

1 **Item 14: Declare functions noexcept if they won't emit ex-** 2 **ceptions.**

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

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

19 Whether a function should be so declared is a matter of interface design. The ex-
20 ception-emitting behavior of a function is of key interest to clients. Callers can
21 query a function's `noexcept` status, and the results of such a query can affect the
22 exception safety or efficiency of the calling code. As such, whether a function is `no-`
23 `except` is as important a piece of information as whether a member function is
24 `const`. Failure to declare a function `noexcept` when you know that it won't emit
25 an exception is simply poor interface specification.

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

```
1 int f(int x) throw(); // no exceptions from f: C++98 style
2 int f(int x) noexcept; // no exceptions from f: C++11 style
```

3 If, at run time, an exception leaves `f`, `f`'s exception specification is violated. With
4 the C++98 exception specification, the call stack is unwound to `f`'s caller, and, after
5 some actions not relevant here, program execution is terminated. With the C++11
6 exception specification, runtime behavior is slightly different: the stack is only *pos-*
7 *sibly* unwound before program execution is terminated.

8 The difference between unwinding the call stack and *possibly* unwinding it has a
9 surprisingly large impact on code generation. In a `noexcept` function, optimizers
10 need not keep the runtime stack in an unwindable state if an exception would
11 propagate out of the function, nor must they ensure that objects in a `noexcept`
12 function are destroyed in the inverse order of construction should an exception
13 leave the function. Functions with “`throw()`” exception specifications lack such
14 optimization flexibility, as do functions with no exception specification at all. The
15 situation can be summarized this way:

```
16 RetType function(params) noexcept; // most optimizable
17 RetType function(params) throw(); // less optimizable
18 RetType function(params); // less optimizable
```

19 This alone is sufficient reason to declare functions `noexcept` whenever you know
20 they won't produce exceptions.

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

```
25 std::vector<Widget> vw;
26 ...
27 Widget w;
28 ... // work with w
29 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 17).

8 When a new element is added to a `std::vector`, it's possible that the
9 `std::vector` lacks space for it, i.e., that the `std::vector`'s size is equal to its ca-
10 pacity. 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 ob-
14 jects 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 old memory were destroyed until all elements had been success-
18 fully 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 unless
30 it's known that the move operations won't emit exceptions. In that case, having

1 moves replace copies would be safe, and the only side effect would be improved
2 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 only if the move operations are known to not emit exceptions. But
9 how can a function know if a move operation won’t produce an exception? The an-
10 swer is obvious: it checks to see if the operation is declared `noexcept`.[†]

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

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

[†] 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 23), depending on whether the type’s move constructor is `noexcept`. In turn, `std::move_if_noexcept` consults `std::is_nothrow_move_constructible`, and the value of this type trait (see Item 9) 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` clauses are `noexcept`. Given two
3 arrays of `Widget`, for example, swapping them is `noexcept` only if swapping indi-
4 vidual elements in the arrays is `noexcept`, i.e., if `swap` for `Widget` is `noexcept`.
5 The author of `Widget`'s `swap` thus determines whether swapping arrays of `Widget`
6 is `noexcept`. That, in turn, determines whether other swaps, such as the one for
7 arrays 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` should motivate you to offer `noexcept` `swap` functions whenever you can.

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, yet 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 kindled your desire to change the imple-
24 mentation 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 exception-neutral function allows the emitted exception to pass
28 through on its way to a handler further up the call chain. Exception-neutral func-
29 tions are never `noexcept`, because they may emit such "just passing through" ex-
30 ceptions. Most functions, therefore, quite properly lack the `noexcept` designation.

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` can
3 have such a significant payoff, it's worth implementing them in a `noexcept` man-
4 ner if at all possible.† When you can honestly say that a function should never emit
5 exceptions, you should definitely declare it `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. For example, callers may have to check for status codes or special return
15 values. The runtime cost of those complications (e.g., extra branches, larger func-
16 tions 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 be poor software engineering.

19 For some functions, being `noexcept` is so important, they're that way by default.
20 In C++98, it was considered bad style to permit the memory deallocation functions
21 (i.e., `operator delete` and `operator delete[]`) and destructors to emit excep-
22 tions, and in C++11, this style rule has been all but upgraded to a language rule. By
23 default, all memory deallocation functions and all destructors—both user-defined
24 and compiler-generated—are implicitly `noexcept`. There's thus no need to declare
25 them `noexcept`. (Doing so doesn't hurt anything, it's just unconventional.) The

† The interface specifications 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 aren't thrown, implementers often declare the operations `noexcept`, even though the Standard does not require them to do so.

1 only time a destructor is not implicitly `noexcept` is when a data member of the
2 class (including inherited members and those contained inside other data mem-
3 bers) is of a type that expressly states that its destructor may emit exceptions (e.g.,
4 declares it “`noexcept(false)`”). Such destructors are uncommon. There are none
5 in the Standard Library, and if the destructor for an object being used by the
6 Standard Library (e.g., because it’s in a container or was passed to an algorithm)
7 emits an exception, the behavior of the program is undefined.

8 It’s worth noting that some library interface designers distinguish functions with
9 *wide contracts* from those with *narrow contracts*. A function with a wide contract
10 has no preconditions. Such a function may be called regardless of the state of the
11 program, and it imposes no constraints on the arguments that callers pass it.†
12 Functions with wide contracts never exhibit undefined behavior.

13 Functions without wide contracts have narrow contracts. For such functions, if a
14 precondition is violated, results are undefined.

15 If you’re writing a function with a wide contract and you know it won’t emit excep-
16 tions, following the advice of this Item and declaring it `noexcept` is easy. For func-
17 tions with narrow contracts, the situation is trickier. For example, suppose you’re
18 writing a function `f` taking a `std::string` parameter, and suppose `f`’s natural
19 implementation never yields an exception. That suggests that `f` should be declared
20 `noexcept`.

21 Now suppose that `f` has a precondition: the length of its `std::string` parameter
22 doesn’t exceed 32 characters. If `f` were to be called with a `std::string` whose
23 length is greater than 32, behavior would be undefined, because a precondition
24 violation *by definition* results in undefined behavior. `f` is under no obligation to
25 check this precondition, because functions may assume that their preconditions

† “Regardless of the state of the program” and “no constraints” doesn’t legitimize programs whose behavior is already undefined. For example, `std::vector::size` has a wide contract, but that doesn’t require that it behave reasonably if you invoke it on a random chunk of memory that you’ve cast to a `std::vector`. The result of the cast is undefined, so there are no behavioral guarantees beyond that point.

1 are satisfied. (Callers are responsible for ensuring that such assumptions are val-
2 id.) Even with a precondition, then, declaring `f noexcept` seems appropriate:

```
3 void f(const std::string& s) noexcept;    // precondition:  
4                                         // s.length() <= 32
```

5 But suppose that `f`'s implementer chooses to check for precondition violations, at
6 least in debug builds. Checking isn't required, but it's also not forbidden, and
7 checking the precondition could be useful during system testing. Debugging an ex-
8 ception that's been thrown is generally easier than trying to track down the cause
9 of undefined behavior. But how should a precondition violation be reported such
10 that a test harness could detect it? A straightforward approach would be to throw
11 a "precondition was violated" exception, but if `f` is declared `noexcept`, that would
12 be impossible; throwing an exception would lead to program termination. For this
13 reason, library designers who distinguish wide from narrow contracts generally
14 reserve `noexcept` for functions with wide contracts.

15 As a final point, let me elaborate on my earlier observation that compilers typically
16 offer no help in identifying inconsistencies between function implementations and
17 their exception specifications. Consider this code, which is perfectly legal:

```
18 void setup();                // functions defined elsewhere  
19 void cleanup();  
  
20 void doWork() noexcept  
21 {  
22     setup();                 // set up work to be done  
23     ...                      // do the actual work  
24     cleanup();              // perform cleanup actions  
25 }
```

26 Here, `doWork` is declared `noexcept`, even though it calls the non-`noexcept` func-
27 tions `setup` and `cleanup`. This seems contradictory, but it could be that `setup`
28 and `cleanup` document that they never emit exceptions, even though they're not
29 declared that way. There could be good reasons for their non-`noexcept` declara-
30 tions. For example, they might be part of a library written in C. (Even functions
31 from the C Standard Library that have been moved into the `std` namespace lack
32 exception specifications, e.g., `std::strlen` isn't declared `noexcept`.) Or they

1 could be part of a C++98 library that decided not to use C++98 exception specifica-
2 tions and hasn't yet been revised for C++11.

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

6 **Things to Remember**

- 7 ♦ `noexcept` is part of a function's interface, and that means that callers may de-
8 pend on it.
- 9 ♦ `noexcept` functions are more optimizable than non-`noexcept` functions.
- 10 ♦ `noexcept` is particularly valuable for the move operations, `swap`, memory
11 deallocation functions, and destructors.
- 12 ♦ Most functions are exception-neutral rather than `noexcept`.

13