

Avoiding Source Code Spoofing

From: Mark Davis, Robin Leroy, Peter Constable, Markus Scherer

Date: 2022-01-05

Summary

This is a proposal to form a working group / task force focusing on providing guidance for dealing with the so-called Trojan Source exploit. While we have long-standing documentation on dealing with bidirectional behavior and confusables, this would focus on source code. This task force would deliver its results to the Properties and Algorithms working group of the Unicode Technical Committee for review, and eventually to the UTC for approval.

We would look to involve people who are experts in the Unicode encoding, bidirectional algorithm, and security concerns/mechanisms, as well as experts in programming language standardization and tooling and security experts. The proposal also provides a (rough draft) sketch of what such types of guidance could look like.

Contents

[Summary](#)

[Contents](#)

[Issue](#)

[Proposed Plan](#)

[Preliminary Ideas](#)

[Examples](#)

[High-Level Linters](#)

[Compilers](#)

[Bidirectional order](#)

[Confusables](#)

[Source Code Editors](#)

[Bidirectional order](#)

[Confusables \(homoglyphs\)](#)

[Examples, Analysis](#)

[Properties](#)

[Stateful Format Characters](#)

[White Space](#)

[Requires Bidi](#)

Issue

There have been [reports](#) recently about problems in review of source code containing non-ASCII Unicode characters. The basic problem is that two *different* lines of code (in memory) can have the *same* (or confusingly similar) appearance on the screen. That is, the actual text is different from what the reader perceives it to be. The person reviewing a submission of code from a contributor could be fooled into thinking that the code was ok, when it was really malicious.

This can result from stateful bidirectional controls (used in Arabic, Hebrew, and other right-to-left scripts) that change the natural ordering of characters.

Misleading text can also result from a string containing Arabic or Hebrew characters, from “hidden” characters (such as a zero-width space), or from confusable characters or sequences of characters (Greek omicron vs Latin o). These are the known consequences of the human writing systems. Using BIDI or confusable characters to mislead users has been documented in a number of specifications and technical reports for many years: [UAX #9](#) (“Bidi”), [UAX #31](#) (“Identifiers”), [UTR #36](#) (“Security”), and [UTS #39](#) (“Security Mechanisms”).

In particular, the stateful bidi controls behavior has been present and documented in the Unicode encoding for quite a while: since the very first version (1991), so it is not new.



<https://xkcd.com/1137/>

2012-11-22

Indeed, the developers of some tools have taken notice of these documents, and mitigated some of the issues described therein:

- The Rust compiler warns against confusable identifiers occurring in the same library¹: [CONFUSABLE IDENTIS in rustc lint::non_ascii_idents - Rust](#). It also warns against scripts introduced into a library solely with the use of characters that are confusable with ones from already-used scripts: [MIXED_SCRIPT_CONFUSABLES in rustc lint::non_ascii_idents - Rust](#). It also warns against characters with Identifier_Status=Restricted, *i.e.*, it follows UTS #39 C1: [UNCOMMON CODEPOINTS in rustc lint::non_ascii_idents - Rust](#)
- An OCaml package that adds support for Unicode identifiers warns not just about confusable identifiers, but about identifiers that mix scripts without separating them by U+005F LOW LINE: [whitequark/ocaml-m17n: Multilingualization for the OCaml source code](#).
- Visual Studio orders tokens² left-to-right³, *i.e.*, it implements UAX #9 HL4, see the discussion above <https://twitter.com/KhaledGhetas/status/993118048575021058>.
- Both the Swift and Rust compilers make use of confusable detection to provide better error messages, *e.g.*, when General Punctuation “quotation marks” are mistakenly used in place of their Basic Latin “counterparts”: <https://bugs.swift.org/browse/SR-331>.

The relative scarcity of these examples shows, however, that we can’t expect developers of compilers or source code editors to find that material and interpolate how it would apply to their specific domain. It is often tricky to find a balance between security measures and usability, when deep knowledge of Unicode is not normally a requirement for compiler developers. So we should do a better job about providing more focused documentation for such developers.

While the main focus is security, this guidance can also be useful in improving the usability of code editors and language tooling when non-ASCII characters can be in strings, comments, and identifiers such as variable names. Some people may be surprised that identifiers are included, but many modern programming languages allow for non-ASCII identifiers, including Java, Rust, Swift, Go, and many [others](#).

