

Introduction

This whitepaper looks at the susceptibility of Arm implementations following research findings from [security researchers, including Google and MIT](#), on new potential cache timing side-channels exploiting processor speculation. This paper also outlines possible mitigations that can be employed for software designed to run on existing Arm processors.

Overview of speculation-based cache timing side-channels

Cache timing side-channels are a well understood concept in the area of security research. As such, this whitepaper will provide a simple conceptual overview rather than an in-depth explanation.

The basic principle behind cache timing side-channels is that the pattern of allocations into the cache, and, in particular, which cache sets have been used for the allocation, can be determined by measuring the time taken to access entries that were previously in the cache, or by measuring the time to access the entries that have been allocated. This then can be used to determine which addresses have been allocated into the cache.

The novelty of speculation-based cache timing side-channels is their use of speculative memory reads. Speculative memory reads are typical of advanced micro-processors and part of the overall functionality which enables very high performance. By performing speculative memory reads to cacheable locations beyond an architecturally unresolved branch (or other change in program flow), the result of those reads can themselves be used to form the addresses of further speculative memory reads. These speculative reads cause allocations of entries into the cache whose addresses are indicative of the values of the first speculative read. This becomes an exploitable side-channel if untrusted code is able to control the speculation in such a way it causes a first speculative read of location which would not otherwise be accessible at that untrusted code. But the effects of the second speculative allocation within the caches can be measured by that untrusted code.

At this time, four variant mechanisms have been identified. Each potentially using the speculation of a processor to influence which cache entries have been allocated in a way to extract some information which would not otherwise be accessible to software.

This paper examines the nature of these four mechanisms, their state of knowledge and potential mitigations for the mechanisms in Arm software.

Variant 1 (CVE-2017-5753): bypassing software checking of untrusted values

Overview of the Mechanism

For any form of supervisory software, it is common for untrusted software to pass a data value to be used as an offset into an array or similar structure that will be accessed by the trusted software. For example, an application (untrusted) may ask for information about an open file, based on the file descriptor ID. Of course, the supervisory software will check that the offset is within a suitable range before its use, so the software for such a paradigm could be written in the form:

```
1 struct array {
2     unsigned long length;
3     unsigned char data[];
4 };
5 struct array *arr = ...;
6 unsigned long untrusted_offset_from_user = ...;
7 if (untrusted_offset_from_user < arr->length) {
8     unsigned char value;
9     value = arr->data[untrusted_offset_from_user];
10    ...
11 }
```

In a modern micro-processor, the processor implementation commonly might perform the data access (implied by line 9 in the code above) speculatively to establish value before executing the branch that is associated with the `untrusted_offset_from_user` range check (implied by line 7). A processor running this code at a supervisory level (such as an OS Kernel or Hypervisor) can speculatively load from anywhere in Normal memory accessible to that supervisory level, determined by an out-of-range value for the `untrusted_offset_from_user` passed by the untrusted software. This is not a problem architecturally as, if the speculation is incorrect, then the value loaded will be discarded by the hardware.

However, advanced processors can use the values that have been speculatively loaded for further speculation. It is this further speculation that is exploited by the speculation-based cache timing side-channels. For example, the previous example might be extended to be of the following form:

```
1 struct array {
2     unsigned long length;
3     unsigned char data[];
4 };
5 struct array *arr1 = ...; /* small array */
6 struct array *arr2 = ...; /*array of size 0x400 */
7 unsigned long untrusted_offset_from_user = ...;
8 if (untrusted_offset_from_user < arr1->length) {
9     unsigned char value;
10    value = arr1->data[untrusted_offset_from_user];
11    unsigned long index2 = ((value&1)*0x100)+0x200;
12    if (index2 < arr2->length) {
13        unsigned char value2 = arr2->data[index2];
14    }
15 }
```

In this example, `value`, which is loaded from memory using an address calculated from `arr1->data` combined with the `untrusted_offset_from_user` (line 10), is then used as the basis of a further memory access (line 13). Therefore, the speculative load of `value2` comes from an address that is derived from the data speculatively loaded for `value`.

If the speculative load of `value2` by the processor cause an allocation into the cache, then part of the address of that load can be inferred using standard cache timing side-channels. Since that address depends on data in `value`, then part of the data of `value` can be inferred using the side-channel. By applying this approach to different bits of `value`, (in a number of speculative executions) the entirety of the data of `value` can be determined.

As shown earlier, the untrusted software can, by providing out-of-range quantities for `untrusted_offset_from_user`, access anywhere accessible to the supervisory software, and as such, this approach can be used by untrusted software to recover the value of any memory accessible by the supervisory software.

Modern processors have multiple different types of caching, including instruction caches, data caches and branch prediction cache. Where the allocation of entries in these caches is determined by the value of any part of some data that has been loaded based on untrusted input, then in principle this side channel could be stimulated.

It should be noted that the above code examples are not the only way of generating sequences that can be speculated over. In particular code where at least one of the following happens is susceptible to Variant 1:

- A data address is determined from a value read from an untrusted offset
- An indirect branch destination is determined from a value read from an untrusted offset
- A branch decision is determined from a value read from an untrusted offset

As a generalization of this mechanism, it should be appreciated that the underlying hardware techniques mean that code past a branch might be speculatively executed, and so any sequence accessing memory after a branch may be executed speculatively. In such speculation, where one value speculatively loaded is then used to construct an address for a second load or indirect branch that can also be performed speculatively, that second load or indirect branch can leave an indication of the value loaded by the first speculative load in a way that could be read using a timing analysis of the cache by code that would otherwise not be able to read that value. This generalization implies that many code sequences commonly generated will leak information into the pattern of cache allocations that could be read by other, less privileged software. The most severe form of this issue is that described earlier in this section, where the less privileged software can select what values are leaked in this way.

