

# Análisis y Gestión del Rendimiento

## Basic Performance Troubleshooting

*Alejandro Solana – Iberia SE Manager*



Confidential

vmware®

© 2009 VMware Inc. All rights reserved

# Agenda

---

- **Troubleshooting de Rendimiento**
  - Common Issues
  - El Rol del HW y el SW
  - Herramientas para chequear el Rendimiento (vCenter, esxtop, vscsistats)
  - Analizar el uso de los Recursos (cpu, mem, net, storage)

# Resources

---

- **VMware's Performance Technology Pages**
  - <http://vmware.com/technical-resources/performance>
- **VMware's Performance Blog – VROOM!**
  - <http://blogs.vmware.com/performance>
- **Performance Community Forum**
  - <http://communities.vmware.com/community/vmtn/general/performance>
- **Mark's Performance Ramblings**
  - <http://www.virtualizationeh.ca>, @vmmarka (Twitter)
- **VMware Education**
  - vSphere: Manage for Performance Class
- **VMware Performance Service Offering**
  - Custom Engagement: Performance Health Check

## ¿Habéis escuchado ...?

---

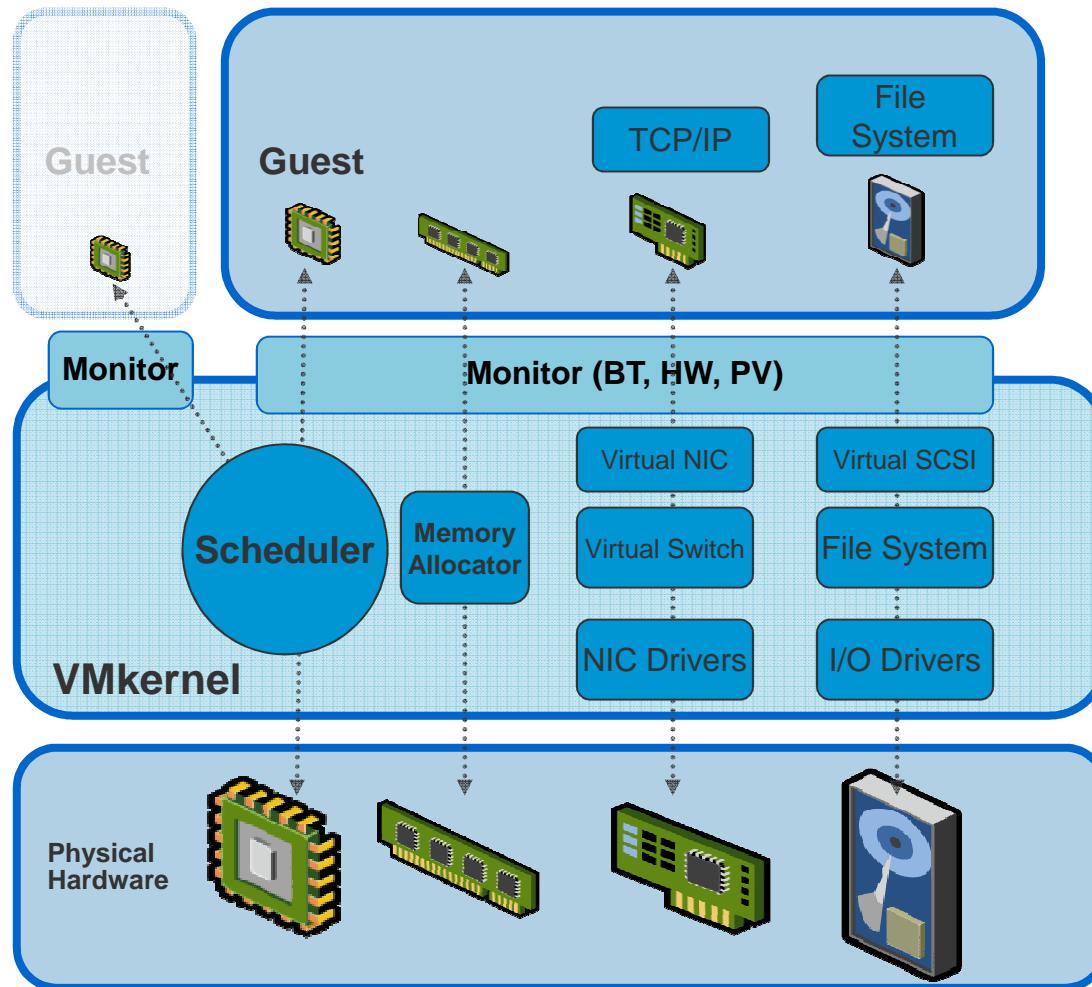
- *Mi VM va muy lenta, no sé que hacer*
- *Esta aplicación “obviamente” no se puede virtualizar*
- *Intenté incrementar la memoria y las CPUs pero no funcionó*
- *Les asigné más Shares pero no ayudó*
- *Los Administradores de la SAN dicen que el almacenamiento está bien*
- *Mis VMs son un 50% más lentas que en físico.*

## Issues de Rendimiento más habituales

---

- #1 No Comparar “manzanas con manzanas”
- #2 Almacenamiento, Almacenamiento, Almacenamiento (85% de los problemas)
- Errores de “planificación” y “gestión” de la capacidad disponible
- Errores de Configuración
  - Server power set to “max performance” in BIOS
  - Not following application best practices documents
- Over-provisioning
  - Poor use of SMP (eg: 8 vCPU on dual socket host)
- Expectativa de que hardware y software antiguo se comporten igual que en físico
  
- Falta de Training

# VMware ESX Architecture



CPU is controlled by scheduler and virtualized by monitor

Monitor supports:

- BT (Binary Translation)
- HW (Hardware assist)
- PV (Paravirtualization)

Memory is allocated by the VMkernel and virtualized by the monitor

Network and I/O devices are emulated and proxied through native device drivers

# El Rol del Hardware

# HW Assist CPU Architecture

---

- Intel
  - Intel-VT
  - Disponible desde 2006
- AMD
  - AMD-V
  - Disponible desde 2006
- **Previamente a estas tecnologías hardware se utilizaba “binary translation”.**
- **Las primeras implementaciones del soporte hardware eran más lentas que la binary translation.**

# Hardware Assist (HWmmu) Arquitecturas de Memoria

---

- **Intel**

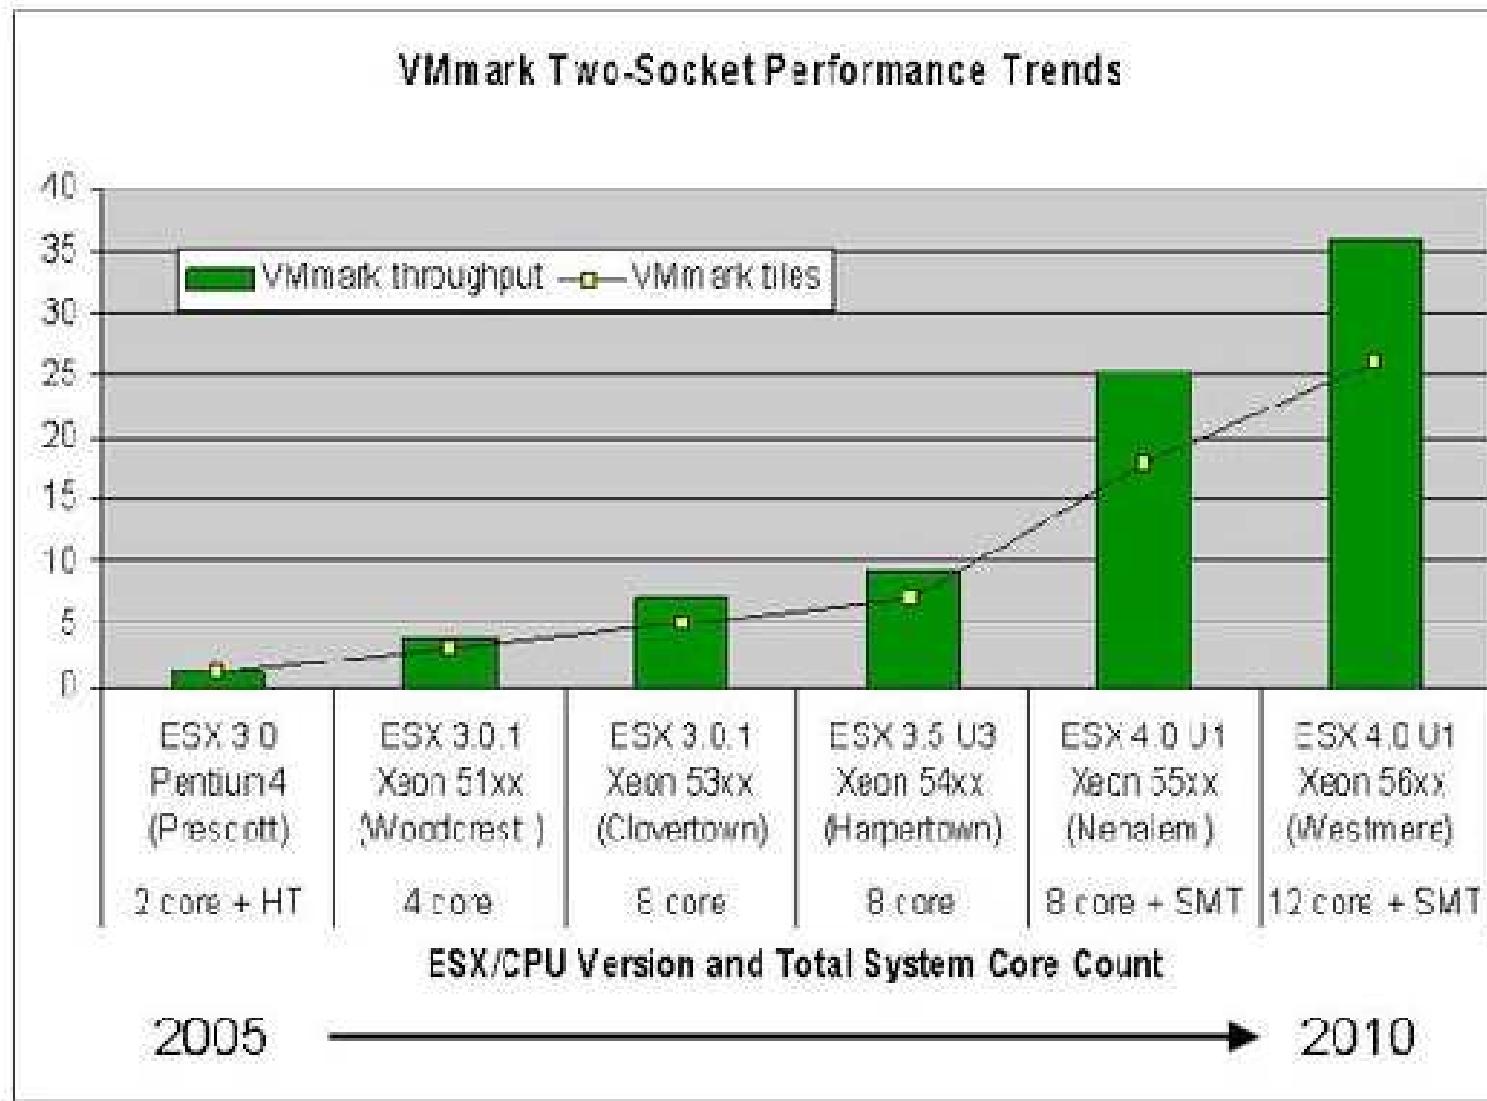
- Extended Page Tables (EPT)
- Disponible desde: 2009
- Supported en ESX4.0 +
- Nehalem or better

