

1 **Item 19: Declare functions noexcept whenever possible.**

2 In C++98, exception specifications were rather temperamental beasts. You had to
3 summarize the exception types a function might emit, and this imposed con-
4 straints on the function's implementation. If the implementation was changed, the
5 exception specification might have to be updated, too, and that not only opened
6 the door to consistency errors (i.e., an exception specification that no longer corre-
7 sponded to the implementation), it also meant that callers might stop compiling,
8 because an exception specification is part of a function's interface. For these rea-
9 sons, C++98 exception specifications never gained much popularity. Developers
10 and libraries generally shied away from them, and some compilers didn't even ful-
11 ly implement them.

12 Over time, a consensus emerged that the only meaningful information about a
13 function's exception-emitting behavior was whether it had any. Black or white:
14 either a function might emit an exception or it guaranteed that callers would never
15 see one. 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 continue to be legal in C++11 and C++14, but they're deprecated.) In C++11,
18 `noexcept` is for functions that guarantee they'll never emit an exception. When
19 you write a function that can make that guarantee, you'll want to use `noexcept`.
20 Why? Because it permits compilers to generate better code.

21 There are two reasons for this, but we'll begin with how C++98's way of saying
22 "this function emits no exceptions" differs from C++11's.

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

```
25 int f(int x) throw();           // C++98 approach: f emits no  
26                               // exceptions  
27 int f(int x) noexcept;        // C++11 approach: f emits no  
28                               // 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, `f` could be imple-
3 mented like this:

```
4 int f(int x) noexcept          // C++98 version would use "throw()"
5 {
6     if (x >= 0) return x * x - 42;          // 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 absurd, but it's perfectly legal C++. Furthermore, looks aren't every-
12 thing. The code here could be a way of enforcing the precondition that `x` must be
13 non-negative. If `f` is called with a legitimate value, it doesn't throw. However, if an
14 invalid value is passed in, the precondition violation causes the function to have
15 undefined behavior, and this implementation uses that freedom—undefined be-
16 havior means that *anything* can happen—to throw an exception.

17 As an aside, 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`, `doubles`, etc.), but the naming
22 convention for those functions, albeit following a consistent pattern, is rather cryp-
23 tic: `stoi`, `stol`, `stod`, etc. The C++11 Standard Library also offers `std::wstring`-
24 based versions of all these functions.

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

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

28 If, at runtime, an exception leaves `f`, `f`'s exception specification is violated. With
29 the C++98 approach, the call stack is unwound to `f`'s caller, then the *unexpected*
30 *handler* function is invoked, and that leads to program termination (typically by
31 calling `std::terminate`). With the C++11 approach, runtime behavior is slightly

1 different: the stack is *possibly* unwound, then program execution is terminated
2 (typically by calling `std::terminate`).

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

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

17 This alone should provide sufficient motivation to declare functions `noexcept`
18 whenever you can. For some functions, however, the case is even stronger.

19 The move operations are the preeminent example. Suppose you have a large in-
20 vestment in a C++98 code base making use of `std::vectors of Widgets`. `Widgets`
21 are added to the `std::vectors` from time to time, perhaps via `push_back`:

```
22 std::vector<Widget> vw;
23 ...
24 Widget w;
25 ...           // put w into proper state
26               // for addition to vw
27 vw.push_back(w);    // add w to vw
28 ...
```

29 Assume this code works fine, and you have no interest in modifying it for C++11.
30 However, you do want to take advantage of the fact that C++11’s move semantics

1 can improve the performance of existing code when move-enabled types are in-
2 volved. You therefore ensure that `Widget` has move operations, either by writing
3 them yourself or by seeing to it that the conditions for their automatic generation
4 are fulfilled (see Item 20).

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

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

24 This is a serious problem, because the behavior of your C++98 code base could de-
25 pend on `push_back`'s strong exception safety guarantee. C++11 implementations
26 therefore can't silently replace copy operations inside `push_back` with moves.
27 They must continue to employ copy operations. *Unless*, that is, it's known that the
28 move operations are guaranteed not to emit exceptions. In that case, replacing el-
29 ement copy operations inside `push_back` with move operations would be safe,
30 and the only side effect would be improved performance.

1 `std::vector::push_back` takes advantage of this “move if you can, but copy if
2 you must” strategy, and it’s not the only function in the Standard Library that does.
3 Other functions sporting the strong exception safety guarantee in C++98 (e.g.,
4 `std::vector::reserve`, `std::deque::insert`, etc.) behave the same way. All
5 these functions replace calls to copy operations in C++98 with calls to move opera-
6 tions in C++11 if (and only if) the move operations are known to never emit excep-
7 tions. But how can a compiler know if a move operation won’t produce an excep-
8 tion? The answer is obvious: it checks to see if the operation is declared `noex-`
9 `cept`.*

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

12 Inside a function like `push_back`, suppose we want to transfer an object from one
13 location in memory to another (i.e., copy or move it, depending on what is appro-
14 priate). Assume we have an iterator, `src`, referring to the object to be transferred
15 and a second iterator, `dest`, referring to where it should go. Our code would have a
16 statement like this:

```
17 *dest = *src;           // transfer *src to *dest;  
18                        // this is incorrect
```

19 The statement would be inside a loop, because we’d ultimately need to transfer all
20 the objects in the container, but understanding how things work for one object is
21 all we need for this discussion.

