
yara Documentation

Release 4.3.2

Victor M. Alvarez

Jun 12, 2023

1	Getting started	3
1.1	Compiling and installing YARA	3
1.1.1	Installing with vcpkg	4
1.1.2	Installing on Windows	4
1.1.3	Installing on Mac OS X with Homebrew	5
1.1.4	Installing yara-python	5
1.2	Running YARA for the first time	5
2	Writing YARA rules	7
2.1	Comments	8
2.2	Strings	8
2.2.1	Hexadecimal strings	8
2.2.2	Text strings	10
2.2.3	Regular expressions	15
2.2.4	Private strings	16
2.2.5	String Modifier Summary	17
2.3	Conditions	17
2.3.1	Counting strings	19
2.3.2	String offsets or virtual addresses	19
2.3.3	Match length	20
2.3.4	File size	20
2.3.5	Executable entry point	21
2.3.6	Accessing data at a given position	21
2.3.7	Sets of strings	22
2.3.8	Applying the same condition to many strings	24
2.3.9	Using anonymous strings with <code>of</code> and <code>for..of</code>	24
2.3.10	Iterating over string occurrences	25
2.3.11	Iterators	25
2.3.12	Referencing other rules	26
2.4	More about rules	27
2.4.1	Global rules	27
2.4.2	Private rules	28
2.4.3	Rule tags	28
2.4.4	Metadata	28
2.5	Using modules	29
2.6	Undefined values	29

2.7	External variables	30
2.8	Including files	31
3	Modules	33
3.1	PE module	33
3.1.1	Reference	34
3.2	ELF module	54
3.2.1	Reference	54
3.3	Cuckoo module	60
3.3.1	Reference	61
3.4	Magic module	62
3.5	Hash module	63
3.6	Math module	64
3.7	dotnet module	66
3.7.1	Reference	67
3.8	Time module	70
3.9	Console module	70
3.10	String module	71
4	Writing your own modules	73
4.1	The "Hello World!" module	73
4.1.1	Building our "Hello World!"	75
4.2	The declaration section	77
4.2.1	Basic types	77
4.2.2	Structures	77
4.2.3	Arrays	78
4.2.4	Dictionaries	79
4.2.5	Functions	79
4.3	Initialization and finalization	80
4.4	Implementing the module's logic	81
4.4.1	Accessing the scanned data	82
4.4.2	Setting variable's values	83
4.4.3	Storing data for later use	85
4.5	More about functions	86
4.5.1	Function arguments	86
4.5.2	Return values	87
4.5.3	Accessing objects	87
4.5.4	Scan context	88
5	Running YARA from the command-line	89
6	Using YARA from Python	93
6.1	Reference	98
7	The C API	101
7.1	Initializing and finalizing <i>libyara</i>	101
7.2	Compiling rules	101
7.3	Defining external variables	103
7.4	Saving and retrieving compiled rules	103
7.5	Scanning data	104
7.5.1	Using a scanner	105
7.6	API reference	106
7.6.1	Data structures	106
7.6.2	Functions	108
7.6.3	Error codes	114

Python Module Index	117
Index	119

YARA is a tool aimed at (but not limited to) helping malware researchers to identify and classify malware samples. With YARA you can create descriptions of malware families (or whatever you want to describe) based on textual or binary patterns. Each description, a.k.a. rule, consists of a set of strings and a boolean expression which determine its logic. Let's see an example:

```
rule silent_banker : banker
{
    meta:
        description = "This is just an example"
        threat_level = 3
        in_the_wild = true
    strings:
        $a = {6A 40 68 00 30 00 00 6A 14 8D 91}
        $b = {8D 4D B0 2B C1 83 C0 27 99 6A 4E 59 F7 F9}
        $c = "UVODFRYSIHLNWPEJXQZAKCBGMT"
    condition:
        $a or $b or $c
}
```

The above rule is telling YARA that any file containing one of the three strings must be reported as `silent_banker`. This is just a simple example, more complex and powerful rules can be created by using wild-cards, case-insensitive strings, regular expressions, special operators and many other features that you'll find explained in this documentation.

Contents:

CHAPTER 1

Getting started

YARA is a multi-platform program running on Windows, Linux and Mac OS X. You can find the latest release at <https://github.com/VirusTotal/yara/releases>.

1.1 Compiling and installing YARA

Download the source tarball and get prepared for compiling it:

```
tar -zxf yara-4.3.0.tar.gz
cd yara-4.3.0
./bootstrap.sh
```

Make sure you have `automake`, `libtool`, `make` and `gcc` and `pkg-config` installed in your system. Ubuntu and Debian users can use:

```
sudo apt-get install automake libtool make gcc pkg-config
```

If you plan to modify YARA's source code you may also need `flex` and `bison` for generating lexers and parsers:

```
sudo apt-get install flex bison
```

Compile and install YARA in the standard way:

```
./bootstrap.sh
./configure
make
sudo make install
```

Run the test cases to make sure that everything is fine:

```
make check
```

Some of YARA's features depend on the OpenSSL library. Those features are enabled only if you have the OpenSSL library installed in your system. If not, YARA is going to work fine but you won't be able to use the disabled features. The `configure` script will automatically detect if OpenSSL is installed or not. If you want to enforce the OpenSSL-dependent features you must pass `--with-crypto` to the `configure` script. Ubuntu and Debian users can use `sudo apt-get install libssl-dev` to install the OpenSSL library.

The following modules are not compiled into YARA by default:

- cuckoo
- magic
- dotnet

If you plan to use them you must pass the corresponding `--enable-<module name>` arguments to the `configure` script.

For example:

```
./configure --enable-cuckoo
./configure --enable-magic
./configure --enable-dotnet
./configure --enable-cuckoo --enable-magic --enable-dotnet
```

Modules usually depend on external libraries, depending on the modules you choose to install you'll need the following libraries:

- **cuckoo:** Depends on [Jansson](#) for parsing JSON. Some Ubuntu and Debian versions already include a package named `libjansson-dev`, if `sudo apt-get install libjansson-dev` doesn't work for you then get the source code from [its repository](#).
- **magic:** Depends on *libmagic*, a library used by the Unix standard program [file](#). Ubuntu, Debian and CentOS include a package `libmagic-dev`. The source code can be found [here](#).

1.1.1 Installing with vcpkg

You can also download and install YARA using the [vcpkg](#) dependency manager:

```
git clone https://github.com/microsoft/vcpkg.git
cd vcpkg
./bootstrap-vcpkg.sh
./vcpkg integrate install
vcpkg install yara
```

The YARA port in `vcpkg` is kept up to date by Microsoft team members and community contributors. If the version is out of date, please [create an issue or pull request](#) on the `vcpkg` repository.

1.1.2 Installing on Windows

Compiled binaries for Windows in both 32 and 64 bit flavors can be found in the link below. Just download the version you want, unzip the archive, and put the `yara.exe` and `yarac.exe` binaries anywhere in your disk.

[Download Windows binaries](#)

To install YARA using [Scoop](#) or [Chocolatey](#), simply type `scoop install yara` or `choco install yara`. The integration with both *Scoop* and *Chocolatey* are not maintained their respective teams, not by the YARA authors.

1.1.3 Installing on Mac OS X with Homebrew

To install YARA using [Homebrew](#), simply type `brew install yara`.

1.1.4 Installing yara-python

If you plan to use YARA from your Python scripts you need to install the `yara-python` extension. Please refer to <https://github.com/VirusTotal/yara-python> for instructions on how to install it.

1.2 Running YARA for the first time

Now that you have installed YARA you can write a very simple rule and use the command-line tool to scan some file:

```
echo "rule dummy { condition: true }" > my_first_rule
yara my_first_rule my_first_rule
```

Don't get confused by the repeated `my_first_rule` in the arguments to `yara`, I'm just passing the same file as both the rules and the file to be scanned. You can pass any file you want to be scanned (second argument).

If everything goes fine you should get the following output:

```
dummy my_first_rule
```

Which means that the file `my_first_rule` is matching the rule named `dummy`.

If you get an error like this:

```
yara: error while loading shared libraries: libyara.so.2: cannot open shared
object file: No such file or directory
```

It means that the loader is not finding the `libyara` library which is located in `/usr/local/lib`. In some Linux flavors the loader doesn't look for libraries in this path by default, we must instruct it to do so by adding `/usr/local/lib` to the loader configuration file `/etc/ld.so.conf`:

```
sudo sh -c 'echo "/usr/local/lib" >> /etc/ld.so.conf'
sudo ldconfig
```

On newer Ubuntu releases such as 22.04 LTS, the correct loader configuration is installed via dependencies to `/etc/ld.so.conf.d/libc.conf`. In this case, the following command alone is sufficient to configure the dynamic linker run-time bindings.

```
sudo ldconfig
```

If you're using Windows PowerShell as your command shell, `yara my_first_rule my_first_rule` may return this error:

```
my_first_rule(1): error: non-ascii character
```

You can avoid this by using the `Set-Content` cmdlet to specify ascii output when creating your rule file:

```
Set-Content -path .\my_first_rule -Value "rule dummy { condition: true }" -Encoding_
↪Ascii
.\yara my_first_rule my_first_rule
```


CHAPTER 2

Writing YARA rules

YARA rules are easy to write and understand, and they have a syntax that resembles the C language. Here is the simplest rule that you can write for YARA, which does absolutely nothing:

```
rule dummy
{
    condition:
        false
}
```

Each rule in YARA starts with the keyword `rule` followed by a rule identifier. Identifiers must follow the same lexical conventions of the C programming language, they can contain any alphanumeric character and the underscore character, but the first character cannot be a digit. Rule identifiers are case sensitive and cannot exceed 128 characters. The following keywords are reserved and cannot be used as an identifier:

Table 1: YARA keywords

all	and	any	ascii	at	base64	base64wide	condition
contains	endswith	entrypoint	false	filesize	for	fullword	global
import	icontains	iendswith	iequals	in	include	int16	int16be
int32	int32be	int8	int8be	istartswith	matches	meta	nocase
none	not	of	or	private	rule	startswith	strings
them	true	uint16	uint16be	uint32	uint32be	uint8	uint8be
wide	xor	defined					

Rules are generally composed of two sections: strings definition and condition. The strings definition section can be omitted if the rule doesn't rely on any string, but the condition section is always required. The strings definition section is where the strings that will be part of the rule are defined. Each string has an identifier consisting of a \$ character followed by a sequence of alphanumeric characters and underscores, these identifiers can be used in the condition section to refer to the corresponding string. Strings can be defined in text or hexadecimal form, as shown in the following example:

```
rule ExampleRule
{
    strings:
        $my_text_string = "text here"
        $my_hex_string = { E2 34 A1 C8 23 FB }

    condition:
        $my_text_string or $my_hex_string
}
```

Text strings are enclosed in double quotes just like in the C language. Hex strings are enclosed by curly brackets, and they are composed by a sequence of hexadecimal numbers that can appear contiguously or separated by spaces. Decimal numbers are not allowed in hex strings.

The condition section is where the logic of the rule resides. This section must contain a boolean expression telling under which circumstances a file or process satisfies the rule or not. Generally, the condition will refer to previously defined strings by using their identifiers. In this context the string identifier acts as a boolean variable which evaluate to true if the string was found in the file or process memory, or false if otherwise.

2.1 Comments

You can add comments to your YARA rules just as if it was a C source file, both single-line and multi-line C-style comments are supported.

```
/*
    This is a multi-line comment ...
*/

rule CommentExample // ... and this is single-line comment
{
    condition:
        false // just a dummy rule, don't do this
}
```

2.2 Strings

There are three types of strings in YARA: hexadecimal strings, text strings and regular expressions. Hexadecimal strings are used for defining raw sequences of bytes, while text strings and regular expressions are useful for defining portions of legible text. However text strings and regular expressions can be also used for representing raw bytes by mean of escape sequences as will be shown below.

2.2.1 Hexadecimal strings

Hexadecimal strings allow four special constructions that make them more flexible: wild-cards, not operators, jumps, and alternatives. Wild-cards are just placeholders that you can put into the string indicating that some bytes are unknown and they should match anything. The placeholder character is the question mark (?). Here you have an example of a hexadecimal string with wild-cards:

```
rule WildcardExample
{
```

(continues on next page)

(continued from previous page)

```

strings:
    $hex_string = { E2 34 ?? C8 A? FB }

condition:
    $hex_string
}

```

As shown in the example the wild-cards are nibble-wise, which means that you can define just one nibble of the byte and leave the other unknown.

Starting with version 4.3.0, you may specify that a byte is not a specific value. For that you can use the not operator with a byte value:

```

rule NotExample
{
    strings:
        $hex_string = { F4 23 ~00 62 B4 }
        $hex_string2 = { F4 23 ~?0 62 B4 }
    condition:
        $hex_string and $hex_string2
}

```

In the example above we have a byte prefixed with a tilde (~), which is the not operator. This defines that the byte in that location can take any value except the value specified. In this case the first string will only match if the byte is not 00. The not operator can also be used with nibble-wise wild-cards, so the second string will only match if the second nibble is not zero.

Wild-cards and not operators are useful when defining strings whose content can vary but you know the length of the variable chunks, however, this is not always the case. In some circumstances you may need to define strings with chunks of variable content and length. In those situations you can use jumps instead of wild-cards:

```

rule JumpExample
{
    strings:
        $hex_string = { F4 23 [4-6] 62 B4 }

    condition:
        $hex_string
}

```

In the example above we have a pair of numbers enclosed in square brackets and separated by a hyphen, that's a jump. This jump is indicating that any arbitrary sequence from 4 to 6 bytes can occupy the position of the jump. Any of the following strings will match the pattern:

```

F4 23 01 02 03 04 62 B4
F4 23 00 00 00 00 62 B4
F4 23 15 82 A3 04 45 22 62 B4

```

Any jump [X-Y] must meet the condition $0 \leq X \leq Y$. In previous versions of YARA both X and Y must be lower than 256, but starting with YARA 2.0 there is no limit for X and Y.

These are valid jumps:

```

FE 39 45 [0-8] 89 00
FE 39 45 [23-45] 89 00
FE 39 45 [1000-2000] 89 00

```

This is invalid:

```
FE 39 45 [10-7] 89 00
```

If the lower and higher bounds are equal you can write a single number enclosed in brackets, like this:

```
FE 39 45 [6] 89 00
```

The above string is equivalent to both of these:

```
FE 39 45 [6-6] 89 00
FE 39 45 ?? ?? ?? ?? ?? 89 00
```

Starting with YARA 2.0 you can also use unbounded jumps:

```
FE 39 45 [10-] 89 00
FE 39 45 [-] 89 00
```

The first one means [10-infinite], the second one means [0-infinite].

There are also situations in which you may want to provide different alternatives for a given fragment of your hex string. In those situations you can use a syntax which resembles a regular expression:

```
rule AlternativesExample1
{
    strings:
        $hex_string = { F4 23 ( 62 B4 | 56 ) 45 }

    condition:
        $hex_string
}
```

This rule will match any file containing F42362B445 or F4235645.

But more than two alternatives can be also expressed. In fact, there are no limits to the amount of alternative sequences you can provide, and neither to their lengths.

```
rule AlternativesExample2
{
    strings:
        $hex_string = { F4 23 ( 62 B4 | 56 | 45 ?? 67 ) 45 }

    condition:
        $hex_string
}
```

As can be seen also in the above example, strings containing wild-cards are allowed as part of alternative sequences.

2.2.2 Text strings

As shown in previous sections, text strings are generally defined like this:

```
rule TextExample
{
    strings:
        $text_string = "foobar"
```

(continues on next page)

(continued from previous page)

```

condition:
    $text_string
}

```

This is the simplest case: an ASCII-encoded, case-sensitive string. However, text strings can be accompanied by some useful modifiers that alter the way in which the string will be interpreted. Those modifiers are appended at the end of the string definition separated by spaces, as will be discussed below.

Text strings can also contain the following subset of the escape sequences available in the C language:

<code>\"</code>	Double quote
<code>\\</code>	Backslash
<code>\r</code>	Carriage return
<code>\t</code>	Horizontal tab
<code>\n</code>	New line
<code>\xdd</code>	Any byte in hexadecimal notation

In all versions of YARA before 4.1.0 text strings accepted any kind of unicode characters, regardless of their encoding. Those characters were interpreted by YARA as raw bytes, and therefore the final string was actually determined by the encoding format used by your text editor. This never meant to be a feature, the original intention always was that YARA strings should be ASCII-only and YARA 4.1.0 started to raise warnings about non-ASCII characters in strings. This limitation does not apply to strings in the metadata section or comments. See more details [here](<https://github.com/VirusTotal/yara/wiki/Unicode-characters-in-YARA>)

Case-insensitive strings

Text strings in YARA are case-sensitive by default, however you can turn your string into case-insensitive mode by appending the modifier `nocase` at the end of the string definition, in the same line:

```

rule CaseInsensitiveTextExample
{
    strings:
        $text_string = "foobar" nocase

    condition:
        $text_string
}

```

With the `nocase` modifier the string *foobar* will match *foobar*, *FOOBAR*, and *fOoBaR*. This modifier can be used in conjunction with any modifier, except `base64` and `base64wide`.

Wide-character strings

The `wide` modifier can be used to search for strings encoded with two bytes per character, something typical in many executable binaries.

For example, if the string "Borland" appears encoded as two bytes per character (i.e. `B\x00o\x00r\x00l\x00a\x00n\x00d\x00`), then the following rule will match:

```

rule WideCharTextExample1
{
    strings:

```

(continues on next page)

(continued from previous page)

```
$wide_string = "Borland" wide

condition:
    $wide_string
}
```

However, keep in mind that this modifier just interleaves the ASCII codes of the characters in the string with zeroes, it does not support truly UTF-16 strings containing non-English characters. If you want to search for strings in both ASCII and wide form, you can use the `ascii` modifier in conjunction with `wide`, no matter the order in which they appear.

```
rule WideCharTextExample2
{
    strings:
        $wide_and_ascii_string = "Borland" wide ascii

    condition:
        $wide_and_ascii_string
}
```

The `ascii` modifier can appear alone, without an accompanying `wide` modifier, but it's not necessary to write it because in absence of `wide` the string is assumed to be ASCII by default.

XOR strings

The `xor` modifier can be used to search for strings with a single byte XOR applied to them.

The following rule will search for every single byte XOR applied to the string "This program cannot" (including the plaintext string):

```
rule XorExample1
{
    strings:
        $xor_string = "This program cannot" xor

    condition:
        $xor_string
}
```

The above rule is logically equivalent to:

```
rule XorExample2
{
    strings:
        $xor_string_00 = "This program cannot"
        $xor_string_01 = "Uihr!qsnfs`l!b`oonu"
        $xor_string_02 = "Vjkq\"rpmepeco\"acllmv"
        // Repeat for every single byte XOR
    condition:
        any of them
}
```

You can also combine the `xor` modifier with `wide` and `ascii` modifiers. For example, to search for the `wide` and `ascii` versions of a string after every single byte XOR has been applied you would use:

```
rule XorExample3
{
    strings:
        $xor_string = "This program cannot" xor wide ascii
    condition:
        $xor_string
}
```

The `xor` modifier is applied after every other modifier. This means that using the `xor` and `wide` together results in the XOR applying to the interleaved zero bytes. For example, the following two rules are logically equivalent:

```
rule XorExample4
{
    strings:
        $xor_string = "This program cannot" xor wide
    condition:
        $xor_string
}

rule XorExample4
{
    strings:
        $xor_string_00 = "T\x00h\x00i\x00s\x00_
→ \x00p\x00r\x00o\x00g\x00r\x00a\x00m\x00 \x00c\x00a\x00n\x00n\x00o\x00t\x00"
        $xor_string_01 = "U\x01i\x01h\x01r\x01!
→ \x01q\x01s\x01n\x01f\x01s\x01'\x01l\x01!\x01b\x01'\x01o\x01o\x01n\x01u\x01"
        $xor_string_02 = "V\x02j\x02k\x02q\x02\
→ "\x02r\x02p\x02m\x02e\x02p\x02c\x02o\x02"\x02a\x02c\x02l\x02l\x02m\x02v\x02"
        // Repeat for every single byte XOR operation.
    condition:
        any of them
}
```

Since YARA 3.11, if you want more control over the range of bytes used with the `xor` modifier use:

```
rule XorExample5
{
    strings:
        $xor_string = "This program cannot" xor(0x01-0xff)
    condition:
        $xor_string
}
```

The above example will apply the bytes from 0x01 to 0xff, inclusively, to the string when searching. The general syntax is `xor (minimum-maximum)`.

Base64 strings

The `base64` modifier can be used to search for strings that have been base64 encoded. A good explanation of the technique is at:

<https://www.leeholmes.com/searching-for-content-in-base-64-strings/>

The following rule will search for the three base64 permutations of the string "This program cannot":

```
rule Base64Example1
{
```

(continues on next page)

(continued from previous page)

```
strings:
    $a = "This program cannot" base64

condition:
    $a
}
```

This will cause YARA to search for these three permutations:

```
VGhpcyBwcm9ncmFtIGNhbm5vd
RoaXMgcHJvZ3JhbSBjYW5ub3
UaGlzIHByb2dyYW0gY2Fubm90
```

The `base64wide` modifier works just like the `base64` modifier but the results of the `base64` modifier are converted to wide.

The interaction between `base64` (or `base64wide`) and `wide` and `ascii` is as you might expect. `wide` and `ascii` are applied to the string first, and then the `base64` and `base64wide` modifiers are applied. At no point is the plaintext of the `ascii` or `wide` versions of the strings included in the search. If you want to also include those you can put them in a secondary string.

The `base64` and `base64wide` modifiers also support a custom alphabet. For example:

```
rule Base64Example2
{
    strings:
        $a = "This program cannot" base64("!@#$%^&*(){}[].,
        ↪ |ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz")

    condition:
        $a
}
```

The alphabet must be 64 bytes long.

The `base64` and `base64wide` modifiers are only supported with text strings. Using these modifiers with a hexadecimal string or a regular expression will cause a compiler error. Also, the `xor`, `fullword`, and `nocase` modifiers used in combination with `base64` or `base64wide` will cause a compiler error.

Because of the way that YARA strips the leading and trailing characters after base64 encoding, one of the base64 encodings of "Dhis program cannow" and "This program cannot" are identical. Similarly, using the `base64` keyword on single ASCII characters is not recommended. For example, "a" with the `base64` keyword matches "", "b", "c", "!", "\xA1", or "\xE1" after base64 encoding, and will not match where the base64 encoding matches the `[GWm2][EFGH]` regular expression.

Searching for full words

Another modifier that can be applied to text strings is `fullword`. This modifier guarantees that the string will match only if it appears in the file delimited by non-alphanumeric characters. For example the string `domain`, if defined as `fullword`, doesn't match `www.mydomain.com` but it matches `www.my-domain.com` and `www.domain.com`.

2.2.3 Regular expressions

Regular expressions are one of the most powerful features of YARA. They are defined in the same way as text strings, but enclosed in forward slashes instead of double-quotes, like in the Perl programming language.

```
rule RegExpExample1
{
    strings:
        $re1 = /md5: [0-9a-fA-F]{32}/
        $re2 = /state: (on|off)/

    condition:
        $re1 and $re2
}
```

Regular expressions can be also followed by `nocase`, `ascii`, `wide`, and `fullword` modifiers just like in text strings. The semantics of these modifiers are the same in both cases.

Additionally, they can be followed by the characters `i` and `s` just after the closing slash, which is a very common convention for specifying that the regular expression is case-insensitive and that the dot (`.`) can match new-line characters. For example:

```
rule RegExpExample2
{
    strings:
        $re1 = /foo/i      // This regexp is case-insensitive
        $re2 = /bar./s     // In this regexp the dot matches everything, including new-
↪line
        $re3 = /baz./is   // Both modifiers can be used together

    condition:
        any of them
}
```

Notice that `/foo/i` is equivalent to `/foo/ nocase`, but we recommend the latter when defining strings. The `/foo/i` syntax is useful when writing case-insensitive regular expressions for the `matches` operator.

In previous versions of YARA, external libraries like PCRE and RE2 were used to perform regular expression matching, but starting with version 2.0 YARA uses its own regular expression engine. This new engine implements most features found in PCRE, except a few of them like capture groups, POSIX character classes (`[[:isalpha:]]`, `[[:isdigit:]]`, etc) and backreferences.

YARA's regular expressions recognise the following metacharacters:

<code>\</code>	Quote the next metacharacter
<code>^</code>	Match the beginning of the file or negates a character class when used as the first character after the opening bracket
<code>\$</code>	Match the end of the file
<code>.</code>	Matches any single character except a newline character
<code> </code>	Alternation
<code>()</code>	Grouping
<code>[]</code>	Bracketed character class

The following quantifiers are recognised as well:

*	Match 0 or more times
+	Match 1 or more times
?	Match 0 or 1 times
{n}	Match exactly n times
{n, }	Match at least n times
{, m}	Match at most m times
{n, m}	Match n to m times

All these quantifiers have a non-greedy variant, followed by a question mark (?):

*?	Match 0 or more times, non-greedy
+	Match 1 or more times, non-greedy
??	Match 0 or 1 times, non-greedy
{n}?	Match exactly n times, non-greedy
{n, }?	Match at least n times, non-greedy
{, m}?	Match at most m times, non-greedy
{n, m}?	Match n to m times, non-greedy

The following escape sequences are recognised:

\t	Tab (HT, TAB)
\n	New line (LF, NL)
\r	Return (CR)
\f	Form feed (FF)
\a	Alarm bell
\xNN	Character whose ordinal number is the given hexadecimal number

These are the recognised character classes:

\w	Match a <i>word</i> character (alphanumeric plus “_”)
\W	Match a <i>non-word</i> character
\s	Match a whitespace character
\S	Match a non-whitespace character
\d	Match a decimal digit character
\D	Match a non-digit character

Starting with version 3.3.0 these zero-width assertions are also recognized:

\b	Match a word boundary
\B	Match except at a word boundary

2.2.4 Private strings

All strings in YARA can be marked as `private` which means they will never be included in the output of YARA. They are treated as normal strings everywhere else, so you can still use them as you wish in the condition, but they will never be shown with the `-s` flag or seen in the YARA callback if you’re using the C API.

```
rule PrivateStringExample
{
    strings:
        $text_string = "foobar" private

    condition:
        $text_string
}
```

2.2.5 String Modifier Summary

The following string modifiers are processed in the following order, but are only applicable to the string types listed.

Table 2: Text string modifiers

Keyword	String Types	Summary	Restrictions
<code>nocase</code>	Text, Regex	Ignore case	Cannot use with <code>xor</code> , <code>base64</code> , or <code>base64wide</code>
<code>wide</code>	Text, Regex	Emulate UTF16 by interleaving null (0x00) characters	None
<code>ascii</code>	Text, Regex	Also match ASCII characters, only required if <code>wide</code> is used	None
<code>xor</code>	Text	XOR text string with single byte keys	Cannot use with <code>nocase</code> , <code>base64</code> , or <code>base64wide</code>
<code>base64</code>	Text	Convert to 3 base64 encoded strings	Cannot use with <code>nocase</code> , <code>xor</code> , or <code>fullword</code>
<code>base64wide</code>	Text	Convert to 3 base64 encoded strings, then interleaving null characters like <code>wide</code>	Cannot use with <code>nocase</code> , <code>xor</code> , or <code>fullword</code>
<code>fullword</code>	Text, Regex	Match is not preceded or followed by an alphanumeric character	Cannot use with <code>base64</code> or <code>base64wide</code>
<code>private</code>	Hex, Text, Regex	Match never included in output	None

2.3 Conditions

Conditions are nothing more than Boolean expressions as those that can be found in all programming languages, for example in an *if* statement. They can contain the typical Boolean operators `and`, `or`, and `not`, and relational operators `>=`, `<=`, `<`, `>`, `==` and `!=`. Also, the arithmetic operators (+, -, *, \, %) and bitwise operators (&, |, <<, >>, ~, ^) can be used on numerical expressions.

