

Re: Optimizations.

Source: <http://unix.derkeiler.com/Mailing-Lists/FreeBSD/performance/2003-05/0011.html>

From: Narvi (narvi_at_haldjas.folklore.ee)

Date: 05/16/03

Date: Fri, 16 May 2003 01:22:57 +0300 (EEST)

To: Poul-Henning Kamp <phk@phk.freebsd.dk>

On Thu, 15 May 2003, Poul-Henning Kamp wrote:

>
> *In message <20030516002105.K40030-100000@haldjas.folklore.ee>, Narvi writes:*
> > *On Thu, 15 May 2003, Marcel Moolenaar wrote:*
> > > *On Thu, May 15, 2003 at 02:30:33PM +0200, Pawel Jakub Dawidek wrote:*
> > > > *Maybe is time to think about some 'optimization team' creation?*
> > > *I think I don't want to see this happen based on professional*
> > > *experience.*
> > *Well, supposedly any such team would need to start by creating a set of*
> > *tools and benchmarks [...]*
>
> *If I have to be honest, I think this is the wrong way to approach the*
> *subject, if on no other ground than on the 3. rule of optimizations*
> *("Don't do it yet").*
>
> *While it would be nice to have a set of "blessed benchmarks" canned*
> *and ready to run, we should learn from the lmbench fiasco in Linux*
> *that such benchmarks can easier mislead than lead.*
>

Microbenchmarks can never lead you – at most they can alert you if something went wrong. Even more comprehensive ones just give you feedback unless what you are doing really is trying to find that one parameter / code sequence you can tune that would result in speedup in a large variety of relevant to the exercise apps (yes, it might be for as boring reason as 'ship criteria' or 'speed up by 10%').

> *My personal professional experience with optimizations or "Performance*
> *management" as it was called, is that you want some very rigid*
> *_functional_ testcases, which must pass at any one time so you don't*
> *unnoticed loose functionality to optimizations.*
>
> *We also know that the main performance issue is Giant, Giant and Giant.*
>

freebsd-performance: Re: Optimizations.

Ah, yes – but I was thinking more in the 1 – 1.5 year from now timeframe, which is a sort of reasonable timeframe if you don't expect miracles or a largish corporate backer employing several people to do this in addition to volunteers and do expect people getting bored, wandering off, having real life issues and work, etcetera to deal with aswell.

- > *So I really think the band of merry men we are talking about, if they*
- > *can be interested, would do much more good if they would start out*
- > *building functional and regression tests for our most critical*
- > *facilities.*
- >
- > *I can't speak for the other heavy-duty guys in the project, but I*
- > *would personally be _really_ _REALLY_ grateful if I could "cd*
- > */usr/src ; make test" and know that a significant fraction of our*
- > *functionality worked if it returned a zero exit code.*
- >

Whichever framework they come up with to produce reliable, repeatable and fair results will very probably allow you to plug in regression tests aswell or be easily modifiable for such.

If it ends up producing additions to the presently pretty small set of profiling tools, so much the better – IMHO.

- >
- > --
- > *Poul-Henning Kamp | UNIX since Zilog Zeus 3.20*
- > *phk@FreeBSD.ORG | TCP/IP since RFC 956*
- > *FreeBSD committer | BSD since 4.3-tahoe*
- > *Never attribute to malice what can adequately be explained by incompetence.*
- >

freebsd-performance@freebsd.org mailing list

<http://lists.freebsd.org/mailman/listinfo/freebsd-performance>

To unsubscribe, send any mail to "freebsd-performance-unsubscribe@freebsd.org"