

D1095R0/N2xxx draft 4: Zero overhead deterministic failure

– A unified mechanism for C and C++

Document #: WG14 N2xxx draft 4
WG21 D1095R0 draft 4
Date: 2018-08-08
Project: WG14 Programming Language C
WG21 Programming Language C++
WG21 Evolution Working Group
WG21 SG14 Low Latency study group
Reply-to: Niall Douglas
<s_sourceforge@nedprod.com>

A proposed universal mechanism for enabling C speaking programming languages to tell C code, and potentially one another, about failure and disappointment.

This proposal is being concurrently submitted to both WG14 (C) and WG21 (C++), as it is proposed that C++'s proposed deterministic exception throws [P0709] be implemented using this universal C mechanism.

Changes since draft 3:

- Dropped section on Rust for brevity.
- Rebased onto `_Fails` syntax as per advice from the WG14 reflector.
- Converted `_Either(A, B)` from a meta-type into an ordinary sum type as per WG14 reflector suggestion.
- Added `_Fails(errno)` to solve the `errno` impurity problem as suggested by the WG14 reflector.

Contents

1	Introduction	2
2	Proposed Design: C	3
2.1	Purifying <code>errno</code> setting functions	5
2.1.1	What happens to <code>errno</code> when a <code>_Fails(T)</code> function calls a <code>_Fails(errno)</code> function?	7
3	Proposed Design: C++	7
3.0.1	Background: C-compatible types in C++	8
3.0.2	Background: How C++ currently implements exceptions	9
3.0.3	Background: <code>std::expected<T, E></code>	10
3.0.4	Background: The current proposed mechanism for deterministic C++ exceptions	10

3.0.5	Background: The alternative proposed mechanism for deterministic C++ exceptions	11
3.0.6	Background: <code>std::error_code</code>	12
3.0.7	Background: <code>std::exception_ptr</code>	13
3.0.8	How C++ function linkage would work using proposed <code>_Fails(X)</code>	14
3.0.9	C++ calling <code>_Fails(errno)</code> C functions	15
4	Frequently Asked Questions	15
4.1	How interchangeable for one another are the proposed C and C++ facilities?	15
5	Acknowledgements	16
6	References	16

1 Introduction

This is an unusual paper, being submitted concurrently to both to WG14 (C programming language) as well as WG21 (C++ programming language). It originated from WG21 SG14 the Low Latency study group, where most of our group would tend to work more frequently with C-ish code than other WG21 members. This has influenced our guidelines for writing fixed latency C++ libraries, see [P1027] *SG14 Design guidelines for latency preserving (standard) libraries*.

C++ is considering making its exception handling implementation value-based instead of type-based in order to make it deterministic, see [P0709] *Zero-overhead deterministic exceptions: Throwing values*. We could come up with some local-to-C++ solution, but I think we can do a lot better by closing the gap between the C++ ABI and the C ABI such that C code will be able to legally call, *without intermediate bindings*, C++ functions capable of throwing value-based exceptions (so long as such functions do not receive nor return C-incompatible types). C code would directly understand that a C++ function it called had failed – moreover, C code would be able to return exception throws to C++ code.

If this coordination can be pulled off, the benefits could be profound for all C speaking programming languages, for example Rust, Python or Fortran.

However there are substantial benefits for C, as well:

1. The use of `errno` to indicate cause of failure introduces a lot of unhelpful side effects both in C and C++, because it prevents some functions (the `<tgmath.h>` functions especially) being marked as pure, which in turn inhibits more aggressive optimisation.
2. It introduces a hard dependency on thread local storage, which in turn forces million thread compute resources to not implement standard C for their compilers, as transporting thread local storage to each CPU is impractical.
3. And finally, this proposed mechanism makes it much easier for C to import Contracts from C++, as a number of people on the WG14 reflector appear keen to do, as the contract violation handler can do something other than abort the program, which is very useful for unit test code ensuring that contract violation is detected.

