



NetSpectre

Read Arbitrary Memory over Network

Michael Schwarz¹

Martin Schwarzl¹


Moritz Lipp¹



Jon Masters²




Daniel Gruss¹

¹Graz University of Technology

²Red Hat









We want to build a Spectre attack which...



We want to build a Spectre attack which...

- is capable of leaking secrets from a remote system



We want to build a Spectre attack which...

- is capable of leaking secrets from a remote system
- has neither physical access nor code execution on system



We want to build a Spectre attack which...

- is capable of leaking secrets from a remote system
- has neither physical access nor code execution on system
- does not rely on software vulnerabilities

CVSS v3 for CVE-2017-5753 (Spectre)

Attack Vector

Network	Adjacent Network	Local	Physical
---------	------------------	-------	----------

CVSS v3 for CVE-2017-5753 (Spectre)

Attack Vector



Attack Complexity



CVSS v3 for CVE-2017-5753 (Spectre)

Attack Vector



Attack Complexity



Privilege Required



CVSS v3 for CVE-2017-5753 (Spectre)

Attack Vector**Attack Complexity****Privilege Required****User Interaction**

`index = 0`

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z


```
if (index < 4)
```

then

else

```
glyph[data[index]]
```

```
{
```



Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

index = 0

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
if (index < 4)
```

then

```
glyph[data[index]]
```


else

Speculate

```
}
```

Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	




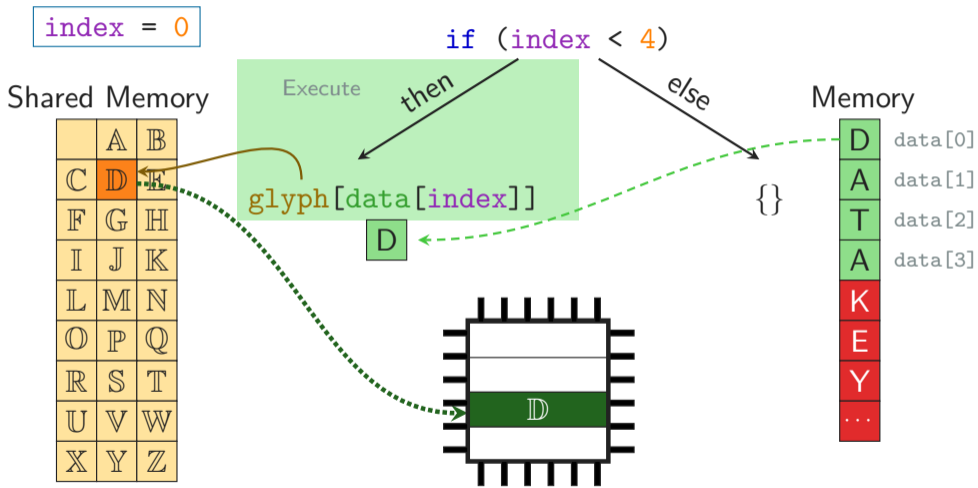
index = 0

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z









`index = 1`

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

index = 1


Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z



Memory



D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

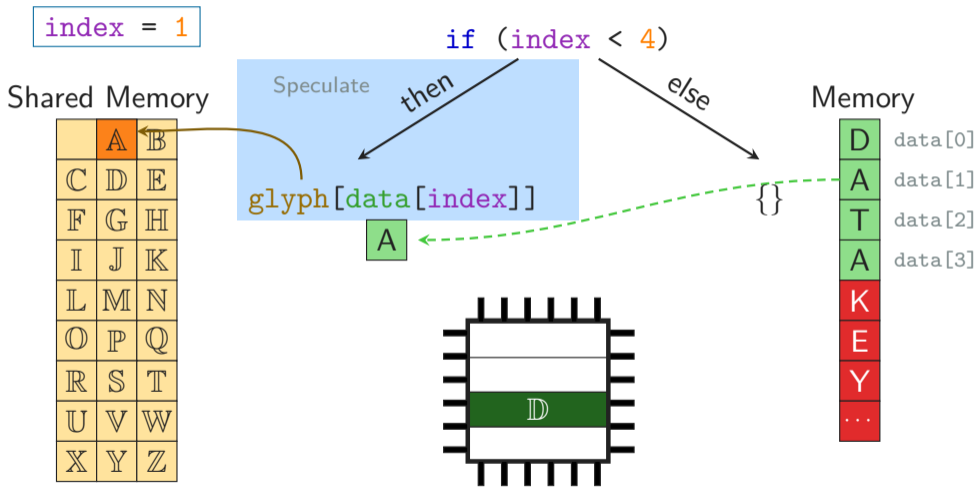


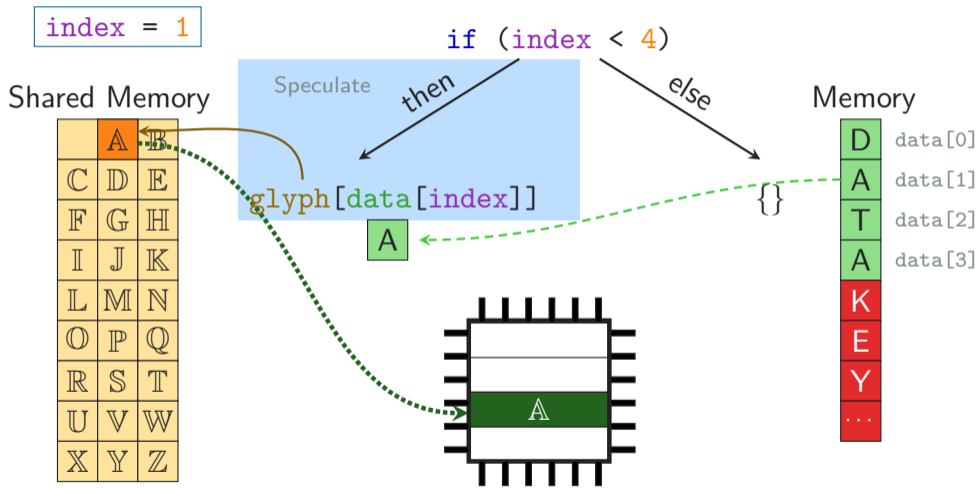
index = 1

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z







index = 1

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
if (index < 4)
```

