Title: Xen Project Spectre/Meltdown FAQ (Jan 22 Update)

On January 3rd, 2018, Google's Project Zero announced several information leak vulnerabilities affecting all modern superscalar processors. Details can be found on their blog, and in the Xen Project Advisory 254. To help our users understand the impact and our next steps forward, we put together the following FAQ. We divided the FAQ into several sections to make it easier to consume:

- General considerations affecting all 3 vulnerabilities
- "Rogue Data Load" (aka SP3, "Variant 3", Meltdown, CVE-2017-5754)
- "Branch Target Injection" (aka SP2, "Variant 2", Spectre CVE-2017-5715)
- "Bounds-check bypass" (aka SP1, "Variant 1", Spectre CVE-2017-5753)

The project has been developing patches in order of exploitability. Our initial focus was on fixes for Meltdown, then on fixes for Spectre Variant 2, and finally on Variant 1.

Generally in the context of Xen based systems, there are many different considerations that have gone into our strategy, such as

- Can a guest (user or kernel space) attack the hypervisor using Meltdown or Spectre?
- Can a guest (user or kernel space) attack another guest (user or kernel space) when running in a Xen VM?

Note that impact and mitigations are specific to CPU architectures (and in some cases models) and may also differ depending on virtualization mode. The below FAQ tries to lay this out clearly, but if you have any questions, please reply to this email thread.

Note that we will update or re-issue the FAQ on this blog as new information surfaces.

1) General Questions related to all 3 vulnerabilities

1.1) Is there any risk of privilege escalation?

Meltdown and Spectre are, by themselves, *only* information leaks. There is no suggestion that speculative execution can be used to modify memory or cause the system to do anything it might not have done already.

1.2) Where can I find more information?

We will update this blog post and Advisory 254 as new information becomes available. Updates will also be published on the xen-announce@ mailing list.

1.3) Where can I ask questions?

This blog post, has been posted in text form on the xen-devel@ mailing list. If you have questions or improvement suggestions, please reply to the email thread here.

1.4) Where does development of mitigations happen?

As information related to Meltdown and Spectre is now public, development will continue in public on xen-devel@ and patches will be posted and attached to Advisory 254 as they become available.

2) SP3, "Variant 3", Meltdown, CVE-2017-5754

2.1) Is Xen impacted by Meltdown (“Variant 3”)?

Only Intel processors are impacted by Meltdown (referred to as SP3 in Advisory 254). On Intel processors, only 64-bit PV mode guests can attack Xen. Guests running in 32-bit PV mode, HVM mode, and PVH mode cannot attack the hypervisor using SP3.

Note that in general, some ARM processors are impacted by Meltdown (see https://developer.arm.com/support/security-update): however these cannot be exploited on Xen.

Guest Type	**Is a user space attack from a guest to Xen possible?**	**Is a kernel space attack from a guest to Xen possible?**	**Available Mitigations**
32 bit PV	No	No	N/A
64 bit PV	Yes	Yes	Several with different trade-offs See Question 2.2
HVM	No	No	N/A
PVH	No	No	N/A
ARM [1]	No	No	N/A

Notes:
[1] ARM’s security update refers to a subvariant of Meltdown called “Variant 3a”. The impact analysis of this variant is not yet fully complete, but we believe that no sensitive data can be leaked to exploit Xen.

2.2) Are there any patches available for Meltdown (“Variant 3”)?

The project has published five different mitigations with Advisory 254 following different mitigation strategies for Meltdown. Two strategies involve switching from PV guests to PVH or HVM guests. The others require application of patches as outlined in Advisory 254:

- **Vixen:** The basic principle is to run PV guests (which can read all of host memory due to Meltdown) as HVM guests (which cannot read memory due to Meltdown) using a hypervisor shim.
- **Comet:** The basic principle is to run PV guests (which can read all of host memory due to the hardware bugs) as PVH guests (which cannot read memory due to Meltdown) using a hypervisor shim.
- **PTI or Xen PTI stage-1:** This solution implements Page Table Isolation (PTI) for Xen.

