

+ 21

A (Short) Tour of C++ Modules

DANIELA ENGERT



20
21 |
October 24-29

A (SHORT) TOUR OF C++ MODULES

Modules demystified and applied

Daniela Engert - CppCon 2021

ABOUT ME

- Electrical engineer
- Build computers and create software for 40 years
- Develop hardware and software in the field of applied digital signal processing for 30 years
- Member of the C++ committee (learning novice)


- employed by  GMH Prüftechnik
GmbH · ND-Testing - Systems - Services


OVERVIEW


- Modules Foundations
 - C++20 Modules, a short recap
 - Module unit types and Module composition
 - Visibility of Identifiers vs. Reachability of Declarations
 - Relationships, linkage and linker symbols
- Modules in practice
 - Moving towards modules (by example)
 - Imports are different!
 - Is it worth it? (a case study)
 - The state of the ecosystem


C++20 MODULES


a short recap











(primary) module interface unit

global module fragment

- no declarations
- only preprocessor directives

exported
"exportedness"
applies to names

module declaration
without a name

global module
default name 'domain'

```
1 module;
2
3 #include <vector>
4
5 export module my.first_module;
6
7 export import other.module;
8 #include "mod.h"
9 constexpr int beast = 666;
10
11 export std::vector<int> frob(S s);
```

module purview

```
1 // mod.h
2 #pragma once;
3
4 struct S {
5     int value = 1;
6 }
```

The name of this module can be referred to only in

module implementation unit

```
1 module;
2
3 #include <vector>
4
5 module my.first_module;
6
7 std::vector<int> frob(S s) {
8     return {s.value + beast};
9 }
```

- module declaration
- import declaration

- not a namespace
- separate name 'domain'

MODULE TU TYPES & FEATURES

| | Defines interface | contributes to interface | implicitly imports interface | part of module purview | part of global module | exports MACROS | creates BMI | contributes to BMI | fully isolated |
|--------------------------|-------------------|--------------------------|------------------------------|------------------------|-----------------------|----------------|-------------|--------------------|----------------|
| Interface unit | ✓ | ✓ | | ✓ | ✗ | | ✓ | ✓ | ✓ |
| Implementation unit | | | ✓ | ✓ | ✗ | | | | ✓ |
| Interface partition | | ✓ | | ✓ | ✗ | | ✓ | ✓ | ✓ |
| Implementation partition | | | | ✓ | ✗ | | ✓ | ✗ | ✓ |
| Private module fragment | | | | ✓ | | | | | ✓ |
| Header unit | ✓ | ✓ | | ✓ | ✓ | ✓ | ✓ | ✓ | |

✓ unconditionally

✗ if a GMF exists in the TU / if TU's BMI is imported into the primary module interface

MODULE COMPOSITION

This zoo of module TU types allow for various module structures:

- simple module: primary module interface unit + 1 ... n module implementation units

```
1 module; // GMF
2
3 #include <vector>
4
5 export module SimpleModule;
6
7 // non-exported declarations
8 struct detail {
9     int answer = 42;
10 };
11
12 export
13 namespace SimpleModule {
14
15 void f(const detail &);
16 std::vector<detail> g();
17
18 }
```

formerly the interface header

```
1 module SimpleModule;
2
3 namespace SimpleModule {
4
5 void f(const detail & D) {
6     // do something with D
7 }
8
9 }
```

formerly implementation sources for f() and g()

```
1 module;
2
3 #include <vector>
4
5 module SimpleModule;
6
7 namespace SimpleModule {
8
9 std::vector<detail> g() {
10     return {{ 42 }, { 43 }};
11 }
12
13 }
```

MODULE COMPOSITION

- large module: primary module interface unit + 1 ... n module partitions

```
1 export module LargeModule;
2
3 export import : iface.f;
4 export import : iface.g;
```

```
1 export
2 module LargeModule : iface.f;
3 import : detail;
4
5 namespace LargeModule {
6
7 export
8 void f(const detail & D);
9
10 }
```

```
1 module LargeModule : impl.f;
2 import : iface.f;
3
4 namespace LargeModule {
5
6 void f(const detail & D) {
7 // do something with D
8 }
9
10 }
```

```
1 module LargeModule : detail;
2
3 // non-exported declarations
4 struct detail {
5     int answer = 42;
6 };
```

```
1 module;
2 #include <vector>
3
4 export
5 module LargeModule : iface.g;
6 import : detail;
7
8 namespace LargeModule {
9
10 export
11 std::vector<detail> g();
12
13 }
```

```
1 module;
2 #include <vector>
3
4 module LargeModule : impl.g;
5 import : iface.g;
6
7 namespace LargeModule {
8
9 std::vector<detail> g() {
10     return {{ 42 }, { 43 }};
11 }
12
13 }
```

MODULE COMPOSITION

- large module: hierarchy of primary module interface unit + 1 ... n related modules

```
1 export module HierModule;
2
3 export import HierModule.f;
4 export import HierModule.g;
```

```
1 export
2 module HierModule.f;
3 import HierModule.detail;
4
5 namespace HierModule {
6
7 export
8 void f(const detail & D);
9
10 }
```

```
1 module HierModule.f;
2 import HierModule.detail;
3
4 namespace HierModule {
5
6 void f(const detail & D) {
7 // do something with D
8 }
9
10 }
```

```
1 export module HierModule.detail;
2
3 struct detail {
4     int answer = 42;
5 };
```

```
1 module;
2 #include <vector>
3
4 export
5 module HierModule.g;
6 import HierModule.detail;
7
8 namespace HierModule {
9
10 export
11 std::vector<detail> g();
12
13 }
```

```
1 module;
2 #include <vector>
3
4 module HierModule.g;
5 import HierModule.detail;
6
7 namespace HierModule {
8
9 std::vector<detail> g() {
10     return {{ 42 }, { 43 }};
11 }
12
13 }
```

