Thou Shalt Not Specialize **std** Function Templates!

Document #: WG21 P0551R0
Date: 2017-02-01
Project: JTC1.22.32 Programming Language C++
Audience: LEWG ⇒ LWG
Reply to: Walter E. Brown <webrown.cpp@gmail.com>

Contents

Abstract

This paper proposes to modify clause [namespace.std] so as (a) to forbid users from specializing standard library function templates and (b) to allow function objects as implementations of standard library facilities specified as function templates.

To create architecture is to put in order. Put what in order? Function and objects.

— LE CORBUSIER, *né* CHARLES-ÉDOUARD JEANNERET

1 Introduction

Specializing function templates has proven problematic in practice; specializing function templates in namespace **std** has proven even more problematic. This paper (a) will cite knowledgable and well-respected colleagues in describing the core language causes of the issues involved, and (b) will then recommend wording adjustments (to subclause [namespace.std]) to address these issues in the context of the standard library.

2 Discussion

2.1 What we say today

Quoted verbatim from [N4618], here are the three paragraphs that constitute the entirety of [namespace.std]:

> 1 The behavior of a C++ program is undefined if it adds declarations or definitions to namespace **std** or to a namespace within namespace **std** unless otherwise specified. A program may add a template specialization for any standard library template to namespace **std** only if the declaration depends on a user-defined type and the specialization meets the standard library requirements for the original template and is not explicitly prohibited.

[Footnote: Any library code that instantiates other library templates must be prepared to work adequately with any user-supplied specialization that meets the minimum requirements of the Standard.]

2 The behavior of a C++ program is undefined if it declares

(2.1) — an explicit specialization of any member function of a standard library class template, or

(2.2) — an explicit specialization of any member function template of a standard library class or class template, or

(2.2) — an explicit or partial specialization of any member class template of a standard library class or class template.

A program may explicitly instantiate a template defined in the standard library only if the declaration depends on the name of a user-defined type and the instantiation meets the standard library requirements for the original template.

3 A translation unit shall not declare namespace **`std`** to be an inline namespace (7.3.1).

2.2 What's wrong with that?

According to several C++ *cognoscenti*, it is a poor C++ programming practice to specialize a function template, especially one in namespace **`std`**. Here are representative explanations and advice:

- Herb Sutter: "specializations don't participate in overloading. . . . If you want to customize a function base template and want that customization to participate in overload resolution (or, to always be used in the case of exact match), make it a plain old function, not a specialization. And, if you do provide overloads, avoid also providing specializations."[1]

- David Abrahams: "it's wrong to use function template specialization [because] it interacts in bad ways with overloads. . . . For example, if you specialize the regular **`std::swap`** for **`std::vector<mytype>&`**, your specialization won't get chosen over the standard's **`vector`**-specific **`swap`**, because specializations aren't considered during overload resolution."[2]

- Howard Hinnant: "this issue has been settled for a long time. . . . Disregard Dave's expert opinion/answer in this area at your own peril."[3]

- Eric Niebler: "[because of] the decidedly wonky way C++ resolves function calls in templates. . . , [w]e make an unqualified call to **`swap`** in order to find an overload that might be defined in . . . associated namespaces. . . , and we do **`using std::swap`** so that, on the off-chance that there is no such overload, we find the default version defined in the **`std`** namespace."[4]

- High Integrity C++ Coding Standard: "Overload resolution does not take into account explicit specializations of function templates. Only after overload resolution has chosen a function template will any explicit specializations be considered."[5]

---

[1] Herb Sutter: " Why Not Specialize Function Templates?" http://www.gotw.ca/publications/mill17.htm, 2009 (retrieved 2016–10–17). Originally published in *C/C++ Users Journal*, 19(7), July 2001.

[2] David Abrahams: Reply to "How to overload **`std::swap()`**." http://stackoverflow.com/questions/11562/how-to-overload-stdswap#comment-5729583, 2011–02–24 (retrieved 2016–10–17).

[3] Howard Hinnant: Reply to "How to overload **`std::swap()`**." http://stackoverflow.com/questions/11562/how-to-overload-stdswap#8439357, 2011–12–08 (retrieved 2016–10–17).

[4] Eric Niebler: "Customization Point Design in C++11 and Beyond." http://ericniebler.com/2014/10/21/customization-point-design-in-c11-and-beyond, 2014–10–21 (retrieved 2016–10–17).

[5] Programming Research Ltd.: "Do not explicitly specialize a function template that is overloaded with other templates," In *High Integrity C++ Coding Standard, version 4.0*, 2013–10–03 (retrieved 2016–10–20). http://www.codingstandard.com/rule/14-2-2-do-not-explicitly-specialize-a-function-template-that-is-overloaded-with-other-templates/.

While this issue has been known for over 15 years, it seems not particulary well known. Moreover, the wording in [namespace.std] is still, even today, being tweaked,[6] and there are papers (e.g., [N4381]) considering further refinements in customizing library-provided function templates.