These sorts of speculative corruptions are effectively a speculative form of a buffer overflow, and some descriptions of these mechanism use the term *speculative buffer overflow*.

Extending this mechanism to cover the Speculative Execution of Stores, also referred to as 'Bounds check bypass stores' (CVE-2018-3693)

It is common in advanced processors for the speculative execution of stores to result in their data being placed into a store buffer before the cache. This store buffer can be used by later speculative reads in the same thread of execution. In such an implementation, the following sequence could be the start of code that is susceptible to Variant 1:

```
1 struct array {
2     unsigned long length;
3     unsigned char data[];
4 };
5 struct array *arr1 = ...; /* small array */
6 struct array *arr2 = ...; /*array of size 0x400 */
7 unsigned long untrusted_offset_from_user = ...;
8 unsigned long untrusted_data_from_user = ...;
9 if (untrusted_offset_from_user < arr1->length) {
10     arr1->data[untrusted_offset_from_user] = untrusted_data_from_user;
11     ...
```

In this sequence, by selecting a suitable value of `untrusted_offset_from_user`, any location in memory can be temporarily speculatively overwritten with the `untrusted_data_from_user`. This could be exploited in a number of different ways to cause cache allocation, based on addresses generated from selected data:

- If this temporarily corrupted location is speculatively used as a return address or a function pointer, then that could allow control of the speculative execution of the processor, allowing the software to speculatively branch to code that reads suitable data and uses this to form an address for a load that causes cache allocation.
- If this temporarily corrupted location is speculatively used as a pointer to a location in memory, then that allows the selection of an arbitrary location to be read; if data read from that arbitrary location is used to form an address for a subsequent load, then this can cause cache allocation based on that data.

Exploiting this mechanism across a network

The pattern of memory access from this exploit are similar to patterns seen in software that processes network packets. Therefore, where such patterns exist, a set of network messages could be constructed to exploit this pattern. Such an exploit could cause a cache allocation of a memory location based on the value of an attacker-selected bit of data. If the response time to a second network message varies significantly as a result of the allocation of that memory location, measurement of the response time to this second message could be used to infer the value of the attacker-selected bit of data.

As a concrete example, the software for handling a network message might take a value (`x`) from the message as an indicator of a bit position in a bitstream being passed by the message. The software might then use this to index a memory location in a buffer of data being passed by the network message. In this case, the software would have the form:

```
1 if (x < bitstream_length) {
2     if (bitstream[x]) {
3         flag = TRUE;
4         ...
5     }
6 }
```

If a second network message is selected, whose response time depends on whether the variable `flag` is present in the cache, and the difference in the response time can be determined (perhaps over multiple observations), then a side-channel for communicating the value of an attacker-selected bit has been created.

Practicality of this side-channel

This side-channel has been demonstrated on several processors using code that is run in kernel space using the eBPF bytecode interpreter or JIT engine contained in the Linux kernel. The code run in this way holds a routine to perform the necessary shifting and dereferencing of the speculatively loaded data. The use of this mechanism has avoided the need to search for suitable routines in kernel space that can be directly exploited.

Note: It should be appreciated that this is one example way of exploiting the speculation. Analysis of code has shown that there are a small number of places where the value loaded using an untrusted offset is itself used to form an address to the extent that meaningful amounts of information can be retrieved using this mechanism.

It is very common that processors will speculate past an unresolved branch, and as such this is likely to be observed on cached Arm processors which perform execution out of order. For Arm processors that perform their execution in-order, there is insufficient speculative execution to allow this approach to be used to cause the necessary allocations into the cache. A definitive list of which Arm-designed processors are potentially susceptible to this issue can be found at www.arm.com/security-update.

The network form of this side-channel has been demonstrated at very low bandwidths over a wired network with relatively low jitter on the network response times. It has not been demonstrated over wireless networks.

Software Mitigations

The practical software mitigation for the scenario where the value being leaked is determined by less privileged software is to ensure that the address that is derived from the `untrusted_offset` is forced to a safe value in a way that the hardware cannot speculate past if the `untrusted_offset` is out of range.

This can be achieved on Arm implementations by using a performance optimized code sequence that mitigates speculation and enforces validation of the limits of the untrusted value. Such code sequences are based around specific data processing operations (for example conditional select or conditional move) and a new barrier instruction (CSDB). The combination of both a conditional select/conditional move and the new barrier are sufficient to address this problem on ALL Arm implementations, both current and future. The details of the new barrier are described later in this section.

It is generally unusual for sequences that allow exploitation of this side-channel to exist in privileged code. However, the compilation of byte-code supplied by a lower level of privilege is an avenue to inject such sequences into privileged software. It is particularly important that just-in-time compilers that compile such byte-code use these mechanisms as part of their compiled sequences. Arm also recommends that the provision of code injection mechanisms of this type (for example eBPF) is disabled in systems where that is practical.

Note: For Android systems, the `bpf()` syscall is not available, and the only BPF available to user space, `seccomp-bpf`, is believed to be insufficient to be able to trigger this issue.

