

## Item 19: Declare functions noexcept whenever possible.

In C++98, exception specifications were <sup>rather</sup> somewhat temperamental beasts. You had to <sup>summarize the</sup> expressly list all exception types a function might emit (~~though the ability to specify a base class for all derived class exception types helped~~), and this imposed constraints on the function's implementation. If the implementation was changed, the exception specification might require modification, too, and that not only opened the door to consistency errors (i.e., an exception specification that no longer corresponds to the implementation), it also meant that callers ~~of the function~~ might be <sup>red stop compiling</sup> broken (i.e., wouldn't compile), because an exception specification is part of a function's interface. For these reasons, C++98 exception specifications were <sup>never gained much popularity</sup> largely ignored. Developers and libraries generally shied away from them, and some compilers didn't even fully implement them.

Over <sup>time</sup> the years, a consensus emerged that the only meaningful information about a function's exception-emitting behavior was whether it had any. Black or white: either a function might emit an exception (~~the type was immaterial~~), or the function guaranteed that callers would never see one. This maybe-or-never dichotomy <sup>essentially replace</sup> forms the basis of C++11's exception specifications, which <sup>and C++14</sup> supplement C++98's. (~~The~~ C++98-style exception specifications continue to be legal in C++11, but they're deprecated.) In C++11, noexcept is for functions that guarantee they'll never emit an exception. When you write a function that can make that guarantee, you'll want to use noexcept.

Why? Because it permits compilers to generate better code (~~i.e., code that's smaller or faster or both~~). There are two reasons for this, but we'll begin with how C++98's way of saying "this function emits no exceptions" differs from C++11's ~~way~~.

Suppose we have a function `f` that promises callers they'll never receive an exception. The C++98 and C++11 ways of expressing that are:

```
void f(int x) throw();    // C++98 approach: f emits no
                          // exceptions
void f(int x) noexcept;   // C++11 approach: f emits no
                          // exceptions
```



1 Perhaps surprisingly, neither C++98 nor C++11 permits compilers to reject code in  
2 f that could violate these exception specifications. As a result, the function could  
3 be implemented like this:

```
4 void f(int x) noexcept      // C++98 version would use "throw()"
5 {
6     if (x >= 0) return x * q42 - b;           // if x >= 0 ...
7     throw std::invalid_argument(           // else throw!
8         "Invalid value for x: " + std::to_string(x)
9     );
10 }
```

11 This may look ~~ridiculous~~<sup>absurd</sup>, but it's perfectly legal C++. Furthermore, looks aren't  
12 everything. The code here could ~~simply~~ be a way of enforcing the precondition that  
13 x must be non-negative. If f is called with a legitimate value, it doesn't throw. If an  
14 invalid value is passed in, however, the precondition violation causes the function  
15 to have undefined behavior, and this implementation uses that freedom—  
16 undefined behavior means that *anything* can happen—to throw an exception.

17 ~~Incidentally~~<sup>As an aside</sup>, note the use of `std::to_string` to produce a textual representation  
18 of the value of x. Among C++11's lesser-known features is a set of overloaded  
19 `std::to_string` functions that produce `std::string` objects from numeric val-  
20 ues. The Standard Library has functions to perform the reverse transformations,  
21 too (i.e., from `std::strings` to ints, unsigneds, floats, ~~doubles~~, etc.), but the  
22 naming convention for those functions, albeit following a consistent pattern, is ra-  
23 ther cryptic: `stoi`, `stol`, ~~`stod`~~, etc. The C++11 Standard Library also offers  
24 `std::wstring`-based versions of all these functions.

25 But back to the difference in meaning between these two declarations:

```
26 void f(int x) throw();           // C++98 approach
27 void f(int x) noexcept;          // C++11 approach
```

28 If f's implementation permits an exception to escape, the function's exception  
29 specification is violated. With the C++98 approach, runtime behavior is to unwind  
30 the call stack to f's caller, then invoke the *unexpected handler* function, which will



1 lead to program termination (typically by calling `std::terminate`).<sup>\*</sup> With the  
2 C++11 approach, runtime behavior is slightly different: *possibly* unwind the stack,  
3 then call `std::terminate`.