Each strategy has different trade-offs and will work well for some use-cases, but not others. A high-level comparison of the different trade-offs for each mitigation, including information about code and documentation can be found in Advisory 254 (under “SP3 MITIGATION OPTIONS SUMMARY TABLE FOR 64-bit X86 PV GUESTS”). Please make sure you carefully read this section and the README files in the advisory.

2.3) How are Xen Guests impacted by Meltdown (“Variant 3”)?

In 32-bit PV mode, HVM mode, and PVH mode, guest user spaces *can* attack guest kernels using SP3; so updating guest kernels is advisable. Interestingly, guest kernels running in 64-bit PV mode are *not* vulnerable to attack using SP3, but attacks on user and kernel spaces of *other* guests are possible.

Guest Type	**Is a user space attack on the guest kernel possible (when running in a Xen VM)?**	**Is a user space attack on other guests possible (when running in a Xen VM)?**	**Is a kernel space attack on other guests possible (when running in a Xen VM)?**
32 bit PV	Yes [1]	No	No
64 bit PV	No [2]	Yes [3]	Yes [3]
HVM	Yes [1]	No	No
PVH	Yes [1]	No	No
ARM	Yes [1]	No	No

Mitigations and notes:

[1] Can be mitigated by the Linux KPTI patch and similar patches for other operating systems

[2] Although, a direct user space attack on the kernel is not possible, user space can indirectly be exploited via [3]. When Vixen and Comet are deployed, all guest memory is mapped by the “shim,” which is itself vulnerable to Meltdown. The Xen PTI patches protect both the hypervisor and the guest kernel from attacks from the guest user (without need for additional guest kernel patches). Note that KPTI is automatically disabled when running in 64 bit PV guests: thus running XPTI together with KPTI should not have any adverse effects.

[3] Mitigated by stage-1 Xen PTI

2.4) What is the long-term plan for Meltdown (“Variant 3”)?

Longer term, we will merge Vixen with Comet and release in suitable Xen Point releases with the codename Rudolph. In addition, we will improve PTI. We will likely backport and release PTI in suitable Xen point releases.

Note that Vixen and Comet will not be released in Xen point releases, but only through Advisory 254.

2.5) Does Xen have any equivalent to Linux's KPTI series?

Linux's KPTI series is designed to address SP3 only. For Xen guests, only 64-bit PV guests are affected by SP3. We have released a PTI (sometimes called XPTI) series, which we will continue to improve over the coming weeks.

3) SP2, "Variant 2", Spectre, CVE-2017-5715

3.1) Is Xen impacted by Spectre ("Variant 2")?

Both Intel and AMD CPUs are vulnerable to Spectre (both variants). Vulnerability of ARM processors to Spectre (both variants) varies by model and manufacturer.

Guest Type	**Is a user space attack from a guest to Xen possible?**	**Is a kernel space attack from a guest to Xen possible?**	**Available Mitigations**
x86	Yes	Yes	See Question 3.4.1
ARM 32 [1]	Yes	Yes	See Question 3.4.2
ARM 64 [1]	Yes	Yes	

Mitigations and notes:

[1] ARM has information on affected models on the following website: https://developer.arm.com/support/security-update. According to Cavium Thunder X1 is not vulnerable to Spectre (both variants).

3.2) How are Xen Guests impacted by Spectre ("Variant 2")?

Both Intel and AMD CPUs are vulnerable to Spectre (both variants). Vulnerability of ARM processors to Spectre (both variants) varies by model and manufacturer.

Guest Type	**Is a user space attack on other user processes or the guest kernel within the same guest possible (when running in a Xen VM)?**	**Is a user space attack on other guests possible (when running in a Xen VM)?**	**Is a kernel space attack on other guests possible (when running in a Xen VM)?**
x86	Yes [2]	Yes [3]	Yes [3]
ARM 32 [1]	Yes [2]	Yes [4]	Yes [4]
ARM 64 [1]	Yes [2]	Yes [5]	Yes [5]

Mitigations and notes:

[1] ARM has information on affected models on the following website: https://developer.arm.com/support/security-update. According to Cavium Thunder X1 is not vulnerable to Spectre (both variants).
[2] Mitigated by retpoline or firmware based Indirect Branch Control mitigations in guest operating systems (see here for Linux Kernel mitigations)
[3] Mitigated by “Intel and AMD CPUs” approach as outlined in question 3.4.1
[4] Mitigated by “Affected ARM CPUs” (64 bit) approach as outlined in question 3.4.2
[5] Mitigated by “Affected ARM CPUs” (32 bit) approach as outlined in question 3.4.2

3.3) Are mitigations for Spectre possible (“Variant 2”)?

