ESX host may not have the desired consequences if the virtual machine is moved to another server and the mix of virtual machines on that server represents different performance loads.

To monitor memory usage with `esxtop`, press m (lowercase only). This gives you real-time statistics about the ESX host's memory usage in the top portion and the virtual machines' memory usage in the lower section. As with CPU statistics, you can press V (uppercase only) to show only virtual machines. This helps you weed out VMkernel resources when you are trying to isolate a problem with a virtual machine. The %ACTV counter, which shows current active guest physical memory, is a useful counter, as are the %ACTVS (slow moving average for long-term estimates), %ACTVF (fast moving average for short-term estimates), %ACTVN (prediction of %ACTV at next sampling), and SWCUR (current swap usage) counters.

To monitor network statistics about the vmnics, individual virtual machines, or VMkernel ports used for iSCSI, VMotion, and NFS, press n (lowercase only). The columns showing network usage include packets transmitted and received and megabytes transmitted and received for each vmnic or port. Also shown in the DNAME column are the vSwitches or dvSwitches and, to the left, what is plugged into them, including virtual machines, VMkernel, and Service Console ports. If a particular virtual machine is monopolizing the vSwitch, you can look at the amount of network traffic on a specific switch and the individual ports to see which virtual machine is the culprit. Unlike other `esxtop` views, you can't use V (uppercase only) here to show only virtual machines.

To monitor disk I/O statistics about each of the SCSI controllers, press d (lowercase only). Like some other views, you can press V (uppercase only) to show only virtual machines. The columns labeled READS/s, WRITES/s, MBREAD/s, and MBWRTN/s are most often used to determine disk loads. Those columns show loads based on reads and writes per second and megabytes read and written per second.

The `esxtop` command also lets you view CPU interrupts by pressing i. This command will show you the device(s) using the interrupt and is a great way to identify VMkernel devices, such as a vmnic, that might be sharing an interrupt with the Service Console. This sort of interrupt sharing can impede performance.

Another great feature of `esxtop` is the ability to capture performance data for a short period of time and then play back that data. Using the command `vm-support`, you can set an interval and duration for the capture.

Perform the following steps to capture data to be played back on `esxtop`:

1. Using PuTTY (Windows) or a terminal window (Mac OS X or Linux), open an SSH session to an ESX host.
2. Enter the `su -` command to assume root privileges.
3. While logged in as root or after switching to the root user, change your working directory to `/tmp` by issuing the command `cd /tmp`.
4. Enter the command `vm-support -S -i 10 -d 180`. This creates an `esxtop` snapshot, capturing data every 10 seconds, for the duration of 180 seconds.
5. The resulting file is a tarball and is gzipped. It must be extracted with the command `tar -xzf esx*.tgz`. This creates a `vm-support` directory that is called in the next command.
6. Run `esxtop -R /vm-support*` to replay the data for analysis.

For command-line junkies, `esxtop` is a great tool. Unfortunately, it's limited to VMware ESX because it relies upon the Linux-based Service Console. However, VMware does have a tool for performing some of the same tasks with ESXi. It's a tool called `resxtop`.

Using resxtop

Because VMware ESXi lacks a user-accessible Service Console where you can execute scripts, you can't use ''traditional'' `esxtop` with VMware ESXi. Instead, you have to use ''remote'' `esxtop`, or `resxtop`. The `resxtop` command is included with the vSphere Management Assistant (vMA), a special virtual appliance available from VMware that provides a command-line interface for managing both VMware ESX and VMware ESXi hosts.

Using `resxtop` is much the same as using `esxtop`. Before you can actually view real-time performance data, though, you first have to tell `resxtop` which remote server you want to use. To launch `resxtop` and connect to a remote server, enter this command:

```
resxtop --server esx1.vmwarelab.net
```

You'll want to replace `esx1.vmwarelab.net` with the appropriate hostname or IP address of the ESX/ESXi host to which you want to connect. When prompted, supply a username and password, and then `resxtop` will launch. Once `resxtop` is running, you can use the same command to switch between the various views.

Now that I've shown you the various tools that you will use to monitor performance in a VMware vSphere environment, let's go through the four major resources—CPU, RAM, network, and disk—and see how to monitor the usage of these resources.

Monitoring CPU Usage

When monitoring a virtual machine, it's always a good starting point to keep an eye on CPU consumption. Many virtual machines started out in life as underperforming physical servers. One of VMware's most successful sales pitches is being able to take all those lackluster physical boxes that are not busy and convert them to virtual machines. Once converted, virtual infrastructure managers tend to think of these virtual machines as simple, lackluster, and low-utilization servers with nothing to worry over or monitor. The truth, though, is quite the opposite.