¹ An exception is made for confusable Basic Latin, *e.g.*, 0 and O, or I, l, and 1. This may mean that an alternate, narrower, definition of confusability is needed for source code, as [suggested in a swift-evolution discussion](#) when confusables were brought up for that language. Such a definition might treat p (Latin) and ρ (Greek) as non-confusable, though it would treat p (Latin) and p (Cyrillic) as confusable. Alternatively, it may mean that some exceptions in the case of ASCII are needed for other reasons. This needs further investigation.

² Technically units smaller than tokens: escape sequences within string literals are isolated and ordered as tokens would be.

³ Right-to-left token order is not supported: <https://docs.microsoft.com/en-us/visualstudio/ide/use-bidirectional-languages#right-to-left-reading-order>.

Proposed Plan

We put together a work group / task force of experts to work on the following tasks. We create a Unicode email group where we can invite security experts and others to participate, and also schedule regular meetings for f2f discussion. It is anticipated that this task force would be of limited duration.

The goals would be to:

- A. Assemble documentation providing guidance for avoiding spoofing issues. Make that available for review and feedback.
- B. Produce draft proposed updates of [UAX #9](#) (“Bidi”, aka UBA), [UAX #31](#) (“Identifiers”), [UTR #36](#) (“Security”), and [UTS #39](#) (“Security Mechanisms”) using the information in #1, and post for comment.
 - We would not be changing the UBA, but otherwise exactly where in these documents additions would be made is as yet unclear. For example, #39 and #36 could have new sections, while the others would point to the new sections at appropriate places.
 - However, we may also add lengthier sections in multiple places, such as a new section of #9 describing issues in the display order of program text. The key will be to have new material in a location where the desired audience could find it.
- C. In ICU, provide code snippets and/or APIs to implement utility functions that could be used directly to help avoid problems. (The implementations could also be ported to other languages.) Some areas are:
 - Determine where the ordering of programming language tokens on a line would not be visually monotonic.
 - Given the visual appearance of a line, determine where tokens would appear to have different boundaries.
 - Given a token, determine when a different memory representation could appear identical (or nearly so). For example, the in-memory token **a21N** looks visually identical to two in-memory tokens: **a2N1** and **aN21**.

Preliminary Ideas

We would expect to have progressive iterations of the documentation for coverage and comprehension among experts in Unicode and security. Here are some preliminary thoughts.

There are three different areas where problems can be addressed:

- **High-level linters** — These are linters that would check the contents of source code files in a repository, but with little to no knowledge of the particular programming language of the file, though they would at least need to know that it is programming language source code.
- **Compilers** — These include compilers and linters that are specific to particular programming languages.
- **Source Code Editors** — These include source-code editors and code review applications.

Some of the solutions can be implemented in multiple areas, not just one. In practice, it is best to implement whatever can be done in *all* of these areas, because source code can be in multiple repositories, be compiled by different compilers (even for the same programming language), and viewed in different source code editors.

Much of the text below focuses on viewing source code where the overall direction of text is left-to-right. Most of the discussion can also be applied where the overall direction of text is right-to-left, *mutatis mutandis*. There are, however, some additional wrinkles for right-to-left text in that most programming languages have ASCII keywords, forcing a larger percentage of text to be in the opposite direction of the overall direction.

However, since the goal is to prevent cases where the memory representation may be not what people expect, we have to take into account the fact that different people have different expectations. A person whose native language is written with right-to-left characters may have quite different expectations of order than one whose native language is written with left-to-right characters.

Notation

The examples are presented in the following format. The first section of three rows shows the in-memory representation of the line of text (also called the ‘logical order’). The *memory* row has the characters, the *index* row has the starting index of each character in memory, and the *token#* row has a token number (a segment of text treated as atomic by a compiler), where the number is the number of the token in memory: 0 is first, 1 is second, and so on.