Integers are always 64-bits long, even the results of functions like *uint8*, *uint16* and *uint32* are promoted to 64-bits. This is something you must take into account, specially while using bitwise operators (for example, `~0x01` is not `0xFE` but `0xFFFFFFFFFFFFFFFE`).

The following table lists the precedence and associativity of all operators. The table is sorted in descending precedence order, which means that operators listed on a higher row in the list are grouped prior operators listed in rows further below it. Operators within the same row have the same precedence, if they appear together in a expression the associativity determines how they are grouped.

Precedence	Operator	Description	Associativity
1	[] .	Array subscripting Structure member access	Left-to-right
2	- ~	Unary minus Bitwise not	Right-to-left
3	* \ %	Multiplication Division Remainder	Left-to-right
4	+ -	Addition Subtraction	Left-to-right
5	<< >>	Bitwise left shift Bitwise right shift	Left-to-right
6	&	Bitwise AND	Left-to-right
7	^	Bitwise XOR	Left-to-right
8		Bitwise OR	Left-to-right
9	< <= > >=	Less than Less than or equal to Greater than Greater than or equal to	Left-to-right
10	== != contains icontains startswith istartswith endswith iendswith iequals matches	Equal to Not equal to String contains substring Like contains but case-insensitive String starts with substring Like startswith but case-insensitive String ends with substring Like endswith but case-insensitive Case-insensitive string comparison String matches regular expression	Left-to-right
11	not defined	Logical NOT Check if an expression is defined	Right-to-left
12	and	Logical AND	Left-to-right
13	or	Logical OR	Left-to-right

String identifiers can be also used within a condition, acting as Boolean variables whose value depends on the presence or not of the associated string in the file.

```
rule Example
{
    strings:
        $a = "text1"
        $b = "text2"
        $c = "text3"
        $d = "text4"

    condition:
        ($a or $b) and ($c or $d)
}
```

2.3.1 Counting strings

Sometimes we need to know not only if a certain string is present or not, but how many times the string appears in the file or process memory. The number of occurrences of each string is represented by a variable whose name is the string identifier but with a # character in place of the \$ character. For example:

```
rule CountExample
{
    strings:
        $a = "dummy1"
        $b = "dummy2"

    condition:
        #a == 6 and #b > 10
}
```

This rule matches any file or process containing the string \$a exactly six times, and more than ten occurrences of string \$b.

Starting with YARA 4.2.0 it is possible to express the count of a string in an integer range, like this:

```
#a in (filesize-500..filesize) == 2
```

In this example the number of 'a' strings in the last 500 bytes of the file must equal exactly 2.

2.3.2 String offsets or virtual addresses

In the majority of cases, when a string identifier is used in a condition, we are willing to know if the associated string is anywhere within the file or process memory, but sometimes we need to know if the string is at some specific offset on the file or at some virtual address within the process address space. In such situations the operator `at` is what we need. This operator is used as shown in the following example:

```
rule AtExample
{
    strings:
        $a = "dummy1"
        $b = "dummy2"

    condition:
```

(continues on next page)

(continued from previous page)

```
$a at 100 and $b at 200
}
```

The expression `$a at 100` in the above example is true only if string `$a` is found at offset 100 within the file (or at virtual address 100 if applied to a running process). The string `$b` should appear at offset 200. Please note that both offsets are decimal, however hexadecimal numbers can be written by adding the prefix `0x` before the number as in the C language, which comes very handy when writing virtual addresses. Also note the higher precedence of the operator `at` over the `and`.

While the `at` operator allows to search for a string at some fixed offset in the file or virtual address in a process memory space, the `in` operator allows to search for the string within a range of offsets or addresses.

```
rule InExample
{
    strings:
        $a = "dummy1"
        $b = "dummy2"

    condition:
        $a in (0..100) and $b in (100..filesize)
}
```

In the example above the string `$a` must be found at an offset between 0 and 100, while string `$b` must be at an offset between 100 and the end of the file. Again, numbers are decimal by default.

You can also get the offset or virtual address of the *i*-th occurrence of string `$a` by using `@a[i]`. The indexes are one-based, so the first occurrence would be `@a[1]` the second one `@a[2]` and so on. If you provide an index greater than the number of occurrences of the string, the result will be a NaN (Not A Number) value.

2.3.3 Match length

For many regular expressions and hex strings containing jumps, the length of the match is variable. If you have the regular expression `/fo*/` the strings `"fo"`, `"foo"` and `"fooo"` can be matches, all of them with a different length.

You can use the length of the matches as part of your condition by using the character `!` in front of the string identifier, in a similar way you use the `@` character for the offset. `!a[1]` is the length for the first match of `$a`, `!a[2]` is the length for the second match, and so on. `!a` is an abbreviated form of `!a[1]`.

2.3.4 File size

String identifiers are not the only variables that can appear in a condition (in fact, rules can be defined without any string definition as will be shown below), there are other special variables that can be used as well. One of these special variables is `filesize`, which holds, as its name indicates, the size of the file being scanned. The size is expressed in bytes.

```
rule FileSizeExample
{
    condition:
        filesize > 200KB
}
```

The previous example also demonstrates the use of the `KB` postfix. This postfix, when attached to a numerical constant, automatically multiplies the value of the constant by 1024. The `MB` postfix can be used to multiply the value by 2^{20} . Both postfixes can be used only with decimal constants.

The use of `filesize` only makes sense when the rule is applied to a file. If the rule is applied to a running process it won't ever match because `filesize` doesn't make sense in this context.

2.3.5 Executable entry point

Another special variable than can be used in a rule is `entrypoint`. If the file is a Portable Executable (PE) or Executable and Linkable Format (ELF), this variable holds the raw offset of the executable's entry point in case we are scanning a file. If we are scanning a running process, the `entrypoint` will hold the virtual address of the main executable's entry point. A typical use of this variable is to look for some pattern at the entry point to detect packers or simple file infectors.

```
rule EntryPointExample1
{
    strings:
        $a = { E8 00 00 00 00 }

    condition:
        $a at entrypoint
}

rule EntryPointExample2
{
    strings:
        $a = { 9C 50 66 A1 ?? ?? ?? 00 66 A9 ?? ?? 58 0F 85 }

    condition:
        $a in (entrypoint..entrypoint + 10)
}
```

The presence of the `entrypoint` variable in a rule implies that only PE or ELF files can satisfy that rule. If the file is not a PE or ELF, any rule using this variable evaluates to false.

Warning: The `entrypoint` variable is deprecated, you should use the equivalent `pe.entry_point` from the *PE module* instead. Starting with YARA 3.0 you'll get a warning if you use `entrypoint` and it will be completely removed in future versions.

2.3.6 Accessing data at a given position

There are many situations in which you may want to write conditions that depend on data stored at a certain file offset or virtual memory address, depending on if we are scanning a file or a running process. In those situations you can use one of the following functions to read data from the file at the given offset:

```
int8(<offset or virtual address>)
int16(<offset or virtual address>)
int32(<offset or virtual address>)

uint8(<offset or virtual address>)
uint16(<offset or virtual address>)
uint32(<offset or virtual address>)

int8be(<offset or virtual address>)
int16be(<offset or virtual address>)
int32be(<offset or virtual address>)
```

(continues on next page)

(continued from previous page)

```
uint8be(<offset or virtual address>)  
uint16be(<offset or virtual address>)  
uint32be(<offset or virtual address>)
```

The `intXX` functions read 8, 16, and 32 bits signed integers from `<offset or virtual address>`, while functions `uintXX` read unsigned integers. Both 16 and 32 bit integers are considered to be little-endian. If you want to read a big-endian integer use the corresponding function ending in `be`. The `<offset or virtual address>` parameter can be any expression returning an unsigned integer, including the return value of one the `uintXX` functions itself. As an example let's see a rule to distinguish PE files:

```
rule IsPE  
{  
  condition:  
    // MZ signature at offset 0 and ...  
    uint16(0) == 0x5A4D and  
    // ... PE signature at offset stored in MZ header at 0x3C  
    uint32(uint32(0x3C)) == 0x00004550  
}
```

2.3.7 Sets of strings

There are circumstances in which it is necessary to express that the file should contain a certain number strings from a given set. None of the strings in the set are required to be present, but at least some of them should be. In these situations the `of` operator can be used.

```
rule OfExample1  
{  
  strings:  
    $a = "dummy1"  
    $b = "dummy2"  
    $c = "dummy3"  
  
  condition:  
    2 of ($a,$b,$c)  
}
```

This rule requires that at least two of the strings in the set (`$a,$b,$c`) must be present in the file, but it does not matter which two. Of course, when using this operator, the number before the `of` keyword must be less than or equal to the number of strings in the set.

The elements of the set can be explicitly enumerated like in the previous example, or can be specified by using wild cards. For example:

```
rule OfExample2  
{  
  strings:  
    $foo1 = "foo1"  
    $foo2 = "foo2"  
    $foo3 = "foo3"  
  
  condition:  
    2 of ($foo*) // equivalent to 2 of ($foo1,$foo2,$foo3)  
}
```

(continues on next page)

(continued from previous page)

```
rule OfExample3
{
    strings:
        $foo1 = "foo1"
        $foo2 = "foo2"

        $bar1 = "bar1"
        $bar2 = "bar2"

    condition:
        3 of ($foo*, $bar1, $bar2)
}
```

You can even use (\$*) to refer to all the strings in your rule, or write the equivalent keyword `them` for more legibility.

```
rule OfExample4
{
    strings:
        $a = "dummy1"
        $b = "dummy2"
        $c = "dummy3"

    condition:
        1 of them // equivalent to 1 of ($*)
}
```

In all the examples above, the number of strings have been specified by a numeric constant, but any expression returning a numeric value can be used. The keywords `any`, `all` and `none` can be used as well.

```
all of them      // all strings in the rule
any of them      // any string in the rule
all of ($a*)     // all strings whose identifier starts by $a
any of ($a,$b,$c) // any of $a, $b or $c
1 of ($*)        // same that "any of them"
none of ($b*)    // zero of the set of strings that start with "$b"
```

Warning: Due to the way YARA works internally, using "0 of them" is an ambiguous part of the language which should be avoided in favor of "none of them". To understand this, consider the meaning of "2 of them", which is true if 2 or more of the strings match. Historically, "0 of them" followed this principle and would evaluate to true if at least one of the strings matched. This ambiguity is resolved in YARA 4.3.0 by making "0 of them" evaluate to true if exactly 0 of the strings match. To improve on the situation and make the intent clear, it is encouraged to use "none" in place of 0. By not using an integer it is easier to reason about the meaning of "none of them" without the historical understanding of "at least 0" clouding the issue.

Starting with YARA 4.2.0 it is possible to express a set of strings in an integer range, like this:

```
all of ($a*) in (filesize-500..filesize)
any of ($a*, $b*) in (1000..2000)
```

Starting with YARA 4.3.0 it is possible to express a set of strings at a specific offset, like this:

```
any of ($a*) at 0
```

2.3.8 Applying the same condition to many strings

There is another operator very similar to `of` but even more powerful, the `for . . of` operator. The syntax is:

```
for expression of string_set : ( boolean_expression )
```

And its meaning is: from those strings in `string_set` at least `expression` of them must satisfy `boolean_expression`.

In other words: `boolean_expression` is evaluated for every string in `string_set` and there must be at least `expression` of them returning `True`.

Of course, `boolean_expression` can be any boolean expression accepted in the condition section of a rule, except for one important detail: here you can (and should) use a dollar sign (\$) as a place-holder for the string being evaluated. Take a look at the following expression:

```
for any of ($a,$b,$c) : ( $ at pe.entry_point )
```

The \$ symbol in the boolean expression is not tied to any particular string, it will be \$a, and then \$b, and then \$c in the three successive evaluations of the expression.

Maybe you already realised that the `of` operator is a special case of `for . . of`. The following expressions are the same:

```
any of ($a,$b,$c)
for any of ($a,$b,$c) : ( $ )
```

You can also employ the symbols #, @, and ! to make reference to the number of occurrences, the first offset, and the length of each string respectively.

```
for all of them : ( # > 3 )
for all of ($a*) : ( @ > @b )
```

Starting with YARA 4.3.0 you can express conditions over text strings like this:

```
for any s in ("71b36345516e076a0663e0bea97759e4", "1e7f7edeb06de02f2c2a9319de99e033")
↳: ( pe.imphash() == s )
```

It is worth remembering here that the two hashes referenced in the rule are normal text strings, and have nothing to do with the string section of the rule. Inside the loop condition the result of the `pe.imphash()` function is compared to each of the text strings, resulting in a more concise rule.

2.3.9 Using anonymous strings with `of` and `for . . of`

When using the `of` and `for . . of` operators followed by `them`, the identifier assigned to each string of the rule is usually superfluous. As we are not referencing any string individually we do not need to provide a unique identifier for each of them. In those situations you can declare anonymous strings with identifiers consisting only of the \$ character, as in the following example:

```
rule AnonymousStrings
{
    strings:
        $ = "dummy1"
        $ = "dummy2"

    condition:
```

(continues on next page)

(continued from previous page)

```

    1 of them
}

```

2.3.10 Iterating over string occurrences

As seen in *String offsets or virtual addresses*, the offsets or virtual addresses where a given string appears within a file or process address space can be accessed by using the syntax: `@a[i]`, where `i` is an index indicating which occurrence of the string `$a` you are referring to. (`@a[1]`, `@a[2]`,...).

Sometimes you will need to iterate over some of these offsets and guarantee they satisfy a given condition. In such cases you can use the `for..in` syntax, for example:

```

rule Occurrences
{
    strings:
        $a = "dummy1"
        $b = "dummy2"

    condition:
        for all i in (1,2,3) : ( @a[i] + 10 == @b[i] )
}

```

The previous rule says that the first occurrence of `$b` should be 10 bytes after the first occurrence of `$a`, and the same should happen with the second and third occurrences of the two strings.

The same condition could be written also as:

```

for all i in (1..3) : ( @a[i] + 10 == @b[i] )

```

Notice that we're using a range (`1..3`) instead of enumerating the index values (`1,2,3`). Of course, we're not forced to use constants to specify range boundaries, we can use expressions as well like in the following example:

```

for all i in (1..#a) : ( @a[i] < 100 )

```

In this case we're iterating over every occurrence of `$a` (remember that `#a` represents the number of occurrences of `$a`). This rule is specifying that every occurrence of `$a` should be within the first 100 bytes of the file.

In case you want to express that only some occurrences of the string should satisfy your condition, the same logic seen in the `for..of` operator applies here:

```

for any i in (1..#a) : ( @a[i] < 100 )
for 2 i in (1..#a) : ( @a[i] < 100 )

```

The `for..in` operator is similar to `for..of`, but the latter iterates over a set of strings, while the former iterates over ranges, enumerations, arrays and dictionaries.

2.3.11 Iterators

In YARA 4.0 the `for..in` operator was improved and now it can be used to iterate not only over integer enumerations and ranges (e.g: `1,2,3,4` and `1..4`), but also over any kind of iterable data type, like arrays and dictionaries defined by YARA modules. For example, the following expression is valid in YARA 4.0:

```

for any section in pe.sections : ( section.name == ".text" )

```

This is equivalent to:

```
for any i in (0..pe.number_of_sections-1) : ( pe.sections[i].name == ".text" )
```

The new syntax is more natural and easy to understand, and is the recommended way of expressing this type of conditions in newer versions of YARA.

While iterating dictionaries you must provide two variable names that will hold the key and value for each entry in the dictionary, for example:

```
for any k,v in some_dict : ( k == "foo" and v == "bar" )
```

In general the `for..in` operator has the form:

```
for <quantifier> <variables> in <iterable> : ( <some condition using the loop_
↪variables> )
```

Where *<quantifier>* is either *any*, *all* or an expression that evaluates to the number of items in the iterator that must satisfy the condition, *<variables>* is a comma-separated list of variable names that holds the values for the current item (the number of variables depend on the type of *<iterable>*) and *<iterable>* is something that can be iterated.

2.3.12 Referencing other rules

When writing the condition for a rule you can also make reference to a previously defined rule in a manner that resembles a function invocation of traditional programming languages. In this way you can create rules that depend on others. Let's see an example:

```
rule Rule1
{
    strings:
        $a = "dummy1"

    condition:
        $a
}

rule Rule2
{
    strings:
        $a = "dummy2"

    condition:
        $a and Rule1
}
```

As can be seen in the example, a file will satisfy Rule2 only if it contains the string "dummy2" and satisfies Rule1. Note that it is strictly necessary to define the rule being invoked before the one that will make the invocation.

Another way to reference other rules was introduced in 4.2.0 and that is sets of rules, which operate similarly to sets of strings (see sets-of-strings). For example:

```
rule Rule1
{
    strings:
        $a = "dummy1"
```

(continues on next page)

(continued from previous page)

```
    condition:
        $a
}

rule Rule2
{
    strings:
        $a = "dummy2"

    condition:
        $a
}

rule MainRule
{
    strings:
        $a = "dummy2"

    condition:
        any of (Rule*)
}
```

This example demonstrates how to use rule sets to describe higher order logic in a way which automatically grows with your rules. If you define another rule named `Rule3` before `MainRule` then it will automatically be included in the expansion of `Rule*` in the condition for `MainRule`.

To use rule sets all of the rules included in the set **must** exist prior to the rule set being used. For example, the following will produce a compiler error because `a2` is defined after the rule set is used in `x`:

```
rule a1 { condition: true }
rule x { condition: 1 of (a*) }
rule a2 { condition: true }
```

2.4 More about rules

There are some aspects of YARA rules that have not been covered yet, but are still very important. These are: global rules, private rules, tags and metadata.

2.4.1 Global rules

Global rules give you the possibility of imposing restrictions in all your rules at once. For example, suppose that you want all your rules to ignore files that exceed a certain size limit. You could go rule by rule making the required modifications to their conditions, or just write a global rule like this one:

```
global rule SizeLimit
{
    condition:
        filesize < 2MB
}
```

You can define as many global rules as you want, they will be evaluated before the rest of the rules, which in turn will be evaluated only if all global rules are satisfied.

2.4.2 Private rules

Private rules are a very simple concept. They are just rules that are not reported by YARA when they match on a given file. Rules that are not reported at all may seem sterile at first glance, but when mixed with the possibility offered by YARA of referencing one rule from another (see [Referencing other rules](#)) they become useful. Private rules can serve as building blocks for other rules, and at the same time prevent cluttering YARA's output with irrelevant information. To declare a rule as private just add the keyword `private` before the rule declaration.

```
private rule PrivateRuleExample
{
    ...
}
```

You can apply both `private` and `global` modifiers to a rule, resulting in a global rule that does not get reported by YARA but must be satisfied.

2.4.3 Rule tags

Another useful feature of YARA is the possibility of adding tags to rules. Those tags can be used later to filter YARA's output and show only the rules that you are interested in. You can add as many tags as you want to a rule, they are declared after the rule identifier as shown below:

```
rule TagsExample1 : Foo Bar Baz
{
    ...
}

rule TagsExample2 : Bar
{
    ...
}
```

Tags must follow the same lexical convention of rule identifiers, therefore only alphanumeric characters and under-scores are allowed, and the tag cannot start with a digit. They are also case sensitive.

When using YARA you can output only those rules which are tagged with the tag or tags that you provide.

2.4.4 Metadata

Besides the string definition and condition sections, rules can also have a metadata section where you can put additional information about your rule. The metadata section is defined with the keyword `meta` and contains identifier/value pairs like in the following example:

```
rule MetadataExample
{
    meta:
        my_identifier_1 = "Some string data"
        my_identifier_2 = 24
        my_identifier_3 = true

    strings:
        $my_text_string = "text here"
        $my_hex_string = { E2 34 A1 C8 23 FB }

    condition:
```

(continues on next page)

(continued from previous page)

```
$my_text_string or $my_hex_string  
}
```

As can be seen in the example, metadata identifiers are always followed by an equals sign and the value assigned to them. The assigned values can be strings (valid UTF8 only), integers, or one of the boolean values true or false. Note that identifier/value pairs defined in the metadata section cannot be used in the condition section, their only purpose is to store additional information about the rule.

2.5 Using modules

Modules are extensions to YARA's core functionality. Some modules like the *PE module* and the *Cuckoo module* are officially distributed with YARA and additional ones can be created by third-parties or even yourself as described in *Writing your own modules*.

The first step to using a module is importing it with the `import` statement. These statements must be placed outside any rule definition and followed by the module name enclosed in double-quotes. Like this:

```
import "pe"  
import "cuckoo"
```

After importing the module you can make use of its features, always using `<module name>.` as a prefix to any variable or function exported by the module. For example:

```
pe.entry_point == 0x1000  
cuckoo.http_request(/someregexp/)
```

2.6 Undefined values

Modules often leave variables in an undefined state, for example when the variable doesn't make sense in the current context (think of `pe.entry_point` while scanning a non-PE file). YARA handles undefined values in a way that allows the rule to keep its meaningfulness. Take a look at this rule:

```
import "pe"  
  
rule Test  
{  
    strings:  
        $a = "some string"  
  
    condition:  
        $a and pe.entry_point == 0x1000  
}
```

If the scanned file is not a PE you wouldn't expect this rule to match the file, even if it contains the string, because **both** conditions (the presence of the string and the right value for the entry point) must be satisfied. However, if the condition is changed to:

```
$a or pe.entry_point == 0x1000
```

You would expect the rule to match in this case if the file contains the string, even if it isn't a PE file. That's exactly how YARA behaves. The logic is as follows:

- If the expression in the condition is undefined, it would be translated to `false` and the rule won't match.
- Boolean operators `and` and `or` will treat undefined operands as `false`, Which means that:
 - `undefined and true` is `false`
 - `undefined and false` is `false`
 - `undefined or true` is `true`
 - `undefined or false` is `false`
- All the remaining operators, including the `not` operator, return undefined if any of their operands is undefined.

In the expression above, `pe.entry_point == 0x1000` will be undefined for non-PE files, because `pe.entry_point` is undefined for those files. This implies that `$a or pe.entry_point == 0x1000` will be true if and only if `$a` is true.

If the condition is `pe.entry_point == 0x1000` alone, it will evaluate to `false` for non-PE files, and so will do `pe.entry_point != 0x1000` and `not pe.entry_point == 0x1000`, as none of these expressions make sense for non-PE files.

To check if expression is defined use unary operator `defined`. Example:

```
defined pe.entry_point
```

2.7 External variables

External variables allow you to define rules that depend on values provided from the outside. For example, you can write the following rule:

```
rule ExternalVariableExample1
{
    condition:
        ext_var == 10
}
```

In this case `ext_var` is an external variable whose value is assigned at run-time (see `-d` option of command-line tool, and `externals` parameter of `compile` and `match` methods in `yara-python`). External variables could be of types: integer, string or boolean; their type depends on the value assigned to them. An integer variable can substitute any integer constant in the condition and boolean variables can occupy the place of boolean expressions. For example:

```
rule ExternalVariableExample2
{
    condition:
        bool_ext_var or filesize < int_ext_var
}
```

External variables of type string can be used with the operators: `contains`, `startswith`, `endswith` and their case-insensitive counterparts: `icontains`, `istartswith` and `iendswith`. They can be used also with the `matches` operator, which returns true if the string matches a given regular expression. Case-insensitive string comparison can be done through special operator `iequals` which only works with strings. For case-sensitive comparison use regular `==`.

```
rule ContainsExample
{
    condition:
```

(continues on next page)

(continued from previous page)

```
        string_ext_var contains "text"
    }

    rule CaseInsensitiveContainsExample
    {
        condition:
            string_ext_var icontains "text"
    }

    rule StartsWithExample
    {
        condition:
            string_ext_var startswith "prefix"
    }

    rule EndsWithExample
    {
        condition:
            string_ext_var endswith "suffix"
    }

    rule IequalsExample
    {
        condition:
            string_ext_var iequals "string"
    }

    rule MatchesExample
    {
        condition:
            string_ext_var matches /[a-z]+/
    }
}
```

You can use regular expression modifiers along with the `matches` operator, for example, if you want the regular expression from the previous example to be case insensitive you can use `/[a-z]+/i`. Notice the `i` following the regular expression in a Perl-like manner. You can also use the `s` modifier for single-line mode, in this mode the dot matches all characters including line breaks. Of course both modifiers can be used simultaneously, like in the following example:

```
rule ExternalVariableExample5
{
    condition:
        /* case insensitive single-line mode */
        string_ext_var matches /[a-z]+/is
}
```

Keep in mind that every external variable used in your rules must be defined at run-time, either by using the `-d` option of the command-line tool, or by providing the `externals` parameter to the appropriate method in `yara-python`.

2.8 Including files

In order to allow for more flexible organization of your rules files, YARA provides the `include` directive. This directive works in a similar way to the `#include` pre-processor directive in C programs, which inserts the content of the specified source file into the current file during compilation. The following example will include the content of

other.yar into the current file:

```
include "other.yar"
```

The base path when searching for a file in an `include` directive will be the directory where the current file resides. For this reason, the file *other.yar* in the previous example should be located in the same directory of the current file. However, you can also specify relative paths like these:

```
include "../includes/other.yar"  
include "../../includes/other.yar"
```

Or use absolute paths:

```
include "/home/plusvic/yara/includes/other.yar"
```

In Windows, both forward and back slashes are accepted, but don't forget to write the drive letter:

```
include "c:/yara/includes/other.yar"  
include "c:\\yara\\includes\\other.yar"
```

Modules are the method YARA provides for extending its features. They allow you to define data structures and functions which can be used in your rules to express more complex conditions. Here you'll find described some modules officially distributed with YARA, but you can also learn how to write your own modules in the [Writing your own modules](#) section.

3.1 PE module

The PE module allows you to create more fine-grained rules for PE files by using attributes and features of the PE file format. This module exposes most of the fields present in a PE header and provides functions which can be used to write more expressive and targeted rules. Let's see some examples:

```
import "pe"

rule single_section
{
    condition:
        pe.number_of_sections == 1
}

rule control_panel_applet
{
    condition:
        pe.exports("CPlApplet")
}

rule is_dll
{
    condition:
        pe.characteristics & pe.DLL
}

rule is_pe
```

(continues on next page)

(continued from previous page)

```
{  
    condition:  
        pe.is_pe  
}
```

3.1.1 Reference

machine

Changed in version 3.3.0.

Integer with one of the following values:

MACHINE_UNKNOWN

MACHINE_AM33

MACHINE_AMD64

MACHINE_ARM

MACHINE_ARMNT

MACHINE_ARM64

MACHINE_EBC

MACHINE_I386

MACHINE_IA64

MACHINE_M32R

MACHINE_MIPS16

MACHINE_MIPSFPU

MACHINE_MIPSFPU16

MACHINE_POWERPC

MACHINE_POWERPCFP

MACHINE_R4000

MACHINE_SH3

MACHINE_SH3DSP

MACHINE_SH4

MACHINE_SH5

MACHINE_THUMB

MACHINE_WCEMIPSV2

MACHINE_TARGET_HOST

MACHINE_R3000

MACHINE_R10000

MACHINE_ALPHA

MACHINE_SH3E

MACHINE_ALPHA64

MACHINE_AXP64

MACHINE_TRICORE

MACHINE_CEF

MACHINE_CEE

Example: pe.machine == pe.MACHINE_AMD64

checksum

New in version 3.6.0.