When the server was physical, it had an entire box to itself. Now it must share its resources with many other workloads. In aggregate, they represent quite a load, and if some or many of them become somewhat busy, they contend with each other for the finite capabilities of the ESX/ESXi host on which they run. Of course, they don't know they are contending for resources, but the VMkernel hypervisor tries to placate them. Virtual CPUs need to be scheduled, and ESX/ESXi does a remarkable job given that there are more virtual machines than physical processors most of the time. Still, the hypervisor can do only so much with the resources it has, and invariably there comes a time when the applications running in that virtual machine need more CPU time than the host can give.

When this happens, it's usually the application owner who notices first and raises the alarm with the system administrators. Now the vSphere administrators have the task of determining why this virtual machine is underperforming. Fortunately, vCenter Server provides a number of tools that make monitoring and analysis easier. These are the tools you've already seen: alarms, performance graphs, and command-line utilities.

Let's begin with a hypothetical scenario. A help desk ticket has been submitted indicating that an application owner isn't getting the expected level of performance on a particular server, which in this case is a virtual machine. As the vSphere administrator, you need to first delve deeper into the problem and ask as many questions as necessary to discover what the application owner needs to be satisfied with performance. Some performance issues are subjective, meaning some users might complain about the slowness of their applications, but they have no objective benchmark for such a claim. Other times, this is reflected in a specific benchmark, such as the number of transactions by a database server or throughput for a web server. In this case, our issue revolves around benchmarking CPU usage, so our application is CPU intensive when it does its job.

Assessments, Expectations, and Adjustments

If an assessment was done prior to virtualizing a server, there might be hard numbers to look at to give some details as to what was expected with regard to minimum performance or a service-level agreement (SLA). If not, the vSphere administrator needs to work with the application's owner to make more CPU resources available to the virtual machine when needed.

vCenter Server's graphs, which you have explored in great detail, are the best way to analyze usage, both short- and long-term. In this case, let's assume the help desk ticket describes a slowness issue in the last hour. As you've already seen, you can easily create a custom performance graph to show CPU usage over the last hour for a particular virtual machine or ESX/ESXi host.

Perform the following steps to create a CPU graph that shows data for a virtual machine from the last hour:

1. Connect to a vCenter Server instance with the vSphere Client.
2. Navigate to the Hosts And Clusters or VMs And Templates inventory view.
3. In the inventory tree, select a virtual machine.
4. Select the Performance tab from the content pane on the right, and then change the view to Advanced.
5. Click the Chart Options link.

6. In the Customize Performance Chart dialog box, select CPU from the resource type list. Select the Custom interval.
7. Near the bottom of the Chart Options section, change the interval to Last 1 Hours.
8. Set the chart type to Line graph.
9. Select the virtual machine itself from the list of objects.
10. From the list of counters, select CPU Usage In MHz (Average) and CPU Ready. This shows you how much processor is actually being used and how long it's taking to schedule the VM on a physical processor.
11. Click OK to apply the chart settings.

CPU READY

CPU Ready shows how long a virtual machine is waiting to be scheduled on a physical processor. A virtual machine waiting many thousands of milliseconds to be scheduled on a processor might indicate that the ESX/ESXi host is overloaded, a resource pool has too tight a limit, or the virtual machine has too few CPU shares (or, if no one is complaining, nothing at all). Be sure to work with the server or application owner to determine an acceptable amount of CPU Ready for any CPU-intensive virtual machine.

This graph shows CPU utilization for the selected virtual machine, but it won't necessarily help you get to the bottom of why this particular virtual machine isn't performing as well as expected. In this scenario, I would fully expect the CPU Usage in MHz (Average) counter to be high; this simply tells you that the virtual machine is using all the CPU cycles it can get. Unless the CPU Ready counters are also high, indicating that the virtual machine is waiting on the host to schedule it onto a physical processor, you still haven't uncovered the cause of the slowness that triggered the help desk ticket. Instead, you'll need to move to monitoring host CPU usage.

Monitoring a host's overall CPU usage is fairly straightforward. Keep in mind that other factors usually come into play when looking at spare CPU capacity. Add-ons such as VMotion, VMware DRS, and VMware HA directly impact whether there is enough spare capacity on a server or a cluster of servers. Compared to earlier versions of ESX, the Service Console will usually not be as competitive for processor 0 because there are fewer processes to consume CPU time. Agents installed on the Service Console will have some impact, again on processor 0.