Another area that could be subject to this issue is where there are software-enforced privilege boundaries within a single exception level, as may occur with JavaScript interpreters or Java runtimes. For example, in an interpreter, a key element of the software enforcement of privilege involves the sort of sanitization of untrusted values seen in this example, so potentially giving examples of this mechanism. Similarly, the sequences generated by a run-time compilation of Java byte-code may need to incorporate the work-around in their generated sequences.

Where it is impractical to insert this barrier, an alternative approach of inserting the combination of an DSB SYS and an ISB can be inserted to prevent speculation, but this is likely to have a much greater performance effect than using the conditional select/conditional move and CSDB barrier, and so should only be used where the conditional select/conditional move and CSDB cannot be inserted due to challenges with code generation.

Where mitigations against the network exploitation of this variant are required, the equivalent data processing and CSDB barrier, or DSB SYS + ISB, should be added to network handling code that is processing the handling of the network supplied offset.

Details of the CSDB barrier

The new barrier is called CSDB, and has the following encodings:

A64:
1101_0101_0000_0011_0010_0010_100_11111
A32:
1110_0011_0010_0000_1111_0000_0001_0100
T32:
1111_0011_1010_1111_1000_0000_0001_0100

The semantics of the barrier are:

AArch64:

1. No instruction, other than a branch instruction, appearing in program order after the CSDB can be speculatively executed using the results of any:
 - data value predictions of any instructions, or
 - PSTATE.NZCV predictions of any instructions other than conditional branch instructions, or predictions of SVE predication state for any SVE instructions appearing in program order before the CSDB that have not been architecturally resolved.

Note: For purposes of the definition of CSDB, PSTATE.NZCV or SVE prediction registers are not considered a data value. This definition permits:

- Control flow speculation before and after the CSDB.
- Speculative execution of conditional data processing instructions after the CSDB, unless they use the results of data value or PSTATE.NZCV predictions of instructions appearing in program order before the CSDB that have not been architecturally resolved.

AArch32:

1. No instruction, other than a branch instruction or other instruction that writes to the PC, appearing in program order after the CSDB can be speculatively executed using the results of any:
 - data value predictions of any instructions, or
 - PSTATE.NZCV predictions of any instructions other than conditional branch instructions or conditional instructions that write to the PC appearing in program order before the CSDB that have not been architecturally resolved.

Note: For purposes of the definition of CSDB, PSTATE.NZCV is not considered a data value. This definition permits:

- Control flow speculation before and after the CSDB.
- Speculative execution of conditional data processing instructions after the CSDB, unless they use the results of data value or PSTATE.NZCV predictions of instructions appearing in program order before the CSDB that have not been architecturally resolved.

Note on differences between current CSDB definition with v1.1 of this document.

The definition of the CSDB instruction in this document is significantly different from that presented in v1.1 (and earlier) of this document.

However, any mitigations based on the previous definition of the CSDB will be equally effective under the current definition.

Use of the Barrier

These examples show how we expect the barrier to be used in the assembly code executed on the processor.

The CSDB instruction prevents an implementation from using hardware data value prediction to speculate the result of a conditional select.

Taking the example shown previously:

```
struct array {
    unsigned long length;
    unsigned char data[];
};
struct array *arr1 = ...; /* small array */
struct array *arr2 = ...; /* array of size 0x400 */
unsigned long untrusted_offset_from_user = ...;
if (untrusted_offset_from_user < arr1->length) {
    unsigned char value;
    value = arr1->data[untrusted_offset_from_user];
    unsigned long index2 = ((value&1)*0x100)+0x200;
    if (index2 < arr2->length) {
        unsigned char value2 = arr2->data[index2];
    }
}
```

This example would typically be compiled into assembly of the following (simplified) form in AArch64:

```
LDR X1, [X2] ; X2 is a pointer to arr1->length
CMP X0, X1 ; X0 holds untrusted_offset_from_user
BGE out_of_range
LDRB W4, [X5,X0] ; X5 holds arr1->data base
AND X4, X4, #1
LSL X4, X4, #8
ADD X4, X4, #0x200
CMP X4, X6 ; X6 holds arr2->length
BGE out_of_range
LDRB X7, [X8, X4] ; X8 holds arr2->data base
out_of_range
```

The side-channel can be mitigated in this case by changing this code to be:

```
LDR X1, [X2] ; X2 is a pointer to arr1->length
CMP X0, X1 ; X0 holds untrusted_offset_from_user
BGE out_of_range
CSEL X0, XZR, X0, GE
CSDB ; this is the new barrier
LDRB W4, [X5,X0] ; X5 holds arr1->data base
AND X4, X4, #1
LSL X4, X4, #8
ADD X4, X4, #0x200
CMP X4, X6 ; X6 holds arr2->length
BGE out_of_range
LDRB X7, [X8, X4] ; X8 holds arr2->data base
out_of_range
```


For AArch32, the equivalent code is as follows:

Original code:

```
LDR R1, [R2] ; R2 is a pointer to arr1->length
CMP R0, R1 ; R0 holds untrusted_offset_from_user
BGE out_of_range
LDRB R4, [R5,R0] ; R5 holds arr1->data base
AND R4, R4, #1
LSL R4, R4, #8
ADD R4, R4, #0x200
CMP R4, R6 ; R6 holds arr2->length
BGE out_of_range
LDRB R7, [R8, R4]; R8 holds arr2->data base
out_of_range
```