SP2 can be mitigated in two ways, both of which essentially prevent speculative execution of indirect branches. The first is to flush the branch prediction logic on entry into the hypervisor. This requires microcode updates, which Intel and AMD are in the process of preparing, as well as patches to the hypervisor which are also in process and should be available soon. On ARM, firmware updates are required (see here).

The second is to do indirect jumps in a way that is not subject to speculative execution (this approach is called Retpoline). This requires the hypervisor to be recompiled with a compiler that contains special new features. These new compiler features are also in the process of being prepared for both GCC (see here and here) and clang, and should be available soon.

3.4) What is our plan for Spectre (“Variant 2”)?

3.4.1 Intel and AMD CPUs:

We have developed prototype patches for x86 CPUs. These patches depend on firmware updates. Our prototype patches were developed against pre-released versions of MSR specifications and are currently being reviewed for correctness against recently published MSR specifications (see here). This may require changes to our patches. There have also been reports of issues with some published firmware updates (see here) leading to frequent reboots of systems where these have been deployed. We are currently evaluating the situation to verify whether Xen based systems with mitigations are affected.

Once this work has been completed, we will publish Variant 2 mitigations via Advisory 254. More information on ongoing development can be found on relevant xen-devel@ discussions which are linked to from here.

3.4.2 Affected ARM CPUs:

A framework to mitigate Spectre Variant 2 has been developed (for 64 bit only) and is currently undergoing testing and backporting. A first 32 bit version of this framework has been posted for initial review. CPU vendors, will be able to add support for specific CPUs to the framework.

The framework and vendor specific mitigations will be published via Advisory 254. More information on ongoing development can be found on relevant xen-devel@ discussions which are linked to from here.

4) SP1, "Variant 1", Spectre, CVE-2017-5753

4.1) Is Xen impacted by Spectre (“Variant 1”)?

Both Intel and AMD CPUs are vulnerable to Spectre (both variants). Vulnerability of ARM processors to Spectre (both variants) varies by model and manufacturer.

Guest Type	**Is a user space attack from a guest to Xen possible?**	**Is a kernel space attack from a guest to Xen possible?**	**Available Mitigations**
x86	Yes	Yes	See Question 4.3
ARM	Yes	Yes	

Mitigations and notes:
[1] ARM has information on affected models on the following website: https://developer.arm.com/support/security-update. According to Cavium Thunder X1 is not vulnerable to Spectre (both variants).

4.2) How are Xen Guests impacted by Spectre (“Variant 1”)?

Both Intel and AMD CPUs are vulnerable to Spectre (both variants). Vulnerability of ARM processors to Spectre (both variants) varies by model and manufacturer.

Guest Type	**Is a user space attack on other user processes or the guest kernel within the same guest possible (when running in a Xen VM)?**	**Is a user space attack on other guests possible (when running in a Xen VM)?**	**Is a kernel space attack on other guests possible (when running in a Xen VM)?**
x86	Yes [2]	Yes [3]	Yes [3]
ARM [1]	Yes [2]	Yes [3]	Yes [3]

Mitigations and notes:
[1] ARM has information on affected models on the following website: https://developer.arm.com/support/security-update. According to Cavium Thunder X1 is not vulnerable to Spectre (both variants).
[2] Please refer to guest operating specific mitigations (see here for Linux Kernel mitigations)
[3] See question 4.3

4.3) Are mitigations for Spectre possible (“Variant 1”)?

Spectre Variant 1 is much more difficult to mitigate. We have some ideas we’re exploring, but they’re still at the design stage at this point.