SERVICE CONSOLE STUCK ON 0

The Service Console, as noted, uses processor 0, but it will use processor 0 only. The Service Console does not get migrated to other processors even in the face of heavy contention.

Perform the following steps to create a real-time graph for a host's CPU usage:

1. Launch the vSphere Client if it is not already running, and connect to a vCenter Server instance.
2. Navigate to the Hosts And Clusters or VMs And Templates inventory view.

3. In the inventory tree, select a host. This shows you the Summary tab.
4. Click the Performance tab, and switch to Advanced view.
5. Click the Chart Options link.
6. In the Customize Performance Chart dialog box, select the CPU resource type and the Real-Time display interval.
7. Set Chart Type to Stacked Graph (Per VM).
8. Select all objects. You should see a separate object for each VM hosted on the selected ESX/ESXi host.
9. Select the CPU Usage (Average) performance counter.
10. Click OK to apply the chart settings and return to the Performance tab.

This chart shows the usage of all the virtual machines on the selected ESX/ESXi host in a stacked fashion. From this view, you should be able to determine whether there is a specific virtual machine or group of virtual machines that are consuming abnormal amounts of CPU capacity.

VMkernel Balancing Act

Always remember that on an oversubscribed ESX/ESXi host the VMkernel will load balance the virtual machines based on current loads, reservations, and shares represented on individual virtual machines and/or resource pools.

In this artificial scenario, I identified the application within the virtual machine as CPU-bound, so these two performance charts should clearly identify why the virtual machine isn't performing well. In all likelihood, the ESX/ESXi host on which the virtual machine is running doesn't have enough CPU capacity to satisfy the requests of all the virtual machines. Your solution, in this case, would be to use the resource allocation tools described in Chapter 10 to ensure that this specific application receives the resources it needs to perform at acceptable levels.

Monitoring Memory Usage

Monitoring memory usage, whether on a host or a virtual machine, can be challenging. The monitoring itself is not difficult; it's the availability of the physical resource that can be a challenge. Of the four resources, memory can be oversubscribed without much effort. Depending on the physical form factor chosen to host VMware ESX/ESXi, running out of physical RAM is easy to do. Although the blade form factor creates a very dense consolidation effort, the blades are sometimes constrained by the amount of physical memory and network adapters that can be installed. But even with other regular form factors, having enough memory installed comes down to how much the physical server can accommodate and your budget.

If you suspect that memory usage is a performance issue, the first step is to isolate whether this is a memory shortage affecting the host (you've oversubscribed physical memory and need to add more memory) or whether this is a memory limit affecting only that virtual machine (meaning you need to allocate more memory to this virtual machine or change resource allocation policies). Normally, if the ESX/ESXi host is suffering from high memory utilization, the predefined vCenter

Server alarm will trigger and alert the vSphere administrator. However, the alarm doesn't allow you to delve deeper into the specifics of how the host is using memory. For that, you'll need a performance graph.

Perform the following steps to create a real-time graph for a host's memory usage:

1. Connect to a vCenter Server instance with the vSphere Client.
2. Navigate to the Hosts And Clusters inventory view.
3. In the inventory tree, click an ESX/ESXi host. This shows you the Summary tab.
4. Click the Performance tab, and switch to Advanced view.
5. Click the Chart Options link.
6. In the Customize Performance Chart dialog box, select the Memory resource type and the Real-Time display interval.
7. Select Line Graph as the chart type. The host will be selected as the only available object.
8. In the Counters area, select the Memory Usage (Average), Memory Overhead (Average), Memory Active (Average), Memory Consumed (Average), Memory Used by VMkernel, and Memory Swap Used (Average). This should give you a fairly clear picture of how memory is being used by the ESX/ESXi host.

COUNTERS, COUNTERS, AND MORE COUNTERS

As with virtual machines, a plethora of counters can be utilized with a host for monitoring memory usage. Which ones you select will depend on what you're looking for. Straight memory usage monitoring is common, but don't forget that there are other counters that could be helpful, such as Ballooning, Unreserved, VMkernel Swap, and Shared, just to name a few. The ability to assemble the appropriate counters for finding the right information comes with experience and depends on what is being monitored.