Code with the mitigation added:

```
LDR R1, [R2] ; R2 is a pointer to arr1->length
CMP R0, R1 ; R0 holds untrusted_offset_from_user
BGE out_of_range
MOVGE R0, #0
CSDB
LDRB R4, [R5,R0] ; R5 holds arr1->data base
AND R4, R4, #1
LSL R4, R4, #8
ADD R4, R4, #0x200
CMP R4, R6 ; R6 holds arr2->length
BGE out_of_range
LDRB R7, [R8, R4]; R8 holds arr2->data base
out_of_range
```

Similarly for the store case (bounds check bypass stores), taking the original code:

```
1 struct array {
2     unsigned long length;
3     unsigned char data[];
4 };
5 struct array *arr1 = ...; /* small array */
6 struct array *arr2 = ...; /*array of size 0x400 */
7 unsigned long untrusted_offset_from_user = ...;
8 unsigned long untrusted_data_from_user = ...;
9 if (untrusted_offset_from_user < arr1->length) {
10     arr1->data[untrusted_offset_from_user] = untrusted_data_from_user;
11     ...
}
```

This example would typically be compiled into assembly of the following (simplified) form in AArch64:

```
LDR X1, [X2] ; X2 is a pointer to arr1->length
CMP X0, X1 ; X0 holds untrusted_offset_from_user
BGE out_of_range
STR X4, [X5,X0] ; X5 holds arr1->data base; X4 holds untrusted_data_from_user
```

The side-channel can be mitigated in this case by changing this code to be:

```
LDR X1, [X2] ; X2 is a pointer to arr1->length
CMP X0, X1 ; X0 holds untrusted_offset_from_user
BGE out_of_range
CSEL X0, XZR, X0, GE
CSDB ; this is the new barrier
STR X4, [X5,X0] ; X5 holds arr1->data base; X4 holds untrusted_data_from_user
```

For AArch32, the equivalent code is as follows:

Original code:

```
LDR R1, [R2] ; R2 is a pointer to arr1->length
CMP R0, R1 ; R0 holds untrusted_offset_from_user
BGE out_of_range
STR R4, [R5,R0] ; R5 holds arr1->data base; R4 holds untrusted_data_from_user
```

Code with the mitigation added:

```
LDR R1, [R2] ; R2 is a pointer to arr1->length
CMP R0, R1 ; R0 holds untrusted_offset_from_user
BGE out_of_range
MOVGE R0, #0
CSDB
STR R4, [R5,R0] ; R5 holds arr1->data base; R4 holds untrusted_data_from_user
```

In order to prevent this side-channel from being created in data caches, instruction caches or branch prediction caches, this mitigation approach should be used when:

- A data address is determined from a value read from an untrusted offset
- An indirect branch destination is determined from a value read from an untrusted offset
- A branch decision is determined from a value read from an untrusted offset
- The address of a store is determined by an untrusted offset

When applied to a particular code sequence involving the use of an untrusted value, this mitigation will prevent that code sequence from being able to be used to exploit this side-channel to access any data.

For some, but not all, Arm implementations, mapping particularly important secrets, such as Cryptographic keys, in Device memory will prevent their being allocated into a cache. Mapping such data in this way, where it is feasible under the operating system, could be used as an additional safeguard for those implementations, albeit at significantly increased performance cost.

Note: Arm has investigated using Device memory under Linux in this way for bulk memory, and does not believe that this additional safeguard is practical to deploy.

Tooling to help with the Software Mitigation

Details of tooling to support the software mitigation can be found at <https://developer.arm.com/compiler-support>.

Variant 2 (CVE-2017-5715): forcing privileged speculation by training branch predictors

Overview of the Method

All modern processors, including those from Arm, have a variety of different mechanisms for branch prediction that cause the processor to speculatively change the instruction stream in response to predictions of the directions and targets of future branches. The forms of such branch predictors are not described by the architecture, and implementations can employ a variety of different mechanisms to speculate the changes of instruction stream.

In order to give high-performance execution, these predictors are designed to use the history of previous branches to speculate the change of instruction stream. The resulting speculation can take considerable time to be resolved. This delay in resolution can result in the processor performing speculative memory accesses, and so cause allocation into the caches.

In some implementations, including many of those from Arm, the history of previous branches used to drive the speculation is not filtered by the exception level that the processor was in. Therefore, it is possible for the code running at one exception level to train the branch predictors in a manner that causes other exception levels (or other contexts) to perform speculative memory accesses. This can then be used to stimulate the speculation-based cache timing side-channel by having a lower exception level train the branch predictors to influence the speculative instruction stream of a higher exception level, or in a different context, to read data otherwise inaccessible at the lower exception level, and additionally to allocate items speculatively into the caches based on that data. Code running at the lower exception level can then examine the impact of cache allocations, so exploiting the cache timing side-channel.

As advanced implementations can typically have multiple outstanding speculative changes of address stream caused by branch prediction, in principle it is possible to string together a number of different pieces of privileged code – in effect to create a string of *speculation gadgets*. These gadgets are strung together using the trained predictions of the branch predictor – to construct sequences to read arbitrary data and use this data to form the addresses to allocate into the caches.

