
Subject: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Walter F.J. Müller](#) on Thu, 24 Mar 2005 09:35:58 GMT

[View Forum Message](#) <> [Reply to Message](#)

Soon we have to decide whether this years upgrade of the batch farm is done with 32bit or 64bit systems, whether we stick with Xeon's or switch to Opteron's.

First test results were reported in the last 'DV Ausschuss', see Report of W. Schön.
(accessible only within GSI)

On the HEPix Autumn Meeting Stephan Wiesand summarized the experience gained at DESY, see his HEPix presentation.

Further input, comments, and suggestions are very welcome.

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Anar Manafov](#) on Fri, 25 Mar 2005 01:21:52 GMT

[View Forum Message](#) <> [Reply to Message](#)

Walter F.J. Müller wrote on Thu, 24 March 2005 10:35: Soon we have to decide whether this years upgrade of the batch farm is done with 32bit or 64bit systems, whether we stick with Xeon's or switch to Opteron's.

First test results were reported in the last 'DV Ausschuss', see Report of W. Schön.

On the HEPix Autumn Meeting Stephan Wiesand summarized the experience gained at DESY, see his HEPix presentation.

Further input, comments, and suggestions are very welcome.

Good day, Walter!

This subject is quite engaging.

I looked for the documents you provided.

I only cannot reach Walter's report. Is it possible to get access to the report of Walter Schön (from DV Ausschuss)? It seems the document is in the restricted area.

May be you could provide a copy, if it is not really restricted.

I wonder, do we have any machine for test already?

I would love to see at least some test machines and play with them; I have a bit of code (my test 64b development in C++ and Java) to test on 64b stones.

There will be better 64b integer arithmetic. Pointers can address up to 2^{64} bytes (theoretically), but still there is Linux kernel restriction, as far as I know, it can't address more than 2^{42} bytes. In anyway it is better than 2^{32} .

Then we can work with large files more efficient... Try real 64b software (so much advertised)...

Walter, do you know any details on the specification of machines with 64b stone, which GSI probably wants to order?

(vender, memory size, disc, bus and so on...)

From my point of view, I would order only view machines for DB management and WEB servers, maybe some for J2EE app-server if GSI has a development on it (I read very good test reports on that) and couple for test purpose.

Concerning, GSI batch farm, IMHO, it is too early to switch it to 64b.

BTW, the complete hardcopy issue of Dr.Dobb's journal (March 2005; www.ddj.com) covers the 64-bit computing. I have successfully proposed to order DDJ for GSI library two month ago. There are 2 (3) issues already in the library.

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Walter F.J. Müller](#) on Fri, 25 Mar 2005 19:21:47 GMT

[View Forum Message](#) <> [Reply to Message](#)

Anar Manafov wrote on Fri, 25 March 2005 02:21...

The 'DV Ausschuss' material is 'considered GSI internal' and only accessible from within GSI.

Anar Manafov wrote on Fri, 25 March 2005 02:21...

Concerning, GSI batch farm, IMHO, it is too early to switch it to 64b.

Could you be more specific ? Isn't computational price/performance an essential argument ?

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Anar Manafov](#) on Fri, 25 Mar 2005 21:41:01 GMT

[View Forum Message](#) <> [Reply to Message](#)

Walter F.J. Müller wrote on Fri, 25 March 2005 20:21Anar Manafov wrote on Fri, 25 March 2005 02:21...

I only cannot reach Walter's report.

The 'DV Ausschuss' material is 'considered GSI internal' and only accessible from within GSI.

Thx! I got it now.

Walter F.J. Müller wrote on Fri, 25 March 2005 20:21Anar Manafov wrote on Fri, 25 March 2005 02:21...

Concerning, GSI batch farm, IMHO, it is too early to switch it to 64b.

Could you be more specific ? Isn't computational price/performance an essential argument ?

Walter, It is just my IMHO.

The 64b instructions are good and much better than 32b and you are 100% RIGTH when saying that "computational price/performance" is essential argument! Here nobody, I think, can argue. It is obvious and even obvious that at this moment AMD is the best with Opteron. But this works only when you are using it, using advantages of 64b! Get it to use as a 32b CPU, I think is not a very good idea. I would better get new Intel Pentiom4 (by the way also with 64b instructions; unfortunately with less performance)

As I mentioned in my previous post, the DB servers, WEB servers, Application-ser (for instance, J2EE app-server). this is where we can use 64b with guaranty of saving money and gaining valuable performance. Concerning desktops and GSI batch farm I would wait a lit bit. Wait until more 64b application and Development toolkits are out and tested. Wait until developers and users can learn to use the advantages of 64b processing in their software. Wait until it became cheaper and with higher performance. I think GSI using the same environment on all of the machines, this is like GSI policy. So it will be not an advantage to use Opteron 64b on the same environment as other 32b have. I think, it would be like using rocket machine in zone 30!

