



Mosaic's IT Director Series

VMware Basics

Capacity Planning and Performance Monitoring

VMware Basics – Capacity Planning and Performance Monitoring

Introduction	3
VMware Complexity	5
VMware Virtual Infrastructure Overview	5
Physical Objects	6
ESX Server Hosts	6
VirtualCenter Server	7
Virtual Objects	7
Virtual Machines	7
DataCenters	8
Datastores	9
Clusters	9
Resource Pools	9
Folders	11
Performance Monitoring and Capacity Planning In VMware Environments	12
Capacity Planning	12
VMware Performance Monitoring	13
Virtual Machines	13
ESX Hosts	13
Clusters	13
Resource Pools	13
VMware’s Monitoring Capability	14
Finding VMware’s Goldilocks Zone with v Charter Pro	14
vCharter Pro	15
VM status summary view	16
ESX Host Server View	17
Cluster View	18
Resource Pools View	19
Summary	20
About Mosaic Technology	20

Finding VMware's Goldilocks Zone

*"This porridge is too hot," Goldilocks exclaimed.
So she tasted the porridge from the second bowl.
"This porridge is too cold."
So she tasted the last bowl of porridge.
"Ahhh, this porridge is just right!" she said happily.
And she ate it all up.*

In NASA and Astronomy terms the Goldilocks Zone is an area of space in which a planet is just the right distance from its home star so that its surface is neither too hot nor too cold.

In VMware environments ESX Hosts' Goldilocks Zone is where resource utilization levels are: Not too high. Not too low. Just right.

~

Introduction

We virtualize servers to consolidate and maximize resource utilization. A successful virtualization is in the Goldilocks Zone when it:

- Takes full advantage of available resources without hitting use levels that degrade performance of one or more virtual machines (too high)
- And uses available resources at levels that don't dip below minimum established utilization rates (too low)

You can achieve consistent VMware Goldilocks Zone performance with sound capacity planning and performance monitoring. And therein lays the problem.

Virtualization changes everything. Traditional approaches to infrastructure management – while sound in concept – no longer work. As companies race to take advantage of the well known benefits of virtualization they frequently are not prepared to deal with realities of a virtualized world.

In a typical scenario IT management decides to invest in virtualization. VMware is purchased, an administrator is sent to VMware training, and the virtualization train is off and running.

There is an underlying assumption that the methods and procedures in place can be easily adapted to the new virtual arena. However, with virtualization, instead of managing a series of stand-alone servers in terms of performance and resource utilization you now need to continually monitor multiple servers sharing defined resource pools.

There's a complexity inherent in virtualization that is often overlooked. Traditional one-box-at-a-time management will not work. Even the best designed and architected environment managed with a traditional server and/or OS centric administration approach will eventually deteriorate.

That's because the complexities of virtualization – and VMware in particular – demand on-going performance and capacity planning in order to be successful. Without it you can survive with short term fixes and lengthy trouble shooting sessions.

If you want a successful virtualization environment you need a good method of capacity planning and on-going performance evaluation.

In this paper we discuss virtualization in terms of VMware the current market leader. However, the concepts we use will be applicable to most virtual environments.

VMware Complexity

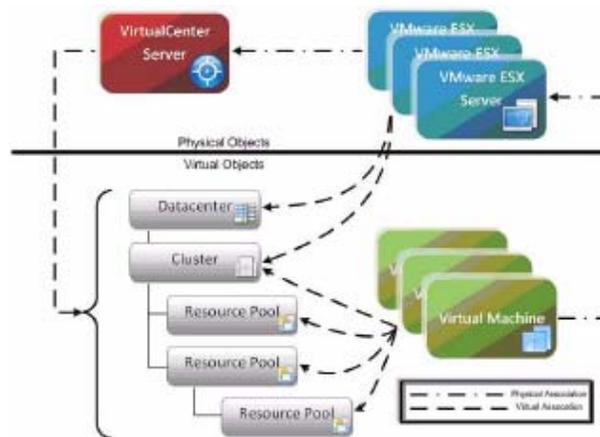
While simple in concept, VMware is complex. It contains components that make virtualization possible and quite manageable in limited configurations. However when you begin to expand into multiple ESX servers and multiple virtual machines competing for resources – effective management becomes an issue.