2 Proposed Design: C

It is proposed that C add the following:

- A new function declaration modifier `_Fails(X)` where `X` is either a type, or `errno`. This modifies the type of the function to indicate that it has a different calling convention to ordinary C functions, specifically that now two channels of function return are possible, *successful* and *failure*. Successful is the ordinary return type of the function (if any). Failure is the type `X`.
If `errno` is also specified in the failure declaration, the reading and writing of `errno` does not affect the thread locally stored `errno`. See section below specifically dedicated to this topic.
- A `_Failure(expr)` used to tell the compiler to return the expression via the failure return channel, rather than via the success return channel.
- A sum type `_Either(A, B)` which stores either an `A` type in `.left`, or a `B` type in `.right`, and has a boolean `.positive` which is true if the sum type has an `A` type stored¹.
- A `_Catch(expr)` which converts the result of executing the expression into an `_Either(A, B)` where an `A` is stored if the expression evaluated successfully to a value of type `A`, or a `B` is stored if any part of the expression failed with a value of type `B`².
- A `_Try(expr)` is a convenience piece of boilerplate equivalent to:

```
1  _Either(A, B) __unique_temporary = _Catch(expr);
2  if(!__unique_temporary.positive)
3  {
4      return _Failure(__unique_temporary.right);
5  }
```

Writing the above boilerplate is very common in code with success and failure return channels, and saving having to write it out by hand every time is user friendly.

One might wonder why not return `_Either(A, B)` directly from the function? One can do this of course, however the `_Fails` approach has four advantages:

1. The C calling convention has the caller allocate the space for the returned value before calling a function. `_Either(A, B)` can store `A` and `B` in the same union storage, however an additional *discriminant* must also be stored, which could take up to eight additional bytes or worse depending on packing.

For `_Fails` returns, it is proposed for at least AArch64, ARM, x86 and x64, that the discriminant be the CPU's carry flag. This is because compilers can often fold the setting or clearing of the CPU's carry flag into the ordering of other operations, thus making this a zero runtime overhead choice of discriminant³. On other architectures such as RISC-V (which has no status

¹This is the opposite orientation to the sum type in functional languages, but the same orientation to C++'s `expected<T, E>` and Rust's `result<T, E>`. I chose this orientation as it seems to be the de facto for systems languages.

²If the expression calls functions which return dissimilar failure types, the failure type of the first function evaluated is chosen. The failure types of all other functions must be convertible into this type.

³As it is a single bit being branched upon, status register update pipeline stalls should not occur.

register), an extra register would make more sense. It doesn't matter what an architecture chooses, so long as it is *consistent* across all compilers.

2. We can do magic with `errno`, see later section.
3. Except for `_Fails(errno)`, calling a `_Fails(X)` function without wrapping it in either `_Try(expr)` or `_Catch(expr)` is a compile time error. This requires the caller to explicitly specify how to handle failure. We could not be so enforcing with returns of `_Either(A, B)`.
4. Returning an `_Either(A, B)` from a function capable of failure will likely be common enough in real world code. One could still do `_Either(_Either(A, B), E) func();` of course, but the key thing to note here is that `_Either(A, B)` is a generalised C sum type with little to do, per se, with handling failure. I don't wish to mention the 'm word'⁴ in this paper, but in a lot of code there is A or B or abort function call outcomes.

An example of the proposed syntax:

```
1 int some_function(int x) _Fails(float)
2 {
3     // Return failure if x is zero
4     return (x != 0) ? 5 : _Failure(2.0);
5 }
6
7 const char *some_other_function(int x) _Fails(float)
8 {
9     // If calling some_function() fails, return its failure immediately
10    // as if by return _Failure(some_function(x).right)
11    int v = _Try(some_function(x));
12
13    return (v == 5) ? "Yes" : "No";
14 }
15
16 int main(int argc, char *argv[])
17 {
18     if(argc < 2)
19         abort();
20
21     _Either(const char *, float) v = _Catch(some_other_function(atoi(argv[1])));
22
23     if(v.positive)
24     {
25         printf("v is a successful %s\n", v.left);
26     }
27     else
28     {
29         printf("v is failure %f\n", v.right);
30     }
31     return 0;
32 }
```

And running this program would yield:

```
ned@lyta:~/windocs/boostish/wg21$ ./test 0
```

⁴monad.

```
v is a failure 2.000000
ned@lyta:~/windocs/boostish/wg21$ ./test 1
v is a successful Yes
```

2.1 Purifying `errno` setting functions

A long time ago POSIX chose to return the cause of failure of a function by setting a then-global variable `errno`. This has since become a thread-local variable, but the side effects of calling functions which modify `errno` impedes optimisation, and makes impossible the use of standard C in very large CPU core machines.