4 The fact that, with `noexcept`, the call stack only *might* be unwound turns out to  
5 make a big difference during code generation. Optimizers are no longer con-  
6 strained to keep the runtime stack in an unwindable state if an exception would  
7 propagate out of the function, nor must they ensure that objects in a `noexcept`  
8 function are destroyed in the inverse order of construction should an exception  
9 leave the function. The result is greater opportunities for optimization, not only  
10 within the body of a `noexcept` function, but also at call sites to the function. This  
11 degree of flexibility is present only for `noexcept` functions. Functions with  
12 "throw()" exception specifications lack it, as do functions with no exception speci-  
13 fication at all. The situation can be summarized this way (where it doesn't make  
14 any difference what func does):

```
15 RetType func(parameters) noexcept;    // more optimizable
16 RetType func(parameters) throw();     // less optimizable
17 RetType func(parameters);             // less optimizable
```

18 This alone should provide sufficient motivation to declare functions `noexcept`  
19 whenever you can. For some functions, however, the case is even stronger. The  
20 move operations are the preeminent example.

21 → Suppose you have a large investment in a C++98 code base making use of  
22 `std::vectors` of Widgets. Naturally, Widgets are added to the `std::vector`  
23 from time to time, perhaps via `push_back`:

---

<sup>\*</sup> By default, the unexpected handler function is `std::unexpected`, and, by default, this calls the *terminate handler* function, which, by default, is `std::terminate`. The unexpected and terminate handler functions may be replaced via calls to `std::set_unexpected` and `std::set_terminate`, but there are constraints on the behavior of replacement functions, and in the example we're considering, program execution must terminate. `std::unexpected` and `std::terminate`, being part of the C++98 approach to exception handling, are deprecated in C++11.



```
1  std::vector<Widget> vw;
2  ...
3  Widget w;
4  ...                // put w into proper state
5  ...                // for addition to vw
6  vw.push_back(w);    // add w to vw
7  ...
```

8 Assume this code works fine, and you have no interest in modifying it for C++11.  
9 However, you do want to take advantage of the fact that C++11's move semantics  
10 can improve the performance of existing code when move-enabled types are in-  
11 volved. You therefore ensure that `Widget` has move operations, either by writing  
12 them yourself or by seeing to it that the conditions for their automatic generation  
13 are fulfilled (see Item 20).

14 When a new element is added to a `std::vector` via `push_back`, it's possible that  
15 the vector lacks space for it, i.e., that the vector's size is equal to its capacity. When  
16 that happens, the vector allocates a new, larger, chunk of memory to hold its ele-  
17 ments, and it transfers the elements from the existing chunk of memory to the new  
18 one. In C++98, the transfer was accomplished by copying each element from the  
19 old memory into the new memory, then destroying the copies in the old memory.  
20 This approach enabled `push_back` to offer the strong exception safety guarantee:  
21 if an exception was thrown during the copying of the elements, the state of the vec-  
22 tor remained unchanged, because none of the elements in the original memory  
23 was destroyed until all elements had been successfully copied into the new  
24 memory.

25 In C++11, a natural optimization would be to replace the copying of vector ele-  
26 ments with moves. Unfortunately, blindly doing this runs the risk of violating  
27 `push_back`'s exception safety guarantee. If  $n$  elements have been moved from old  
28 memory to new and ~~then~~ an exception is thrown moving element  $n+1$ , the  
29 `push_back` operation can't run to completion. But the original vector has been  
30 modified:  $n$  of its elements have been moved from. Restoring their original state

everywhere



1 may not be possible, because attempting to move each object back into the original  
2 memory may itself yield an exception.

3 This is a serious problem, because the behavior of your C++98 code base could de-  
4 pend on `push_back`'s strong exception safety guarantee. C++11 ~~compilers~~ there-  
5 fore can't silently replace copy operations inside `push_back` with moves. They  
6 must continue to employ copy operations. *Unless*, that is, it's known that the move  
7 operations are guaranteed not to emit exceptions. In that case, replacing element  
8 copy operations inside `push_back` with move operations would be safe, and the  
9 only side effect would be improved performance.