In the following, the characters in token #4 (indices 4-6) are Hebrew characters in memory order (Arabic would also work, except that the cursive nature of the characters makes it a bit less obvious what is going on). Highlighting is used to show areas of interest.

memory:	s		=		κ	ב	λ	-	1	0	0	;
index:	0	1	2	3	4	5	6	7	8	9	10	11
token#:	0	1	2	3	4			5	6		7	

visual:	s		=		1	0	0	-	א	ב	א	;
index:	0	1	2	3	8	9	10	7	6	5	4	11
token#:	0	1	2	3		6		5		4		7

The second section shows the visual appearance. The visual row shows the order the user would see. The index row again shows the character indexes in memory : so the Hebrew character א again has the index 4. The *token#* row also shows the token with their in-memory order.

Notice that not only are the Hebrew characters reversed, but the tokens are also reordered on the line. Token 6 is coming right after token 3, then token 5, then token 4.

High-Level Linters

There are some changes that could be made at a very high level, perhaps without even knowing what the program language is. This includes raising errors on broad classes of characters present in the source files or in easily-identified sections thereof.

The goal of such high-level linters is to be able to quickly add support for many languages, making it possible to process large repositories in short order.

While it should be possible to have any Unicode character in source code, a high-level linter can require some of those characters to be escaped in programming language source code, which prevents their use for spoofing. For example, the stateful bidirectional characters are invisible, but affect the order of text around them.

Alternatively, a linter could disallow them except in end-of-line comments, where the reordering would be confined to the comment; this would make it possible to use those characters to get comments with bidirectional text to render as desired.

[TBD Add example.]

One point many reports have overlooked is that stateful bidi controls are *not* the only cause of reordering issues — they do not have to be present for bidi reordering problems to occur! Problems can occur without them, as the examples under [Compilers](#) below illustrate. However, requiring the stateful controls to be escaped takes care of a large swath of possible ordering problems.

Importantly, the stateful bidi controls and other characters below are not in Unicode on a whim; they are needed to properly display plain text. So the escaping approach *must not* be applied to files that don't contain source code, such as plaintext or markdown.

The following are candidates for required escaping:

- [Deprecated characters](#)
- [Controls \(except common whitespace\)](#)
- [Format characters \(with certain exceptions\)](#)
- [Private use](#)

- [Surrogates](#)
- [Unassigned code points](#) (including non-characters)
- [Misleading whitespace](#), including the “blank” Hangul fillers

Most of these are orthogonal to the reordering problems, but can also contribute to reducing confusables.⁴ Note that only linters that are updated regularly to the latest version of Unicode should disallow unassigned code points.

There is a broad spectrum of levels of understanding of the programming language which can be used in such a linter.

The better the understanding of the language, the more actual issues can be detected, and the more remedies can be implemented without harming usability; some potentially interesting levels are:

1. Knowing nothing about the structure of the file, except that it is in a programming language for which this kind of linting is appropriate.
 - By definition, this is the minimum level of understanding at which something could be done.
2. Identifying comments and string literals;
 - A linter at this level is needed to diagnose situations where reordering causes comments or strings to look like executable code, and vice versa;
 - it could forbid more characters outside of comments or strings, while not getting in the way of legitimate usage of plain text within them.
 - This level remains very quick to implement, as only two types of tokens, both delimited, need to be taken into account for each language.
3. Identifying comments, string literals, and identifiers;
 - A linter at this level could start considering mixed-script or confusable detection;
 - it could, in most languages, check the exit directionality of all strongly directional tokens, avoiding token reordering (see next section).
 - Implementing a linter at this level requires implementing the definition of identifiers in each language, whether based on UAX#31 or something else, *e.g.*, general categories.
4. Full lexing.
 - Linters at that level (and above, if parsing or even semantic analysis are also used) are discussed in the next section.

Compilers

This level includes compilers, linters, and other programming language tooling that has a deep understanding of particular source code.

⁴ Note that restricting some of these characters in comments may make it more difficult to use natural text in right-to-left scripts, such as described in [how-to-comment-in-a-right-to-left-language-in-visual-studio-ide](#). For example, users may need to reword and split comments to replace the use of RLE..PDF ranges.

The advantages of changes on this level is that the help they provide is not limited to particular IDEs or editors. They can require escaping (as with [High-Level Linters](#)), but then can also apply some additional tests since they are aware of the particular syntax of the programming language in question.

