

kAFL: Hardware-Assisted Feedback Fuzzing for OS Kernels

Sergej Schumilo¹, Cornelius Aschermann¹, Robert Gawlik¹, Sebastian Schinzel², Thorsten Holz¹

¹Ruhr-Universität Bochum, ²Münster University of Applied Sciences

Motivation

IJG jpeg libjpeg-turbo libpng libtiff mozjpeg PHP Mozilla Firefox Internet Explorer PCRE sqlite OpenSSL LibreOffice poppler freetype GnuTLS GnuPG PuTTY ntpd nginx bash tcpdump JavaScriptCore pdfium ffmpeg libmatroska libarchive ImageMagick BIND QEMU lcms Adobe Flash Oracle BerkeleyDB Android libstagefright iOS ImageIO FLAC audio library libsndfile less lesspipe strings file dpkg rcs systemd-resolved libyaml Info-Zip unzip libtasn1OpenBSD pfctl NetBSD bpf man mandocIDA Pro clamav libxml2glibc clang llvmnasm ctags mutt procmail fontconfig pdksh Qt wavpack OpenSSH redis lua-cmsgpack taglib privoxy perl libxmp radare2 SleuthKit fwknop X.Org exifprobe jhead capnproto Xerces-C metacam djvulibre exiv Linux btrfs Knot DNS curl wpa_supplicant Apple Safari libde265 dnsmasq libbpg lame libwmf uudecode MuPDF iplib2 libraw libbson libsass yara W3C tidy-html5 VLC FreeBSD syscons John the Ripper screen tmux mosh UPX indent openjpeg MMIX OpenMPT rxvt dhcpcd Mozilla NSS Nettle mbed TLS Linux netlink Linux ext4 Linux xfs botan expat Adobe Reader libav libical OpenBSD kernel collectd libidn MatrixSSL jasperMaraDNS w3m Xen OpenH232 irssi cmark OpenCV Malheur gstreamer Tor gdk-pixbuf audiofilezstd lz4 stb cJSON libpcre MySQL gnulib openexr libmad ettercap lrzip freetds Asterisk ytnefraptor mpg123 exempi libgmime pev v8 sed awk make m4 yacc PHP ImageMagick freedesktop.org patch libtasn1 libvorbis zsh lua ninja ruby busybox gcrypt vim Tor poppler libopus BSD sh gcc qemu w3m zsh dropbear wireshark libtorrent git rust gravity e2fsprogs parrot lodepng json-glib cabextract libmspack qprint gpsbabel dmg2img antiword arj unrar unace zoo rzip lrzip libiso libtta duktape splint zpaq assimp cppcheck fasm catdoc pngcrush cmark p7zip libjbig2 aaphoto t1utils apngopt sqlparser mdp libtinyxml freexl bgpparser testdisk photorec btcd gumbo chaiscript teseq colcrt pttbbs capstone dex2oat pillow elftoolchain aribas universal-ctags uriparser jq lha xdelta gnuplot libwpd teseq cimg libiberty policycoreutils libsemanage renoise metapixel openclone mp3splt podofo Apache httpd glslang UEFITool libcbor lldpd pngquant muparserx mochilo pyhocon sysdig Overpass-API fish-shell gumbo-parser mapbox-gl-native rapidjson libjson FLIF MultiMarkdown astyle pax-utils zziplib PyPDF spiffing apk pgpdump icoutils msitools dosfstools

Motivation

Internet Explorer

IJG jpeg libjpeg-turbo libpng libtiff mozjpeg PHP Mozilla Emacs Erlang OpenSSL LibreOffice poppler freetype GnuTLS GnuPG PuTTY ntpd nginx bash tcpdump JavaScriptCore pdrium ffmpeg libmatroska libarchive ImageMagick BIND QEMU lcms Adobe Flash Oracle BerkeleyDB Android libstagefright iOS ImageIO FLAC audio library libsndfile less lesspipe strings file dpkg rcs systemd-resolved libyaml Info-Zip unzip libtasn1OpenBSD pfctl NetBSD bpf man mandocIDA Pro clamav libxml2glibc clang llvmnasm ctags mutt procmail fontconfig pdksh Qt wavpack OpenSSH redis lua-cmsgpack taglib privoxy perl libxmp radare2 SleuthKit fwknop X.Org exifprobe jhead capnproto Xerces-C metacam djvulibre exiv Linux btrfs Knot DNS curl wpa_supplicant Apple Safari libde265 dnsmasq libbpg lame libwmf uudecode MuPDF imlib2 libraw libbson libsass yara W3C tidy-html5 VLC FreeBSD syscons John the Ripper screen tmux mosh UPX indent openjpeg MMIX OpenMPT rxvt dhcpcd Mozilla NSS Nettle mbed TLS Linux netlink Linux ext4 Linux xfs botan expat Adobe Reader libav libical OpenBSD kernel collectd libidn MatrixSSL jasperMaraDNS w3m Xen OpenH232 irssi cmark OpenCV Malheur gstreamer Tor gdk-pixbuf audiofilezstd lz4 stb cJSON libpcre MySQL gnulib openexr libmad ettercap lrzip freetds Asterisk ytnefraptor mpg123 exempi libgmime pev v8 sed awk make m4 yacc PHP ImageMagick freedesktop.org patch libtasn1 libvorbis zsh lua ninja ruby busybox gcrypt vim Tor poppler libopus BSD sh gcc qemu w3m zsh dropbear wireshark libtorrent git rust gravity e2fsprogs parrot lodepng json-glib cabextract libmspack qprint gpsbabel dmg2img antiword arj unrar unace zoo rzip lrzip libiso libtta duktape splint zpaq assimp cppcheck fasm catdoc pngcrush cmark p7zip libjbig2 aaphoto t1utils apngopt sqlparser mdp libtinyxml freexl bgpparser testdisk photorec btcd gumbo chaiscript teseq colcrt pttbbs capstone dex2oat pillow elftoolchain aribas universal-ctags uriparser jq lha xdelta gnuplot libwpd teseq cimg libiberty policycoreutils libsemanage renoise metapixel openclone mp3splt podofo Apache httpd glslang UEFITool libcbor lldpd pngquant muparserx mochilo pyhocon sysdig Overpass-API fish-shell gumbo-parser mapbox-gl-native rapidjson libjson FLIF MultiMarkdown astyle pax-utils zziplib PyPDF spiffing apk ppgdump icoutils msitools dosfstools

Motivation

Internet Explorer

Adobe Flash

OpenSSH

Apple Safari

Motivation

Internet Explorer

Adobe Flash

OpenSSH

Apple Safari

Adobe Reader

Motivation

Internet Explorer

Adobe Flash

OpenSSH

Apple Safari

Adobe Reader

Wireshark

Motivation

Adobe Flash

Internet Explorer

Apple Safari

OpenSSH

Adobe Reader

Wireshark

Apache httpd

Motivation

Stephens, Nick, et al. "Driller: Augmenting Fuzzing Through Selective Symbolic Execution." Proceedings of the Network and Distributed System Security Symposium (NDSS). 2016.

Böhme, Marcel, et al. "Coverage-based greybox fuzzing as markov chain." Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. ACM. 2016.

Rawat, Sanjay, et al. "Vuzzer: Application-aware evolutionary fuzzing." Proceedings of the Network and Distributed System Security Symposium (NDSS). 2017.

Motivation

What about Kernel Security?

Related Work

	Fast	Crash Tolerant	OS Independent	Binary Only

Related Work

	Fast	Crash Tolerant	OS Independent	Binary Only
TriforceAFL (Jesse Hertz & Tim Newsham, NCC Group)	X	✓	~	✓

Related Work

	Fast	Crash Tolerant	OS Independent	Binary Only
TriforceAFL (Jesse Hertz & Tim Newsham, NCC Group)	X	✓	~	✓
Syzkaller (Dmitry Vyukov)	✓	✓	X	X

Related Work

	Fast	Crash Tolerant	OS Independent	Binary Only
TriforceAFL (Jesse Hertz & Tim Newsham, NCC Group)	X	✓	~	✓
Syzkaller (Dmitry Vyukov)	✓	✓	X	X
AFL Filesystem Fuzzer (Vegard Nossum & Quentin Casanovas, Oracle)	✓	~	X	X

Related Work

	Fast	Crash Tolerant	OS Independent	Binary Only
TriforceAFL (Jesse Hertz & Tim Newsham, NCC Group)	X	✓	~	✓
Syzkaller (Dmitry Vyukov)	✓	✓	X	X
AFL Filesystem Fuzzer (Vegard Nossum & Quentin Casanovas, Oracle)	✓	~	X	X
PT Kernel Fuzzer (Richard Johnson, Talos)	✓	X	X	✓

Goal

Build a **kernel** fuzzer that is all of this:

- Fast
- Reliable
- (mostly) OS independent
- No source level access required

Fuzzing in a nutshell

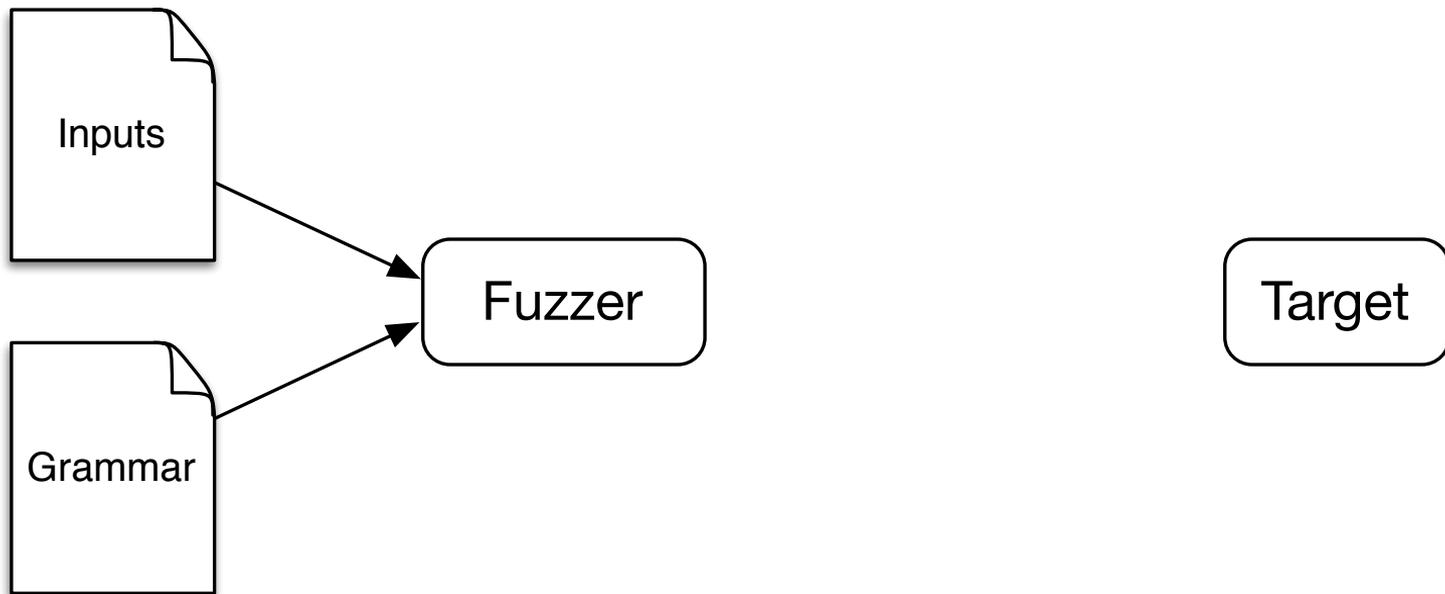
Fuzzing in a nutshell

Fuzzer

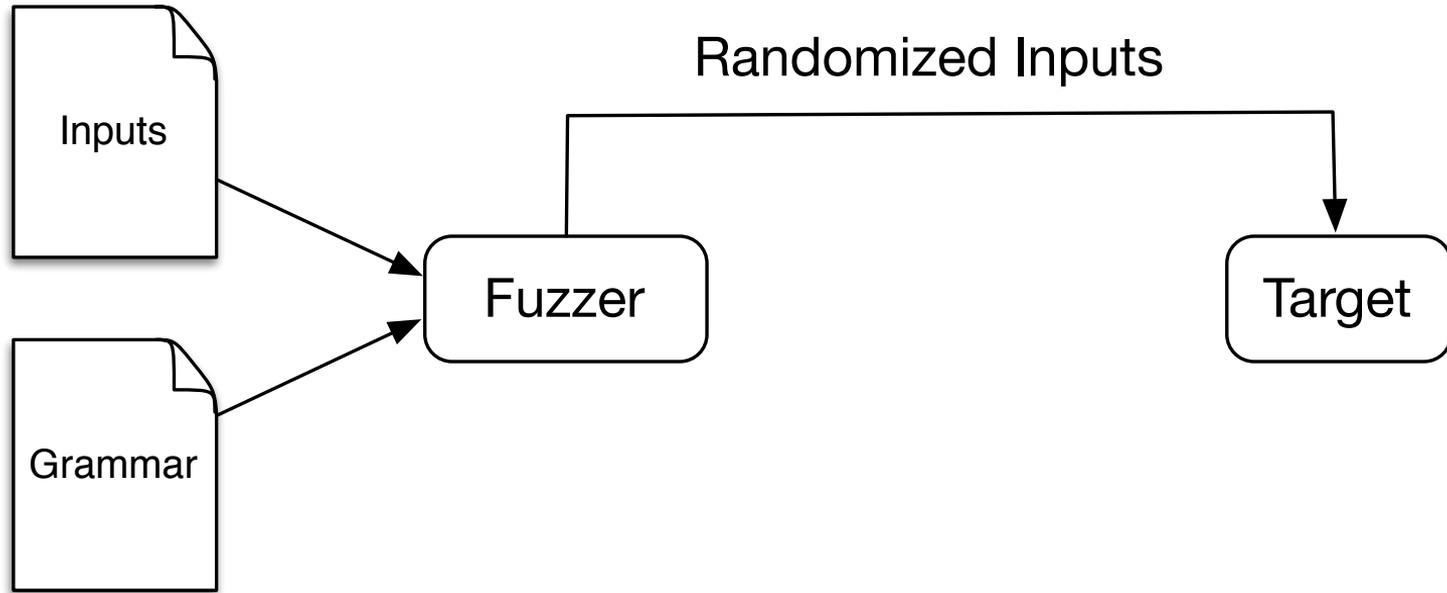
The diagram consists of two rounded rectangular boxes. The first box on the left contains the word 'Fuzzer'. The second box on the right contains the word 'Target'. There are no arrows or other graphical elements connecting the two boxes.