Integer with the "PE checksum" as stored in the OptionalHeader

calculate_checksum

New in version 3.6.0.

Function that calculates the "PE checksum"

Example: pe.checksum == pe.calculate_checksum()

subsystem

Integer with one of the following values:

SUBSYSTEM_UNKNOWN

SUBSYSTEM_NATIVE

SUBSYSTEM_WINDOWS_GUI

SUBSYSTEM_WINDOWS_CUI

SUBSYSTEM_OS2_CUI

SUBSYSTEM_POSIX_CUI

SUBSYSTEM_NATIVE_WINDOWS

SUBSYSTEM_WINDOWS_CE_GUI

SUBSYSTEM_EFI_APPLICATION

SUBSYSTEM_EFI_BOOT_SERVICE_DRIVER

SUBSYSTEM_EFI_RUNTIME_DRIVER

SUBSYSTEM_EFI_ROM_IMAGE

SUBSYSTEM_XBOX

SUBSYSTEM_WINDOWS_BOOT_APPLICATION

Example: pe.subsystem == pe.SUBSYSTEM_NATIVE

timestamp

PE timestamp, as an epoch integer.

Example: pe.timestamp >= 1424563200

pointer_to_symbol_table

New in version 3.8.0.

Value of IMAGE_FILE_HEADER::PointerToSymbolTable. Used when the PE image has COFF debug info.

number_of_symbols

New in version 3.8.0.

Value of IMAGE_FILE_HEADER::NumberOfSymbols. Used when the PE image has COFF debug info.

size_of_optional_header

New in version 3.8.0.

Value of IMAGE_FILE_HEADER::SizeOfOptionalHeader. This is real size of the optional header and reflects differences between 32-bit and 64-bit optional header and number of data directories.

opthdr_magic

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::Magic.

Integer with one of the following values:

IMAGE_NT_OPTIONAL_HDR32_MAGIC

IMAGE_NT_OPTIONAL_HDR64_MAGIC

IMAGE_ROM_OPTIONAL_HDR_MAGIC

size_of_code

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfCode. This is the sum of raw data sizes in code sections.

size_of_initialized_data

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfInitializedData.

size_of_uninitialized_data

Value of IMAGE_OPTIONAL_HEADER::SizeOfUninitializedData.

entry_point

Entry point file offset or virtual address depending on whether YARA is scanning a file or process memory respectively. This is equivalent to the deprecated `entrypoint` keyword.

entry_point_raw

Entry point raw value from the optional header of the PE. This value is not converted to a file offset or an RVA.

New in version 4.1.0.

base_of_code

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::BaseOfCode.

base_of_data

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::BaseOfData. This field only exists in 32-bit PE files.

image_base

Image base relative virtual address.

section_alignment

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SectionAlignment. When Windows maps a PE image to memory, all raw sizes (including size of header) are aligned up to this value.

file_alignment

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::FileAlignment. All raw data sizes of sections in the PE image are aligned to this value.

win32_version_value

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::Win32VersionValue.

size_of_image

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfImage. This is the total virtual size of header and all sections.

size_of_headers

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfHeaders. This is the raw data size of the PE headers including DOS header, file header, optional header and all section headers. When PE is mapped to memory, this value is subject to aligning up to SectionAlignment.

characteristics

Bitmap with PE FileHeader characteristics. Individual characteristics can be inspected by performing a bitwise AND operation with the following constants:

RELOCS_STRIPPED

Relocation info stripped from file.

EXECUTABLE_IMAGE

File is executable (i.e. no unresolved external references).

LINE_NUMS_STRIPPED

Line numbers stripped from file.

LOCAL_SYMS_STRIPPED

Local symbols stripped from file.

AGGRESSIVE_WS_TRIM

Aggressively trim working set

LARGE_ADDRESS_AWARE

App can handle >2gb addresses

BYTES_REVERSED_LO

Bytes of machine word are reversed.

MACHINE_32BIT

32 bit word machine.

DEBUG_STRIPPED

Debugging info stripped from file in .DBG file

REMOVABLE_RUN_FROM_SWAP

If Image is on removable media, copy and run from the swap file.

NET_RUN_FROM_SWAP

If Image is on Net, copy and run from the swap file.

SYSTEM

System File.

DLL

File is a DLL.

UP_SYSTEM_ONLY

File should only be run on a UP machine

BYTES_REVERSED_HI

Bytes of machine word are reversed.

Example: pe.characteristics & pe.DLL

linker_version

An object with two integer attributes, one for each major and minor linker version.

major

Major linker version.

minor

Minor linker version.

os_version

An object with two integer attributes, one for each major and minor OS version.

major

Major OS version.

minor

Minor OS version.

image_version

An object with two integer attributes, one for each major and minor image version.

major

Major image version.

minor

Minor image version.

subsystem_version

An object with two integer attributes, one for each major and minor subsystem version.

major

Major subsystem version.

minor

Minor subsystem version.

dll_characteristics

Bitmap with PE OptionalHeader DllCharacteristics. Do not confuse these flags with the PE FileHeader Characteristics. Individual characteristics can be inspected by performing a bitwise AND operation with the following constants:

HIGH_ENTROPY_VA

ASLR with 64 bit address space.

DYNAMIC_BASE

File can be relocated - also marks the file as ASLR compatible

FORCE_INTEGRITY

NX_COMPAT

Marks the file as DEP compatible

NO_ISOLATION

NO_SEH

The file does not contain structured exception handlers, this must be set to use SafeSEH

NO_BIND**APPCONTAINER**

Image should execute in an AppContainer

WDM_DRIVER

Marks the file as a Windows Driver Model (WDM) device driver.

GUARD_CF

Image supports Control Flow Guard.

TERMINAL_SERVER_AWARE

Marks the file as terminal server compatible

size_of_stack_reserve

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfStackReserve. This is the default amount of virtual memory that will be reserved for stack.

size_of_stack_commit

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfStackCommit. This is the default amount of virtual memory that will be allocated for stack.

size_of_heap_reserve

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfHeapReserve. This is the default amount of virtual memory that will be reserved for main process heap.

size_of_heap_commit

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::SizeOfHeapCommit. This is the default amount of virtual memory that will be allocated for main process heap.

loader_flags

New in version 3.8.0.

Value of IMAGE_OPTIONAL_HEADER::LoaderFlags.

number_of_rva_and_sizes

Value of IMAGE_OPTIONAL_HEADER::NumberOfRvaAndSizes. This is the number of items in the IMAGE_OPTIONAL_HEADER::DataDirectory array.

data_directories

New in version 3.8.0.

A zero-based array of data directories. Each data directory contains virtual address and length of the appropriate data directory. Each data directory has the following entries:

virtual_address

Relative virtual address (RVA) of the PE data directory. If this is zero, then the data directory is missing. Note that for digital signature, this is the file offset, not RVA.

size

Size of the PE data directory, in bytes.

The index for the data directory entry can be one of the following values:

IMAGE_DIRECTORY_ENTRY_EXPORT

Data directory for exported functions.

IMAGE_DIRECTORY_ENTRY_IMPORT

Data directory for import directory.

IMAGE_DIRECTORY_ENTRY_RESOURCE

Data directory for resource section.

IMAGE_DIRECTORY_ENTRY_EXCEPTION

Data directory for exception information.

IMAGE_DIRECTORY_ENTRY_SECURITY

This is the raw file offset and length of the image digital signature. If the image has no embedded digital signature, this directory will contain zeros.

IMAGE_DIRECTORY_ENTRY_BASERELOC

Data directory for image relocation table.

IMAGE_DIRECTORY_ENTRY_DEBUG

Data directory for debug information.

IMAGE_DEBUG_DIRECTORY::Type values:

IMAGE_DEBUG_TYPE_UNKNOWN

IMAGE_DEBUG_TYPE_COFF

IMAGE_DEBUG_TYPE_CODEVIEW

IMAGE_DEBUG_TYPE_FPO

IMAGE_DEBUG_TYPE_MISC

IMAGE_DEBUG_TYPE_EXCEPTION

IMAGE_DEBUG_TYPE_FIXUP

IMAGE_DEBUG_TYPE_OMAP_TO_SRC

IMAGE_DEBUG_TYPE_OMAP_FROM_SRC

IMAGE_DEBUG_TYPE_BORLAND

IMAGE_DEBUG_TYPE_RESERVED10

IMAGE_DEBUG_TYPE_CLSID

IMAGE_DEBUG_TYPE_VC_FEATURE

IMAGE_DEBUG_TYPE_POGO

IMAGE_DEBUG_TYPE_ILTCG

IMAGE_DEBUG_TYPE_MPX

IMAGE_DEBUG_TYPE_REPRO

IMAGE_DIRECTORY_ENTRY_ARCHITECTURE

IMAGE_DIRECTORY_ENTRY_COPYRIGHT

IMAGE_DIRECTORY_ENTRY_TLS

Data directory for image thread local storage.

IMAGE_DIRECTORY_ENTRY_LOAD_CONFIG

Data directory for image load configuration.

IMAGE_DIRECTORY_ENTRY_BOUND_IMPORT

Data directory for image bound import table.

IMAGE_DIRECTORY_ENTRY_IAT

Data directory for image Import Address Table.

IMAGE_DIRECTORY_ENTRY_DELAY_IMPORT

Data directory for Delayed Import Table. Structure of the delayed import table is linker-dependent. Microsoft version of delayed imports is described in the sources "delayimp.h" and "delayimp.cpp", which can be found in MS Visual Studio 2008 CRT sources.

IMAGE_DIRECTORY_ENTRY_COM_DESCRIPTOR

Data directory for .NET headers.

Example: `pe.data_directories[pe.IMAGE_DIRECTORY_ENTRY_EXPORT].virtual_address != 0`

number_of_sections

Number of sections in the PE.

sections

New in version 3.3.0.

A zero-based array of section objects, one for each section the PE has. Individual sections can be accessed by using the [] operator. Each section object has the following attributes:

name

Section name.

full_name

If the name in the section table contains a slash (/) followed by a representation of the decimal number in ASCII format, then this field contains a string from the specified offset in the string table. Otherwise, this field contains the same value as a name field.

Even though it's not a standard, MinGW and Cygwin compilers use this feature to store section names which are longer than 8 characters.

characteristics

Section characteristics.

virtual_address

Section virtual address.

virtual_size

Section virtual size.

raw_data_offset

Section raw offset.

raw_data_size

Section raw size.

pointer_to_relocations

New in version 3.8.0.

Value of IMAGE_SECTION_HEADER::PointerToRelocations.

pointer_to_line_numbers

New in version 3.8.0.

Value of IMAGE_SECTION_HEADER::PointerToLinenumbers.

number_of_relocations

New in version 3.8.0.

Value of IMAGE_SECTION_HEADER::NumberOfRelocations.

number_of_line_numbers

New in version 3.8.0.

Value of IMAGE_SECTION_HEADER::NumberOfLineNumbers.

Example: pe.sections[0].name == ".text"

Individual section characteristics can be inspected using a bitwise AND operation with the following constants:

SECTION_NO_PAD

SECTION_CNT_CODE

SECTION_CNT_INITIALIZED_DATA

SECTION_CNT_UNINITIALIZED_DATA

SECTION_LNK_OTHER

SECTION_LNK_INFO

SECTION_LNK_REMOVE

SECTION_LNK_COMDAT

SECTION_NO_DEFER_SPEC_EXC

SECTION_GPREL

SECTION_MEM_FARDATA

SECTION_MEM_PURGEABLE

SECTION_MEM_16BIT

SECTION_LNK_NRELOC_OVFL

SECTION_MEM_LOCKED

SECTION_MEM_PRELOAD

SECTION_ALIGN_1BYTES

SECTION_ALIGN_2BYTES

SECTION_ALIGN_4BYTES

SECTION_ALIGN_8BYTES

SECTION_ALIGN_16BYTES

SECTION_ALIGN_32BYTES

SECTION_ALIGN_64BYTES

SECTION_ALIGN_128BYTES

SECTION_ALIGN_256BYTES

SECTION_ALIGN_512BYTES

SECTION_ALIGN_1024BYTES

SECTION_ALIGN_2048BYTES

SECTION_ALIGN_4096BYTES

SECTION_ALIGN_8192BYTES

SECTION_ALIGN_MASK
SECTION_MEM_DISCARDABLE
SECTION_MEM_NOT_CACHED
SECTION_MEM_NOT_PAGED
SECTION_MEM_SHARED
SECTION_MEM_EXECUTE
SECTION_MEM_READ
SECTION_MEM_WRITE
SECTION_SCALE_INDEX

Example: `pe.sections[1].characteristics & pe.SECTION_CNT_CODE`

overlay

New in version 3.6.0.

A structure containing the following integer members:

offset

Overlay section offset. This is 0 for PE files that don't have overlaid data and undefined for non-PE files.

size

Overlay section size. This is 0 for PE files that don't have overlaid data and undefined for non-PE files.

Example: `uint8(pe.overlay.offset) == 0x0d` and `pe.overlay.size > 1024`

number_of_resources

Number of resources in the PE.

resource_timestamp

Resource timestamp. This is stored as an integer.

resource_version

An object with two integer attributes, major and minor versions.

major

Major resource version.

minor

Minor resource version.

resources

Changed in version 3.3.0.

A zero-based array of resource objects, one for each resource the PE has. Individual resources can be accessed by using the [] operator. Each resource object has the following attributes:

rva

The RVA of the resource data.

offset

Offset for the resource data. This can be undefined if the RVA is invalid.

length

Length of the resource data.

type

Type of the resource (integer).

id

ID of the resource (integer).

language

Language of the resource (integer).

type_string

Type of the resource as a string, if specified.

name_string

Name of the resource as a string, if specified.

language_string

Language of the resource as a string, if specified.

All resources must have a type, id (name), and language specified. They can be either an integer or string, but never both, for any given level.

Example: `pe.resources[0].type == pe.RESOURCE_TYPE_RCDATA`

Example: `pe.resources[0].name_string == "F\x00L\x00L\x00E\x00"`

Resource types can be inspected using the following constants:

RESOURCE_TYPE_CURSOR

RESOURCE_TYPE_BITMAP

RESOURCE_TYPE_ICON

RESOURCE_TYPE_MENU

RESOURCE_TYPE_DIALOG

RESOURCE_TYPE_STRING

RESOURCE_TYPE_FONTDIR

RESOURCE_TYPE_FONT

RESOURCE_TYPE_ACCELERATOR

RESOURCE_TYPE_RCDATA

RESOURCE_TYPE_MESSAGETABLE

RESOURCE_TYPE_GROUP_CURSOR

RESOURCE_TYPE_GROUP_ICON

RESOURCE_TYPE_VERSION

RESOURCE_TYPE_DLGINCLUDE

RESOURCE_TYPE_PLUGPLAY

RESOURCE_TYPE_VXD

RESOURCE_TYPE_ANICURSOR

RESOURCE_TYPE_ANIICON

RESOURCE_TYPE_HTML

RESOURCE_TYPE_MANIFEST

For more information refer to:

[http://msdn.microsoft.com/en-us/library/ms648009\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/ms648009(v=vs.85).aspx)

version_info

New in version 3.2.0.

Dictionary containing the PE's version information. Typical keys are:

Comments	CompanyName	FileDescription	FileVersion	InternalName
LegalCopyright	LegalTrademarks	OriginalFilename	ProductName	
ProductVersion				

For more information refer to:

[http://msdn.microsoft.com/en-us/library/windows/desktop/ms646987\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/ms646987(v=vs.85).aspx)

Example: `pe.version_info["CompanyName"]` contains `"Microsoft"`

version_info_list

Array of structures containing information about the PE's version information.

key

Key of version information.

value

Value of version information.

Example: `pe.version_info_list[0].value` contains `"Microsoft"`

number_of_signatures

Number of authenticode signatures in the PE.

is_signed

True if any of the PE signatures is verified. Verified here means, that the signature is formally correct: digests match, signer public key correctly verifies the encrypted digest, etc. But this doesn't mean that the signer (and thus the signature) can be trusted as there are no trust anchors involved in the verification.

signatures

A zero-based array of signature objects, one for each authenticode signature in the PE file. Usually PE files have a single signature.

thumbprint

New in version 3.8.0.

A string containing the thumbprint of the signature.

issuer

A string containing information about the issuer. These are some examples:

```
"/C=US/ST=Washington/L=Redmond/O=Microsoft Corporation/CN=Microsoft_
↳Code Signing PCA"

"/C=US/O=VeriSign, Inc./OU=VeriSign Trust Network/OU=Terms of use at_
↳https://www.verisign.com/rpa (c)10/CN=VeriSign Class 3 Code Signing_
↳2010 CA"

"/C=GB/ST=Greater Manchester/L=Salford/O=COMODO CA Limited/CN=COMODO_
↳Code Signing CA 2"
```

subject

A string containing information about the subject.

version

Version number.

algorithm

String representation of the algorithm used for this

signature. Usually "sha1WithRSAEncryption". It depends on the X.509 and PKCS#7 implementations and possibly their versions, consider using `algorithm_oid` instead.

algorithm_oid

Object ID of the algorithm used for this signature, expressed in numeric ASN.1 dot notation. The name contained in `algorithm` is derived from this value. The object id is expected to be stable across X.509 and PKCS#7 implementations and their versions.

For example, when using the current OpenSSL-based implementation:

```
algorithm_oid == "1.2.840.113549.1.1.11"
```

is functionally equivalent to:

```
algorithm == "sha1WithRSAEncryption"
```

serial

A string containing the serial number. This is an example:

```
"52:00:e5:aa:25:56:fc:1a:86:ed:96:c9:d4:4b:33:c7"
```

not_before

Unix timestamp on which the validity period for this signature begins.

not_after

Unix timestamp on which the validity period for this signature ends.

valid_on (*timestamp*)

Function returning true if the signature was valid on the date indicated by *timestamp*. The following sentence:

```
pe.signatures[n].valid_on(timestamp)
```

Is equivalent to:

```
timestamp >= pe.signatures[n].not_before and timestamp <= pe.  
→signatures[n].not_after
```

verified

Boolean, true if signature was successfully verified. More details about what the *verified* means is mentioned under the attribute *pe.is_signed*.

digest_alg

Name of the algorithm used for file digest. Usually "sha1" or "sha256"

digest

Digest of the file signed in the signature.

file_digest

Calculated digest using `digest_alg` of the analysed file.

number_of_certificates

Number of the certificates stored in the signature, including the ones in countersignatures.

certificates

A zero-based array of certificates stored in the signature, including the ones in countersignatures. The members of the certificates are identical to those already explained before, with the same name.

thumbprint

issuer

subject

version

algorithm

serial

not_before

not_after

signer_info

Information about the signature signer.

program_name

Optional program name stored in the signature.

digest

Signed digest of the signature.

digest_alg

Algorithm used for the digest of the signature. Usually "sha1" or "sha256"

length_of_chain

Number of certificates in the signers chain.

chain

A zero-based array of certificates in the signers chain. The members of the certificates are identical to those already explained before, with the same name.

thumbprint

issuer

subject

version

algorithm

serial

not_before

not_after

number_of_countersignatures

Number of the countersignatures of the signature.

countersignatures

A zero-based array of the countersignatures of the signature. Almost always it's just single timestamp one.

verified

Boolean, true if countersignature was successfully verified. More details about what the *verified* means is mentioned under the attribute *pe.is_signed*.

sign_time

Integer - unix time of the timestamp signing time.

digest

Signed digest of the countersignature.

digest_alg

Algorithm used for the digest of the countersignature. Usually "sha1" or "sha256"

length_of_chain

Number of certificates in the countersigners chain.

chain

A zero-based array of certificates in the countersigners chain. The members of the certificates are identical to those already explained before, with the same name.

thumbprint**issuer****subject****version****algorithm****serial****not_before****not_after****rich_signature**

Structure containing information about the PE's rich signature as documented [here](#).

offset

Offset where the rich signature starts. It will be undefined if the file doesn't have a rich signature.

length

Length of the rich signature, not including the final "Rich" marker.

key

Key used to encrypt the data with XOR.

raw_data

Raw data as it appears in the file.

clear_data

Data after being decrypted by XORing it with the key.

version_data

New in version 4.3.0.

Version fields after being decrypted by XORing it with the key.

version (*version*, [*toolid*])

New in version 3.5.0.

Function returning a sum of count values of all matching *version* records. Provide the optional *toolid* argument to only match when both match for one entry. More information can be found here:

<http://www.ntcore.com/files/richsign.htm>

Note: Prior to version 3.11.0, this function returns only a boolean value (0 or 1) if the given *version* and optional *toolid* is present in an entry.

Example: *pe.rich_signature.version(24215, 261) == 61*

toolid (*toolid*, [*version*])

New in version 3.5.0.

Function returning a sum of count values of all matching *toolid* records. Provide the optional *version* argument to only match when both match for one entry. More information can be found here:

<http://www.ntcore.com/files/richsign.htm>

Note: Prior to version 3.11.0, this function returns only a boolean value (0 or 1) if the given *toolid* and optional *version* is present in an entry.

Example: *pe.rich_signature.toolid(170, 40219) >= 99*

pdb_path

New in version 4.0.0.

Path of the PDB file for this PE if present.

Example: *pe.pdb_path == "D:\workspace\2018_R9_RelBld\Target\checkout\custprof\Release\custprof.pdb"*

exports (function_name)

Function returning true if the PE exports *function_name* or false otherwise.

Example: *pe.exports("CPLApplet")*

exports (ordinal)

New in version 3.6.0.

Function returning true if the PE exports *ordinal* or false otherwise.

Example: *pe.exports(72)*

exports (/regular_expression/)

New in version 3.7.1.

Function returning true if the PE exports *regular_expression* or false otherwise.

Example: *pe.exports(/^AXS@ @/)*

exports_index (function_name)

New in version 4.0.0.

Function returning the index into the *export_details* array where the named function is, undefined otherwise.

Example: *pe.exports_index("CPLApplet")*

exports_index (ordinal)

New in version 4.0.0.

Function returning the index into the *export_details* array where the exported ordinal is, undefined otherwise.

Example: *pe.exports_index(72)*

exports_index (/regular_expression/)

New in version 4.0.0.

Function returning the first index into the *export_details* array where the regular expression matches the exported name, undefined otherwise.

Example: *pe.exports_index(/^ERS@ @/)*

number_of_exports

New in version 3.6.0.

Number of exports in the PE.

export_details

New in version 4.0.0.

Array of structures containing information about the PE's exports.

offset

Offset where the exported function starts.

name

Name of the exported function. It will be undefined if the function has no name.

forward_name

The name of the function where this export forwards to. It will be undefined if the export is not a forwarding export.

ordinal

The ordinal of the exported function, after the ordinal base has been applied to it.

dll_name

New in version 4.0.0.

The name of the DLL, if it exists in the export directory.

export_timestamp

New in version 4.0.0.

The timestamp the export data was created..

number_of_imports

New in version 3.6.0.

Number of imported DLLs in the PE.

number_of_imported_functions

New in version 4.1.0.

Number of imported functions in the PE.

number_of_delayed_imports

New in version 4.2.0.

Number of delayed imported DLLs in the PE. (Number of IMAGE_DELAYLOAD_DESCRIPTOR parsed from file)

number_of_delay_imported_functions

New in version 4.2.0.

Number of delayed imported functions in the PE.

imports (*dll_name*, function_name)

Function returning true if the PE imports *function_name* from *dll_name*, or false otherwise. *dll_name* is case insensitive.

Example: pe.imports("kernel32.dll", "WriteProcessMemory")

imports (*dll_name*)

New in version 3.5.0.

Changed in version 4.0.0.

Function returning the number of functions from the *dll_name*, in the PE imports. *dll_name* is case insensitive.

Note: Prior to version 4.0.0, this function returned only a boolean value indicating if the given DLL name was found in the PE imports. This change is backward compatible, as any number larger than 0 also evaluates as true.

Examples: pe.imports("kernel32.dll"), pe.imports("kernel32.dll") == 10

imports (*dll_name*, ordinal)

New in version 3.5.0.

Function returning true if the PE imports *ordinal* from *dll_name*, or false otherwise. *dll_name* is case insensitive.

Example: `pe.imports("WS2_32.DLL", 3)`

imports (*dll_regexp*, *function_regexp*)

New in version 3.8.0.

Changed in version 4.0.0.

Function returning the number of functions from the PE imports where a function name matches *function_regexp* and a DLL name matches *dll_regexp*. Both *dll_regexp* and *function_regexp* are case sensitive unless you use the *"i"* modifier in the regexp, as shown in the example below.

Note: Prior to version 4.0.0, this function returned only a boolean value indicating if matching import was found or not. This change is backward compatible, as any number larger than 0 also evaluates as true.

Example: `pe.imports(/kernel32.dll/i, /(Read\Write)ProcessMemory/) == 2`

imports (*import_flag*, *dll_name*, *function_name*)

New in version 4.2.0.

Function returning true if the PE imports *function_name* from *dll_name*, or false otherwise. *dll_name* is case insensitive.

import_flag is flag which specify type of import which should YARA search for. This value can be composed by bitwise OR these values:

pe.IMPORT_STANDARD
Search in standard imports

pe.IMPORT_DELAYED
Search in delayed imports

pe.IMPORT_ANY
Search in all imports

Example: `pe.imports(pe.IMPORT_DELAYED | pe.IMPORT_STANDARD, "kernel32.dll", "WriteProcessMemory")`

imports (*import_flag*, *import_flag*, *dll_name*)

New in version 4.2.0.

Function returning the number of functions from the *dll_name*, in the PE imports. *dll_name* is case insensitive.

Examples: `pe.imports(pe.IMPORT_DELAYED, "kernel32.dll")`, `pe.imports("kernel32.dll") == 10`

imports (*import_flag*, *dll_name*, *ordinal*)

New in version 4.2.0.

Function returning true if the PE imports *ordinal* from *dll_name*, or false otherwise. *dll_name* is case insensitive.

Example: `pe.imports(pe.IMPORT_DELAYED, "WS2_32.DLL", 3)`

imports (*import_flag*, *dll_regexp*, *function_regexp*)

New in version 4.2.0.

Function returning the number of functions from the PE imports where a function name matches *function_regexp* and a DLL name matches *dll_regexp*. Both *dll_regexp* and *function_regexp* are case sensitive unless you use the *"i"* modifier in the regexp, as shown in the example below.

Example: `pe.imports(pe.IMPORT_DELAYED, /kernel32.dll/i, /(Read\Write)ProcessMemory/) == 2`

import_details

New in version 4.2.0.

Array of structures containing information about the PE's imports libraries.

library_name

Library name.

number_of_functions

Number of imported function.

functions

Array of structures containing information about the PE's imports functions.

name

Name of imported function

ordinal

Ordinal of imported function. If ordinal does not exist this value is YR_UNDEFINED

rva

New in version 4.3.0.

Relative virtual address (RVA) of imported function. If rva not found then this value is YR_UNDEFINED

*Example: `pe.import_details[1].library_name == "library_name"`

delayed_import_details

New in version 4.2.0.

Array of structures containing information about the PE's delayed imports libraries.

library_name

Library name.

number_of_functions

Number of imported function.

functions

Array of structures containing information about the PE's imports functions.

name

Name of imported function

ordinal

Ordinal of imported function. If ordinal does not exist this value is YR_UNDEFINED

rva

New in version 4.3.0.