Before we discuss performance monitoring specifics a brief review of the key VMware components makes sense. As you review this information keep in mind that the interrelationship among ESX Hosts, VMs, Clusters, and Resource Pools make effective monitoring a must for successful VMware management.

VMware Virtual Infrastructure Overview

VMware Virtual Infrastructure 3 (VI3) provides an innovative mechanism for organizing and viewing any virtual infrastructure built on its platform. Using a combination of physical and logical components, VI3 effectively and efficiently fulfills VMware’s vision of the modern virtual infrastructure.

Figure 1 identifies components or objects that make up a typical VI3 implementation. This diagram makes a clear distinction between objects that exist in the physical world and those that are virtual.



VMware’s VirtualCenter Server enables the configuration of a hierarchical organizational structure that resides primarily within the virtual domain.

This lets you easily configure physical VMware ESX Servers and virtual machines to reside in logical groups that dictate various aspects of the virtual infrastructure (physical object location, resource allocations and limitations for virtual machines, and high availability settings for physical and virtual components).

In order to manage VMware better we need to understand the different roles the various physical and virtual objects play within the overall virtual infrastructure.

The VirtualCenter Server and VMware ESX Servers provide the physical foundation for the V13 infrastructure.

For management and monitoring purposes, virtual machines are classified as virtual components even though they have many characteristics of physical systems (like direct network and storage access). At any given time, a virtual machine must be contained within a single VMware ESX Server. The particular ESX Server in which a given virtual machine is contained may change over the lifetime of the virtual machine through the use of VMotion or VMware High Availability (VMware HA).

Physical Objects

Physical objects within VMware virtual infrastructure are those you can physically interact with. Virtual components or objects that make up a virtual environment cannot exist without underlying physical components.

A VMware ESX Server is an example of a physical component. Each ESX Server used to run virtual machines must have its own managing VirtualCenter Agent installed on it.

ESX Server Hosts

An ESX Server Host -- the single physical component required to begin building a virtual infrastructure -- provides hypervisor based architecture for controlling and managing resources for virtual machines that run on it.

Virtual machines running on a host share resources it provides. If resources become over-committed, the ESX Server hypervisor determines which virtual machines have priority access to shared resources (based on manual virtual machine configurations) and distributes available resources accordingly.

Each ESX Server is managed by a single VirtualCenter Server instance, and can be configured to exist logically within either a Datacenter or Cluster virtual object within the overall virtual infrastructure.

VirtualCenter Server

Although a VirtualCenter Server can technically exist as a virtual machine, it is considered a physical component within the VMware virtual infrastructure.

VMware's VirtualCenter is software that manages virtual environments built on the VMware virtualization platform. VirtualCenter creates a hierarchical structure of virtual objects that lets system administrators logically lay out virtual infrastructure configurations. VirtualCenter also enables other VMware functionality such as Distributed Resource Scheduling (DRS), VMotion, and High Availability (HA).

A single VirtualCenter Server can monitor approximately 100 VMware ESX Servers and 1500 virtual machines before performance and scalability challenges require a second VirtualCenter Server. Multiple VirtualCenter instances can be disbursed geographically to localize the management of large, distributed VI3 implementations.

Virtual Objects

Virtual objects can exist only within the confines of the virtual infrastructure. With the exception of virtual machines, virtual objects are logical and are used for organizing VMware ESX Servers and virtual machines, geographically or by function. In addition, virtual objects enable advanced configuration of resource management and high availability settings.

Virtual Machines

Creation and subsequent use of virtual machines is the primary purpose of building and maintaining a virtual infrastructure. Virtual machines share many of the characteristics of physical systems (like storage and network interaction). They do not have direct access to hardware used to process information *and* are considered virtual components within the virtual infrastructure.

