Minutes: x86 Community Call April 2018

I moved some of the sections around to better reflect dependencies and where work items are blocked to ensure that the May agenda can run more smoothly. New actions are labelled ***ACTION (April) in brown****.*

Attendees

- Lars Kurth (chair)
- Chao Gao, Luwei Kang , Yang Zhong, Yu Zhang, Zhang Yi (Intel)
- Juergen Gross, Jan Beulich (Suse)
- Daniel P Smith, Jason Andryuk, Christopher Clark, Rich Persaud (OpenXT)
- Julien Grall (ARM)
- Tamas Lengyel (Zentific, AIS)
- Brian Woods (AMD)
- Andrew Cooper, George Dunlap, Paul Durrant, Sergey Dyasli, Roger Pau Monné, Wei Liu (Citrix)

CC the following people who did not attend (due to ACTIONs)

- Suravee Suthikulpanit (AMD)
- Razvan Cojocaru (Bitdefender)

AOB

Need to have a list of dial-in numbers, such that people can call in without using GTM. Lars can't see the list of numbers (sigh), but people who have joined the meeting as guests can.

ACTION (April - DONE): George to send list to Lars
ACTION (April): Lars to include in invite

Agreed to set up a side IRC channel for discussions
ACTION (April): Lars to include blurb into invite on using #xendevel - or do we want to use something ad-hoc (e.g. #xendevel-communitycall)

Project Management stuff to keep the Momentum going

I will keep these in the agenda for progress tracking, but do not expect a discussion unless someone requests so. Open ACTIONs marked in RED, closed in blue.

[RFC XEN PATCH v4 00/41] Add vNVDIMM support to HVM domains

https://xen.markmail.org/thread/6uzmarrlws73mq5d

Plan forward

Royger to work with Zhang: write down the updated design first. Then resolve the difficult outstanding issues either by mail or if this doesn't work in a meeting.

ACTION (only sent to Royger, not on list): Haozhong Zhang to update the design doc and include it into the next version of the series (1st patch of series).

There were several off-line discussions and a whiteboard discussion. This seems even more complex than we thought. This looks like it is progressing, but Royger/George have come across gaps in relevant sections of ACPI / UEFI / NVDIMM specs, which make resolving this hard.

George: working on it, moving forward.

Yu Zhang: Intel is following up internally. Yi Zhang, will pick up development for this feature from Haozhong Zhang. There were quite a few discussions off-list, and we still need time to understand the big picture. Should have something more in 2 weeks.

Yu Zhang: Roger has proposed two different stages, but not yet sure whether this will cause any difficulties.

Royger: I introduced two stages to make it easier to review / commit.

Yu Zhang: Agrees

ACTION (April): Lars will move this into a new section called "Progressing - No further action needed" for the May meeting.

ACTION (March): Haozhong Zhang to drop the RFC and CC George and Roger
ACTION (March): Royger will help and give feedback. George will also be involved as he needs to review the memory side of the series. He will
ACTION (March): If needed - we can set up a meeting between Zhang and other stakeholders. Lars and John to take over an admin role to make sure developers can focus on the substance.

[PATCH RFC 00/10] x86 passthrough code cleanup

Sent in for meeting agenda by Wei

https://lists.xenproject.org/archives/html/xen-devel/2018-02/msg01939.html

ACTION (DONE): John - ask Kevin Tian to give a clear go/no-go decision about the direction of this series
~~**ACTION:** Janakarajan Natarajan (AMD) to follow up within AMD~~

Wei: have not heard from AMD. Got an ACK from our Xen maintainers.
Andrew: This is primarily Common code clean-up
Wei: There is dead code deletions

Julien: Had a similar issues with patches that are blocked on Suravee Suthikulpanit (AMD)

ACTION (April - DONE): Brian Wood to ping Suravee Suthikulpanit (AMD)

There was a little bit of a discussion on that Suravee is primarily working on Linux, while Brian is not yet fully up-to-speed.

