

Re: Beryl options and nvidia

Source: <http://linux.derkeiler.com/Mailing-Lists/Fedora/2007-01/msg00786.html>

- *From:* Marc Schwartz <marc_schwartz@xxxxxxxxxxx>
 - *Date:* Fri, 05 Jan 2007 20:57:25 -0600
-

Lonni, sorry for the delay in my reply here, I got pulled into a meeting this afternoon and am just finally getting back to this.

Lonni J Friedman wrote:

On 1/5/07, Marc Schwartz <marc_schwartz@xxxxxxxxxxx> wrote:

On Fri, 2007-01-05 at 12:19 -0800, Lonni J Friedman wrote:
> On 1/5/07, Marc Schwartz <marc_schwartz@xxxxxxxxxxx> wrote:
>> Gérard Milmeister wrote:
>>> I have beryl running with nvidia prop. drivers with the following
>>> options:
>>> beryl --use-cow --force-aiglx --skip-gl-yield
>>> There are now problems with black windows etc... that may appear
with
>>> other options. The X server is running with increased priority (set by
>>> using schedtool), so there responsiveness is good running tvtime
>>> smoothly even under system load.
>>> There is one small issue however: Before a menu window is filled with
>>> content, there may appear some artifacts in the window background.
These
>>> disappear if I use the --force-nvidia option, but then UI
responsiveness
>>> is distinctly slower and tvtime does not run smoothly anymore.
>>> Is this a known issue?
>>
>> There are known problems with the current nVidia drivers and the use of
>> Compiz/Beryl, secondary to the current incomplete implementation of
>> texture_from_pixmap (TFP) required for compositing. Based upon the
>
> incomplete in what way?

Quoting from the 9629 Release Highlights last November 7:

"Added initial support for GLX_EXT_texture_from_pixmap."

The word "initial" suggests to any reasonable reader that the implementation is not complete or at least not fully stable and certainly experience over the two months since then has born that out.

Re: Beryl options and nvidia

There has been no further mention of TFP in any of the similar highlight notes for any of the subsequent stable releases, despite hundreds of posts on nvNews, here and elsewhere relative to ongoing problems with running any of the compositing managers.

Sorry if the word "initial" was misleading. It was meant as an indication that the 1.0-9629 beta driver was the first driver to provide support for TFP, not that the support was incomplete. The support that was added to a driver from November was complete with respect to the texture_from_pixmap specification. As you didn't actually reference any explicit missing functionality with respect to the TFP specification, it seems that you're just equating compositing bugs with incomplete support.

I couldn't reference anything specifically given the lack of details provided by nVidia. I made a reasonable interpretation of the wording, which is all one can do.

I do however appreciate your clarification on the point.

Just to reinforce the uniqueness of the phrasing with respect to the use of the word "initial", compare the wording for the recent driver version, where support for your latest cards was added:

"Added support for GeForce 8800 GTS and GeForce 8800 GTX boards"

The use of the words "Added" or "Adds" is typical of prior releases when new functionality is implemented. Also words like "Improved" or "Fixed" are used when there are enhancements of prior implementations.

> > present timeline since nVidia was made aware of these issues (> 2
> > months), a fix appears to be a low priority for them and their official
> > line is that they are still "investigating the problem". Taken
> > literally, that would suggest that they have not yet even identified the
> > problem, as opposed to "we are investigating possible solutions".
>
> And you know this how?

Well, as paid employee of nVidia, perhaps you will do me the courtesy of formally correct me? That has been your stock communication on the nvNews Linux forum for weeks without further elaboration or any sign of incremental progress.

Note that I used the words `_incremental progress_`, not "final, stable, complete solution". There has been no confirmation of anyone actually working on the problem other than trying to read between the lines of the above statement. It does not mean that someone (or even more than one person) is even working on the problem full time given the other priorities that seem to be implicit in other statements also made on

Re: Beryl options and nvidia

nvNews.

Ignoring your misconceptions of how software development works for any large codebase, I can tell you that the root cause of the issue is currently fully understood.

Well, thank you Lonni. This is the first time to my knowledge that there has been any official clarification from nVidia regarding the current status of the bug.

THAT is the type of incremental progress that I was referring to. Progress sometimes just means communications back to the customer so that the customer's expectations can be appropriately set. It is the lack of communication (or lack of deviation from the stock line above) over the past two months that has led to an increasing level of frustration expressed in multiple fora.