Virtual machines (VMs) encompass more than a guest operating system such as Microsoft Windows. VMs also contain specific configurations that help define it, such as number of processors and amount of memory it can

leverage. All resource utilization for any VM on a VMware ESX Server is scheduled through that Server's hypervisor.

The efficient tracking and analysis of this scheduling of resources at both the virtual machine and the ESX Server Host level is an area that requires intensive analysis.

At any given time a VM must reside on a single VMware ESX Server. However, it can be moved across physical ESX Servers, without downtime, using vMotion. vMotion proactively moves a VM from one ESX Server to another while avoiding downtime that can arise from events such as patching a physical host server. vMotion also gives system administrator a manual method to balance virtual machine workloads based on resource utilization trends.

VMware VirtualCenter offers additional features including VMware Distributed Resource Scheduling (DRS) for automating the process of balancing VMware ESX Server utilization and VMware High Availability (HA) feature for recovering from Host failure within a Cluster.

DataCenters

A Datacenter is the topmost virtual object within a VirtualCenter Server implementation. It is required before any VMware ESX Server Hosts can be added to a VirtualCenter.

A Datacenter identifies physical boundaries within which an ESX Server Host can exist. In most implementations these boundaries constitute a single physical location that contains a large number of ESX Server Hosts. While there is no hard and fast rule that says a Datacenter must exist entirely at one physical location, other Datacenter implementations are uncommon.

Within the boundaries of a Datacenter, objects of the same type cannot have the same name. For example, it is not possible to configure two ESX Server Hosts with the same name to reside within the same Datacenter. The same is true for virtual machines, Clusters, Resource Pools and any other objects that can be created and configured to reside within a Datacenter. Objects of the same type can have identical names as long as they are located in different Datacenters.

Datstores

Datstore management is conducted at both Datacenter and ESX Server levels. Each Datstore is contained within a Datacenter. It must be uniquely named within that Datacenter.

Datstores represent storage locations for VM files. The storage location can be a local file system path, a Virtual Machine File System Storage (VMFS) volume, or a Network Attached Storage directory.

You can configure ESX Server Hosts to mount a set of network drives (or Datstores). For each storage location within a Datacenter there is only one Datstore, so multiple Hosts may be configured to point to the same Datstore. Whenever an ESX Server Host accesses a virtual machine or file within a Datacenter it must use the appropriate Datstore path.

Each Datstore object keeps a record of ESX Server Hosts that have mounted it, and a Datstore object can be removed only if no Hosts are currently mounting that Datstore. Datstores are host-independent *and* platform-independent. Therefore, they do not change in any way when the virtual machines contained within them are moved from one ESX Server to another.

Clusters

A Cluster object is a group of VMware ESX Servers that share common storage resources and network configurations. A Cluster represents a pool of combined resources of all ESX Server Hosts assigned to the Cluster.

For example, if you add four ESX Servers to a Cluster and each ESX Server has 2x2 GHz processors with 4 GB of memory, the Cluster shows a pool of 16 GHz of CPU processing power and 16 GB of memory available for use by virtual machines.

A Cluster also serves as the boundary for virtual machine migration activity through the VMware VMotion or VMware HA features. When using either of these technologies for virtual machine migration it is critical that participating ESX Server Hosts have identical storage resource and network configurations -- this is guaranteed by the very definition of a Cluster.

Resource Pools

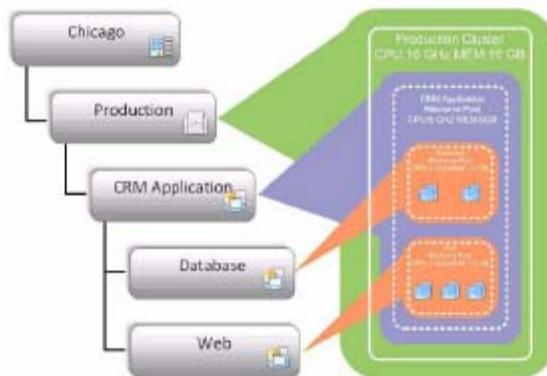
Resource Pools let administrators fine tune resource allocations within a Cluster. You can configure a Resource Pool to "own" a portion of the overall

available resources within a Cluster and then assign virtual machines to that Resource Pool. This lets you prioritize virtual machines— either limit or guarantee certain resources to a particular virtual machine or group of virtual machines.

