

1 **Item 16: Declare functions noexcept whenever possible.**

2 In C++98, exception specifications were rather temperamental creatures. You had
3 to summarize the exception types a function might emit, so if the function's im-
4 plementation was modified, the exception specification might need revision, too.
5 Changing an exception specification could break client code, because callers might
6 be dependent on the original exception specification. Compilers typically offered
7 no help in maintaining consistency among function implementations, exception
8 specifications, and client code. Most programmers ultimately decided that C++98
9 exception specifications weren't worth the trouble.

10 Interest in the idea of exception specifications remained strong, however, and as
11 work on C++ progressed, a consensus emerged that the truly meaningful infor-
12 mation about a function's exception-emitting behavior was whether it had any.
13 Black or white, either a function might emit an exception or it guaranteed that it
14 wouldn't. This maybe-or-never dichotomy forms the basis of C++11's exception
15 specifications, which essentially replace C++98's. (C++98-style exception specifica-
16 tions remain valid, but they're deprecated.) In C++11, `noexcept` is for functions
17 that guarantee they won't emit an exception.

18 Whether a function should be so declared is fundamentally a matter of interface
19 design. The exception-emitting behavior of a function is of key interest to clients.
20 Callers can query a function's `noexcept` status, and the results of such a query can
21 affect the exception safety or efficiency of the calling code. As such, whether a func-
22 tion is `noexcept` is as important a piece of information as whether a member func-
23 tion is `const`. Failure to declare a function `noexcept` when you know that it will
24 never emit an exception is simply poor interface specification.

25 But there's an additional incentive to apply `noexcept` to functions that won't pro-
26 duce exceptions: it permits compilers to generate better object code. To under-
27 stand why, it helps to examine the difference between the C++98 and C++11 ways
28 of saying that a function won't emit exceptions. Consider a function `f` that promis-
29 es callers they'll never receive an exception. The two ways of expressing that are:

```
30 int f(int x) throw(); // no exceptions from f: C++98 style
```

```
1 int f(int x) noexcept; // no exceptions from f: C++11 style
```

2 If, at runtime, an exception leaves `f`, `f`'s exception specification is violated. With
3 the C++98 approach, the call stack is unwound to `f`'s caller, and, after some actions
4 not relevant here, program execution is terminated. With the C++11 approach,
5 runtime behavior is a bit different: the stack is only *possibly* unwound before pro-
6 gram execution is terminated.

7 The difference between unwinding the call stack and *possibly* unwinding it has a
8 surprisingly large impact on code generation. In a `noexcept` function, optimizers
9 need not keep the runtime stack in an unwindable state if an exception would
10 propagate out of the function, nor must they ensure that objects in a `noexcept`
11 function are destroyed in the inverse order of construction should an exception
12 leave the function. The result is more opportunities for optimization, not only
13 within the body of a `noexcept` function, but also at sites where the function is
14 called. Such flexibility is present only for `noexcept` functions. Functions with
15 “`throw()`” exception specifications lack it, as do functions with no exception speci-
16 fication at all. The situation can be summarized this way:

```
17 RetType function(params) noexcept; // most optimizable
```

```
18 RetType function(params) throw(); // less optimizable
```

```
19 RetType function(params); // less optimizable
```

20 This alone should provide sufficient motivation to declare functions `noexcept`
21 whenever you can.

22 For some functions, the case is even stronger. The move operations are the
23 preeminent example. Suppose you have a C++98 code base making use of
24 `std::vector`s of `Widget`s. `Widget`s are added to the `std::vector`s from time to
25 time, perhaps via `push_back`:

```
26 std::vector<Widget> vw;
```

```
27 ...
```

```
28 Widget w;
```

```
29 ... // work with w
```

```
30 vw.push_back(w); // add w to vw
```

1 ...

2 Assume this code works fine, and you have no interest in modifying it for C++11.
3 However, you do want to take advantage of the fact that C++11's move semantics
4 can improve the performance of legacy code when move-enabled types are in-
5 volved. You therefore ensure that `Widget` has move operations, either by writing
6 them yourself or by seeing to it that the conditions for their automatic generation
7 are fulfilled (see Item 19).

8 When a new element is added to a `std::vector` via `push_back`, it's possible that
9 the `std::vector` lacks space for it, i.e., that the `std::vector`'s size is equal to its
10 capacity. When that happens, the `std::vector` allocates a new, larger, chunk of
11 memory to hold its elements, and it transfers the elements from the existing chunk
12 of memory to the new one. In C++98, the transfer was accomplished by copying
13 each element from the old memory to the new memory, then destroying the origi-
14 nals in the old memory. This approach enabled `push_back` to offer the strong ex-
15 ception safety guarantee: if an exception was thrown during the copying of the el-
16 ements, the state of the `std::vector` remained unchanged, because none of the
17 elements in the original memory was destroyed until all elements had been suc-
18 cessfully copied into the new memory.