Extension of this mechanism to return stacks

One of the common mechanisms that can be used for branch prediction is a *return stack*. A return stack predicts the indirect branches that form a function return. This structure acts as a predictive stack, where a branch that calls a function pushes the address of the instruction onto the return stack, and the function return branch pops the prediction off the stack. The return address predicted is used to control speculative execution while the prediction is checked by the hardware against the architectural return information. If the return stack gets out of synchronisation with the architectural return addresses, this can lead to mis-predictions and mis-speculation.

If an attacker at one level of privilege, including a software or language managed sandbox (such as the running of JavaScript), can force the mis-synchronisation of the return stack, then that attacker may be able to determine the path speculation and choose suitable speculation gadgets to perform reads of arbitrary data. This data can then be used to form the addresses to allocate into the caches. Within JavaScript, mis-predictions from the return stack could stimulate type confusion, where register or memory contents that are expected to contain pointers to objects actually contain attacker supplied data. This potentially permits speculative memory accesses to break out of the JavaScript sandbox and access arbitrary memory that is accessible to that process. Such speculatively-accessed data

could then be used to form an address for speculative cache allocation, allowing the channel to leak the information through cache timing.

Practicality of this side-channel

This side-channel has been demonstrated on some processors by the training of indirect branches to allow the extraction of data from a KVM based hypervisor (though the same approach could also be used for an operating system kernel, or in principle for a different application or virtual machine). This demonstration used the branch predictor to speculatively enter the back-end of the eBPF bytecode interpreter contained in the host Linux kernel to run some user-space held byte-codes to perform speculatively the necessary shifting of the speculatively loaded data from the address space of the hypervisor. It should be noted that this speculative use of the eBPF bytecode interpreter does not rely on the bytecode interpreter being enabled, but can be used simply if the bytecode interpreter is present in the hypervisor (or operating system kernel) image and is marked as executable.

The use of this mechanism has avoided the need to string together speculation gadgets as part of an initial proof of concept, but it is possible that the more general approach of using speculation gadgets may be developed over time.

Most current out-of-order Arm processors have branch predictors of the form that allow training from one exception level to influence the execution at other exception levels or in other contexts. The exact mechanism of the branch predictor training varies between different implementations, and so significant reverse engineering of the branch prediction algorithms would be necessary to achieve this training. Current cached Arm processors which perform their execution in order do not exhibit sufficient speculative execution for this approach to be used to extract useful information.

The use of a mis-synchronized return stack has been shown to be a route to malicious processor speculation. A full attack for leaking data out of JavaScript requires a strong understanding of the inner workings of the JavaScript compiler and has not been fully demonstrated at the time of writing.

Arm has also added a new field to the v8.5-A architecture. This field (CSV2, which is part of the ID_AA64PFR0_EL1 register) will indicate whether the CPU has hardware updates applied to mitigate the effect of Variant 2. That field is described in the documentation for affected CPUs and will be described in a future release of the Arm Architecture Reference Manual.

A definitive list of which Arm-designed processors are potentially susceptible to this issue, as well as processors that have been updated, can be found at www.arm.com/security-update.

Software Mitigations

For Arm implementations, there is no generic mitigation available that applies for all Arm processors. However, many Arm processors have implementation specific controls that can be used to either disable branch prediction or to provide mechanisms to invalidate the branch predictor. Where an ability to invalidate the branch predictor exists, it should be used on a context switch. It should also be used on an exception entry from an exception level where it is judged that code might be used to attack a higher level. Invalidations on exception entry will likely have a nontrivial performance impact.

As many Arm implementations will not speculatively execute code at EL1 (or EL2 when HCR_EL2.E2H==1) that has been fetched from memory that is marked as Privileged Execute Never, this fact can be used in workarounds to avoid the need to invalidate branch predictors on exception entry to those exception levels. Similarly, where an

implementation has a capability to disable branch prediction, then this should be invalidated for exception levels that are judged to be particularly vulnerable to attack.

In addition, for working around the issues caused by malicious mis-synchronization of the return stack by JavaScript or other language-managed sandboxes:

1. Where possible, using site-isolation for JavaScript (where each JavaScript routine from a particular source is run in its own process) means that the JavaScript sandbox is enforced by the translation system. This means it would not be able to access any other data.
2. If site isolation is not possible, the compilers for these languages should take the following steps:
 - a. Forbid callee-saved registers.
 - b. Do the following at the return landing site (without the use of branches to avoid other Spectre issues):
 - i. Compare `lr` with `pc`.
 - ii. Clear frame and stack pointer if not equal.

Variant 3 (CVE-2017-5754): using speculative reads of inaccessible data

Overview of the Mechanism

In some, but not all, Arm implementations, a processor that speculatively performs a read from an area of memory with a permission fault (or additionally in AArch32, a domain fault) will actually access the associated location, and return a speculative register value that can be used as an address in subsequent speculative load or indirect branch instructions. If the speculation is not correct, then the results of the speculation will be discarded, so there is no architectural revelation of the data accessed at the permission faulting location. However, on some implementations, the data returned from the speculative load can be used to perform further speculation. It is this further speculation that is exploited by the speculation-based cache timing side-channels.