Execute

then


```
glyph[data[index]]
```

else

```
{
```

Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



`index = 2`


Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	





index = 2

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z









index = 2

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

`index = 3`


Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 3

Shared Memory


	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z





A

Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	









index = 3

Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

`index = 4`


Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 4

Shared Memory


	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




K

Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	






index = 4


Shared Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z



Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	





Spectre **without code execution** is complicated



Spectre **without code execution** is complicated

- Which **branch** can be exploited



Spectre **without code execution** is complicated

- Which **branch** can be exploited
- Cannot observe the **cache state**



Spectre **without code execution** is complicated

- Which **branch** can be exploited
- Cannot observe the **cache state**
- Spectre **gadgets** will be different



Spectre **without code execution** is complicated

- Which **branch** can be exploited
- Cannot observe the **cache state**
- Spectre **gadgets** will be different
- No **timing** measurement on the attacked system



Spectre **without code execution** is complicated

- Which **branch** can be exploited
- Cannot observe the **cache state**
- Spectre **gadgets** will be different
- No **timing** measurement on the attacked system
- How to select the **data** to leak



- No code can be injected



- No code can be injected
- Public interface (**API**) accessing data



- No code can be injected
- Public interface (**API**) accessing data
- Branches in API can be mistrained remotely



- No code can be injected
- Public interface (**API**) accessing data
- Branches in API can be mistrained remotely
- Attacker only calls the API via **network requests**


```
def check_user_privileges(user_id):  
    [...]  
    if user_id < len(users):  
        if test_bit(privileges, user_id) == True:  
            admin = True  
  
    return SUCCESS
```

```
def check_user_privileges(user_id):  
    [...]  
    if user_id < len(users): Bounds check  
        if test_bit(privileges, user_id) == True:  
            admin = True  
  
    return SUCCESS
```

```
def check_user_privileges(user_id):  
    [...]  
    if user_id < len(users):  
        if test_bit(privileges, user_id) == True:  
            admin = True  
  
    return SUCCESS
```

Bounds check

Speculative
out-of-bounds
read

```
def is_admin():  
    return admin
```

```
def is_admin():  
    return admin
```



```
def is_admin():  
    return admin
```



- If **bit** in array was **set** → admin is **cached**

```
def is_admin():  
    return admin
```



- If **bit** in array was **set** → admin is **cached**
- If **bit** was **not set** → admin is **not cached**

```
def is_admin():  
    return admin
```



- If **bit** in array was **set** → admin is **cached**
- If **bit** was **not set** → admin is **not cached**
- Observe cache state via function execution time



- Cannot measure time directly on the attacked system




- Cannot measure time directly on the attacked system
- Network **latency** depends on API **execution time**



- Cannot measure time directly on the attacked system
 - Network **latency** depends on API **execution time**
- Measure the network **roundtrip time**



- Cannot measure time directly on the attacked system
 - Network **latency** depends on API **execution time**
- Measure the network **roundtrip time**
- Reveals whether the variable is cached





- After **measuring** variable is always **cached**




- After **measuring** variable is always **cached**
- How do we **evict** the variable?