19 In C++11, a natural optimization would be to replace the copying of `std::vector`
20 elements with moves. Unfortunately, doing this runs the risk of violating
21 `push_back`'s exception safety guarantee. If n elements have been moved from the
22 old memory and an exception is thrown moving element $n+1$, the `push_back` op-
23 eration can't run to completion. But the original `std::vector` has been modified:
24 n of its elements have been moved from. Restoring their original state may not be
25 possible, because attempting to move each object back into the original memory
26 may itself yield an exception.

27 This is a serious problem, because the behavior of legacy code could depend on
28 `push_back`'s strong exception safety guarantee. Therefore, C++11 implementa-
29 tions can't silently replace copy operations inside `push_back` with moves. They
30 must continue to employ copy operations. *Unless*, that is, it's known that the move
31 operations are guaranteed not to emit exceptions. In that case, replacing element

1 copy operations inside `push_back` with move operations would be safe, and the
2 only side effect would be improved performance.

3 `std::vector::push_back` takes advantage of this “move if you can, but copy if
4 you must” strategy, and it’s not the only function in the Standard Library that does.
5 Other functions sporting the strong exception safety guarantee in C++98 (e.g.,
6 `std::vector::reserve`, `std::deque::insert`, etc.) behave the same way. All
7 these functions replace calls to copy operations in C++98 with calls to move opera-
8 tions in C++11 if (and only if) the move operations are known to not emit excep-
9 tions. But how can a function know if a move operation won’t produce an excep-
10 tion? The answer is obvious: it checks to see if the operation is declared `noex-`
11 `cept`.*

12 swap functions comprise another case where `noexcept` is particularly desirable.
13 `swap` is a key component of many STL algorithm implementations, and it’s com-
14 monly employed in copy assignment operators, too. Its widespread use renders
15 the optimizations that `noexcept` affords especially worthwhile. Furthermore,
16 whether swaps in the Standard Library are `noexcept` is sometimes dependent on
17 whether user-defined swaps are `noexcept`. For example, the declarations for the
18 Standard Library’s swaps for arrays and for `std::pair` are:

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

* The checking is typically rather roundabout. Functions like `std::vector::push_back` call `std::move_if_noexcept`, a variation of `std::move` that conditionally casts to an rvalue (see Item 28), depending on whether the type’s move constructor is `noexcept`. In turn, `std::move_if_noexcept` calls `std::is_nothrow_move_constructible`, and the value of this type trait is set by compilers, based on whether the move constructor has a `noexcept` (or `throw()`) designation.

1 These functions are *conditionally noexcept*: whether they are `noexcept` depends
2 on whether the expressions inside the `noexcept`s are `noexcept`. Given two arrays
3 of `Widget`, for example, swapping them is `noexcept` only if swapping individual
4 elements from the arrays is `noexcept`, i.e., if `swap` for `Widget` is `noexcept`. The
5 author of `Widget`'s `swap` thus determines whether swapping arrays of `Widget` is
6 `noexcept`. That, in turn, determines whether other swaps, such as the one for ar-
7 rays of arrays of `Widget`, are `noexcept`. Similarly, whether swapping two
8 `std::pair` objects containing `Widgets` is `noexcept` depends on whether `swap` for
9 `Widgets` is `noexcept`. The fact that swapping higher-level data structures can
10 generally be `noexcept` only if swapping their lower-level constituents is `noex-`
11 `cept` is the reason why you should strive to offer `noexcept` swap functions.

12 By now, I hope you're excited about the optimization opportunities that `noexcept`
13 affords. Alas, I must temper your enthusiasm. Optimization is important, but cor-
14 rectness is more important. I noted at the beginning of this Item that `noexcept` is
15 part of a function's interface, so you should declare a function `noexcept` only if
16 you are willing to commit to a `noexcept` implementation over the long term. If
17 you declare a function `noexcept` and later regret that decision, your options are
18 bleak. You can remove `noexcept` from the function's declaration (i.e., change its
19 interface), thus running the risk of breaking client code. You can change the im-
20 plementation such that an exception could escape, but keep the original (now in-
21 correct) exception specification. If you do that, your program will be terminated if
22 an exception tries to leave the function. Or you can resign yourself to your existing
23 implementation, abandoning whatever motivated your desire to change the im-
24 plementation in the first place. None of these options is appealing.