9. Click OK to apply the chart options and return to the Performance tab.

These counters, in particular the Memory Swap Used (Average) counter, will give you an idea of whether the ESX/ESXi host is under memory pressure. If the ESX/ESXi host is not suffering from memory pressure and you still suspect a memory problem, then the issue likely lies with the virtual machine.

Perform the following steps to create a real-time graph for a virtual machine's memory usage:

1. Use the vSphere client to connect to a vCenter Server instance.
2. Navigate to either the Hosts And Clusters or the VMs And Templates inventory view.
3. In the inventory tree, click a virtual machine. This shows you the Summary tab.
4. Click the Performance tab, and switch to the Advanced view.
5. Click the Chart Options link.
6. In the Customize Performance Chart dialog box, select the Memory resource type and the Real-Time display interval.

7. Select Line Graph as the chart type.
8. In the list of counters, select to show the Memory Usage (Average), Memory Overhead (Average), Memory Consumed (Average), and Memory Granted (Average) counters. This shows memory usage, including usage relative to the amount of memory configured for the virtual machine.
9. Click OK to apply the chart options and return to the Performance tab.

From this performance graph, you will be able to tell how much of the memory configured for the virtual machine is actually being used. This might reveal to you that the applications running inside that virtual machine need more memory than the virtual machine has been assigned and that adding more memory to the virtual machine—assuming that there is sufficient memory at the host level—might improve performance.

Memory, like CPU, is just one of several different factors that can impact virtual machine performance. Network usage is another area that can impact performance, especially perceived performance.

Monitoring Network Usage

vCenter Server's graphs provide a wonderful tool for measuring a virtual machine's or a host's network usage.

Monitoring network usage requires a slightly different approach than monitoring CPU or memory. With either CPU or memory, reservations, limits, and shares can dictate how much of these two resources can be consumed by any one virtual machine. Network usage cannot be constrained by these mechanisms. Since virtual machines plug into a virtual machine port group, which is part of a vSwitch on a single host, how the virtual machine interacts with the vSwitch can be manipulated by the virtual switch's or port group's policy. For instance, if you need to restrict a virtual machine's overall network output, you would configure traffic shaping on the port group to restrict the virtual machine to a specific amount of outbound bandwidth. Unless you are using vNetwork Distributed Switches or the Nexus 1000V third-party distributed virtual switch, there is no way to restrict virtual machine inbound bandwidth on ESX/ESXi hosts.

Virtual Machine Isolation

Certain virtual machines may indeed need to be limited to a specific amount of outbound bandwidth. Servers such as FTP, file and print, or web and proxy servers, or any server whose main function is to act as a file repository or connection broker, may need to be limited or traffic shaped to an amount of bandwidth that allows it to meet its service target but not monopolize the host it runs on. Isolating any of these virtual machines to a vSwitch of its own is more likely a better solution, but it requires the appropriate hardware configuration.

To get an idea of how much network traffic is actually being generated, you can measure a virtual machine's or a host's output or reception of network traffic using the graphs in vCenter Server. The graphs can provide accurate information on the actual usage or ample information that a particular virtual machine is monopolizing a virtual switch, especially using the Stacked Graph chart type.

Perform the following steps to create a real-time graph for a stacked graph of transmitted network usage by each virtual machine on an ESX/ESXi host:

1. Launch the vSphere Client if it is not already running, and connect to a vCenter Server instance.
2. Navigate to either the Hosts And Clusters inventory view or the VMs And Templates inventory view.
3. In the inventory tree, click an ESX/ESXi host. This shows you the Summary tab.
4. Click the Performance tab, and switch to Advanced view.
5. Click the Chart Options link.
6. From the Customize Performance Chart dialog box, select the Network resource type and the Real-Time display interval in the Chart Options area.
7. Select a chart type of Stacked Graph (Per VM).
8. In the objects list, be sure all the virtual machines are selected.
9. In the list of counters, select the Network Data Transmit Rate counter. This gives you an idea of how much network bandwidth each virtual machine is consuming outbound on this ESX/ESXi host.
10. Click OK to apply the changes and return to the Performance tab.

What if you wanted a breakdown of traffic on each of the network interface cards (NICs) in the ESX/ESXi host, instead of by virtual machine? That's fairly easily accomplished by another trip back to the Customize Performance Chart dialog box.

