

# Contracts on lambdas

Timur Doumler ([papers@timur.audio](mailto:papers@timur.audio))

Document #: P2890R2  
Date: 2023-12-13  
Project: Programming Language C++  
Audience: SG21

## Abstract

This paper proposes to allow `pre` and `post` specifiers on lambda expressions. The proposal is straightforward except for the semantics of lambda captures triggered by the odr-use of entities inside a contract predicate. We consider several alternatives for how to resolve this issue.

## 1 Proposal

In the Contracts MVP (see [P2900R2]), `contract_assert` is already allowed inside a lambda expression (because it is allowed anywhere an expression of type `void` is allowed), but `pre` and `post` are only allowed on the declaration of ordinary functions. However, there is no a priori reason why we should not allow them to appear also on lambda expressions. We propose that the following code should be well-formed:

```
auto f = [] (int i)
  pre (i > 0)      // proposed; ill-formed in MVP
  post (r: r > i) // proposed; ill-formed in MVP
{
  contract_assert (s > 0); // already allowed in MVP
  return i + s;
};
```

The proposed syntax grammar for adding `pre` and `post` has already been specified in [P2961R2]: the sequence of `pre` and `post` specifiers is always the last part of the declaration before the opening brace of the lambda body.

Regarding name lookup and access control for the predicate in `pre` and `post` on a lambda, we propose that it should follow the existing MVP rule for functions (see [P2900R2] section 2.3.2): it should work as-if these predicates were at the start of the lambda's body. This rule is well-motivated and applies equally to this case.

## 2 History

C++20 Contracts [P0542R5] did not allow `pre` and `post` on lambda expressions because it was using attribute syntax, and attributes could not be meaningfully applied to lambdas at the time:

Contracts attributes (as any other attribute in that syntactical location) appertain to the function type. However, they are not part of the function type. Note that this does not solve the issue of being able to use attributes on lambda expressions (see Core issue 2097). In fact, until that issue is resolved it will not be possible to specify preconditions and postconditions for lambda expressions.

The cited issue was later resolved by adopting [P2173R1] for C++23, and is no longer relevant for the Contracts MVP because we are no longer using attribute syntax. The more recent [P2388R4] also did not allow `pre` and `post` on lambda expressions, but for another reason:

These features are deferred due to unresolved issues: [...] a way to express preconditions and postconditions for lambdas: name lookup is already problematic in lambdas in the face of lambda captures. This problem is pursued in [P2036R1], and until it has been solved we see no point in delaying the minimum contract support proposal.

The cited issue was later resolved by adopting [P2036R3] for C++23, but in fact is also not relevant. The issue solved by [P2036R3] was about name lookup in the trailing return type of a lambda. Name lookup in `pre` and `post` is fundamentally different, because identifiers in these predicates should refer to captured entities accessible in the lambda's body, not to entities in the scope containing the declaration of the lambda: those entities will no longer be available by the time the lambda is called and the preconditions and postconditions are checked. The existing MVP rule for functions (as-if at the start of the body) is straightforward and gives us the desired semantics for lambdas.

### 3 Captures

Entities in the predicate of a contract annotation are unconditionally odr-used (see [P2900R2]), regardless of whether the contract annotation is checked or ignored. This must be so because the contract semantic is in general unknown at compile time (see [P2877R0]).

Such odr-use can trigger lambda captures. Therefore, if we let these language features compose naturally and do not introduce any special rules, then a contract annotation on a lambda can trigger a lambda capture if it odr-uses an entity not odr-used anywhere else. This can have observable effects both at runtime and at compile time. Here is an example of a runtime effect (taken from [P2932R2]):

```
std::function<int()> g(const std::vector<S>& v)
{
    int ndx = pickIndexAtRandom(v);
    return [=]()
        pre (0 <= ndx && ndx < v.size()) // needs to capture v
        {
            return ndx; // Obviously we intend this to capture ndx by value.
        };
}
```

Here, the lambda capture triggered by the precondition predicate triggers an expensive copy of the whole vector `v`. Without the precondition specifier being present, the copy would not happen.