This particularly affects the C math functions. Let us take a simple example, making an integer positive:

```
1 int abs(int x)
2 {
3     if(x == INT_MIN)
4     {
5         errno = ERANGE;
6         return 0;
7     }
8     return (x < 0) ? -x : x;
9 }
```

This sets `errno` if it fails with the cause of failure, which prevents us from marking this functions with `__attribute__((pure))`, or even `__attribute__((const))`. This means that the compiler must assume that calling these functions has side effects on global state. That in turn greatly limits the ability of the optimiser to perform common subexpression elimination and loop optimization on code containing math functions, which is unfortunate.

I propose to fix this via the `_Fails(errno)` modifier, so:

```
1 int abs(int x) _Fails(errno) __attribute__((const))
2 {
3     if(x == INT_MIN)
4     {
5         errno = ERANGE;
6         return INT_MIN;
7     }
8     return (x < 0) ? -x : x;
9 }
```

This is transformed by the compiler into as if:

```
1 int abs(int x) _Fails(struct { int /*errno*/, int /*failure return value*/}) __attribute__((const))
2 {
3     if(x == INT_MIN)
4     {
5         // Do not set errno, return what it would be set to
6         return _Failure({ERANGE, INT_MIN});
7     }
8     return (x < 0) ? -x : x;
```

```
9 }
```

Callers of this function get transformed into a form which only sets `errno` lazily. Original code:

```
1 errno = 0;
2 int v = abs(x);
3 if(errno != 0)
4 {
5     fprintf(stderr, "abs(x) failed with code %d\n", errno);
6 }
```

Transformed code is as if:

```
1 _Either(int, {int, int}) __unique_temporary = _Catch(abs(x));
2 int v;
3 if(__unique_temporary.positive)
4 {
5     v = __unique_temporary.left;
6 }
7 else
8 {
9     // As we don't know if fprintf() reads errno, cannot set errno
10    // any later than this unfortunately.
11    errno = __unique_temporary.right.__errno;
12    fprintf(stderr, "abs(x) failed with code %d\n", errno);
13 }
```

But if `fprintf()` were also marked `_Fails(errno)` ...

```
1 extern int fprintf(FILE *, const char *, ...) _Fails(errno);
2 ...
3 _Either(int, {int, int}) __unique_temporary1 = _Catch(abs(x));
4 int v;
5 if(__unique_temporary1.positive)
6 {
7     v = __unique_temporary1.left;
8 }
9 else
10 {
11     _Either(int, {int, int}) __unique_temporary2 = fprintf(stderr, "abs(x) failed with code %d\n",
12     __unique_temporary1.right.__errno);
13     // Original code did not test for success/failure ...
14 }
15 ...
16 // Somewhere later in the function, set errno as late as
17 // is possible to not violate current errno semantics
18 errno = __unique_temporary2.right.__errno;
```

In other words, we are hoisting the setting of real `errno` out of the callee and into the caller, which in turn may get hoisted even further up the call stack, ideally if possible to the `main()` function such that real `errno` almost never gets read nor written at all⁵.

⁵And if this feature sees widespread adoption, I look forward to the day that POSIX officially marks `errno` as deprecated, as the legacy global variable `environ` now is.

A question might be why not default all functions in the C standard library (and POSIX, while we are at it) to `_Fails(errno)`?

Ideally, this is where one would want to reach if this is adopted into C22, and with consultation with the Austin Working Group. But a standard library function being considered for conversion to `_Fails(errno)` would need to be audited for any reads of `errno` before writing it anywhere in the functions called during implementation. To help this conversion, I would propose that the following be a compile error:

```
1 int func(...) _Fails(errno)
2 {
3     // "Illegal to read an errno before writing it in a _Fails(errno) function"
4     if(errno == 0) ...
5 }
```

2.1.1 What happens to `errno` when a `_Fails(T)` function calls a `_Fails(errno)` function?

It would seem reasonable that a function marked `_Fails(T)` represents the failure by `errno` of any functions it calls within its `T`, and so real `errno` ought to not be modified in a `_Fails(T)` function. However I find that dangerous: the lazy setting of real `errno` ought to be avoidable using some common construct which is usable in any code, including non-`_Fails(T)` C functions.

My proposal is that for this pattern of code, real `errno` modification is guaranteed elided by the compiler:

```
1 extern int fprintf(FILE *, const char *, ...) _Fails(errno);
2
3 int func(...)
4 {
5     int old_errno = errno; // save errno on entry
6     ...
7     errno = 0; // reset before call to errno modifying function
8     fprintf(stderr, "Hi!\n");
9     if(errno != 0)
10    {
11        ...
12    }
13    ...
14    errno = old_errno; // restore errno on exit
15    return ...;
16 }
```

3 Proposed Design: C++

This section is necessarily quite involved, as the C mechanism just proposed is calling convention compatible with how I propose C++ implements [P0709] *Zero-overhead deterministic exceptions: Throwing values*.