Perform the following steps to create a real-time graph for a host's transmitted network usage by NIC:

1. Connect to a vCenter Server instance with the vSphere Client.
2. Navigate to the Hosts And Clusters inventory view.
3. In the inventory tree, select an ESX host. This will show you the Summary tab in the Details section on the right.
4. Select the Performance tab, and switch to Advanced view.
5. Click the Chart Options link.
6. Under Chart Options in the Customize Performance Chart dialog box, select the Network resource type and the Real-Time display interval.
7. Set the chart type to Line Graph.
8. In the objects list, select the ESX/ESXi host as well as all the specific NICs.
9. Select the Network Data Transmit Rate and Network Packets Transmitted counters.
10. Click OK to apply the changes and return to the Performance tab.

Very much like the earlier example for a virtual machine, these two counters will give you a window into how much network activity is occurring on this particular host in the outbound

direction for each physical NIC. This is especially relevant if you want to see different rates of usage for each physical network interface, which, by definition, represents different virtual switches.

Now that you've examined how to monitor CPU, memory, and network usage, there's only one major area left: monitoring disk usage.

Monitoring Disk Usage

Monitoring a host's controller or virtual machine's virtual disk usage is similar in scope to monitoring network usage. This resource, which represents a controller or the storing of a virtual machine's virtual disk on a type of supported storage, isn't restricted by CPU or memory mechanisms like reservations, limits, or shares. The only way to restrict a virtual machine's disk activity is to assign shares on the individual virtual machine, which in turn may have to compete with other virtual machines running from the same storage volume. vCenter Server's graphs come to our aid again in showing actual usage for both ESX/ESXi hosts and virtual machines.

Perform the following steps to create a host graph showing disk controller utilization:

1. Use the vSphere Client to connect to a vCenter Server instance.
2. Navigate to the Hosts And Clusters inventory view.
3. In the inventory tree, select an ESX/ESXi host. This shows you the Summary tab in the Details section on the right.
4. Select the Performance tab, and switch to the Advanced view.
5. Click the Chart Options link. This opens the Customize Performance Chart dialog box.
6. Under Chart Options, choose the Real-Time display interval for the Disk resource type.
7. Set the chart type to Line Graph.
8. Selecting an object or objects—in this case a controller—and a counter or counters lets you monitor for activity that is interesting or necessary to meet service levels. Select the objects that represent the ESX/ESXi host and one of the disk controllers.
9. In the counters list, select Disk Read Rate, Disk Write Rate, and Disk Usage (Average/Rate) to get an overall view of the activity for the selected controller.
10. Click OK to return to the Performance tab.

This performance graph will give you an idea of the activity on the selected disk controller. But what if you want to see disk activity for the entire host by each VM? In this case, a Stacked Graph view can show you what you need.

STACKED VIEWS

A stacked view is very helpful in identifying whether one particular virtual machine is monopolizing a volume. Whichever virtual machine has the tallest stack in the comparison may be degrading the performance of other virtual machines' virtual disks.

Now let's switch to the virtual machine view. Looking at individual virtual machines for insight into their disk utilization can lead to some useful conclusions. File and print virtual machines, or any server that provides print queues or database services, will generate some disk-related I/O that needs to be monitored. In some cases, if the virtual machine is generating too much I/O, it may degrade the performance of other virtual machines running out of the same volume. Let's take a look at a virtual machine's graph.

Perform the following steps to create a virtual machine graph showing real-time disk controller utilization:

1. Launch the vSphere Client if it is not already running, and connect to a vCenter Server instance.
2. Navigate to either the Hosts And Clusters view or the VMs And Templates inventory view.
3. In the inventory tree, click a virtual machine. This shows you the Summary tab in the Details section on the right.
4. Select the Performance tab, and switch to Advanced view.
5. Click the Chart Options link to open the Customize Performance Chart dialog box.
6. Under Chart Options, select the Disk resource type and the Real-Time display interval.
7. Set the chart type to Line Graph.
8. Set both objects listed in the list of objects.
9. In the list of counters, select Disk Read, Disk Write, and Disk Usage (Average/Rate).
10. Click OK to apply these changes and return to the Performance tab.

With this graph, you should have an informative picture of this virtual machine's disk I/O behavior. This virtual machine is busy at work generating reads and writes for its application. Does the graph show enough I/O to meet a service-level agreement, or does this virtual machine need some help? The graphs allow administrators to make informed decisions, usually working with the application owners, so that any adjustments to improve I/O will lead to satisfied virtual machine owners.

