Cutting-edge technology for the desktop

Egg, December 1, 2016: The new BigFoot series for ArchivistaVM Server features up to 18 CPU cores and a maximum memory capacity of 64 TB. Two 10 Gbit and 40 Gbit network cards respectively ensure that data can be simultaneously stored in real-time on multiple machines. This blog post presents the new flagship and, on the basis of a conversation held at this year's LinuxDay, explains why it is worth taking a close(r) look when it comes to server virtualisation.



SSDs and the fastest processors

The virtualised guests in ArchivistaVM clusters are always saved three-fold and in real time on at least two computers. Whereas 10 Gbit network cards have up to now been regarded as being the measure of all things, two fast solid-state drives push them plus/minus to the limits of their capacity. Just four SSDs require a throughput of approx. 2 GB per second to achieve the redundancy at full speed.

Although such a throughput may not, at first sight, seem to be absolutely imperative, since the cost of the technology has meanwhile dropped to a moderate level, why should one not take advantage of it. Prices for the **BigFoot series start at less than 5,000 Swiss francs / euros per machine.** Although with their three nodes this is still a low five-figure amount, in contrast to other solutions no further costs are incurred with BigFoot systems as they already contain all the components. And maintenance for the first year is included in the price as well.

Server infrastructure on the desktop

Almost even more important is the fact that BigFoot servers explicitly do not require

server rooms with the high follow-on costs that they involve. Thanks to their water cooling system, BigFoot servers stay agreeably cool even with the most powerful Intel CPUs. The power requirements per computer lie (even under load) at barely over 100W, or 70W when idle. The desktop design can accommodate a maximum of 16 drives per machine with a total capacity of 64 TB. Just 12 SSDs and a three-fold redundancy produce a throughput of 1.8 GB/s when simultaneously writing multiple jobs (the standard task of any server virtualisation) with a RAID 10 array and concurrent data transfer to the second computer with DRBD, or even 2,1 GB/s when reading:

hdparm -tT /dev/md8 Timing cached reads: 13256 MB in 2.00 seconds = 6637.05 MB/sec Timing buffered disk reads: 6374 MB in 3.00 seconds = 2124.66 MB/sec

A rack system can hold up to 24 SSDs with a maximum capacity of 96 TB. A throughput of up to 4 GB/s can be realised with this configuration, although this exhausts the performance potential of the 40 Gbit network cards. While 100 Gbit cards would produce a higher performance, at the moment they are not practicable, quite simply because the RAM cannot (yet) cope with this throughput.

With the BigFoot series, it is possible to realise ArchivistaVM clusters with the most powerful server platform that is currently available (LGA 2011-3). Thanks to flexible motherboards, Intel processors of all price and performance classes can be used, while the BigFoot models are also extremely adaptable when it comes to the main memory. Between 16 and 512 GB RAM can be installed, while between 6 and 18 processors are available for the CPU cores. Even systems with a two to four-fold performance (36 cores and 2 TB RAM) are possible with a rack design.

Open source and standard components



The BigFoot series impressively demonstrates what can be achieved today using open source and standard components, even though some IT professionals still believe that this is only possible using the most expensive server components. At this point we would like to mention a rather long discussion held at **LinuxDay** in Austria. Two smart computer scientists examined the ArchivistaVM series. What's the configuration? Xeon-D with up to 16 CPUs. Astounded, they then asked whether a hardware RAID controller was installed, since without one the system would be pointless anyway, and on top of that it would have to be a full-height controller. A throughput of 800 MB is pretty good, but only the RAID controller from manufacturer X could handle a throughput of 2 GB/s.

Of course that's not true, but the discussion now revolves around redundant power

supply units which they could obtain from manufacturer X for 30 percent of the retail price. Isn't a margin of 70 percent rather dubious, was the next question. Not at all, even at a price of 4000 euros customers still get very affordable servers. When asked about the configuration, it is at least admitted that the price does not include the setup (about 1000 euros), the software (for up to 6 sockets, software Y is practically given away, likewise the backup software), or the maintenance.

Keyword "open source". The question is justified inasmuch as it is obvious from their shirts that the guys are advertising (and probably work) for a company that has been sponsoring **LinuxDay** for several years. Well, they used to do some XEN, but Y, even if it is a proprietary solution, meanwhile costs almost nothing for low-volume customers.

The next question is whether the guests are run redundantly on multiple machines. That's not necessary for lower volume customers, and 40 Gbit isn't an issue either. So how are the 2 GB/s sent down the line, then? Accumulating and then doubling up isn't necessary with the backup software. Confirmation that this isn't a backup, but it's about live operation, and possible failures of a computer and how to intercept such cases. For this, they could fall back on the supplier's 24-hour service.

Of course, such concepts are quite practical for computer service providers: a 70% margin on the components, a day's work setting up each server, delegation of the hardware deployment to service providers, and if the installation grows to cover multiple servers, there's also the sharp rise in the revenue curve thanks to the licence fees. *"True, we don't do anything special. We use the same raw materials as well, just like everyone else: the hardware and software that's available today".* Although that doesn't stand on the smart gentlemen's shirts, it can be found on the provider's website.



Accepted, nobody does anything special. That's why the BigFoot series also uses 40 Gbit network cards, CPUs from the market leader, and powerful SSDs. Nevertheless, significantly more can usually be achieved using standard components and open source with much less effort, especially where server virtualisation is concerned.

That is proved by the infrastructure solutions from internet players who realise their server infrastructures using standard hardware and open source. Cluster solutions for thousands of servers are simply too complex for SME companies (take a look at the **Ganeti project,** for instance). Yet it isn't necessary to forego standard hardware and open source just for this reason – especially where ArchivistaVM is concerned.

The BigFoot series can be checked out both here in Egg or in an on-premise installation, because, as we know, anybody who doesn't do anything special hasn't got anything to hide, either.