

Virtualization with entry-level servers by IBM and Sun

AT YOUR SERVICE



Crawford Del Prete, Senior Vice President with the market researchers IDC, believes that every decade has its own IT revolution. The sixties were the heyday of the mainframe, which was ousted by the mini-computer ten years later. The client/server revolution was launched in 1985, and the nineties saw the breakthrough of the Internet.

Today's number one topic is utility computing, that is, the demand-driven computing resources provided as utilities just like electricity or water. Virtualization, which is one major precondition for utility computing, is the field with the strongest growth, says IDC. The analysts predict a market volume of US\$ 2.2 billion by 2007 [1].

Success stories from major software vendors appear to back up this prediction: VMware says that it is the fastest growing enterprise in the whole software industry, and SWSOft, whose virtual Virtuozzo servers are available for all major 32 and 64 bit platforms, reported 160 percent growth and doubled their staff last year. The free derivative, OpenVZ [2], is not the only contender in the open

The top players on the Unix server scene all include virtualization support with their operating systems. We compared two candidates from IBM and Sun.

BY JENS-CHRISTOPH BRENDEL

source sector; Xen [3] is also looking to make it into the Linux kernel. And there are many alternatives to these products.

Little wonder that traditional suppliers of Unix servers such as IBM, Sun, or HP, are making sure they aren't left behind, and presenting entry-level solutions to secure their share of the cake. Linux Magazine investigated two test candidates.

IBM P5-505

The first machine to reach our editorial offices was an IBM System P5-505. The machine came preinstalled with the AIX 5L 5.3 operating system and the so called Virtual I/O Server (VIOS), the central entity that corners and manages all of the system's I/O resources.

VIOS maps physical disks, network interfaces, or optical drives to resources for the virtual machines which are known as logical partitions, or LPARs for short. Besides this, of course, the LPARs need RAM, and either a dedicated CPU or an entitlement to a guaranteed minimum quota of CPU cycles from a shared CPU pool, which administrators can allocate in steps of 0.1 CPUs.

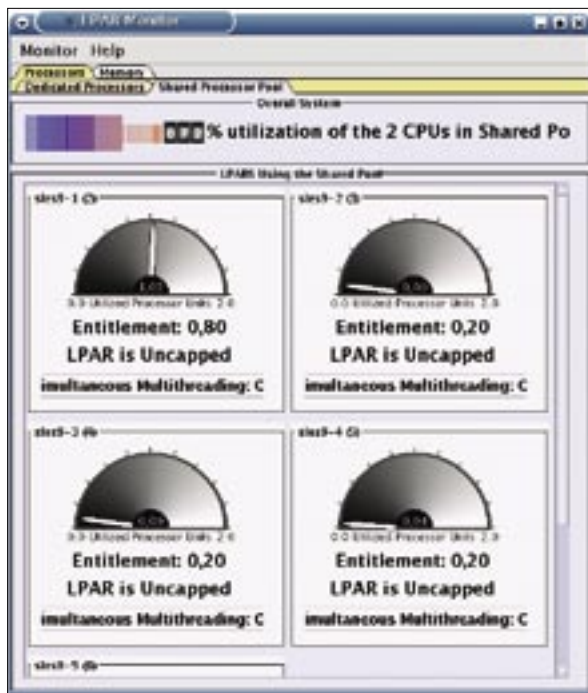


Figure 1: LPARMon visualizes the load on the various virtual partitions. In this scenario, the first instance, which is subject to more load, has been assigned more CPU performance.

All of these allocations are static. In other words, the resources assigned to one LPAR are not available to any other partitions, and changes on the fly are not supported. This said, to change allocations, administrators only need to disable the partition in question temporarily; there is no need to reboot the whole server.

VIOS also provides the user interface in the form of the Integrated Virtualization Manager (IVM) [4]. In this case, IVM replaces the Hardware Management Console (HMC) typically deployed by IBM, which would normally require an additional machine. Besides a command line interface, IVM also comes with an intuitive web GUI that allows administrators to set up and configure virtual partitions. A wizard is available to lend a hand for most jobs, and the default values also make sense, removing the need for expert-level skills.

Virtual Quintet

For our first test, we configured five logical partitions and installed Novell's Enterprise Linux SLES9 [5] on each. A web server was set up on each of the virtual Linux machines, and an external client was set up to bombard each server with requests. We used the Siege HTTP regression testing tool for this purpose.

The idea was to use parallel instances of the benchmarks to request a maximum number of pages from one, two, three, four, and finally all five virtual servers (Figure 1).

This approach makes it necessary to allocate resources to the consumers based on weighting (the original setting here was 0.2 CPU entitlements per LPAR).