Given the WG14 audience, I will go into more detail than I normally would, for a purely WG21 targeted paper, regarding how C++ 20 currently works for those not familiar with modern C++. Those not interested in C++ particulars can feel free to skip this section, though I have started from first principles in order to ensure that no knowledge of the C++ standard is required to understand this section.

3.0.1 Background: C-compatible types in C++

Before the C++ 20 standard retired the trait, there was a concept that C-compatible types in C++ matched a concept called Plain Old Data (POD). C++ 11 added a trait `std::is_pod<T>`, but ‘POD types’ were well established into the C++ programmer’s vernacular long before C++ 11.

In C++ 11 parlance, POD types are those which are (i) trivial (`std::is_trivial<T>`) and (ii) have standard layout (`std::is_standard_layout<T>`).

Trivial types, in turn, are those which are (i) trivially copyable (`std::is_trivially_copyable<T>`) and (ii) have a trivial default constructor (`std::is_trivially_default_constructible<T>`).

These subcategories of C-compatible types are important to understand, because while you may not be able to construct many C++ types in C, you are able to legally pass a larger subset of C++-only types *through* C. Such C++ types may thus be transported by C code via the proposed `C_Fails(T)`.

There are two specific subcategories of C++ types which C can transport safely:

1. Trivially copyable

- Every copy constructor is trivial or deleted.
- Every move constructor is trivial or deleted.
- Every copy assignment operator is trivial or deleted.
- Every move assignment operator is trivial or deleted.
- At least one copy constructor, move constructor, copy assignment operator, or move assignment operator is non-deleted.
- Trivial non-deleted destructor.

As trivially copyable types must have a trivial destructor, C can copy such types without concern, as destruction has no side effects.

2. Move relocating (if [P1029] *SG14* `[[move_relocates]]`, or an equivalent, is accepted)

- Copies can be non-trivial.
- Moves can be non-trivial, so long as moves are equal in effect to a memory copy of bytes from source to destination, followed by memory copy of default constructed instance to source.
- Destructor can be non-trivial, but must have no side effects when executed on a default constructed instance.

As move relocating types have a destructor which is guaranteed by the programmer to have no side effects when called on a default constructed instance, and that a moved-from instance has the same representation as a default constructed instance, C can copy such types without concern, even though it will not overwrite the source with a default constructed instance. This is because C would never call C++ destructors in any case, so it does not matter if the source is simply discarded.

The caveat here is that while move relocating types may pass *through* C code, they must always return to C++ code before their life ends, otherwise it is undefined behaviour.

3.0.2 Background: How C++ currently implements exceptions

C++ exception throws take the form of a `throw expr`. This causes the stack to be unwound until a `try` wrapped statement block has a `catch(type)` matching the type of the value thrown. Types are matched using a Run Time Type Information (RTTI) search, making them non-deterministic as the RTTI in a process is unknowable in advance (and can change over time in a given running process as shared libraries are loaded and unloaded).

Back in the 1990s, C++ exceptions were implemented by setting an unwind handler on entry to each stack frame. When one threw an exception, the handler for each stack frame on the stack would be executed in order to unwind the stack.

This setting up of unwind handlers per stack frame executed added a fair bit of overhead on the CPUs of the late 1990s and early 2000s, so stack frame based handling was generally replaced with table based handling instead, which has remained until now. At compile time, the compiler generates a set of lookup tables for all C++ code compiled. When an exception is thrown, a runtime library routine is called which uses the current instruction pointer and the stack frames on the stack to look up in the EH tables which unwind handlers to invoke. This enables the successful code path to execute with effectively zero runtime overhead, but at the cost of bloating the executable binary with EH tables, and making the failure code path execution time highly non-deterministic.

However, CPU technology has marched onwards. The cost relative to CPU speed of scanning EH tables during a stack unwind has risen exponentially over what it was in the early 2000s, as such tables are inevitably in cold memory out of cache (and sometimes even paged out, on disc), and thus each table entry access costs at best hundreds of CPU cycles. One naïve benchmark puts the average *lower* bound cost at approximately 2,000 CPU cycles per stack frame unwound. Upper bounds tend into the hundreds of milliseconds on a machine with limited free RAM, due to page faulting in the EH tables.