- After **measuring** variable is always **cached**
- How do we **evict** the variable?
- Constantly evict the cache via a **file download**




- After **measuring** variable is always **cached**
- How do we **evict** the variable?
- Constantly evict the cache via a **file download**
- Thrash+Reload → crude form of Evict+Reload

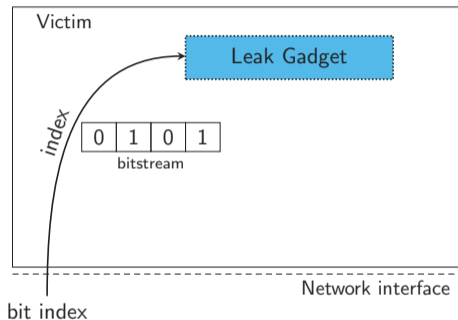





```
if (x < bitstream_length)
    if(bitstream[x])
        flag = true
```



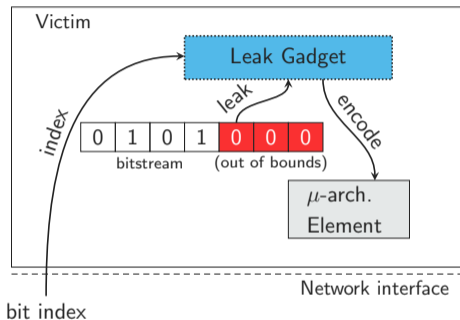
```
if (x < bitstream_length)
  if(bitstream[x])
    flag = true
```




```
if (x < bitstream_length)
  if(bitstream[x])
    flag = true
```



```
if (x < bitstream_length)
  if(bitstream[x])
    flag = true
```




```
if (x < bitstream_length)
  if(bitstream[x])
    flag = true
```




```
if (x < bitstream_length)
  if(bitstream[x])
    flag = true
```

```
send(flag)
```





```


if (x < bitstream_length)
    if(bitstream[x])
        flag = true
        send(flag)
    
```




- **Mistrain** branch predictor with in-bounds requests



- **Mistrain** branch predictor with in-bounds requests
- **Evict** everything from cache via file download



- **Mistrain** branch predictor with in-bounds requests
- **Evict** everything from cache via file download
- **Leak** a bit: do nothing ('0') or cache a memory location ('1')




- **Mistrain** branch predictor with in-bounds requests
- **Evict** everything from cache via file download
- **Leak** a bit: do nothing ('0') or cache a memory location ('1')
- **Measure** function latency which uses the memory location




Leaking byte 'd' (0)




Leaking byte 'd' (01)




Leaking byte 'd' (011)




Leaking byte 'd' (0110)




Leaking byte 'd' (01100)



Leaking byte 'd' (011001)



Leaking byte 'd' (0110010)



Leaking byte 'd' (01100100)

- Several possible attack targets

- Several possible attack targets
- Different impacts depending on target

- Several possible attack targets
- Different impacts depending on target



Web/FTP Servers
(user gadget)

- Several possible attack targets
- Different impacts depending on target



Web/FTP Servers
(user gadget)



SSH Daemons
(user gadget)

- Several possible attack targets
- Different impacts depending on target



Web/FTP Servers
(user gadget)



SSH Daemons
(user gadget)

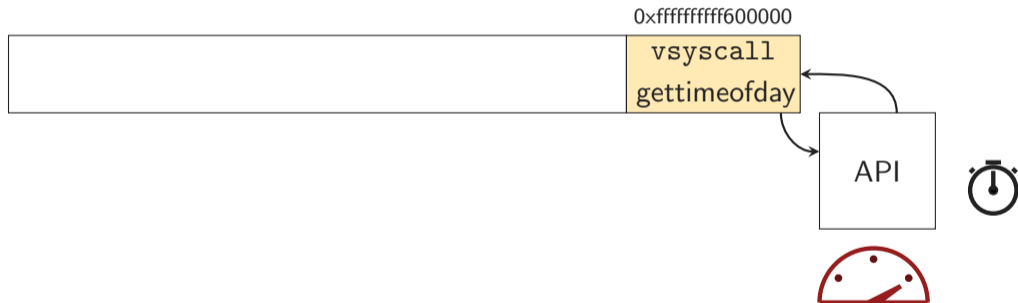



Network Drivers
(kernel gadget)

