Monday, September 14. 2009

### The impact of virtualisation

My colleague Jan searched a little bit in the large heap of benchmarking certifications of SAP and found something really interesting. The benchmarks certified in document number 2009034 (Sun Fire X4270, 2 processors /8 cores / 16 threads, Intel Xeon Processor X5570, 2.93 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 8 MB L3 cache per processor, 48 GB main memory) yielded 15,320 SAPS on 8 vCPU. On the other side an Fujitsu PRIMERGY Model RX300 S5 (2 processors / 8 cores / 16 threads, Intel Xeon Processor X5570, 2.93 GHz, 64 KB L1 cache and 256 KB L2 cache per core, 8 MB L3 cache per processor, 96 GB main memory) yielded 11,230 SAPS (as certified by 2009029). I should add something ... in my opinion it's not the system, that was responsible for this large gap. So what's the difference?

Both benchmarks included an virtualisation layer. Both systems were limited to use 8 vCPU for the central server. And here is the difference: The Sun system used Containers on Solaris 10 with Oracle 10g and started the central server in a non-global zone. The Fujitsu system used SLES 10 with MaxDB as a VMware ESX 4.0 VM. Based on this benchmark, it looks like that the containerized system is able to yield 136 percent of the performance of the VMware-virtualized system. 35 percent performance are a quite significant difference.

Posted by Joerg Moellenkamp in English, Oracle, Solaris, Technology, The IT Business at 16:09

Interesting, yet the Sun platform is at least 20% more expensive. Obviously containers have less overhead than VMs, but the technology to migrate VMs without dropping connections is well developed. VMs have less downtime than sun servers. Additionally the Oracle licenses are far more expensive than maxDB licenses. Consequently, the second configuration provides more results per dollar and has less downtime.
Anonymous on Sep 14 2009, 17:38

Well, the Sun Server and the Fujitsu-Server are comparable, even in the aspect of hardware costs, but the VMWare hypervisor & SuSE Linux together are more expensive than a Solaris License. On top, running the DB of a productive SAP system in a VM is only supported for Solaris Container. Vmotion is not supported by SAP (See note 1374671).
Anonymous on Sep 14 2009, 18:44

1. A higher availability of VMs would imply that the Hypervisor can predict hardware damages or outages. There is just one case, where there is an reduction of downtime and this is the case of the planed downtime for the hypervisor. I just consider your statement as largely unfounded. When you work in your operating system you have the same procedures in a VM than on bare metal. Live migration doesn't help you with updates of your OS or your application. No difference here. Furthermore every moving part in a Sun Server is hot-swapable. Thus the story of "live migration to swap a fan" doesn't catch, as you could do it anyway.

From my perspective the impact to reduce unplanned downtime is rather low, as the system has to detect the upcoming failure early enough to initiate live migration. With 96 (or just 48 GBytes) the migration can take a while The HA functionality of VMware is similiar provided to the one you find in SunCluster, it isn't stateless. The FT lockstepping functionality of VMware can't be used here, as FT just supports VM with 1 vCPUs.

2. The database has a low impact to the SAPS result as stated earlier in other articles in this blog. You would get similar results with the MaxDB.

3. The Fujitsu needs the additional SLES licenses and VMware licenses ans Support. Those aren't cheap as well. Solaris support is included in the SunHW support you would need for the VMware/sles version at any case.
Anonymous on Sep 14 2009, 19:51

Hm.. I wonder.. what happens when a non-moving part in the Sun server has a problem, however?

An example would be a memory chip or CPU.
Even if the hot swap is possible, it could be risky; a human making a mistake during a hardware swap procedure is possible, e.g. power cable to PSU A accidentally ejected while swapping PSU B.

I'd be hesitant to draw conclusions on this question without empircal data comparing actual availability of VMware VM vs Sun container...
Anonymous on Sep 15 2009, 05:57

1. A Sun server has enough LED on the outside and the inside for a christmas tree (even the DIMM slots have LEDs). Someone who swaps the wrong part should schedule a meeting with her or his ophthalmologist.

2. Regarding the non-movable parts failures. Solaris has the Fault Management Framework. It monitors the components of the system and take action in the case the errors start to reach certain thresholds: CPUs or cores are offlined when they start to throw to many correctable errors or a memory will be evacuated and retired, when it throws too many errors. It prints a nice message in your log, it throws an SNMP trap and you can schedule the repair at your convenience.

Anonymous on Sep 15 2009, 08:13