

Agenda and Minutes: x86 Community Call July 2018

No new items were added to the agenda. *Minutes are added in blue.*

Attendees

Lars Kurth, Citrix
Roger Pau Monne, Citrix
Juergen Gross, Suse
Jan Beulich, Suse
Christopher Clark, OpenXT
Janakarajan Natarajan, AMD
Brian Woods, AMD
Rich Persaud, OpenXT
George Dunlap, Citrix
Wei, Andy, Paul - Citrix

Release Cadence for Xen 4.12

Following the release cadence session at the developer summit (see <https://lists.xenproject.org/archives/html/xen-devel/2018-07/threads.html#00166> & <https://docs.google.com/document/d/1W7OulSUau-FtPG6tlinD4GXYFb-hKDjaqTj84pogNrA/edit>) we have to make a decision whether

- Go on as we are for 4.12
- Move to 9 months, until we fixed the underlying issues as outlined in the thread and write-up: the problem is that unless we get some sort of commitment to address the issues, just changing the release cadence will not make a difference
- Skip a release as a one-off: Set ourselves some goals that must be achieved in this cycle around testing - this will need some commitment from vendors

I was planning to allocate up to 30 minutes to this discussion

Juergen: raises the point that keeping the release cadence at 6 months is very unfair on Jan who has raised many times that the workload resulting from having to maintain so many release branches would be too high. After running 6 monthly releases for some time, this has in fact come true, when at the time Jan's concerns were dismissed. The overhead breaks down into backporting fixes, backporting security fixes and dealing with the release mechanics.

Jan: raised the point that hardly anyone responds to calls for back-ports and if so, only send change-sets and let Jan do the backporting. Jan also says he suspects that people may not respond to backport requests, because that would require them to backport the patches.

George: points out that unless he remembers at the time he writes or reviews a patch, whether it is back-port worthy.

George and Andrew raised the idea that we could maintain a list of pending backports and assign backport tasks to people.

Jan: maintaining releases as a single person is the most efficient way of doing it. A single person doing all trees is most efficient, but then we need to restrict the number of trees. And 2 releases per year are too many.

Andrew: suggests that an even/odd releases model with different support cycles would solve this. By doing this, we would retain the discipline of doing releases.

Juergen: this would however impose the release overhead

Andrew: agrees that we need to reduce our release overhead regardless, but this issue is orthogonal from the release cadence.

Staying at 6 months we would either have to find someone who would like to carry the maintenance load, or move to a longer cadence. Also we need to make it clear that reducing the release overhead is independent from release cadence and process. We should be doing this irrespective depending on the cadence.

Juergen: We could **look at 8 months (instead of 9)it is better from a scheduling perspective (working around public holidays)**. With an 8 month release cycle, the release occurs at only 3 different dates during the calendar year, rather than the 4 dates with a 9 month cycle. This makes planning easier for selecting dates that avoid public holidays. 8 months is also closer to the 6 month cycle for those preferring shorter cadence. An 8 month cycle would not increase the number of concurrently supported branches when compared with a 9 month cycle.

ACTION: George will put together a survey for the committers outlining the issue and trade-offs and then go from there

Project Management stuff to keep the Momentum going

We have made significant progress on design related questions at the developer summit. Although not all the notes for these have been published (SGX and NVDIMM are missing, the former are on my plate). The series, which have been discussed at the summit and where I believe that good progress has been made were.

In other words, we should expect new versions of these series

Series	Stakeholders
--------	--------------

<p>Add vNVDIMM support to HVM domains</p> <p><i>As far as I understand a simple and clean way to implement this has been found, but the design session notes are still missing</i></p> <p><i>We spent almost two days on NVDIMM related discussions: we have something that should be fairly simple and easy to implement. Dan Williams is happy to take changes into upstream as long as they are sensible.</i></p> <p><i>George: the key behind the discussion was to be able to deliver a functional solution soon. We can make it nicer incrementally.</i></p> <p><i>ACTION: George will update and re-submit the NVDIMM doc (he didn't take any notes during the discussion - we are going to have to reconstruct some of the discussion)</i></p> <p><i>Andrew: Yi & Yu were taking notes in the meeting</i></p> <p><i>ACTION: Lars to reach out to Yi & Yu and see what they have</i></p>	<p>Zhang Yi, Intel Zhang Yu, Intel George Dunlap, Citrix</p>
<p>Intel Processor Trace virtualization enabling</p> <p><i>Partly blocked on CPUID & MSR</i></p> <p><i>Discussed the corner cases - these are in a PPT from Intel which Lars is waiting for. There was an open question re nested virt and a recognition that both cannot co-exist.</i></p>	<p>Luwei Kang, Intel</p>
<p>Extend resources to support more vcpus in single VM</p> <p><i>Also depends on the topology work IOREQ work needs another iteration Virtual IOMMU needs to be done</i></p>	<p>Chao Gao, Intel</p>
<p>EPT-Based Sub-page Write Protection Support</p> <p><i>Intel posted series and doesn't know what to do next due to lack of feedback. We were also lacking a plausible use-case: Intel and BitDefender are talking together to clarify the use-case. Still largely blocked on reviews.</i></p>	<p>Zhang Yi, Intel</p>
<p>SGX Virtualization design and draft patches</p> <p><i>Kai sent Lars some notes, which are published here.</i></p>	<p>Kai Huang, Intel</p>

<i>Partly blocked on CPUID & MSR</i>	
5 Level Paging <i>XPTI would become very problematic with 5 level paging. Currently Intel's lowest priority.</i>	

Then there were series which were blocked on CPUID and related work

Series	Stakeholders
Add guest CPU topology support <i>[PATCH 00/13] x86: CPUID and MSR policy marshalling support</i> has been posted on which this series depends on, but it is only covering 1/3 of the needed patches and requires some fixes. Sergey is working on the libxc side and Andrew on the hypervisor auditing/checking. Roger is working on topology support, which depends on the other three pieces.	Zhang Yi, Intel Andrew Cooper, Citrix Sergey Dyasli, Citrix Roger Pau Monne, Citrix

And other series, which are moving forward

Series	Stakeholders
paravirtual IOMMU interface <i>v2 posted recently</i>	Paul Durrant, Citrix
x86/cpuid: enable new cpu features <i>Waiting for v5</i>	Yang Zhong, Intel
add vIOMMU support with irq remapping function of virtual VT-d <i>Waiting for v5</i>	Chao Gao, Intel
AMD Avic Series <i>Waiting for v3</i>	Janakarajan Natarajan, AMD

MSR Spec Support for AMD speculative store bypass mitigations <i>Work has just started</i>	Brian Woods, AMD
Dom B <i>Waiting for Christopher's reply</i>	Christopher Clark, OpenXT
XSM Daniel De Graf on sabbatical - not sure for how long ACTION: Rich to follow up with committers@xenproject.org	