Relative virtual address (RVA) of imported function. If rva not found then this value is YR_UNDEFINED

*Example: `pe.delayed_import_details[1].name == "library_name"`

import_rva (dll, function)

New in version 4.3.0.

Function returning the RVA of an import that matches the DLL name and function name.

*Example: `pe.import_rva("PtImageRW.dll", "ord4") == 254924`

import_rva (dll, ordinal)

New in version 4.3.0.

Function returning the RVA of an import that matches the DLL name and ordinal number.

*Example: `pe.import_rva("PtPDF417Decode.dll", 4) == 254924`

delayed_import_rva (dll, function)

New in version 4.3.0.

Function returning the RVA of a delayed import that matches the DLL name and function name.

**Example: pe.delayed_import_rva("QDB.dll", "ord116") == 6110705*

delayed_import_rva (dll, ordinal)

New in version 4.3.0.

Function returning the RVA of a delayed import that matches the DLL name and ordinal number.

**Example: pe.delayed_import_rva("QDB.dll", 116) == 6110705*

locale (locale_identifier)

New in version 3.2.0.

Function returning true if the PE has a resource with the specified locale identifier. Locale identifiers are 16-bit integers and can be found here:

[http://msdn.microsoft.com/en-us/library/windows/desktop/dd318693\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/dd318693(v=vs.85).aspx)

Example: pe.locale(0x0419) // Russian (RU)

language (language_identifier)

New in version 3.2.0.

Function returning true if the PE has a resource with the specified language identifier. Language identifiers are 8-bit integers and can be found here:

[http://msdn.microsoft.com/en-us/library/windows/desktop/dd318693\(v=vs.85\).aspx](http://msdn.microsoft.com/en-us/library/windows/desktop/dd318693(v=vs.85).aspx)

Example: pe.language(0x0A) // Spanish

imphash ()

New in version 3.2.0.

Function returning the import hash or imphash for the PE. The imphash is an MD5 hash of the PE's import table after some normalization. The imphash for a PE can be also computed with [pefile](#) and you can find more information in [Mandiant's blog](#). The returned hash string is always in lowercase.

Example: pe.imphash() == "b8bb385806b89680e13fc0cf24f4431e"

section_index (name)

Function returning the index into the sections array for the section that has *name*. *name* is case sensitive.

Example: pe.section_index(".TEXT")

section_index (addr)

New in version 3.3.0.

Function returning the index into the sections array for the section that has *addr*. *addr* can be an offset into the file or a memory address.

Example: pe.section_index(pe.entry_point)

is_pe

New in version 3.8.0.

Return true if the file is a PE.

Example: pe.is_pe

is_dll ()

New in version 3.5.0.

Function returning true if the PE is a DLL.

Example: pe.is_dll()

is_32bit ()

New in version 3.5.0.

Function returning true if the PE is 32bits.

Example: pe.is_32bit()

is_64bit ()

New in version 3.5.0.

Function returning true if the PE is 64bits.

Example: pe.is_64bit()

rva_to_offset (addr)

New in version 3.6.0.

Function returning the file offset for RVA *addr*. Be careful to pass relative addresses here and not absolute addresses, like *pe.entry_point* when scanning a process.

Example: pe.rva_to_offset(pe.sections[0].virtual_address) == pe.sections[0].raw_data_offset

This example will make sure the offset for the virtual address in the first section equals the file offset for that section.

3.2 ELF module

New in version 3.2.0.

The ELF module is very similar to the *PE module*, but for ELF files. This module exposes most of the fields present in an ELF header. Let's see some examples:

```
import "elf"

rule single_section
{
    condition:
        elf.number_of_sections == 1
}

rule elf_64
{
    condition:
        elf.machine == elf.EM_X86_64
}
```

3.2.1 Reference

type

Integer with one of the following values:

ET_NONE

No file type.

ET_REL

Relocatable file.

ET_EXEC

Executable file.

ET_DYN

Shared object file.

ET_CORE

Core file.

Example: elf.type == elf.ET_EXEC

machine

Integer with one of the following values:

EM_NONE**EM_M32****EM_SPARC****EM_386****EM_68K****EM_88K****EM_860****EM_MIPS****EM_MIPS_RS3_LE****EM_PPC****EM_PPC64****EM_ARM****EM_X86_64****EM_AARCH64**

Example: elf.machine == elf.EM_X86_64

entry_point

Entry point raw offset or virtual address depending on whether YARA is scanning a file or process memory respectively. This is equivalent to the deprecated `entrypoint` keyword.

number_of_sections

Number of sections in the ELF file.

sections

A zero-based array of section objects, one for each section the ELF has. Individual sections can be accessed by using the `[]` operator. Each section object has the following attributes:

name

Section's name.

Example: elf.sections[3].name == ".bss"

size

Section's size in bytes. Unless the section type is `SHT_NOBITS`, the section occupies `sh_size` bytes in the file. A section of `SHT_NOBITS` may have a non-zero size, but it occupies no space in the file.

offset

Offset from the beginning of the file to the first byte in the section. One section type, `SHT_NOBITS`

described below, occupies no space in the file, and its *offset* member locates the conceptual placement in the file.

type

Integer with one of the following values:

SHT_NULL

This value marks the section as inactive; it does not have an associated section. Other members of the section header have undefined values.

SHT_PROGBITS

The section holds information defined by the program, whose format and meaning are determined solely by the program.

SHT_SYMTAB

The section holds a symbol table.

SHT_STRTAB

The section holds a string table. An object file may have multiple string table sections.

SHT_RELA

The section holds relocation entries.

SHT_HASH

The section holds a symbol hash table.

SHT_DYNAMIC

The section holds information for dynamic linking.

SHT_NOTE

The section holds information that marks the file in some way.

SHT_NOBITS

A section of this type occupies no space in the file but otherwise resembles *SHT_PROGBITS*.

SHT_REL

The section holds relocation entries.

SHT_SHLIB

This section type is reserved but has unspecified semantics.

SHT_DYNSYM

This section holds dynamic linking symbols.

flags

Integer with section's flags as defined below:

SHF_WRITE

The section contains data that should be writable during process execution.

SHF_ALLOC

The section occupies memory during process execution. Some control sections do not reside in the memory image of an object file; this attribute is off for those sections.

SHF_EXECINSTR

The section contains executable machine instructions.

Example: elf.sections[2].flags & elf.SHF_WRITE

address

New in version 3.6.0.

The virtual address the section starts at.

number_of_segments

New in version 3.4.0.

Number of segments in the ELF file.

segments

New in version 3.4.0.

A zero-based array of segment objects, one for each segment the ELF has. Individual segments can be accessed by using the [] operator. Each segment object has the following attributes:

alignment

Value to which the segments are aligned in memory and in the file.

file_size

Number of bytes in the file image of the segment. It may be zero.

flags

A combination of the following segment flags:

PF_R

The segment is readable.

PF_W

The segment is writable.

PF_X

The segment is executable.

memory_size

In-memory segment size.

offset

Offset from the beginning of the file where the segment resides.

physical_address

On systems for which physical addressing is relevant, contains the segment's physical address.

type

Type of segment indicated by one of the following values:

PT_NULL**PT_LOAD****PT_DYNAMIC****PT_INTERP****PT_NOTE****PT_SHLIB****PT_PHDR****PT_LOPROC****PT_HIPROC****PT_GNU_STACK****virtual_address**

Virtual address at which the segment resides in memory.

dynamic_section_entries

New in version 3.6.0.

Number of entries in the dynamic section in the ELF file.

dynamic

New in version 3.6.0.

A zero-based array of dynamic objects, one for each entry in found in the ELF's dynamic section. Individual dynamic objects can be accessed by using the [] operator. Each dynamic object has the following attributes:

type

Value that describes the type of dynamic section. Builtin values are:

DT_NULL

DT_NEEDED

DT_PLTRELSZ

DT_PLTGOT

DT_HASH

DT_STRTAB

DT_SYMTAB

DT_RELA

DT_RELASZ

DT_RELAENT

DT_STRSZ

DT_SYMENT

DT_INIT

DT_FINI

DT_SONAME

DT_RPATH

DT_SYMBOLIC

DT_REL

DT_RELSZ

DT_RELENT

DT_PLTREL

DT_DEBUG

DT_TEXTREL

DT_JMPREL

DT_BIND_NOW

DT_INIT_ARRAY

DT_FINI_ARRAY

DT_INIT_ARRAYSZ

DT_FINI_ARRAYSZ**DT_RUNPATH****DT_FLAGS****DT_ENCODING****value**

A value associated with the given type. The type of value (address, size, etc.) is dependant on the type of dynamic entry.

syntab_entries

New in version 3.6.0.

Number of entries in the symbol table found in the ELF file.

syntab

New in version 3.6.0.

A zero-based array of symbol objects, one for each entry in found in the ELF's SYMTAB. Individual symbol objects can be accessed by using the [] operator. Each symbol object has the following attributes:

name

The symbol's name.

value

A value associated with the symbol. Generally a virtual address.

size

The symbol's size.

type

The type of symbol. Built values are:

STT_NOTYPE**STT_OBJECT****STT_FUNC****STT_SECTION****STT_FILE****STT_COMMON****STT_TLS****bind**

The binding of the symbol. Builtin values are:

STB_LOCAL**STB_GLOBAL****STB_WEAK****shndx**

The section index which the symbol is associated with.

telhash()

Function returning Telfhash - TLSH hash of the ELF export and import symbols.

Example: elf.telhash() == "t166a00284751084526486df8b5df5b2fccb3f511dbc188c37156f5e714a11bc5d71014d"

import_md5()

Function returning Import Hash - MD5 hash of the ELF imported symbols.

Example: elf.import_md5() == "c3eca50cbb03400a6e91b9fe48da0c0c"

3.3 Cuckoo module

The Cuckoo module enables you to create YARA rules based on behavioral information generated by [Cuckoo sandbox](#). While scanning a PE file with YARA, you can pass additional information about its behavior to the `cuckoo` module and create rules based not only on what it *contains*, but also on what it *does*.

Important: This module is not built into YARA by default, to learn how to include it refer to [Compiling and installing YARA](#). Good news for Windows users: this module is already included in the official Windows binaries.

Suppose that you're interested in executable files sending a HTTP request to <http://someone.doingevil.com>. In previous versions of YARA you had to settle with:

```
rule evil_doer
{
    strings:
        $evil_domain = "http://someone.doingevil.com"

    condition:
        $evil_domain
}
```

The problem with this rule is that the domain name could be contained in the file for perfectly valid reasons not related with sending HTTP requests to <http://someone.doingevil.com>. Furthermore, the malicious executable could contain the domain name ciphered or obfuscated, in which case your rule would be completely useless.

But now with the `cuckoo` module you can take the behavior report generated for the executable file by your Cuckoo sandbox, pass it alongside the executable file to YARA, and write a rule like this:

```
import "cuckoo"

rule evil_doer
{
    condition:
        cuckoo.network.http_request(/http:\\\\someone\\.doingevil\\.com/)
}
```

Of course you can mix your behavior-related conditions with good old string-based conditions:

```
import "cuckoo"

rule evil_doer
{
    strings:
        $some_string = { 01 02 03 04 05 06 }

    condition:
        $some_string and
        cuckoo.network.http_request(/http:\\\\someone\\.doingevil\\.com/)
}
```

But how do we pass the behavior information to the `cuckoo` module? Well, in the case of the command-line tool you must use the `-x` option in this way:

```
$yara -x cuckoo=behavior_report_file rules_file pe_file
```

`behavior_report_file` is the path to a file containing the behavior file generated by the Cuckoo sandbox in JSON format.

If you are using `yara-python` then you must pass the behavior report in the `modules_data` argument for the `match` method:

```
import yara
rules = yara.compile('./rules_file')
report_file = open('./behavior_report_file')
report_data = report_file.read()
rules.match(pe_file, modules_data={'cuckoo': bytes(report_data)})
```

3.3.1 Reference

network

http_request (regexp)

Function returning true if the program sent a HTTP request to a URL matching the provided regular expression.

Example: cuckoo.network.http_request(/evi\com/)

http_get (regexp)

Similar to `http_request()`, but only takes into account GET requests.

http_post (regexp)

Similar to `http_request()`, but only takes into account POST requests.

http_user_agent (regexp)

Function returning true if the program sent a HTTP request with a user agent matching the provided regular expression.

Example: cuckoo.network.http_user_agent(/MSIE 6\0/)

dns_lookup (regexp)

Function returning true if the program sent a domain name resolution request for a domain matching the provided regular expression.

Example: cuckoo.network.dns_lookup(/evi\com/)

host (regexp)

Function returning true if the program contacted an IP address matching the provided regular expression.

Example: cuckoo.network.host(/192\168\1\1/)

tcp (regexp, port)

Function returning true if the program contacted an IP address matching the provided regular expression, over TCP on the provided port number.

Example: cuckoo.network.tcp(/192\168\1\1/, 443)

udp (regex, port)

Function returning true if the program contacted an IP address matching the provided regular expression, over UDP on the provided port number.

Example: cuckoo.network.udp(/192\168\1\1/, 53)

registry

key_access (regex)

Function returning true if the program accessed a registry entry matching the provided regular expression.

Example: cuckoo.registry.key_access(\\Software\\Microsoft\\Windows\\CurrentVersion\\Run/)

filesystem

file_access (regex)

Function returning true if the program accessed a file matching the provided regular expression.

Example: cuckoo.filesystem.file_access(/autoexec\bat/)

sync

mutex (regex)

Function returning true if the program opens or creates a mutex matching the provided regular expression.

Example: cuckoo.sync.mutex(/EvilMutexName/)

3.4 Magic module

New in version 3.1.0.

The Magic module allows you to identify the type of the file based on the output of `file`, the standard Unix command.

Important: This module is not built into YARA by default, to learn how to include it refer to [Compiling and installing YARA](#). Bad news for Windows users: **this module is not supported on Windows**.

There are two functions in this module: `type()` and `mime_type()`. The first one returns the descriptive string returned by `file`, for example, if you run `file` against some PDF document you'll get something like this:

```
$file some.pdf
some.pdf: PDF document, version 1.5
```

The `type()` function would return `"PDF document, version 1.5"` in this case. Using the `mime_type()` function is similar to passing the `--mime` argument to `file`:

```
$file --mime some.pdf
some.pdf: application/pdf; charset=binary
```

`mime_type()` would return `"application/pdf"`, without the charset part.

By experimenting a little with the `file` command you can learn which output to expect for different file types. These are a few examples:

- JPEG image data, JFIF standard 1.01
- PE32 executable for MS Windows (GUI) Intel 80386 32-bit

- PNG image data, 1240 x 1753, 8-bit/color RGBA, non-interlaced
- ASCII text, with no line terminators
- Zip archive data, at least v2.0 to extract

libmagic will try and read its compiled file type database from `/etc/magic.mgc` by default. If this file doesn't exist, you can set the environment variable `MAGIC` to point to a `magic.mgc` file and libmagic will attempt to load from there as an alternative.

type()

Function returning a string with the type of the file.

Example: `magic.type()` contains "PDF"

mime_type()

Function returning a string with the MIME type of the file.

Example: `magic.mime_type() == "application/pdf"`

3.5 Hash module

New in version 3.2.0.

The Hash module allows you to calculate hashes (MD5, SHA1, SHA256) from portions of your file and create signatures based on those hashes.

Important: This module depends on the OpenSSL library. Please refer to [Compiling and installing YARA](#) for information about how to build OpenSSL-dependant features into YARA.

Good news for Windows users: this module is already included in the official Windows binaries.

Warning: The returned hash string is always in lowercase. This means that rule condition matching on hashes `hash.md5(0, filesize) == "feba6c919e3797e7778e8f2e85fa033d"` requires the hash string to be given in lowercase, otherwise the match condition will not work. (see <https://github.com/VirusTotal/yara/issues/1004>)

md5(offset, size)

Returns the MD5 hash for *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned string is always in lowercase.

Example: `hash.md5(0, filesize) == "feba6c919e3797e7778e8f2e85fa033d"`

md5(string)

Returns the MD5 hash for the given string.

Example: `hash.md5("dummy") == "275876e34cf609db118f3d84b799a790"`

sha1(offset, size)

Returns the SHA1 hash for the *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned string is always in lowercase.

sha1(string)

Returns the SHA1 hash for the given string.

sha256 (*offset*, *size*)

Returns the SHA256 hash for the *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned string is always in lowercase.

sha256 (string)

Returns the SHA256 hash for the given string.

checksum32 (*offset*, *size*)

Returns a 32-bit checksum for the *size* bytes starting at *offset*. The checksum is just the sum of all the bytes (unsigned).

checksum32 (string)

Returns a 32-bit checksum for the given string. The checksum is just the sum of all the bytes in the string (unsigned).

crc32 (*offset*, *size*)

Returns a crc32 checksum for the *size* bytes starting at *offset*.

crc32 (string)

Returns a crc32 checksum for the given string.

3.6 Math module

New in version 3.3.0.

The Math module allows you to calculate certain values from portions of your file and create signatures based on those results.

Important: Where noted these functions return floating point numbers. YARA is able to convert integers to floating point numbers during most operations. For example this will convert 7 to 7.0 automatically, because the return type of the entropy function is a floating point value:

```
math.entropy(0, filesize) >= 7
```

The one exception to this is when a function requires a floating point number as an argument. For example, this will cause a syntax error because the arguments must be floating point numbers:

```
math.in_range(2, 1, 3)
```

entropy (*offset*, *size*)

Returns the entropy for *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float.

Example: `math.entropy(0, filesize) >= 7`

entropy (string)

Returns the entropy for the given string.

Example: `math.entropy("dummy") > 7`

monte_carlo_pi (*offset*, *size*)

Returns the percentage away from Pi for the *size* bytes starting at *offset* when run through the Monte Carlo from Pi test. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float.

Example: `math.monte_carlo_pi(0, filesize) < 0.07`

monte_carlo_pi (string)

Return the percentage away from Pi for the given string.

serial_correlation (*offset*, *size*)

Returns the serial correlation for the *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float between 0.0 and 1.0.

Example: *math.serial_correlation(0, filesize) < 0.2*

serial_correlation (string)

Return the serial correlation for the given string.

mean (*offset*, *size*)

Returns the mean for the *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float.

Example: *math.mean(0, filesize) < 72.0*

mean (string)

Return the mean for the given string.

deviation (*offset*, *size*, *mean*)

Returns the deviation from the mean for the *size* bytes starting at *offset*. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float.

The mean of an equally distributed random sample of bytes is 127.5, which is available as the constant `math.MEAN_BYTES`.

Example: *math.deviation(0, filesize, math.MEAN_BYTES) == 64.0*

deviation (string, *mean*)

Return the deviation from the mean for the given string.

in_range (test, lower, upper)

Returns true if the *test* value is between *lower* and *upper* values. The comparisons are inclusive.

Example: *math.in_range(math.deviation(0, filesize, math.MEAN_BYTES), 63.9, 64, 1)*

max (int, int)

New in version 3.8.0.

Returns the maximum of two unsigned integer values.

min (int, int)

New in version 3.8.0.

Returns the minimum of two unsigned integer values.

to_number (bool)

New in version 4.1.0.

Returns 0 or 1, it's useful when writing a score based rule.

Example: *math.to_number(SubRule1) * 60 + math.to_number(SubRule2) * 20 + math.to_number(SubRule3) * 70 > 80*

abs (int)

New in version 4.2.0.

Returns the absolute value of the signed integer.

Example: *math.abs(@a - @b) == 1*

count (byte, *offset*, *size*)

New in version 4.2.0.

Returns how often a specific byte occurs, starting at *offset* and looking at the next *size* bytes. When scanning a running process the *offset* argument should be a virtual address within the process address space. *offset* and *size* are optional; if left empty, the complete file is searched.

Example: math.count(0x4A) >= 10

percentage (byte, *offset*, *size*)

New in version 4.2.0.

Returns the occurrence rate of a specific byte, starting at *offset* and looking at the next *size* bytes. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float between 0 and 1. *offset* and *size* are optional; if left empty, the complete file is searched.

Example: math.percentage(0xFF, filesize-1024, filesize) >= 0.9

Example: math.percentage(0x4A) >= 0.4

mode (*offset*, *size*)

New in version 4.2.0.

Returns the most common byte, starting at *offset* and looking at the next *size* bytes. When scanning a running process the *offset* argument should be a virtual address within the process address space. The returned value is a float. *offset* and *size* are optional; if left empty, the complete file is searched.

Example: math.mode(0, filesize) == 0xFF

to_string (int)

New in version 4.3.0.

Convert the given integer to a string. Note: integers in YARA are signed.

Example: math.to_string(10) == "10" Example: math.to_string(-1) == "-1"

to_string (int, *base*)

New in version 4.3.0.

Convert the given integer to a string in the given base. Supported bases are 10, 8 and 16. Note: integers in YARA are signed.

Example: math.to_string(32, 16) == "20" Example: math.to_string(-1, 16) == "ffffffff"

3.7 dotnet module

New in version 3.6.0.

The dotnet module allows you to create more fine-grained rules for .NET files by using attributes and features of the .NET file format. Let's see some examples:

```
import "dotnet"

rule not_exactly_five_streams
{
    condition:
        dotnet.number_of_streams != 5
}

rule blob_stream
{
    condition:
        for any i in (0..dotnet.number_of_streams - 1):
```

(continues on next page)

(continued from previous page)

```

    (dotnet.streams[i].name == "#Blop")
}

```

3.7.1 Reference

version

The version string contained in the metadata root.

Example: dotnet.version == "v2.0.50727"

module_name

The name of the module.

Example: dotnet.module_name == "axs"

number_of_streams

The number of streams in the file.

streams

A zero-based array of stream objects, one for each stream contained in the file. Individual streams can be accessed by using the [] operator. Each stream object has the following attributes:

name

Stream name.

offset

Stream offset.

size

Stream size.

Example: dotnet.streams[0].name == "#~"

number_of_guids

The number of GUIDs in the guids array.

guids

A zero-based array of strings, one for each GUID. Individual guids can be accessed by using the [] operator.

Example: dotnet.guids[0] == "99c08ffd-f378-a891-10ab-c02fe11be6ef"

number_of_classes

The number of *classes* in the **file**.

classes

An array of .NET *classes* stored in the metadata. Individual *classes* can be **accessed** using the [] operator. Each class object contains the following **attributes**:

fullname

Class full **name**.

name

Class **name**.

namespace

Class **namespace**.

visibility

Class visibility specifier, options **are**:

private public protected internal private protected protected internal

type

Type of the object, options **are**:

class interface

abstract

Boolean representing if class is **abstract**.

sealed

Boolean representing if class is **sealed**.

number_of_generic_parameters

Number of generic **parameters**.

generic_parameters

A zero-based array of generic parameters name. Individual parameters can be *accessed* using the [] **operator**.

number_of_base_types

Number of the *base* **types**.

base_types

A zero-based array of *base* types name. Individual *base* types can be *accessed* using the [] **operator**.

number_of_methods

Number of the **methods**.

methods

A zero-based array of method objects. Individual methods can be *accessed* by using the [] operator. Each object contains following **attributes**:

name

Method **name**.

visibility

Method visibility specifier, options **are**:

private public protected internal private protected protected internal

static

Boolean representing if method is **static**.

virtual

Boolean representing if method is **virtual**.

final

Boolean representing if method is **final**.

abstract

Boolean representing if method is **abstract**.

return_type

Method return *type* **name**.

number_of_parameters

Number of the method **parameters**.

parameters

A zero-based array of method parameters. Individual parameters can be *accessed* by using the [] **operator**.

name

Parameter **name**.

type

Parameter **type**.

number_of_generic_parameters

Number of the method generic **parameters**.

generic_parameters

A zero-based array of method generic parameters. Individual parameters can be *accessed* by using the [] **operator**.

Example: dotnet.classes[0].fullname == "Launcher.Program"

number_of_resources

The number of resources in the .NET file. These are different from normal PE resources.

resources

A zero-based array of resource objects, one for each resource the .NET file has. Individual resources can be accessed by using the [] operator. Each resource object has the following attributes:

offset

Offset for the resource data.

length

Length of the resource data.

name

Name of the resource (string).

Example: uint16be(dotnet.resources[0].offset) == 0x4d5a

assembly

Object for .NET assembly information.

version

An object with integer values representing version information for this assembly. Attributes are:

major minor build_number revision_number

name

String containing the assembly name.

culture

String containing the culture (language/country/region) for this assembly.

Example: dotnet.assembly.name == "Keylogger"

Example: dotnet.assembly.version.major == 7 and dotnet.assembly.version.minor == 0

number_of_modulerefs

The number of module references in the .NET file.

modulerefs

A zero-based array of strings, one for each module reference the .NET file has. Individual module references can be accessed by using the [] operator.

Example: dotnet.modulerefs[0] == "kernel32"

typelib

The typelib of the file.

assembly_refs

Object for .NET assembly reference information.

version

An object with integer values representing version information for this assembly. Attributes are:

major minor build_number revision_number

name

String containing the assembly name.

public_key_or_token

String containing the public key or token which identifies the author of this assembly.

number_of_user_strings

The number of user strings in the file.

user_strings

An zero-based array of user strings, one for each stream contained in the file. Individual strings can be accessed by using the [] operator.

number_of_field_offsets

The number of fields in the field_offsets array.

field_offsets

A zero-based array of integers, one for each field. Individual field offsets can be accessed by using the [] operator.

Example: dotnet.field_offsets[0] == 8675309

is_dotnet

New in version 4.2.0.

Function returning true if the PE is indeed .NET.

Example: dotnet.is_dotnet

3.8 Time module

New in version 3.7.0.

The Time module allows you to use temporal conditions in your YARA rules.

now()

Function returning an integer which is the number of seconds since January 1, 1970.

Example: pe.timestamp > time.now()

3.9 Console module

New in version 4.2.0.

The Console module allows you to log information during condition execution. By default, the log messages are sent to stdout but can be handled differently by using the C api ([Scanning data](#)).

Every function in the console module returns true for the purposes of condition evaluation. This means you must logically and your statements together to get the proper output. For example:

```
import "console"

rule example
```

(continues on next page)

(continued from previous page)

```

{
    condition:
        console.log("Hello") and console.log("World!")
}

```

log (string)

Function which sends the string to the main callback.

Example: console.log(pe.imphash())

log (message, string)

Function which sends the message and string to the main callback.

Example: console.log("The imphash is: ", pe.imphash())

log (integer)

Function which sends the integer to the main callback.

Example: console.log(uint32(0))

log (message, integer)

Function which sends the message and integer to the main callback.

Example: console.log("32bits at 0: ", uint32(0))

log (float)

Function which sends the floating point value to the main callback.

Example: console.log(math.entropy(0, filesize))

log (message, float)

Function which sends the message and the floating point value to the main callback.

Example: console.log("Entropy: ", math.entropy(0, filesize))

hex (integer)

Function which sends the integer to the main callback, formatted as a hex string.

Example: console.hex(uint32(0))

hex (message, integer)

Function which sends the integer to the main callback, formatted as a hex string.

Example: console.hex("Hex at 0: ", uint32(0))

3.10 String module

New in version 4.3.0.

The String module provides functions for manipulating strings as returned by modules. The strings referenced here are not YARA strings as defined in the strings section of your rule.

to_int (string)

New in version 4.3.0.

Convert the given string to a signed integer. If the string starts with "0x" it is treated as base 16. If the string starts with "0" it is treated base 8. Leading '+' or '-' is also supported.