In particular, they can apply tests in the following areas:

Bidirectional order

1. Tokens out of order

memory:	s		=		א	ב	ג	-	1	0	0	;
index:	0	1	2	3	4	5	6	7	8	9	10	11
token#:	0	1	2	3	4			5	6			7
visual:	s		=		1	0	0	-	א	ב	א	;
index:	0	1	2	3	8	9	10	7	6	5	4	11
token#:	0	1	2	3	6			5	4			7

- In memory, the Hebrew identifier (token #4) is before 100 (token #6), but visually it appears *after*. That is, tokens 4 and 6 are swapped.
- For identifiers, this only affects programming languages that allow non-ASCII identifiers. However, the same effect can be produced with strings:

memory:	s		=		"	א	ב	ג	"		+		"	1	0	0	"	;
index:	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
token#:	0	1	2	3	4					5	6	7	8					9
visual:	s		=		"	1	0	0	"		+		"	א	ב	א	"	;
index:	0	1	2	3	4	13	14	15	12	11	10	9	8	7	6	5	16	17
token#:	0	1	2	3	4	8				7	6	5	4				8	9

- In this case, tokens #4-#8 are reversed *and* tokens 4 and 8 are split

It is recommended that programming languages allow for insertion of LRM (*left-right-mark*) characters in whitespace (as some already do), see [UAX31-R3 “Pattern White Space and Pattern Syntax Characters”](#). That provides a mechanism for users to correct the visual display of RTL text, and thus avoid raising errors in these cases (and look correct even when pasted into email, etc.). A sufficiently smart editor could handle this automatically.

2. Token with ambiguous memory order: two tokens appear identical visually, but are ordered differently in memory.

- Line 1 — no reordering happens

memory:	s		=		a	1	א	;
---------	---	--	---	--	---	---	---	---

index:	0	1	2	3	4	5	6	7
token#:	0	1	2	3	4	4	4	5
visual:	s		=		a	1	κ	;
index:	0	1	2	3	4	5	6	7
token#:	0	1	2	3	4	4	4	5

- Line 2 — the last two characters of an identifier are swapped in visual order

memory:	s		=		a	κ	1	;
index:	0	1	2	3	4	5	6	7
token#:	0	1	2	3	4	4	4	5
visual:	s		=		a	1	κ	;
index:	0	1	2	3	4	6	5	7
token#:	0	1	2	3	4	4	4	5

- Note that what appears to be the same token visually (an identifier a1κ on both visual lines) has two different memory representations (a1κ and aκ1)

3. **Token with unexpected boundaries:** A token encompasses more or fewer characters than expected.

memory:	p	r	i	n	t	f	(R	"	(κ	(υ)	"	?	p	a	s	s	w	o	r	d	:	")	")	;
index:	0	1	2	3	4	5	6	7	8	9	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2
											0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
token#:	0	0	0	0	0	0	1	2	3							4	5							6	7		8	9		
visual:	p	r	i	n	t	f	(R	"	()	υ	(κ	"	?	p	a	s	s	w	o	r	d	:	")	")	;
index:	0	1	2	3	4	5	6	7	8	9	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2
											3	2	1	0	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
token#:	0	0	0	0	0	0	1	2	3							4	5							6	7		8	9		

- This uses the ‘raw string’ construct in C++, R"(...)”, but in the visual row it appears that the raw string goes out to position 27, printing a long literal. In fact, the raw string only goes out to position 14, so what it actually prints out is contents of the variable password: see <https://gcc.gnu.org/z/Y8dPhP49Y>

Code to detect tokens out of order is relatively straightforward to implement using an API that handles the Unicode Bidirectional Algorithm, such as the Bidi class in ICU. Code to detect a token with ambiguous memory order or with unexpected boundaries is trickier. In all cases we should supply utilities in ICU so that developers don’t need to learn the intricacies of bidirectional ordering. Those utilities can also serve as a pattern for constructing similar utilities in other programming languages.

Confusables

[TBD We already have descriptions and ICU code and data to support confusable detection, mixed scripts, and so on. We need to have a description of the application to source code, with examples. At a compiler level there is enough context that the compiler knows the identifiers that are in scope at any particular location, and can test whether a new identifier declaration is confusable with one of the previous ones.]

Source Code Editors

Spoofing problems can also be addressed at the editing stage. This also applies to viewing applications, such as with code review diffs. But for brevity we'll subsume those under the term 'code editor'.