25 The fact of the matter is that most functions are *exception-neutral*. Such functions
26 throw no exceptions themselves, but functions they call might emit one. When that
27 happens, the calling function allows the emitted exception to pass through on its
28 way to a handler further up the call chain. Exception-neutral functions are never
29 `noexcept`, because they may emit such "just passing through" exceptions. Most
30 functions, therefore, quite properly lack the `noexcept` designation.

Comment [sdm1]: Style: Code + Term Intro-
duction.

1 Some functions, however, have natural implementations that emit no exceptions,
2 and for a few more—notably the move operations and `swap`—being `noexcept` has
3 such a significant payoff, it's worth implementing them in a `noexcept` manner if at
4 all possible.† **When you can honestly say that a function should never emit excep-**
5 **tions, you should definitely declare it `noexcept`.**

Comment [sdm2]: Should I say that `constexpr` functions are typically good candidates for `noexcept`?

6 Please note that I said some functions have *natural* `noexcept` implementations.
7 Twisting a function's implementation to permit a `noexcept` declaration is the tail
8 wagging the dog. Is putting the cart before the horse. Is not seeing the forest for
9 the trees. Is...choose your favorite metaphor. If a straightforward function imple-
10 mentation might yield exceptions (e.g., by invoking a function that might throw),
11 the hoops you'll jump through to hide that from callers (e.g., catching all excep-
12 tions and replacing them with status codes or special return values) will not only
13 complicate your function's implementation, it will typically complicate code at call
14 sites, too (e.g., code there may have to check for status codes or special return val-
15 ues). The runtime cost of those complications (e.g., extra branches, larger functions
16 that put more pressure on instruction caches, etc.) could exceed any speedup
17 you'd hope to achieve via `noexcept`, plus you'd be saddled with source code that's
18 more difficult to comprehend and maintain. That'd hardly be exemplary software
19 engineering. As a general rule, the only time it makes sense to actively search for a
20 `noexcept` algorithm is when you're implementing the move functions or `swap`.

21 Two more points about `noexcept` functions are worth mentioning. First, in C++98,
22 it was considered bad style to permit the memory deallocation functions (i.e., op-
23 erator `delete` and operator `delete[]`) and destructors to emit exceptions, and
24 in C++11, this style rule has been all but upgraded to a language rule. By default, all
25 memory deallocation functions and all destructors—both user-defined and com-

† The prescribed declarations for move operations on containers in the Standard Library lack `noexcept`. However, implementers are permitted to strengthen exception specifications for Standard Library functions, and, in practice, it is common for at least some container move operations to be declared `noexcept`. That practice exemplifies this Item's advice. Having found that it's possible to write container move operations such that exceptions never need to be emitted, implementers often declare the operations `noexcept`, even though the Standard does not require them to do so.

1 piler-generated—are implicitly `noexcept`. There’s thus no need to declare them
2 `noexcept`. (Doing so doesn’t hurt anything, it’s just unconventional.) The only
3 time a destructor is not implicitly `noexcept` is when a data member of the class
4 (including inherited members and those contained inside other data members) is
5 of a type that expressly states that its destructor may emit exceptions (e.g., de-
6 clares it “`noexcept(false)`”). Such destructors are uncommon. There are none in
7 the Standard Library.

8 Second, let me elaborate on my earlier observation that compilers typically offer
9 no help in identifying inconsistencies between function implementations and their
10 exception specifications. Consider this code, which is perfectly legal:

```
11 void setup();           // functions defined elsewhere
12 void cleanup();
13 void doWork() noexcept
14 {
15     setup();           // set up work to be done
16     ...               // do the actual work
17     cleanup();        // perform cleanup actions
18 }
```

19 Here, `doWork` is declared `noexcept`, even though it calls the non-`noexcept` func-
20 tions `setup` and `cleanup`. This seems contradictory, but it could be that `setup`
21 and `cleanup` document that they never emit exceptions, even though they’re not
22 declared that way. There could be good reasons for their non-`noexcept` declara-
23 tions. For example, they might be part of a library written in C. (Even functions
24 from the C Standard Library that have been moved into the `std` namespace lack
25 exception specifications, e.g., `std::strlen` isn’t declared `noexcept`.) Or they
26 could be part of a C++98 library that decided not to use C++98 exception specifica-
27 tions and hasn’t yet been revised for C++11.

28 Because there are legitimate reasons for `noexcept` functions to rely on code lack-
29 ing the `noexcept` guarantee, C++ permits such code, and compilers generally don’t
30 issue warnings about it.

1 **Things to Remember**

- 2 ♦ `noexcept` is part of a function's interface, so callers may depend on it.
- 3 ♦ `noexcept` functions are more optimizable than non-`noexcept` functions.
- 4 ♦ `noexcept` is particularly valuable for the move operations and for `swap`.
- 5 ♦ Most functions are exception-neutral rather than `noexcept`.

6