Resource Pools can be configured in many ways, from simple to complex.

- **Simple example** -- Two Resource Pools are configured within a Cluster. One is named Production Virtual Machines, the other Development Virtual machines. The Production Resource Pool is configured with a "High" share priority and the Development Resource Pool is configured with the default "Normal" share priority. In this case any virtual machine residing in the Production Resource Pool automatically gets twice the priority -- in terms of access to system resources during periods of contention -- of any virtual machine in the Development Resource Pool.
- **Advanced example** -- Four ESX Servers are added to a Cluster and each ESX Server has 2x2 GHz processors with 4 GB of memory. The Cluster therefore has a pool of 16 GHz of CPU processing power and 16 GB of memory available to virtual machines. Figure 2 illustrates that the Production Cluster resource that resides in the Chicago Datacenter has 16 GHz of processing power and 16 GB of memory.

A Resource Pool is created for a CRM Application with access to 8 GHz of the Cluster's total CPU resources and 6 GB of the Cluster's total memory. Drilling down further you see that within the CRM Application Resource Pool there are two more Resource Pools (Database and Web). The existence of the Database Resource Pool ensures that key database virtual machines have access to resources necessary to perform their high transactional operations. The web servers have access to a smaller portion of the overall resources—just enough to provide the necessary end user responsiveness from a web transaction perspective without impacting the key backend database infrastructure.



Folders

Folders are hierarchical components within VirtualCenter that let administrators easily organize and manage virtual environments. Three different types of folders can exist within the various layers of the virtual infrastructure. The following table lists available types of folders, and explains levels at which they can exist and the objects they can contain.

Folder Type	Level at Which It Can Exist	Objects IT Can Contain
Datacenter	VirtualCenter Root	Datacenters
Virtual Machine	Datacenter	Virtual Machines and Templates
Compute Resources	Datacenter	Hosts and Clusters

Folders may contain nested folders of the same type, but not other types. For example you cannot create a Virtual Machine folder within a Datacenter folder.

Folders are provided strictly for organizational and management purposes. They offer administrators a way to classify objects that is not tied to (and therefore bound by) the virtual/physical relationship framework.

For example, two Datacenter folders are created at a VirtualCenter root; one is labeled Primary Datacenters, the other Disaster Recovery Datacenters. An administrator can:

1. Configure multiple Primary Datacenters containing production ESX Servers
2. Place those Datacenters in the Primary Datacenters folder
3. Assign the necessary permissions to that folder to allow standard users to perform management tasks for the entire primary virtual infrastructure.

The administrator can then:

1. Configure multiple Disaster Recovery Datacenters containing disaster recovery ESX Servers,
2. Place those Datacenters in the Disaster Recovery Datacenters folder
3. Assign a different set of permissions to that folder.

This prevents standard users from building virtual machines that may take over resources that are necessarily dedicated to HA-configured disaster failover virtual infrastructure components.

Performance Monitoring and Capacity Planning In VMware Environments

Finding the “Goldilocks” zone -- Not too high. Not too low. Just right.

An interesting facet of virtualization is that the very reason you virtualized – to maximize resource utilization – creates an environment that must be managed more aggressively than stand-alone servers. While it is ok to let a stand-alone server run at 20% capacity – the objective of virtualization is to keep resource utilization in the “Goldilocks” range which could easily be 70% to 80% capacity.

At issue is that multiple VMs on a physical server compete for resources. You need to account for and manage this contention. As we showed above VMware is complex and requires a new way of measuring and management.

Capacity Planning

Most everyone does some form of capacity planning prior to installing VMware. This initial effort defines how you configure your ESX servers and begins to set limits on virtualized processes. If done well it also gives you baseline performance data that provides concrete insight into utilization peaks and valleys. This previews how one VM will impact others and leads to logical groupings of VM on appropriate physical machines.