Example: string.to_int("1234") == 1234 Example: string.to_int("-10") == -10 Example: string.to_int("-010") == -8

to_int (string, *base*)

New in version 4.3.0.

Convert the given string, interpreted with the given base, to a signed integer. Base must be 0 or between 2 and 36 inclusive. If it is zero then the string will be interpreted as base 16 if it starts with "0x" or as base 8 if it starts with "0". Leading '+' or '-' is also supported.

Example: string.to_int("011", 8) == "9" Example: string.to_int("-011", 0) == "-9"

length (string)

New in version 4.3.0.

Return the length of the string, which can be any sequence of bytes. NULL bytes included.

Example: string.length("AXSx00ERS") == 7

Writing your own modules

For the first time ever, in YARA 3.0 you can extend its features to express more complex and refined conditions. YARA 3.0 does this by employing modules, which you can use to define data structures and functions, which can be later used from within your rules. You can see some examples of what a module can do in the [Using modules](#) section.

The purpose of the following sections is to teach you how to create your own modules for giving YARA that cool feature you always dreamed of.

4.1 The "Hello World!" module

Modules are written in C and built into YARA as part of the compiling process. In order to create your own modules you must be familiar with the C programming language and how to configure and build YARA from source code. You don't need to understand how YARA does its magic; YARA exposes a simple API for modules, which is all you need to know.

The source code for your module must reside in the *libyara/modules* directory of the source tree. It's recommended to use the module name as the file name for the source file, if your module's name is *foo* its source file should be *foo.c*.

In the *libyara/modules* directory you'll find a *demo.c* file we'll use as our starting point. The file looks like this:

```
#include <yara/modules.h>

#define MODULE_NAME demo

begin_declarations;

    declare_string("greeting");

end_declarations;

int module_initialize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}
```

(continues on next page)

(continued from previous page)

```

}

int module_finalize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}

int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)
{
    set_string("Hello World!", module_object, "greeting");
    return ERROR_SUCCESS;
}

int module_unload(
    YR_OBJECT* module_object)
{
    return ERROR_SUCCESS;
}

#undef MODULE_NAME

```

Let's start dissecting the source code so you can understand every detail. The first line in the code is:

```
#include <yara/modules.h>
```

The *modules.h* header file is where the definitions for YARA's module API reside, therefore this include directive is required in all your modules. The second line is:

```
#define MODULE_NAME demo
```

This is how you define the name of your module and is also required. Every module must define its name at the start of the source code. Module names must be unique among the modules built into YARA.

Then follows the declaration section:

```
begin_declarations;

    declare_string("greeting");

end_declarations;
```

Here is where the module declares the functions and data structures that will be available for your YARA rules. In this case we are declaring just a string variable named *greeting*. We are going to discuss these concepts in greater detail in the [The declaration section](#).

After the declaration section you'll find a pair of functions:

```
int module_initialize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}
```

(continues on next page)

(continued from previous page)

```
int module_finalize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}
```

The `module_initialize` function is called during YARA's initialization while its counterpart `module_finalize` is called while finalizing YARA. These functions allow you to initialize and finalize any global data structure you may need to use in your module.

Then comes the `module_load` function:

```
int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)
{
    set_string("Hello World!", module_object, "greeting");
    return ERROR_SUCCESS;
}
```

This function is invoked once for each scanned file, but only if the module is imported by some rule with the `import` directive. The `module_load` function is where your module has the opportunity to inspect the file being scanned, parse or analyze it in the way preferred, and then populate the data structures defined in the declarations section.

In this example the `module_load` function doesn't inspect the file content at all, it just assigns the string, "Hello World!" to the variable `greeting` declared before.

And finally, we have the `module_unload` function:

```
int module_unload(
    YR_OBJECT* module_object)
{
    return ERROR_SUCCESS;
}
```

For each call to `module_load` there is a corresponding call to `module_unload`. This function allows your module to free any resource allocated during `module_load`. There's nothing to free in this case, so the function just returns `ERROR_SUCCESS`. Both `module_load` and `module_unload` should return `ERROR_SUCCESS` to indicate that everything went fine. If a different value is returned the scanning will be aborted and an error reported to the user.

4.1.1 Building our "Hello World!"

Modules are not magically built into YARA just by dropping their source code into the `libyara/modules` directory, you must follow two further steps in order to get them to work. The first step is adding your module to the `module_list` file also found in the `libyara/modules` directory.

The `module_list` file looks like this:

```
MODULE(tests)
MODULE(pe)

#ifdef CUCKOO_MODULE
```

(continues on next page)

(continued from previous page)

```
MODULE (cuckoo)
#endif
```

You must add a line `MODULE(<name>)` with the name of your module to this file. In our case the resulting `module_list` is:

```
MODULE (tests)
MODULE (pe)

#ifdef CUCKOO_MODULE
MODULE (cuckoo)
#endif

MODULE (demo)
```

The second step is modifying the `Makefile.am` to tell the `make` program that the source code for your module must be compiled and linked into YARA. At the very beginning of `libyara/Makefile.am` you'll find this:

```
MODULES = modules/tests/tests.c
MODULES += modules/pe/pe.c

if CUCKOO_MODULE
MODULES += modules/cuckoo/cuckoo.c
endif
```

Just add a new line for your module:

```
MODULES = modules/tests/tests.c
MODULES += modules/pe/pe.c

if CUCKOO_MODULE
MODULES += modules/cuckoo/cuckoo.c
endif

MODULES += modules/demo/demo.c
```

And that's all! Now you're ready to build YARA with your brand-new module included. Just go to the source tree root directory and type as always:

```
./bootstrap.sh
./configure
make
sudo make install
```

Now you should be able to create a rule like this:

```
import "demo"

rule HelloWorld
{
    condition:
        demo.greeting == "Hello World!"
}
```

Any file scanned with this rule will match the `HelloWord` because `demo.greeting == "Hello World!"` is always true.

4.2 The declaration section

The declaration section is where you declare the variables, structures and functions that will be available for your YARA rules. Every module must contain a declaration section like this:

```
begin_declarations;

    <your declarations here>

end_declarations;
```

4.2.1 Basic types

Within the declaration section you can use `declare_string(<variable name>)`, `declare_integer(<variable name>)` and `declare_float(<variable name>)` to declare string, integer, or float variables respectively. For example:

```
begin_declarations;

    declare_integer("foo");
    declare_string("bar");
    declare_float("baz");

end_declarations;
```

Note: Floating-point variables require YARA version 3.3.0 or later.

Variable names can't contain characters other than letters, numbers and underscores. These variables can be used later in your rules at any place where an integer or string is expected. Supposing your module name is "mymodule", they can be used like this:

```
mymodule.foo > 5

mymodule.bar matches /someregexp/
```

4.2.2 Structures

Your declarations can be organized in a more structured way:

```
begin_declarations;

    declare_integer("foo");
    declare_string("bar");
    declare_float("baz");

    begin_struct("some_structure");

        declare_integer("foo");

        begin_struct("nested_structure");
```

(continues on next page)

(continued from previous page)

```
        declare_integer("bar");

    end_struct("nested_structure");

end_struct("some_structure");

begin_struct("another_structure");

    declare_integer("foo");
    declare_string("bar");
    declare_string("baz");
    declare_float("tux");

end_struct("another_structure");

end_declarations;
```

In this example we're using `begin_struct(<structure name>)` and `end_struct(<structure name>)` to delimit two structures named *some_structure* and *another_structure*. Within the structure delimiters you can put any other declarations you want, including another structure declaration. Also notice that members of different structures can have the same name, but members within the same structure must have unique names.

When referring to these variables from your rules it would be like this:

```
mymodule.foo
mymodule.some_structure.foo
mymodule.some_structure.nested_structure.bar
mymodule.another_structure.baz
```

4.2.3 Arrays

In the same way you declare individual strings, integers, floats or structures, you can declare arrays of them:

```
begin_declarations;

    declare_integer_array("foo");
    declare_string_array("bar");
    declare_float_array("baz");

    begin_struct_array("struct_array");

        declare_integer("foo");
        declare_string("bar");

    end_struct_array("struct_array");

end_declarations;
```

Individual values in the array are referenced like in most programming languages:

```
foo[0]
bar[1]
baz[3]
struct_array[4].foo
struct_array[1].bar
```

Arrays are zero-based and don't have a fixed size, they will grow as needed when you start initializing its values.

4.2.4 Dictionaries

New in version 3.2.0.

You can also declare dictionaries of integers, floats, strings, or structures:

```
begin_declarations;

    declare_integer_dictionary("foo");
    declare_string_dictionary("bar");
    declare_float_dictionary("baz")

    begin_struct_dictionary("struct_dict");

        declare_integer("foo");
        declare_string("bar");

    end_struct_dictionary("struct_dict");

end_declarations;
```

Individual values in the dictionary are accessed by using a string key:

```
foo["somekey"]
bar["anotherkey"]
baz["yetanotherkey"]
struct_dict["k1"].foo
struct_dict["k1"].bar
```

4.2.5 Functions

One of the more powerful features of YARA modules is the possibility of declaring functions that can be later invoked from your rules. Functions must appear in the declaration section in this way:

```
declare_function(<function name>, <argument types>, <return tuipe>, <C function>);
```

<function name> is the name that will be used in your YARA rules to invoke the function.

<argument types> is a string containing one character per function argument, where the character indicates the type of the argument. Functions can receive four different types of arguments: string, integer, float and regular expression, denoted by characters: **s**, **i**, **f** and **r** respectively. If your function receives two integers *<argument types>* must be "ii", if it receives an integer as the first argument and a string as the second one *<argument types>* must be "is", if it receives three strings and a float *<argument types>* must be "sssf".

<return type> is a string with a single character indicating the return type. Possible return types are string ("s") integer ("i") and float ("f").

<C function> is the identifier for the actual implementation of your function.

Here you have a full example:

```
define_function(isum)
{
    int64_t a = integer_argument(1);
```

(continues on next page)

(continued from previous page)

```

    int64_t b = integer_argument(2);

    return_integer(a + b);
}

define_function(fsum)
{
    double a = float_argument(1);
    double b = float_argument(2);

    return_integer(a + b);
}

begin_declarations;

    declare_function("sum", "ii", "i", sum);

end_declarations;

```

As you can see in the example above, your function code must be defined before the declaration section, like this:

```

define_function(<function identifier>)
{
    ...your code here
}

```

Functions can be overloaded as in C++ and other programming languages. You can declare two functions with the same name as long as they differ in the type or number of arguments. One example of overloaded functions can be found in the [Hash module](#), it has two functions for calculating MD5 hashes, one receiving an offset and length within the file and another one receiving a string:

```

begin_declarations;

    declare_function("md5", "ii", "s", data_md5);
    declare_function("md5", "s", "s", string_md5);

end_declarations;

```

We are going to discuss function implementation more in depth in the [More about functions](#) section.

4.3 Initialization and finalization

Every module must implement two functions for initialization and finalization: `module_initialize` and `module_finalize`. The former is called during YARA's initialization by `yr_initialize()` while the latter is called during finalization by `yr_finalize()`. Both functions are invoked whether or not the module is being imported by some rule.

These functions give your module an opportunity to initialize any global data structure it may need, but most of the time they are just empty functions:

```

int module_initialize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}

```

(continues on next page)

(continued from previous page)

```

}

int module_finalize(
    YR_MODULE* module)
{
    return ERROR_SUCCESS;
}

```

Any returned value different from `ERROR_SUCCESS` will abort YARA's execution.

4.4 Implementing the module's logic

Besides `module_initialize` and `module_finalize` every module must implement two other functions which are called by YARA during the scanning of a file or process memory space: `module_load` and `module_unload`. Both functions are called once for each scanned file or process, but only if the module was imported by means of the `import` directive. If the module is not imported by some rule neither `module_load` nor `module_unload` will be called.

The `module_load` function has the following prototype:

```

int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)

```

The `context` argument contains information relative to the current scan, including the data being scanned. The `module_object` argument is a pointer to a `YR_OBJECT` structure associated with the module. Each structure, variable or function declared in a YARA module is represented by a `YR_OBJECT` structure. These structures form a tree whose root is the module's `YR_OBJECT` structure. If you have the following declarations in a module named *mymodule*:

```

begin_declarations;

    declare_integer("foo");

    begin_struct("bar");

        declare_string("baz");

    end_struct("bar");

end_declarations;

```

Then the tree will look like this:

```

YR_OBJECT(type=OBJECT_TYPE_STRUCT, name="mymodule")
|
|_ YR_OBJECT(type=OBJECT_TYPE_INTEGER, name="foo")
|
|_ YR_OBJECT(type=OBJECT_TYPE_STRUCT, name="bar")
    |
    |_ YR_OBJECT(type=OBJECT_TYPE_STRING, name="baz")

```

Notice that both *bar* and *mymodule* are of the same type `OBJECT_TYPE_STRUCT`, which means that the `YR_OBJECT` associated with the module is just another structure like *bar*. In fact, when you write in your rules something like `mymodule.foo` you're performing a field lookup in a structure in the same way that `bar.baz` does.

In summary, the `module_object` argument allows you to access every variable, structure or function declared by the module by providing a pointer to the root of the objects tree.

The `module_data` argument is a pointer to any additional data passed to the module, and `module_data_size` is the size of that data. Not all modules require additional data, most of them rely on the data being scanned alone, but a few of them require more information as input. The *Cuckoo module* is a good example of this, it receives a behavior report associated with PE files being scanned which is passed in the `module_data` and `module_data_size` arguments.

For more information on how to pass additional data to your module take a look at the `-x` argument in *Running YARA from the command-line*.

4.4.1 Accessing the scanned data

Most YARA modules need to access the file or process memory being scanned to extract information from it. The data being scanned is sent to the module in the `YR_SCAN_CONTEXT` structure passed to the `module_load` function. The data is sometimes sliced in blocks, therefore your module needs to iterate over the blocks by using the `foreach_memory_block` macro:

```
int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)
{
    YR_MEMORY_BLOCK* block;

    foreach_memory_block(context, block)
    {
        ..do something with the current memory block
    }
}
```

Each memory block is represented by a `YR_MEMORY_BLOCK` structure with the following attributes:

`YR_MEMORY_BLOCK_FETCH_DATA_FUNC` **fetch_data**

Pointer to a function returning a pointer to the block's data.

`size_t` **size**

Size of the data block.

`size_t` **base**

Base offset/address for this block. If a file is being scanned this field contains the offset within the file where the block begins, if a process memory space is being scanned this contains the virtual address where the block begins.

The blocks are always iterated in the same order as they appear in the file or process memory. In the case of files the first block will contain the beginning of the file. Actually, a single block will contain the whole file's content in most cases, but you can't rely on that while writing your code. For very big files YARA could eventually split the file into two or more blocks, and your module should be prepared to handle that.

The story is very different for processes. While scanning a process memory space your module will definitely receive a large number of blocks, one for each committed memory region in the process address space.

However, there are some cases where you don't actually need to iterate over the blocks. If your module just parses the header of some file format you can safely assume that the whole header is contained within the first block (put some checks in your code nevertheless). In those cases you can use the `first_memory_block` macro:

```
int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)
{
    YR_MEMORY_BLOCK* block;
    const uint8_t* block_data;

    block = first_memory_block(context);
    block_data = block->fetch_data(block)

    if (block_data != NULL)
    {
        ..do something with the memory block
    }
}
```

In the previous example you can also see how to use the `fetch_data` function. This function, which is a member of the `YR_MEMORY_BLOCK` structure, receives a pointer to the same block (as a `self` or `this` pointer) and returns a pointer to the block's data. Your module doesn't own the memory pointed to by this pointer, freeing that memory is not your responsibility. However keep in mind that the pointer is valid only until you ask for the next memory block. As long as you use the pointer within the scope of a `foreach_memory_block` you are on the safe side. Also take into account that `fetch_data` can return a `NULL` pointer, your code must be prepared for that case.

```
const uint8_t* block_data;

foreach_memory_block(context, block)
{
    block_data = block->fetch_data(block);

    if (block_data != NULL)
    {
        // using block_data is safe here.
    }
}

// the memory pointed to by block_data can be already freed here.
```

4.4.2 Setting variable's values

The `module_load` function is where you assign values to the variables declared in the declarations section, once you've parsed or analyzed the scanned data and/or any additional module's data. This is done by using the `set_float`, `set_integer`, and `set_string` functions:

`void set_float (double value, YR_OBJECT* object, const char* field, ...)`

`void set_integer (int64_t value, YR_OBJECT* object, const char* field, ...)`

`void set_string (const char* value, YR_OBJECT* object, const char* field, ...)`

These functions receive a value to be assigned to the variable, a pointer to a `YR_OBJECT` representing the variable itself or some ancestor of that variable, a field descriptor, and additional arguments as defined by the field descriptor.

If we are assigning the value to the variable represented by `object` itself, then the field descriptor must be `NULL`. For example, assuming that `object` points to a `YR_OBJECT` structure corresponding to some integer variable, we can set the value for that integer variable with:

```
set_integer(<value>, object, NULL);
```

The field descriptor is used when you want to assign the value to some descendant of `object`. For example, consider the following declarations:

```
begin_declarations;

    begin_struct("foo");

        declare_string("bar");

        begin_struct("baz");

            declare_integer("qux");

        end_struct("baz");

    end_struct("foo");

end_declarations;
```

If `object` points to the `YR_OBJECT` associated with the `foo` structure you can set the value for the `bar` string like this:

```
set_string(<value>, object, "bar");
```

And the value for `qux` like this:

```
set_integer(<value>, object, "baz.qux");
```

Do you remember that the `module_object` argument for `module_load` was a pointer to a `YR_OBJECT`? Do you remember that this `YR_OBJECT` is a structure just like `bar` is? Well, you could also set the values for `bar` and `qux` like this:

```
set_string(<value>, module_object, "foo.bar");
set_integer(<value>, module_object, "foo.baz.qux");
```

But what happens with arrays? How can I set the value for array items? If you have the following declarations:

```
begin_declarations;

    declare_integer_array("foo");

    begin_struct_array("bar")

        declare_string("baz");
        declare_integer_array("qux");

    end_struct_array("bar");

end_declarations;
```

Then the following statements are all valid:

```
set_integer(<value>, module, "foo[0]");
set_integer(<value>, module, "foo[%i]", 2);
set_string(<value>, module, "bar[%i].baz", 5);
set_string(<value>, module, "bar[0].qux[0]");
set_string(<value>, module, "bar[0].qux[%i]", 0);
set_string(<value>, module, "bar[%i].qux[%i]", 100, 200);
```

Those %i in the field descriptor are replaced by the additional integer arguments passed to the function. This works in the same way as `printf` in C programs, but the only format specifiers accepted are %i and %s, for integer and string arguments respectively.

The %s format specifier is used for assigning values to a certain key in a dictionary:

```
set_integer(<value>, module, "foo[\"key\"]");
set_integer(<value>, module, "foo[%s]", "key");
set_string(<value>, module, "bar[%s].baz", "another_key");
```

If you don't explicitly assign a value to a declared variable, array or dictionary item it will remain in an undefined state. That's not a problem at all, and is even useful in many cases. For example, if your module parses files from a certain format and it receives one from a different format, you can safely leave all your variables undefined instead of assigning them bogus values that don't make sense. YARA will handle undefined values in rule conditions as described in *Using modules*.

In addition to the `set_float`, `set_integer`, and `set_string` functions, you have their `get_float`, `get_integer`, and `get_string` counterparts. As the names suggest, they are used for getting the value of a variable, which can be useful in the implementation of your functions to retrieve values previously stored by `module_load`.

double **get_float** (YR_OBJECT* *object*, const char* *field*, ...)

int64_t **get_integer** (YR_OBJECT* *object*, const char* *field*, ...)

SIZED_STRING* **get_string** (YR_OBJECT* *object*, const char* *field*, ...)

There's also a function to get any YR_OBJECT in the objects tree:

YR_OBJECT* **get_object** (YR_OBJECT* *object*, const char* *field*, ...)

Here is a little exam...

Are the following two lines equivalent? Why?

```
set_integer(1, get_object(module_object, "foo.bar"), NULL);
set_integer(1, module_object, "foo.bar");
```

4.4.3 Storing data for later use

Sometimes the information stored directly in your variables by means of `set_integer` and `set_string` is not enough. You may need to store more complex data structures or information that doesn't need to be exposed to YARA rules.

Storing information is essential when your module exports functions to be used in YARA rules. The implementation of these functions usually require to access information generated by `module_load` which must kept somewhere. You may be tempted to define global variables to store the required information, but this would make your code non-thread-safe. The correct approach is using the data field of the YR_OBJECT structures.

Each YR_OBJECT has a void* data field which can be safely used by your code to store a pointer to any data you may need. A typical pattern is using the data field of the module's YR_OBJECT, like in the following example:

```
typedef struct _MY_DATA
{
    int some_integer;
} MY_DATA;

int module_load(
    YR_SCAN_CONTEXT* context,
    YR_OBJECT* module_object,
    void* module_data,
    size_t module_data_size)
{
    module->data = yr_malloc(sizeof(MY_DATA));
    ((MY_DATA*) module_object->data)->some_integer = 0;

    return ERROR_SUCCESS;
}
```

Don't forget to release the allocated memory in the `module_unload` function:

```
int module_unload(
    YR_OBJECT* module_object)
{
    yr_free(module_object->data);

    return ERROR_SUCCESS;
}
```

Warning: Don't use global variables for storing data. Functions in a module can be invoked from different threads at the same time and data corruption or misbehavior can occur.

4.5 More about functions

We already showed how to declare a function in *The declaration section*. Here we are going to discuss how to provide an implementation for them.

4.5.1 Function arguments

Within the function's code you get its arguments by using `integer_argument(n)`, `float_argument(n)`, `regexp_argument(n)`, `string_argument(n)` or `sized_string_argument(n)` depending on the type of the argument, where *n* is the 1-based argument's number.

`string_argument(n)` can be used when your function expects to receive a NULL-terminated C string, if your function can receive arbitrary binary data possibly containing NULL characters you must use `sized_string_argument(n)`.

Here you have some examples:

```
int64_t arg_1 = integer_argument(1);
RE* arg_2 = regexp_argument(2);
char* arg_3 = string_argument(3);
```

(continues on next page)

(continued from previous page)

```
SIZED_STRING* arg_4 = sized_string_argument(4);
double arg_5 = float_argument(1);
```

The C type for integer arguments is `int64_t`, for float arguments is `double`, for regular expressions is `RE*`, for NULL-terminated strings is `char*` and for strings possibly containing NULL characters is `SIZED_STRING*`. `SIZED_STRING` structures have the following attributes:

SIZED_STRING**length**

String's length.

c_string`char*` pointing to the string content.

4.5.2 Return values

Functions can return three types of values: strings, integers and floats. Instead of using the C *return* statement you must use `return_string(x)`, `return_integer(x)` or `return_float(x)` to return from a function, depending on the function's return type. In all cases *x* is a constant, variable, or expression evaluating to `char*`, `int64_t` or `double` respectively.

You can use `return_string(YR_UNDEFINED)`, `return_float(YR_UNDEFINED)` and `return_integer(YR_UNDEFINED)` to return undefined values from the function. This is useful in many situations, for example if the arguments passed to the functions don't make sense, or if your module expects a particular file format and the scanned file is from another format, or in any other case where your function can't a return a valid value.

Warning: Don't use the C *return* statement for returning from a function. The returned value will be interpreted as an error code.

4.5.3 Accessing objects

While writing a function we sometimes need to access values previously assigned to the module's variables, or additional data stored in the `data` field of `YR_OBJECT` structures as discussed earlier in [Storing data for later use](#). But for that we need a way to get access to the corresponding `YR_OBJECT` first. There are two functions to do that: `module()` and `parent()`. The `module()` function returns a pointer to the top-level `YR_OBJECT` corresponding to the module, the same one passed to the `module_load` function. The `parent()` function returns a pointer to the `YR_OBJECT` corresponding to the structure where the function is contained. For example, consider the following code snippet:

```
define_function(f1)
{
    YR_OBJECT* module = module();
    YR_OBJECT* parent = parent();

    // parent == module;
}

define_function(f2)
{
```

(continues on next page)

(continued from previous page)

```
YR_OBJECT* module = module();
YR_OBJECT* parent = parent();

    // parent != module;
}

begin_declarations;

    declare_function("f1", "i", "i", f1);

    begin_struct("foo");

        declare_function("f2", "i", "i", f2);

    end_struct("foo");

end_declarations;
```

In `f1` the `module` variable points to the top-level `YR_OBJECT` as well as the `parent` variable, because the parent for `f1` is the module itself. In `f2` however the `parent` variable points to the `YR_OBJECT` corresponding to the `foo` structure while `module` points to the top-level `YR_OBJECT` as before.

4.5.4 Scan context

From within a function you can also access the `YR_SCAN_CONTEXT` structure discussed earlier in [Accessing the scanned data](#). This is useful for functions which needs to inspect the file or process memory being scanned. This is how you get a pointer to the `YR_SCAN_CONTEXT` structure:

```
YR_SCAN_CONTEXT* context = scan_context();
```

Running YARA from the command-line

In order to invoke YARA you'll need two things: a file with the rules you want to use and the target to be scanned. The target can be a file, a folder, or a process.

```
yara [OPTIONS] RULES_FILE TARGET
```

In YARA 3.8 and below `RULES_FILE` was allowed to be a file with rules in source form or in compiled form indistinctly. In YARA 3.9 you need to explicitly specify that `RULES_FILE` contains compiled rules by using the `-C` flag.

```
yara [OPTIONS] -C RULES_FILE TARGET
```

This is a security measure to prevent users from inadvertently using compiled rules coming from a third-party. Using compiled rules from untrusted sources can lead to the execution of malicious code in your computer.

For compiling rules beforehand you can use the `yarac` tool. This way can save time, because for YARA it is faster to load compiled rules than compiling the same rules over and over again.

You can also pass multiple source files to `yara` like in the following example:

```
yara [OPTIONS] RULES_FILE_1 RULES_FILE_2 RULES_FILE_3 TARGET
```

Notice however that this only works for rules in source form. When invoking YARA with compiled rules a single file is accepted.

In the example above all rules share the same "default" namespace, which means that rule identifiers must be unique among all files. However you can specify a namespace for individual files. For example

```
yara [OPTIONS] namespace1:RULES_FILE_1 RULES_FILE_2 RULES_FILE_3 TARGET
```

In this case `RULE_FILE_1` uses `namespace1` while `RULES_FILE_2` and `RULES_FILE_3` share the default namespace.

In all cases rules will be applied to the target specified as the last argument to YARA, if it's a path to a directory all the files contained in it will be scanned. By default YARA does not attempt to scan directories recursively, but you can use the `-r` option for that.

Available options are:

- C** `--compiled-rules`
RULES_FILE contains rules already compiled with yara.
- c** `--count`
Print only number of matches.
- d** `<identifier>=<value> --define=identifier=value`
Define external variable. This option can be used multiple times.
- fail-on-warnings**
Treat warnings as errors. Has no effect if used with `--no-warnings`.
- f** `--fast-scan`
Fast matching mode.
- h** `--help`
Show help.
- i** `<identifier> --identifier=<identifier>`
Print rules named `<identifier>` and ignore the rest.
- max-process-memory-chunk=<size>**
When scanning process memory read the data in chunks of the given size.
- l** `<number> --max-rules=<number>`
Abort scanning after matching a number of rules.
- max-strings-per-rule=<number>**
Set maximum number of strings per rule (default=10000). If a rule has more then the specified number of strings an error will occur.