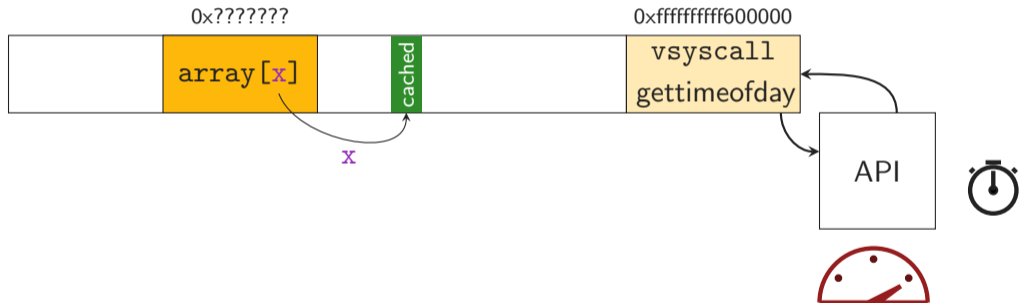
- No indirection, **simple array access**


- No indirection, **simple array access**


```
if (x < array_length)
    y = array[x];
```

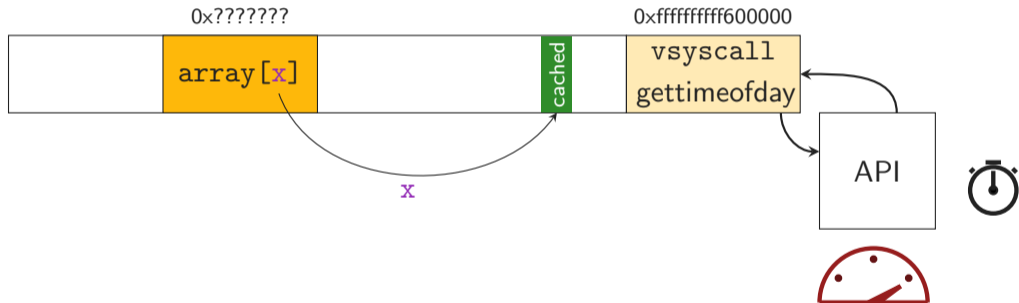


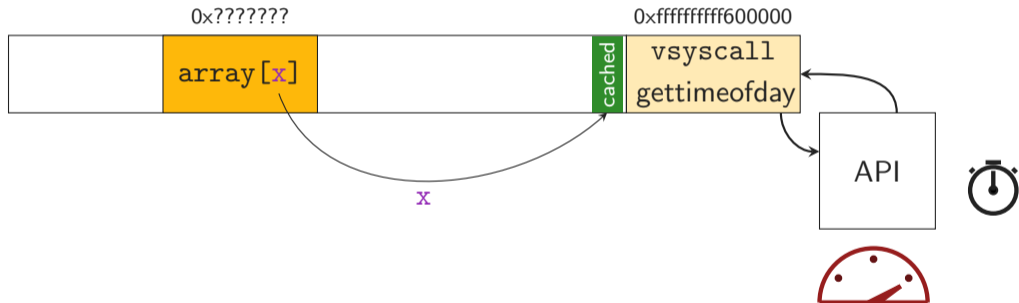

















- 256-bit instructions need a lot of **power**



- 256-bit instructions need a lot of **power**
 - On Intel, **disabled by default**, enabled on first use



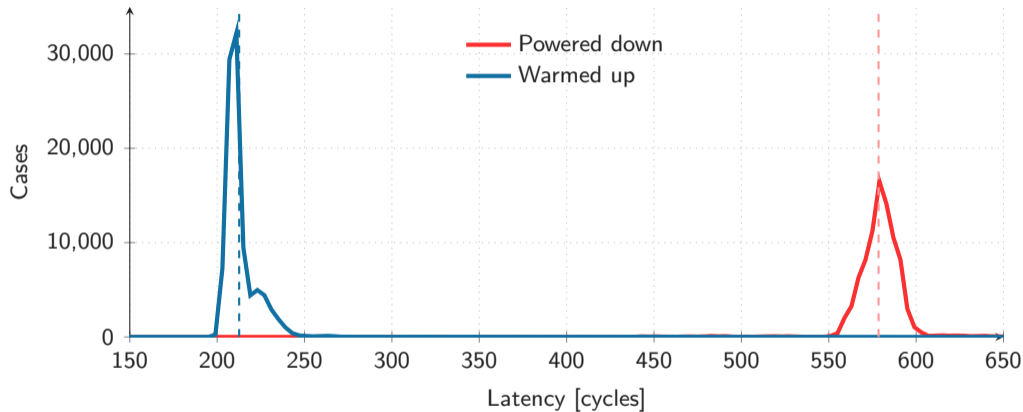
- 256-bit instructions need a lot of **power**
 - On Intel, **disabled by default**, enabled on first use
- Requires some time to power up