As you would expect, the response times of the individual virtual servers increase with the number of virtual servers handling requests, whereas the throughput or transaction rate drops more or less inversely proportionally (Figure 2). Adding the transaction rate or

throughput results for all servers returns figures similar to the values for a single instance, leading us to conclude that the virtualization overhead is more or less negligible.

No Competition

In another test, the benchmark waited for a random period of time, but less than one second, before requesting each page. This meant less overlap when accessing system resources on the individual, virtual servers, and thus a more re-

alistic load distribution scenario than the full load test.

In this scenario, the number of parallel web servers had practically no influence on the transaction rate. The CPU load measured for the host machine on the AIX-based virtual server host machine simply shows that the CPU idle time drops when more web servers are enabled (Figure 3). This demonstrates that multiple LPARs can mobilize unused power reserves without getting in each other's way – assuming a task that places maximum load on the LPARs simultaneously.

Sun Fire 4100

The second test candidate was a Sun entry-level model. The Sun Fire 4100 machine also had 2 CPUs (dual core Opteron 280s in this case), and 8GBytes RAM. Solaris 10, with its zone-based virtualization technology, was preinstalled. At first sight, this concept is not much different from IBM's logical partitions: in both cases, specific applications can be completely isolated, assigned to independent execution environments, and allocated a predefined quantity of system resources.

At second glance, some obvious differences become visible. The most glaring difference is that Solaris zones automatically use the Solaris instance running on the host machine as their operating system. Other versions, or other operating systems, are not supported – at least not at this time of writing – an extension in the form of the BrandZ framework has




Figure 2: The performance of the virtual machines on the IBM and Sun entry-level systems was similar, and response times increased proportionally to the number of concurrent instances under full load.

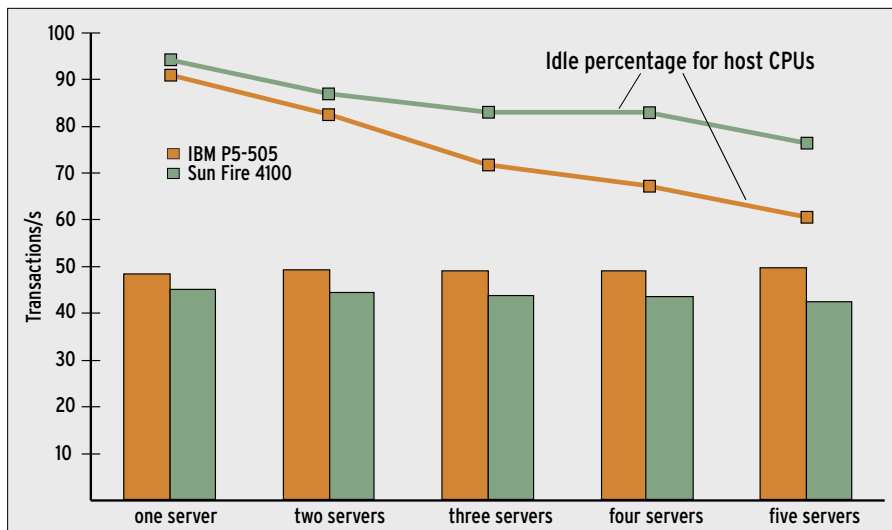


Figure 3: In a less competitive scenario, the performance of each virtual server does not depend to any extent on the number of concurrent virtual peers - what you have less of is CPU idle time.

already been released for OpenSolaris systems.

In addition to this, the resource pool, which groups the operative resources for a zone and isolates them from other zones, only supports one resource type, the CPU. In contrast to IBM's approach, there is no way of allocating RAM to a zone. This said, you can set a threshold for a specific application's RAM consumption within a zone, although in a fairly roundabout way. To set a threshold for an application's RAM consumption, you need to define projects to leverage Solaris' resource management feature, as projects can be allocated to zones.

A similar approach lets you set the guaranteed minimum CPU performance for a zone. This assignment of a guaran-

teed minimum can be modified dynamically at any time.

Less Breaks, More Work

Admins need to use the command line to set up and launch zones on Solaris; this is quite a simple task, despite the lack of a GUI or a wizard. Just like in our previous lab, we configured five zones and enabled a web server in each one. After doing so, we ran the benchmarks described earlier.

Of course, the absolute benchmark values do not allow a comprehensive judgment on the system performance - but this was not the aim of the test. What we can say is that the benchmark results were similar for both systems.

