

January 6, 2009

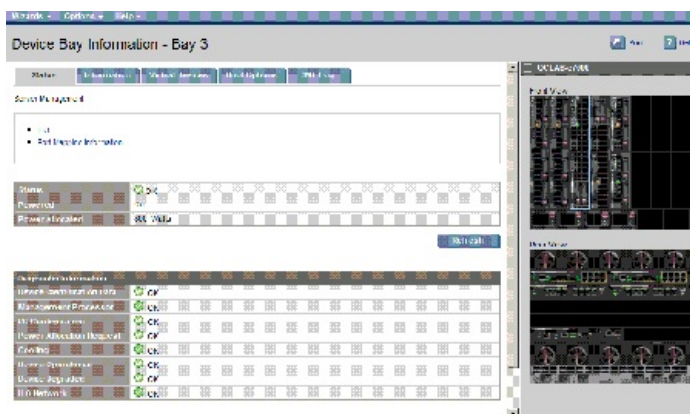
## BIOS settings for VMware ESX servers

Filed under: [bladesystem.esx.hp.vmware](#) — emac318 @ 11:32 am

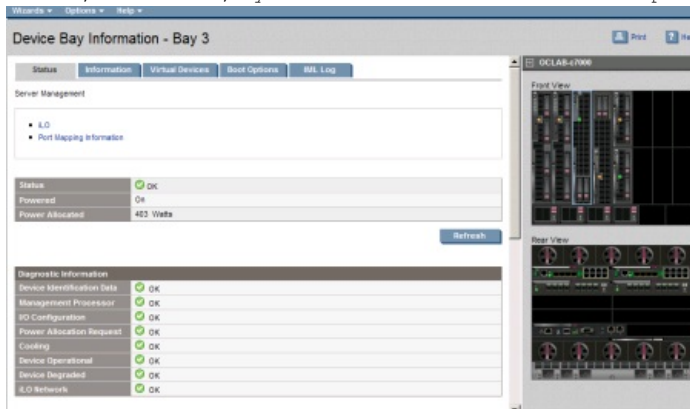
Dave, AKA [The VMguy](#), made a post about [BIOS recommendations for IBM servers](#). I thought I'd make a few comments on his recommendations:

Dave recommends turning off "PowerNow" or "Cool'N'Quiet" and points to a KB article about how this may affect timekeeping. However, the KB article does state that this does not apply to ESX server, but only to the various hosted products such as VMware (GSX) Server, Workstation, etc.

I would recommend doing the opposite, at least for HP servers. I have seen situations in our lab where a server with power management turned off will use more power than necessary. For example, with HP Dynamic Power Savings off, one of my test blades showed this:



However, with DPS on, my blade was down to almost half that power usage:



That's a pretty substantial power savings right there! I've also never seen a timekeeping caused by this setting. With IBM servers, your mileage may vary, however.

The last point I would add is that while Dave says you should turn on the "No Execute" or "Execute Disable" (the NX or XD bit, depending if you're in AMD or Intel land) just to fix an issue with a specific BIOS revision of an IBM server, I would argue that one should always turn on NX or XD on your systems. Why? Because [Enhanced VMotion Compatibility requires it!](#)

**You must ensure the BIOS settings for these processors enable Hardware Virtualization (if available) and Execute Protection. Default BIOS settings may not always enable these features. Hardware Virtualization is Intel VT on Intel processors and AMD-V on (supported) AMD processors. Execute Protection is Intel eXecute Disable (XD) on Intel processors and AMD No eXecute (NX) on AMD processors.**

Another setting that wasn't mentioned was that on AMD systems, there's generally an option involving NUMA called "Node Interleaving". Here's where it is on an HP system:

DISCLAIMER  
THIS BLOG IS  
AFFILIATED  
COMPANY OR  
IT'S JUST ME  
MOUTH WITH

RECENT TWEETS

Met @HPStorageGuy and @hermod today! Both awesome guys. Can't wait for Steve's keynote tomorrow! #Nth2010 5 days ago

New Blog Post: vCenter 4.1 upgrades & the LocalSystem account: Greeting everyone. Been very busy prep... <http://bit.ly/cl2glQ> from VMJunkie 1 week ago

New Blog Post: HP Client Virtualization Reference Architectures: One of the great things that came out... <http://bit.ly/aUBy5l> from VMJunkie 1 month ago

New Blog Post: The latest jerk move by Oracle: I don't really have a very high opinion of Oracle as a... <http://bit.ly/asYPtj> from VMJunkie 1 month ago

New Blog Post: Yay: I'm a vExpert! Shame: I haven't updated in too long: So it was with quite a lot o... <http://bit.ly/bB8c1L> from VMJunkie 2 months ago

ARCHIVES

July 2010  
June 2010  
April 2010  
March 2010  
February 2010  
December 2009  
November 2009  
October 2009  
September 2009  
August 2009  
July 2009  
June 2009  
May 2009  
April 2009  
March 2009  
February 2009  
January 2009  
December 2008