C++ has also moved closer to the hardware within the technology stack since the early 2000s. It is no longer used as a general purpose application programming language in new projects where indeterminacy might be more acceptable; for new projects it is increasingly going into use cases where determinacy for both success and failure is paramount (reliability critical systems such as car self driving), and where the executable bloat generated by the EH tables is considered unacceptable (the embedded domain is the obvious one).

Finally, modern CPUs now tend to speculatively execute out-of-order rather than the in-order mechanism of earlier CPUs. Adding explicit checks for whether an exception has been thrown to

the successful code path is now usually free of cost on such CPUs, as stalls in other parts of the pipeline will block the CPU, thus leaving idle spare CPU execution units. As much as it might be tempting to return to the stack frame based handling approach, modern CPUs prefer potential jumps to known branches, rather than potential jumps to unknown installed unwind handlers, as the former suit their branch predictors better. Thus was born [P0709] *Zero-overhead deterministic exceptions: Throwing values*, which recognises that the hardware has changed, and the typical use cases for C++ in new code have changed.

In P0709, lightweight exception throwing functions are separated from legacy exception throwing functions via a `throws` modifier. Here are some examples:

```
1 // May throw legacy type based exceptions
2 extern int func(int);
3
4 // Never throws any exception (noexcept was added in C++ 11)
5 extern int func(int) noexcept;
6
7 // May throw new value based exceptions, proposed by P0709
8 extern int func(int) throws;
```

3.0.3 Background: `std::expected<T, E>`

[P0323] `std::expected` is in the final stages of standardisation, and it proposes a standard A-or-B sum type for C++. This vocabulary type provides:

- An observer (`.has_error()`) to check whether it contains a value (T) or an error (E).
- An observer to access the value, if present (`.value()`).
- An observer to access the error, if present (`.error()`).

[P0786] *ValuedOrError and ValueOrNone types* proposes a C++ Concept `ValueOrError` which matches all types providing at least the observer interface listed above. The expectation is that types able to consume `ValueOrError` concept matching types would provide constructors which do so. One such type, obviously enough, would be `std::expected<T, E>` itself which would construct from any `ValueOrError<T, E>` concept match.

If C adopts the generalised `_Either(A, B)` sum type proposed above, I would propose that `std::expected<A, B>` would gain a constructor overload implicitly constructing from it.

3.0.4 Background: The current proposed mechanism for deterministic C++ exceptions

It is legal in the current type throwing system to throw a non-copyable, non-movable type, and during unwind to throw lots more of them, which forces compiler implementers into some amazing gymnastics in order to meet the requirements of the C++ standard that this works correctly. Compiler implementers therefore must use extremely conservative implementations which are non-deterministic (e.g. dynamically allocating memory), in order to cope with what is a very rare, and barely used, use case.

One of the key proposals of P0709 is that the new exception throw mechanism throw *values*, not types, and that the only type throwable under the new lightweight mechanism would be a `std::error` which is defined to be no more than two CPU registers in size, and either move relocating or trivially copyable. This guarantees to the compiler that thrown values can be hoisted up the stack frame during an unwind exclusively using CPU registers, something which enormously simplifies implementation for compiler vendors, and which ought to lend itself well to aggressive optimisation.

P0709 currently considers `std::error` to be a *token* to opaque, possibly reference counted or thread locally stored, state containing the ‘real’ exception state. This would permit `std::uncaught_exceptions()` to work as at present, and for throws of type based exceptions to use the value based unwind mechanism, rather than EH tables.

This has obvious attractions for source code backwards compatibility – one is effectively throwing around `std::exception_ptr` instances which are well understood, and existing code is already using (see below for more detail).

However, in my opinion, a token-based deterministic exceptions implementation is insufficiently zero overhead or deterministic to warrant the title of the proposal paper. I therefore give additional background on SG14 low latency study group’s alternative, hard deterministic proposed mechanism.

3.0.5 Background: The alternative proposed mechanism for deterministic C++ exceptions

P0709 does suggest an alternative mechanism to throwing tokens to opaque exception state, and that is to throw the whole exception state in the value with no linked opaque data at all. Under this scheme, `std::error` is merely the *default* exception state type consisting of two CPU register sized members (the first is a pointer to a constexpr explanatory *domain* which inlines itself readily, the second is a code with meaning to that domain). If one annotates a function with merely `throws`, then that function deterministically throws `std::error` and nothing else. If one however annotates a function with `throws(T)`, then that function deterministically throws `T` instead.