A model of capacity planning for virtualized environments generally includes:

- Profile virtual machines: Identify by type (database, webserver, etc) or application (email, finance, CRM, etc.). Should include all relevant information from OS to configurations – the more detail the better.
- Profile host servers: Identify by type (ESX, Virtual Server Center, etc). Include hardware specs and workload/utilization projections.
- Identify storage requirements with projected growth requirements.
- Define “Goldilocks” thresholds – upper utilization limits that define critical use situations.
- Understand demand – normal and peak. Data collected over time gives insight into virtual resource utilization and allows for rational growth projections and subsequent infrastructure sizing.

The activities of capacity planning are essentially the same required in performance monitoring – it’s just done in an on-going and consistent manner.

VMware Performance Monitoring

In VMware environments there are four key areas you need to track and monitor:

- Virtual Machines
- ESX hosts
- Clusters
- Resource Pools

Virtual Machines

As we said earlier, the beauty of VMs is the way they let you maximize resources. But with that comes increased tracking and monitoring to keep resource utilization in the “Goldilocks” zone.

Part of VMs flexibility is the ability to grab resources such as memory from adjacent machines to temporarily meet increased demand. In a best practices model you would have your VMs set up with optimal sharing of resource pools to avoid those temporary events.

To achieve that you need to have on-going and consistent VM monitoring of CPU, memory, disk and network metrics.

ESX Hosts

The ESX host is the foundation of your virtual infrastructure. If its performance lags (or fails) all the resident VMs and their associated applications will pay the price. So you need robust monitoring here as well – on the physical CPU, disk, memory and network.

Clusters

Clusters behave like and can be managed like ESX hosts. Loss of performance in a cluster could be an indicator of an underlying infrastructure problem. So a consistent monitoring of cluster services, CPU, and memory is recommended.

Resource Pools

Resource Pools let you fine tune resource allocations within an ESX or Cluster. You can configure a Resource Pool to “own” a portion of the overall available resources within a Cluster and then assign virtual machines to that Resource Pool. This lets you prioritize virtual machines— either limit or guarantee certain resources to a particular virtual machine or group of virtual

machines. From a performance monitoring perspective you should be tracking memory and CPU usage.

VMware's Monitoring Capability

From whatever angle you look, VMware monitoring is complex and could be a daunting task. VMware Virtual Center does provide detailed reporting on most aspects of the virtualized environment. However, the information is generally silo based – it looks at independent data points and does not present information in an integrated time-line focused fashion. Hence a lot of manual collating and interpretation are needed.

The good news is that there is an emerging class of products designed specifically for virtual environment capacity planning and performance monitoring – Vizioncore, VKernel, Platespin, and Nimsoft among others.

For purposes of discussion we'll be using Vizioncore's vCharter as an example of a sound automated Performance Monitoring/Capacity Planning solution.

(Editor's note: Mosaic Technology and Vizioncore have a formal business partner relationship)

Finding VMware's Goldilocks Zone with v Charter Pro

With the explosion of server virtualization technology, the number of virtual machines (VMs) running in the typical datacenter has grown exponentially. New VMs can be deployed easily, and VMs can also be moved about the datacenter just as easily. While this flexibility resolves some issues for IT managers, it brings up a whole host of other concerns.

You need to track new VMs, know where they moved and where they currently reside. And as we've said previously, with much higher utilization per server and multiple guest VMs sharing resources of a single physical host, performance of these interrelated components are closely linked and interdependent.

With companies running hundreds to thousands of VMs and multiple VirtualCenters, it is imperative to deploy a centralized and scalable monitoring solution. This lets you quickly pinpoint trouble spots and easily track the impact of one VM on other VMs. It also lets you follow VMs as they move through the datacenter.

vCharter Pro helps users prevent unwanted downtime. It helps you maintain all your VMware environments in the Goldilocks Zone. It optimizes resource

allocation among VMs – a major consideration for companies deploying small to large-scale virtualization implementations.

vCharter Pro

vCharter Pro monitors your entire VMware enterprise through a single “pane-of-glass” approach. It gives you top-down dashboard views with specific rule and metrics for each level. It’s scalable and fully customizable to any environment.