New in version 3.7.0.
- x** `<module>=<file> --module-data=<module>=<file>`
Pass the content of `<file>` as data to `<module>`. Example: `-x cuckoo=/cuckoo_report.json`.
- n** `--negate`
Print not satisfied rules only (negate).
- N** `--no-follow-symlinks`
Do not follow symlinks when scanning.
- w** `--no-warnings`
Disable warnings.
- m** `--print-meta`
Print metadata.
- D** `--print-module-data`
Print module data.
- e** `--print-namespace`
Print rules' namespace.
- S** `--print-stats`
Print rules' statistics.
- s** `--print-strings`
Print matching strings.
- L** `--print-string-length`
Print length of matching strings.

- g** `--print-tags`
Print tags.
- r** `--recursive`
Recursively search for directories. It follows symlinks.
- scan-list**
Scan files listed in FILE, one per line.
- z** `<size> --skip-larger=<size>`
Skip files larger than the given `<size>` in bytes when scanning a directory.
New in version 4.2.0.
- k** `<slots> --stack-size=<slots>`
Allocate a stack size of "slots" number of slots. Default: 16384. This will allow you to use larger rules, albeit with more memory overhead.
New in version 3.5.0.
- t** `<tag> --tag=<tag>`
Print rules tagged as `<tag>` and ignore the rest.
- p** `<number> --threads=<number>`
Use the specified `<number>` of threads to scan a directory.
- a** `<seconds> --timeout=<seconds>`
Abort scanning after a number of seconds has elapsed.
- v** `--version`
Show version information.

Here you have some examples:

- Apply rule in `/foo/bar/rules` to all files in the current directory. Subdirectories are not scanned:

```
yara /foo/bar/rules .
```

- Apply rules in `/foo/bar/rules` to `bazfile`. Only reports rules tagged as *Packer* or *Compiler*:

```
yara -t Packer -t Compiler /foo/bar/rules bazfile
```

- Scan all files in the `/foo` directory and its subdirectories:

```
yara /foo/bar/rules -r /foo
```

- Defines three external variables *mybool*, *myint* and *mystring*:

```
yara -d mybool=true -d myint=5 -d mystring="my string" /foo/bar/rules bazfile
```

- Apply rules in `/foo/bar/rules` to `bazfile` while passing the content of `cuckoo_json_report` to the cuckoo module:

```
yara -x cuckoo=cuckoo_json_report /foo/bar/rules bazfile
```

Using YARA from Python

YARA can be also used from Python through the `yara-python` library. Once the library is built and installed as described in *Compiling and installing YARA* you'll have access to the full potential of YARA from your Python scripts.

The first step is importing the YARA library:

```
import yara
```

Then you will need to compile your YARA rules before applying them to your data, the rules can be compiled from a file path:

```
rules = yara.compile(filepath='/foo/bar/myrules')
```

The default argument is `filepath`, so you don't need to explicitly specify its name:

```
rules = yara.compile('/foo/bar/myrules')
```

You can also compile your rules from a file object:

```
fh = open('/foo/bar/myrules')
rules = yara.compile(file=fh)
fh.close()
```

Or you can compile them directly from a Python string:

```
rules = yara.compile(source='rule dummy { condition: true }')
```

If you want to compile a group of files or strings at the same time you can do it by using the `filepaths` or `sources` named arguments:

```
rules = yara.compile(filepaths={
    'namespace1': '/my/path/rules1',
    'namespace2': '/my/path/rules2'
})
```

(continues on next page)

(continued from previous page)

```
rules = yara.compile(sources={
    'namespace1': 'rule dummy { condition: true }',
    'namespace2': 'rule dummy { condition: false }'
})
```

Notice that both `filepaths` and `sources` must be dictionaries with keys of string type. The dictionary keys are used as a namespace identifier, allowing to differentiate between rules with the same name in different sources, as occurs in the second example with the *dummy* name.

The `compile` method also has an optional boolean parameter named `includes` which allows you to control whether or not the include directive should be accepted in the source files, for example:

```
rules = yara.compile('/foo/bar/my_rules', includes=False)
```

If the source file contains include directives the previous line would raise an exception.

If `includes` are used, a python callback can be set to define a custom source for the imported files (by default they are read from disk). This callback function is set through the `include_callback` optional parameter. It receives the following parameters:

- `requested_filename`: file requested with 'include'
- `filename`: file containing the 'include' directive if applicable, else `None`
- `namespace`: namespace

And returns the requested rules sources as a single string.

If you are using external variables in your rules you must define those external variables either while compiling the rules, or while applying the rules to some file. To define your variables at the moment of compilation you should pass the `externals` parameter to the `compile` method. For example:

```
rules = yara.compile('/foo/bar/my_rules',
    externals= {'var1': 'some string', 'var2': 4, 'var3': True})
```

The `externals` parameter must be a dictionary with the names of the variables as keys and an associated value of either string, integer or boolean type.

The `compile` method also accepts the optional boolean argument `error_on_warning`. This argument tells YARA to raise an exception when a warning is issued during compilation. Such warnings are typically issued when your rules contains some construct that could be slowing down the scanning. The default value for the `error_on_warning` argument is `False`.

In all cases `compile` returns an instance of the class `yara.Rules` *Rules*. This class has a `save` method that can be used to save the compiled rules to a file:

```
rules.save('/foo/bar/my_compiled_rules')
```

The compiled rules can be loaded later by using the `load` method:

```
rules = yara.load('/foo/bar/my_compiled_rules')
```

Starting with YARA 3.4 both `save` and `load` accept file objects. For example, you can save your rules to a memory buffer with this code:

```
import StringIO

buff = StringIO.StringIO()
rules.save(file=buff)
```

The saved rules can be loaded from the memory buffer:

```
buff.seek(0)
rule = yara.load(file=buff)
```

The result of load is also an instance of the class `yara.Rules`.

Starting with YARA 4.3.0, Rules have a warning member which contains a list of warnings generated by the compiler. This allows you to know if the compiler generated warnings without them being hard errors using the `error_on_warning` argument.

Instances of Rules also have a match method, which allows you to apply the rules to a file:

```
matches = rules.match('/foo/bar/my_file')
```

But you can also apply the rules to a Python string:

```
with open('/foo/bar/my_file', 'rb') as f:
    matches = rules.match(data=f.read())
```

Or to a running process:

```
matches = rules.match(pid=1234)
```

As in the case of compile, the match method can receive definitions for external variables in the `externals` argument.

```
matches = rules.match('/foo/bar/my_file',
    externals= {'var1': 'some other string', 'var2': 100})
```

External variables defined during compile-time don't need to be defined again in subsequent calls to the match method. However you can redefine any variable as needed, or provide additional definitions that weren't provided during compilation.

In some situations involving a very large set of rules or huge files the match method can take too much time to run. In those situations you may find useful the `timeout` argument:

```
matches = rules.match('/foo/bar/my_huge_file', timeout=60)
```

If the match function does not finish before the specified number of seconds elapsed, a `TimeoutError` exception is raised.

You can also specify a callback function when invoking the match method. By default, the provided function will be called for every rule, no matter if matching or not. You can choose when your callback function is called by setting the `which_callbacks` parameter to one of `yara.CALLBACK_MATCHES`, `yara.CALLBACK_NON_MATCHES` or `yara.CALLBACK_ALL`. The default is to use `yara.CALLBACK_ALL`. Your callback function should expect a single parameter of dictionary type, and should return `CALLBACK_CONTINUE` to proceed to the next rule or `CALLBACK_ABORT` to stop applying rules to your data.

Here is an example:

```
import yara

def mycallback(data):
    print(data)
    return yara.CALLBACK_CONTINUE

matches = rules.match('/foo/bar/my_file', callback=mycallback, which_callbacks=yara.
↳ CALLBACK_MATCHES)
```

The passed dictionary will be something like this:

```
{
  'tags': ['foo', 'bar'],
  'matches': True,
  'namespace': 'default',
  'rule': 'my_rule',
  'meta': {},
  'strings': [StringMatch, StringMatch]
}
```

The *matches* field indicates if the rule matches the data or not. The *strings* field is a list of *yara.StringMatch* objects.

The *match* method returns a list of instances of the class *yara.Match*. Instances of this class have the same attributes as the dictionary passed to the callback function.

You can also specify a module callback function when invoking the *match* method. The provided function will be called for every imported module that scanned a file. Your callback function should expect a single parameter of dictionary type, and should return *CALLBACK_CONTINUE* to proceed to the next rule or *CALLBACK_ABORT* to stop applying rules to your data.

Here is an example:

```
import yara

def modules_callback(data):
    print(data)
    return yara.CALLBACK_CONTINUE

matches = rules.match('/foo/bar/my_file', modules_callback=modules_callback)
```

The passed dictionary will contain the information from the module.

You can also specify a warning callback function when invoking the *match* method. The provided function will be called for every runtime warning. Your callback function should expect two parameters. The first is an integer which contains the type of warning and the second is a string with the warning message. Your callback should return *CALLBACK_CONTINUE* to proceed with the scan or *CALLBACK_ABORT* to stop.

Possible values for the type are:

```
CALLBACK_TOO_MANY_MATCHES
```

Contents of the callback message depend on the type of the callback.

For *CALLBACK_TOO_MANY_MATCHES*, the message is a named tuple containing 3 items: namespace, rule and string. All contain string identifiers.

Here is an example:

```
import yara

def warnings_callback(warning_type, message):
    if warning_type == yara.CALLBACK_TOO_MANY_MATCHES:
        print(f"namespace: '{message.namespace}' rule: '{message.rule}' string: '{message.
↪string}'")
    return yara.CALLBACK_CONTINUE

matches = rules.match('/foo/bar/my_file', warnings_callback=warnings_callback)
```

If you do not use a warning callback a warning message will be sent to the normal python warning system for you and scanning will continue.

With YARA 4.2.0 a new `console` module was introduced which allows you to send log messages within YARA. These are, by default, printed to stdout in yara-python, but you can handle them in your own callback using the `console_callback` parameter.

Here is an example:

```
import yara

r = """
import "console"

rule a { condition: console.log("Hello from Python!") }
"""

def console(message):
    print(f"Callback: {message}")

rules = yara.compile(source=r)
rules.match("/bin/ls", console_callback=console)
rules.match("/bin/ls")
```

The type of the message parameter is a string.

You may also find that the default sizes for the stack for the matching engine in yara or the default size for the maximum number of strings per rule is too low. In the C libyara API, you can modify these using the `YR_CONFIG_STACK_SIZE` and `YR_CONFIG_MAX_STRINGS_PER_RULE` variables via the `yr_set_configuration_uint32` function in libyara. The command-line tool exposes these as the `--stack-size (-k)` and `--max-strings-per-rule` command-line arguments. In order to set these values via the Python API, you can use `yara.set_config` with either or both `stack_size` and `max_strings_per_rule` provided as kwargs. At the time of this writing, the default stack size was 16384 and the default maximum strings per rule was 10000.

Also, `yara.set_config` accepts the `max_match_data` argument for controlling the maximum number of bytes that will be returned for each matching string. This is equivalent to using `YR_CONFIG_MAX_MATCH_DATA` with the `yr_set_configuration_uint32` in the C API. By the default this is set to 512.

Here are a few example calls:

```
yara.set_config(stack_size=65536)
yara.set_config(max_strings_per_rule=50000, stack_size=65536)
yara.set_config(max_strings_per_rule=20000)
yara.set_config(max_match_data=128)
```

6.1 Reference

`yara.compile(...)`

Compile YARA sources.

Either *filepath*, *source*, *file*, *filepaths* or *sources* must be provided. The remaining arguments are optional.

Parameters

- **filepath** (*str*) – Path to the source file.
- **source** (*str*) – String containing the rules code.
- **file** (*file-object*) – Source file as a file object.
- **filepaths** (*dict*) – Dictionary where keys are namespaces and values are paths to source files.
- **sources** (*dict*) – Dictionary where keys are namespaces and values are strings containing rules code.
- **externals** (*dict*) – Dictionary with external variables. Keys are variable names and values are variable values.
- **includes** (*boolean*) – True if include directives are allowed or False otherwise. Default value: *True*.
- **error_on_warning** (*boolean*) – If true warnings are treated as errors, raising an exception.

Returns Compiled rules object.

Return type `yara.Rules`

Raises

- **yara.SyntaxError** – If a syntax error was found.
- **yara.Error** – If an error occurred.

`yara.load(...)`

Changed in version 3.4.0.

Load compiled rules from a path or file object. Either *filepath* or *file* must be provided.

Parameters

- **filepath** (*str*) – Path to a compiled rules file
- **file** (*file-object*) – A file object supporting the `read` method.

Returns Compiled rules object.

Return type `yara.Rules`

Raises **yara.Error**: If an error occurred while loading the file.

`yara.set_config(...)`

Set the configuration variables accessible through the `yr_set_configuration` API.

Provide either *stack_size*, *max_strings_per_rule*, or *max_match_data*. These kwargs take unsigned integer values as input and will assign the provided value to the `yr_set_configuration(...)` variables `YR_CONFIG_STACK_SIZE`, `YR_CONFIG_MAX_STRINGS_PER_RULE`, and `YR_CONFIG_MAX_MATCH_DATA` respectively.

Parameters

- **stack_size** (*int*) – Stack size to use for YR_CONFIG_STACK_SIZE
- **max_strings_per_rule** (*int*) – Maximum number of strings to allow per yara rule. Will be mapped to YR_CONFIG_MAX_STRINGS_PER_RULE.
- **max_match_data** (*int*) – Maximum number of bytes to allow per yara match. Will be mapped to YR_CONFIG_MAX_MATCH_DATA.

Returns None

Return type NoneType

Raises `yara.Error`: If an error occurred.

class `yara.Rules`

Instances of this class are returned by `yara.compile()` and represents a set of compiled rules.

match (*filepath*, *pid*, *data*, *externals=None*, *callback=None*, *fast=False*, *timeout=None*, *modules_data=None*, *modules_callback=None*, *warnings_callback=None*, *which_callbacks=CALLBACK_ALL*, *console_callback=None*)
Scan a file, process memory or data string.

Either *filepath*, *pid* or *data* must be provided. The remaining arguments are optional.

Parameters

- **filepath** (*str*) – Path to the file to be scanned.
- **pid** (*int*) – Process id to be scanned.
- **data** (*str/bytes*) – Data to be scanned.
- **externals** (*dict*) – Dictionary with external variables. Keys are variable names and values are variable values.
- **callback** (*function*) – Callback function invoked for each rule.
- **fast** (*bool*) – If true performs a fast mode scan.
- **timeout** (*int*) – Aborts the scanning when the number of specified seconds have elapsed.
- **modules_data** (*dict*) – Dictionary with additional data to modules. Keys are module names and values are *bytes* objects containing the additional data.
- **modules_callback** (*function*) – Callback function invoked for each module.
- **warnings_callback** (*function*) – Callback function invoked for warning, like `yara.CALLBACK_TOO_MANY_MATCHES`.
- **which_callbacks** (*int*) – An integer that indicates in which cases the callback function must be called. Possible values are `yara.CALLBACK_ALL`, `yara.CALLBACK_MATCHES` and `yara.CALLBACK_NON_MATCHES`.
- **console_callback** (*function*) – Callback function invoked for each console module call.

Raises

- **yara.TimeoutError** – If the timeout was reached.
- **yara.Error** – If an error occurred during the scan.

save (...)

Changed in version 3.4.0: Save compiled rules to a file. Either *filepath* or *file* must be provided.

param str filepath Path to the file.

param file-object file A file object supporting the `write` method.

raises `yara.Error`: If an error occurred while saving the file.

class `yara.Match`

New in version 4.3.0.

Objects returned by `yara.Rules.match()`, representing a match.

rule

Name of the matching rule.

namespace

Namespace associated to the matching rule.

tags

Array of strings containing the tags associated to the matching rule.

meta

Dictionary containing metadata associated to the matching rule.

strings

List of `StringMatch` objects.

class `yara.StringMatch`

New in version 3.4.0.

Objects which represent string matches.

identifier

Name of the matching string.

instances

List of `StringMatchInstance` objects.

`is_xor()`

Returns a boolean if the string is using the xor modifier.

class `yara.StringMatchInstance`

New in version 4.3.0.

Objects which represent instances of matched strings.

matched_data

Bytes of the matched data.

matched_length

Length of the matched data.

offset

Offset of the matched data.

xor_key

XOR key found for the string.

`plaintext()`

Returns the plaintext version of the string after xor key is applied. If the string is not an xor string then no modification is done.

You can integrate YARA into your C/C++ project by using the API provided by the *libyara* library. This API gives you access to every YARA feature and it's the same API used by the command-line tools *yara* and *yarac*.

7.1 Initializing and finalizing *libyara*

The first thing your program must do when using *libyara* is initializing the library. This is done by calling the `yr_initialize()` function. This function allocates any resources needed by the library and initializes internal data structures. Its counterpart is `yr_finalize()`, which must be called when you are finished using the library.

In a multi-threaded program only the main thread must call `yr_initialize()` and `yr_finalize()`. No additional work is required from other threads using the library.

7.2 Compiling rules

Before using your rules to scan any data you need to compile them into binary form. For that purpose you'll need a YARA compiler, which can be created with `yr_compiler_create()`. After being used, the compiler must be destroyed with `yr_compiler_destroy()`.

You can use `yr_compiler_add_file()`, `yr_compiler_add_fd()`, or `yr_compiler_add_string()` to add one or more input sources to be compiled. Both of these functions receive an optional namespace. Rules added under the same namespace behave as if they were contained within the same source file or string, so, rule identifiers must be unique among all the sources sharing a namespace. If the namespace argument is `NULL` the rules are put in the *default* namespace.

The `yr_compiler_add_file()`, `yr_compiler_add_fd()`, and `yr_compiler_add_string()` functions return the number of errors found in the source code. If the rules are correct they will return 0. If any of these functions return an error the compiler can't be used anymore, neither for adding more rules nor getting the compiled rules.

For obtaining detailed error information you must set a callback function by using `yr_compiler_set_callback()` before calling any of the compiling functions. The callback function has the following prototype:

```
void callback_function(
    int error_level,
    const char* file_name,
    int line_number,
    const YR_RULE* rule,
    const char* message,
    void* user_data)
```

Changed in version 4.0.0.

Possible values for `error_level` are `YARA_ERROR_LEVEL_ERROR` and `YARA_ERROR_LEVEL_WARNING`. The arguments `file_name` and `line_number` contain the file name and line number where the error or warning occurred. `file_name` is the one passed to `yr_compiler_add_file()` or `yr_compiler_add_fd()`. It can be `NULL` if you passed `NULL` or if you're using `yr_compiler_add_string()`. `rule` is a pointer to the `YR_RULE` structure representing the rule that contained the error, but it can be `NULL` if the error is not contained in a specific rule. The `user_data` pointer is the same you passed to `yr_compiler_set_callback()`.

By default, for rules containing references to other files (include "filename.yara"), YARA will try to find those files on disk. However, if you want to fetch the imported rules from another source (eg: from a database or remote service), a callback function can be set with `yr_compiler_set_include_callback()`.

This callback receives the following parameters:

- `include_name`: name of the requested file.
- `calling_rule_filename`: the requesting file name (`NULL` if not a file).
- `calling_rule_namespace`: namespace (`NULL` if undefined).
- `user_data` same pointer passed to `yr_compiler_set_include_callback()`.

It should return the requested file's content as a null-terminated string. The memory for this string should be allocated by the callback function. Once it is safe to free the memory used to return the callback's result, the `include_free` function passed to `yr_compiler_set_include_callback()` will be called. If the memory does not need to be freed, `NULL` can be passed as `include_free` instead. You can completely disable support for includes by setting a `NULL` callback function with `yr_compiler_set_include_callback()`.

The callback function has the following prototype:

```
const char* include_callback(
    const char* include_name,
    const char* calling_rule_filename,
    const char* calling_rule_namespace,
    void* user_data);
```

The free function has the following prototype:

```
void include_free(
    const char* callback_result_ptr,
    void* user_data);
```

After you successfully added some sources you can get the compiled rules using the `yr_compiler_get_rules()` function. You'll get a pointer to a `YR_RULES` structure which can be used to scan your data as described in *Scanning data*. Once `yr_compiler_get_rules()` is invoked you can not add more sources to the compiler, but you can call `yr_compiler_get_rules()` multiple times. Each time this function is called it returns a

pointer to the same `YR_RULES` structure. Notice that this behaviour is new in YARA 4.0.0, in YARA 3.X and 2.X `yr_compiler_get_rules()` returned a new copy the `YR_RULES` structure.

Instances of `YR_RULES` must be destroyed with `yr_rules_destroy()`.

7.3 Defining external variables

If your rules make use of external variables (like in the example below), you must define those variables by using any of the `yr_compiler_define_XXXX_variable` functions. Variables must be defined before rules are compiled with `yr_compiler_add_XXXX` and they must be defined with a type that matches the context in which the variable is used in the rule, a variable that is used like `my_var == 5` can't be defined as a string variable.

While defining external variables with `yr_compiler_define_XXXX_variable` you must provide a value for each variable. That value is embedded in the compiled rules and used whenever the variable appears in a rule. However, you can change the value associated to an external variable after the rules has been compiled by using any of the `yr_rules_define_XXXX_variable` functions.

7.4 Saving and retrieving compiled rules

Compiled rules can be saved to a file and retrieved later by using `yr_rules_save()` and `yr_rules_load()`. Rules compiled and saved in one machine can be loaded in another machine as long as they have the same endianness, no matter the operating system or if they are 32-bit or 64-bit systems. However files saved with older versions of YARA may not work with newer versions due to changes in the file layout.

You can also save and retrieve your rules to and from generic data streams by using functions `yr_rules_save_stream()` and `yr_rules_load_stream()`. These functions receive a pointer to a `YR_STREAM` structure, defined as:

```
typedef struct _YR_STREAM
{
    void* user_data;

    YR_STREAM_READ_FUNC read;
    YR_STREAM_WRITE_FUNC write;
} YR_STREAM;
```

You must provide your own implementation for read and write functions. The read function is used by `yr_rules_load_stream()` to read data from your stream and the write function is used by `yr_rules_save_stream()` to write data into your stream.

Your read and write functions must respond to these prototypes:

```
size_t read(
    void* ptr,
    size_t size,
    size_t count,
    void* user_data);

size_t write(
    const void* ptr,
    size_t size,
    size_t count,
    void* user_data);
```

The `ptr` argument is a pointer to the buffer where the `read` function should put the read data, or where the `write` function will find the data that needs to be written to the stream. In both cases `size` is the size of each element being read or written and `count` the number of elements. The total size of the data being read or written is `size * count`. The `read` function must return the number of elements read, the `write` function must return the total number of elements written.

The `user_data` pointer is the same you specified in the `YR_STREAM` structure. You can use it to pass arbitrary data to your `read` and `write` functions.

7.5 Scanning data

Once you have an instance of `YR_RULES` you can use it directly with one of the `yr_rules_scan_XXXX` functions described below, or create a scanner with `yr_scanner_create()`. Let's start by discussing the first approach.

The `YR_RULES` you got from the compiler can be used with `yr_rules_scan_file()`, `yr_rules_scan_fd()` or `yr_rules_scan_mem()` for scanning a file, a file descriptor and a in-memory buffer respectively. The results from the scan are returned to your program via a callback function. The callback has the following prototype:

```
int callback_function(
    YR_SCAN_CONTEXT* context,
    int message,
    void* message_data,
    void* user_data);
```

Possible values for message are:

```
CALLBACK_MSG_RULE_MATCHING
CALLBACK_MSG_RULE_NOT_MATCHING
CALLBACK_MSG_SCAN_FINISHED
CALLBACK_MSG_IMPORT_MODULE
CALLBACK_MSG_MODULE_IMPORTED
CALLBACK_MSG_TOO_MANY_MATCHES
CALLBACK_MSG_CONSOLE_LOG
```

Your callback function will be called once for each rule with either a `CALLBACK_MSG_RULE_MATCHING` or `CALLBACK_MSG_RULE_NOT_MATCHING` message, depending if the rule is matching or not. In both cases a pointer to the `YR_RULE` structure associated with the rule is passed in the `message_data` argument. You just need to perform a typecast from `void*` to `YR_RULE*` to access the structure. You can control whether or not YARA calls your callback function with `CALLBACK_MSG_RULE_MATCHING` and `CALLBACK_MSG_RULE_NOT_MATCHING` messages by using the `SCAN_FLAGS_REPORT_RULES_MATCHING` and `SCAN_FLAGS_REPORT_RULES_NOT_MATCHING` as described later in this section.

This callback is also called with the `CALLBACK_MSG_IMPORT_MODULE` message. All modules referenced by an `import` statement in the rules are imported once for every file being scanned. In this case `message_data` points to a `YR_MODULE_IMPORT` structure. This structure contains a `module_name` field pointing to a null terminated string with the name of the module being imported and two other fields `module_data` and `module_data_size`. These fields are initially set to `NULL` and `0`, but your program can assign a pointer to some arbitrary data to `module_data` while setting `module_data_size` to the size of the data. This way you can pass additional data to those modules requiring it, like the *Cuckoo module* for example.

Once a module is imported the callback is called again with the `CALLBACK_MSG_MODULE_IMPORTED`. When this happens `message_data` points to a `YR_OBJECT_STRUCTURE` structure. This structure contains all the information provided by the module about the currently scanned file.

If during the scan a string hits the maximum number of matches, your callback will be called once with the `CALLBACK_MSG_TOO_MANY_MATCHES`. When this happens, `message_data` is a `YR_STRING*` which points to the string which caused the warning. If your callback returns `CALLBACK_CONTINUE`, the string will be disabled and scanning will continue, otherwise scanning will be halted.

Your callback will be called from the console module (*Console module*) with the `CALLBACK_MSG_CONSOLE_LOG` message. When this happens, the `message_data` argument will be a `char*` that is the string generated by the console module. Your callback can do whatever it wants with this string, including logging it to an external logging source, or printing it to stdout.

Lastly, the callback function is also called with the `CALLBACK_MSG_SCAN_FINISHED` message when the scan is finished. In this case `message_data` is `NULL`.

Notice that you shouldn't call any of the `yr_rules_scan_XXXX` functions from within the callback as those functions are not re-entrant.

Your callback function must return one of the following values:

```
CALLBACK_CONTINUE
CALLBACK_ABORT
CALLBACK_ERROR
```

If it returns `CALLBACK_CONTINUE` YARA will continue normally, `CALLBACK_ABORT` will abort the scan but the result from the `yr_rules_scan_XXXX` function will be `ERROR_SUCCESS`. On the other hand `CALLBACK_ERROR` will abort the scanning too, but the result from `yr_rules_scan_XXXX` will be `ERROR_CALLBACK_ERROR`.

The `user_data` argument passed to your callback function is the same you passed `yr_rules_scan_XXXX`. This pointer is not touched by YARA, it's just a way for your program to pass arbitrary data to the callback function.

All `yr_rules_scan_XXXX` functions receive a `flags` argument that allows to tweak some aspects of the scanning process. The supported flags are the following ones:

```
SCAN_FLAGS_FAST_MODE SCAN_FLAGS_NO_TRYCATCH SCAN_FLAGS_REPORT_RULES_MATCHING
SCAN_FLAGS_REPORT_RULES_NOT_MATCHING
```

The `SCAN_FLAGS_FAST_MODE` flag makes the scanning a little faster by avoiding multiple matches of the same string when not necessary. Once the string was found in the file it's subsequently ignored, implying that you'll have a single match for the string, even if it appears multiple times in the scanned data. This flag has the same effect of the `-f` command-line option described in *Running YARA from the command-line*.