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

* 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.

```
1  *dest = std::move(*src);           // transfer *src to *dest;
2                                     // this is incorrect
```

3 This is also incorrect, because now we're moving `*src`, regardless of whether its
4 move assignment operator is `noexcept`. As we've discussed, doing that would
5 prevent `push_back` from maintaining its strong exception safety guarantee, and
6 maintaining that guarantee is essential for ensuring that code written under
7 C++98 continues to function properly.

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

```
10 *dest = std::move_if_noexcept(*src); // transfer *src to *dest;
11                                     // this is correct
```

12 Conceptually, `std::move_if_noexcept` causes `*src` to be moved if its move as-
13 signment operator is `noexcept`, and otherwise it causes `*src` to be copied. That's
14 exactly what we want, and that's why this code is correct. Uses of
15 `std::move_if_noexcept` are scattered throughout strongly exception safe func-
16 tions in the Standard Library, and that's why you have a special incentive to de-
17clare your move operations `noexcept`: it enables their use inside such functions.

18 The conceptual description of `std::move_if_noexcept` deviates from its true
19 behavior in two small ways. First, if `std::move_if_noexcept` is invoked on an
20 object of a move-only type, a move will be performed, even if it might yield an ex-
21 ception. This is understandable: what else can `std::move_if_noexcept` do, giv-
22 en that the type can't be copied? Besides, this behavior can't break legacy code,
23 because there was no such thing as a move-only type in C++98.

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

30 Accompanying the move operations on the list of functions where a `noexcept` dec-
31 laration is especially beneficial is `swap`. Being a heavily-used function, the optimi-

1 zation opportunities that `noexcept` affords are unusually worthwhile. (Many algo-
2 rithms rely on `swap`, as do implementations of many copy assignment operators.)
3 Furthermore, whether particular versions of `swap` in the Standard Library are `no-`
4 `except` is sometimes dependent on whether user-defined type-specific swaps are
5 `noexcept`. For example, the declarations for the Standard Library's swaps for ar-
6 rays and for `std::pair` are:

```
7 template <class T, size_t N>  
8 void swap(T (&a)[N],  
9           T (&b)[N]) noexcept(noexcept(swap(*a, *b)));  
  
10 template <class T1, class T2>  
11 struct pair {  
12     ...  
13     void swap(pair& p) noexcept(noexcept(swap(first, p.first)) &&  
14                                   noexcept(swap(second, p.second)));  
15     ...  
16 };
```

17 These functions are *conditionally* `noexcept`: whether they are `noexcept` depends
18 on whether the expressions inside the `noexcept`s are. Given two arrays of `Widget`,
19 for example, swapping them is `noexcept` only if swapping individual elements
20 from the arrays is `noexcept`, i.e., if `swap` for `Widget` is `noexcept`. The author of
21 `Widget`'s `swap` thus determines whether swapping arrays of `Widget` is `noexcept`
22 (which, in turn, could determine whether other swaps are `noexcept`, e.g., `swap` for
23 arrays of arrays of `Widget`). Similarly, whether swapping two `std::pair` objects
24 containing `Widgets` is `noexcept` depends on whether `swap` for `Widgets` is `noex-`
25 `cept`.

26 The fact that swapping higher-level data structures can generally be `noexcept`
27 only if swapping their lower-level constituents is `noexcept` is the reason why you
28 should strive to offer `noexcept` swap functions.

29 Because `noexcept` is part of a function's interface, you should declare a function
30 `noexcept` only if you are willing to commit to a `noexcept` implementation over
31 the long term. If you declare a function `noexcept` and implement it accordingly,
32 then later decide you wish you hadn't made the `noexcept` promise, your options
33 are bleak. You can remove `noexcept` from the function's declaration, thus running
34 the risk of breaking arbitrary amounts of client code. You can retain the `noexcept`

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

1 declaration, but change the implementation such that an exception could actually
2 escape. In that case, if an exception did escape at runtime, your program would be
3 terminated. Or you can retain your existing implementation, thus abandoning
4 whatever motivated you to want to change the implementation in the first place.
5 None of these options is appealing.

6 Most functions are *exception-neutral*: they don't throw exceptions themselves, but
7 if a function they call produces one (directly or indirectly), the exception causes no
8 harm as it passes through on its way to a handler in a different function. Exception-
9 neutral functions aren't `noexcept`, because exceptions may pass through them.
10 Most functions, therefore, aren't `noexcept`.

11 Some functions, however, are naturally `noexcept`, and for a few more—notably
12 the move operations and `swap`—being `noexcept` has such a significant payoff, it's
13 worth implementing them in a `noexcept` manner if at all possible. When you can
14 honestly say that a function should never emit exceptions, you should definitely
15 declare it `noexcept`.

16 **Things to Remember**

- 17 ♦ `noexcept` functions offer more optimization opportunities than non-
18 `noexcept` functions.
- 19 ♦ C++98 functions offering the strong exception safety guarantee may internally
20 call `std::move_if_noexcept` instead of `std::move`.
- 21 ♦ Strive to declare the move operations and `swap` `noexcept`.

22