



# YARA User's Manual Ver. 2.0

Víctor Manuel Álvarez  
[vmalvarez@virustotal.com](mailto:vmalvarez@virustotal.com)

<b>YARA in a nutshell</b>	<b>3</b>
<b>Writing rules</b>	<b>3</b>
<i>2.1 Comments</i>	<i>5</i>
<b>Strings</b>	<b>5</b>
<i>3.1 Hexadecimal strings</i>	<i>5</i>
<i>3.2 Text strings</i>	<i>7</i>
<i>3.2.1 Case-insensitive strings</i>	<i>8</i>
<i>3.2.2 Wide-character strings</i>	<i>8</i>
<i>3.2.3 Searching for full words</i>	<i>9</i>
<i>3.3 Regular expressions</i>	<i>9</i>
<b>Conditions</b>	<b>11</b>
<i>4.1 Counting strings</i>	<i>12</i>
<i>4.3 File size</i>	<i>13</i>
<i>4.4 Executable entry point</i>	<i>14</i>
<i>4.5 Accessing data at a given position</i>	<i>14</i>
<i>4.6 Sets of strings</i>	<i>15</i>
<i>4.7 Applying the same condition to many strings</i>	<i>16</i>
<i>4.8 Using anonymous strings with "of" and "for..of"</i>	<i>17</i>
<i>4.9 Iterating over string occurrences</i>	<i>17</i>
<i>4.9 Referencing other rules</i>	<i>18</i>
<b>More about rules</b>	<b>19</b>
<i>5.1 Private rules</i>	<i>19</i>
<i>5.3 Rule tags</i>	<i>20</i>
<i>5.4 Metadata</i>	<i>20</i>
<i>5.5 External variables</i>	<i>21</i>
<b>Includes</b>	<b>22</b>
<b>Using YARA from command-line</b>	<b>23</b>
<b>Using YARA from Python</b>	<b>24</b>

# 1. YARA in a nutshell

YARA is a tool aimed at helping malware researchers to identify and classify malware families. With YARA you can create descriptions of malware families based on textual or binary information contained on samples of those families. These descriptions, a.k.a rules, consist patterns and a boolean expression which determines its logic. Rules can be applied to files or running processes in order to determine if it belongs to the described malware family.

Let's explain it with an example. Suppose that we have a malware family with two variants, one of them downloads a malicious file from *http://foo.com/badfile1.exe*, the other downloads a file from *http://bar.com/badfile2.exe*, the URLs are hardcoded into the malware code. Both variants drops the downloaded file with the name *win.exe*, which also appears hardcoded into the samples. For this hypothetical family we can create a rule like this:

```
rule BadBoy
{
    strings:
        $a = "win.exe"
        $b = "http://foo.com/badfile1.exe"
        $c = "http://bar.com/badfile2.exe"

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

The rule above instructs YARA that those files or processes containing the string *win.exe* and any of the two URLs must be reported as BadBoy.

This is just a simple example, but more complex and powerful rules can be created by using binary strings with wild-cards, case-insensitive text strings, regular expressions, and many other features provided by YARA that will be covered in this manual.

## 2. Writing rules

YARA rules are easy to write and understand, and they have a syntax that resembles in some way a C struct declaration. 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 can not 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:

<code>all</code>	<code>in</code>	<code>private</code>
<code>and</code>	<code>include</code>	<code>rule</code>
<code>any</code>	<code>index</code>	<code>rva</code>
<code>ascii</code>	<code>indexes</code>	<code>section</code>
<code>at</code>	<code>int8</code>	<code>strings</code>
<code>condition</code>	<code>int16</code>	<code>them</code>
<code>contains</code>	<code>int32</code>	<code>true</code>
<code>entrypoint</code>	<code>matches</code>	<code>uint8</code>
<code>false</code>	<code>meta</code>	<code>uint16</code>
<code>filesize</code>	<code>nocase</code>	<code>uint32</code>
<code>fullword</code>	<code>not</code>	<code>wide</code>
<code>for</code>	<code>or</code>	
<code>global</code>	<code>of</code>	

Rules are generally composed of two sections: strings definition and condition, although the strings definition section can be omitted if the rule doesn't rely on any string. 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 in 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 on 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 the string identifier. 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 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 an dummy rule, don't do this  
}
```

## 3. 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.

### 3.1 Hexadecimal strings

Hexadecimal strings allow three special constructions that make them more flexible: wild-cards, 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  
{  
    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.

Wild-cards 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.