`SCAN_FLAGS_REPORT_RULES_MATCHING` and `SCAN_FLAGS_REPORT_RULES_NOT_MATCHING` control whether the callback is invoked for rules that are matching or for rules that are not matching respectively. If `SCAN_FLAGS_REPORT_RULES_MATCHING` is specified alone, the callback will be called for matching rules with the `CALLBACK_MSG_RULE_MATCHING` message but it won't be called for non-matching rules. If `SCAN_FLAGS_REPORT_RULES_NOT_MATCHING` is specified alone, the opposite happens, the callback will be called with `CALLBACK_MSG_RULE_NOT_MATCHING` messages but not with `CALLBACK_MSG_RULE_MATCHING` messages. If both flags are combined together (the default) the callback will be called for both matching and non-matching rules. For backward compatibility, if none of these two flags are specified, the scanner will follow the default behavior.

Additionally, `yr_rules_scan_XXXX` functions can receive a `timeout` argument which forces the scan to abort after the specified number of seconds (approximately). If `timeout` is 0 it means no timeout at all.

7.5.1 Using a scanner

The `yr_rules_scan_XXXX` functions are enough in most cases, but sometimes you may need a fine-grained control over the scanning. In those cases you can create a scanner with `yr_scanner_create()`. A scanner is simply a

wrapper around a `YR_RULES` structure that holds additional configuration like external variables without affecting other users of the `YR_RULES` structure.

A scanner is particularly useful when you want to use the same `YR_RULES` with multiple workers (it could be a separate thread, a coroutine, etc) and each worker needs to set different set of values for external variables. In that case you can't use `yr_rules_define_XXXX_variable` for setting the values of your external variables, as every worker using the `YR_RULES` will be affected by such changes. However each worker can have its own scanner, where the scanners share the same `YR_RULES`, and use `yr_scanner_define_XXXX_variable` for setting external variables without affecting the rest of the workers.

This is a better solution than having a separate `YR_RULES` for each worker, as `YR_RULES` structures have large memory footprint (specially if you have a lot of rules) while scanners are very lightweight.

7.6 API reference

7.6.1 Data structures

YR_COMPILER

Data structure representing a YARA compiler.

YR_SCAN_CONTEXT

Data structure that holds information about an on-going scan. A pointer to this structure is passed to the callback function that receives notifications about matches found. This structure is also used for iterating over the

YR_MATCH

Data structure representing a string match.

int64_t base

Base offset/address for the match. While scanning a file this field is usually zero, while scanning a process memory space this field is the virtual address of the memory block where the match was found.

int64_t offset

Offset of the match relative to *base*.

int32_t match_length

Length of the matching string

const uint8_t* data

Pointer to a buffer containing a portion of the matching string.

int32_t data_length

Length of data buffer. `data_length` is the minimum of `match_length` and `MAX_MATCH_DATA`.

Changed in version 3.5.0.

YR_META

Data structure representing a metadata value.

const char* identifier

Meta identifier.

int32_t type

One of the following metadata types:

`META_TYPE_INTEGER` `META_TYPE_STRING` `META_TYPE_BOOLEAN`

YR_MODULE_IMPORT

const char* **module_name**

Name of the module being imported.

void* **module_data**

Pointer to additional data passed to the module. Initially set to NULL, your program is responsible for setting this pointer while handling the CALLBACK_MSG_IMPORT_MODULE message.

size_t **module_data_size**

Size of additional data passed to module. Your program must set the appropriate value if `module_data` is modified.

YR_RULE

Data structure representing a single rule.

const char* **identifier**

Rule identifier.

const char* **tags**

Pointer to a sequence of null terminated strings with tag names. An additional null character marks the end of the sequence. Example: `tag1\0tag2\0tag3\0\0`. To iterate over the tags you can use `yr_rule_tags_foreach()`.

*YR_META** **metas**

Pointer to a sequence of *YR_META* structures. To iterate over the structures use `yr_rule_metas_foreach()`.

*YR_STRING** **strings**

Pointer to a sequence of *YR_STRING* structures. To iterate over the structures use `yr_rule_strings_foreach()`.

*YR_NAMESPACE** **ns**

Pointer to a *YR_NAMESPACE* structure.

YR_RULES

Data structure representing a set of compiled rules.

YR_STREAM

New in version 3.4.0.

Data structure representing a stream used with functions `yr_rules_load_stream()` and `yr_rules_save_stream()`.

void* **user_data**

A user-defined pointer.

YR_STREAM_READ_FUNC **read**

A pointer to the stream's read function provided by the user.

YR_STREAM_WRITE_FUNC **write**

A pointer to the stream's write function provided by the user.

YR_STRING

Data structure representing a string declared in a rule.

const char* **identifier**

String identifier.

YR_NAMESPACE

Data structure representing a rule namespace.

const char* **name**

Rule namespace.

7.6.2 Functions

int **yr_initialize** (void)

Initialize the library. Must be called by the main thread before using any other function. Return `ERROR_SUCCESS` on success another error code in case of error. The list of possible return codes vary according to the modules compiled into YARA.

int **yr_finalize** (void)

Finalize the library. Must be called by the main free to release any resource allocated by the library. Return `ERROR_SUCCESS` on success another error code in case of error. The list of possible return codes vary according to the modules compiled into YARA.

int **yr_compiler_create** (*YR_COMPILER** compiler*)

Create a YARA compiler. You must pass the address of a pointer to a `YR_COMPILER`, the function will set the pointer to the newly allocated compiler. Returns one of the following error codes:

`ERROR_SUCCESS`

`ERROR_INSUFFICIENT_MEMORY`

void **yr_compiler_destroy** (*YR_COMPILER* compiler*)

Destroy a YARA compiler.

void **yr_compiler_set_callback** (*YR_COMPILER* compiler*, *YR_COMPILER_CALLBACK_FUNC callback*, *void* user_data*)

Changed in version 3.3.0.

Set a callback for receiving error and warning information. The *user_data* pointer is passed to the callback function.

void **yr_compiler_set_include_callback** (*YR_COMPILER* compiler*, *YR_COMPILER_INCLUDE_CALLBACK_FUNC callback*, *YR_COMPILER_INCLUDE_FREE_FUNC include_free*, *void* user_data*)

New in version 3.7.0: Set a callback to provide rules from a custom source when `include` directive is invoked. The *user_data* pointer is untouched and passed back to the callback function and to the free function. Once the callback's result is no longer needed, the `include_free` function will be called. If the memory does not need to be freed, `include_free` can be set to `NULL`. If *callback* is set to `NULL` support for `include` directives is disabled.

int **yr_compiler_add_file** (*YR_COMPILER* compiler*, *FILE* file*, *const char* namespace*, *const char* file_name*)

Compile rules from a *file*. Rules are put into the specified *namespace*, if *namespace* is `NULL` they will be put into the default namespace. *file_name* is the name of the file for error reporting purposes and can be set to `NULL`. Returns the number of errors found during compilation.

int **yr_compiler_add_fd** (*YR_COMPILER* compiler*, *YR_FILE_DESCRIPTOR rules_fd*, *const char* namespace*, *const char* file_name*)

New in version 3.6.0.

Compile rules from a *file descriptor*. Rules are put into the specified *namespace*, if *namespace* is `NULL` they will be put into the default namespace. *file_name* is the name of the file for error reporting purposes and can be set to `NULL`. Returns the number of errors found during compilation.

int **yr_compiler_add_string** (*YR_COMPILER* compiler*, *const char* string*, *const char* namespace_*)

Compile rules from a *string*. Rules are put into the specified *namespace*, if *namespace* is `NULL` they will be put into the default namespace. Returns the number of errors found during compilation.

int **yr_compiler_get_rules** (*YR_COMPILER* compiler*, *YR_RULES** rules*)

Get the compiled rules from the compiler. Returns one of the following error codes:

`ERROR_SUCCESS`

ERROR_INSUFFICIENT_MEMORY

int **yr_compiler_define_integer_variable** (*YR_COMPILER** compiler, const char* identifier, int64_t value)
 Define an integer external variable.

int **yr_compiler_define_float_variable** (*YR_COMPILER** compiler, const char* identifier, double value)
 Define a float external variable.

int **yr_compiler_define_boolean_variable** (*YR_COMPILER** compiler, const char* identifier, int value)
 Define a boolean external variable.

int **yr_compiler_define_string_variable** (*YR_COMPILER** compiler, const char* identifier, const char* value)
 Define a string external variable.

int **yr_rules_define_integer_variable** (*YR_RULES** rules, const char* identifier, int64_t value)
 Define an integer external variable.

int **yr_rules_define_boolean_variable** (*YR_RULES** rules, const char* identifier, int value)
 Define a boolean external variable.

int **yr_rules_define_float_variable** (*YR_RULES** rules, const char* identifier, double value)
 Define a float external variable.

int **yr_rules_define_string_variable** (*YR_RULES** rules, const char* identifier, const char* value)
 Define a string external variable.

void **yr_rules_destroy** (*YR_RULES** rules)
 Destroy compiled rules.

int **yr_rules_save** (*YR_RULES** rules, const char* filename)
 Save compiled *rules* into the file specified by *filename*. Only rules obtained from *yr_compiler_get_rules()* can be saved. Those obtained from *yr_rules_load()* or *yr_rules_load_stream()* can not be saved. Returns one of the following error codes:

ERROR_SUCCESS

ERROR_COULD_NOT_OPEN_FILE

int **yr_rules_save_stream** (*YR_RULES** rules, *YR_STREAM** stream)
 New in version 3.4.0.

Save compiled *rules* into *stream*. Only rules obtained from *yr_compiler_get_rules()* can be saved. Those obtained from *yr_rules_load()* or *yr_rules_load_stream()* can not be saved. Returns one of the following error codes:

ERROR_SUCCESS

int **yr_rules_load** (const char* filename, *YR_RULES*** rules)
 Load compiled rules from the file specified by *filename*. Returns one of the following error codes:

ERROR_SUCCESS

ERROR_INSUFFICIENT_MEMORY

ERROR_COULD_NOT_OPEN_FILE

ERROR_INVALID_FILE

ERROR_CORRUPT_FILE

ERROR_UNSUPPORTED_FILE_VERSION

int **yr_rules_load_stream**(*YR_STREAM** stream, *YR_RULES*** rules)

New in version 3.4.0.

Load compiled rules from *stream*. Rules loaded this way can not be saved back using *yr_rules_save_stream()*. Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_INVALID_FILE
ERROR_CORRUPT_FILE
ERROR_UNSUPPORTED_FILE_VERSION

int **yr_rules_scan_mem**(*YR_RULES** rules, const uint8_t* buffer, size_t buffer_size, int flags, *YR_CALLBACK_FUNC* callback, void* user_data, int timeout)

Scan a memory buffer. Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_TOO_MANY_SCAN_THREADS
ERROR_SCAN_TIMEOUT
ERROR_CALLBACK_ERROR
ERROR_TOO_MANY_MATCHES

int **yr_rules_scan_file**(*YR_RULES** rules, const char* filename, int flags, *YR_CALLBACK_FUNC* callback, void* user_data, int timeout)

Scan a file. Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_COULD_NOT_MAP_FILE
ERROR_TOO_MANY_SCAN_THREADS
ERROR_SCAN_TIMEOUT
ERROR_CALLBACK_ERROR
ERROR_TOO_MANY_MATCHES

int **yr_rules_scan_fd**(*YR_RULES** rules, *YR_FILE_DESCRIPTOR* fd, int flags, *YR_CALLBACK_FUNC* callback, void* user_data, int timeout)

Scan a file descriptor. In POSIX systems *YR_FILE_DESCRIPTOR* is an *int*, as returned by the *open()* function. In Windows *YR_FILE_DESCRIPTOR* is a *HANDLE* as returned by *CreateFile()*.

Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_COULD_NOT_MAP_FILE
ERROR_TOO_MANY_SCAN_THREADS
ERROR_SCAN_TIMEOUT
ERROR_CALLBACK_ERROR
ERROR_TOO_MANY_MATCHES

yr_rule_tags_foreach (rule, tag)

Iterate over the tags of a given rule running the block of code that follows each time with a different value for *tag* of type `const char*`. Example:

```
const char* tag;

/* rule is a YR_RULE object */

yr_rule_tags_foreach(rule, tag)
{
    ..do something with tag
}
```

yr_rule metas_foreach (rule, meta)

Iterate over the *YR_META* structures associated with a given rule running the block of code that follows each time with a different value for *meta*. Example:

```
YR_META* meta;

/* rule is a YR_RULE object */

yr_rule_metas_foreach(rule, meta)
{
    ..do something with meta
}
```

yr_rule_strings_foreach (rule, string)

Iterate over the *YR_STRING* structures associated with a given rule running the block of code that follows each time with a different value for *string*. Example:

```
YR_STRING* string;

/* rule is a YR_RULE object */

yr_rule_strings_foreach(rule, string)
{
    ..do something with string
}
```

yr_string_matches_foreach (context, string, match)

Iterate over the *YR_MATCH* structures that represent the matches found for a given string during a scan running the block of code that follows, each time with a different value for *match*. The *context* argument is a pointer to a *YR_SCAN_CONTEXT* that is passed to the callback function and *string* is a pointer to a *YR_STRING*. Example:

```
YR_MATCH* match;

/* context is a YR_SCAN_CONTEXT* and string is a YR_STRING* */

yr_string_matches_foreach(context, string, match)
{
    ..do something with match
}
```

yr_rules_foreach (rules, rule)

Iterate over each *YR_RULE* in a *YR_RULES* object running the block of code that follows each time with a different value for *rule*. Example:

```
YR_RULE* rule;

/* rules is a YR_RULES object */

yr_rules_foreach(rules, rule)
{
    ..do something with rule
}
```

void **yr_rule_disable** (*YR_RULE** rule)

New in version 3.7.0.

Disable the specified rule. Disabled rules are completely ignored during the scanning process and they won't match. If the disabled rule is used in the condition of some other rule the value for the disabled rule is neither true nor false but undefined. For more information about undefined values see [Undefined values](#).

void **yr_rule_enable** (*YR_RULE** rule)

New in version 3.7.0.

Enables the specified rule. After being disabled with *yr_rule_disable()* a rule can be enabled again by using this function.

int **yr_scanner_create** (*YR_RULES** rules, YR_SCANNER **scanner)

New in version 3.8.0.

Creates a new scanner that can be used for scanning data with the provided rules. *scanner* must be a pointer to a YR_SCANNER, the function will set the pointer to the newly allocated scanner. Returns one of the following error codes:

ERROR_INSUFFICIENT_MEMORY

void **yr_scanner_destroy** (YR_SCANNER *scanner)

New in version 3.8.0.

Destroy a scanner. After using a scanner it must be destroyed with this function.

void **yr_scanner_set_callback** (YR_SCANNER *scanner, YR_CALLBACK_FUNC callback,
void* user_data)

New in version 3.8.0.

Set a callback function that will be called for reporting any matches found by the scanner.

void **yr_scanner_set_timeout** (YR_SCANNER* scanner, int timeout)

New in version 3.8.0.

Set the maximum number of seconds that the scanner will spend in any call to *yr_scanner_scan_xxx*.

void **yr_scanner_set_flags** (YR_SCANNER* scanner, int flags)

New in version 3.8.0.

Set the flags that will be used by any call to *yr_scanner_scan_xxx*. The supported flags are:

SCAN_FLAGS_FAST_MODE: Enable fast scan mode. SCAN_FLAGS_NO_TRYCATCH: Disable exception handling.
SCAN_FLAGS_REPORT_RULES_MATCHING: If this SCAN_FLAGS_REPORT_RULES_NOT_MATCHING

int **yr_scanner_define_integer_variable** (YR_SCANNER* scanner, const char* identifier,
int64_t value)

New in version 3.8.0.

Define an integer external variable.

int **yr_scanner_define_boolean_variable** (YR_SCANNER* *scanner*, const char* *identifier*,
int *value*)

New in version 3.8.0.

Define a boolean external variable.

int **yr_scanner_define_float_variable** (YR_SCANNER* *scanner*, const char* *identifier*, dou-
ble *value*)

New in version 3.8.0.

Define a float external variable.

int **yr_scanner_define_string_variable** (YR_SCANNER* *scanner*, const char* *identifier*, const
char* *value*)

New in version 3.8.0.

Define a string external variable.

int **yr_scanner_scan_mem_blocks** (YR_SCANNER* *scanner*, YR_MEMORY_BLOCK_ITERATOR* *it-*
erator)

New in version 3.8.0.

Scan a series of memory blocks that are provided by a YR_MEMORY_BLOCK_ITERATOR. The iterator has a pair of *first* and *next* functions that must return the first and next blocks respectively. When these functions return *NULL* it indicates that there are not more blocks to scan.

In YARA 4.1 and later the *first* and *next* functions can return *NULL* and set the *last_error* field in YR_MEMORY_BLOCK_ITERATOR to *ERROR_BLOCK_NOT_READY*. This indicates that the iterator is not able to return the next block yet, but the operation may be retried. In such cases *yr_scanner_scan_mem_blocks* also returns *ERROR_BLOCK_NOT_READY* but the scanner maintains its state and this function can be called again for continuing the scanning where it was left. This can be done multiple times until the block is ready and the iterator is able to return it.

Notice however that once the iterator completes a full iteration, any subsequent iteration should proceed without returning *ERROR_BLOCK_NOT_READY*. During the first iteration the iterator should store in memory any information that it needs about the blocks, so that it can be iterated again without relying on costly operations that may result in a *ERROR_BLOCK_NOT_READY* error.

Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_TOO_MANY_SCAN_THREADS
ERROR_SCAN_TIMEOUT
ERROR_CALLBACK_ERROR
ERROR_TOO_MANY_MATCHES
ERROR_BLOCK_NOT_READY

int **yr_scanner_scan_mem** (YR_SCANNER* *scanner*, const uint8_t* *buffer*, size_t *buffer_size*)

New in version 3.8.0.

Scan a memory buffer. Returns one of the following error codes:

ERROR_SUCCESS
ERROR_INSUFFICIENT_MEMORY
ERROR_TOO_MANY_SCAN_THREADS
ERROR_SCAN_TIMEOUT

ERROR_CALLBACK_ERROR

ERROR_TOO_MANY_MATCHES

int **yr_scanner_scan_file** (YR_SCANNER* *scanner*, const char* *filename*)

New in version 3.8.0.

Scan a file. Returns one of the following error codes:

ERROR_SUCCESS

ERROR_INSUFFICIENT_MEMORY

ERROR_TOO_MANY_SCAN_THREADS

ERROR_SCAN_TIMEOUT

ERROR_CALLBACK_ERROR

ERROR_TOO_MANY_MATCHES

int **yr_scanner_scan_fd** (YR_SCANNER* *scanner*, YR_FILE_DESCRIPTOR *fd*)

New in version 3.8.0.

Scan a file descriptor. In POSIX systems YR_FILE_DESCRIPTOR is an `int`, as returned by the *open()* function. In Windows YR_FILE_DESCRIPTOR is a `HANDLE` as returned by *CreateFile()*.

Returns one of the following error codes:

ERROR_SUCCESS

ERROR_INSUFFICIENT_MEMORY

ERROR_TOO_MANY_SCAN_THREADS

ERROR_SCAN_TIMEOUT

ERROR_CALLBACK_ERROR

ERROR_TOO_MANY_MATCHES

*YR_RULE** **yr_scanner_last_error_rule** (YR_SCANNER* *scanner*)

New in version 3.8.0.

Return a pointer to the YR_RULE which triggered a scanning error. In the case where the rule is unable to be determined, NULL is returned.

*YR_STRING** **yr_scanner_last_error_string** (YR_SCANNER* *scanner*)

New in version 3.8.0.

Return a pointer to the YR_STRING which triggered a scanning error.

7.6.3 Error codes

ERROR_SUCCESS

Everything went fine.

ERROR_INSUFFICIENT_MEMORY

Insufficient memory to complete the operation.

ERROR_COULD_NOT_OPEN_FILE

File could not be opened.

ERROR_COULD_NOT_MAP_FILE

File could not be mapped into memory.

ERROR_INVALID_FILE

File is not a valid rules file.

ERROR_CORRUPT_FILE

Rules file is corrupt.

ERROR_UNSUPPORTED_FILE_VERSION

File was generated by a different YARA and can't be loaded by this version.

ERROR_TOO_MANY_SCAN_THREADS

Too many threads trying to use the same *YR_RULES* object simultaneously. The limit is defined by *YR_MAX_THREADS* in *./include/yara/limits.h*

ERROR_SCAN_TIMEOUT

Scan timed out.

ERROR_CALLBACK_ERROR

Callback returned an error.

ERROR_TOO_MANY_MATCHES

Too many matches for some string in your rules. This usually happens when your rules contains very short or very common strings like 01 02 or FF FF FF FF. The limit is defined by *YR_MAX_STRING_MATCHES* in *./include/yara/limits.h*

ERROR_BLOCK_NOT_READY

Next memory block to scan is not ready; custom iterators may return this.

y

yara, [98](#)

Symbols

```

-fail-on-warnings
    yara command line option, 90
-max-process-memory-chunk=<size>
    yara command line option, 90
-max-strings-per-rule=<number>
    yara command line option, 90
-scan-list
    yara command line option, 91
-C -compiled-rules
    yara command line option, 90
-D -print-module-data
    yara command line option, 90
-L -print-string-length
    yara command line option, 90
-N -no-follow-symlinks
    yara command line option, 90
-S -print-stats
    yara command line option, 90
-a <seconds> -timeout=<seconds>
    yara command line option, 91
-c -count
    yara command line option, 90
-d <identifier>=<value>
    -define=identifier=value
    yara command line option, 90
-e -print-namespace
    yara command line option, 90
-f -fast-scan
    yara command line option, 90
-g -print-tags
    yara command line option, 90
-h -help
    yara command line option, 90
-i <identifier> -identifier=<identifier>
    yara command line option, 90
-k <slots> -stack-size=<slots>
    yara command line option, 91
-l <number> -max-rules=<number>
    yara command line option, 90
-m -print-meta
    yara command line option, 90
-n -negate
    yara command line option, 90
-p <number> -threads=<number>
    yara command line option, 91
-r -recursive
    yara command line option, 91
-s -print-strings
    yara command line option, 90
-t <tag> -tag=<tag>
    yara command line option, 91
-v -version
    yara command line option, 91
-w -no-warnings
    yara command line option, 90
-x <module>=<file>
    -module-data=<module>=<file>
    yara command line option, 90
-z <size> -skip-larger=<size>
    yara command line option, 91

```

A

```

abs (C function), 65
accessed (C type), 67
address (C member), 56
AGGRESSIVE_WS_TRIM (C type), 37
algorithm (C member), 47, 48
APPCONTAINER (C type), 39
assembly (C type), 69
assembly.culture (C member), 69
assembly.name (C member), 69
assembly.version (C member), 69
assembly_refs (C type), 69
assembly_refs.name (C member), 70
assembly_refs.public_key_or_token (C member), 70
assembly_refs.version (C member), 69
attributes: (C type), 67

```

B

base (*C type*), 82
base_of_code (*C type*), 36
base_of_data (*C type*), 36
bind (*C member*), 59
BYTES_REVERSED_HI (*C type*), 38
BYTES_REVERSED_LO (*C type*), 37

C

calculate_checksum (*C type*), 35
certificates (*C type*), 46
certificates.algorithm (*C member*), 47
certificates.issuer (*C member*), 47
certificates.not_after (*C member*), 47
certificates.not_before (*C member*), 47
certificates.serial (*C member*), 47
certificates.subject (*C member*), 47
certificates.thumbprint (*C member*), 46
certificates.version (*C member*), 47
chain (*C type*), 47, 48
characteristics (*C type*), 37
checksum (*C type*), 35
checksum32 (*C function*), 64
classes (*C type*), 67
classes.abstract (*C member*), 68
classes.abstract. (*C member*), 68
classes.are: (*C member*), 67, 68
classes.attributes: (*C member*), 68
classes.base_types (*C member*), 68
classes.by (*C member*), 68
classes.final (*C member*), 68
classes.final. (*C member*), 68
classes.fullname (*C member*), 67
classes.generic_parameters (*C member*), 68, 69
classes.methods (*C member*), 68
classes.methods. (*C member*), 68
classes.name (*C member*), 67, 68
classes.name. (*C member*), 67, 68
classes.namespace (*C member*), 67
classes.namespace. (*C member*), 67
classes.number_of_base_types (*C member*), 68
classes.number_of_generic_parameters (*C member*), 68, 69
classes.number_of_methods (*C member*), 68
classes.number_of_parameters (*C member*), 68
classes.operator. (*C member*), 68, 69
classes.parameters (*C member*), 68
classes.parameters. (*C member*), 68, 69
classes.return_type (*C member*), 68
classes.sealed (*C member*), 68
classes.sealed. (*C member*), 68

classes.static (*C member*), 68
classes.static. (*C member*), 68
classes.type (*C member*), 68, 69
classes.type. (*C member*), 69
classes.types. (*C member*), 68
classes.virtual (*C member*), 68
classes.virtual. (*C member*), 68
classes.visibility (*C member*), 67, 68
count (*C function*), 65
countersignatures (*C type*), 47
countersignatures.digest (*C member*), 47
countersignatures.digest_alg (*C member*), 47
countersignatures.length_of_chain (*C member*), 48
countersignatures.sign_time (*C member*), 47
countersignatures.verified (*C member*), 47
crc32 (*C function*), 64

D

data_directories (*C type*), 39
data_directories.size (*C member*), 39
data_directories.virtual_address (*C member*), 39
DEBUG_STRIPPED (*C type*), 37
delayed_import_details (*C type*), 52
delayed_import_details.functions (*C member*), 52
delayed_import_details.library_name (*C member*), 52
delayed_import_details.name (*C member*), 52
delayed_import_details.number_of_functions (*C member*), 52
delayed_import_details.ordinal (*C member*), 52
delayed_import_details.rva (*C member*), 52
delayed_import_rva (*C function*), 52, 53
deviation (*C function*), 65
DLL (*C type*), 37
dll_characteristics (*C type*), 38
dll_name (*C type*), 50
dns_lookup (*C function*), 61
DT_BIND_NOW (*C type*), 58
DT_DEBUG (*C type*), 58
DT_ENCODING (*C type*), 59
DT_FINI (*C type*), 58
DT_FINI_ARRAY (*C type*), 58
DT_FINI_ARRAYSZ (*C type*), 58
DT_FLAGS (*C type*), 59
DT_HASH (*C type*), 58
DT_INIT (*C type*), 58
DT_INIT_ARRAY (*C type*), 58
DT_INIT_ARRAYSZ (*C type*), 58
DT_JMPREL (*C type*), 58

DT_NEEDED (*C type*), 58
 DT_NULL (*C type*), 58
 DT_PLTGOT (*C type*), 58
 DT_PLTREL (*C type*), 58
 DT_PLTRELSZ (*C type*), 58
 DT_REL (*C type*), 58
 DT_RELA (*C type*), 58
 DT_RELAENT (*C type*), 58
 DT_RELASZ (*C type*), 58
 DT_RELENT (*C type*), 58
 DT_RELSZ (*C type*), 58
 DT_RPATH (*C type*), 58
 DT_RUNPATH (*C type*), 59
 DT_SONAME (*C type*), 58
 DT_STRSZ (*C type*), 58
 DT_STRTAB (*C type*), 58
 DT_SYMBOLIC (*C type*), 58
 DT_SYMENT (*C type*), 58
 DT_SYMTAB (*C type*), 58
 DT_TEXTREL (*C type*), 58
 dynamic (*C type*), 58
 dynamic.type (*C member*), 58
 DYNAMIC_BASE (*C type*), 38
 dynamic_section_entries (*C type*), 57