implementations

10 `std::vector::push_back` takes advantage of this "move if you can, but copy if  
11 you must" strategy, and it's not the only function in the Standard Library that does.  
12 Other functions sporting the strong exception safety guarantee in C++98 (e.g.,  
13 `std::vector::reserve`, `std::deque::insert`, etc.) behave the same way. All  
14 these functions replace calls to copy operations in C++98 with calls to move opera-  
15 tions in C++11 if (and only if) the move operations are known to never emit excep-  
16 tions. But how does a compiler know if a move operation won't produce an excep-  
17 tion? The answer ~~should be~~ obvious: it checks to see if the operation is declared  
18 `noexcept`.<sup>is</sup>

19 And yet, it's not quite that simple. Popping the hood and peeking inside to see how  
20 things work is instructive, so here we go.

21 Inside a function like `push_back`, suppose we want to transfer an object (i.e., copy  
22 or move it, depending on what is appropriate) from one place to another. Assume  
23 we have an iterator, `src`, referring to the object to be transferred and a second it-  
24 erator, `dest`, referring to where it should be transferred. So we'd have a statement  
25 something like this:

---

\* Alternatively, the function could have a C++98-style empty exception specification (i.e., `"throw()"`), but the only reason I can imagine why a move operation—something that didn't exist in C++98 and therefore can't be part of a legacy code base—would employ `throw()` instead of `noexcept` would be to accommodate compilers with incomplete C++11 support, i.e., compilers where move operations are supported, but `noexcept` isn't. ~~Sadly, such compilers do exist.~~



```
1  *dest = *src;                // transfer *src to *dest
2                                // (incorrect version 1)
```

3 The statement would be inside a loop, because we'd ultimately need to transfer all  
4 the objects in the container, but understanding how things work for one object is  
5 all we need here.

6 As the comment indicates, the code is incorrect. The problem is that \*src is an  
7 lvalue, so this statement would unconditionally copy \*src to \*dest. That'd have  
8 been fine in C++98, but in C++11, we want to do a move if we can. The usual way to  
9 move an lvalue is to apply `std::move` to it:

```
10 *dest = std::move(*src);      // transfer *src to *dest
11                               // (incorrect version 2)
```

12 This is also incorrect, because now we're moving \*src, regardless of whether it's  
13 move assignment operator is `noexcept`. As we've discussed, doing that would  
14 prevent `push_back` from maintaining its strong exception safety guarantee, and  
15 maintaining that guarantee is essential for ensuring that code written under  
16 C++98 continues to function correctly.

17 The correct code takes advantage of `std::move`'s poorly publicized cousin,  
18 `std::move_if_noexcept`:

```
19 *dest = std::move_if_noexcept(*src); // transfer *src to *dest
20                                     // (correct version)
```

21 Conceptually, `std::move_if_noexcept` causes \*src to be moved if its move as-  
22 signment operator is `noexcept`, and otherwise it causes \*src to be copied. That's  
23 exactly what we want, and that's why this code is correct. ○ Uses of  
24 `std::move_if_noexcept` are scattered throughout strongly exception safe func-  
25 tions in the Standard Library, and that's why you have a special incentive to de-  
26 clare your move operations `noexcept`: it enables their use inside such functions.

27 The conceptual description of `std::move_if_noexcept` deviates from its true  
28 behavior in two small ways. First, if `std::move_if_noexcept` is invoked on an  
29 object of a move-only type, a move will be performed, even if it might yield an ex-  
30 ception. This is understandable: what else can `std::move_if_noexcept` do, giv-  
31 en that the type can't be copied? Anyway, this behavior can't break any C++98 leg-

Besides

; this is incorrect



1 acy code, because <sup>there was</sup> there's no such thing as a move-only type in C++98. Moreover,  
2 ~~your chances of encountering a move-only type with non-noexcept move opera-~~  
3 ~~tions are quite small. In fact, your chances of encountering any non-noexcept~~  
4 ~~move operation are small. Most move operations have an implementation that~~  
5 ~~naturally consist of statements where exceptions don't arise.~~

6 Second, `std::move_if_noexcept`, like `std::move`, doesn't actually move any-  
7 thing. Rather, it performs a cast to an rvalue that, through overloading resolution,  
8 can cause a move assignment operator or a move constructor to be invoked. It's  
9 these functions that actually move values around. For details on the relationship  
10 among `std::move` (and `std::move_if_noexcept`), casting to rvalues, move op-  
11 erations, and overload resolution, consult Item 21.