Using role based administration, vCharter Pro automates metrics monitoring and storage. And aggregates data from VirtualCenter for simplified analysis and planning.

vCharter Pro is easy to use and install. You download it and install in minutes. Turn it on and start monitoring. It has a web based interface – no client needs to be installed.

vCharter Pro takes you out of VMware’s spreadsheet-centric data focus to consolidated information presented through an easy to understand and act on GUI.

vCharter Pro VC agents collect raw metrics about the following levels:

- VirtualCenter
- Datacenters
- Resource Pools
- Clusters
- ESX Servers
- Virtual Machines
- Datastores

The data is stored in a database over time to provide history. You use this data for trend analysis (when will I need to buy new storage, etc.)

In the absence of a tool such as vCharter – you can still get this information from VMware. However it is not aggregated and/or consolidated. Many smart and creative VMware administrators have built spreadsheet based work-arounds. This provides an effective day to day management system that keeps the system running. However it is labor intensive and may not be the most effective way to find the Goldilocks Zone.

The following pages give you a view of the four key areas we discussed earlier; VMs, ESX Hosts, Clusters, and Resource Pools. The objective is to give you a sample of the types of information you can have at your fingertips to help effectively manage your virtual environment. It is not intended as a comprehensive capabilities discussion.

VM status summary view

Provides comprehensive view of VM for quick analysis and problem identification.

Detailed display of key metrics

- CPU (%Ready / % Used)
- Network I/O
- Disk I/O
- Memory

Detailed Network and Storage

- I/O metrics

Change Control log to track changes to VM

vmExplorer
Virtual Machine

Summary (vRanger Pro 3.2.1 20699) | Datastores

Summary

Connection Status: connected
 Type: Virtual Machine
 Current State: Powered On
 Uptime: 100 %
 VMware Tools: Running, Up to Date
 OS: Microsoft Windows Server 2003, Standard Edition (32-bit)

Resources

Processors: 1 CPUs : CPU Shares ()
 Memory Capacity: 256.0 MB : Memory Shares ()
 Network Interfaces: 1
 Storage Devices: 1

Virtual Machine Utilizations

CPU 0
54.0 MHz
Composite CPU Utilization

Network Summary

Data Sent: 0.0 B/s
 Packets Sent: 0.0
 Data Received: 0.0 B/s
 Packets Received: 133.0

NIC	Send	Receive
1	0.0 B/s	0.0 B/s

Storage Summary

Swap In: 0.0 B/s
 Swap Out: 0.0 B/s
 Write Rate: 0
 Writes: 0
 Read Rate: 0
 Reads: 0

Memory Utilization

Granted	Balloon Target	Swap Target	Zero Page
256.0 MB	0.0 KB	0.0 KB	48.7 MB
Consumed	Balloon Actual	Swap Used	
178.0 MB	0.0 KB	0.0 KB	

Virtual Machine Server and Resource Pool Changes

Change	Time	Details
	3/14/08 10:43 AM	Virtual Machine vRanger Pro 3.2.1 20699 ESX: Server Change from unknown to vc-esx3-suplab03.vizioncore.int
	3/14/08 10:43 AM	Virtual Machine vRanger Pro 3.2.1 20699 Resource Pool Change from unknown to Resources.

ESX Host Server View

Provides detailed view of performance metrics. Show individual VM status. Alarms identify problems based on your criteria.