Target

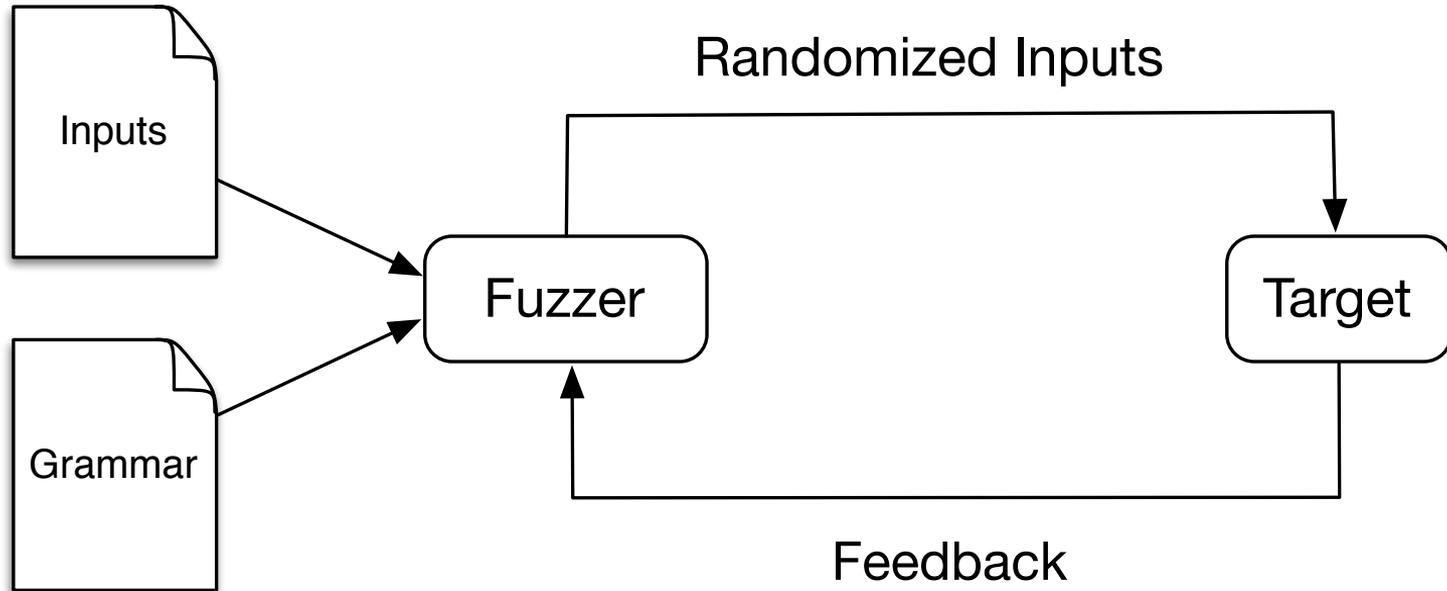
Fuzzing in a nutshell



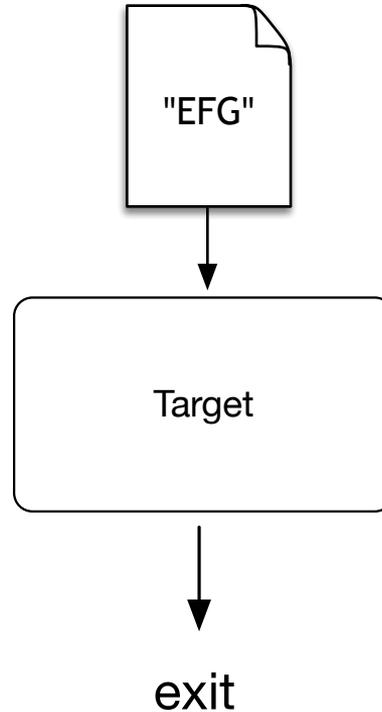
Fuzzing in a nutshell



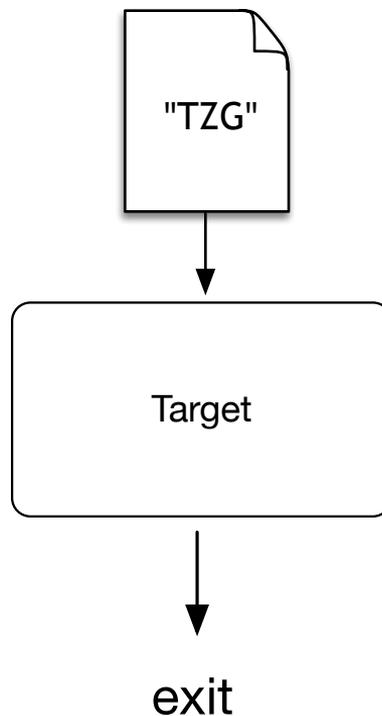
Fuzzing in a nutshell



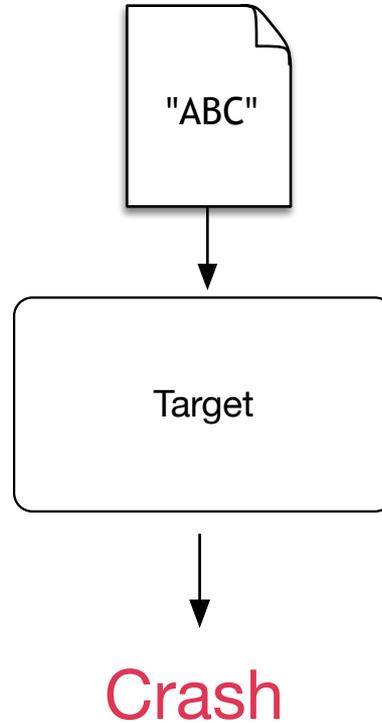
Blackbox Fuzzing



Blackbox Fuzzing

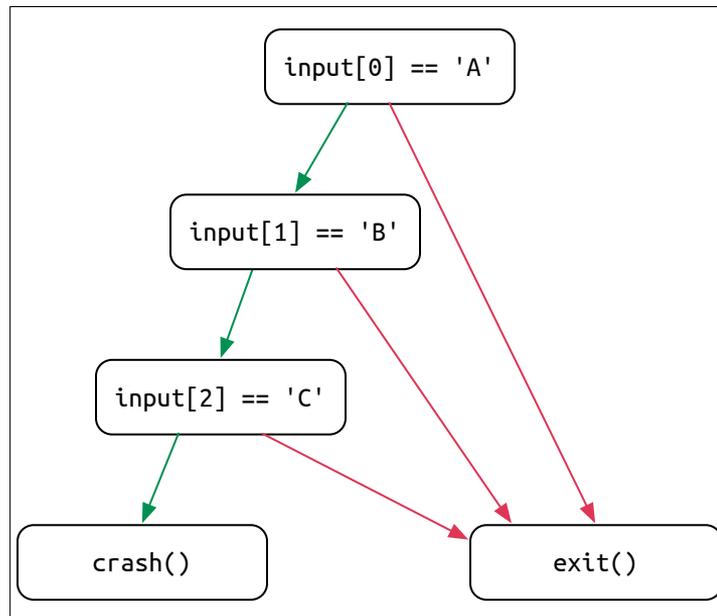


Blackbox Fuzzing

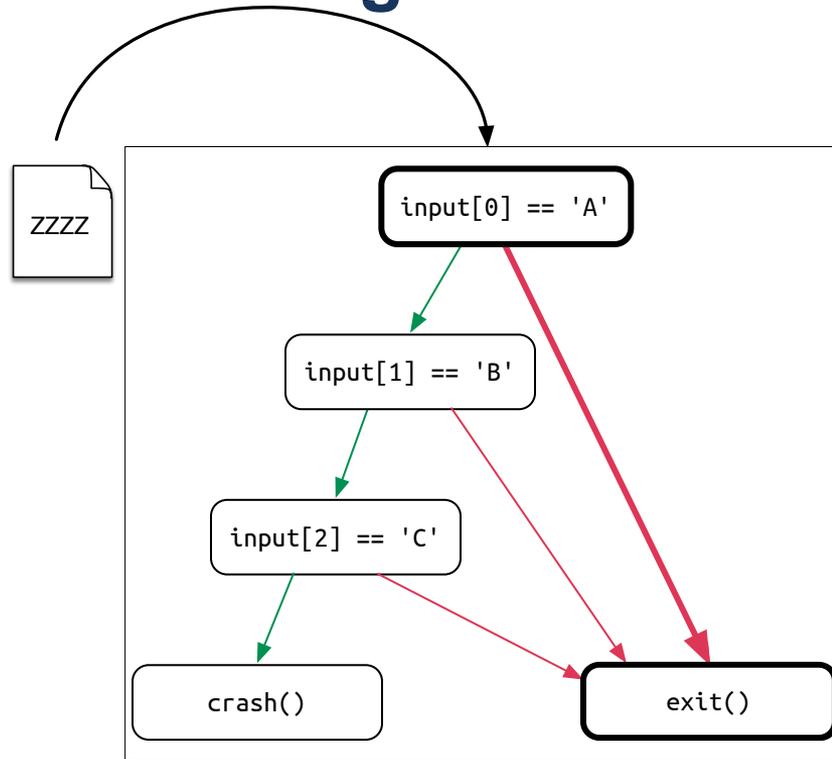


Coverage-Guided Fuzzing

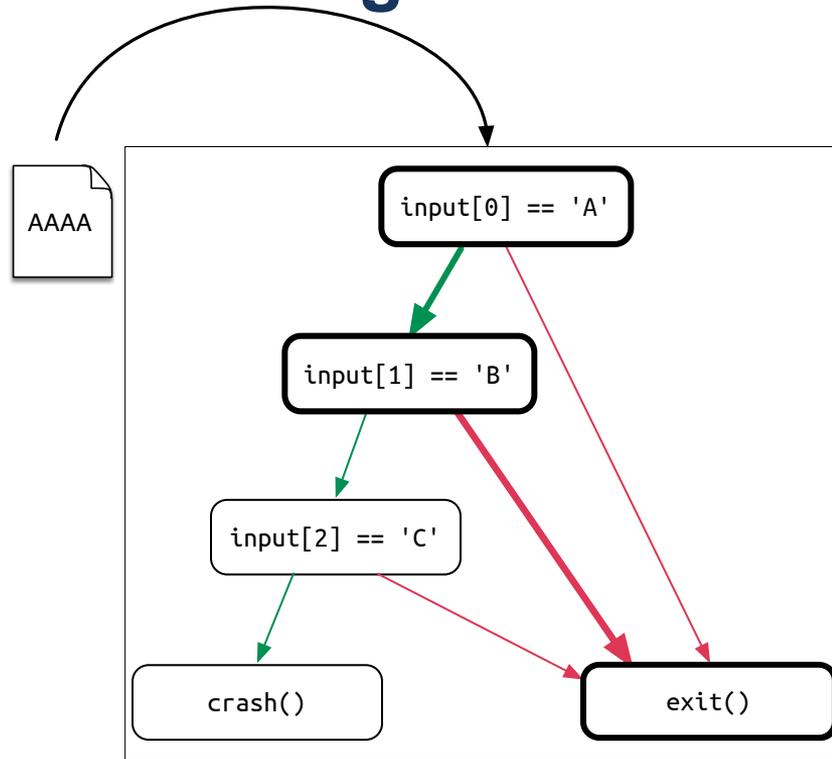
Coverage-Guided Fuzzing



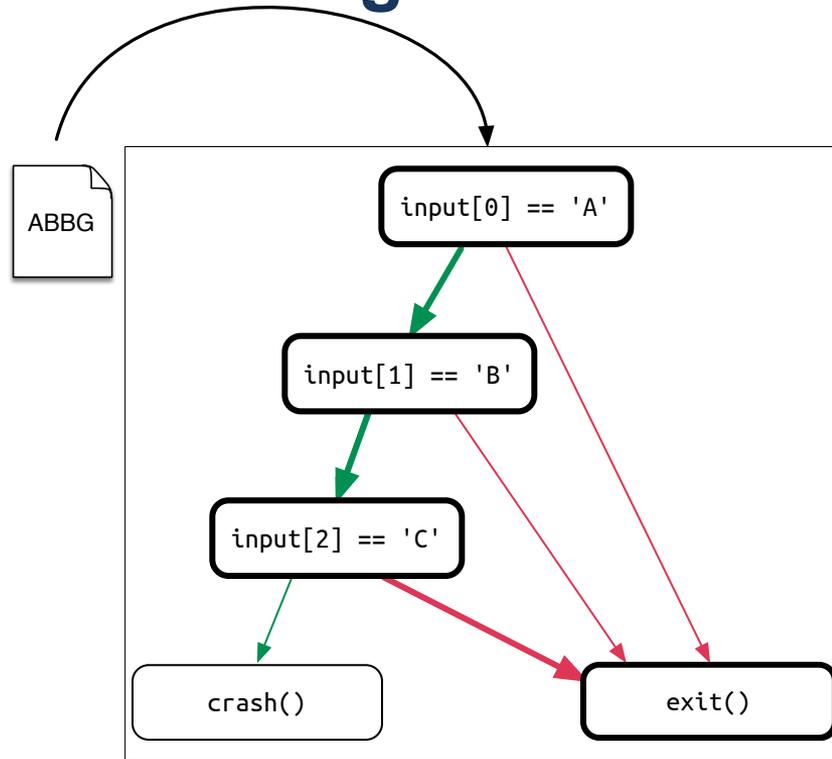
Coverage-Guided Fuzzing



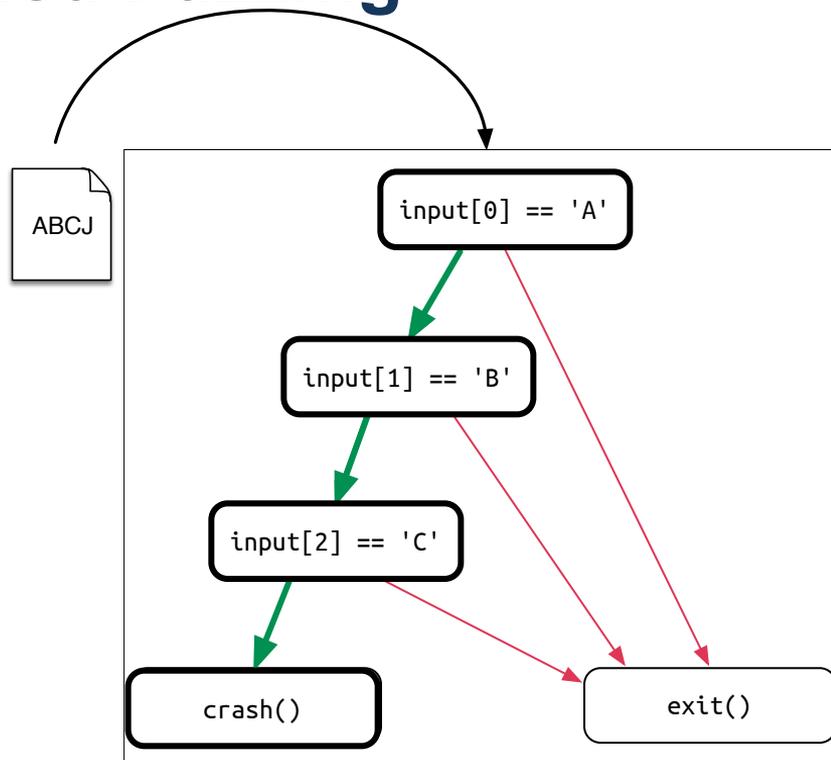
Coverage-Guided Fuzzing



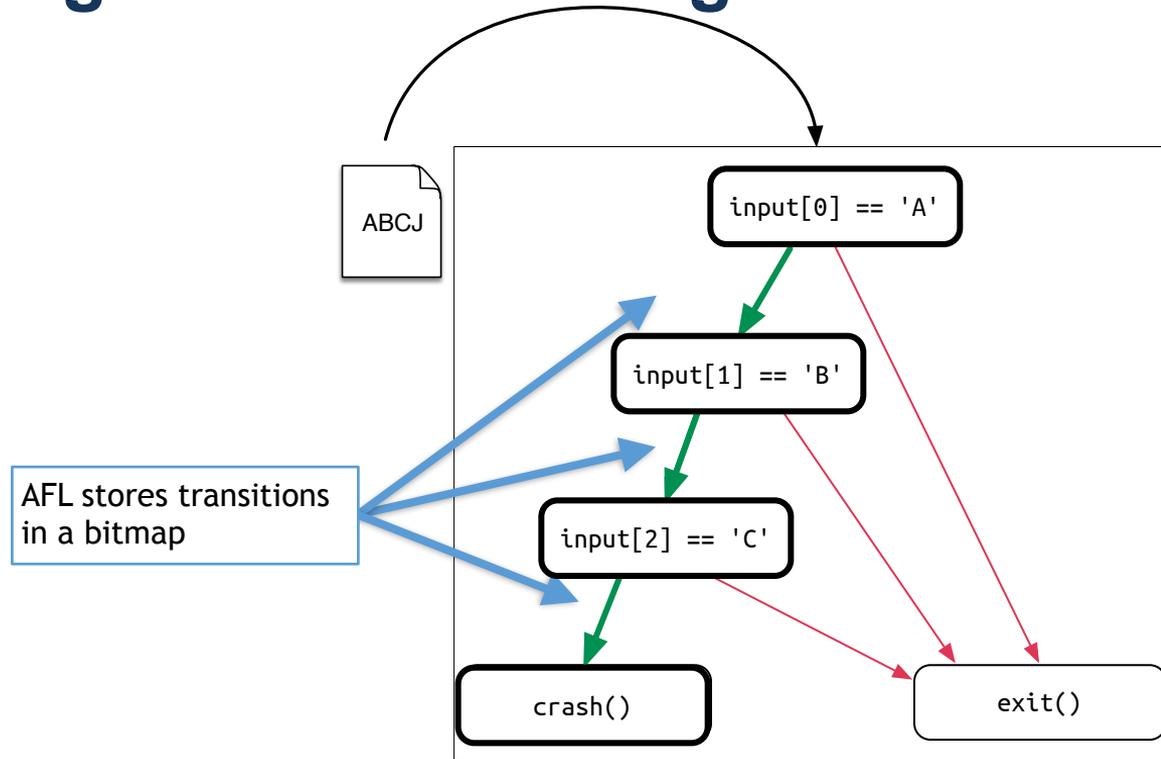
Coverage-Guided Fuzzing



Coverage-Guided Fuzzing



Coverage-Guided Fuzzing



Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++
Static Rewriting	✓	-	X	++

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++
Static Rewriting	✓	-	X	++
Dynamic Binary Instrumentation	✓	-*	✓	-

* Peter Feiner, et al., DRK: DynamoRIO as a Linux Kernel Module

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++
Static Rewriting	✓	-	X	++
Dynamic Binary Instrumentation	✓	-*	✓	-
Emulation	✓	✓	✓	--

* Peter Feiner, et al., DRK: DynamoRIO as a Linux Kernel Module

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++
Static Rewriting	✓	-	X	++
Dynamic Binary Instrumentation	✓	-*	✓	-
Emulation	✓	✓	✓	--
Intel Branch Trace Store	✓	✓	✓	+

* Peter Feiner, et al., DRK: DynamoRIO as a Linux Kernel Module

Feedback Mechanism