Here is an example of a compile-time effect due to a precondition predicate on a lambda triggering a lambda capture:

```
constexpr auto f(int i) {
    return sizeof( [=] pre (i > 0) {} ); // captures i by value
}
```

The same effect would occur with an assertion inside the lambda body:

```
constexpr auto f(int i) {
    return sizeof( [=] { contract_assert (i > 0); } ); // captures i by value
}
```

In both examples above, without the contract annotation, `f` would return 1, but with the contract annotation added, `f` will return `sizeof(int)`, even if the contract semantic is *ignore*. We can construct a case where merely adding the contract annotation to a lambda triggers the layout of a class to change:

```
struct X {
    char data[f(0)];
};
```

We can even construct an example where the lambda capture triggered by the contract annotation may change which function is selected by overload resolution (taken from [P2932R2]):

```
template <typename T>
std::true_type f(T t);

template <typename T>
std::false_type f(T t)
requires std::is_trivially_copyable_v<T>;

int main() {
    std::shared_ptr<int> p;

    auto f1 = [=]{ return true; };
    static_assert(decltype(f(x))::value == false); // passes

    auto f2 = [=]{ contract_assert(p); return true; }; // just add an assertion
    static_assert(decltype(f(x))::value == false); // fails
}
```

Strictly speaking, it would be conforming for the first `static_assert` to fail, because it is not guaranteed that the lambda is trivially copyable, however it is trivially copyable on every major C++ compiler. With the `contract_assert` added, the lambda *must* be non-trivially copyable if `p` is captured, changing overload resolution. Should we allow this to happen?

## Option 1: Contract predicates can trigger captures

The first option is to not treat contract predicates differently from any other expressions in C++ for the purpose of triggering lambda captures, i.e. to allow the implicit captures to happen in all cases shown above. This behaviour is arguably the most straightforward, most consistent with the rest of the language, and least surprising to the user: it simply falls out of the current rules for odr-use and lambda captures. It is also consistent with assumptions, which behave in the same way:

```
constexpr auto f(int i) {
    return sizeof( [=] { [[assume (i > 0)]]; } ); // captures i by value
}
```

This phenomenon has been extensively discussed when assumptions were standardised for C++23 (see [P1774R8] section 4.4). While there were concerns that `[[assume]]` should not be able to change the layout of a class, ultimately EWG decided it was too much of an edge case unlikely to cause problems in real code to justify complicating the rules of the language to make such a capture ill-formed. The same reasoning can be applied to contract annotations.

In the case where the lambda capture from a contract predicate triggers an expensive copy of the vector, there is a straightforward workaround to avoid the expensive copy: do not use a default-capture, but only capture `v.size()` explicitly, or better still, perform the check outside of the lambda. In the case where the lambda capture from a contract predicate changes the layout of

a class, the layout of the class was non-portable in the first place, so user code should not rely on it. Similarly, user code that relies on a lambda *with a default capture* being trivially copyable is doing something that the Standard explicitly calls out as something no code should rely on (see [\[expr.prim.lambda.closure\]](#) p. 3). In all these cases, the user wrote bad code and they got what they asked for. It is completely reasonable to say that we do not want to bend over backwards and make the language more complicated in order to make such cases ill-formed.

## Option 2: Contract predicates do not trigger captures

[\[P2932R2\]](#) proposes that the odr-use of a local entity in a contract predicate should *not* implicitly capture that entity. It argues that allowing the capture would violate the so-called *zero overhead principle*: adding a contract annotation should never lead to a different branch being taken at compile time. Otherwise, this could lead to “heisenbugs” (a program contains a bug, we add a contract annotation to find the bug, but the contract annotation causes the program to take another branch where the bug does not exist) and measurable performance degradations. The paper argues that the existence of cases where the addition of a contract annotation can cause such “heisenbugs” and performance degradations, no matter how obscure these cases are, would be a major disincentive for the adoption of Contracts in C++.

Saying that the addition of contract annotations to a program should never modify the compile-time semantics of the program, such as overload resolution — effectively checking a different program — is a principled and reasonable approach. However, specifying that the capture simply does not happen, as proposed in [\[P2932R2\]](#), leads to the surprising consequence that the wrong variable might end up being used:

```
static int i = 0;

void test() {
    int i = 1;
    auto f = [=] { contract_assert(i > 0); }; // which i is referenced here?
}
```

With the proposal in [\[P2932R2\]](#), the parameter `i` would not be captured, which means that the `i` in the contract predicate would refer to the outer `i` variable which has static storage duration and therefore does not need to be captured in order to be odr-used inside the lambda. We believe that allowing the above behaviour would be very confusing to the user, and that therefore the proposal in [\[P2932R2\]](#) is not viable.

## Option 3: Triggering a capture from a contract predicate is ill-formed

The better alternative to option 2, avoiding the problem of unintentionally referring to the wrong variable inside the lambda, is to say that if the predicate of a contract annotation triggers a lambda capture of an entity not otherwise captured, the program is unconditionally ill-formed:

```
int i = 0;
auto f(int i) {
    return sizeof( [=] { contract_assert(i > 0); } ); // Error: cannot capture i here
}
```

This satisfies the zero overhead principle for contracts on lambdas and fixes the issue with option 2. We therefore consider option 3 viable. The downsides compared to option 1 are that carving out such an exception for how lambda captures work increases the complexity of the language, and that the resulting compile error might not be obvious for a novice user. Further, if we choose this option, we strongly recommend that we introduce the same rule for assumptions, as a DR against C++23, otherwise option 3 would introduce an unfortunate inconsistency.