### 3.2 Text strings

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

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

    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>\t</code>	Horizontal tab
<code>\xdd</code>	Any byte in hexadecimal notation

### 3.2.1 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 other modifier.

### 3.2.2 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.

00 00 00 00	00 00 00 AA	02 00 00 00	00 53 00 74	00 72 00 69	00 6E 00 67	00 46 00 69	00	.....S.t.r.i.n.g.F.i.
6C 00 65 00	49 00 6E 00	66 00 6F 00	00 00 86 02	00 00 00 00	30 00 34 00	30 00 39 00	30	l.e.I.n.f.o.....0.4.0.9.0
00 34 00 45	00 34 00 00	00 5C 00 1E	00 01 00 43	00 6F 00 6D	00 70 00 61	00 6E 00 79	00	.4.E.4...\.....C.o.m.p.a.n.y.
4E 00 61 00	6D 00 65 00	00 00 00 00	42 00 6F 00	72 00 6C 00	61 00 6E 00	64 00 20 00	53	N.a.m.e....B.o.r.l.a.n.d..S
00 6F 00 66	00 74 00 77	00 61 00 72	00 65 00 20	00 43 00 6F	00 72 00 70	00 6F 00 72	00	.o.f.t.w.a.r.e..C.o.r.p.o.r.
61 00 74 00	69 00 6F 00	6E 00 00 00	00 00 5E 00	1B 00 01 00	46 00 69 00	6C 00 65 00	44	a.t.i.o.n....^.....F.i.l.e.D
00 65 00 73	00 63 00 72	00 69 00 70	00 74 00 69	00 6F 00 6E	00 00 00 00	00 42 00 6F	00	.e.s.c.r.i.p.t.i.o.n....B.o.
72 00 6C 00	61 00 6E 00	64 00 20 00	43 00 6F 00	6D 00 70 00	6F 00 6E 00	65 00 6E 00	74	r.l.a.n.d..C.o.m.p.o.n.e.n.t
00 20 00 50	00 61 00 63	00 68 00 61	00 67 00 65	00 00 00 00	00 00 00 34	00 0A 00 01	00	. .P.a.c.k.a.g.e.....4.....
46 00 69 00	6C 00 65 00	56 00 65 00	72 00 73 00	69 00 6F 00	6E 00 00 00	00 00 36 00	2E	F.i.l.e.V.e.r.s.i.o.n.....6..

In the above figure the string "Borland" appears encoded as two bytes per character, therefore the following rule will match:



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

    condition:
        $wide_and_ascii_string
}
```

The `ascii` modifier can appear along, 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.

### 3.2.3 Searching for full words

Another modifier that can be applied to text strings is `fullword`. This modifier guarantee 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`, don't matches "www.mydomain.com" but it matches "www.my-domain.com" and "www.domain.com".

## 3.3 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 backslashes instead of double-quotes, like in the Perl programming language.

```
rule RegExpExample1
{
    strings:
        $re1 = /md5: [0-9a-zA-Z]{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.

In previous versions of YARA externals 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 and backreferences.

YARA's regular expressions recognize the following metacharacters:

\	Quote the next metacharacter
^	Match the beginning of the file
\$	Match the end of the file
	Alternation
()	Grouping
[]	Bracketed character class

The following quantifiers are recognized 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 0 to 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 0 to m times, non-greedy
{n,m}?	Match n to m times, non-greedy

The following escape sequences are recognized:

\t	Tab (HT, TAB)
\n	New line (LF, NL)
\r	Return (CR)
\f	Form feed (FF)
\a	Alarm bell (BEL)
\x00	Character whose ordinal number is the given hexadecimal number

These are the recognized character classes:

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

## 4. 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.

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)
}

```

## 4.1 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`.

## 4.2 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:
        $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.

### 4.3 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 especial variables is `filesize`, which holds, as its name indicates, the size of the file being analyzed. The size is expressed in bytes.

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

The previous example also demonstrate 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 won't never match because `filesize` doesn't make sense in this context.

#### 4.4 Executable entry point

Another special variable than can be used on a rule is `entrypoint`. If 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 scanning a file. If we are scanning a running process `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.

#### 4.5 Accessing data at a given position

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

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 bits integer are considered to be little-endian. 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
}
```

## 4.6 Sets of strings

There are circumstances in which 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 operator `of` come into help.

```
rule OfExample1
{
    strings:
        $a = "dummy1"
        $b = "dummy2"
        $c = "dummy3"

    condition:
        2 of ($a,$b,$c)
}
```

What this rule says is that at least two of the strings in the set (`$a`, `$b`, `$c`) must be present on the file, no matter which. Of course, when using this operator, the number before the `of` keyword must be equal to or less than 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) */
}
```

```

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 above examples the number of strings have been specified by a numeric constant, but any expression returning a numeric value can be used. The keywords *any* and *all* 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" */

```

## 4.7 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 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 to the following expression:

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

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 realized that the *of* operator is an 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 and the first offset of each string respectively.

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

## 4.8 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 don't need to provide a unique identifier for each of them. In those situations you can declare anonymous strings with identifiers consisting only in the \$ character, as in the following example:

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

    condition:
        1 of them
}
```

## 4.9 Iterating over string occurrences

As seen in section 4.2, 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 are you referring to. (*@a[1]*, *@a[2]*,...).

Sometimes you will need to iterate over some of these offsets and guarantee they satisfy a given condition. 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 tells that the first three occurrences of `$b` should be 10 bytes away from the first three occurrences of `$a`.