- **AMD**

- Rapid Virtualization Indexing (RVI)
- Disponible desde: 2008
- Supported en ESX3.5 +
- Shanghai or better

- **Anteriormente a estas tecnologías hardware – se utilizaba “shadow paging” (o Swmmu).**

# Tendencias de Rendimiento



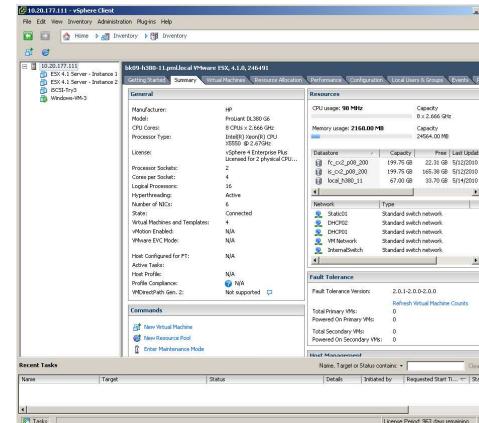
Source: <http://blogs.vmware.com/performance/2010/05/surveying-virtualization-performance-trends-with-vmmark.html>

# Herramientas de Troubleshooting

# Herramientas de Troubleshooting Tools

## ■ vCenter (Alto Nivel)

- Datos de Rendimiento histórico (\*comprobar los niveles de las estadísticas)
- Consolidated metrics for all hosts / datastores in environment



## ■ vscsiStats (info de almacenamiento)

- Métricas detalladas
- Seek distance, IO size, Latency
- Displays Histograms

```
[root@bk09-h380-11:/usr/lib/vmware/bin]# ./vscsiStats -w 4968 -i 8208 -p seekDistance
Histogram: distance (in LBNs) between successive commands for virtual machine worldGroupID : 4968, virtual disk handleID : 8208 {
min : -16729160
max : 16388622
mean : 728
count : 18014
{
4222 (<= -500000)
217 (<= -100000)
27 (<= -5000)
22 (<= -10000)
1 (<= -5000)
2 (<= -1000)
0 (<= -500)
0 (<= -128)
0 (<= -64)
1 (<= -32)
0 (<= -16)
0 (<= -8)
0 (<= -6)
0 (<= -4)
0 (<= -2)
```

## ■ esxtop / resxtop (tactico)

- A nivel de Host ESX/ESXi
- Datos de Rendimiento en tiempo real

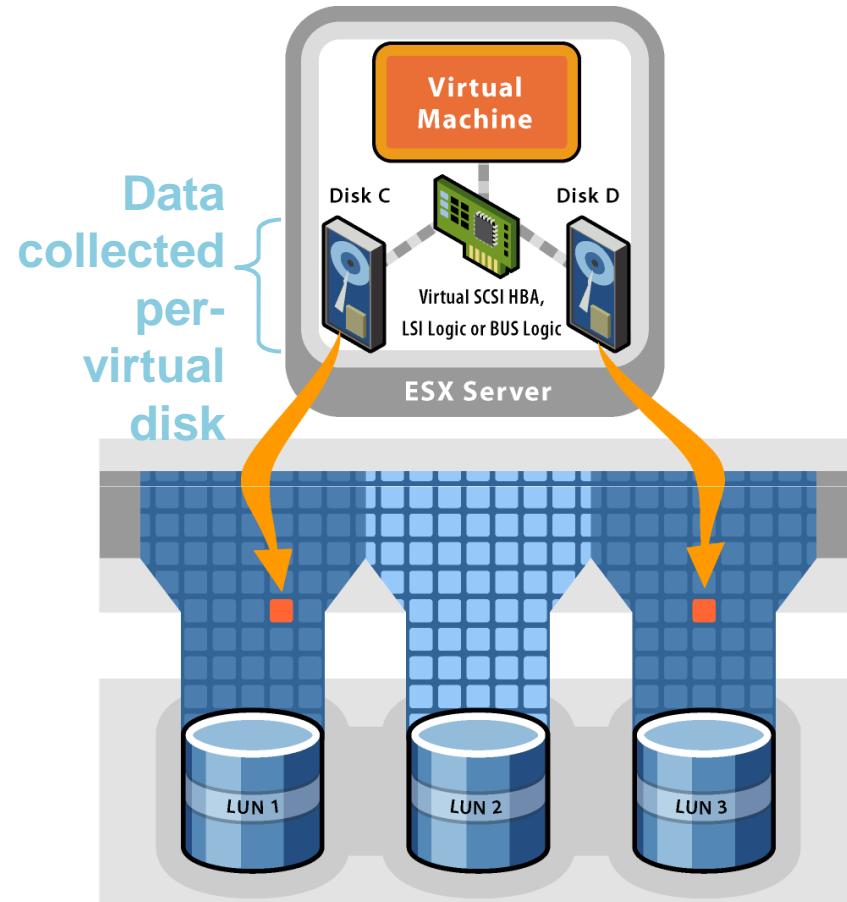
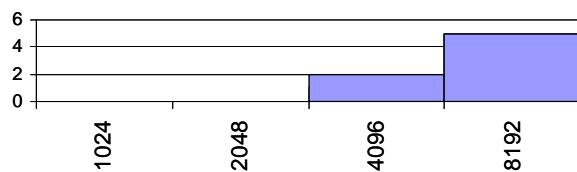
```
3:31:26am up 1 day 23:52, 148 worlds; CPU load average: 0.00, 0.00, 0.00
PCPU USED(%): 58.0 1.0 0.0 0.0 0.0 0.0 0.3 0.5 0.1 0.0 0.2 0.0 0.0 AVG: 3.7
PCPU UTIL(%): 100 100 0.1 100 0.8 100 0.6 0.7 1.7 100 1.0 100 100 100 100 AVG: 62
CORE UTIL(%): 100 0.2 100 0.7 1.4 1.7 0.2 100 AVG: 37
GCPU(%): 0 us, 100 sys, 0 id, 0 wa; cs/sec: 427

```

ID	GID	NAME	NWLD	%USED	%RUN	%SYN	%WAIT	%RDY	%IDLE	%WRLP
1	1	idle	16	635.95	800.00	0.00	0.00	5.20	0.00	0.89
11	11	console	1	78.56	90.31	0.02	0.00	0.00	0.00	0.12
50	50	storageRM.4295	1	1.80	2.51	0.02	87.76	0.02	0.00	0.02
67	67	Windows-VM-3	5	0.96	2.70	0.16	448.26	0.01	89.27	0.01
7	7	helper	80	0.03	0.09	0.00	722.64	0.05	0.00	0.00
2	2	system	8	0.00	0.00	0.00	722.41	0.00	0.00	0.00
8	8	drivers	10	0.00	0.00	0.00	905.32	0.00	0.00	0.00
9	9	vmotion	4	0.00	0.00	0.00	361.00	0.00	0.00	0.00
13	13	tcpipmod	10	0.00	0.00	0.00	1.68	0.00	0.00	0.00
48	48	FT	1	0.00	0.00	0.00	90.29	0.00	0.00	0.00
49	49	vobd.4294	6	0.00	0.00	0.00	541.74	0.00	0.00	0.00
51	51	sensord.4296	1	0.00	0.00	0.00	90.28	0.00	0.00	0.00
52	52	net-cdp.4302	1	0.00	0.00	0.00	90.30	0.00	0.00	0.00
53	53	net-lbt.4303	1	0.00	0.00	0.00	90.22	0.00	0.00	0.00
56	56	vmkiscsifd.4305	2	0.00	0.00	0.00	180.39	0.00	0.00	0.00
57	57	vmware-vmkauthd	1	0.00	0.00	0.00	90.20	0.00	0.00	0.00

