

Wednesday, October 15, 2008

## **Analysing a so-called "Comparison" about Virtualisation at IBM Developerworks**

Whenever you want to dismiss the claims of a competitor or want to set your own or preferred technology in a better light, you should do some research on your topic. Otherwise you may end up with a document that's outright ridiculous.

I found a really strange piece of "comparison". It's called "A comparison of virtualization features of HP-UX, Solaris, and AIX". It's written by Mr. Ken Milberg. And I wasn't able to stop my shaking the head in disbelief. This text reinforces my personal impression, that this author is just a hired gun to publish claims even IBM doesn't want to make. But let's dissect his newest blurb. You shouldn't read it ... it's just a really abysmal document. I've sacrificed my time to do it for you, so don't waste your own

I will start with just four sentences of Mr. Milberg's document:

Scalability -- Only eight CPUs and 64GB RAM on one machine

Server-line -- Only low-end Sparc servers are supported

Limited micro-partitioning -- Four partitions on T1, Eight on T2

No Dynamic allocation between partitions

The truth is:  
Scalability: 4 sockets, 32 cores, 64 pipelines, 256 threads, 512 GB of memory

Server-Line: I would call the Sun SPARC Enterprise T5440 not really low-end SPARC services

Micro-Partitioning: Up to 128 LDOMs on a T2+ system, Up to 64 on a T2 and 32 on a T1.

Of course you can resize the LDOM without rebooting the system. At the end this is the way, you initially configure the system: At the beginning all CPUs and memory resources belong to the control domain, you take them away from this domain to give them to the guests.

Four claims, four times utter bullshit. I could stop the dissection now, but the article starts to get even more funny. So let's go ahead: Sun also offers hardware partitioning, which allows their high-end servers to be divided into four-process partitions. These are referred to as Sun DSD's. In many ways this technology is similar to IBM logical partitioning, which was introduced in 2001, with no real virtualization capabilities. DSD was the name of the technology in the Sun Fire Enterprise line. You were able to split those systems at the granularity of system board, thus the granularity was 4 sockets per domain minimum. That's correct.

But the Sun Fire Enterprise almost reached the end of its lifetime and now we sell the M-Class systems. There is something called quad-XSB mode. The documentation for Dynamic Reconfiguration on the M4000/M5000/M8000/M9000 states: SPARC Enterprise M4000/M5000/M8000/M9000 servers have a unique partitioning feature that can divide one physical system board (PSB) into one logical board (undivided status) or four logical boards. A PSB that is logically divided into one board (undivided status) is called a Uni-XSB, whereas a PSB that is logically divided into four boards is called a Quad-XSB. Each composition of physical unit of the divided PSB is called an XtendedSystemBoard(XSB). These XSBs can be combined freely to create domains. In a M4000 you can create 2 partitions, in a M9000 you can create up to 24.

Despite the statements of Mr. Milberg you are able dynamically move resources from one domain to another. This is a really old trick. I've done this on one of my E10K in 2000 and the system with the capability of creating 16 DSD. And this system was introduced in March 1997. So Mr. Milberg's comment "... IBM logical partitioning, which was introduced in 2001 ..." is good for some amusement.

At the end of the Sun part of the document he even starts to celebrate the advantages of WPARs without mentioning the disadvantages.

To close my article: The whole article is an insult to the real meaning of the sentence "This article explores all of these topics in detail." I'm not really sure how such an article was able to pass the editorial quality control at IBM.

PS: Okay, when you really want to read this botch job ... here is the URL: A comparison of virtualization features of HP-UX, Solaris, and AIX

Posted by Joerg Moellenkamp in English, The IT Business at 22:46

I think "dynamic" does not refer to reconfigure manually, but it talks about sharing resources in pools with weights and min-guarantee.  
Anonymous on Oct 16 2008, 12:22

## Blog Export: c0t0d0s0.org, <http://www.c0t0d0s0.org/>

Apart from missing on the latest CMT-offerings(the 5440 in particular) I think the IBM article is spot on.

While IBMs virtualization is mature, feature rich and constant across the pSeries server line, Sun virtualization is a mess. The Uni/Quad XSB partitioning on M-series line is no real virtualization by any definition. It is just a crude hardware partitioning.

For CMT the virtualization is in its infancy and certainly not enterprise ready. No redundant IO-servers. Memory resizing requires a reboot of the partition.

I haven't tried the container feature in AIX 6.1. However it seems more functional than Solaris containers already in its first release. Containers are by no means any replacement for physical virtualisation but has its place where applicable.

Anonymous on Nov 6 2008, 12:16

An mailaddress at googlemail just for lpars ??? Looks like a hired gun or an otherwise professional interested person

Could you elaborate on your comment, that the container feature is more elaborate in AIX? I don't think so, and i have done some research on both technology (Container by job, WPARS by interest)

BTW: LPARS are not a physical virtualisation. Uni/Quad-XSB are such mechanisms. You talk about hypervisor based virtualisation. That's isn't psysical at all ... the software just sits at a different place

And my personal opinion: A unix system doesn't need virtualisation in any way, as Unix is virtualisation. Virtualisation in the sense of LPARS/Containers/LDOMS is a security feature to generate a multitude of administrative domains. So i prefer containers instead of even LDOMS for most tasks.

Anonymous on Nov 6 2008, 13:26