ACTION (April): Lars to propose fixing CC issue in xen.git:MAINTAINERS copying the L and R section entries from Linux.git:MAINTAINERS (will need changes to get_maintainers.pl also)
ACTION (April): Lars to have a back-channel discussion with AMD on how to solve this.
Note: Brian agreed to be CC'ed

[PATCH 0/7] paravirtual IOMMU interface

Sent in for meeting agenda by George
https://marc.info/?l=xen-devel&m=151843249327749
https://xen.markmail.org/thread/kmxk4hoj2ao65qsa

ACTION (March): Paul to resend the series with a clear problem statement. It may also make sense for Andy, Paul and George to sit together

Paul: Waiting for 2+ week before this happens

ACTION (April): Lars will move this into a new section called “Waiting for Contributor Section” for the May meeting, unless there is a new development.

[PATCH v4 0/4] x86/cpuid: enable new cpu features

From: Yang Zhong
Link: https://lists.xen.org/archives/html/xen-devel/2018-01/msg00049.html

ACTION (DONE): Lars to point to the existing tool
See http://xenbits.xen.org/gitweb/?p=xen.git;a=tree;f=tools/tests/x86_emulator
Also, if you look in Jan's emulator series, most patches touch both the hypervisor and that test logic, e.g. many patches in https://xen.markmail.org/thread/roukz6r3gcuhxinn

ACTION (March): John will make sure that Yang is following up on this.

Jan: There was private follow-up

Agreement: park until patches arrive

ACTION (April): Lars will move this into a new section called “Waiting for Contributor Section” for the May meeting, unless there is a new development.

UPDATES / Design Discussions

PVHv2 Status (Royger)

Royger will give an update on PVHv2 work.

Royger:
DomU interface is stable
Dom0 is missing PCI passthrough
Initial Dom0 submitted marked as experimental

Issue: No-one working on PVHv2 Linux
Coordinating with Juergen/Boris.
Royger can to do it (but would be good to have someone else to get an independent view)

ACTION (April): Royger to add a bit more detail

PCI Emulation - Future Direction (Royger, Stefano, Julien)

Discuss PCI emulation and our future direction. Our current hybrid with QEMU is becoming increasingly problematic.

I would propose that we **cut this off at say 35 minutes after the hour**, if we do not make progress then we set up a separate forum for this, but we should get some of the issues out into the open. We should also propose a Design Session at summit, but may try to progress before by other means.

ACTION (April): Lars to set up separate meeting. Announce on list and CC Royger, Christopher, Rich, Paul, Julien, Daniel P Smith, Alexey ??? (I believe Royger mentioned his name)
Note: need e-mail addresses from Daniel P Smith, Alexey

Blocked Series

AMD AVIC Series

Andrew: looks in good state

Will get reviewed in due course

[PATCH RFC 00/14] EPT-Based Sub-page Write Protection Support (Zhang Yi)

RFC posted by Zhang Yi Oct 19, 2017
https://marc.info/?l=xen-devel&m=150840502417156
https://xen.markmail.org/thread/m75h6b2aiwk5h7fx

No acks, reviews only by memaccess maintainers / developers
Issues: Use case for the feature is still not clear and needs discussion

Lars: who needs to review?

Andrew: mainly George, Tamas, Razvan - major changes to ept2pm structure (different page table structure) and Andrew.

Andrew: did take a quick look. Nothing egregious in it. Did not have a lot of free time to look at it.

George: was going to look at it, but focus on NVDIMM first and EPT stuff second.

Tamas: provides write protection on the sub-page - my tools don't have a use-case for this. Also not sure how this would integrate into alt2pm. Razvan may have a use-case.

Andrew: no interaction with alt2pm, write protection 128 byte aligned granularity instead of 4K for write tracking a substructure within a page

Zhang Yi: I have a change in the patch which he wants to look at

George: can look at modifying that or store the type information elsewhere

Andrew: Have to shuffle the bits in the EPT tree

Summary: the issue isn't really about use-cases, but primarily about prioritizing George's bandwidth

ACTION (April): Andrew will poke Razvan (give some indication interface ios going)
ACTION (April): George will pick this up after NVDIMM has made some progress

[RFC Patch v4 0/8] Extend resources to support more vcpus in single VM

Sent in by George
RFC v3 by Lan Tianyu: https://marc.info/?l=xen-devel&m=150530044827940 (Sep 17)
RFC v4 re-posted by Chao Gao: https://xen.markmail.org/thread/tlto7b3fadp7kkw6 (Dec 17)

From: Chao Gao
Number of ACKs: 2
Quite a bit of feedback on v4 from a few people up to Feb 28th

Dependencies: Virtual interrupt remapping of virtual VT-d and Changes to IOREQ server is based on Paul Durrant's "x86: guest resource mapping".