In addition, by looking at longer intervals of time to gain a historical perspective, you may find that a virtual machine has become busier or fallen off its regular output. If the amount of I/O is just slightly impaired, then adjusting the virtual machine's shares may be a way to prioritize its disk I/O ahead of other virtual machines sharing the volume. The administrator may be forced to move the virtual machine's virtual disk(s) to another volume or LUN if share adjustments don't achieve the required results. You can use Storage VMotion, described in Chapter 6, to perform this sort of LUN-based load balancing without any disruption to the end users.

When evaluating disk utilization for NFS-based datastores, you won't see any statistics or performance information in the Disk counters. To see information on NFS datastores, you'll have to look at the Network counters; specifically, you'll need to look at vmknic usage.

Real World Scenario

PERFORMANCE MONITORING FROM THE INSIDE AND THE OUTSIDE

It's important to remember that the very nature of how virtualization operates means that it is impossible to use performance metrics from within a guest operating system as an indicator of overall resource utilization. Here's why.

In a virtualized environment, each guest operating system "sees" only its slice of the hardware as presented by the VMkernel. A guest operating system that reports 100 percent CPU utilization isn't reporting that it's using 100 percent of the physical server's CPU, but rather that it's using 100 percent of the *CPU capacity given to it by the hypervisor*. A guest operating system that is reporting 90 percent RAM utilization is really only using 90 percent of the *RAM made available to it by the hypervisor*.

Does this mean that performance metrics gathered from within a guest operating system are useless? No, but these metrics cannot be used to establish overall resource usage—only relative resource usage. You must combine any performance metrics gathered from within a guest operating system with matching metrics gathered outside the guest operating system. By combining the metrics from within the guest operating system with metrics outside the guest operating system, you can create a more complete view of how a guest operating system is using a particular type of resource and therefore get a better idea of what steps should be taken to resolve any resource constraints.

For example, if a guest operating system is reporting high memory utilization but the vCenter Server resource management tools are showing that the physical system has plenty of memory available, this tells you that the guest operating system is using everything available to it and might perform better with more memory allocated to it.

Monitoring resources can be tricky, and it requires a good knowledge of the applications running in the virtual machines in your environment. If you are a new vSphere administrator, it's worth it to spend some time using vCenter Server's performance graphs to establish some baseline behaviors. This helps you become much more familiar with the "normal" operation of the virtual machines so that when something unusual or out of the ordinary does occur, you'll be more likely to spot it.

The Bottom Line

Use alarms for proactive monitoring. vCenter Server offers extensive alarms for alerting vSphere administrators to excessive resource consumption or potentially negative events. You can create alarms on virtually any type of object found within vCenter Server, including datacenters, clusters, ESX/ESXi hosts, and virtual machines. Alarms can monitor for resource consumption or for the occurrence of specific events. Alarms can also trigger actions, such as running a script, migrating a virtual machine, or sending a notification email.

Master It What are the questions a vSphere administrator should ask before creating a custom alarm?

Work with performance graphs. vCenter Server's detailed performance graphs are the key to unlocking the information necessary to determine why an ESX/ESXi host or virtual machine is performing poorly. The performance graphs expose a large number of performance counters across a variety of resource types, and vCenter Server offers functionality to save customized chart settings, export performance graphs as graphic figures or Excel workbooks, or view performance graphs in a separate window.

Master It You find yourself using the Chart Options link in the Advanced view of the Performance tab to frequently set up the same graph over and over again. Is there a way to save yourself some time and effort so that you don't have to keep re-creating the custom graph?

Gather performance information using command-line tools. VMware supplies a few command-line tools that are useful in gathering performance information. For VMware ESX hosts, `esxtop` provides real-time information about CPU, memory, network, or disk utilization. For both VMware ESX as well as VMware ESXi, `resxtop` can display the same information. Finally, the `vm-support` tool can gather performance information that can be played back later using `esxtop`.

Master It Compare and contrast the `esxtop` and `resxtop` utilities.

Monitor CPU, memory, network, and disk usage by both ESX/ESXi hosts and virtual machines. Monitoring usage of the four key resources—CPU, memory, network, and disk—can be difficult at times. Fortunately, the various tools supplied by VMware within vCenter Server can lead the vSphere administrator to the right solution. In particular, using customized performance graphs can expose the right information that will help a vSphere administrator uncover the source of performance problems.

Master It A junior vSphere administrator is trying to resolve a performance problem with a virtual machine. You've asked this administrator to see whether it is a CPU problem, and the junior administrator keeps telling you that the virtual machine needs more CPU capacity because the CPU utilization is high within the virtual machine. Is the junior administrator correct, based on the information available to you?