Again, it is just my IMHO

I think the environment is not ready for such a change, yet.

Anyway, as I mentioned in my previous post, I still think that we need small amount of 64b machines (MP, I would suggest) for testing and for moving development to it, for adaptation, but only very few machines with dedicated queue.

About putting 64b on the desk I wouldn't even think, IMHO. At least now.

BTW, in the parallel computing Xeon vs Opteron; the advantages are not so much visible: [The test where made in Lomonosov Moscow State University, laboratory of parallel computing]

AMD on Linux SuSE SLES-8 (AMD64) VERSION 8.1, kernel-2.4.19-SMP x86_64

Intel on Linux RedHat 7.3, kernel-2.4.25

AMD Opteron:

With 1CPU, 2G of memory (partial use of memory. matrix 15000x15000) = 3.38 Gflop/s (of 84.5% pick efficiency)

With 4CPU, 2G of memory (use of full available matrix 29000x29000) = 12.6 Gflop/s (of 79% pick efficiency)

Intel Xeon

With 1CPU, 2G of memory (partial use of memory. matrix 15000x15000) = 3.73 Gflop/s (of 72% pick efficiency)

With 4CPU, 2G of memory (use of full available matrix 21000x21000) = 12.55 Gflop/s (of 60.3% pick efficiency)

Even conserning this „Im Vergleich zu den schnellsten Xeon Rechnern gewinnen die Opteron je nach Anwendung zwischen einem Faktor 3 und 5 in der Rechenleistung.“ (from Walter's report)

Is it REALLY that much???

Surely I believe to Walter and even if it is that much differ!

It is only because of 64b integer arithmetic and nothing to do with ONLY Opteron! New Intel 64 also has 64b integer arithmetic!!!

Concerning servers I have no doubt, the OPTERON is only choice now.

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Christopher Huhn](#) on Tue, 29 Mar 2005 15:26:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

By the end of the year all high performance CPUs by either AMD or Intel will be 64bit-enabled.
The only questions to answer then will be:

Run a 32 or 64bit OS
Buy Intel or AMD

We are currently running two dual Opteron test machines 2GHz with 4 and 8 GB of memory under Debian Sarge.
Anyone interested in testing may contact me or Walter Schön to get access to the boxes.

Maybe the Wiki might be a good place for a compilation of the results.
I'd be especially interested in apps that do not run 64bit or do not gain performance.

Regards,
Christopher

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Anar Manafov](#) on Tue, 29 Mar 2005 20:56:49 GMT
[View Forum Message](#) <> [Reply to Message](#)

Christopher Huhn wrote on Tue, 29 March 2005 17:26
The only questions to answer then will be:

Run a 32 or 64bit OS
Buy Intel or AMD

I think they are not only questions, because it is quite easy to find 64b OS.
The main questions would be
When experiments and software will be ready to use 64b? (With everything comes after it)
When developers will learn to program for 64b? (Because it is far from just switch the compiler option)

As a test, just try (one day) to switch to gcc 3.4 or 4.1 as a default compiler (on current 32b OS) and to check how many libraries and software packages of all of the GSI exp. will be still in a good condition. How many programs will be compiled...?! Even this would be a very big change in the environment! And this will make people be busy for some time. OR This idea!
As I mentioned in my previous post, I think it is very early to even think about 64b!
I am not sure that there is even a FACTOR of better performance (in 32b mode) in comparing with late 32b CPUs.

The programming for 64b is quite different than current 32b programming schema. You can argue that high level programming languages are safer. Right! But Even with Java you will not get a big (THE factor!) performance on 64b in comparing to 32b (based on the latest test from DDJ).

And I think if we switch the farm to 64b that means we need to switch (change) OS and hardware on all of the GSI machines or it will be a major split of supported environment?

But maybe I don't know something, may be we are in a hurry to switch to 64?
Anyway, I would be a bit patient to start to use 64b in our cluster, IMHO.

Christopher Huhn wrote on Tue, 29 March 2005 17:26

Anyone interested in testing may contact me or Walter Schön to get access to the boxes.

Maybe the Wiki might be a good place for a compilation of the results.
I'd be especially interested in apps that do not run 64bit or do not gain performance.

Thx for the offer! This is very convenient! I' would love to test my stuff.

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Christopher Huhn](#) on Wed, 30 Mar 2005 08:35:26 GMT

[View Forum Message](#) <> [Reply to Message](#)

Anar Manafov wrote on Tue, 29 March 2005 22:56

The main questions would be

When experiments and software will be ready to use 64b? (With everything comes after it)

When developers will learn to program for 64b? (Because it is far from just switch the compiler option)

The AMD64 Linux kernel can execute i386 binaries, so it's no problem to run your old binaries under a 64bit OS. The only problem is to keep the shared libs for 32bit and 64bit in distinct places.