Bidirectional order

There are a few techniques that can be applied. Most of these involve knowledge of the particular programming language that is being edited or used.

Syntax highlighting can help to make clear when the visual appearance is misleading. For example, in the following example, part of a comment is reordered to appear as if it were code.

[TBD example]

However, it cannot be the only solution, because there are reordering examples where two strings are simply swapped, and plain syntax highlighting would not help for those. Special highlighting could expose the reordering issues discussed under [Compilers](#), but can also be confusing (and color can cause accessibility problems for the color-blind).

[TBD example]

Visual escaping is where a character is displayed as a "chit" instead of its normal appearance. That chit can be a hex representation (eg `202A`) or a mnemonic (eg `LRE`). In either case, it is helpful to have a mouse hover or other mechanism that can have a longer description (eg "left-to-right embedding: changes the text order around it"). The text should be ordered as if the chit were a single neutral character, thus preventing reordering via stateful controls. Visual escaping can help to expose 'hidden' characters like the stateful bidi controls, and reduce the opportunity for reordering. It does require extra work in the code editor, but many code editors provide a 'show hidden' mode which can be extended.

[TBD example]

However, visual escaping of individual characters is not sufficient. In the example below, there are reordering problems despite there being no 'hidden' characters.

[TBD example]

A fuller solution involves lexing the line for tokens, and presenting those tokens in a monotonic order. Typically that is left-to-right, but specialized or more sophisticated editors would also provide for a

right-to-left mode. Note that something that does enough syntax highlighting most likely already does the bulk of the work of this approach. And if it is rendering it in HTML, it just needs a bunch of `dir=ltr` spans to achieve the desired behaviour.

[TBD example]

That reduces the attack surface dramatically, but not completely, because there can still be tokens that are visual spoofs of others.

[TBD Example]

Confusables (homoglyphs)

[TBD As above, we have descriptions and ICU code and data to support confusable detection, mixed scripts, and so on. Need to have a description of the application to source code, with examples.]

Properties

There are a few cases where we might consider some changes to properties.



Stateful Format Characters

Stateful controls or format characters have an extended visual effect between them (or until the end of a paragraph).

There is currently no property that indicates which characters are stateful, which is important information. We could consider an additional binary property for that, such as the following:

Stateful_Format

U+202A	LEFT-TO-RIGHT EMBEDDING
U+202B	RIGHT-TO-LEFT EMBEDDING
U+202C	POP DIRECTIONAL FORMATTING
U+202D	LEFT-TO-RIGHT OVERRIDE
U+202E	RIGHT-TO-LEFT OVERRIDE
U+2066	LEFT-TO-RIGHT ISOLATE
U+2067	RIGHT-TO-LEFT ISOLATE
U+2068	FIRST STRONG ISOLATE
U+2069	POP DIRECTIONAL ISOLATE
U+206A	🚫 INHIBIT SYMMETRIC SWAPPING
U+206B	🚫 ACTIVATE SYMMETRIC SWAPPING
U+206C	🚫 INHIBIT ARABIC FORM SHAPING
U+206D	🚫 ACTIVATE ARABIC FORM SHAPING
U+206E	🚫 NATIONAL DIGIT SHAPES
U+206F	🚫 NOMINAL DIGIT SHAPES
U+FFF9	INTERLINEAR ANNOTATION ANCHOR
U+FFFB	INTERLINEAR ANNOTATION TERMINATOR
U+13437	EGYPTIAN HIEROGLYPH BEGIN SEGMENT
U+13438	EGYPTIAN HIEROGLYPH END SEGMENT
U+1D173	MUSICAL SYMBOL BEGIN BEAM
U+1D174	MUSICAL SYMBOL END BEAM
U+1D175	MUSICAL SYMBOL BEGIN TIE
U+1D176	MUSICAL SYMBOL END TIE
U+1D177	MUSICAL SYMBOL BEGIN SLUR
U+1D178	MUSICAL SYMBOL END SLUR

13282  EGYPTIAN HIEROGLYPH O033A • end of serekh enclosure → 13258  egyptian hieroglyph o006a	
---	--

The suggested `Stateful_Format` is a binary property, but an even more useful alternative would be to add two **string** properties that associate the character(s) that end a state with the ones that begin it, such as:

```
\p{EndState:\u202A} == [\u202C]
```

...

```
\p{StartState:\u202C} == [\u202A\u202B]
```

These would be more useful for *fixing* broken text, since these supply the missing characters that could be used to form a balanced pair.