This may seem like reintroduction of the hated dynamic exception specifiers, which were only just recently excised from the language [N3051]. These caused the program to terminate if the throw of an exception reached an exception specification boundary where the type of that throw was not listed. As one of the main use cases for the C++ exception mechanism is to throw more expressive locally defined types refined from public types, and for catch clauses to catch the most expressive type it knows about, this made dynamic exception specifiers useless in real world code. They also added extra runtime overhead, as they were checked at runtime. For all these reasons, they were removed from the C++ 17 standard, and rightly so.

The difference with `throws(T)` is that `T`’s compatibility is checked at *compile time*, not at run time. It would be required that if `foo() throws(X)` calls `boo() throws(Y)`, that `Y` be convertible to `X`, otherwise it will not compile. This, in combination with the value-based rather than type-based polymorphism explained below, makes these ‘static exception specifications’ a very different thing to dynamic exception specifications.

SG14’s proposed design for `std::error` in [P1028] *SG14 status_code and standard error object*

for P0709 *Zero-overhead deterministic exceptions* leverages these ‘static exception specifications’. Under this, `std::error` is merely one of a family of possible `status_code` types, all of whom can speak to one another, and which can implicitly decay themselves into a `std::error` on demand. Thus, if this true value-based direction were chosen by WG21 instead of the token-based mechanism, all types `T` in a custom `throws(T)` would be required to be convertible into the caller’s `throws(T)` (which would usually be `std::error`), and be trivially copyable or move relocating.

One might wonder what the gain would be to permit custom `throws(T)` when it is required that `T` will decay to (usually) `std::error` on demand? The main use case is for very high performance code where constructing a `std::error`, with its probable unavoidable indeterminacy⁶, is too unpredictable or expensive. A piece of code may therefore use a custom `throws(T)` perhaps with added (low latency) payload locally, and if the failure is not handled locally, then and only then is a `std::error` lazily constructed, perhaps requiring dynamic memory allocation to store the added payload. This keeps non-deterministic failure handling away from high performance code. It is analogous to how `_Fails(errno)` described earlier works by hoisting the expensive operation up to the caller or the caller’s caller.

Another key part of P1028 is *semantic comparison* where it is possible to compare any arbitrary `std::error` value to any other, and equivalence is true when they semantically rather than literally match. This allows an implementation to fail with some platform-specific error code, perhaps even with a non-public internal domain, retaining all the original information unmodified and untranslated, yet code further up the call stack can compare the `std::error` value to say `std::errc::no_such_file_or_directory`, and if the platform-specific failure is one of not finding a file or directory, the comparison will return true. Usefully, the programmer may decide to dynamically allocate memory for the payload and store the pointer to that allocation in the code field, and use a custom explanatory domain which ‘unpacks’ that payload as needed in the semantic comparison. Thanks to semantic comparison, one doesn’t need to care about implementation details – one just compares `std::error` values, and it all ‘just works’.

As one can infer, this is value-based polymorphism rather than type-based polymorphism, and P1028’s semantic comparison mechanism always executes in constant time rather than the unknowable time of a RTTI search. Instead of placing indeterminism into opaque data structures, it exposes the full implementation of deterministic exceptions to the developer, permitting them to choose where and when to introduce dynamic memory allocation or other sources of non-determinism. This is why this alternative mechanism is SG14’s preferred fixed latency choice of implementation for P0709.

3.0.6 Background: `std::error_code`

If the P1028 alternative mechanism sounds a bit familiar to some people, it is because it is a superset of the very well understood `std::error_code` which has shipped in the C++ standard library since C++ 11, and well before that in the form of `Boost.System`. P1028 `status_code` has the following improvements over `std::error_code`:

⁶In order to implement `std::uncaught_exceptions()`, one would have to increment a cold cache thread locally stored reference count. Debuggers probably would also need to be triggered, as all mainstream debuggers offer the ability to break on exception throw. Some runtimes may also wish to capture a stack backtrace.