# Detalles de vscsiStats

- **vscsiStats caracteriza el IO por cada disco virtual**
  - Permite separar cada tipo de carga en su propio contenedor y observar tendencias
- **Histograms solo se recogen si se activa; para evitar overhead**
- **Tecnica:**
  - Por cada IO request vamos insertando valores en los histogramas
  - E.g., size of I/O request → 4KB



## Datos vscsiStats

---

- Las Distribuciones Read/Write están disponibles en los histogramas  
Distributions are available for our histograms
  - Ratio promedio de Read/Write?
  - Son las Escrituras más pequeñas o mayores que las lecturas para una determinada carga?
  - Son las Lecturas más secuenciales que las lecturas?
  - Que tipo de I/O está incurriendo en más latencia?
- En realidad, el problema no es saber que preguntar
  - Recoger los datos, y analizarlos
- I/O Size
  - All, Reads, Writes
- Seek Distance
  - All, Reads, Writes
- Seek Distance Shortest Among Last 16
- Outstanding IOs
  - All, Reads, Writes
- I/O Interarrival Times
  - All, Reads, Writes
- Latency
  - All, Reads, Write

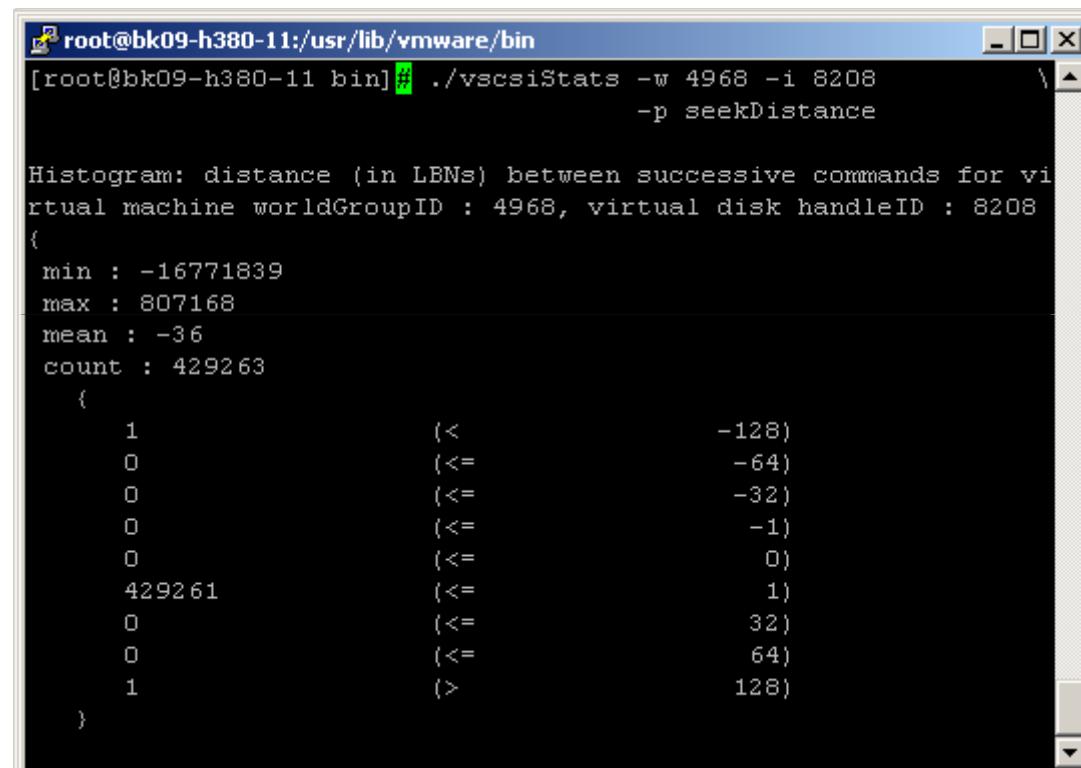
# vscsiStats in use

- vscsiStats is on the ESX console at /usr/lib/vmware/bin

- List Virtual Machines

- Start vscsiStats

- Display Histograms



```
root@bk09-h380-11:/usr/lib/vmware/bin
[root@bk09-h380-11 bin]# ./vscsiStats -w 4968 -i 8208
                                     -p seekDistance

Histogram: distance (in LBNs) between successive commands for virtual machine worldGroupID : 4968, virtual disk handleID : 8208
{
    min : -16771839
    max : 807168
    mean : -36
    count : 429263
    {
        1           (<          -128)
        0           (<=         -64)
        0           (<=         -32)
        0           (<=          -1)
        0           (<=           0)
        429261     (<=           1)
        0           (<=          32)
        0           (<=          64)
        1           (>          128)
    }
}
```



# **Analizando la dimensión de los Recursos**

# CPU: Common Issues

---

- Una VM demasiado “grande” en una plataforma demasiado “pequeña”
  - Estimación: 1-4 vCPU en dual socket hosts, 4-8 vCPU en quad socket hosts.  
Cuantas más oportunidades de planificación mejor
  - Con 8 vCPU, ensure vSphere 4.1, check HIMP (changes in MN)
- Limits/Reservations Accidentales
  - Siempre comprobar los reservations/limits en las VMs y resource pools: as many people “forget” or too many cooks in the kitchen
- Esperar obtener ratios de consolidación elevados con cargas exigentes
  - Según virtualizamos cargas más exigentes que utilizan una mayor capacidad de proceso debemos establecer expectativas razonables en el ratio de consolidación.

# CPU: Key Indicators

---

## ■ Ready (%RDY)

- % intervalo que una vCPU esta preparada para ser planificada en un procesador fisico, pero no pudo serlo debido a la contención
- Investigación a partir de: 10%

## ■ Co-Stop (%CSTP)

- % inteverlo en el que una vCPU en una VM SMP “deja” de ejecutar, para que otra vCPU en la misma VM pueda ejecutar para “actualizarse” y asegurar que las dos vCPUs to estén demasiado desincronizadas
- Investigación a partir: 3%

## ■ Max Limited (%MLMTD)

- % intervalo que una VM estaba preparada para ejecutar pero no fue planificada al violar los límites de CPU establecidos.
- Investigación a partir de: 0%

# Memory: Common Issues

---

## ■ Utilizando SWmmu

- Si hay opción utilizar hardware memory assist (HWmmu) para reducir el overhead con reduces virtualization overhead for memory intensive workloads (ej: Tier 1)

## ■ No monitorizar el ballooning o el swapping de un host

- Ballooning es un aviso de que se puede empezar a paginar
- La paginación es un problema de rendimiento.

## ■ No monitorizar el swapping en la VM

- Asignar poca memoria a la VM

## ■ TPS y Large Memory Pages

- Los datos de resumen de vCenter muestran una utilización elevada de memoria ya que la memoria en la VM es respaldada por 2Mb physical mem pages 1st, y luego TPS cuando se requiere
- Bad capacity planning metric

## ■ Eliminar el balloon driver y deshabilitar TPS

# Memory: Key Indicators

---

## ■ **Balloon driver size (MCTLsz)**

- La cantidad total de memoria física que la VM reclama a través del balloon driver
- Investigación a partir de: 1

## ■ **Swapping (SWCUR)**

- La cantidad actual de memoria física que la VM hace swap out al fichero de swap del kernel de ESX.
- Investigación a partir de: 1

## ■ **Swap Reads/sec (SWR/s)**

- La tasa de transferencia en el que la memoria de la máquina lee del swap de disco.
- Investigación a partir de: 1

## ■ **Swap Writes/sec (SWW/s)**

- La tasa de transferencia en el que la memoria de la máquina hace swap a disco.

## Network: Common Issues & Key Indicators

---

- **Transmit Dropped Packets (%DRPTX)**
  - El porcentaje de paquetes enviados “dropped”.
  - Investigación a partir de: 1
- **Receive Dropped Packets (%DRPRX)**
  - El porcentaje de paquetes recibidos “dropped”.
  - Investigacion a partir de: 1