White Space

The following characters are categorized as `gc=Lo`, but have no appearance and are typically displayed as whitespace or invisible. Here an appearance is shown in «...», but the precise appearance depends on the font and rendering system: when not part of a Hangul syllable they could be invisible, look like a space, or have a special appearance.

«»	U+115F	HANGUL CHOSEONG FILLER
«»	U+1160	HANGUL JUNGSEONG FILLER
« »	U+3164	HANGUL FILLER
« »	U+FFA0	HALFWIDTH HANGUL FILLER

Typically characters like these are either `Cf` or `Z`, so people overlook them when considering security issues. There may be others of this type as well.

We should consider how we can categorize these characters in a way that highlights their nature.

Options:

1. Have a special property for them?
2. Re categorize them as `Cf`? (And make sure they work properly in word-break, etc.)
3. Other options?

Requires Bidi

An important optimization is to quickly check for the following set of characters; if none are found then implementations don't have to invoke any special machinery for bidirectional handling.

`[\p{bc=AL}\p{bc=AN}\p{bc=LRE}\p{bc=RLE}\p{bc=LRO}\p{bc=RLO}\p{bc=PDF}\p{bc=FSI}\p{bc=RLI}\p{bc=LRI}\p{bc=PDI}\p{bc=R}]`

However, this expression is easy for people to get wrong, so we might consider whether there is some way to make that easier and less error-prone.

Option 1: Document this (and its inverse) explicitly in UAX #9 and also in #39.

Option 2: Provide a new property `RequiresBidi` that has the contents above.

Reports

Some of the reports of problems include:

- <https://www.trojansource.codes/>
- <https://www.kb.cert.org/vuls/id/999008>
- <https://www.rapid7.com/blog/post/2021/11/04/trojan-source-cve-2021-42572/>
- <https://www.python.org/dev/peps/pep-0672/>
- And [others](#)

More examples

Robin has done some in-depth analysis of examples and behavior in different programming languages, and we could add more.

Usability

The use of strongly right-to-left text in identifiers or string literals can easily make source code illegible.

C# 1.0/1.2 (ISO/IEC 23270:2003) or later:

```
Console.WriteLine("הודעת", "{1} {0} :השתבש", this);
```

parses—and is typed—as

```
Console.WriteLine("Error: {0} ({1})", message, this);
```

Python 3.0 (2008) or later:

```
return אינטגרל(lambda 1=ל, 0=ח, 2 ** א :א)
```

parses—and is typed—as

```
return integral(lambda a: a ** 2, from_=0, to=1)
```

C++11 (ISO/IEC 14882:2011) or later:

```
std::vector<مِواءَ < سَنُور>
```

parses—and is typed—as

```
std::vector<meow> cat;
```

In the same language,

```
return u8"مِواءَ"; // رسالة العنصر النائب
```

parses—and is typed—as

```
return u8"meow"; // Placeholder message.
```

Rust 1.53.0 (2021) or later:

```
fn אינטגרל פונקציה >: Fn(f64) -> f64 <(  
    קטע : פונקציה, קטע : std::ops::Range<f64>) -> f64 {
```

parses—and is typed—as

```
fn integral<Function: Fn(f64) -> f64>(  
    integrand: Function, interval: std::ops::Range<f64>) -> f64 {
```

Security

The examples in this section don't simply seem confusing or illegible as the ones above do; instead the reader would be expected to think that they are programs with different semantics. While we provide examples involving literals or identifiers in strongly right-to-left scripts, we also provide some that have neither (and instead use the invisible implicit directional marks), as these may not look suspicious even in a codebase that does not use right-to-left scripts in identifiers.

C++98 or later:

```
std::cerr << "encountered" << (errors == 0 ? " 0 " : "")  
    << "errors";
```

Will print “encountered errors” if and only if errors ≠ 0, and “encountered 0 errors” otherwise.

Rust 1.9 (2016) or later:

```
return x >> 8;
```

While this looks like a right shift by eight bits, it is a left shift by eight bits.