In the virtualization stakes, the results for Sun machine were what we ex-

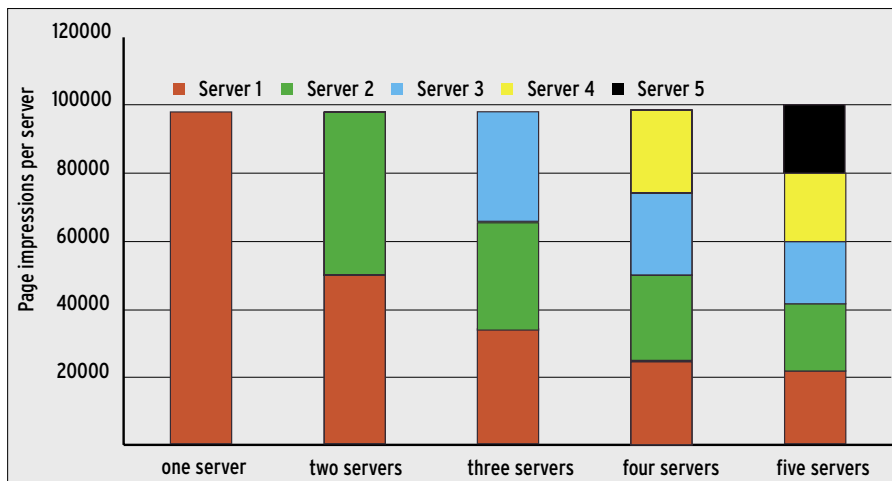


Figure 4: The sum of the benchmark results for all candidates remained constant: adding more Solaris zones obviously has very little effect on the system's performance.

pected: maximum load forces distribution of resources over the benchmark candidates.

The sum of the individual performance results remains more or less constant, again demonstrating that virtualization does not cause notable overhead (Figure 4). Mitigating the concurrency issues - and this is more of a true reflection of a production scenario - shows that the machine is capable of supporting far more than just a handful of virtual servers.

On Solaris, the CPU idle time dropped by just ten percent, although you have to bear in mind that more virtual servers would not only consume more CPU time, but also need more RAM, hard disk space, and interfaces.

Conclusions

Whereas virtualization on Linux or Windows means adding third-party software, many traditional Unix technologies integrate this capability. Virtualization causes very little overhead in this scenario, and it integrates seamlessly

INFO

- [1] IDC Special Report Data Center Virtualization:
http://www.sinaimedia.com/board/file/IDG_Special_Report_DC_Virtualization.pdf
- [2] OpenVZ:
<http://openvz.org>
- [3] Xen:
<http://www.cl.cam.ac.uk/Research/SRG/netos/xen/index.html>
- [4] IVM:
<http://www.redbooks.ibm.com/abstracts/redp4061.html?Open>
- [5] SLES9:
www.novell.com/products/linuxenterpriseserver

THE AUTHOR

As the editor of our sister publication Linux Magazine Germany, Jens-Christoph Brendel, is responsible for topics such as data center applications, databases, networks and security.

In his free time, Jens-Christoph enjoys playing everything from Bach to the Beatles on his acoustic guitar.



with familiar system administration.

In combination with high-performance hardware, and given the fact that an entry-level high-performance system is not much more expensive than a well-equipped PC nowadays, virtualization technologies provide a convincing alternative for many networks.

The differences between the Sun and IBM options are hidden in the details. In our

comparative test the IBM system scored a few extra points thanks to its support for more flexible use of the virtual machines and simpler resource management. But for experienced Solaris system administrators, who can handle all of their users' demands on a single platform, loss-free virtualization without third-party tools, based on the Solaris zone approach, definitely has some significant advantages. ■

IBM System P5-505



Name: IBM System P5-505

CPU: one or two 64 Bit POWER5 CPUs (2 in the test system)

Cache: 1.9MB L2 Cache, and 36MB L3 Cache

RAM: max. 32 GByte (8 Gbyte in the test machine)

Disks: 2x Ultra320 SCSI, max. 600 GByte internally

Slots: 2x PCI (1x long, 1x low profile)

Interfaces: 2x Ethernet (10/100/1000), 2x USB,

optional: Fibre Channel, 10 Gigabit Ethernet, Infiniband

System connectors: 2x HMC

Operating systems: AIX 5L or Linux (RHEL AS 4, or SLES 9)

Power supply: 2 redundant, hot pluggable power supplies

Form factor: 19 inch, single height

Sun Fire 4100



Name: Sun Fire 4100

CPU: one or two 64 Bit AMD or Opteron CPUs (single or dual core: 2 dual core Opteron 280s, 2.4 MHz in the test machine)

Cache: 1 MB L2 Cache per core

RAM: max. 16 GByte (8 GByte in the test machine)

Disks: max. 4x 2.5 inch Serial Attached SCSI disks internally, hot plug

Slots: 2x PCI (low profile,

1x100 MHz, 1x133 MHz)

Interfaces: 4x Ethernet (10/100/1000), 3x USB 1.1

System connectors: Service Processor (Ethernet)

Operating systems: Solaris, Linux (RHEL 3/4, SLES9), Windows Server 2003

Power supply: 2 redundant, hot pluggable power supplies, 550 Watt

Form factor: 19 inch, single height