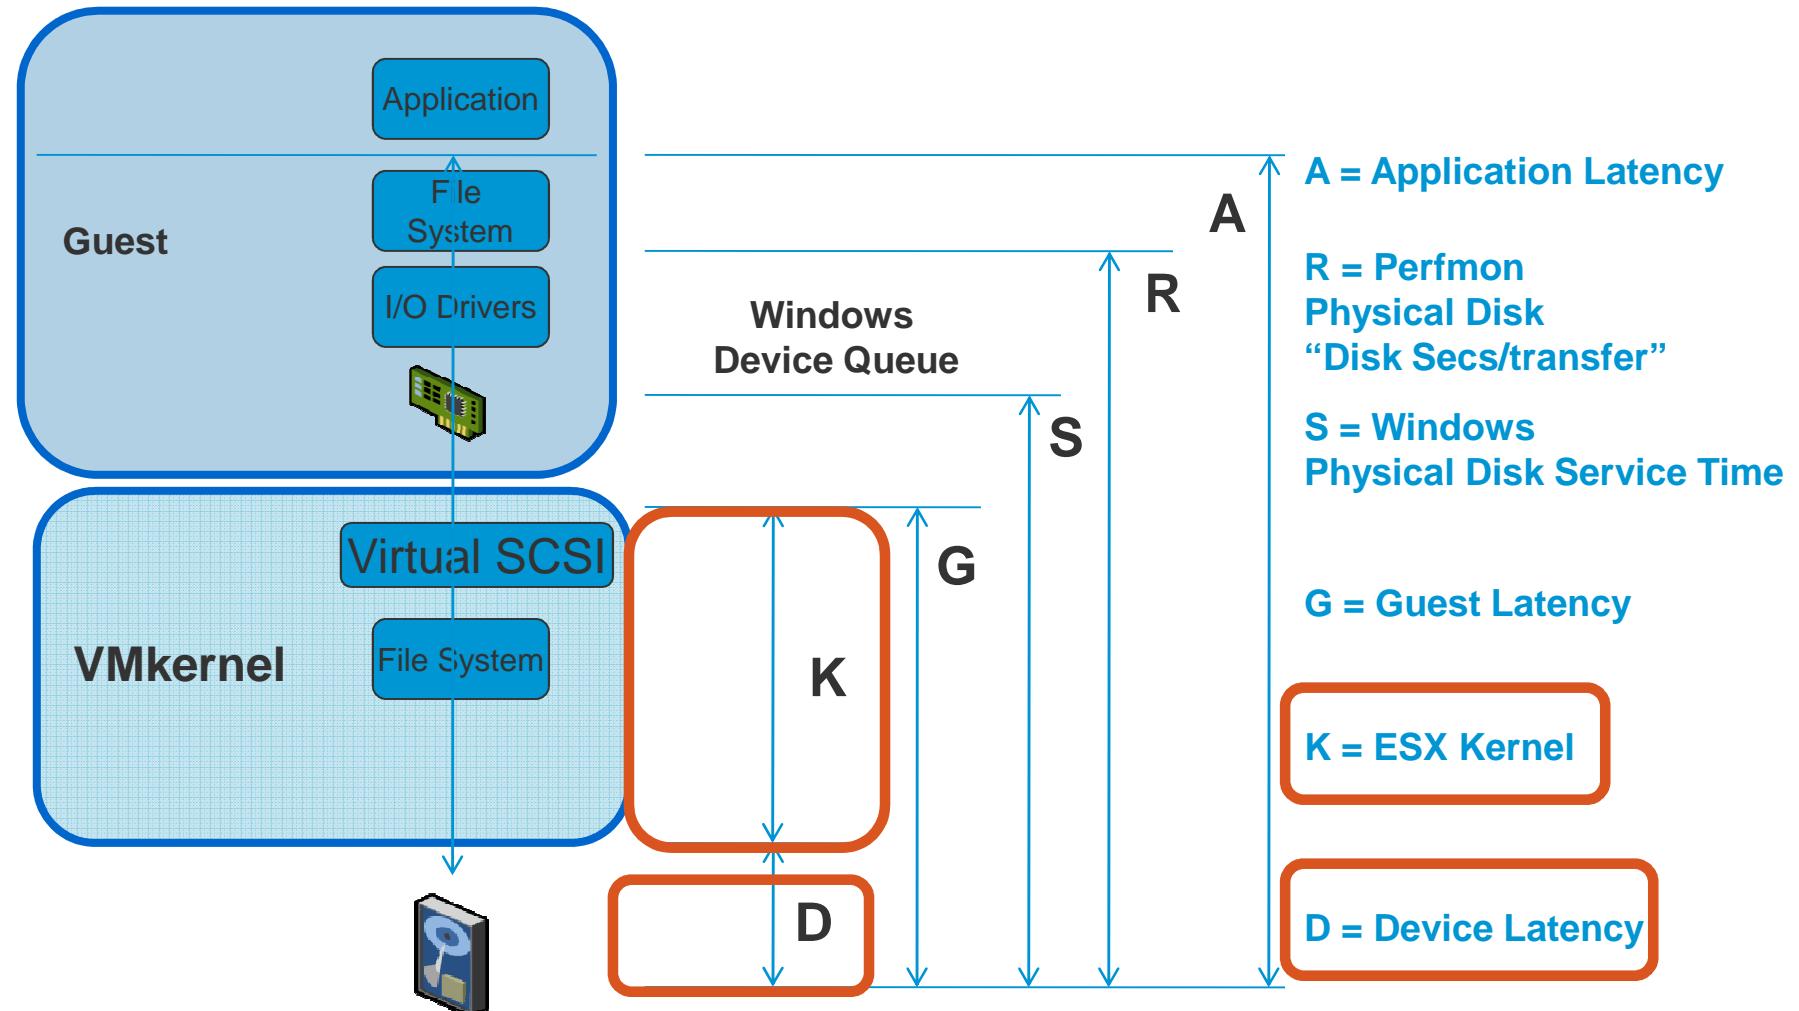
La red es raramente un recurso que cause problemas de rendimiento a las aplicaciones

# Storage: Common Issues

---

- **Demasiado foco en la capacidad**
  - IOPS & latencia de disco son importantes – la capacidad no tanto
- **No hablar en “almacenamiento” al equipo SAN**
  - VM Admins deben aprender a hablar de terminología de almacenamiento para romper silos (y vice-versa en el SAN team)
- **Sin coordinación entre los equipos SAN/VM teams**
  - Se requiere cooperación para crear/proveer la configuración adecuada de disco a cada carga. Trabajar a ciegas puede servir en cargas poco críticas pero no en Tier1.
- **Disk Alignment**

# Storage: Key Indicators



# Storage: Key Indicators

---

## ■ Kernel Latency Average (KAVG)

- Este contador controla la latencia de IO a través del Kernel
- Investigar a partir de: 1ms, típicamente microsegundos

## ■ Device Latency Average (DAVG)

- Esta es la latencia que se observa a nivel de driver de dispositivo. Incluye el roundtrip entre la HBA y el almacenamiento.
- Investigación a partir de: 15-20ms, lower is better, some spikes okay

## ■ Aborts (ABRT/s)

- Número de comandos abortados por segundo.
- Investigación a partir de: 1

**So far... so good?**

---

**Y SI PUDIERAMOS GESTIONAR EL RENDIMIENTO Y  
LA CONFIGURACION DE UNA FORMA MAS  
AUTOMÁTICA?**

# VMware vCenter Operations Standard

*Real-time Performance Management for VMware Administrators*



Confidential

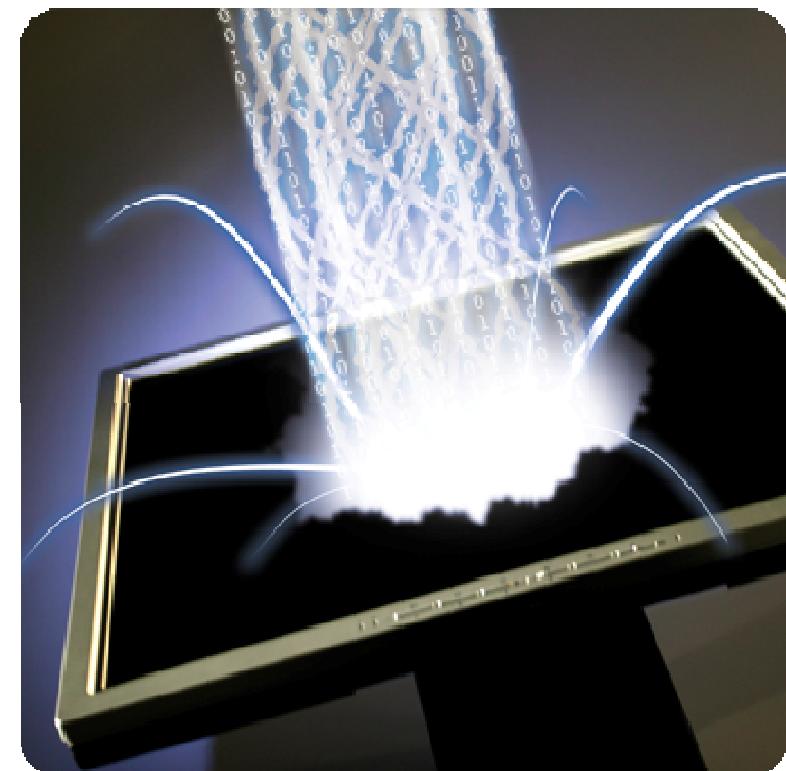
vmware®

© 2009 VMware Inc. All rights reserved

# Why vCenter Operations Standard?

---

- **80% of VMware admin time spent isolating performance problems**
  - “1st generation” green-yellow-red static threshold reporting insufficient and too complex to use
  - Point solutions only address a subset of issues
  
- **VMware administrators have two conflicting goals**
  - Maximize ROI by increasing VM density
  - Ensure required capacity for business growth and other changes in real-time
  - Ensure that virtual component performance supports required application performance



# VMware vCenter Operations Standard Basics

---

- **Clear and quick way to identify VMware performance problems**
- **Easy to use for VMware Administrators**
  - Deeply integrated as a vCenter pane
  - Intuitive screens guide users to issues needing attention
  - Automatically collects data from vCenter
    - Time-series performance data, topological relationships and configuration change events
- **VMware vCenter Operations Standard business benefits**
  - Increased performance for end users of business applications and services
  - Reduced infrastructure costs through increased VM to ESX density
  - Reduced VM administration costs and optimized VMware admin productivity

# Understanding your Virtual Environment - Workload

---

## ■ Workload Measures

- Demand for resources vs. Resources currently used
- Result is a percentage of Workload
  - Low number is Good – Object has the resources it needs
  - Can go above 100% - Object is “Starving”



## ■ Workload summarized across critical resources

- CPU
- Storage I/O
- Network I/O
- Memory (VM and ESX Allocation)

## ■ Workload Details View

- Detailed understanding of the lacking resource and associated metrics
- View the state of the Peer and Parent Objects and troubleshoot
  - Am I a victim or a villain?
  - Is this a population problem?
  - Should we move the VM?
  - A Configuration issue?
  - Lack of resources?
  - Virtual infrastructure is fine. OS or application issue?

# Understanding your Virtual Environment - Health

---

## ■ Health Measures

- How normal is this object behaving: 0-100 (Higher is Healthier)
- Learns dynamic ranges of “Normal” for each metric
- Learns patterns of behavior and identifies metric abnormalities
- Lower the health the more abnormalities



## ■ Once a virtual element Health problem is identified

- Single screen provides details on problem based on behavioral understanding of the element
- Points to the Root Cause metrics to help you troubleshoot
- Eliminates 100s of clicks and memorization of many metric behaviors that 1st generation monitoring tools require

## Important Note

Low Health does not imply a problem. It tells you that the object is acting differently than normal.