For example, in AArch64, a piece of EL0 code could be constructed with the form:

```
1 LDR X1, [X2] ; arranged to miss in the cache
2 CBZ X1, over ; This will be taken but
3           ; is predicted not taken
4 LDR X3, [X4] ; X4 points to some EL1 memory
5 LSL X3, X3, #imm
6 AND X3, X3, #0xFC0
7 LDR X5, [X6,X3] ; X6 is an EL0 base address
8 over
```

where:

- EL1 memory is something mapped as Kernel-only in the page table
- EL0 base address is the address of a User accessible array in memory used for the subsequent timing readings on the cache

The perturbation of the cache by the `LDR X5, [X6,X3]` (line 7) can be subsequently measured by the EL0 code for different values of the shift amount `imm` (line 5). This gives a mechanism to establish the value of the EL1 data at the address pointed to by `X4`, so leaking data that should not be accessible to EL0 code.

The equivalent situation can be used for AArch32, for PL0 code attempting to access PL1 memory:

```
LDR R1, [R2] ; arranged to miss in the cache
CMP R1, #0
BEQ over ; This will be taken but
           ; is predicted not taken
LDR R3, [R4] ; R4 points to some PL1 memory
LSL R3, R3, #imm
AND R3, R3, #0xFC0
LDR R5, [R6,R3] ; R6 is an PL0 base address
over
```

Practicality of this side-channel

For some implementations where a speculative load to a permission faulting (or in AArch32 domain faulting) memory location returns data that can be used for further speculation, this side-channel has been demonstrated to allow the leakage of EL1-only accessible memory to EL0 software. This then means that malicious EL0 applications could be written to exploit this side-channel.

A definitive list of which Arm-designed processors are potentially susceptible to this issue can be found at www.arm.com/security-update.

It is believed that at least some Arm processors designed by Arm and its architecture partners are susceptible to this side-channel, and so Arm recommends that the software mitigations described in this whitepaper are deployed where protection against malicious applications is required.

Software Mitigations

For Arm software, the best mitigation for the memory leakage with this mechanism is to ensure that when running at EL0, there are minimal mappings pointing to Kernel-only data mapped in page tables, or present in the TLB. This is done in preference to the common technique of having the Kernel-only data mapped in the translation tables but with EL1-only access permissions.

A patch to the Linux kernel for AArch64 is available from Arm to perform this mitigation, using two different ASID values for each application to prevent any TLB maintenance entailed when switching between user and privileged execution.

For information on the latest Linux Kernel patches related to this issue, please go to www.arm.com/security-update.

For support with this mitigation on other operating systems, please contact support@arm.com.

Subvariant 3a (CVE-2018-3640): using speculative reads of inaccessible data

Overview of the Mechanism

In much the same way as with the main Variant 3, in a small number of Arm implementations, a processor that speculatively performs a read of a system register that is not accessible at the current exception level, will actually access the associated system register (provided that it is a register that can be read without side-effects). This access will return a speculative register value that can be used in subsequent speculative load instructions. If the speculation is not correct, then the results of the speculation will be discarded, so there is no architectural revelation of the data from the inaccessible system register. However, on such implementations, the data returned from the inaccessible system register can be used to perform further speculation. It is this further speculation that is exploited by the speculation-based cache timing side-channels.

For example, in AArch64, a piece of EL0 code could be constructed with the form:

```
1 LDR X1, [X2] ; arranged to miss in the cache
2 CBZ X1, over ; This will be taken
3 MRS X3, TTBR0_EL1;
4 LSL X3, X3, #imm
5 AND X3, X3, #0xFC0
6 LDR X5, [X6,X3] ; X6 is an EL0 base address
7 over
```

where:

- ELO base address is the address of a User accessible array in memory used for the subsequent timing readings on the cache.

The perturbation of the cache by the `LDR X5, [X6,X3]` (line 6) can be subsequently measured by the ELO code for different values of the shift amount `imm` (line 4). This gives a mechanism to establish the value held in the `TTBR0_EL1` register so leaking data that should not be accessible to ELO code.

The equivalent situation can be used for AArch32, for PL0 code attempting to access say the `TTBR0` under a 32-bit Kernel:

```
LDR R1, [R2] ; arranged to miss in the cache
CMP R1, #0
BEQ over ; This will be taken
MRC p15, 0, R3, c2, c0, 0 ; read of TTBR0
LSL R3, R3, #imm
AND R3, R3, #0xFC0
LDR R5, [R6,R3] ; R6 is an PL0 base address
over
```

Practicality of this side-channel

This side-channel can be used to determine the values held in system registers that should not be accessible. While it is undesirable for lower exception levels to be able to access these data values, for the majority of system registers, the leakage of this information is not material.

Note: It is believed that there are no implementations of Arm processors which are susceptible to this mechanism that also implement the Pointer Authentication Mechanism introduced as part of Armv8.3-A, where there are keys held in system registers.

Arm has also added a new field to the v8.5-A architecture. This field (CSV3, which is part of the `ID_AA64PFR0_EL1` register) will indicate whether the CPU has hardware updates applied to mitigate the effect of Variant 3. That field is described in the documentation for affected CPUs, and will be described in a future release of the Arm Architecture Reference Manual.

A definitive list of which Arm-designed processors are potentially susceptible to this issue, as well as processors that have been updated, can be found at www.arm.com/security-update.

Software Mitigations

Although Arm is addressing this issue in new CPU releases, in general, it is not believed that software mitigations for this issue are necessary.

For system registers that are not in use when working at a particular exception level and which are felt to be sensitive, it would in principle be possible for the software of a higher exception level to substitute in dummy values into the system registers while running at that exception level. In particular, this mechanism could be used in conjunction with the mitigation for Variant 3 to ensure that the location of the `VBAR_EL1` while running at ELO is not indicative of the virtual address layout of the EL1 memory, so preventing the leakage of information useful for compromising KASLR.