1. Can represent warnings and informational codes, as well as failure codes.
2. Ambiguities and some legacy cruft in some parts of `std::error_code` have been eliminated e.g. `std::error_code` provides both semantic and literal comparisons via `operator==`, and it is not obvious from inspection which is in play in a given piece of code.
3. Safe to use custom domains in header only libraries which are included into shared libraries, unlike custom categories for `std::error_code` where semantic comparison randomly stops working silently.
4. Domain defined code and payload type, rather than the sole `int` which `std::error_code` provides.
5. Reduced runtime overhead through maximum use of `constexpr`, including the new `constexpr` virtual functions added to C++ 20.
6. Backwards compatible with `std::error_code`, in that any `std::error_code` can be wrapped into a `status_code` with complete preservation of the original code, and with its semantic comparison mapped into `status_code`'s semantic comparison.
7. Including its implementation does not drag in most of the standard library as a dependency, as is currently the case for `std::error_code`.
8. Prewritten domains are supplied for generic, POSIX, Win32, NT kernel, Microsoft COM and `std::error_code` codes with pregenerated semantic comparison mappings between each.

3.0.7 Background: `std::exception_ptr`

C++ 11 introduced a smart pointer `std::exception_ptr` which refers to a previously thrown C++ exception object. It can be rethrown, examined later, etc. Under deterministic exceptions, if implemented via the SG14 proposed `std::error`, `std::exception_ptr` would become a pointer to a type-based exception throw only⁷, and could be used to convert a non-deterministic type-based exception throw into a deterministic value-based exception throw via wrapping the legacy throw into a `std::exception_ptr`, and transporting that via a throw of `std::error`.

P1028, when combined with P1029 *SG14 [[move_relocates]]*, permits `std::exception_ptr` instances to be thrown directly, wrapped into a custom domain which provides the semantic mapping of whatever is stored by the exception pointer against other `std::error` instances. Thus, a legacy exception throw of a complex type could be erased into a `std::exception_ptr`, thrown upwards unwinding the stack using the value-based lightweight system, and caught in a catch handler exactly as at present. The semantic comparison feature works exactly as expected with `std::exception_ptr` instances, if they are an instance of a standard library exception type. This can provide perfect source compatibility with existing code, but eliminating EH table bloat entirely in executables for those C++ users happy to forgo binary compatibility with binaries built by earlier C++ compilers. It is expected that most compiler vendors would provide a compile time switch which lets end users

⁷There is no point erasing types compatible with SG14 proposed `std::error` into an `std::exception_ptr`, as they are not an exception type. Throwing a type convertible into SG14 proposed `std::error` is actually *control flow*, and only the throw of `std::error` itself is an exception throw. I will try to write a paper specifically on this topic soon, hopefully before the Kona meeting.

choose between EH table bloat with backwards binary compatibility, or no EH table bloat and loss of backwards binary compatibility.

3.0.8 How C++ function linkage would work using proposed `_Fails(X)`

We have established that with any mechanism, the proposed `std::error` object would be trivially copyable or move relocating, and fit into no more than two CPU registers. As explained above, this means that a `std::error` instance can safely pass through C code, and may often be entirely legal for C to construct and manipulate, depending on its domain.

I therefore propose that for all C++ functions with a `throws(X)` modifier indicating that they throw *values* of type X decayable to their caller's `throws(Y)` using the lightweight mechanism, these would map onto a function with a C representation and calling convention of `_Fails(X)`.

If the modifier is without type i.e. `throws`, that would map onto a function with a C representation of `_Fails(struct cxx_std_error)`, which would be defined as:

```
1 struct cxx_std_error
2 {
3     void *domain;
4     intptr_t code;
5 };
```

The above structure is exactly the layout and composition of SG14's proposed `std::error` (which has standard layout, and so can be cast to the above C structure legally in C++).

C++ functions marked `noexcept` would always map onto non-`_Fails` C representations.

And finally, C++ functions with legacy C++ throw potential would be down to compiler switches to choose form of calling convention, as per the discussion of potential implementation options in P0709.

As one can see, this makes the only difference between C `_Fails(X)` and C++ `throws(X)` one of *auto-propagation*. Not explicitly handling calling a C `_Fails(X)` function by using either `_Try(expr)` or `_Catch(expr)` is a compile time error. Not explicitly handling calling a C++ `throws(X)` function means that the compiler silently inserts an as-if `_Try(expr)`, thus causing unhandled failure to auto-propagate to the caller.

It is proposed that in C++, you can call a C++ `throws(X)` function with C `_Try(expr)` or C `_Catch(expr)`, if you would like.

It is proposed that you can also call a C `_Fails(X)` function from C++ *without* a C `_Try(expr)` or C `_Catch(expr)`. This would cause failure by a C function to manifest as if a deterministic C++ exception of type X had been thrown, where the type X must be convertible into the calling function's `throws(Y)` (see below for how to handle calls of `_Fails(errno)` C functions from C++).