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 telling 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 )

```

Resuming, the syntax of this operator is:

```

for expression identifier in indexes : ( boolean_expression )

```

## 4.9 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 depends 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 satisfy Rule1. Note that is strictly necessary to define the rule being invoked before the one that will make the invocation.

## 5. More about rules

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

### 5.1 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 ignoring those files that exceed certain size limit, you could go rule by rule doing 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.

### 5.1 Private rules

Private rules are a very simple concept. That 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 section 4.5) 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. For declaring 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 a global rule that does not get reported by YARA but must be satisfied.

### 5.3 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 interesting 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 underscores are allowed, and the tag cannot start with a digit. They are also case sensitive.

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

### 5.4 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:
        $my_text_string or $my_hex_string
}

```

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

## 5.5 External variables

External variables allow you to define rules which depends 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 strings can be used with operators `contains` and `matches`. The `contains` operator returns true if the string contains the specified substring. The operator `matches` returns true if the string matches the given regular expression.

```

rule ExternalVariableExample3
{
    condition:
        string_ext_var contains "text"
}

rule ExternalVariableExample4
{
    condition:
        string_ext_var matches /[a-z]+/
}

```

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.

## 6. Includes

In order to allow you a 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 your 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 that 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 ones:

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

And you can also use absolute paths:

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

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

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

## 7. Using YARA from command-line

In order to invoke YARA you'll need two things: a file with the rules you want to use (either in source code or compiled form) and the target to be scanned. The target can be a file, a folder, or a process.

```
usage: yara [OPTION]... RULES_FILE FILE | PID
options:
  -t <tag>           print rules tagged as <tag> and ignore the rest.
  -i <identifier>     print rules named <identifier> and ignore the rest.
  -n                 print only not satisfied rules (negate).
  -g                 print tags.
  -m                 print metadata.
  -s                 print matching strings.
  -p <number>         use the specified <number> of threads to scan a directory.
  -l <number>         abort scanning after matching a number of rules.
  -a <seconds>        abort scanning after a number of seconds has elapsed.
  -d <identifier>=<value> define external variable.
  -r                 recursively search directories.
  -f                 fast matching mode.
  -w                 disable warnings.
  -v                 show version information.
```

Rule files can be passed directly in source code form, or can be previously compiled with the `yarac` tool. You may prefer to use your rules in compiled form if you are going to invoke YARA multiple times with the same rules. This way you'll save time, because for YARA is faster to load compiled rules than compiling the same rules over and over again.

The 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 to do it.

The `-t` option allows you to specify one or more tags that will act as filters to YARA's output. If you use this option, only those rules tagged as you specified will be shown. The `-i` option has a similar behavior, filtering all rules except the one having the given identifier. You can also use the `-n` modifier to print those rules that are not satisfied by the files.

The `-l` argument allows to stop applying rules after a given number of them has already matched. The `-a` argument stops the scanning after the given amount of seconds has elapsed.

The `-d` is used to define external variables. For example:

```
-d flag=true
-d beast=666
-d name="James Bond"
```

## 8. Using YARA from Python

YARA can be also invoked from your own Python scripts. The `yara-python` extension is provided in order to make YARA functionality available to Python users. Once `yara-python` is built and installed on your system you can use it as shown below:

```
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'
})

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 have 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 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.

In all cases ***compile*** returns an instance of the class ***Rules***. This class have 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 YARA's ***load*** method:

```
rules = yara.load('/foo/bar/my_compiled_rules')
```

The result of ***load*** is also an instance of the class ***Rules***.

Instances of ***Rules*** also have a ***match*** method, which allows 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:

```
f = fopen('/foo/bar/my_file', 'rb')
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*** parameter.

```
matches = rules.match('/foo/bar/my_file',
    externals= {'var1': 'some other string', 'var2': 100})
```

Externals variables defined during compile-time don't need to be defined again in subsequent invocations of ***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 hasn't completed before the specified number of seconds elapsed, a **TimeoutError** exception will be raised.

You can also specify a callback function when invoking **match** method. The provided function will be called for every rule, no matter if matching or not. 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
    yara.CALLBACK_CONTINUE

matches = rules.match('/foo/bar/my_file', callback=mycallback)
```

The passed dictionary will be something like this:

```
{
    'tags': ['foo', 'bar'],
    'matches': True,
    'namespace': 'default',
    'rule': 'my_rule',
    'meta': {},
    'strings': [(81L, '$a', 'abc'), (141L, '$b', 'def')]
}
```

The **matches** field indicates if the rules matches the data or not.

The **strings** fields is a list of matching strings, with vectors of the form:

```
(<offset>, <string identifier>, <string data>)
```

The **match** method returns a list of instances of the class **Match**. The instances of this class can be treated as text strings containing the name of the matching rule. For example you can print them:

```
for m in matches:
    print "%s" % m
```

In some circumstances you may need to explicitly convert the instance of **Match** to string, for example when comparing it with another string:

```
if str(matches[0]) == 'SomeRuleName':  
    ...
```

The ***Match*** class has the same attributes as the dictionary passed to the callback function:

- rule
- namespace
- meta
- tags
- strings