Variant 4 (CVE-2018-3639): Speculative bypassing of stores by younger loads despite the presence of a dependency

Overview of the method

In many modern high-performance processors, a performance optimization is made whereby a load to an address will speculatively bypass an earlier store whose target address is not yet known by the hardware, but is actually the same as the address of the load. When this happens, the load will speculatively read an earlier value of the data at that address than the value written by the store. That speculatively loaded value can then be used for subsequent speculative memory accesses that will cause allocations into the cache, and the timing of those allocations can be used as an observation side-channel for the data values selected as an address.

In principle, in an advanced out-of-order processor, in any code sequence of the form:

```
STR X1, [X2]
...
LDR X3, [X4] ; X4 contains the same address as X2
<arbitrary data processing of X3>
LDR X5, [X6, X3]
```

then the second load in this sequence might be performed speculatively, using a value for X3 that was derived from the speculatively value returned in X3 from the first load. That speculatively loaded value could be taken from a value held at the first address that was from earlier in the execution of the program than the STR that overwrote that value. Any cache allocation generated by the speculative execution of the second load will reveal some information about this earlier data speculatively loaded into X3. This could be used by an attacker to circumvent situations where a store is overwriting some earlier data in order to prevent the discovery of that value.

This speculative bypassing approach be extended through a chain of speculative loads such that in this case:

```
STR X1, [X2]
...
LDR X3, [X4] ; X4 contains the same address as X2
<arbitrary data processing of X3>
LDR X5, [X6, X3]
<arbitrary data processing of X5>
LDR X7, [X8, X5]
```

then the second and third loads in this sequence might be performed speculatively, using a value for X3 that has been taken from a value held at the first address that was from earlier in the execution than the STR that overwrote that value. Any cache allocation generated by the speculative execution of the third load will reveal some information about the data in X5. In this case, if an attacker has control of the previous value held at the address pointed to be X2 and X4, then it can influence the subsequent speculation, allowing the selection of data by the second speculative load, and the revealing of the selected data by examination of the cache allocations caused by the third load.

Where the store and the first load are to the same virtual and physical address, this sort of speculative re-ordering can only occur within a single exception level.

Where the store and the first load are to different virtual addresses, but to the same physical address, the speculative re-ordering can occur between code at different exception levels, such that in this case:

```
STR X1, [X2]
...
ERET ; exception return to a lower level
...
LDR X3, [X4] ; X4 contains a different virtual address as X2, but the same physical address
<arbitrary data processing of X3>
LDR X5, [X6, X3]
```

The location loaded speculatively into the cache using X3 as an offset can be indicative of the previous data value that was at the physical address pointed to by X2 and X4.

In modern high-performance processors, it is relatively straightforward to exhibit the reordering of a store and a subsequent load to the same address, and the speculative reading of older data by such a load, if the address of the store is delayed in its availability, for example as a result of a cache miss in the generation of the store address, relative to the availability of the address of the load.

Where the store and the load use the same registers to convey the address, the processor will not commonly speculatively execute a load ahead an earlier store to the same address. However, in some micro-architecturally specific cases, it is in principle possible on some implementations. The exact conditions for this re-ordering is typically a complex function of the delays of previous memory accesses being handled by the processor.

A particular concern of this mechanism would be where the Store and the first Load are accesses onto the stack (either using the stack pointer or other registers that have the same addresses), as this is a relatively common pattern in code. In principle, this could provide a mechanism by which an earlier value that was on the stack, but has been overwritten, will control the subsequent speculation of the processor. For a privileged stack, the earlier value that was on the stack might actually be under the control of less privileged execution.

In the following sequence:

```
STR X1, [SP]
...
LDR X3, [SP] ;
<arbitrary data processing of X3>
LDR X5, [X6, X3]
<arbitrary data processing of X5>
LDR X7, [X8, X5]
```

this could then give a control channel for less privileged code having determined the value that was held on the stack before the store (perhaps as a result of a system call requesting the processing of some data) to direct the speculative load of data anywhere in the more privileged address space addresses of the processor using the second load. The result of that second load is then made observable by the fact it is used to form the address of the third load, which causes a cache allocation. The presence of that cache allocation can be detected by a classic cache timing analysis, in the same way as applies to all these side channels. In principle, this could allow the reading of arbitrary privileged data by less privileged code using the timing side-channel.

Similarly the stack could be reused with a function pointer so allowing the selection of arbitrary code to be run speculatively in the more privileged address space, as shown in this example:

```
STR X1, [SP]
...
LDR X3, [SP] ;
...
BLR X3
```

In principle, this would allow the selection of a speculation gadget to reveal interesting data.

A further form of this behavior that might be exhibited on at least some implementations is where an out-of-order processor can have a load speculatively return data from a later store in the instruction stream, as might be seen in this sequence:

```
...
LDR X3, [X4]. ;
<arbitrary data processing of X3>
LDR X5, [X6, X3]
...
STR X1, [X2] ; X2 contains the same address as X4
```

Where this occurs, the allocations in the cache by the second load could give rise to the observation of the later stored value by the use of the cache timing side-channel.