#### Option 4: Triggering a capture from a contract predicate emits a warning

Another alternative to address this issue is to not make triggering a lambda capture from a contract predicate ill-formed, but to issue a compiler warning. In C++, there are many other cases where this approach is taken. Consider:

```
std::map<int, Widget> map = { /* ... */ };
for (const std::pair<int, Widget>& elem : map)
    // do something with elem
```

In this case, the user got the element type of `std::map` wrong (which is `std::pair<const int, Widget>` rather than `std::pair<int, Widget>`); this generates an unintended implicit conversion, which in turn yields a temporary object that is lifetime-extended by the `const&`. This code compiles and works, but has a silent performance degradation due to the unnecessary conversion and object creation on every iteration of the loop. Such inefficiency is unfortunate; however, we do not add special cases to basic language rules such as range-based `for` loops, implicit conversions, or reference semantics to make these cases ill-formed. Instead, the user gets what they get, a quality compiler or static analysis tool will issue a warning, and several straightforward fixes are available (use the correct type, or just use `auto`), just like in the lambda case we are concerned with here.

We could even go further and make the warning mandated by the standard, rather than just a matter of QoI; we created a precedent for having well-formed programs that nevertheless require a diagnostic by standardising `#warning`, and we might standardise another such case if we adopt the compile-time-*observe* semantic from [P2894R1].

However, triggering implicit captures from contract predicates is arguably somewhat different from the `std::map` example. As we saw above, triggering a capture from a contract predicate is *always* suboptimal and questionable code, not just in certain cases. It therefore seems that, if we do not want to allow the user to just “get what they get” (option 1) in those cases, then making them ill-formed (option 3) is a better approach than issuing a warning (option 4). Option 3 is a more effective and more teachable deterrent to writing such code and we do not have to worry about false positives.

#### Option 5: odr-using an entity in a contract predicate that is not otherwise odr-used is ill-formed

Triggering lambda captures is actually not the only way in which the odr-use of an entity in a contract predicate can violate the zero overhead principle from [P2932R2]. Another way is through triggering a template instantiation, which in turn can change observable state via friend injection. This case might be negligible because it is an obscure wart of the C++ language that we would like to make ill-formed (see [CWG2118]). But there is a third and much more plausible case where odr-use of an entity in a contract annotation can violate the zero overhead principle: a template instantiation can trigger dynamic initialisation of a static class member, which in turn can have arbitrary effects on the observable behaviour of the program:

```
struct X {
    X() { std::cout << "This is a bug!\n"; }
    operator bool() const;
};

template <typename T>
struct Y {
    static const X x;
};

template <typename T>
const X Y<T>::x = X{};
```

```

void f()
  pre(Y<int>::x); // the addition of this contract annotation triggers the bug!

int main() {
  f();
}

```

Therefore, making it ill-formed for a contract predicate to trigger a lambda capture (option 3) would not actually be a comprehensive fix that ensures odr-use of entities inside a contract predicate never violate the zero overhead principle. Such a comprehensive fix would be to make any program ill-formed where a contract predicate odr-uses an entity not otherwise odr-used (and not just if that odr-use triggers an implicit lambda capture). This is our option 5.

We believe that this option 5 is impractical, as it would prohibit the user from e.g. using STL algorithms like `std::for_each` in a contract annotation, unless the exact same instantiation of `std::for_each` is already used somewhere else in the program. Further, it is not obvious how this option could be implemented as such odr-use might occur in a different translation unit.

Note that triggering a template instantiation, unlike triggering a lambda capture, while arguably a violation of the zero overhead principle, does not seem to pass the litmus test in Appendix A of [P2932R2] as it cannot affect overload resolution. It therefore seems like a less serious violation of the zero overhead principle than the lambda capture case.

### Option 6: Do not allow contracts on lambdas with default captures

Since the issue discussed here only occurs with lambda expressions that have default captures, we could also say that contract annotations should not be allowed on such lambdas. Note that this would affect not only `pre` and `post`, for which the possibility to add them to lambdas is being proposed in this paper, but also any `contract_assert` in the body of the lambda, which the Contracts MVP permits today.

We believe that this solution is too drastic as it would take away too much value for the user, therefore we do not consider this solution viable.

### Option 7: Do not allow contracts on lambdas

As the most radical solution to the problem, we could say that contract annotations are not allowed on any lambdas (`pre` and `post`) or inside the body of any lambdas (`contract_assert`). We do not consider this solution viable for the same reason as option 6.