May I respectfully suggest that you give serious consideration to posting a clarification message on this at nvNews? I think that you will find many grateful folks.

BTW, while I may not work for a large commercial hardware/firmware vendor such as nVidia, I have worked in the healthcare technology field for over 20 years, which is a highly regulated environment. So I do have an appreciation for such issues. I can tell you however, that as I reference above, communications between vendor and client are critical to maximizing client satisfaction. When client expectations get out of whack, it is usually because the vendor has failed to effectively communicate and by doing so, failed to properly set and manage the client's expectations.

However, as with any bug, its prioritization is based on a number of factors which include, but are not limited to, the severity of the issue, the percentage of users/customers experiencing the issue, and the available workarounds for the issue.

No disagreement. Again, this is an expectation management issue. However, I might challenge the perception of how many people are affected by this particular issue versus say the number of people affected by the lack of Linux support for one of your latest boards, irrespective of how much folks may have paid for them. That seems to be the basis of arguments made by some of the presumptive gamers on nvNews that their bugs are more important than the rest of us because they have the latest and greatest hardware, which many purchased without even considering whether there was Linux support available.

If their business is that critical to nVidia such that it influences the prioritization of bug fixes, then perhaps nVidia needs to consider

Re: Beryl options and nvidia

Re: Beryl options and nvidia

dedicating resources to such issues in a fashion that does not hinder the ability to address the issues that affect a larger number of existing customers.

The old 80/20 rule comes into play here when it comes to making such business decisions. Even in my company, we have dedicated devs to particular "important" clients, even going as far as putting them onsite with the client. In such a fashion, we do not hinder our ability to satisfy the service requirements of the greater number of clients in our portfolio.

If one chases new business at the expense of existing customers, the business model is destined to fail, well before you begin to approach maximal market share.

That all being said, I am not naive enough to think that any of us within the Linux community generate revenue for nVidia at a level that would warrant such consideration. The likelihood is that your key executives received more in total compensation last year than what you could reasonably identify as Linux related revenues. A figure BTW, which as a public company, is readily available. Given that the word Linux does not appear in a search of any of nVidia's recent annual or quarterly SEC filings, that tends to support the notion that Linux revenue is not material to the company, now or based upon near future projections.

The compositing "black windows" bug impacts a relatively small percentage of people, has a clear & simple workaround, and therefore the time being allocated to providing a solution is less than for bugs which have a higher priority.

Sure, as you noted in another reply here, the workaround is to not use compositing.

Given that one of the new 28 key functions targeted for FC7 is to include the Nouveau drivers:

<http://www.redhat.com/archives/fedora-devel-list/2007-January/msg00091.html>

perhaps this will all become a moot issue in April.

- > And i'm not sure why you'd want a beta
- > driver when it was replaced with a newer non-beta several weeks ago.

Because, believe it or not Lonnie, there are lots of people here (me included) who would gladly help nVidia test beta versions of your drivers to help move things along here. But none have been forthcoming, which is why I pointed out the beta page above and why I do check it frequently looking for a new beta version to test.

That is understood, and why beta versions of the driver are posted when available. Currently the latest driver version is not a beta. Believe it or not, people also complain quite vocally when the latest

Re: Beryl options and nvidia

released driver is a beta versus a non-beta. Not everyone can be pleased all of the time. Additional beta drivers will be posted when appropriate.

Again, that is a communication issue. If you are inconsistent in how and when you openly offer beta releases versus stable releases, customers are bound to get pissed off on both sides of the issue. Their perception is that you expended resources in releasing one, when you could have allocated the same resources to the other. Unless you establish some form of release policy covering when and how you release betas publicly (versus presumably having outside testers under NDAs), you will simply reinforce the tension.

Beta and even alpha testing are common for folks within this community and is something that I have done periodically with RH/FC over the past several years dating back to the late RH 8.0 betas and for other FOSS projects that I participate actively in.

Sure, but as you know the NVIDIA driver isn't a FOSS project. Applying FOSS project standards to non-FOSS software (or vice-versa) will always result in disappointment.

Communications.

Ball's in your court.

I'm not sure what that means, but sure.

<http://www.bartleby.com/68/2/702.html>

Regards,

Marc

--

fedora-list mailing list

fedora-list@xxxxxxxxxx

To unsubscribe: <https://www.redhat.com/mailman/listinfo/fedora-list>