2.3 What should we do?

We propose a few related measures to address the present uncomfortable situation regarding user customization of function templates in the standard library. Let **`F`** denote an arbitrary standard library facility that is specified as a non-member function template. (Prominent examples of such an **`F`** include **`begin`**, **`swap`**, and **`forward`**.) Then:

1. Prohibit programs from *specializing* any such **`F`**. (This is in addition to the existing prohibition against *overloading* any such **`F`** in namespace **`std`**.)
2. Despite **`F`**'s specification as a function template, grant implementations license to implement **`F`** in the form of an instantiated function object that has the specified template parameters, function parameters, and return type.

These changes will allow standard library implementors to provide *customization points*[7] that will smoothly interoperate with overloads provided by users in their own namespaces, thus avoiding surprises due to (common) misunderstandings of interactions of specialization and overloading.

3 An alternate approach

The Ranges TS [N4620] specifies customization point functionality that overlaps what we propose above, but that also appears to go into far greater detail about implementation techniques. We are uncertain that all these implementation details are truly necessary to its specification of customization points.

For example, the following extensive details (cross-references elided) are provided as a "convention" in the TS subclause "Customization Point Objects" [customization.point.object]:

> 1 A *customization point object* is a function object with a literal class type that interacts with user-defined types while enforcing semantic requirements on that interaction.
>
> 2 The type of a customization point object shall satisfy Semiregular.
>
> 3 All instances of a specific customization point object type shall be equal.
>
> 4 The type of a customization point object **`T`** shall satisfy **`Invocable<const T,Args...>()`** when the types of **`Args...`** meet the requirements specified in that customization point object's definition. Otherwise, **`T`** shall not have a function call operator that participates in overload resolution.
>
> 5 Each customization point object type constrains its return type to satisfy a particular concept.
>
> 6 The library defines several named customization point objects. In every translation unit where such a name is defined, it shall refer to the same instance of the customization point object.

---

[6] See, for example, LWG Issue 2139, "What is a *user-defined* type?". In *C++ Standard Library Active Issues List (Revision D012)*, revised 2016–12–18 at 14:12:32 UTC (retrieved 2016–12–20), http://cplusplus.github.io/LWG/lwg-active.html#2139.

[7] According to [N4381], "A customization point . . . is a function used by the Standard Library that can be overloaded on user-defined types in the user's namespace and that is found by argument-dependent lookup." Less formally, Eric Niebler defines customization points as "hooks used by generic code that end-users can specialize to customize the behavior for their types." See "Customization Point Design in C++11 and Beyond." 2014–10–21 (retrieved 2017–01–26). http://ericniebler.com/2014/10/21/customization-point-design-in-c11-and-beyond/.

> 7 [Note: Many of the customization point objects in the library evaluate function call expressions with an unqualified name which results in a call to a user-defined function found by argument dependent name lookup. To preclude such an expression resulting in a call to unconstrained functions with the same name in namespace **`std`**, customization point objects specify that lookup for these expressions is performed in a context that includes deleted overloads matching the signatures of overloads defined in namespace **`std`**. When the deleted overloads are viable, user-defined overloads must be more specialized or more constrained to be used by a customization point object. — end note]

In addition to the above "convention," each customization point specification carries significant additional verbiage regarding its implementation. For example, here is the further specification of **`swap`** from the TS's [utility.swap]:

> 1 The name **`swap`** denotes a customization point object. The effect of the expression **`ranges::swap(E1,E2)`** for some expressions **`E1`** and **`E2`** is equivalent to:
>
> (1.1) — **`(void)swap(E1,E2)`**, if that expression is valid, with overload resolution performed in a context that includes the declarations
>
> ```
> template <class T> void swap(T&, T&) = delete;
> template <class T, size_t N> void swap(T(&)[N], T(&)[N]) = delete;
> ```
>
> and does not include a declaration of **`ranges::swap`**. If the function selected by overload resolution does not exchange the values denoted by **`E1`** and **`E2`**, the program is ill-formed with no diagnostic required.
>
> (1.2) — Otherwise, **`(void)swap_ranges(E1,E2)`** if **`E1`** and **`E2`** are lvalues of array types of equal extent and **`ranges::swap(*(E1),*(E2))`** is a valid expression, except that **`noexcept(ranges::swap(E1,E2))`** is equal to **`noexcept(ranges::swap(*(E1),*(E2)))`**.
>
> (1.3) — Otherwise, if **`E1`** and **`E2`** are lvalues of the same type **`T`** which meets the syntactic requirements of **`MoveConstructible<T>()`** and **`Assignable<T&,T>()`**, exchanges the denoted values. **`ranges::swap(E1,E2)`** is a constant expression if the constructor selected by overload resolution for **`T{std::move(E1)}`** is a constexpr constructor and the expression **`E1 = std::move(E2)`** can appear in a constexpr function. **`noexcept(ranges::swap(E1,E2))`** is equal to **`is_nothrow_move_constructible<T>::value && is_nothrow_move_assignable<T>::value`**. If either **`MoveConstructible`** or **`Assignable`** is not satisfied, the program is ill-formed with no diagnostic required.
>
> (1.4) — Otherwise, **`ranges::swap(E1,E2)`** is ill-formed.
>
> 2 *Remark:* Whenever **`ranges::swap(E1,E2)`** is a valid expression, it exchanges the values denoted by **`E1`** and **`E2`** and has type **`void`**.