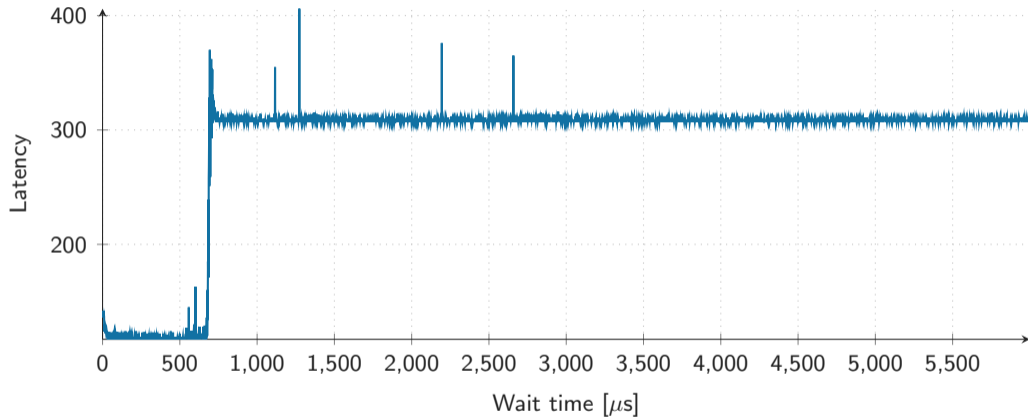
- 256-bit instructions need a lot of **power**
 - On Intel, **disabled by default**, enabled on first use
- Requires some time to power up
- Measure execution time of AVX instruction




- 256-bit instructions need a lot of **power**
 - On Intel, **disabled by default**, enabled on first use
 - Requires some time to power up
 - Measure execution time of AVX instruction
- **Leak** timing information




```
if (x < bitstream_length)
    if(bitstream[x])
        _mm256_instruction();
```







1. **Mistrain** branch predictor with in-bounds requests




1. **Mistrain** branch predictor with in-bounds requests
2. **Wait** for AVX unit to power off (1ms)



1. **Mistrain** branch predictor with in-bounds requests
2. **Wait** for AVX unit to power off (1ms)
3. **Leak** a bit: do nothing ('0') or power AVX unit ('1')



1. **Mistrain** branch predictor with in-bounds requests
2. **Wait** for AVX unit to power off (1ms)
3. **Leak** a bit: do nothing ('0') or power AVX unit ('1')
4. **Measure** function latency which uses AVX instruction



- Local Network (1 000 000 measurements/bit)



- **Local** Network (1 000 000 measurements/bit)



30 min/byte

- **Local** Network (1 000 000 measurements/bit)



30 min/byte



8 min/byte

- **Local** Network (1 000 000 measurements/bit)



30 min/byte



8 min/byte

- **Cloud** (20 000 000 measurements/bit)



- **Local** Network (1 000 000 measurements/bit)



30 min/byte



8 min/byte

- **Cloud** (20 000 000 measurements/bit)



1 h/bit





- NetSpectre requires a **fast and stable network** connection



- NetSpectre requires a **fast and stable network** connection
 - Local networks



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)
- Internet speeds improve (e.g., fiber, 5G)



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)
 - Internet speeds improve (e.g., fiber, 5G)
- possible in the **near future**?



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)
 - Internet speeds improve (e.g., fiber, 5G)
- possible in the **near future**?
- Attack speeds can be drastically **improved**



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)
 - Internet speeds improve (e.g., fiber, 5G)
- possible in the **near future**?
- Attack speeds can be drastically **improved**
 - Better signal processing/filtering



- NetSpectre requires a **fast and stable network** connection
 - Local networks
 - Data centers (VM to VM attack)
 - Internet speeds improve (e.g., fiber, 5G)
- possible in the **near future**?
- Attack speeds can be drastically **improved**
 - Better signal processing/filtering
 - Dedicated measuring hardware



- Speculative execution **leaks** secrets **without** exploiting **bugs**
- Spectre attacks are **not limited** to **local** attackers
- Spectre attacks have a **larger impact** than assumed



NetSpectre

Read Arbitrary Memory over Network

Michael Schwarz¹

Martin Schwarzl¹

Moritz Lipp¹

Jon Masters²

Daniel Gruss¹

¹Graz University of Technology

²Red Hat