12 Accompanying the move operations on the podium for functions that <sup>where</sup> especially  
13 benefit from a noexcept declaration is `swap`. The justification for `swap`'s presence  
14 is different from that for the move operations. ~~First~~ being a heavily-used function  
15 (many algorithms rely on `swap`, as do implementations of many copy assignment  
16 operators), the optimization opportunities that `noexcept` affords are unusually  
17 worthwhile. ~~Second~~, whether particular versions of `swap` in the Standard Library  
18 are `noexcept` is sometimes dependent on whether user-defined type-specific  
19 swaps are `noexcept`. For example, the declarations for the Standard Library's  
20 swaps for arrays and for `std::pair` are:

```
21 template <class T, size_t N>  
22 void swap(T (&a)[N],  
23          T (&b)[N]) noexcept(noexcept(swap(*a, *b)));  
  
24 template <class T1, class T2>  
25 struct pair {  
26     ...  
27     void swap(pair& p) noexcept(noexcept(swap(first, p.first)) &&  
28                                   noexcept(swap(second, p.second)));  
29     ...  
30 };
```

31 These functions are *conditionally noexcept*: whether they are `noexcept` depends  
32 on whether the expression inside the outer `noexcept` is. Given two arrays of  
33 `Widget`, for example, swapping them is `noexcept` only if swapping an element  
34 from each array is `noexcept`, i.e., if `swap` for `Widget` is `noexcept`. The author of

especially  
is usually beneficial

Furthermore

Comment [sdm1]: Font should be both code and Term Introduction.



1 ~~swap~~<sup>is swap</sup> for Widget thus determines whether swapping arrays of Widget is noex-  
2 cept (which, in turn, could determine whether other swaps are noexcept, e.g.,  
3 swap for arrays of arrays of Widget). Similarly, whether swapping two `std::pair`  
4 objects containing Widgets is noexcept depends on whether swap for Widgets is  
5 ~~noexcept~~. The fact that swapping higher-level data structures can generally be  
6 noexcept only if swapping their lower-level constituents is noexcept is the rea-  
7 son why you should strive to offer noexcept swap functions.

8 ~~Of course~~<sup>Because</sup>, noexcept is part of a function's interface, ~~so~~ you should declare a func-  
9 tion noexcept only if you are willing to commit to a noexcept implementation  
10 over the long term. If you declare a function noexcept and implement it accord-  
11 ingly, then later decide you wish you hadn't made the noexcept promise, your op-  
12 tions are bleak. You can remove noexcept from the function's declaration, and in  
13 so doing break arbitrarily amounts of client code. You can retain the noexcept  
14 declaration, but change the implementation such that an exception could actually  
15 escape. <sup>the function</sup> In that case, if an exception did escape at runtime, your program would be  
16 terminated. Or you can retain your existing implementation, thus ~~defeating~~<sup>rejecting</sup> what-  
17 ever motivation you ~~had~~<sup>ed</sup> for wanting ~~to~~<sup>to</sup> change the implementation in the first  
18 place. None of these options is appealing.

19 Most functions are *exception-neutral*: they don't throw exceptions themselves, but  
20 if a function they call produces one (directly or indirectly), it causes no harm as it  
21 passes through on its way to an ~~eventual~~<sup>eventual</sup> handler in a different function. Exception-  
22 neutral functions aren't noexcept, because exceptions may pass through them. ~~AA~~

23 ~~Some~~<sup>A</sup> functions, however, are naturally noexcept, and for a few more—notably  
24 the move operations and swap—being noexcept has such a significant payoff, it's  
25 worth implementing them in a noexcept manner if at all possible. When you can  
26 honestly say that a function should never emit exceptions, you should definitely  
27 declare it noexcept.

## 28 Things to Remember

- 29 • noexcept functions offer more optimization opportunities than non-noexcept  
30 functions.



- 1 ♦ C++98 functions offering the strong exception safety guarantee may internally
- 2 call `std::move_if_noexcept` instead of `std::move`.
- 3 ♦ Strive to declare the move operations and `swap` `noexcept`.
- 4