There is a similarly detailed amount of additional specification for each of the other customization points in the TS: **`iter_move`**, **`iter_swap`**, **`begin`**, **`end`**, **`cbegin`**, **`cend`**, **`rbegin`**, **`rend`**, **`crbegin`**, **`crend`**, **`size`**, **`empty`**, **`data`**, and **`cdata`**. Despite all this bulk, the TS does not speak to the fundamental issue we seek to address, namely to forbid user code from providing inconsistent reinterpretations of standard library features. It seems plausible that some hybrid of the two approaches may prove beneficial, but this paper proposes a minimalist approach in order to provide a contrasting viewpoint.

4 Proposed wording[8]

4.1 Adjust [namespace.std] as shown:

1 Unless otherwise specified, ~~T~~the behavior of a C++ program is undefined if it adds declarations or definitions to namespace **`std`** or to a namespace within namespace **`std`** ~~unless otherwise specified~~.

2 Unless explicitly prohibited, ~~A~~a program may add a template specialization for any standard library class template or variable template to namespace **`std`** ~~only if~~provided that (a) the added declaration depends on a user-defined type and (b) the specialization meets the standard library requirements for the original template ~~and is not explicitly prohibited~~. [Footnote: Any library code that instantiates other library templates must be prepared to work adequately with any user-supplied specialization that meets the minimum requirements of the Standard.]

~~2~~3 The behavior of a C++ program is undefined if it declares

(~~2~~3.1) — an explicit specialization of any member function of a standard library class template, or

(~~2~~3.2) — an explicit specialization of any member function template of a standard library class or class template, or

(~~2~~3.2) — an explicit or partial specialization of any member class template of a standard library class or class template.

4 A program may explicitly instantiate a class template defined in the standard library only if the declaration depends on the name of a user-defined type and the instantiation meets the standard library requirements for the original template.

5 A program may provide (in a namespace of its own) an overload for any library function template designated as a *customization point*, provided that (a) the overload's declaration depends on a user-defined type and (b) the overload meets the standard library requirements for the customization point. [Note: this permits a (qualified or unqualified) call to the customization point to invoke the most appropriate overload for the given arguments.] [Footnote: Any library customization point must be prepared to work adequately with any user-supplied overload that meets the minimum requirements of the Standard. Therefore an implementation may elect, under the as-if rule ([intro.execution]), to provide any customization point in the form of an instantiated function object ([function.objects]) even though the customization point's specification is in the form of a function template. The template parameters of each such function object and the function parameters of its **`operator()`** must match those of the corresponding customization point's specification.]

~~3~~6 A translation unit shall not declare namespace **`std`** to be an inline namespace (7.3.1).

4.2 Where and as shown, designate the following standard library functionality as customization points: (a) **`swap`**; (b) the range access algorithms **`begin`**, **`end`**, and their variants; and (c) the container access algorithms **`size`**, **`empty`**, and **`data`**.

[utility.swap] 1 *Remarks:* This function is a designated customization point ([namespace.std]) and shall not participate in overload resolution unless . . . .

[iterator.range] 1 In addition to being available via inclusion of the **`<iterator>`** header, the function templates in 24.7 are available when any of the following headers are included: **`<array>`**, **`<deque>`**, **`<forward_list>`**, . . . , and **`<vector>`**. Each of these templates is a designated customization point ([namespace.std]).

[8] All proposed additions and ~~deletions~~ are relative to the post-Issaquah Working Draft [N4618]. Editorial notes are displayed against a gray background.

[iterator.container] 1 In addition to being available via inclusion of the **`<iterator>`** header, the function templates in 24.8 are available when any of the following headers are included: **`<array>`**, **`<deque>`**, **`<forward_list>`**, . . . , and **`<vector>`**. Each of these templates is a designated customization point ([namespace.std]).

5 Acknowledgments

Many thanks, for their thoughtful comments, to Andrey Semashev and the other readers of prepublication drafts of this paper.

6 Bibliography

[N4381] Eric Niebler: "Suggested Design for Customization Points." ISO/IEC JTC1/SC22/WG21 document N4381 (pre-Lenexa mailing), 2015–03–11. http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2015/n4381.html.

[N4618] Richard Smith: "Working Draft, Standard for Programming Language C++," ISO/IEC JTC1/SC22/WG21 document N4618 (post-Issaquah mailing), 2016–11–28. http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/n4618.pdf.

[N4620] Eric Niebler and Casey Carter: "Working Draft, C++ Extensions for Ranges." ISO/IEC JTC1/SC22/WG21 document N4620 (post-Issaquah mailing), 2016–11–27. http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/n4620.pdf.

7 Document history

Version	Date	Changes
0	2017-02-01	• Published as P0551R0.