Cluster View

Single view of all Alarms in Cluster

- By Infrastructure level or severity

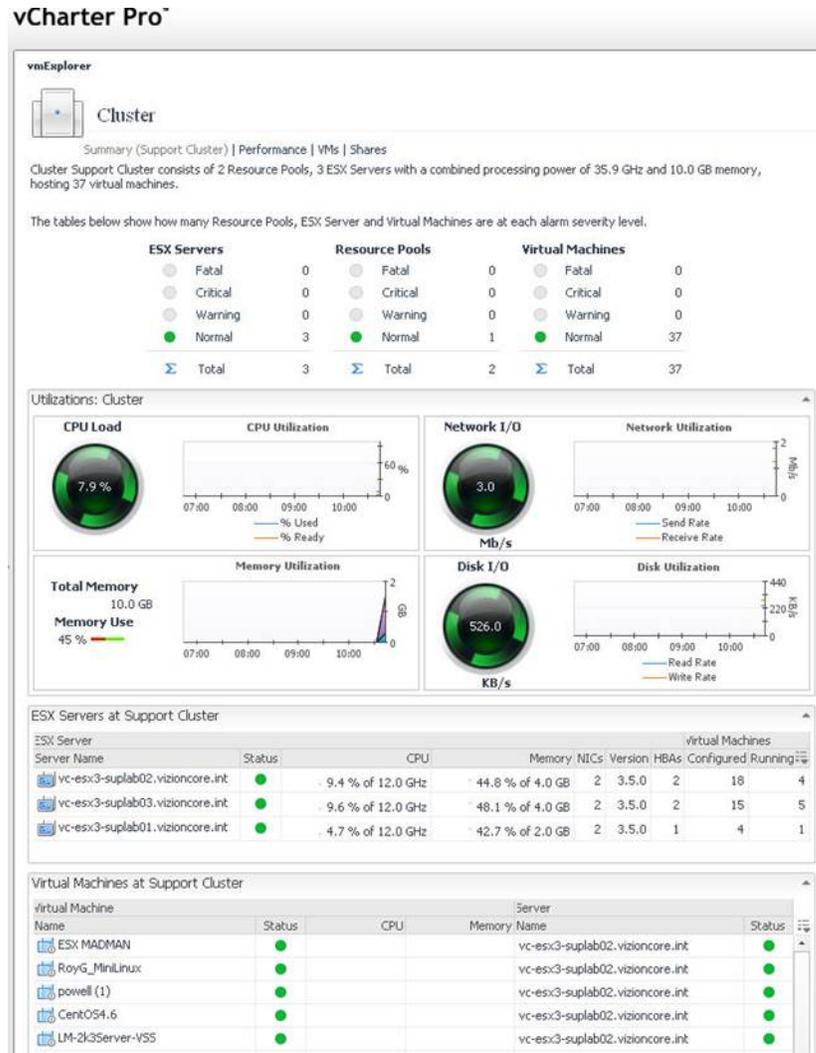
Single view of key performance metrics aggregated at Cluster Level

- CPU (%Ready / % Used)
- Network I/O
- Disk I/O
- Memory

Summary View of Cluster objects

- Resource Pools
- Hosts
- VMs

Unique Drill-Down ability to quickly isolate issues.