UPDATE: Sufficient Patches are now in 4.11 to unblock that
Andrew: This also unblocks some of the introspection work

ACTION (April - DONE): Lars move into a different section

Longer Term - No Code Reviews yet

[PATCH RESEND v1 0/7] Intel Processor Trace virtulization enabling

v1.1 Posted by Kang, Luwei on 15 January 2018.
https://marc.info/?l=xen-devel&m=151608947805423
https://xen.markmail.org/thread/rbaf7cxh2a7wwchf

Issue: No feedback.

Andrew: looked at the Intel manual, which is fairly complex. There are a number of cases that are not clear (e.g. the Hypervisor should emulate X without code example). Thus, I was not in a great position to comment.

Andrew: Linux has some code now, which can be looked at as reference. This should help unblock this.

Jan: Comments that it is rather difficult to look at some series, which also require reading large specs.

George: Should we assign someone to shepherd this through?
General agreement!

George: Anyone volunteer to look at it?
No

Wei: it's in Wei's list, which is rather big
George: We had an overview at the summit, but that didn't help with implementation details

Status: This could now be looked at, but unless someone but Wei steps up, it will take some time. Realistically nothing will happen in the next month or so.

ACTION (April): Lars to talk Intel, George & Wei with a view of whether a design/discussion could be arranged at the summit and whether it makes sense. It will need some preparation by the reviewer though.

Longer Term - Waiting on dependency

[RFC PATCH 0/8] Add guest CPU topology support

Sent in for meeting agenda by George

https://marc.info/?l=xen-devel&m=151538433419631
https://xen.markmail.org/thread/od46uc5nwhshnluz

Some feedback from Andrew Cooper and Daniel De Graaf

Dependencies: Andrew's CPUID work. Currently, this version doesn't have any dependency. But Andrew thought it was on the wrong direction. So Chao decided to wait for Andrew’s work to finish and rework based on CPUID.

Andrew: CPUID work going on at the moment. Patches should arrive fairly soon (2-3 weeks). Thus, this should be unblocked soon.

ACTION (April): Lars will move this into a new section called “Recently unblocked dependencies” for the May meeting.

[RFC PATCH v2 00/17] RFC: SGX Virtualization design and draft patches

Latest Posting Date: Mon, 4 Dec 2017
Link: https://lists.xen.org/archives/html/xen-devel/2017-12/msg00104.html
From: Kai Huang
Number of ACKs: 0

Issue: No feedback.

Andrew: SGX is more complicated than Processor Trace virtualization. Does not yet understand how to understand/create an enclave.

Andrew: SGX case has a server use-case, but it is more complex. It’s a massive area with incomplete information. Can’t currently for example create an enclave without the signing infrastructure (which is not yet in place). There is also no ecosystem for native OSes as reference.

Andrew: But we do have some Hardware

Jan: Saw a large number of errata

Andrew: All of these have been fixed in uCode or ME updates

Andrew: Will need to read specs for 1-2 weeks to be able to start reviewing (due to complexity)

Rich: The signing infrastructure isn’t ready, which makes it hard to understand and experiment

ACTION (April): Lars to think about this, discuss with Intel (CC kai.huang@intel.com). Then get back to the group and discuss. Issue is about understanding how all the pieces fit together.

Dependencies: some areas of this are blocked on CPUID work (which should be resolved shortly)

Longer Term - Agreed to Pause

[PATCH v4 00/28] add vIOMMU support with irq remapping function of virtual VT-d

Sent in for meeting agenda by George
v3 posted by Lan Tianyu on 22 September 2017: marc.info/?l=xen-devel&m=150607140722407
v4 posted by Chao Gao: https://xen.markmail.org/thread/wfyorbn3nzsio6s7

Seems to have had review by Roger Pau Monne (1 ACK)
No issues

Primarily needs George as reviewer
Agreed to park this, because NVDIMM work is more important