	Closed-Source	Kernel	Stable	Fast
Compile-Time Instrumentation	X	✓	✓	++
Static Rewriting	✓	-	X	++
Dynamic Binary Instrumentation	✓	-*	✓	-
Emulation	✓	✓	✓	--
Intel Branch Trace Store	✓	✓	✓	+
Intel Processor Trace	✓	✓	✓	+++

* Peter Feiner, et al., DRK: DynamoRIO as a Linux Kernel Module

Intel Processor Trace

Intel Processor Trace

Instruction

Intel PT Packet

Intel Processor Trace

Instruction

Intel PT Packet

jmp/call loc_b

Inferable from disassembly

Intel Processor Trace

Instruction

Intel PT Packet

`jmp/call loc_b`

Inferable from disassembly

`jnz loc_a`

Taken / Not Taken

Intel Processor Trace

Instruction

Intel PT Packet

`jmp/call loc_b`

Inferable from disassembly

`jnz loc_a`

Taken / Not Taken

`jmp/call [eax]`

Target IP

Intel Processor Trace

Instruction

Intel PT Packet

`jmp/call loc_b`

Inferable from disassembly

`jnz loc_a`

Taken / Not Taken

`jmp/call [eax]`

Target IP

`ret`

Target IP (if required)

Intel Processor Trace

Intel PT Data

Not Taken

Target IP (0x1009)

Target IP (0x1055)

Intel Processor Trace

Intel PT Data

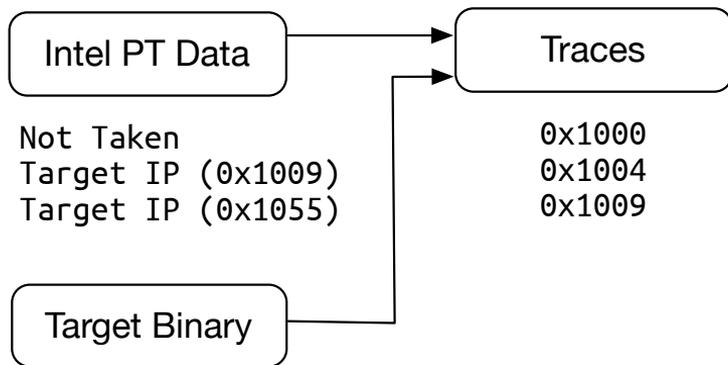
Not Taken

Target IP (0x1009)

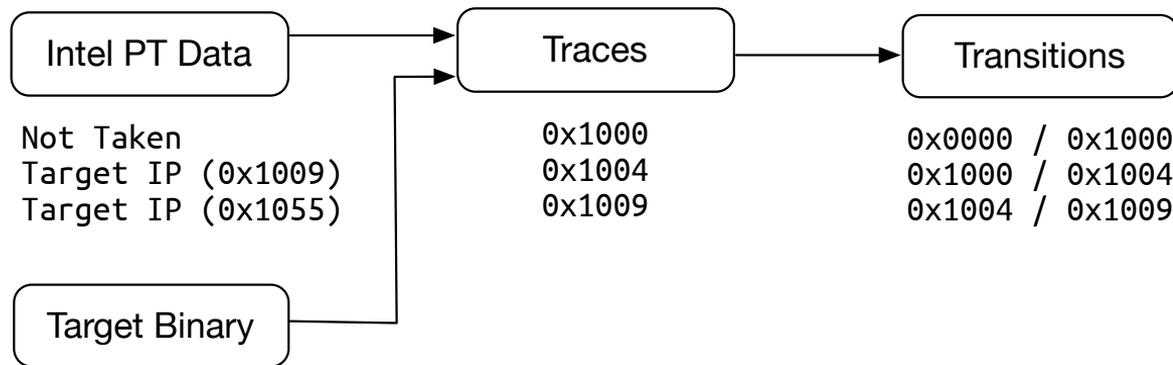
Target IP (0x1055)

Target Binary

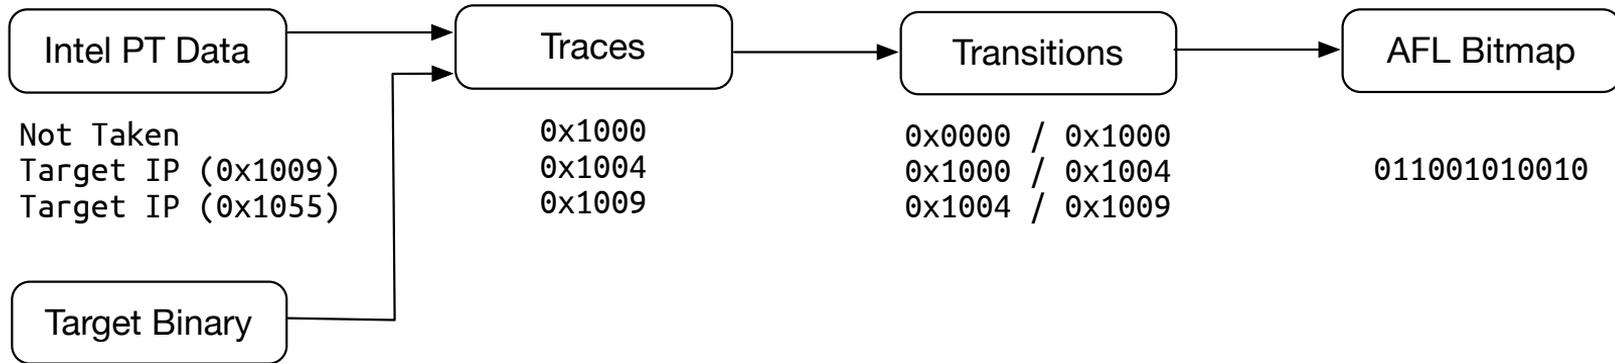
Intel Processor Trace



Intel Processor Trace

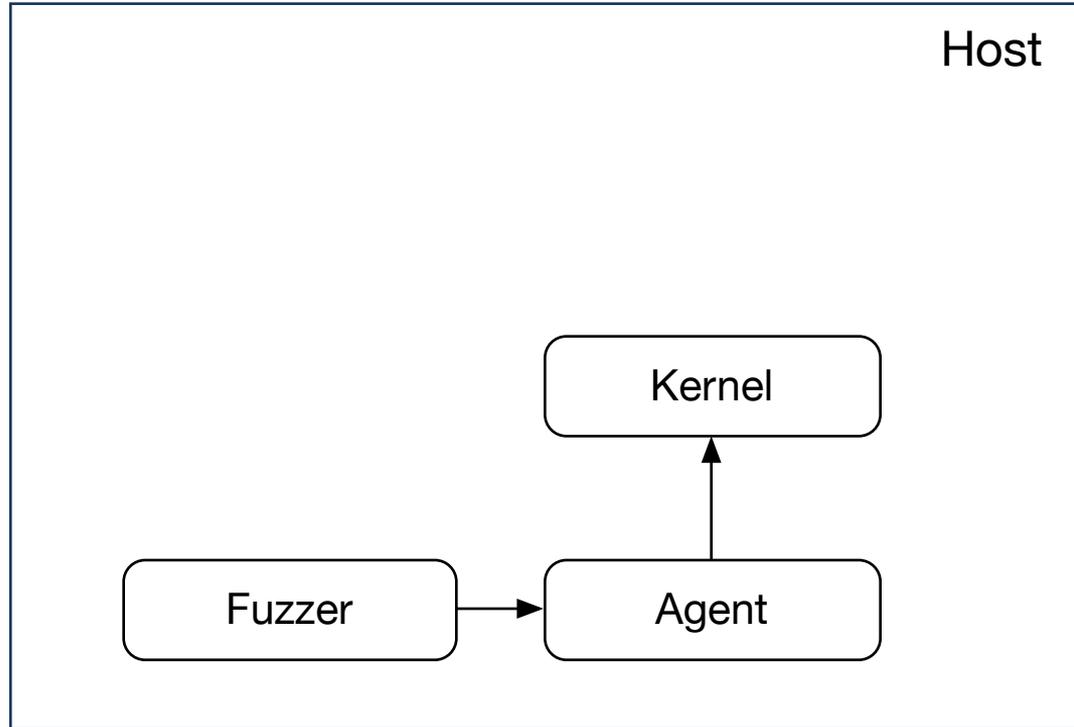


Intel Processor Trace



Architecture

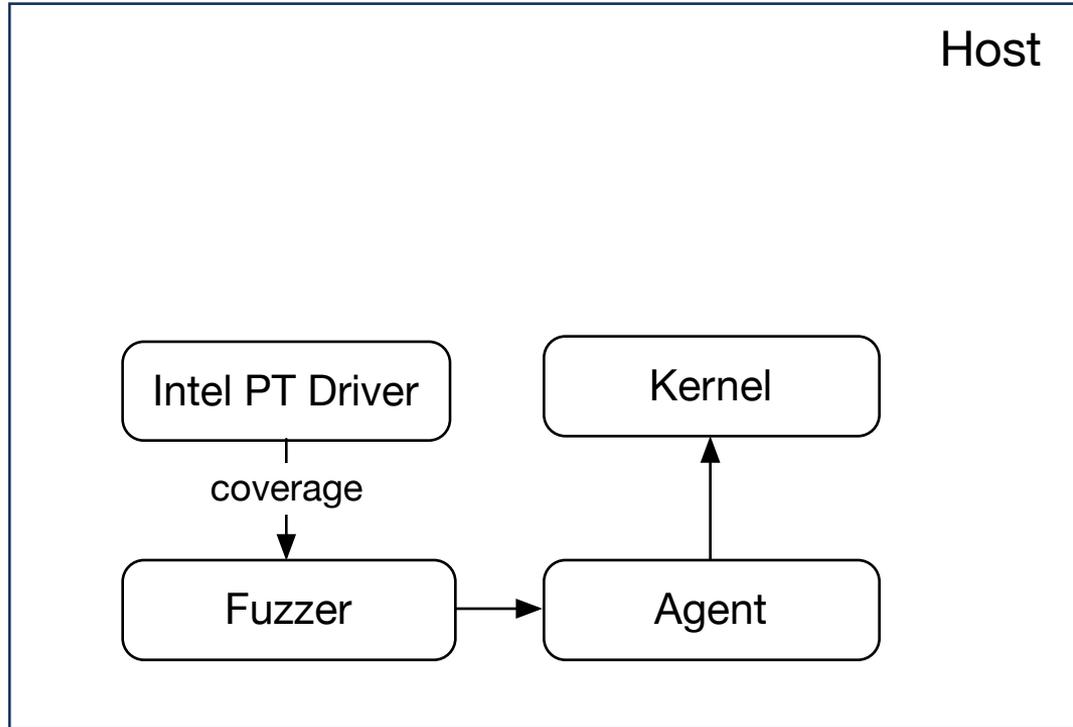
Architecture



Architecture

Benefits:

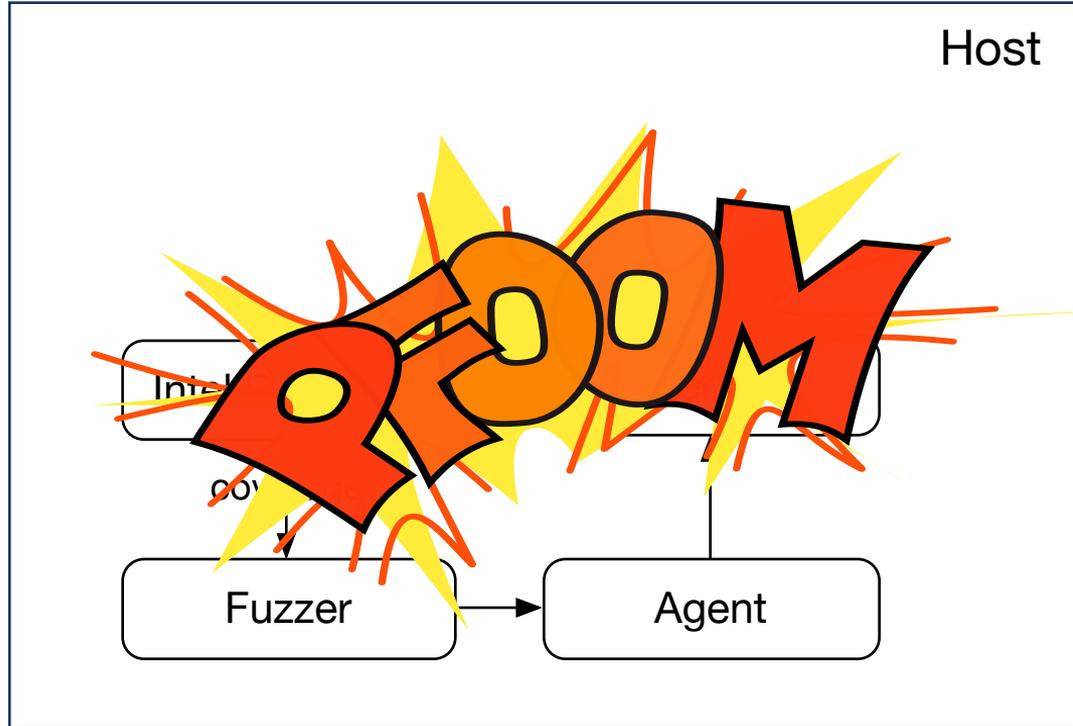
✓ Coverage



Architecture

Benefits:

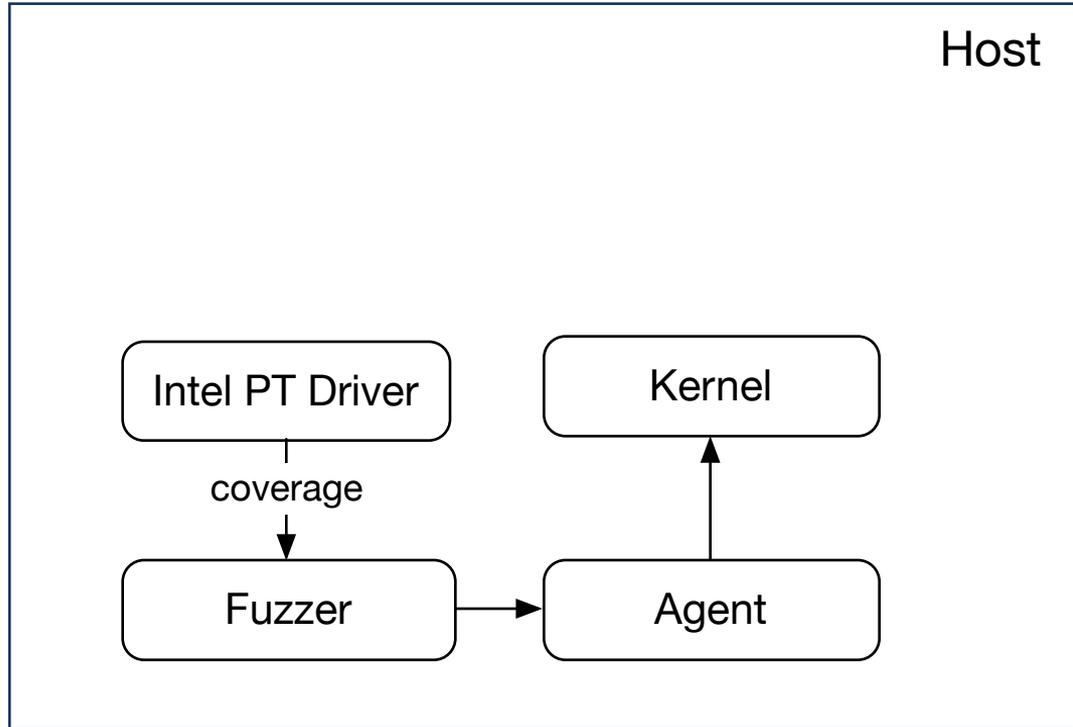
✓ Coverage



Architecture

Benefits:

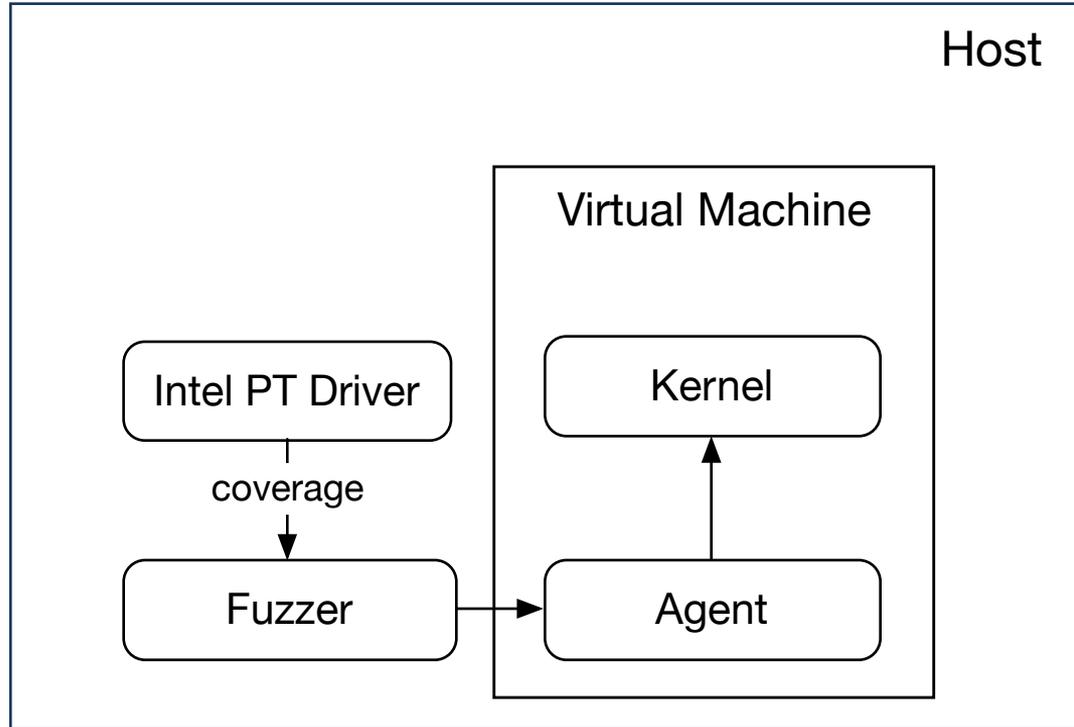
✓ Coverage



Architecture

Benefits:

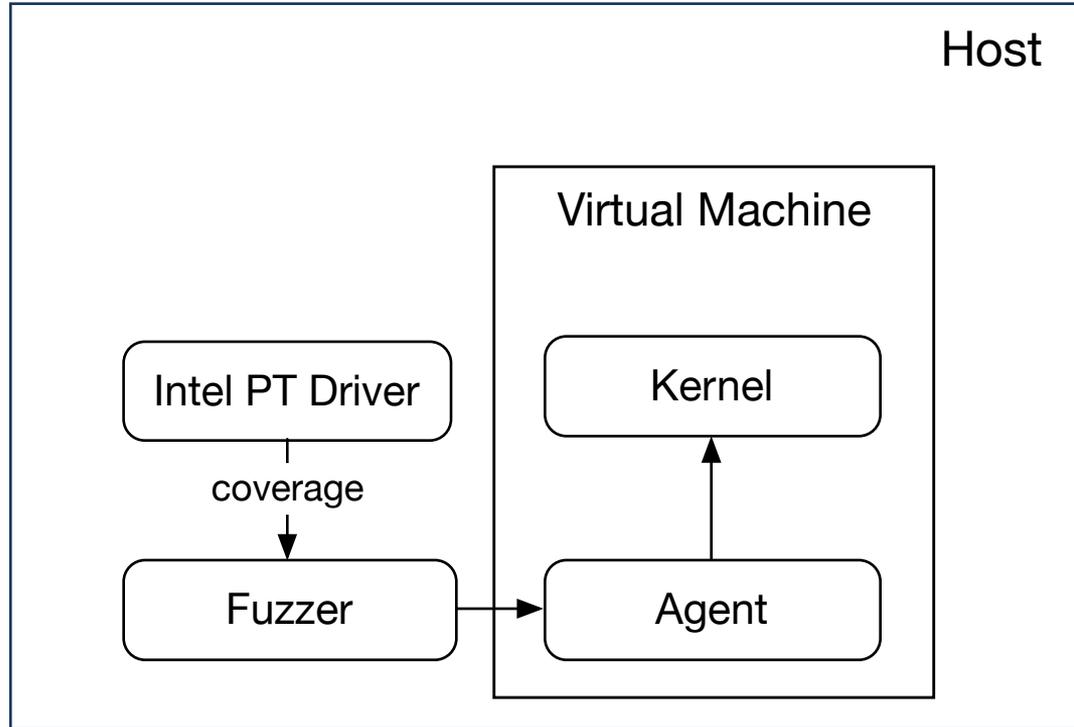
- ✓ Coverage
- ✓ Crash Tolerance



Architecture

Benefits:

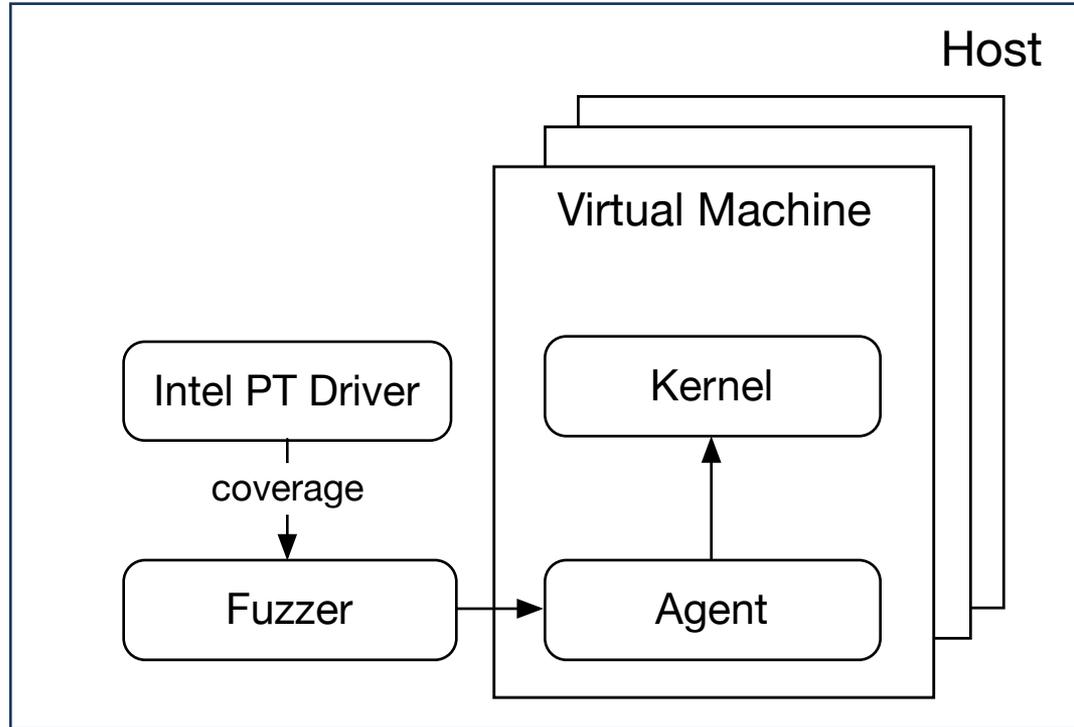
- ✓ Coverage
- ✓ Crash Tolerance
- ✓ OS Independence



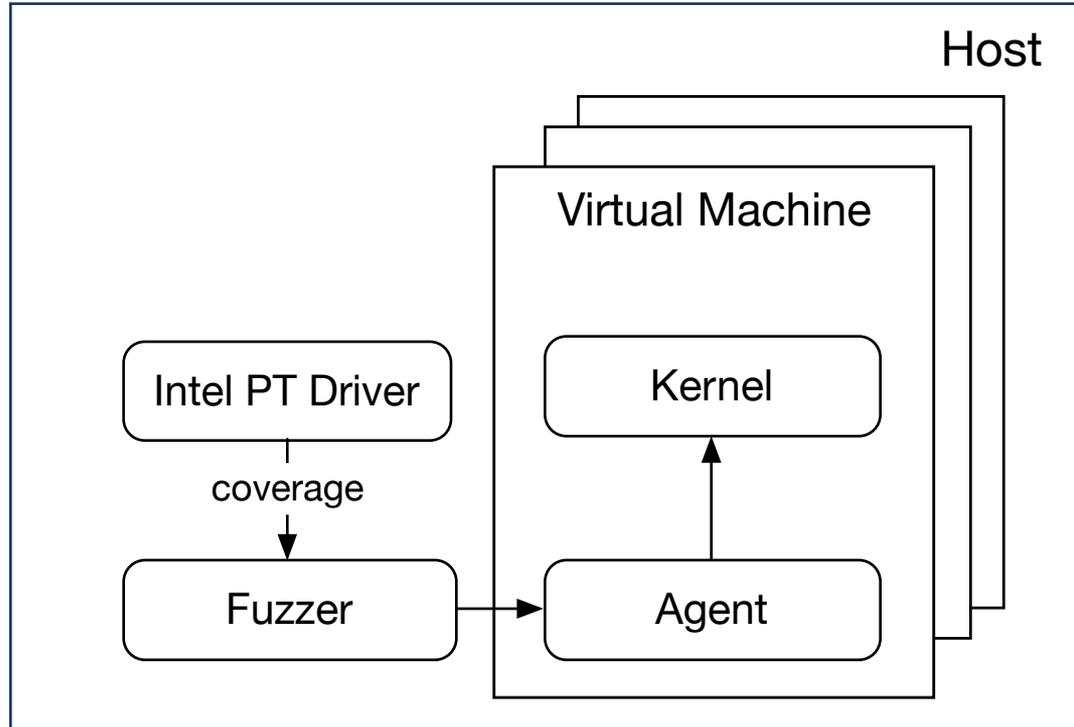
Architecture

Benefits:

- ✓ Coverage
- ✓ Crash Tolerance
- ✓ OS Independence
- ✓ Scalable



Trace Filtering

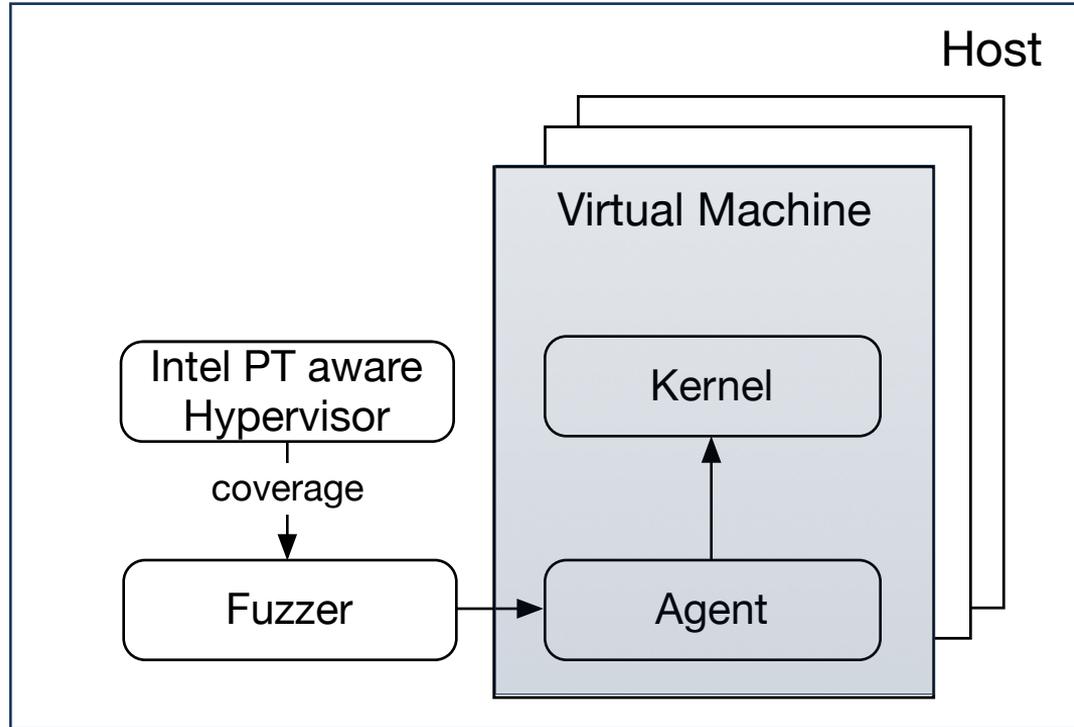


Trace Filtering

vCPU Tracing

Filter-Mechanisms:

✓ vCPU

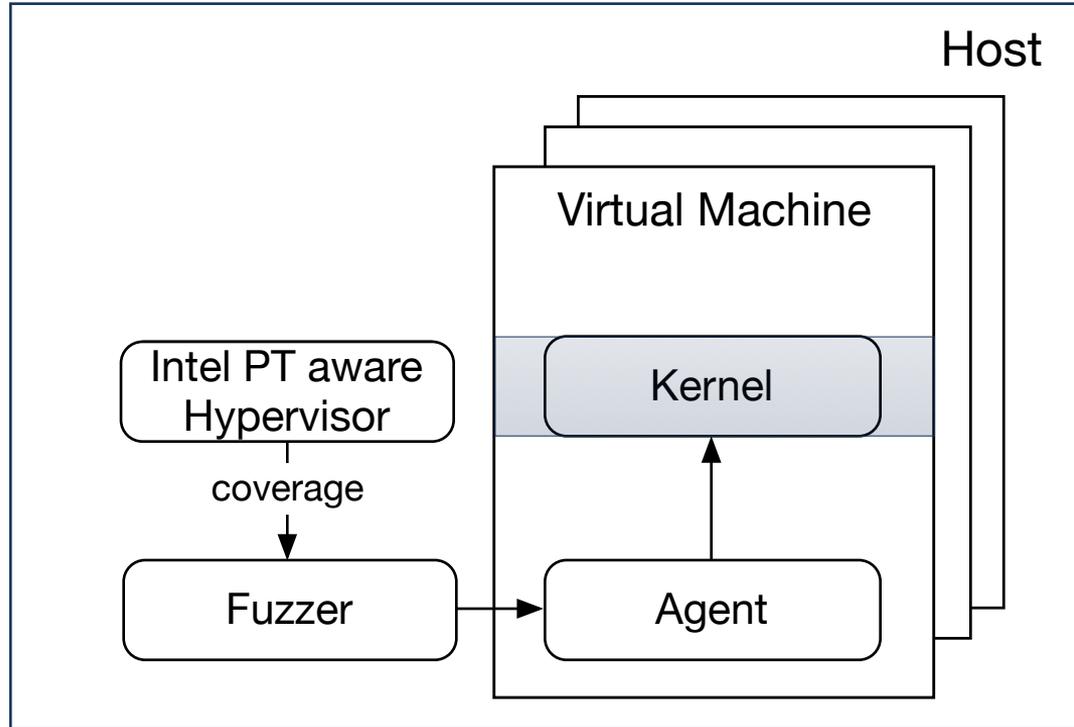


Trace Filtering

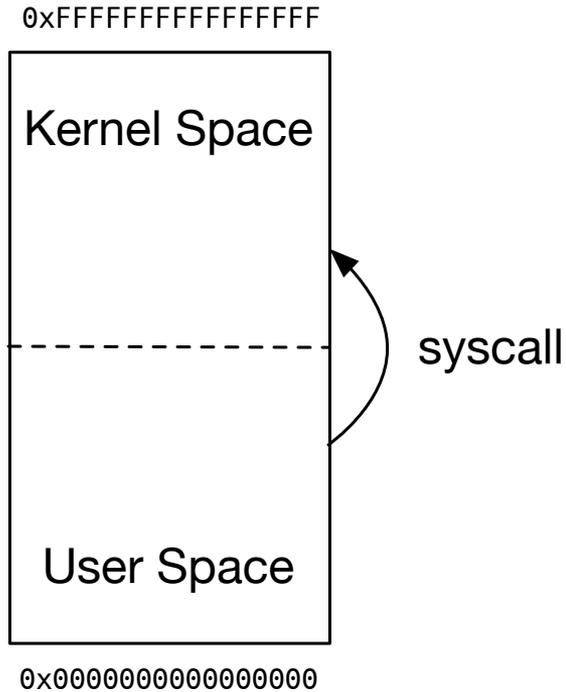
Guest Ring-0 Tracing

Filter-Mechanisms:

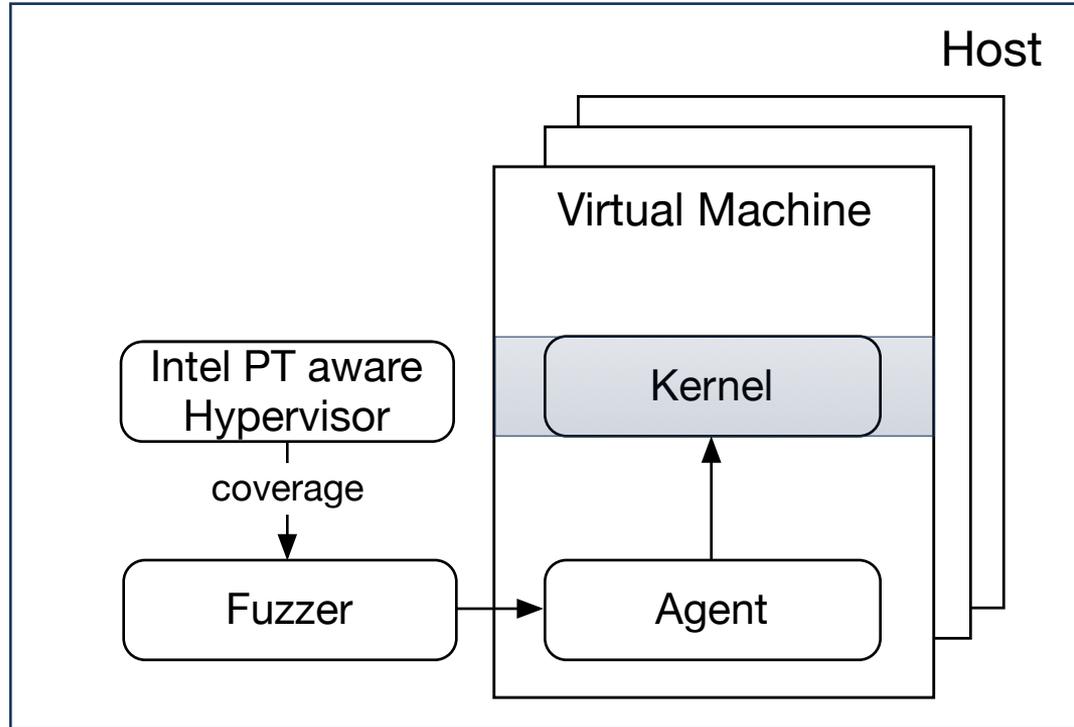
- ✓ vCPU
- ✓ Supervisor



Trace Filtering



Guest Ring-0 Tracing

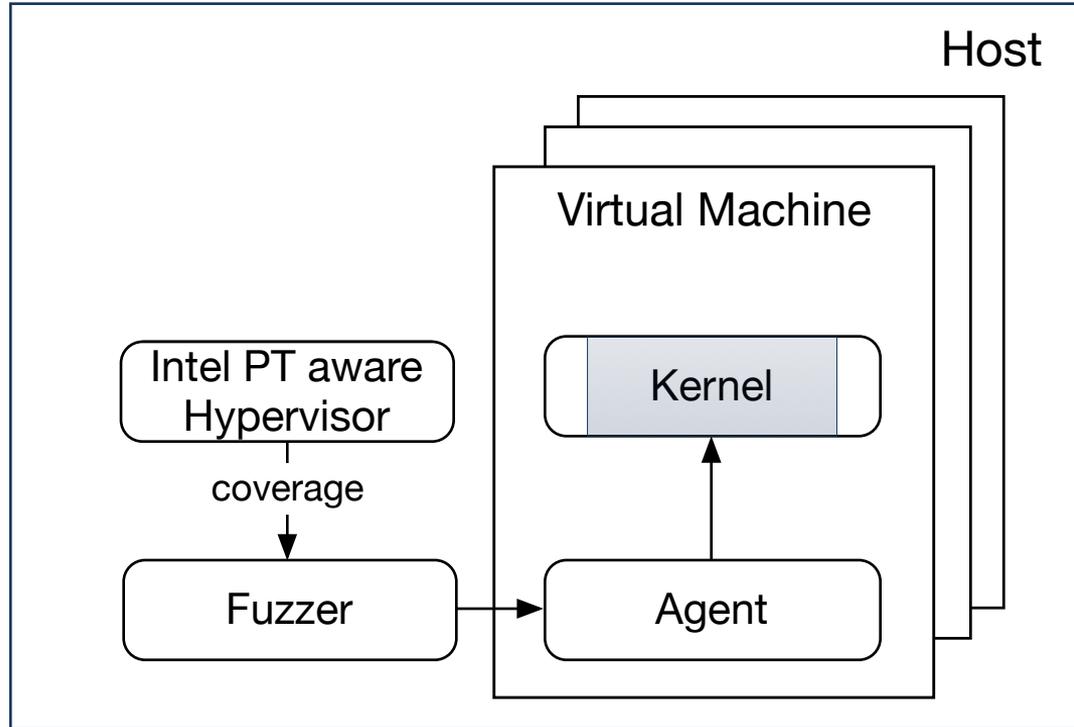


Trace Filtering

Guest Ring-0 Tracing (Fuzzing-Process)

Filter-Mechanisms:

- ✓ vCPU
- ✓ Supervisor
- ✓ CR3

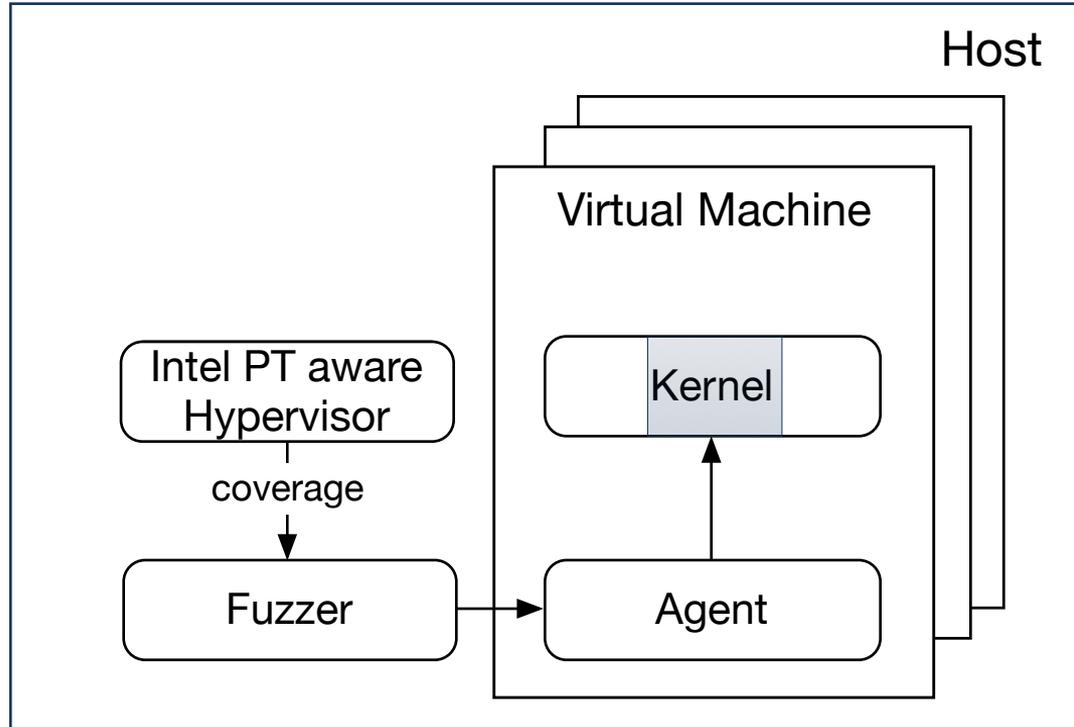


Trace Filtering

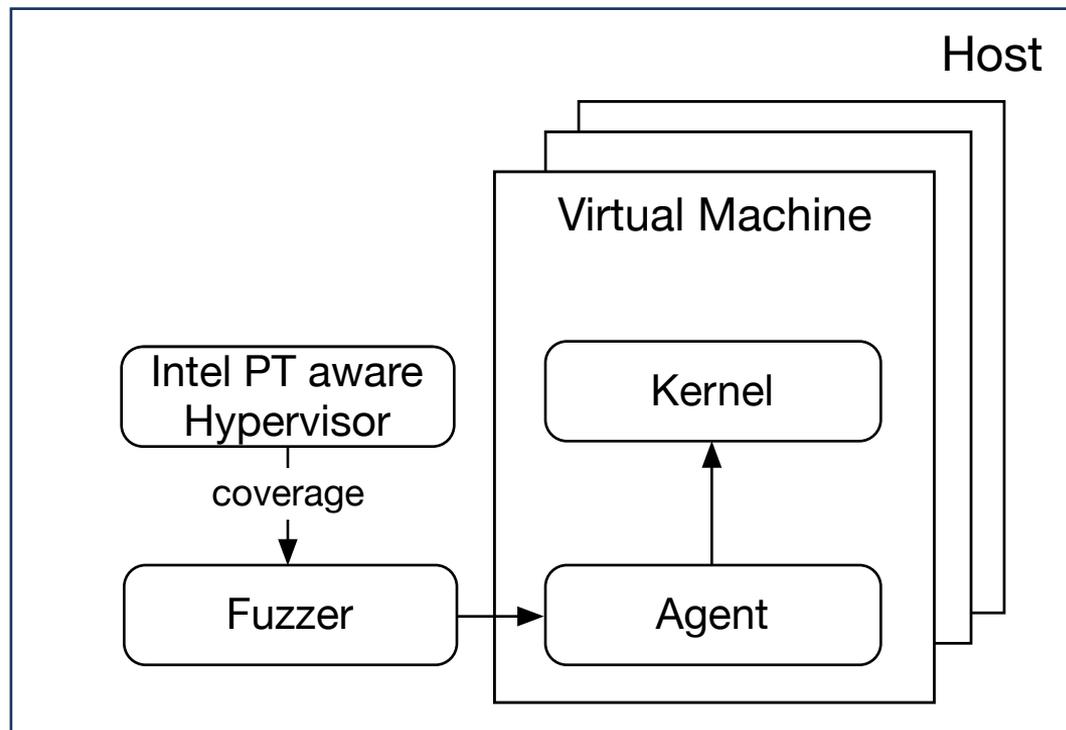
Filter-Mechanisms:

- ✓ vCPU
- ✓ Supervisor
- ✓ CR3
- ✓ IP-Range

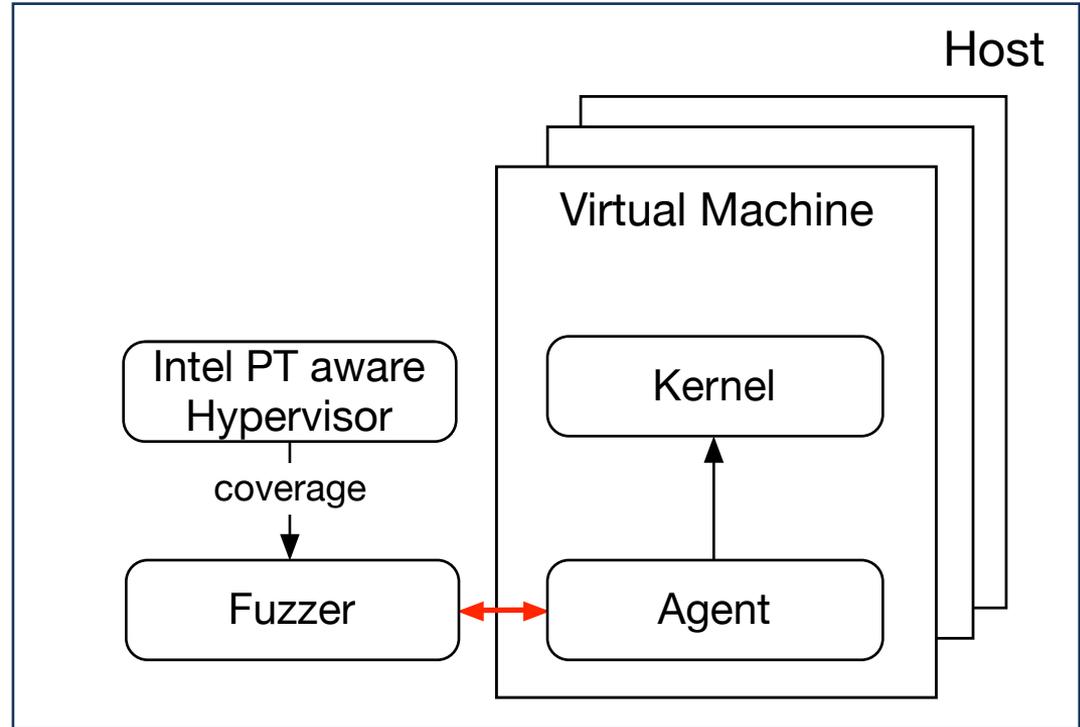
Guest Ring-0 Tracing (Fuzzing-Process & Target Range)



Inter-VM Communication

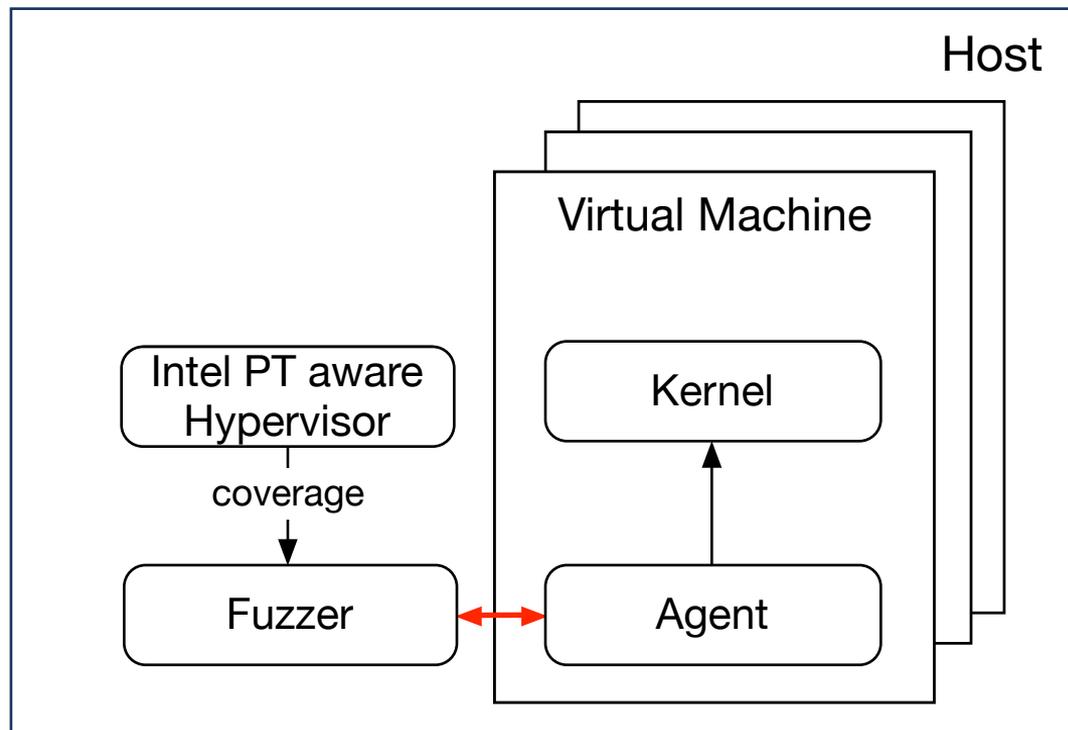


Inter-VM Communication



Inter-VM Communication

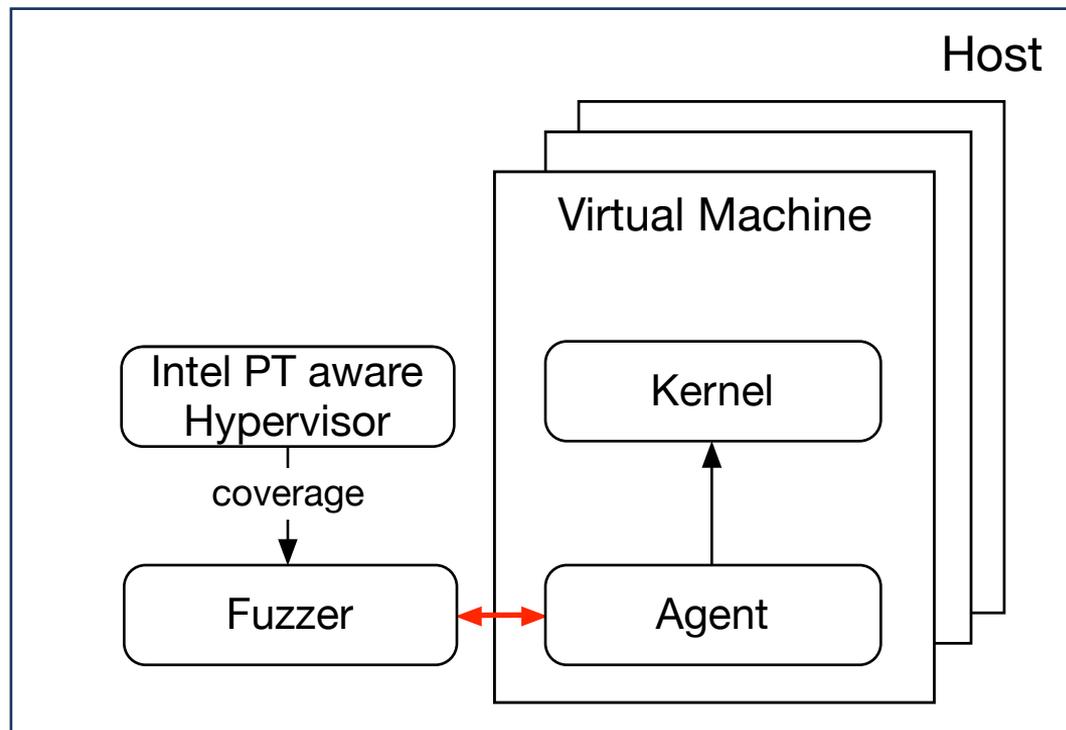
Guest to Host:



Inter-VM Communication

Guest to Host:

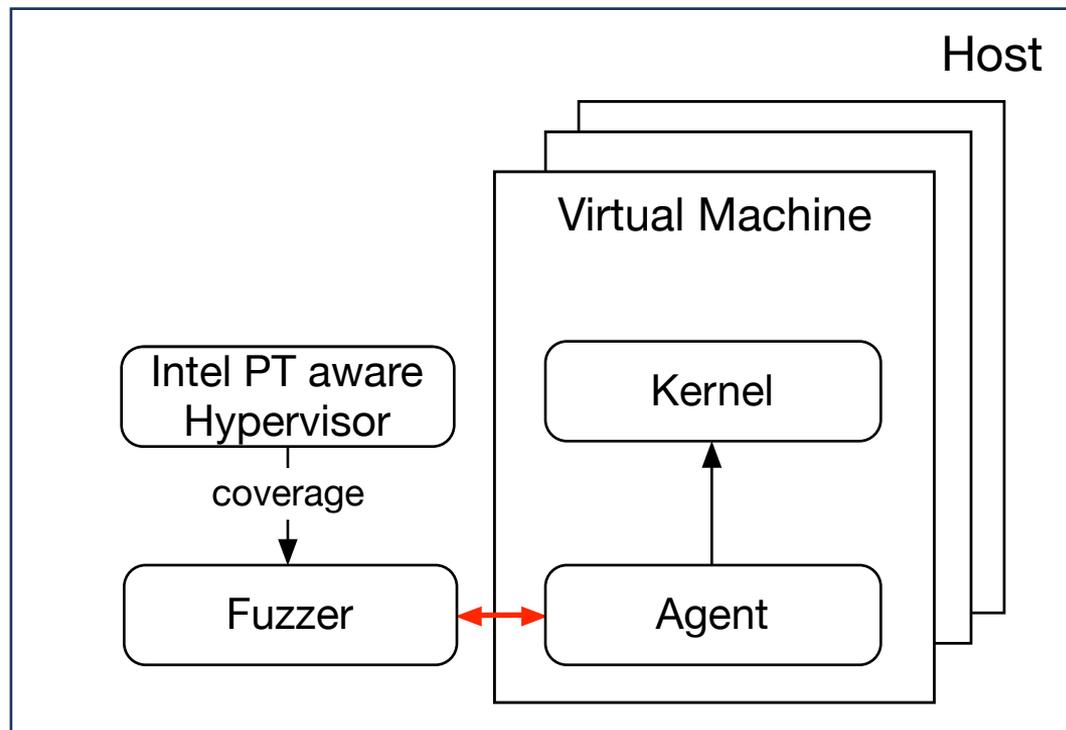
- Synchronization



Inter-VM Communication

Guest to Host:

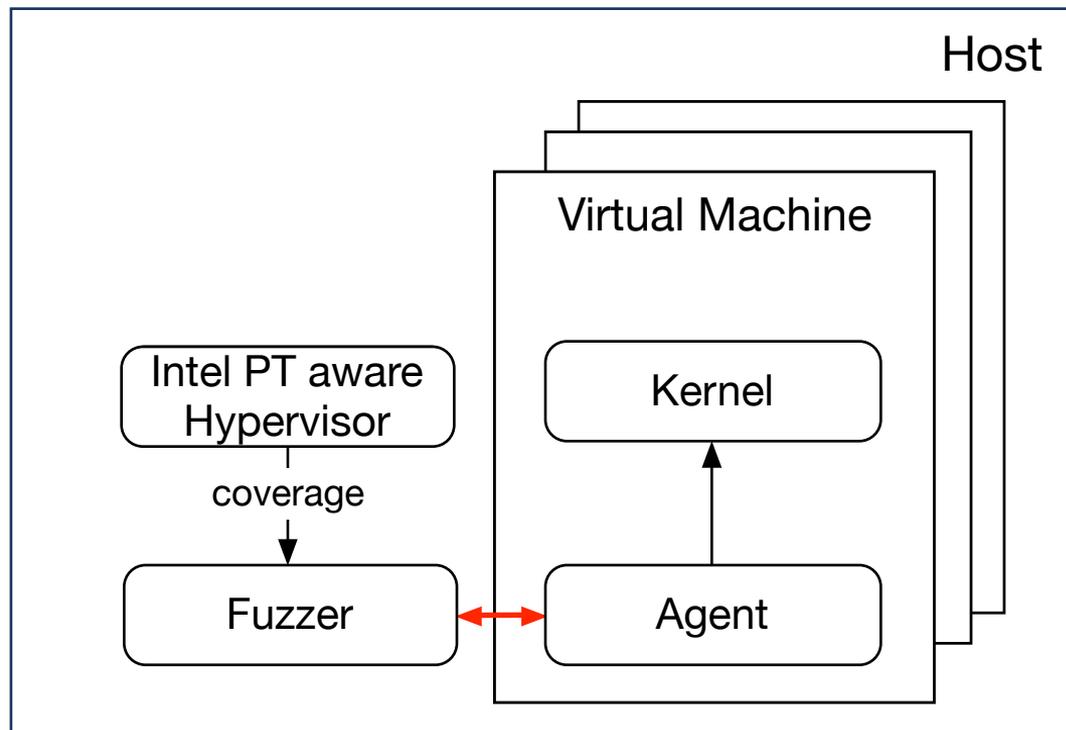
- Synchronization
- Next payload



Inter-VM Communication

Guest to Host:

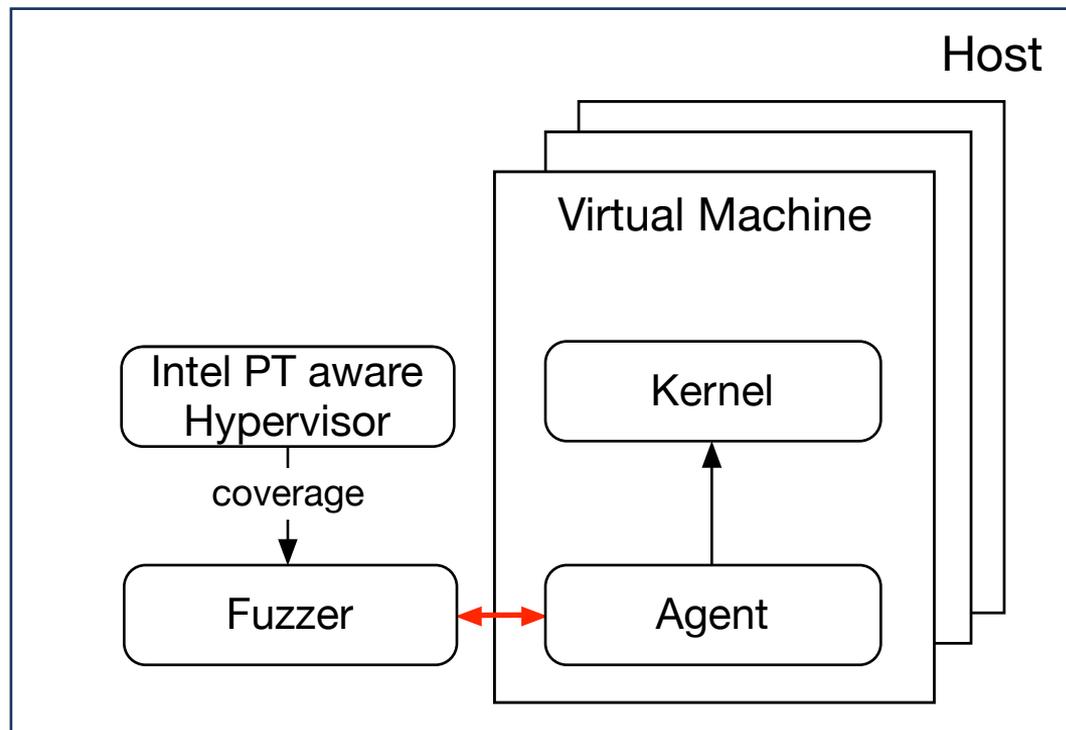
- Synchronization
- Next payload
- Disclose CR3 value



Inter-VM Communication

Guest to Host:

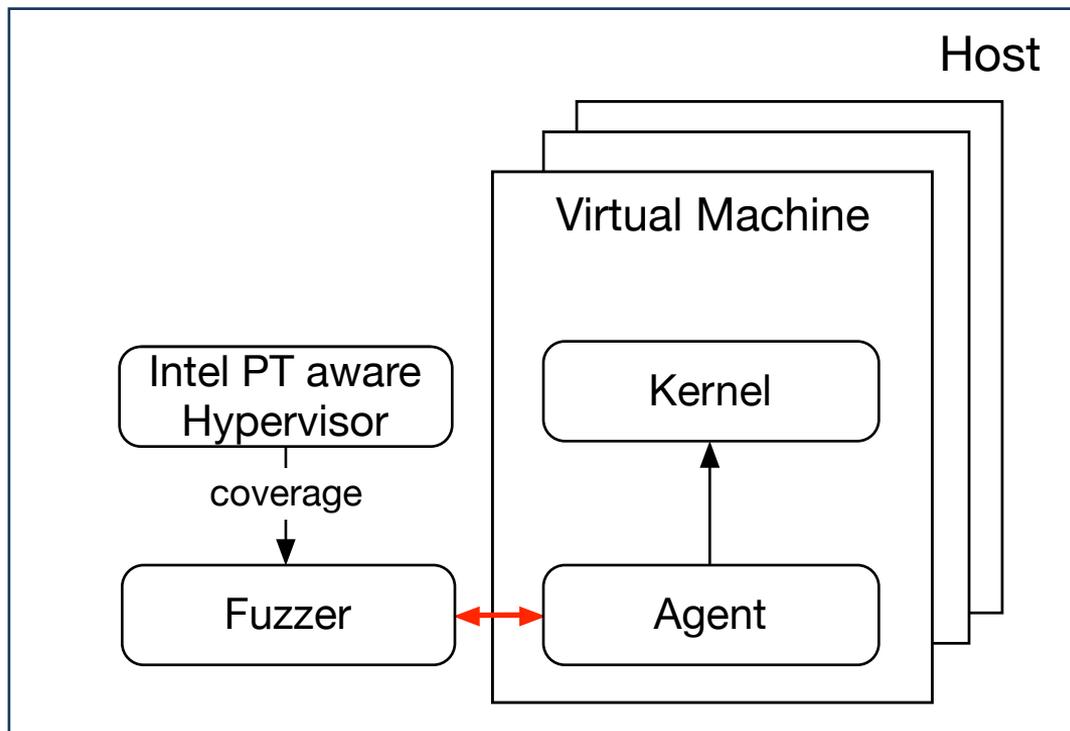
- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address



Inter-VM Communication

Guest to Host:

- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address
- Signal kernel panic

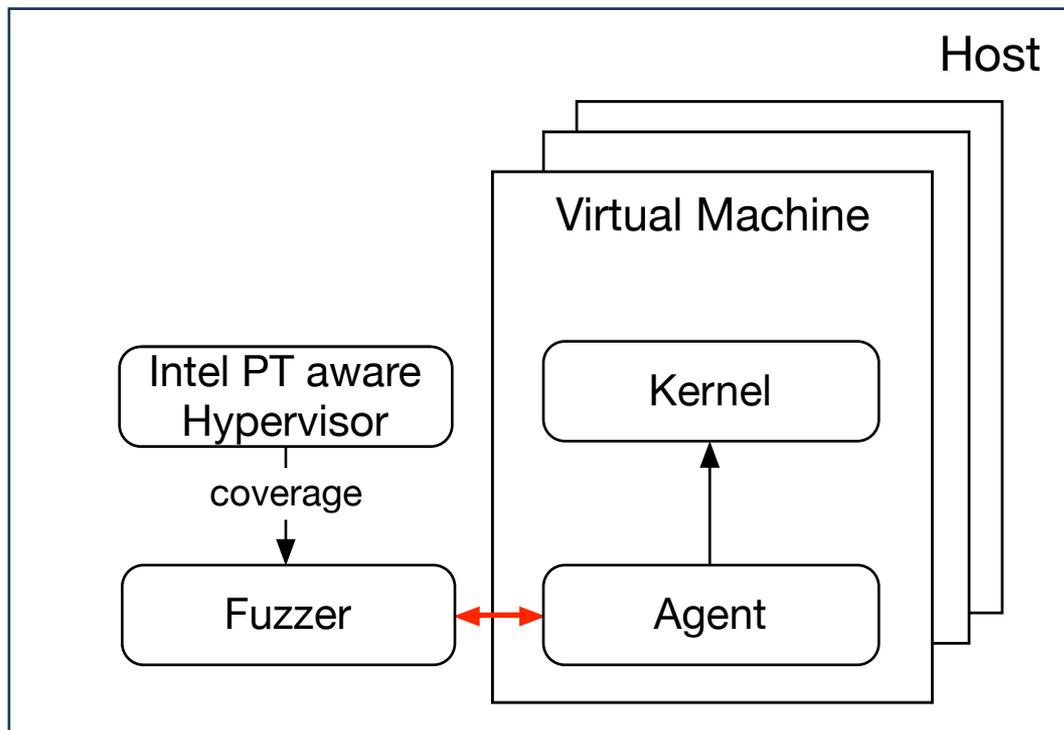


Inter-VM Communication

Guest to Host:

- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address
- Signal kernel panic

Host to Guest:



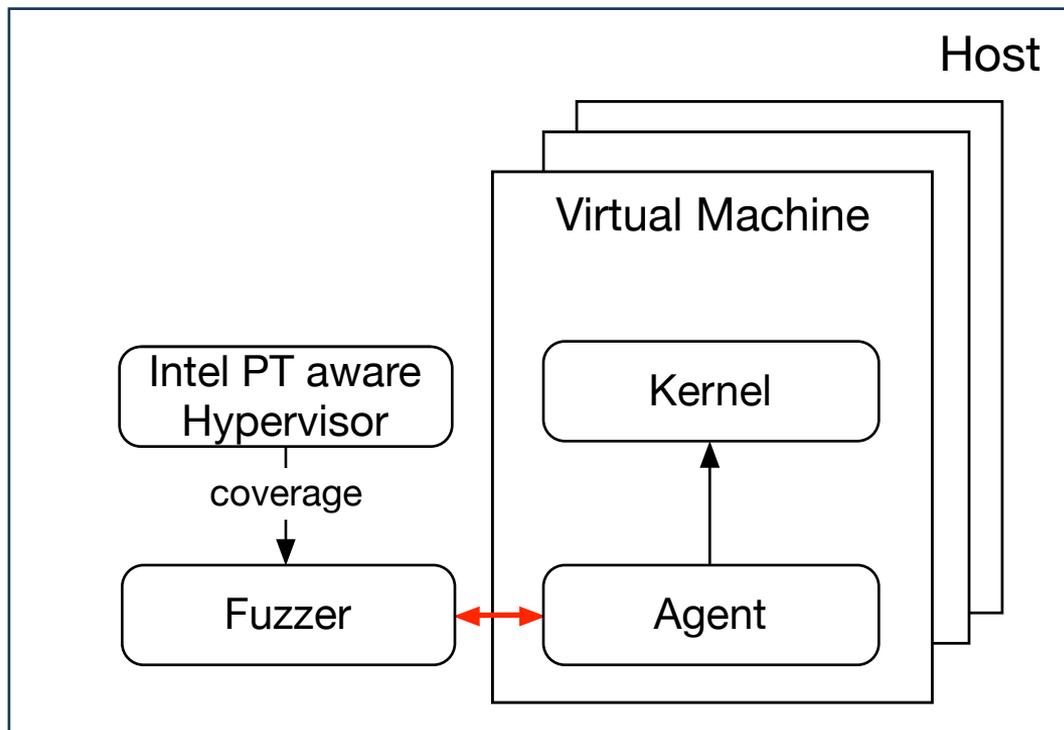
Inter-VM Communication

Guest to Host:

- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address
- Signal kernel panic

Host to Guest:

- Agent



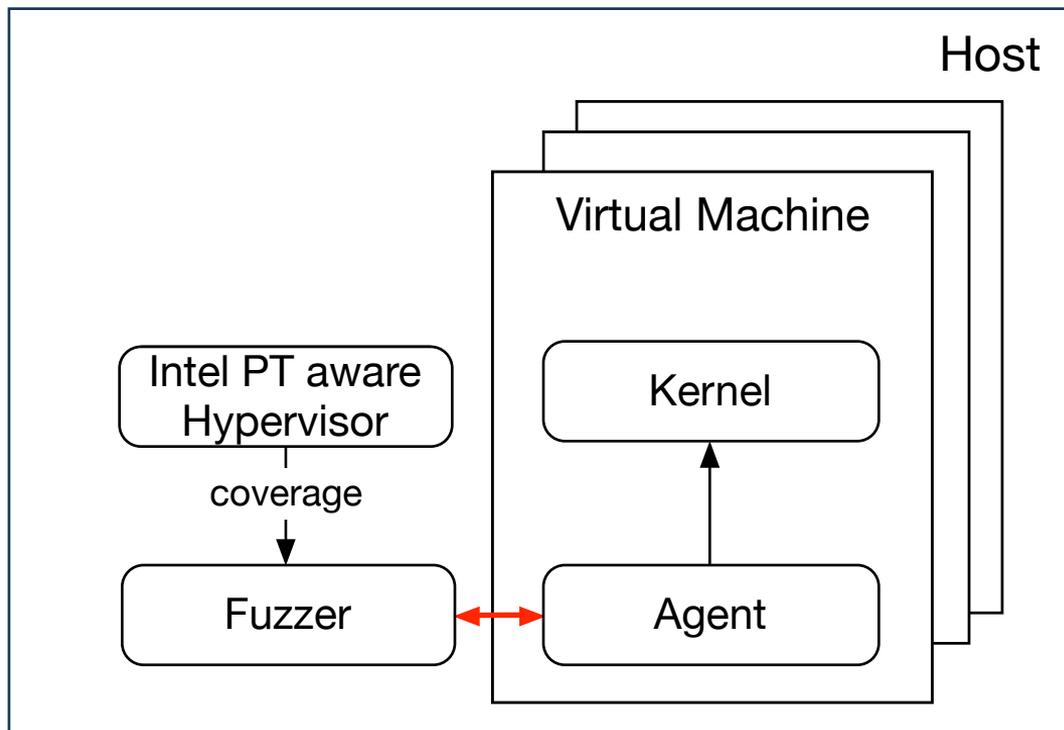
Inter-VM Communication

Guest to Host:

- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address
- Signal kernel panic

Host to Guest:

- Agent
- Payloads



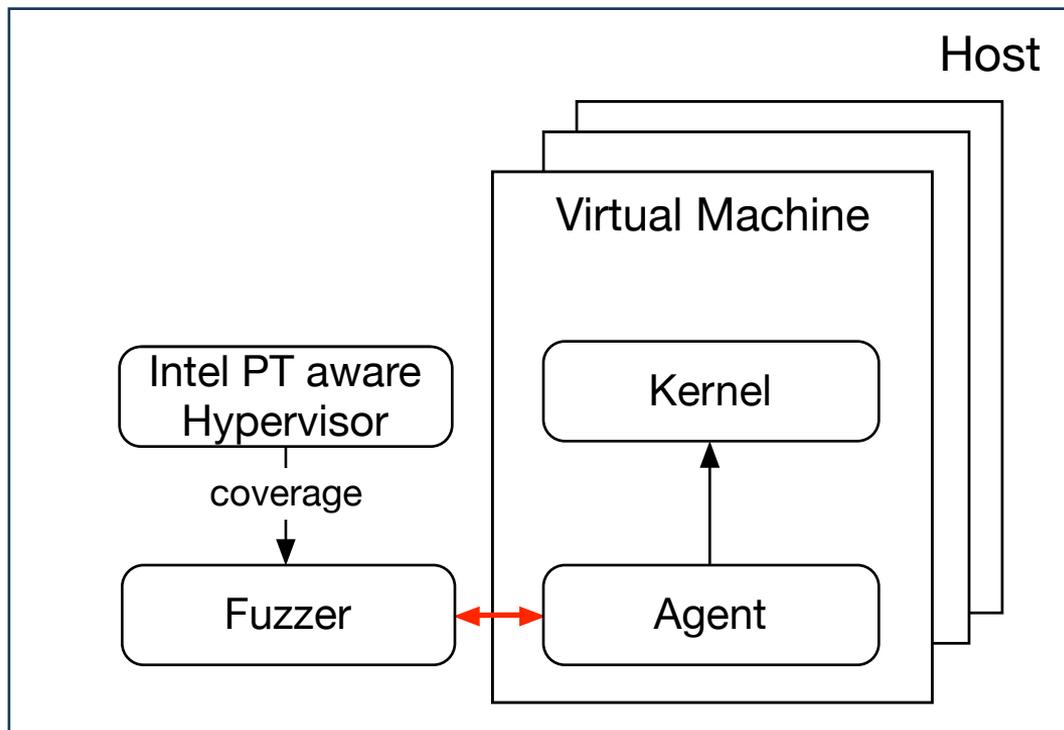
Inter-VM Communication

Guest to Host:

- Synchronization
- Next payload
- Disclose CR3 value
- Panic handler address
- Signal kernel panic

Host to Guest:

- Agent
- Payloads
- Overwrite panic handler



Implementation

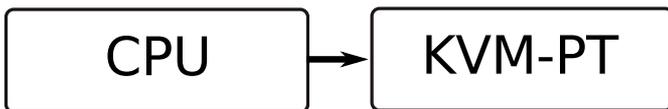
Implementation



CPU:

- generates raw Intel-PT data
- writes data to main memory

Implementation



KVM-PT:

- configures Intel PT via MSR
- enables tracing during VM-Entry transition
- disables tracing during VM-Exit transition

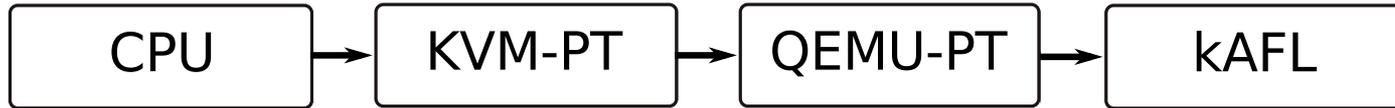
Implementation



QEMU-PT:

- usermode counterpart
- decodes Intel PT data on-the-fly

Implementation



kAFL Fuzzer:

- generates new fuzz payloads
- detects new behavior

Evaluation

New Vulnerabilities

Windows:

- NTFS (DoS)

macOS:

- HFS (DoS, Memory Corruption)
- APFS (Memory Corruption)

Linux:

- EXT4 (DoS, Memory Corruption)
- Keyctl (Nullpointer Dereference)

Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

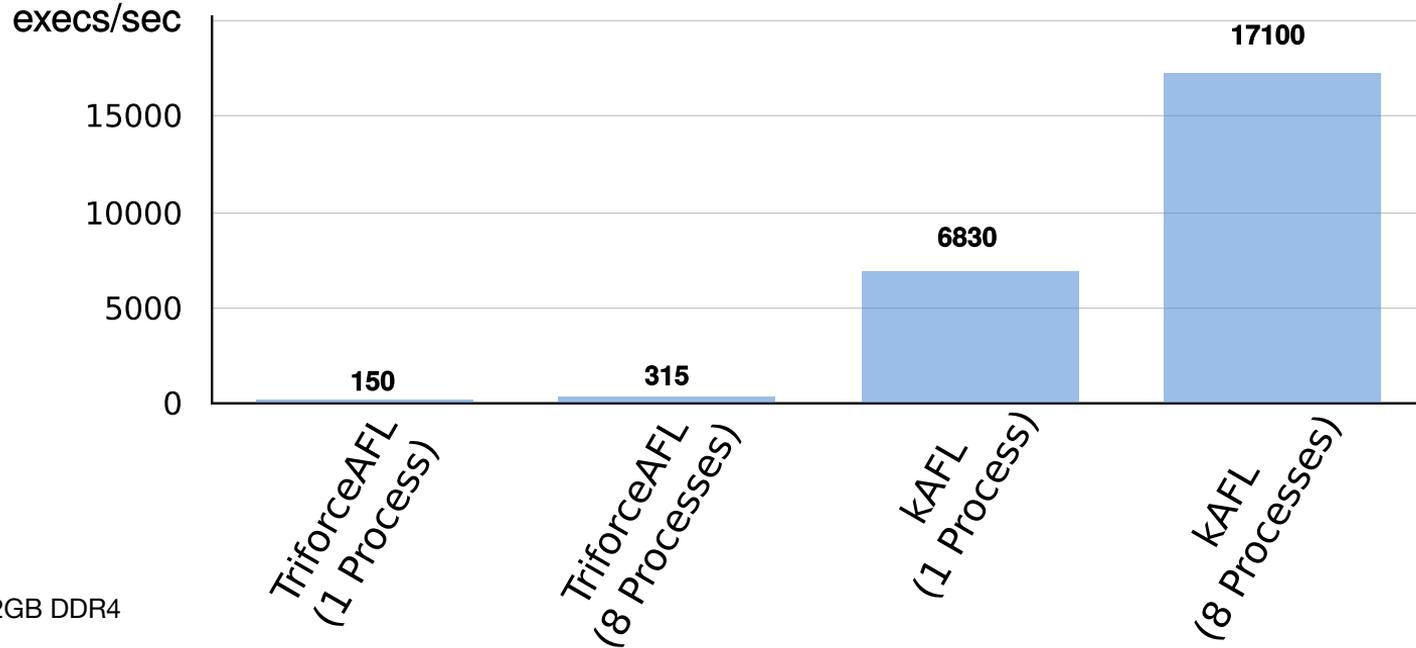
Evaluation Driver

```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Evaluation Driver

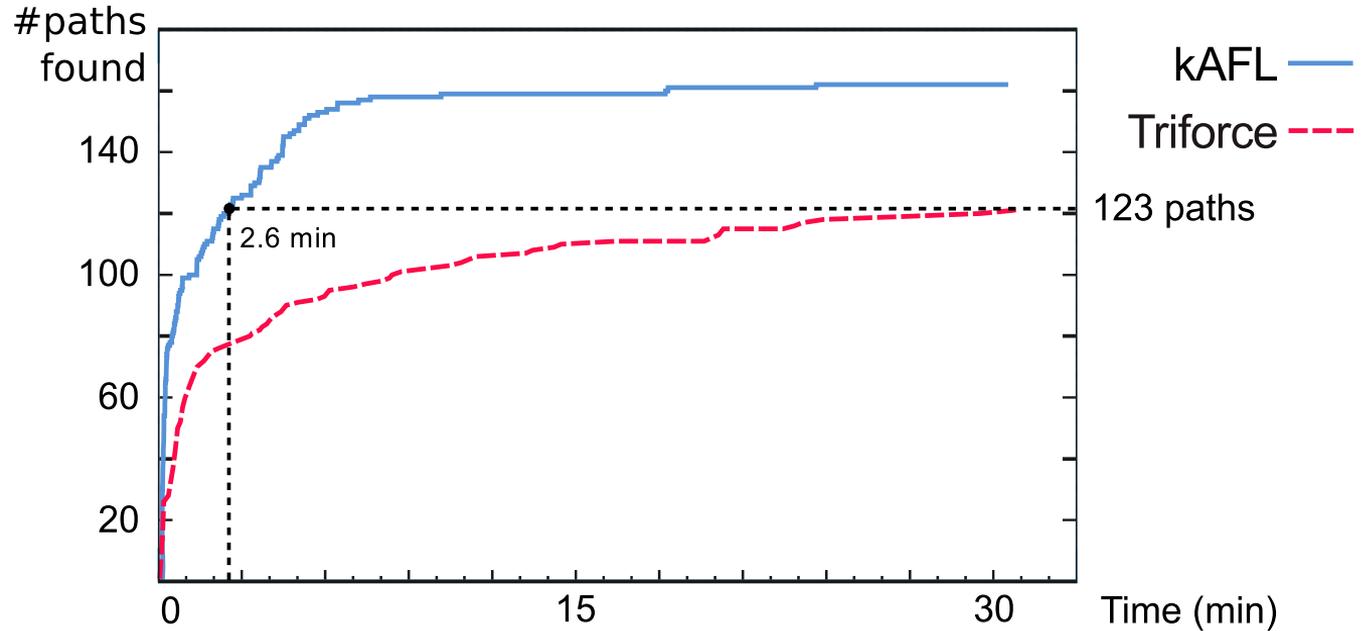
```
1  jsmn_parser parser;
2  jsmntok_t tokens[5];
3  jsmn_init(&parser);
4
5  int res=jsmn_parse(&parser, input, size, tokens, 5);
6  if(res >= 2){
7      if(tokens[0].type == JSMN_STRING){
8          int json_len = tokens[0].end - tokens[0].
              start;
9          int s = tokens[0].start;
10         if(json_len > 0 && input[s+0] == 'K'){
11             if(json_len > 1 && input[s+1] == 'A'){
12                 if(json_len > 2 && input[s+2] == 'F'){
13                     if(json_len > 3 && input[s+3] == 'L'){
14                         panic(KERN_INFO "KAFL...\n");
15                     }
16                 }
17             }
18         }
19     }
```

Performance



Intel i7-6700 / 32GB DDR4

Coverage



Intel i7-6700 / 32GB DDR4

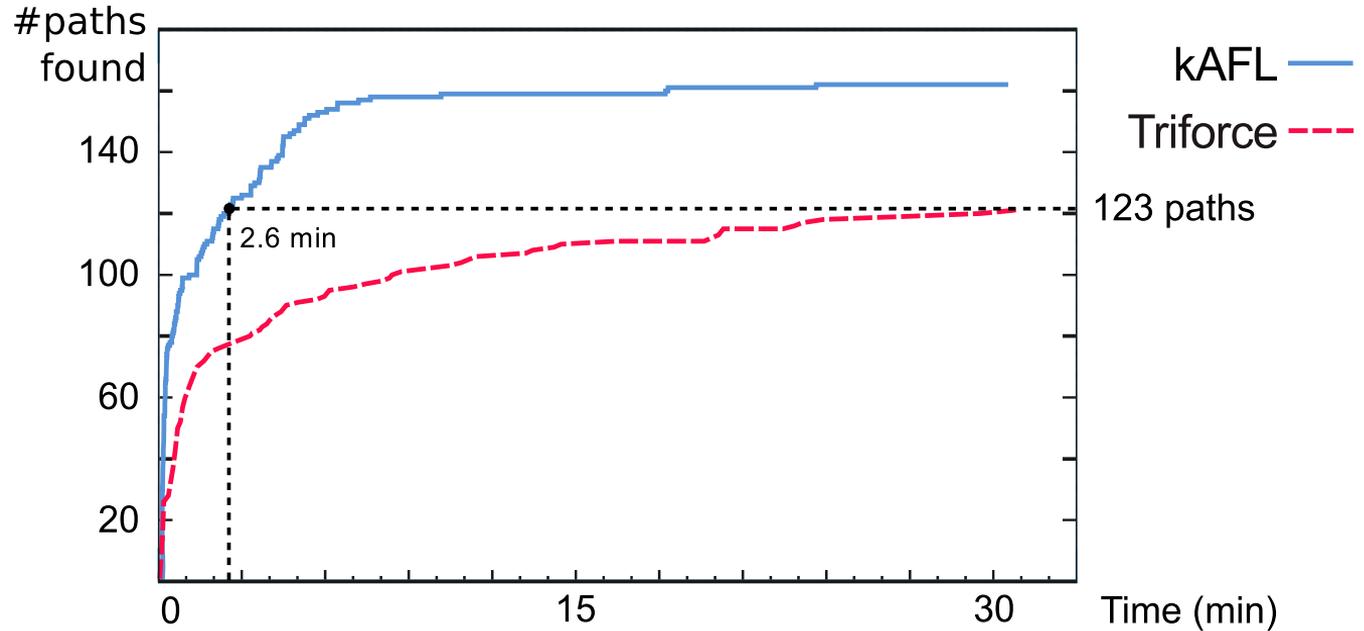
Coverage

kAFL:

■ 5-7 minutes

TriforceAFL:

■ ~2 hours



Intel i7-6700 / 32GB DDR4

Conclusion

- Intel PT and virtualization for feedback fuzzing
 - Fast
 - OS independence (x86-64)
 - Reliable and long-term
 - Fully extensible
- High bug yield for kernel fuzzing
 - Opportunities for further fuzzing

<https://github.com/RUB-SysSec/kAFL>