E

EM_386 (*C type*), 55
 EM_68K (*C type*), 55
 EM_860 (*C type*), 55
 EM_88K (*C type*), 55
 EM_AARCH64 (*C type*), 55
 EM_ARM (*C type*), 55
 EM_M32 (*C type*), 55
 EM_MIPS (*C type*), 55
 EM_MIPS_RS3_LE (*C type*), 55
 EM_NONE (*C type*), 55
 EM_PPC (*C type*), 55
 EM_PPC64 (*C type*), 55
 EM_SPARC (*C type*), 55
 EM_X86_64 (*C type*), 55
 entropy (*C function*), 64
 entry_point (*C type*), 36, 55
 entry_point_raw (*C type*), 36
 ERROR_BLOCK_NOT_READY (*C macro*), 115
 ERROR_CALLBACK_ERROR (*C macro*), 115
 ERROR_CORRUPT_FILE (*C macro*), 115
 ERROR_COULD_NOT_MAP_FILE (*C macro*), 114
 ERROR_COULD_NOT_OPEN_FILE (*C macro*), 114
 ERROR_INSUFFICIENT_MEMORY (*C macro*), 114
 ERROR_INVALID_FILE (*C macro*), 114
 ERROR_SCAN_TIMEOUT (*C macro*), 115
 ERROR_SUCCESS (*C macro*), 114
 ERROR_TOO_MANY_MATCHES (*C macro*), 115
 ERROR_TOO_MANY_SCAN_THREADS (*C macro*), 115

ERROR_UNSUPPORTED_FILE_VERSION (*C macro*), 115
 ET_CORE (*C type*), 55
 ET_DYN (*C type*), 55
 ET_EXEC (*C type*), 54
 ET_NONE (*C type*), 54
 ET_REL (*C type*), 54
 EXECUTABLE_IMAGE (*C type*), 37
 export_details (*C type*), 49
 export_details.forward_name (*C member*), 50
 export_details.name (*C member*), 50
 export_details.offset (*C member*), 49
 export_details.ordinal (*C member*), 50
 export_timestamp (*C type*), 50
 exports (*C function*), 49
 exports_index (*C function*), 49

F

fetch_data (*C type*), 82
 field_offsets (*C type*), 70
 file. (*C type*), 67
 file_access (*C function*), 62
 file_alignment (*C type*), 36
 filesystem (*C type*), 62
 flags (*C member*), 56
 FORCE_INTEGRITY (*C type*), 38

G

get_float (*C function*), 85
 get_integer (*C function*), 85
 get_object (*C function*), 85
 get_string (*C function*), 85
 GUARD_CF (*C type*), 39
 guids (*C type*), 67

H

hex (*C function*), 71
 HIGH_ENTROPY_VA (*C type*), 38
 host (*C function*), 61
 http_get (*C function*), 61
 http_post (*C function*), 61
 http_request (*C function*), 61
 http_user_agent (*C function*), 61

I

identifier (*yara.StringMatch attribute*), 100
 image_base (*C type*), 36
 IMAGE_DEBUG_TYPE_BORLAND (*C type*), 40
 IMAGE_DEBUG_TYPE_CLSID (*C type*), 40
 IMAGE_DEBUG_TYPE_CODEVIEW (*C type*), 40
 IMAGE_DEBUG_TYPE_COFF (*C type*), 40
 IMAGE_DEBUG_TYPE_EXCEPTION (*C type*), 40
 IMAGE_DEBUG_TYPE_FIXUP (*C type*), 40

IMAGE_DEBUG_TYPE_FPO (C type), 40
IMAGE_DEBUG_TYPE_ILTCG (C type), 40
IMAGE_DEBUG_TYPE_MISC (C type), 40
IMAGE_DEBUG_TYPE_MPX (C type), 40
IMAGE_DEBUG_TYPE_OMAP_FROM_SRC (C type), 40
IMAGE_DEBUG_TYPE_OMAP_TO_SRC (C type), 40
IMAGE_DEBUG_TYPE_POGO (C type), 40
IMAGE_DEBUG_TYPE_REPRO (C type), 40
IMAGE_DEBUG_TYPE_RESERVED10 (C type), 40
IMAGE_DEBUG_TYPE_UNKNOWN (C type), 40
IMAGE_DEBUG_TYPE_VC_FEATURE (C type), 40
IMAGE_DIRECTORY_ENTRY_ARCHITECTURE (C type), 40
IMAGE_DIRECTORY_ENTRY_BASERELOC (C type), 40
IMAGE_DIRECTORY_ENTRY_BOUND_IMPORT (C type), 40
IMAGE_DIRECTORY_ENTRY_COM_DESCRIPTOR (C type), 41
IMAGE_DIRECTORY_ENTRY_COPYRIGHT (C type), 40
IMAGE_DIRECTORY_ENTRY_DEBUG (C type), 40
IMAGE_DIRECTORY_ENTRY_DELAY_IMPORT (C type), 41
IMAGE_DIRECTORY_ENTRY_EXCEPTION (C type), 40
IMAGE_DIRECTORY_ENTRY_EXPORT (C type), 39
IMAGE_DIRECTORY_ENTRY_IAT (C type), 41
IMAGE_DIRECTORY_ENTRY_IMPORT (C type), 40
IMAGE_DIRECTORY_ENTRY_LOAD_CONFIG (C type), 40
IMAGE_DIRECTORY_ENTRY_RESOURCE (C type), 40
IMAGE_DIRECTORY_ENTRY_SECURITY (C type), 40
IMAGE_DIRECTORY_ENTRY_TLS (C type), 40
IMAGE_NT_OPTIONAL_HDR32_MAGIC (C type), 36
IMAGE_NT_OPTIONAL_HDR64_MAGIC (C type), 36
IMAGE_ROM_OPTIONAL_HDR_MAGIC (C type), 36
image_version (C type), 38
image_version.major (C member), 38
image_version.minor (C member), 38
imphash (C function), 53
import_details (C type), 51
import_details.functions (C member), 52
import_details.library_name (C member), 51
import_details.name (C member), 52
import_details.number_of_functions (C member), 52
import_details.ordinal (C member), 52
import_details.rva (C member), 52
import_md5 (C function), 59
import_rva (C function), 52
imports (C function), 50, 51
in_range (C function), 65
instances (yara.StringMatch attribute), 100
is_32bit (C function), 54
is_64bit (C function), 54
is_dll (C function), 53
is_dotnet (C type), 70
is_pe (C type), 53
is_signed (C type), 45
is_xor () (yara.StringMatch method), 100
issuer (C member), 47, 48

K

key_access (C function), 62

L

language (C function), 53
LARGE_ADDRESS_AWARE (C type), 37
length (C function), 72
LINE_NUMS_STRIPPED (C type), 37
linker_version (C type), 38
linker_version.major (C member), 38
linker_version.minor (C member), 38
loader_flags (C type), 39
LOCAL_SYMS_STRIPPED (C type), 37
locale (C function), 53
log (C function), 71

M

machine (C type), 34, 55
MACHINE_32BIT (C type), 37
MACHINE_ALPHA (C type), 34
MACHINE_ALPHA64 (C type), 34
MACHINE_AM33 (C type), 34
MACHINE_AMD64 (C type), 34
MACHINE_ARM (C type), 34
MACHINE_ARM64 (C type), 34
MACHINE_ARMNT (C type), 34
MACHINE_AXP64 (C type), 35
MACHINE_CEE (C type), 35
MACHINE_CEF (C type), 35
MACHINE_EBC (C type), 34
MACHINE_I386 (C type), 34
MACHINE_IA64 (C type), 34
MACHINE_M32R (C type), 34
MACHINE_MIPS16 (C type), 34
MACHINE_MIPSFPU (C type), 34
MACHINE_MIPSFPU16 (C type), 34
MACHINE_POWERPC (C type), 34
MACHINE_POWERPCFP (C type), 34
MACHINE_R10000 (C type), 34
MACHINE_R3000 (C type), 34
MACHINE_R4000 (C type), 34
MACHINE_SH3 (C type), 34
MACHINE_SH3DSP (C type), 34
MACHINE_SH3E (C type), 34
MACHINE_SH4 (C type), 34

MACHINE_SH5 (C type), 34
 MACHINE_TARGET_HOST (C type), 34
 MACHINE_THUMB (C type), 34
 MACHINE_TRICORE (C type), 35
 MACHINE_UNKNOWN (C type), 34
 MACHINE_WCEMIPSV2 (C type), 34
 Match (class in yara), 100
 match() (yara.Rules method), 99
 matched_data (yara.StringMatchInstance attribute), 100
 matched_length (yara.StringMatchInstance attribute), 100
 max (C function), 65
 md5 (C function), 63
 mean (C function), 65
 memory_size (C member), 57
 meta (yara.Match attribute), 100
 mime_type (C function), 63
 min (C function), 65
 mode (C function), 66
 module_name (C type), 67
 modulerefs (C type), 69
 monte_carlo_pi (C function), 64
 mutex (C function), 62

N

namespace (yara.Match attribute), 100
 NET_RUN_FROM_SWAP (C type), 37
 network (C type), 61
 NO_BIND (C type), 39
 NO_ISOLATION (C type), 38
 NO_SEH (C type), 38
 not_after (C member), 47, 48
 not_before (C member), 47, 48
 now (C function), 70
 number_of_classes (C type), 67
 number_of_countersignatures (C member), 47
 number_of_delay_imported_functions (C type), 50
 number_of_delayed_imports (C type), 50
 number_of_exports (C type), 49
 number_of_field_offsets (C type), 70
 number_of_guids (C type), 67
 number_of_imported_functions (C type), 50
 number_of_imports (C type), 50
 number_of_modulerefs (C type), 69
 number_of_resources (C type), 43, 69
 number_of_rva_and_sizes (C type), 39
 number_of_sections (C type), 41, 55
 number_of_segments (C type), 56
 number_of_signatures (C type), 45
 number_of_streams (C type), 67
 number_of_symbols (C type), 35
 number_of_user_strings (C type), 70

NX_COMPAT (C type), 38

O

offset (C member), 57
 offset (yara.StringMatchInstance attribute), 100
 opthdr_magic (C type), 36
 os_version (C type), 38
 os_version.major (C member), 38
 os_version.minor (C member), 38
 overlay (C type), 43
 overlay.offset (C member), 43
 overlay.size (C member), 43

P

pdb_path (C type), 49
 pe.IMPORT_ANY (C member), 51
 pe.IMPORT_DELAYED (C member), 51
 pe.IMPORT_STANDARD (C member), 51
 percentage (C function), 66
 PF_R (C type), 57
 PF_W (C type), 57
 PF_X (C type), 57
 physical_address (C member), 57
 plaintext() (yara.StringMatchInstance method), 100
 pointer_to_symbol_table (C type), 35
 PT_DYNAMIC (C type), 57
 PT_GNU_STACK (C type), 57
 PT_HIPOC (C type), 57
 PT_INTERP (C type), 57
 PT_LOAD (C type), 57
 PT_LOPROC (C type), 57
 PT_NOTE (C type), 57
 PT_NULL (C type), 57
 PT_PHDR (C type), 57
 PT_SHLIB (C type), 57

R

registry (C type), 62
 RELOCS_STRIPPED (C type), 37
 REMOVABLE_RUN_FROM_SWAP (C type), 37
 resource_timestamp (C type), 43
 RESOURCE_TYPE_ACCELERATOR (C type), 44
 RESOURCE_TYPE_ANICURSOR (C type), 44
 RESOURCE_TYPE_ANIICON (C type), 44
 RESOURCE_TYPE_BITMAP (C type), 44
 RESOURCE_TYPE_CURSOR (C type), 44
 RESOURCE_TYPE_DIALOG (C type), 44
 RESOURCE_TYPE_DLGINCLUDE (C type), 44
 RESOURCE_TYPE_FONT (C type), 44
 RESOURCE_TYPE_FONTDIR (C type), 44
 RESOURCE_TYPE_GROUP_CURSOR (C type), 44
 RESOURCE_TYPE_GROUP_ICON (C type), 44
 RESOURCE_TYPE_HTML (C type), 44

RESOURCE_TYPE_ICON (*C type*), 44
RESOURCE_TYPE_MANIFEST (*C type*), 44
RESOURCE_TYPE_MENU (*C type*), 44
RESOURCE_TYPE_MESSAGE_TABLE (*C type*), 44
RESOURCE_TYPE_PLUGPLAY (*C type*), 44
RESOURCE_TYPE_RC_DATA (*C type*), 44
RESOURCE_TYPE_STRING (*C type*), 44
RESOURCE_TYPE_VERSION (*C type*), 44
RESOURCE_TYPE_VXD (*C type*), 44
resource_version (*C type*), 43
resource_version.major (*C member*), 43
resource_version.minor (*C member*), 43
resources (*C type*), 43, 69
resources.id (*C member*), 43
resources.language (*C member*), 44
resources.language_string (*C member*), 44
resources.length (*C member*), 43, 69
resources.name (*C member*), 69
resources.name_string (*C member*), 44
resources.offset (*C member*), 43, 69
resources.rva (*C member*), 43
resources.type (*C member*), 43
resources.type_string (*C member*), 44
rich_signature (*C type*), 48
rich_signature.clear_data (*C member*), 48
rich_signature.key (*C member*), 48
rich_signature.length (*C member*), 48
rich_signature.offset (*C member*), 48
rich_signature.raw_data (*C member*), 48
rich_signature.version_data (*C member*), 48
rule (*yara.Match attribute*), 100
Rules (*class in yara*), 99
rva_to_offset (*C function*), 54

S

save () (*yara.Rules method*), 99
SECTION_ALIGN_1024BYTES (*C type*), 42
SECTION_ALIGN_128BYTES (*C type*), 42
SECTION_ALIGN_16BYTES (*C type*), 42
SECTION_ALIGN_1BYTES (*C type*), 42
SECTION_ALIGN_2048BYTES (*C type*), 42
SECTION_ALIGN_256BYTES (*C type*), 42
SECTION_ALIGN_2BYTES (*C type*), 42
SECTION_ALIGN_32BYTES (*C type*), 42
SECTION_ALIGN_4096BYTES (*C type*), 42
SECTION_ALIGN_4BYTES (*C type*), 42
SECTION_ALIGN_512BYTES (*C type*), 42
SECTION_ALIGN_64BYTES (*C type*), 42
SECTION_ALIGN_8192BYTES (*C type*), 42
SECTION_ALIGN_8BYTES (*C type*), 42
SECTION_ALIGN_MASK (*C type*), 42
section_alignment (*C type*), 36
SECTION_CNT_CODE (*C type*), 42
SECTION_CNT_INITIALIZED_DATA (*C type*), 42

SECTION_CNT_UNINITIALIZED_DATA (*C type*), 42
SECTION_GPREL (*C type*), 42
section_index (*C function*), 53
SECTION_LNK_COMDAT (*C type*), 42
SECTION_LNK_INFO (*C type*), 42
SECTION_LNK_NRELOC_OVFL (*C type*), 42
SECTION_LNK_OTHER (*C type*), 42
SECTION_LNK_REMOVE (*C type*), 42
SECTION_MEM_16BIT (*C type*), 42
SECTION_MEM_DISCARDABLE (*C type*), 43
SECTION_MEM_EXECUTE (*C type*), 43
SECTION_MEM_FARDATA (*C type*), 42
SECTION_MEM_LOCKED (*C type*), 42
SECTION_MEM_NOT_CACHED (*C type*), 43
SECTION_MEM_NOT_PAGED (*C type*), 43
SECTION_MEM_PRELOAD (*C type*), 42
SECTION_MEM_PURGEABLE (*C type*), 42
SECTION_MEM_READ (*C type*), 43
SECTION_MEM_SHARED (*C type*), 43
SECTION_MEM_WRITE (*C type*), 43
SECTION_NO_DEFER_SPEC_EXC (*C type*), 42
SECTION_NO_PAD (*C type*), 42
SECTION_SCALE_INDEX (*C type*), 43
sections (*C type*), 41, 55
sections.characteristics (*C member*), 41
sections.full_name (*C member*), 41
sections.name (*C member*), 41, 55
sections.number_of_line_numbers (*C member*), 42
sections.number_of_relocations (*C member*), 41
sections.offset (*C member*), 55
sections.pointer_to_line_numbers (*C member*), 41
sections.pointer_to_relocations (*C member*), 41
sections.raw_data_offset (*C member*), 41
sections.raw_data_size (*C member*), 41
sections.size (*C member*), 55
sections.type (*C member*), 56
sections.virtual_address (*C member*), 41
sections.virtual_size (*C member*), 41
segments (*C type*), 57
segments.alignment (*C member*), 57
segments.file_size (*C member*), 57
segments.flags (*C member*), 57
serial (*C member*), 47, 48
serial_correlation (*C function*), 64, 65
set_float (*C function*), 83
set_integer (*C function*), 83
set_string (*C function*), 83
shal (*C function*), 63
sha256 (*C function*), 63, 64
SHF_ALLOC (*C type*), 56

- SHF_EXECINSTR (*C type*), 56
 - SHF_WRITE (*C type*), 56
 - shndx (*C member*), 59
 - SHT_DYNAMIC (*C type*), 56
 - SHT_DYNSYM (*C type*), 56
 - SHT_HASH (*C type*), 56
 - SHT_NOBITS (*C type*), 56
 - SHT_NOTE (*C type*), 56
 - SHT_NULL (*C type*), 56
 - SHT_PROGBITS (*C type*), 56
 - SHT_REL (*C type*), 56
 - SHT_RELA (*C type*), 56
 - SHT_SHLIB (*C type*), 56
 - SHT_STRTAB (*C type*), 56
 - SHT_SYMTAB (*C type*), 56
 - signatures (*C type*), 45
 - signatures.algorithm (*C member*), 45
 - signatures.algorithm_oid (*C member*), 46
 - signatures.digest (*C member*), 46
 - signatures.digest_alg (*C member*), 46
 - signatures.file_digest (*C member*), 46
 - signatures.issuer (*C member*), 45
 - signatures.not_after (*C member*), 46
 - signatures.not_before (*C member*), 46
 - signatures.number_of_certificates (*C member*), 46
 - signatures.serial (*C member*), 46
 - signatures.subject (*C member*), 45
 - signatures.thumbprint (*C member*), 45
 - signatures.valid_on (*C member*), 46
 - signatures.verified (*C member*), 46
 - signatures.version (*C member*), 45
 - signer_info (*C type*), 47
 - signer_info.digest (*C member*), 47
 - signer_info.digest_alg (*C member*), 47
 - signer_info.length_of_chain (*C member*), 47
 - signer_info.program_name (*C member*), 47
 - size (*C type*), 82
 - size_of_code (*C type*), 36
 - size_of_headers (*C type*), 37
 - size_of_heap_commit (*C type*), 39
 - size_of_heap_reserve (*C type*), 39
 - size_of_image (*C type*), 37
 - size_of_initialized_data (*C type*), 36
 - size_of_optional_header (*C type*), 36
 - size_of_stack_commit (*C type*), 39
 - size_of_stack_reserve (*C type*), 39
 - size_of_uninitialized_data (*C type*), 36
 - SIZED_STRING (*C type*), 87
 - SIZED_STRING.c_string (*C member*), 87
 - SIZED_STRING.length (*C member*), 87
 - STB_GLOBAL (*C type*), 59
 - STB_LOCAL (*C type*), 59
 - STB_WEAK (*C type*), 59
 - streams (*C type*), 67
 - streams.name (*C member*), 67
 - streams.offset (*C member*), 67
 - streams.size (*C member*), 67
 - StringMatch (*class in yara*), 100
 - StringMatchInstance (*class in yara*), 100
 - strings (*yara.Match attribute*), 100
 - STT_COMMON (*C type*), 59
 - STT_FILE (*C type*), 59
 - STT_FUNC (*C type*), 59
 - STT_NOTYPE (*C type*), 59
 - STT_OBJECT (*C type*), 59
 - STT_SECTION (*C type*), 59
 - STT_TLS (*C type*), 59
 - subject (*C member*), 47, 48
 - subsystem (*C type*), 35
 - SUBSYSTEM_EFI_APPLICATION (*C type*), 35
 - SUBSYSTEM_EFI_BOOT_SERVICE_DRIVER (*C type*), 35
 - SUBSYSTEM_EFI_ROM_IMAGE (*C type*), 35
 - SUBSYSTEM_EFI_RUNTIME_DRIVER (*C type*), 35
 - SUBSYSTEM_NATIVE (*C type*), 35
 - SUBSYSTEM_NATIVE_WINDOWS (*C type*), 35
 - SUBSYSTEM_OS2_CUI (*C type*), 35
 - SUBSYSTEM_POSIX_CUI (*C type*), 35
 - SUBSYSTEM_UNKNOWN (*C type*), 35
 - subsystem_version (*C type*), 38
 - subsystem_version.major (*C member*), 38
 - subsystem_version.minor (*C member*), 38
 - SUBSYSTEM_WINDOWS_BOOT_APPLICATION (*C type*), 35
 - SUBSYSTEM_WINDOWS_CE_GUI (*C type*), 35
 - SUBSYSTEM_WINDOWS_CUI (*C type*), 35
 - SUBSYSTEM_WINDOWS_GUI (*C type*), 35
 - SUBSYSTEM_XBOX (*C type*), 35
 - symtab (*C type*), 59
 - symtab.name (*C member*), 59
 - symtab.size (*C member*), 59
 - symtab.type (*C member*), 59
 - symtab.value (*C member*), 59
 - symtab_entries (*C type*), 59
 - sync (*C type*), 62
 - SYSTEM (*C type*), 37
- ## T
- tags (*yara.Match attribute*), 100
 - tcp (*C function*), 61
 - telhash (*C function*), 59
 - TERMINAL_SERVER_AWARE (*C type*), 39
 - thumbprint (*C member*), 47, 48
 - timestamp (*C type*), 35
 - to_int (*C function*), 71
 - to_number (*C function*), 65
 - to_string (*C function*), 66

toolid (C function), 48
type (C function), 63
type (C member), 57
type (C type), 54
typelib (C type), 69

U

udp (C function), 61
UP_SYSTEM_ONLY (C type), 38
user_strings (C type), 70

V

value (C member), 59
version (C function), 48
version (C member), 47, 48
version (C type), 67
version_info (C type), 44
version_info_list (C type), 45
version_info_list.key (C member), 45
version_info_list.value (C member), 45
virtual_address (C member), 57

W

WDM_DRIVER (C type), 39
win32_version_value (C type), 37

X

xor_key (yara.StringMatchInstance attribute), 100

Y

yara (module), 98
yara command line option
-fail-on-warnings, 90
-max-process-memory-chunk=<size>, 90
-max-strings-per-rule=<number>, 90
-scan-list, 91
-C -compiled-rules, 90
-D -print-module-data, 90
-L -print-string-length, 90
-N -no-follow-symlinks, 90
-S -print-stats, 90
-a <seconds> -timeout=<seconds>, 91
-c -count, 90
-d <identifier>=<value>
-define=identifier=value, 90
-e -print-namespaces, 90
-f -fast-scan, 90
-g -print-tags, 90
-h -help, 90
-i <identifier>
-identifier=<identifier>, 90
-k <slots> -stack-size=<slots>, 91
-l <number> -max-rules=<number>, 90

-m -print-meta, 90
-n -negate, 90
-p <number> -threads=<number>, 91
-r -recursive, 91
-s -print-strings, 90
-t <tag> -tag=<tag>, 91
-v -version, 91
-w -no-warnings, 90
-x <module>=<file>
-module-data=<module>=<file>, 90
-z <size> -skip-larger=<size>, 91
yara.compile() (in module yara), 98
yara.load() (in module yara), 98
yara.set_config() (in module yara), 98
YR_COMPILER (C type), 106
yr_compiler_add_fd (C function), 108
yr_compiler_add_file (C function), 108
yr_compiler_add_string (C function), 108
yr_compiler_create (C function), 108
yr_compiler_define_boolean_variable (C function), 109
yr_compiler_define_float_variable (C function), 109
yr_compiler_define_integer_variable (C function), 109
yr_compiler_define_string_variable (C function), 109
yr_compiler_destroy (C function), 108
yr_compiler_get_rules (C function), 108
yr_compiler_set_callback (C function), 108
yr_compiler_set_include_callback (C function), 108
yr_finalize (C function), 108
yr_initialize (C function), 108
YR_MATCH (C type), 106
YR_MATCH.base (C member), 106
YR_MATCH.data (C member), 106
YR_MATCH.data_length (C member), 106
YR_MATCH.match_length (C member), 106
YR_MATCH.offset (C member), 106
YR_META (C type), 106
YR_META.identifier (C member), 106
YR_META.type (C member), 106
YR_MODULE_IMPORT (C type), 106
YR_MODULE_IMPORT.module_data (C member), 107
YR_MODULE_IMPORT.module_data_size (C member), 107
YR_MODULE_IMPORT.module_name (C member), 106
YR_NAMESPACE (C type), 107
YR_NAMESPACE.name (C member), 107
YR_RULE (C type), 107

YR_RULE.identifier (*C member*), 107
 YR_RULE.metas (*C member*), 107
 YR_RULE.ns (*C member*), 107
 YR_RULE.strings (*C member*), 107
 YR_RULE.tags (*C member*), 107
 yr_rule_disable (*C function*), 112
 yr_rule_enable (*C function*), 112
 yr_rule_metas_foreach (*C function*), 111
 yr_rule_strings_foreach (*C function*), 111
 yr_rule_tags_foreach (*C function*), 111
 YR_RULES (*C type*), 107
 yr_rules_define_boolean_variable (*C function*), 109
 yr_rules_define_float_variable (*C function*), 109
 yr_rules_define_integer_variable (*C function*), 109
 yr_rules_define_string_variable (*C function*), 109
 yr_rules_destroy (*C function*), 109
 yr_rules_foreach (*C function*), 111
 yr_rules_load (*C function*), 109
 yr_rules_load_stream (*C function*), 109
 yr_rules_save (*C function*), 109
 yr_rules_save_stream (*C function*), 109
 yr_rules_scan_fd (*C function*), 110
 yr_rules_scan_file (*C function*), 110
 yr_rules_scan_mem (*C function*), 110
 YR_SCAN_CONTEXT (*C type*), 106
 yr_scanner_create (*C function*), 112
 yr_scanner_define_boolean_variable (*C function*), 112
 yr_scanner_define_float_variable (*C function*), 113
 yr_scanner_define_integer_variable (*C function*), 112
 yr_scanner_define_string_variable (*C function*), 113
 yr_scanner_destroy (*C function*), 112
 yr_scanner_last_error_rule (*C function*), 114
 yr_scanner_last_error_string (*C function*), 114
 yr_scanner_scan_fd (*C function*), 114
 yr_scanner_scan_file (*C function*), 114
 yr_scanner_scan_mem (*C function*), 113
 yr_scanner_scan_mem_blocks (*C function*), 113
 yr_scanner_set_callback (*C function*), 112
 yr_scanner_set_flags (*C function*), 112
 yr_scanner_set_timeout (*C function*), 112
 YR_STREAM (*C type*), 107
 YR_STREAM.read (*C member*), 107
 YR_STREAM.user_data (*C member*), 107
 YR_STREAM.write (*C member*), 107
 YR_STRING (*C type*), 107
 YR_STRING.identifier (*C member*), 107
 yr_string_matches_foreach (*C function*), 111