Practicality of the Side-channel

A simple proof of concept has been demonstrated on some Arm implementations, where the store has its address delayed relative to a later load to the same address, leading to later speculative memory accesses of the type described above. Those speculative memory accesses cause allocations in the cache that can, using timing side channels, reveal the value of data selected by the determination of the earlier value held in the memory location being stored to and loaded from. This was demonstrated using bespoke code to prove the concept.

The more general case of this form of bypassing, particularly where the store address is available before, or at the same time as, the load address, as typically occurs when accessing the stack, has not been demonstrated, and it would be very hard for user code to guarantee the necessary complex conditions for delaying previous memory accesses to cause the processor to the necessary re-ordering to leak such data. However, it is not possible to rule out that this mechanism might be exploitable as a low bandwidth channel to read the data from more privileged memory.

The mechanism of observing a later store by a load has not been demonstrated but is believed to be possible on at least some Arm implementations.

Software Mitigations

Arm has allocated two new barriers for use in software mitigations: SSBB and PSSBB.

Use of the SSBB barrier ensures that any stores before the SSBB using a virtual address will not be bypassed by any speculative executions of a load after the SSBB to the same virtual address. The SSBB barrier also ensures that any loads before the SSBB to a particular virtual address will not speculatively load from a store after the SSBB. This barrier can be used to prevent the speculative loads being exploited using this mechanism in cases of software managed privilege within an exception level.

The SSBB barrier is encoded using the current encoding of DSB #0 (in AArch64) or DSB with an option field of 0 (in AArch32)

Use of the PSSBB barrier ensures that any stores before the PSSBB using a particular physical address will not be bypassed by any speculative executions of a load after the PSSBB to the same physical address. The PSSBB barrier also ensures that any loads before the PSSBB to a particular physical address will not speculatively load from a store after the PSSBB. This barrier can be used to prevent the speculative loads being exploited using this mechanism when entering or leaving an OS kernel, for example.

The PSSBB barrier is encoded using the current encoding of DSB #4 (in AArch64) or DSB with an option field of 4 (in AArch32).

However, it is recognized that it would not be practical to insert an SSBB between every reuse of some memory locations, such as the stack, to mitigate the less proven exploitation of this mechanism on the stack within, say, an OS Kernel, or within a managed translated language such as JavaScript.

In many Arm processors, there is a configuration control available at EL3 that will prevent the re-ordering of stores and loads in a manner that prevent this mechanism. Setting this control will impact the performance of the system to a degree that depends on the individual implementation and workload.

For new releases of ARM CPUs, the configuration control will be architecturally defined with a new bit in PSTATE (for AArch64) or the CPSR register (for AArch32). Along with that bit, some implementations will also have instructions to directly access that bit in PSTATE. There will also be a defined field in the ID_AA64PFR1_EL1 register to indicate whether the PSTATE update has been implemented.

In some implementations, it is expected to be preferable that this configuration control is set from the start of execution to prevent the re-ordering, as in those implementations, the performance cost from such a setting is relatively low relative to the cost of dynamically changing the configuration control.

In other implementations, it is expected to be preferable for the setting of this configuration control to be used dynamically to mitigate against this mechanism for pieces of code, such as a managed language, or the OS kernel, where the risk of the currently unproven exploitation of this mechanism to read arbitrary more privileged data is deemed to be unacceptable. This gives a tradeoff between a risk assessment and performance that can be changed between different pieces of code.

Arm will provide a standard SMC calls for the enabling and disabling of this control dynamically for implementations where the dynamic approach is preferred; the installation of this SMC will require an update to the Trusted Firmware running at EL3.

www.arm.com/security-update

Document history

Version/Issue	Date	Confidentiality	Change
1.0	3 January 2018	Non-Confidential	First release for Arm Internal use only.
1.1	3 January 2018	Non-confidential	First release.
1.1	4 January 2018	Non-confidential	Fixed minor errors in code examples.
1.1	19 January 2018	Non-confidential	Fixed minor typographical error.
1.2	20 February 2018	Non-confidential	Make it clear that examples for Variant 1 are not the only form of source that are susceptible to Variant 1. Update description of CSDB instruction and usage. Update Variant 3 <i>Overview of mechanism</i> to include indirect branches.
1.3	22 March 2018	Non-confidential	Updated the definition of CSDB. Addition to cover the SVE predications. Updated code examples.
2.0	21 May 2018	Non-confidential	Variant 4 addition.
2.1	23 May 2018	Non-confidential	Fixed minor typographical error.
2.2	24 July 2018	Non-confidential	Updated to detail 'speculative execution of stores' subvariant of Variant 1. With thanks to Vladimir Kiriansky of MIT. Updated to detail 'return stacks extension' to Variant 2. With thanks to Giorgi Maisuradze of the University of Saarland and Nael Abu-Ghazaleh at University of California, Riverside.
2.3	27 July 2018	Non-confidential	Updated to detail network exploitation of Variant 1 With thanks to Michael Schwarz, Daniel Gruss, Martin Schwarzl, Moritz Lipp and Stefan Mangard, all of Graz University of Technology.
2.4	10 September 2018	Non-confidential	Updated to refer to v8.5 hardware fixes.
2.4	12 October 2018	Non-confidential	Fixed minor typographical errors.
2.5	08 June 2020	Non-confidential	Updated to highlight Variant 2 (CVE-2017-5715) software workaround mechanism to avoid the need to invalidate branch prediction on exception entry to EL1, and EL2.