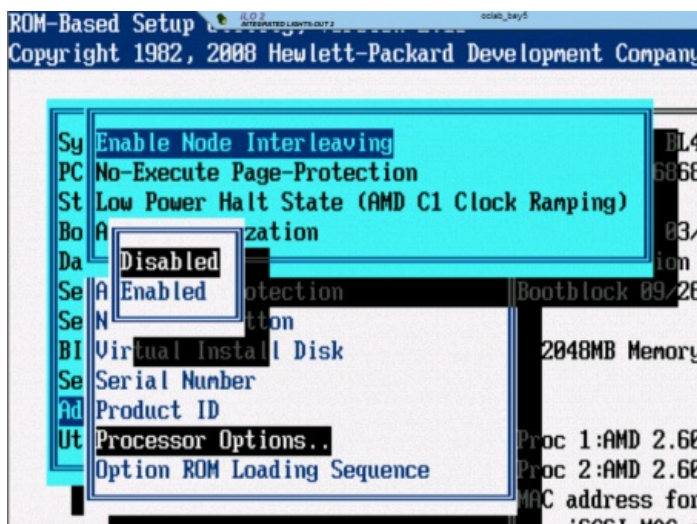
  

RECENT POSTS

VCenter 4.1 Upgrades & The LocalSystem Account  
HP Client Virtualization Reference Architectures  
The Latest Jerk Move By Oracle  
Yay: I'm a VExpert! Shame: I Haven't Updated In Too Long  
At Microsoft Management Summit 2010

TOP POSTS

Using The HP T5545 Thin Client With VMware View



On multi-socket AMD systems, each socket has its own path to memory, and each processor has its own memory bank which it can access the fastest. This means that Opteron systems have a [Non-Uniform Memory Access](#). ESX is NUMA-aware, and so will intelligently place a VM's memory on the same node as the processor its running on (if possible). Enabling Node Interleaving makes the whole memory space appear as a single uniform memory space (even though it isn't) which can negatively affect performance. So make sure to leave Node Interleaving off.

Possibly related posts: (automatically generated)

- [No Title](#)
- [Installing ESX vSphere 4.0 and vCenter 4.0 Best Practices](#)
- [Licensing VMware ESX 4, ESXi and vCenter 4 Video KB article](#)

[Comments \(4\)](#)

#### 4 Comments »

1.



Good info - I'll try it out.

Comment by [symbolik](#) - February 10, 2009 @ [8:08 am](#)

[Reply](#)

2.



HP recommends, when using VMware with DRS, to set the powermode to "OS Controlled" which is the same as Dynamic. See the PDF on page 5.

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c01732803/c01732803.pdf>

In a lab environment i can understand your point of view. However a busy ESX cluster with DRS and DPM enabled is a different situation.

Comment by [Paul Geerlings](#) - February 10, 2010 @ [6:08 am](#)

[Reply](#)

3.



Correction:  
"OS Control mode" is the same as "Static High Performance Mode"

Comment by [Paul Geerlings](#) - February 10, 2010 @ [6:18 am](#)

[Reply](#)

4.



Allocated power is not power usage, it is what the blade enclosure has estimated that the blade "could" use and is vastly different from actual wattage being used. So setting your blade to Balanced is not saving 400 watts of power. It is probably more like 30-60 watts.

If you watch, processes will float around cores, thus as they float, the procs will power/throttle up and down a lot.

To see what what your blade is actually using in watage, you need to go to the ILO on the blade, Power Management tab and then select Power Meter on the left.

Comment by [Garet Jax](#) - April 14, 2010 @ [2:07 pm](#)

[Reply](#)

Making A Thin Client On Fat Hardware: Part 1

Making A Thin Client On Fat Hardware: Part 2

Performing V2P Migrations

Winflp-2009-03-24-22-29-37

VCenter 4.1 Upgrades & The Local System Account

BIOS Settings For VMware ESX Servers

VMware View Article At BrianMadden

Balancing LUN Paths On Your ESX Hosts With PowerShell

Why You May Want To Leave FT Alone For Now...

#### B L O G R O L L

[BrianMadden.Com](#)

[Mike D's Notes From The Field](#)

[Pivot Point](#)

[Scott Lowe's Blog](#)

[The VMGuy](#)

[Vintemals](#)

[Virtualization.Info](#)

[VM /ETC](#)

[Yellow Bricks](#)

#### M Y L I N K S

[My LinkedIn Profile](#)

[Nth](#)

[RSS feed for comments on this post.](#) [TrackBack URI](#)

## Leave a comment

Name (required)

E-mail (required)

Website

Submit Comment

☐ Notify me of follow-up comments via email.

☐ Notify me of new posts via email.

[Blog at WordPress.com.](#)