## ■ Health and Workload together tell you a lot

- Workload High & Health High – Normal Behavior for this timeframe
- Workload High & Health Low – Something is amiss!

# Understanding your Virtual Environment - Capacity

---

## ■ Capacity Measures

- How much time do you have left before an object runs out of resources?
- Based on a 0-100 scale – Higher the number the longer you have
- Thresholds User Configurable
  - 30 Days Left = RED
  - 60 Days Left = Orange
  - Etc.



## ■ Capacity measured for critical resources

- CPU
- Storage I/O
- Network I/O
- Memory

## ■ Capacity Details View

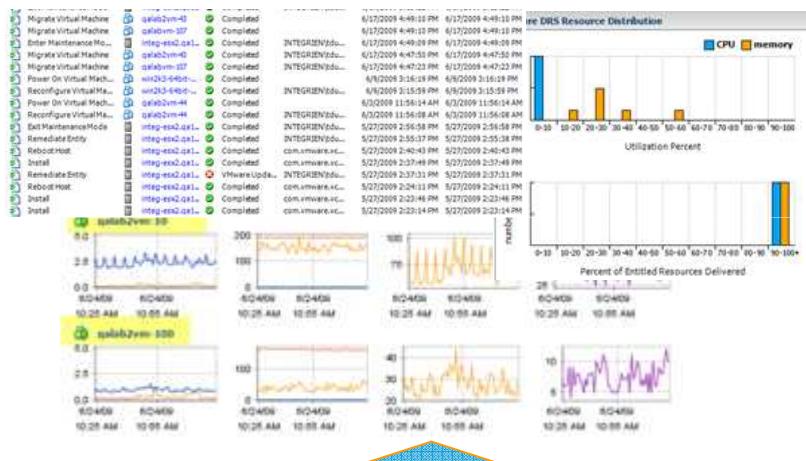
- Shows the chart and trend for each of the above resources
- Denotes current state
- Projected breach point and days left

# **Business Benefits**

# Increased Visibility

P

- Lack of holistic VC environment view
- Can't determine state of all elements (clusters, hosts, guests) at once
- Overwhelming details obscure valuable information.



BEFORE

VALUE

S

- Single pane of glass
- All VC data contextually consolidated
- One click to any detail
- Filters on “all” “normal” and “problem”
- Searches on any string.



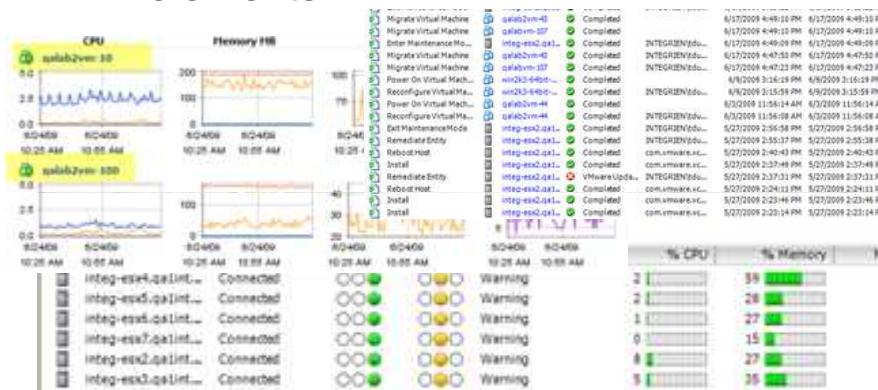
AFTER

- Visibility, comprehension of virtualized environment in one screen
- Better product usability
- Visually isolate problems via a “HUD” for vCenter
- Unnecessary details hidden until necessary.

# Reduced Complexity

P

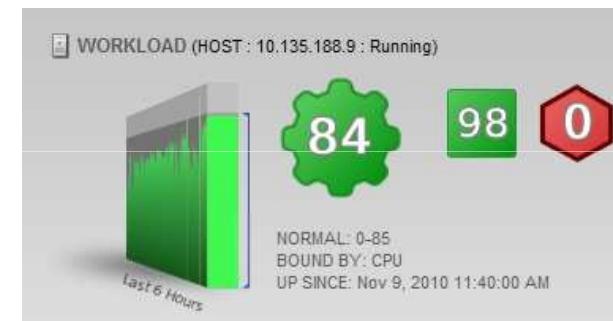
- Administrators blind to brewing problems
- Too much data, too many clicks
- Preset thresholds, many details
- Impossible to understand health of elements



BEFORE

S

- Provide a single measure of normality across all virtualized elements – Health
- Automatically aggregate, correlate states of 100s of metrics into two scores for each element – Health and Workload



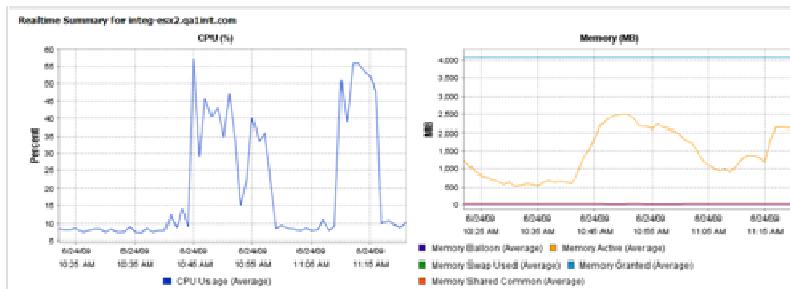
AFTER

- Reduce complexity of usage
- Remove guesswork, provide clarity into the environment
- Speed up MTTR
- Enable administrators to do more with less.

# Understand Normal Metric Behavior

P

- Unable to understand normal range of metrics
  - Is 65% usage normal for an hour, day, week or month?
  - Or, is it the beginning of a problem?



BEFORE

VALUE

AFTER

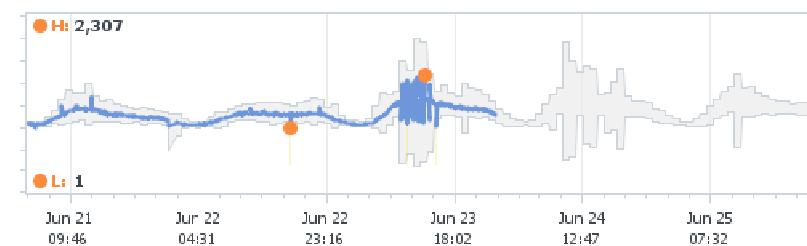
- Understand metric behavior based on history
- Project forward future behavior hours or days in advance
- Remove guess work and confusion, clarify expectations
- Equivalent of 10 people watching, measuring and adjusting system constantly.

S

- Visibility into normal operation of every metric in VC
- Continuous, automatic learning of normal behavior



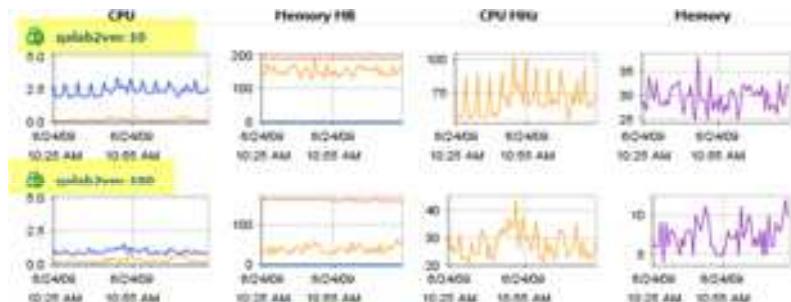
integ-esx3.qa1int.com: Memory Active



# Workload Optimization

P

- VC unaware of affinities and workload profiles of all VMs
- Only understands raw resource consumption



BEFORE

VALUE

- Increase density of VMs per ESX
- Optimize use of resources
- Consistent and maximized ESX workloads

S

- Calculates and stores workload profile of each ESX
- Increase density by matching opposite VM behaviors on an ESX
- Ensure smooth, consistent use of resources



AFTER

# Understand Impact of Change

P

- Change is common and necessary in VM environments
- Change can lead to degradation in performance

S

- Changes and events mashed on health chart for every element
- Easier to see impact of change and before and after performance

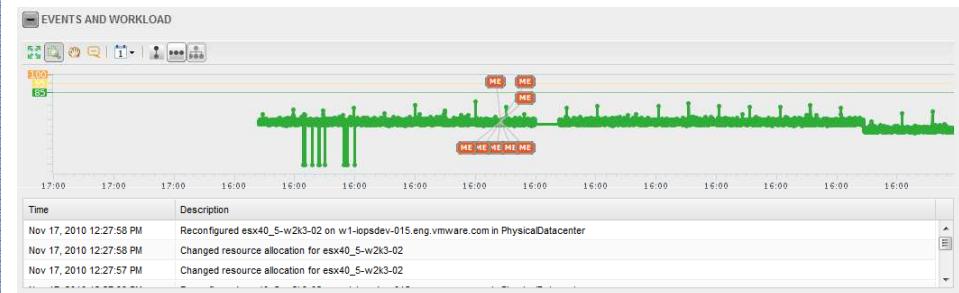
Task Details

Name: **Reconfigure Virtual Machine** Target: [galab2vm-111](#) Initiated by: INTEG

Related Events:

4/13/2009 5:40:24 PM, Reconfigured galab2vm-111 on integ-esx4.qa1int.com in Integri...