## 4 Summary

We propose to allow adding `pre` and `post` specifiers to lambda expressions in the Contracts MVP [P2900R2], with the syntax proposed in [P2961R2]. Regarding name lookup and access control for the predicate in `pre` and `post` on a lambda, we propose that it should follow the existing MVP rule for functions (see [P2900R2] section 2.3.2): it should work as-if these predicates were at the start of the lambda's body.

Entities in a contract predicate are odr-used, even if the contract annotation has the *ignore* semantic. This odr-use inside a lambda can trigger an implicit lambda capture. It follows therefore that adding a contract annotation to a program can trigger a lambda capture that would otherwise not happen, which in turn can change the compile-time semantics of the program. This violates a desirable property called the *zero overhead principle* (see [P2932R2]). To address this issue, we considered the following seven options:

1. Contract predicates can trigger captures (the default if we do not add any special rule).
2. Contract predicates do not trigger captures (proposed by [P2932R2]).
3. Triggering a capture from a contract predicate is ill-formed.
4. Triggering a capture from a contract predicate emits a warning.
5. odr-using an entity in a contract predicate that is not otherwise odr-used is ill-formed.
6. Do not allow any contract annotations on lambdas with default captures
7. Do not allow any contract annotations on any lambdas

For the reasons discussed in this paper, we believe that the only viable options are 1 and 3. Option 1 is the correct choice if we prefer to not add additional complexity to C++, and accept that one can construct cases where the zero overhead principle from [P2932R2] might be violated by captures from entities inside contract predicates (with straightforward fixes available for such cases). Option 3 is the correct choice if there should be *no possible case* where contract predicates would capture entities in a way that violates the zero overhead principle from [P2932R2].

Note that there are other circumstances, unrelated to lambda captures, where odr-use inside a contract predicate can still violate the zero overhead principle, in a way not addressed by option 3. Making these violations impossible would require the more drastic option 5, which would have unfortunate consequences. However, arguably the additional cases addressed by option 5 are less serious violations of the zero overhead principle.

If we choose option 3, we strongly recommend that we introduce the same rule for triggering captures from `[[assume(x)]]`, as a DR against C++23, to be consistent.

## Document history

- **R0**, 2023-03-08: Initial version
- **R1**, 2023-12-07: Considering more options for resolving the captures issue
- **R2**, 2023-12-13: Fixed nonsensical name lookup rules

## Acknowledgements

Thanks to Ville Voutilainen and Joshua Berne for pointing out that the name lookup rules proposed in the previous revision (R1) of this paper were nonsensical.

## References

- [CWG2118] Richard Smith. Core Issue 2118: Stateful metaprogramming via friend injection. <https://cplusplus.github.io/CWG/issues/2118.html>, 2022-06-27.
- [P0542R5] G. Dos Reis, J. D. Garcia, J. Lakos, A. Meredith, N. Myers, and B. Stroustrup. Support for contract based programming in C++. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2018/p0542r5.html>, 2018-06-08.
- [P1774R8] Timur Doumler. Portable assumptions. <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2022/p1774r8.pdf>, 2022-06-14.

- [P2036R1] Barry Revzin. Change scope of lambda *trailing-return-type*. <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2021/p2036r1.html>, 2021-01-13.
- [P2036R3] Barry Revzin. Change scope of lambda *trailing-return-type*. <https://www.open-std.org/jtc1/sc22/wg21/docs/papers/2021/p2036r3.html>, 2021-09-14.
- [P2173R1] Inbal Levi, Daveed Vandevoorde, and Ville Voutilainen. Attributes on Lambda-Expressions. <https://wg21.link/p2173r1>, 2021-12-09.
- [P2388R4] Andrzej Krzemieński and Gašper Ažman. Minimum Contract Support: either *No\_eval* or *Eval\_and\_abort* contracts. <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2021/p2388r4.html>, 2021-11-15.
- [P2877R0] Joshua Berne and Tom Honermann. Contract Build Modes, Semantics, and Implementation Strategies. <https://wg21.link/p2877r0>, 2023-06-09.
- [P2894R1] Timur Doumler. Constant evaluation of Contracts. <https://wg21.link/p2894r1>, 2023-12-04.
- [P2900R2] Joshua Berne, Timur Doumler, and Andrzej Krzemieński. Contracts for C++. <https://wg21.link/p2900r2>, 2023-11-11.
- [P2932R2] Joshua Berne. A Principled Approach to Open Design Questions for Contracts. <https://wg21.link/p2932r2>, 2023-11-14.
- [P2961R2] Timur Doumler and Jens Maurer. A natural syntax for Contracts. <https://wg21.link/p2961r2>, 2023-11-08.