The current proposed solution for Debian is to install a 32b chroot and run your 32bit only apps in there (like openoffice ...), see

<https://alioth.debian.org/docman/view.php/30192/21/debian-amd64-howto.html>. The desired long term solution is multiarch (see <http://www.linuxbase.org/~taggart/multiarch.html>).

Anar Manafov wrote on Tue, 29 March 2005 22:56

Christopher Huhn wrote on Tue, 29 March 2005 17:26

Anyone interested in testing may contact me or Walter Schön to get access to the boxes.

Thx for the offer! This is very convenient! I' would love to test my stuff.

You already have access to one of the machines (lxts09) for a while?!

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Anar Manafov](#) on Wed, 30 Mar 2005 11:18:55 GMT

[View Forum Message](#) <> [Reply to Message](#)

Thanks for the links. It is very interesting to me!

Christo, I still have some questions concerning Debian 64b instillation:

Quote:

This port is actually considered beta.

Quote:

So, when it is released?

The first planned official Debian release of 64bit userland for AMD64 will be Etch (Sarge+1).

The unofficial sarge based on the debian-amd64 port will be ready shortly after the official sarge release, and will have security updates provided by the amd64 porting team.

Quote:

Debian-pure64 is still classified "beta" by the porters, while debian-pure64-3.4 is classified "experimental". The port is being worked at since the beginning of 2004, and many people use it for both workstations and servers. Lacking a stable distribution, the port is not recommended for "mission critical production systems" until sarge released.

There are 64bit related problems in approximately 200 packages, which result either in the package not building at all or simply segfaulting on startup. The latter problem should not show up too often, only in first releases of new upstream versions of some packages which had the problem in the past. The vast majority of packages with known 64bit problems do not build at all.

Is this is a TRUE?

Are we going to install something which is a beta? Doesn't look like GSI policy.

Or the document you pointed is not up to date and the release is done and tested already?

I am sure we will not install beta and going to wait until release, but another question, are we going to install the FIRST 64b release?

I would be very proud to work on the edge of technology, but it is quite unsafe and will require a lot of effort. But you know me, I would be the first to participate and do my best.

If we use 32b chroot, don't it slower down a bit the 32b processes? Redirection and kind of stuff... It looks like the same solution as Microsoft chooses for the Windows 2003 Server (64b edition).

Anyway, as I understand correctly, IT division needs to maintain at least two environments, and several kernels. Since GSI can't get 64b machines for change of Entire Park, right?

I still think that we are too hurry to switch to 64b environment. It is too early. I would give some time, until the software became stable and not the FIRST release, FIRST TRY! Also the prices will be lower for higher performance.

Even if the experiments software will be completely ported (which I am really doubt), we will have some problems to run it on other sites, like CERN. I think CERN is going to keep 32b architecture for some time, and not only the CERN.

BTW, this document (Multiple architecture problem and proposed solution) you pointed is really GOOD! Thx!!!

Christopher Huhn wrote on Wed, 30 March 2005 10:35

You already have access to one of the machines (Ixts09) for a while?!

No, I think I don't have the access. Anyway, I am going to come to you and ask

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Christopher Huhn](#) on Wed, 30 Mar 2005 12:29:45 GMT

[View Forum Message](#) <> [Reply to Message](#)

Anar Manafov wrote on Wed, 30 March 2005 13:18Quote:

The unofficial sarge based on the debian-amd64 port will be ready shortly after the official sarge release, and will have security updates provided by the amd64 porting team.

Is this is a TRUE?

Are we going to install something which is a beta? Doesn't look like GSI policy.

To be honest, Debian "testing" is incredibly stable (as it contains almost only software considered stable "upstream"). The AMD64 port of Sarge will be 99,9% as stable as the i386 Sarge release IMHO.

My long term plans would be to permanently have testing and unstable test clusters running after the upgrade to Sarge "to be on the edge of technology" (even Debian "unstable" is quite conservative IMHO)

Quote:

If we use 32b chroot, don't it slower down a bit the 32b processes?

I don't expect a serious performance impact but we'll have to test.

Quote:Anyway, as I understand correctly, IT division needs to maintain at least two environments, and several kernels. Since GSI can't get 64b machines for change of Entire Park, right?

Absolutely. No hurry, no worries.

Subject: Re: Xeon's, Opteron's, 32 vs 64 bit -- Next steps for the GSI batch farm
Posted by [Anar Manafov](#) on Wed, 30 Mar 2005 12:36:32 GMT

[View Forum Message](#) <> [Reply to Message](#)

All clear for me now, I think.

Thx Christo!