4/13/2009 5:40:22 PM, Task: Reconfigure Virtual Machine



BEFORE

VALUE

AFTER

- Immediate visibility into impact of change
- Visual correlation to component's health
- Admin can immediately determine if change had positive (expected) or negative (unexpected) effect on the element

# Multidimensional Analysis

P

- Which of my many Hosts have high levels of CPU Ready contention but low memory usage?



S

- Slice, dice, visualize entire environment by any of 100s of VC-collected metrics



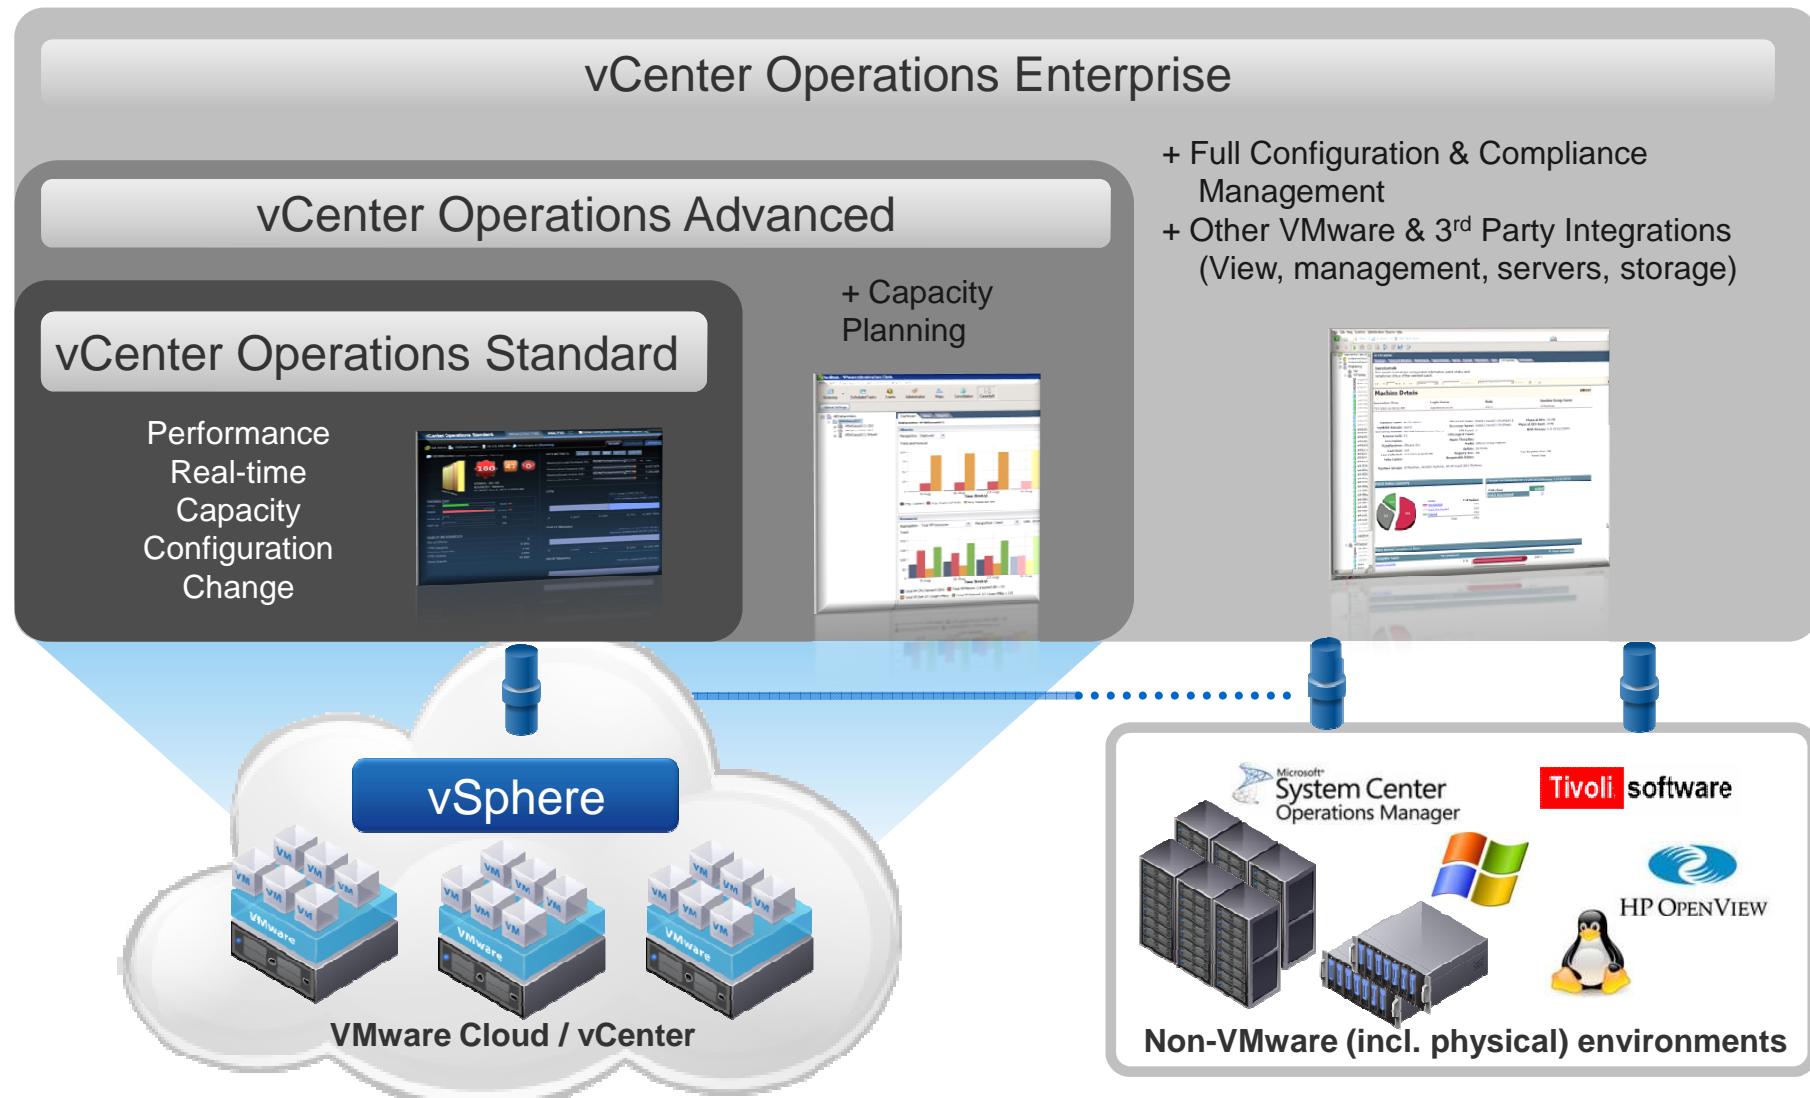
BEFORE

VALUE

AFTER

- Full Business Intelligence like capabilities
- Slice and dice historical collected data across any dimension
- Visualize results in heat maps, single click drill down to resource details.

# VMware vCenter Operations Editions



# Understanding the vCenter Operations Editions

	vCenter Operations Standard Edition	vCenter Operations Enterprise - Stand-Alone
Scope	<b>Data Sources</b> vCenter x 1	<ul style="list-style-type: none"> <li>Any 3<sup>rd</sup> party monitoring tools' time series data</li> <li>Change events</li> <li>Multiple vCenter Servers</li> </ul>
Function	<b>Objects</b> <u>vCenter Objects (i.e.)</u> <ul style="list-style-type: none"> <li>Data Centers</li> <li>Clusters</li> <li>ESX Hosts</li> <li>Datastores</li> <li>VMs x 500</li> </ul>	<u>Unlimited Scope (i.e.)</u> <ul style="list-style-type: none"> <li>Applications</li> <li>Network Infrastructure</li> <li>Storage</li> <li>Hosts (ESX, Win, Linux, etc)</li> <li>VMs</li> </ul>
Scope	<b>Users</b> Infrastructure (e.g. VI Admins)	Operations, Infrastructure, Application Teams, Business Owners, CxOs
Function	<b>Dynamic Thresholds</b> <b>Performance Root Cause</b> <b>Proactive Alerting</b> <b>Customizable Dashboards</b> <b>Notifications</b>	<b>Yes</b> <b>Yes</b> <b>No</b> <b>No</b> <b>No</b>

# Screenshots

# Performance dashboard based on self-learning analytics

The screenshot shows the vCenter Operations Standard interface in Internet Explorer. The top navigation bar includes File, Edit, View, Favorites, Tools, Help, and a Favorites section for vCenter Operations Standard. Below the navigation is a main menu with tabs for INFRASTRUCTURE (selected), ANALYSIS, and Home/Configuration/Help/About/Signout.

A green callout box points to a group of three icons: WORKLOAD (blue gear), HEALTH (green square), and CAPACITY (grey hexagon). A blue callout box points to a row of five yellow/green gear icons with scores: 56 (yellow), 28, 21, 22, and 25 (green). A blue callout box points to a series of gear icons below, highlighting some in orange and grey, with a text overlay: "Highlights resources that are deviating from ‘normal’ behaviour".

On the left sidebar, there are collapsed sections for vCENTER SERVER SYSTEMS (1 of 1), DATACENTERS (5 of 5), CLUSTERS (2 of 2), ESXs (13 of 13), and VMs (74 of 74).

At the bottom right is the VMware logo.

Visualize environment performance in three unique dimensions

Simple, actionable scores that indicate overall performance

Highlights resources that are deviating from “normal” behaviour

# Get “At-a-glance” insights into performance issues

The screenshot shows the vCenter Operations Standard interface in a web browser. The top navigation bar includes File, Edit, View, Favorites, Tools, Help, and a Favorites section for vCenter Operations Standard. The main menu tabs are INFRASTRUCTURE (selected), ANALYSIS, and Home | Configuration | Help | About | Signout.