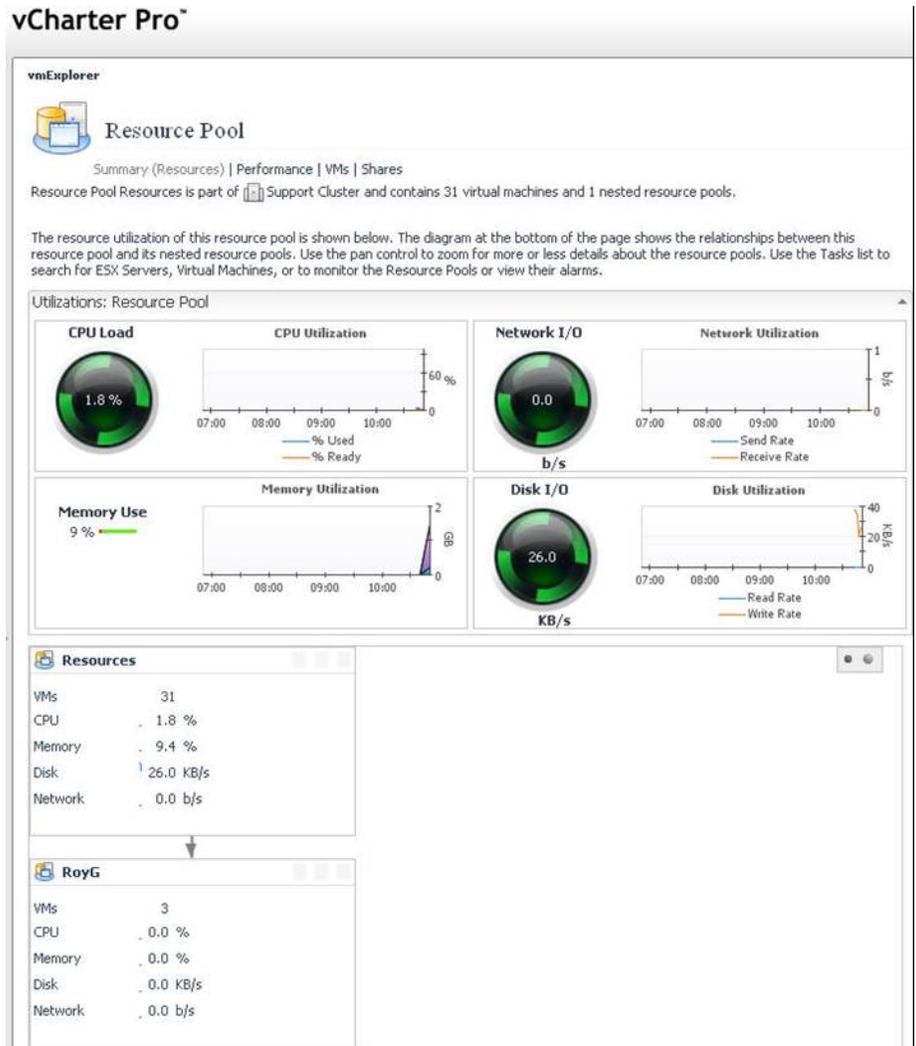
Resource Pools View

Single view of key performance metrics

- CPU (%Ready / % Used)
- Network I/O
- Disk I/O
- Memory

Summary View of Resource Pool Performance

Unique Drill-Down ability to quickly isolate issues.



Consider the value you would get using a centralized monitoring system. Performance metrics and trend analysis at your fingertips lets you plan and manage growth with concrete current data – no guesstimates needed.

Goldilocks threshold levels are set by you and tracked automatically -- with alarms that alert you to problems before they get out of hand.

Plus with most systems you get additional functionality such as migration modeling and charge back capability among others.

Just makes sense if you want to maximize your virtualization investment

Summary

Virtualization while simple in concept is quite complex. It demands rigorous management to optimize your ROI.

Finding and maintaining your VMware Goldilocks Zone in an on-going and consistent way will deliver additional value to your overall operation. While you can do this with existing functionality within VMware – it may not be practical as you expand the number of VMs and ESX Hosts.

Alternatives exist that automate and facilitate successful VMware management. Most – if not all of them -- have non-invasive trial versions available. Finding bandwidth to try a new product is often difficult – yet the rewards of finding one that meets your needs will far exceed the time spent evaluating.

About Mosaic Technology

Mosaic Technology, a privately held company, provides IT infrastructure solutions for companies worldwide through three divisions: Value add solutions, IT Asset Management and Recovery, and Used and Refurbished IT Hardware. Based in Salem, NH, Mosaic has sales and technical offices throughout the United States and associates worldwide.

We have built a Virtualization practice based on a blend of Professional Services, and business relationships that have allowed us to successfully implement end-to-end virtualization projects in companies of all sizes.

Mosaic provides IT infrastructure solutions and Professional Services in Virtualization, Email Management, Data Storage and Data Management. We have well established partnerships with leading technology companies such as VMware, Vizioncore, Exagrid, Data Domain, CommVault, Dell, EqualLogic, Overland, EMC, Hitachi, and Symantec.

Contact:

Don Basler
don_basler@mosaictec.com
603-898-5966

Mosaic Technology Corporation
41 Northwestern Drive
Salem, NH 03079
Mosaic Technology
www.mosaictec.com