Given this equivalence mapping of `_Fails(X)` and `throws(X)` onto one another, we can demonstrate some examples of C++ to C linkage equivalences:

C++	C
<code>extern int func(double) throws;</code>	<code>extern int _Z4funcd(double) _Fails(struct cxx_std_error);</code>
<code>extern int func(double) throws(char *);</code>	<code>extern int _Z4funcd(double) _Fails(char *);</code>
<code>extern double func(int) noexcept;</code>	<code>extern double _Z4funci(int);</code>

3.0.9 C++ calling `_Fails(errno)` C functions

P1028 provides a specific constructor of `std::error` for POSIX error codes, `posix_code(int)`. So I would propose that for C++ functions marked `throws` or `throws(T)` which call a `_Fails(errno)` function which fails, the returned code is automatically propagated as a throw of `std::error` constructed by `posix_code(errno)`.

I would love if type-based exception throwing C++ functions and `noexcept` C++ functions could do the same. However, for existing source compatibility, they cannot. I therefore propose that for those kinds of C++ function, the real `errno` setting avoidance technique described in section 2.1.1 is used.

4 Frequently Asked Questions

4.1 How interchangeable for one another are the proposed C and C++ facilities?

C `_Fails(E)` can be exchanged for C++ `throws(E)`, with the only difference being that calling a `_Fails(T)` function does not auto-propagate failure to its caller, and thus requires wrapping its invoking expression with either `_Catch(expr)` or `_Try(expr)`. `throws(errno)` probably would be a compile time error, unless we decide it aliases `throws(std::error)` given P1028's `posix_code(int)` constructor described earlier.

C `return _Failure(expr);` can be swapped for C++ `throw expr;`, but only in a `_Fails(T)` or `throws(T)` function.

C `_Either(A, B)` is not exchangeable with C++ `std::expected<A, B>`, but it is proposed that the latter can implicitly construct from the former.

C `_Catch(expr)` has no parallel in C++ in the above proposal, but if WG21 feel like it, a more C++ friendly form of `catch(expr)` could be added.

C `_Try(expr)` has no parallel in C++ in the above proposal either, as there is an implicit `_Try(expr)` added by the C++ compiler when the programmer performs no explicit handling in a `throws(T)` C++ function. One is permitted to optionally use C `_Try(expr)` when calling a C++ `throws(T)` function, it is not necessary, but some programmers may find it useful for some kinds of code to indicate exactly where control flow may change.

5 Acknowledgements

On SG14 Low Latency study group, thanks to JF Bastien, Paul Bendixen, Ben Craig, Odium Holmes, Nevin Liber, Vicente J. Botet Escriba, Arthur O'Dwyer, Herb Sutter and Tony van Eerd for feedback and suggestions.

On WG14, thanks to David Keaton, Jacob Navia, Derek M Jones, Jens Gustedt, Florian Weimer, Martin Uecker, Joseph Myers and Doug Gwyn for suggesting a complete rearchitecture of this proposal into its current form.

On std-proposals, thanks to Nicol Bolas, Thiago Macieira, Andrew Giese, Marcin Jacewski, Russell Johnston, Bengt Gustafsson and David Brown for feedback and suggestions.

I apologise if I missed your name in the above. There has been an awful lot of feedback from multiple forums. Omissions gladly rectified upon notification.

6 References

- [N3051] Doug Gregor,
Deprecating Exception Specifications
<https://wg21.link/N3051>
- [P0323] Vicente J. Botet Escribá, JF Bastien,
std::expected
<https://wg21.link/P0323>
- [P0709] Herb Sutter,
Zero-overhead deterministic exceptions: Throwing values
<https://wg21.link/P0709>
- [P0786] Vicente J. Botet Escribá,
ValuedOrError and ValueOrNone types
<https://wg21.link/P0786>
- [P0829] Ben Craig,
Freestanding proposal
<https://wg21.link/P0829>
- [P0939] B. Dawes, H. Hinnant, B. Stroustrup, D. Vandevoorde, M. Wong,
Direction for ISO C++
<http://wg21.link/P0939>
- [P1027] Niall Douglas,
SG14 Design guidelines for future low level libraries
<http://wg21.link/P1027>
- [P1028] Niall Douglas,
SG14 status_code and standard error object for P0709 Zero-overhead deterministic exceptions
<http://wg21.link/P1028>

[P1029] Niall Douglas,
SG14 *[[move_relocates]]*
<http://wg21.link/P1029>