The left sidebar displays navigation links for vCENTER SERVER SYSTEMS (1 of 1), DATACENTERS (5 of 5), CLUSTERS (2 of 2), ESXs (13 of 13), and VMs (74 of 74). The right side features a "Performance scores" summary card with icons for WORKLOAD (71), HEALTH (89), and CAPACITY (100), along with buttons for Details, Scoreboard, and All Metrics. Below this is a "SHOW STATUS" section for the host 10.135.188.119, showing status icons for each category.

A large central area displays a grid of colored gears representing performance metrics. A specific gear with the value "28" is highlighted with a dashed blue border. A callout bubble labeled "Visualize impact" points to this gear. Another callout bubble labeled "Performance scores" points to the summary card. A third callout bubble labeled "Details" for further analysis points to the "Details" button in the status summary.

# Drill down into problem source

The screenshot shows the vCenter Operations Standard interface in a web browser. The main navigation bar includes File, Edit, View, Favorites, Tools, Help, and a Certificate Error notice. Below the navigation bar, there are tabs for vCenter Operations Standard, INFRASTRUCTURE, ANALYSIS, Home, Configuration, Help, About, and Signout.

The main content area displays performance metrics for an ESX host (10.135.188.119). A large orange gear icon indicates a score of 71, with a green hexagon showing 89 and a green hexagon showing 100. A callout bubble says "Stress caused by net I/O".

A section titled "WORKLOAD" shows resource usage: CPU (24%), MEM (53%), DISK I/O (1%), and NET I/O (71%). A callout bubble says "Quickly identify problem source".

The "KEY METRICS" tab is selected, showing various metrics like Memory Usage / Usable(%), Disk/Disk Write Latency (ms), CPU Usage / Core Utilization(%), and CPU Usage / Idle (ms). A callout bubble says "Key metrics of interest based on continuous learning of 'normal' behavior".

Below the metrics are two horizontal bar charts for CPU and Memory usage.

# Correlate cause-and-effect of the problem

The screenshot shows the vCenter Operations Standard interface in a web browser. The main navigation bar includes File, Edit, View, Favorites, Tools, Help, and a Favorites section for vCenter Operations Standard. The top menu bar has links for Home, Configuration, Help, About, and Signout.

The main content area displays the following sections:

- RELATED OBJECTS:** Shows a PARENT OBJECT (DATACENTER : CIQDataCenter) with 28 objects, all marked as NORMAL (6-48) and BOUND BY: Memory.
- PEER OBJECTS OF 10.135.188.119 (4 of 4 ESXs):** Shows four objects, with the first one marked as orange (warning).
- CHILD OBJECTS OF 10.135.188.119 (14 of 14 VMs):** Shows 14 objects, with several marked as orange (warning).

A green callout bubble points to the peer objects section with the text: "Check health of related objects in the hierarchy".

**EVENTS AND WORKLOAD:** A graph shows CPU usage over time from 03:00 to 08:30. Blue boxes labeled 'E' indicate specific events. A green callout bubble points to the graph with the text: "Correlate events that occurred at the same time".

A table below the graph lists events:

Time	Description
8:09:01 AM	Guest Info Tools Status changed from guestToolsNotRunning to guestToolsRunning
8:08:31 AM	Guest Info Tools Status changed from guestToolsRunning to guestToolsNotRunning
7:57:00 AM	Guest Info Tools Status changed from guestToolsNotRunning to guestToolsRunning
7:56:50 AM	Guest Info Tools Status changed from guestToolsRunning to guestToolsNotRunning
5:07:21 AM	Guest Info Tools Status changed from guestToolsNotRunning to guestToolsRunning

# Deep Dive into Disk and Network IO performance

The screenshot shows the vCenter Operations Standard interface in a Windows Internet Explorer browser. The title bar reads "vCenter Operations Standard - Windows Internet Explorer". The address bar shows the URL <https://10.137.158.28/vmMain.action#?opType='ec','id':492,'type':'esx','name>. A "Certificate Error" message is displayed in the address bar.

The menu bar includes File, Edit, View, Favorites, Tools, Help, and a Favorites section containing "vCenter Operations Standard". The toolbar includes Home, Configuration, Help, About, Signout, and other standard icons.

The main navigation tabs are INFRASTRUCTURE and ANALYSIS. The ANALYSIS tab is selected, showing a timeline of events for the host 10.135.188.119 (Running). The timeline shows several status changes related to guest tools running and not running.

Two callout boxes highlight specific features:

- A blue callout box points to the "Disk subsystem performance details by datastores and LUNs" section, which displays metrics for the "Aggregate FC\_LUN5\_RAID10" subsystem, including Total I/O Commands, Total Throughput (KBps), Total Latency (ms), Host I/O Commands, Host Throughput (KBps), and Host Latency (ms) for both READ and WRITE operations.
- A green callout box points to the "Network statistics for every NIC" section, which displays metrics for the "vmnic2" interface, including Request Rate (KBps) and SENT/RECEIVED rates.

The bottom left corner of the interface shows the number 46.

# Identify and isolate KPI metrics

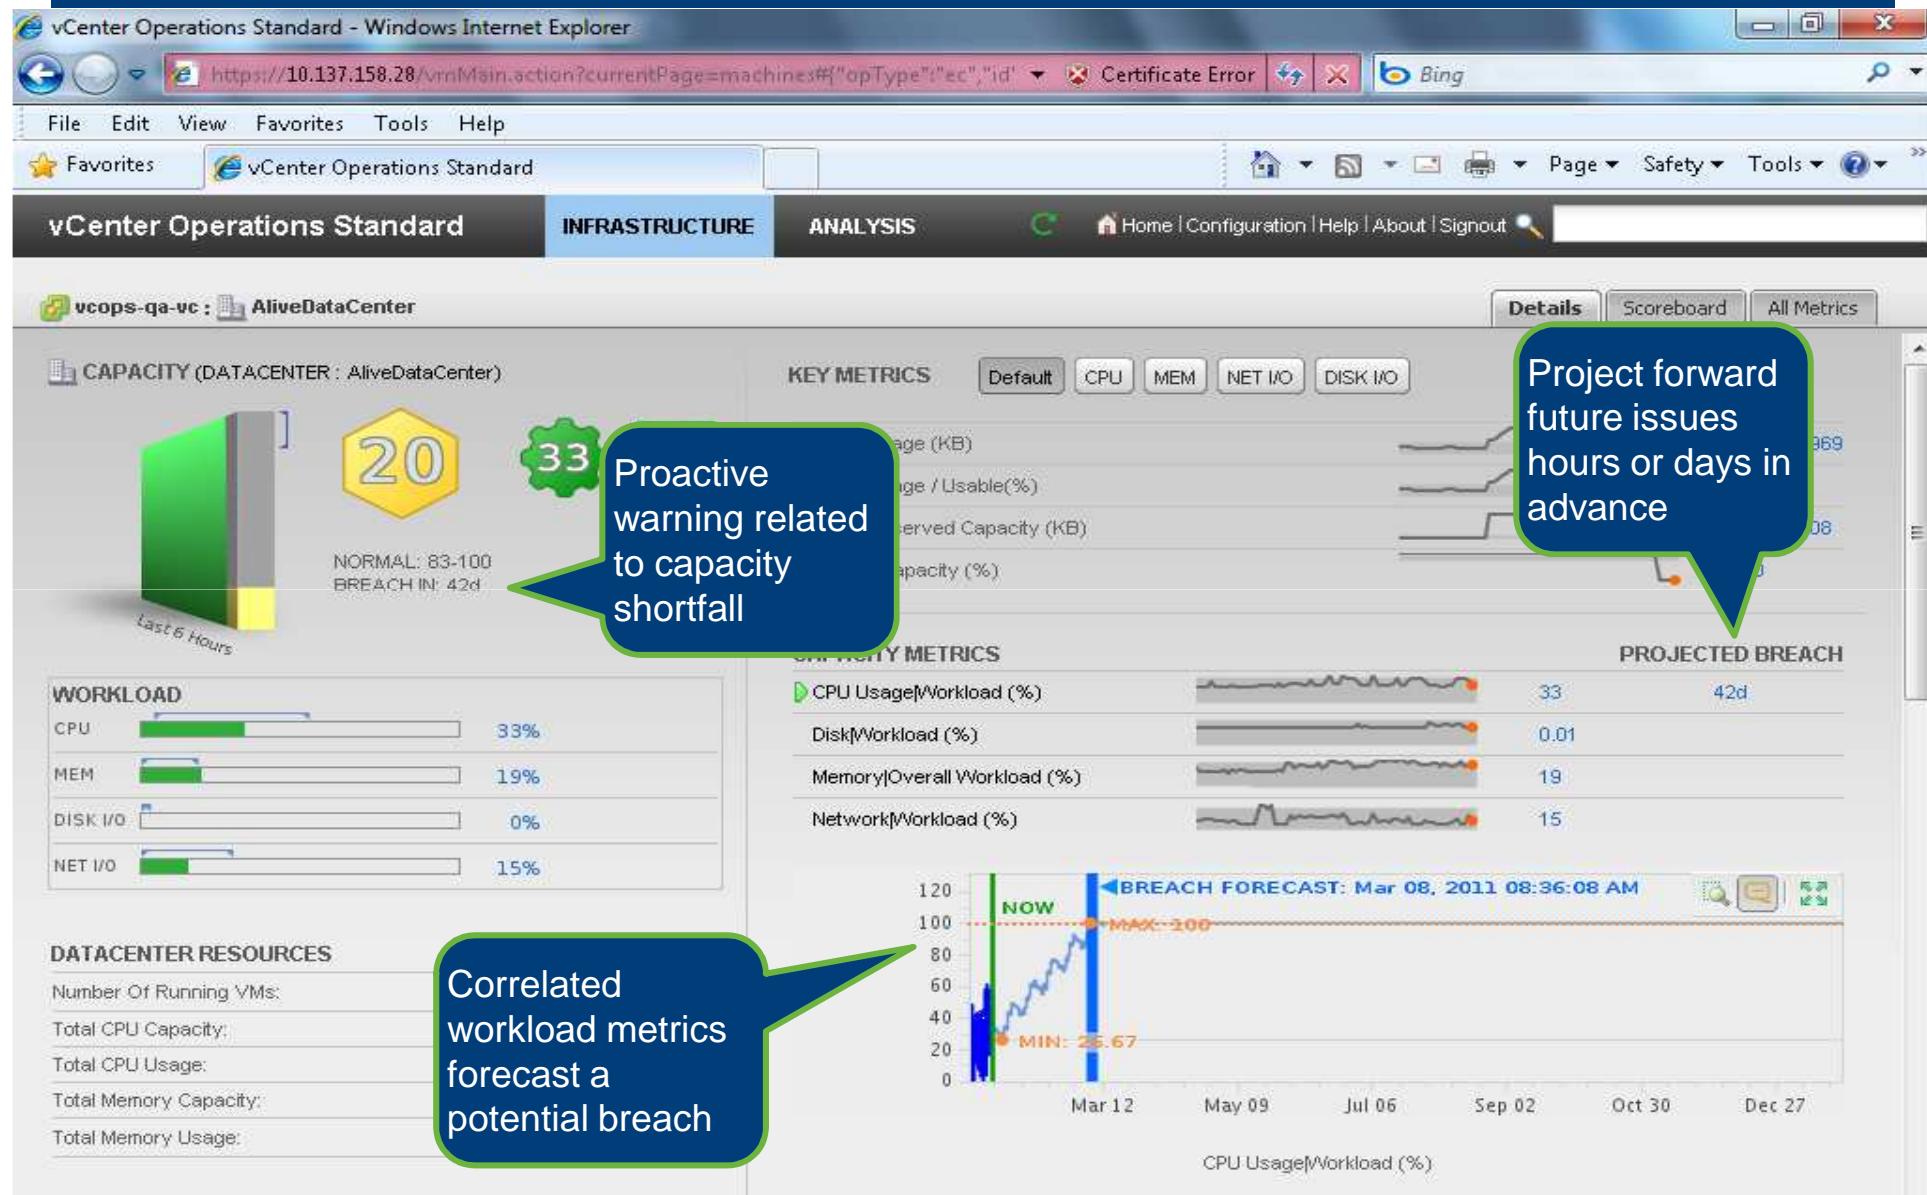
The screenshot shows the vCenter Operations Standard interface in Internet Explorer. The main dashboard displays a green bar chart for 'HEALTH (DATACENTER : AliveDataCenter)' with a value of 90, and three smaller boxes showing 33, 20, and 20. A callout bubble highlights these numbers with the text: "Quickly identify ‘suspect’ performance metric". Below the chart, there's a section for 'WORKLOAD' with CPU, MEM, DISK I/O, and NET I/O metrics. A callout bubble highlights the 'NET I/O' bar at 15% with the text: "KPI history with timestamp to indicate root cause". On the right, the 'KEY METRICS' tab is selected, showing four line charts for Memory Usage (KB), Memory Usage / Usable (%), Memory Reserved Capacity (KB), and Summary Capacity (%). The 'MEMORY' section lists five items all at 100% usage from 5:08AM to 5:19AM. The 'SYSTEM' section lists five items all at 100% usage from 5:14AM to 5:29AM. The 'SUMMARY' section lists three items all at 100% usage from 4:38AM to 5:44AM.

Quickly identify “suspect” performance metric

KPI history with timestamp to indicate root cause

Metric	Value	Timestamp
Memory Usage (KB)	43,518,969	5:19AM
Memory Usage / Usable (%)	92	5:19AM
Memory Reserved Capacity (KB)	1,630,208	5:19AM
Summary Capacity (%)	20	5:19AM
Memory VM Overhead (KB)	100 %	5:08AM
Memory Unreserved (KB)	100 %	5:08AM
Memory Reserved Capacity (KB)	100 %	5:08AM
Memory VMkernel Usage (KB)	100 %	5:08AM
Memory Usage (%)	100 %	5:19AM
System host User Resource Memory	100 %	5:14AM
System host Resource Memory O	100 %	5:14AM
System host User Resource Memory	100 %	5:29AM
System host Resource Memory M	100 %	5:29AM
System host Resource CPU Usage	100 %	5:49AM
Summary Total Number of VMs	100 %	4:38AM
Summary Number of Running VMs	100 %	5:08AM
Summary Capacity (%)	100 %	5:44AM

# Anticipate Capacity Issues Before They Happen



# Opportunities to remediate

vCenter Operations Standard - Windows Internet Explorer  
https://10.137.158.28/analysis.action?currentPage=analysis Certificate Error Bing

File Edit View Favorites Tools Help  
Favorites vCenter Operations Standard

vCenter Operations Standard INFRASTRUCTURE ANALYSIS Home Configuration Help About Signout

HEATMAP GALLERY Focus Area All Smallest Box Shows VM Search Showing 8 of 23 Customize

Description	Color By	Size By	Group By	Smallest Box
VM Health Sized By Workload Grouped By Cluster/Host	System Generated Self - Health Score	Summary Workload (%)	CLUSTER / ESX	VM
VM I/O Contention Sized By I/O Usage Grouped By Datastore	Datastore:<datastore> Total Latency ...	Datastore:<datastore> Usage Rate(K...)	DATA STORE	VM
VM Memory Contention Sized By Memory Grouped By Host	Memory Contention (%)	Memory Usage (KB)	CLUSTER / ESX	VM
VM Sized By Network Grouped By Host	Network:Aggregate of all instances...	Network:Aggregate of all instances...	CLUSTER / ESX	VM
VM Sized By Space Grouped By Host	Datastore:<datastore> Used Space (...	Datastore:<datastore> Used Space (...	DATA STORE	VM

Move VMs to another host?

VM Health Sized By Workload Grouped By Cluster/Host

This host looks healthy...

This host seems to be overloaded!

VM	CLUSTER	ESX	Summary Workload (%)	System Generated Self - Health Score
ST-VCOPS-353132	Other	10.135.188.9	38	14
vcops-347315	Other	10.135.188.9	56	48
ST-VCOPS-353132	Other	10.135.188.9	?	10

Page 2 of 2 | 101 - 161 of 161

# Individual performance metric details

The screenshot shows the vCenter Operations Standard interface in a web browser. The top navigation bar includes File, Edit, View, Favorites, Tools, Help, and a Certificate Error message. The main menu has tabs for vCenter Operations Standard, INFRASTRUCTURE (which is selected), ANALYSIS, and other options. A sidebar on the left shows a HEALTH TREE and a METRIC SELECTOR for the IP 10.135.188.119. The METRIC SELECTOR lists various metrics under categories like CPU Usage, Disk, and Aggregate of all instances. Two callout bubbles highlight specific features: one pointing to the METRIC SELECTOR with the text "Detailed list of all metrics indicating smart alerts", and another pointing to the METRIC CHART with the text "Single view that correlates multiple metrics". The METRIC CHART displays two line graphs for "Datastore|Aggregate of all instances|Write Latency (ms)" and "Disk|Aggregate of all instances|Disk Write Latency (ms)", both showing data from December 25 to January 26, 2010.

Detailed list of all metrics indicating smart alerts

Single view that correlates multiple metrics

vCenter Operations Standard - Windows Internet Explorer  
File Edit View Favorites Tools Help  
Favorites vCenter Operations Standard  
vCenter Operations Standard INFRASTRUCTURE ANALYSIS Signout  
vcops-qa-vc : CIQDataCenter : 10.135.188.119 (Running)  
HEALTH TREE

METRIC SELECTOR (10.135.188.119)

CPU Usage  
CPU Utilization for Resources  
Datastore  
Disk  
Idle (ms)  
Read Rate (KBps)  
Usage Rate(KBps)  
Workload (%)  
Write Rate (KBps)  
Aggregate of all instances  
Bus Resets  
Commands Aborted  
Commands per second  
Demand  
Disk Command Latency (ms)  
Disk Read Latency (ms)  
Disk Write Latency (ms)  
Kernel Disk Command Latency (ms)  
Max Queue Depth  
Number of Outstanding IO Operations  
Physical Device Command Latency (ms)  
Queue Command Latency (ms)  
Queued Operations

METRIC CHART

10.135.188.119: Datastore|Aggregate of all instances|Write Latency (ms)

H: 223.2  
L: 0.2

Dec 25 Dec 27 Dec 29 Dec 31 Jan 02 Jan 05 Jan 07 Jan 09 Jan 11 Jan 13 Jan 15 Jan 17 Jan 19 Jan 22 Jan 24 Jan 26 08:49 12:01 15:13 18:25 21:37 00:49 04:01 07:13 10:25 13:37 16:49 20:01 23:13 02:25 05:37 08:49

10.135.188.119: Disk|Aggregate of all instances|Disk Write Latency (ms)

H: 289.8  
L: 0.2

Dec 25 Dec 27 Dec 29 Dec 31 Jan 02 Jan 05 Jan 07 Jan 09 Jan 11 Jan 13 Jan 15 Jan 17 Jan 19 Jan 22 Jan 24 Jan 26 08:49 12:01 15:13 18:25 21:37 00:49 04:01 07:13 10:25 13:37 16:49 20:01 23:13 02:25 05:37 08:49

# Demo