MODULE COMPOSITION

- small module: only primary module interface unit

```
1 module;
2
3 #include <vector>
4
5 export module SmallModule;
6
7 struct detail {
8     int answer = 42;
9 };
10
11 export
12 namespace SmallModule {
13
14 void f(const detail & D) {
15     // do something with D;
16 }
17
18 std::vector<detail> g() {
19     return {{ 42 }, { 43 }};
20 }
21
22 }
```


MODULE COMPOSITION

- single file module: only primary module interface unit with private partition

```
1 module;
2 #include <vector>
3
4 export module SingleFileModule;
5
6 struct detail { // not exported but reachable
7     int answer = 42;
8 }
9
10 export namespace SingleFileModule {
11 void f(const detail & D);
12 std::vector<detail> g();
13 }
14
15 module : private; // neither exported nor reachable!
16
17 namespace SingleFileModule {
18 void f(const detail & D) {
19     // do something with D;
20 }
21
22 std::vector<detail> g() {
23     return {{ 42 }, { 43 }};
24 }
25 }
```

VISIBILITY

hide and seek



STARTING SIMPLE

global scope

Lookup of entities
at global scope

relative to their
point of declaration

```
1 // translation unit 1
2
3 int i;      // point of declaration (POD), introduces entity 'i'
4
5 int j = i; // POD, introduces entity 'j'
6           // point of look-up (POL), names visible entity 'i'
7
8 int k = l; // POD, introduces entity 'k'
9           // POL, names invisible entity 'l'
10          // entity 'l' is not yet declared
11          // (relative invisibility)
12
13 int l;     // point of declaration (POD), introduces entity 'l'
14
15 int m = n; // POD, introduces entity 'm'
16           // POL, names invisible entity 'n'
17           // entity 'n' is declared in a different TU
18           // (total invisibility)
```

```
1 // translation unit 2
2
3 int n;      // POD, introduces entity 'n'
```

LESS OBVIOUS

class scope

Lookup of entities at
class scope
from **outside** the class

```
1 template <typename T>
2 int foo(T t) {
3     return t.value_; // POL ?, names not yet visible entity 'value'
4 } // at class scope of dependent, visible entity 't'
5
6 struct S {
7     int value_ = 1; // POD, introduces entity 'value'
8 }; // at class scope of 'S'
9
10 int x = S{}.value_; // POL, names visible entity 'value' at
11 // class scope of visible entity 'S'
12
13 x = foo(S{}); // POL !, names visible entity 'value' at
14 // class scope of visible entity 'S'
```

Lookup of entities at
class scope
from **within** the class

```
1 struct S {
2     int foo() {
3         return value; // POL ?, names not yet declared entity 'value'
4     } // at class scope of visible entity 'S'
5
6     int value = 1; // POD, introduces entity 'value'
7     // at class scope of 'S'
8
9 }; // POL !, names visible entity 'value'
10 // at class scope of visible entity 'S'
```

LOOKUP ALGORITHMS

visibility, it depends on *how you look*

Entities may become
hidden (i.e. invisible to
lookup) by

- names introduced in
scopes nearer to the
point of lookup
- **hidden friends**

but become **visible** again
by selecting the
appropriate **lookup**
algorithm

```
1 namespace N {                                // POD, introduces entity 'N::n'
2 int n = 1;                                     // POD, introduces entity 'N::n'
3
4 class S {                                     // FAIL: entity 'S' hides entity 'n' from UL
5     int j_ = 0;                                 // FAIL: entity 'S' is invisible to UL
6     friend int f(S s) { // POD, introduces entity 'N::f', declared as friend from within
7         return s.j_;                            // class scope 'S', → so-called "hidden friend"
8     }
9 };
10
11 namespace M {                                // POD, introduces entity 'M::n'
12 void n();                                     // POD, introduces entity 'M::n'
13
14 int x = n;                                    // FAIL: entity 'n' hides entity 'N::n' from UL
15                                         // FAIL: entity 'n' is invisible to UL
16 auto y = &f;                                  // FAIL: entity 'f' is not member of class S
17                                         // FAIL: entity 'f' was not declared in scope 'N' but scope 'S'
18 auto z = &S::f;                                // FAIL: entity 'f' is not member of class S
19                                         // FAIL: entity 'f' was not declared in scope 'N' but scope 'S'
20 auto z = &N::f;                                // FAIL: entity 'f' is not member of class S
21                                         // FAIL: entity 'f' was not declared in scope 'N' but scope 'S'
22
23 int a = N::n;                                // FAIL: entity 'n' is not member of class S
24                                         // FAIL: entity 'n' was not declared in scope 'N' but scope 'S'
25 } // namespace M
26 } // namespace N
```

EVEN WEIRDER

near total invisibility

```
1 auto foo(int x) {
2     struct S {           // POD, introduces entity 'S' in function block scope
3         int s_;
4     };
5     return S{x};        // POL, names entity 'S', part of the function's signature
6 }
7
8 auto what = foo(1);
9 assert(what.s_ == 1);
10
11 static_assert(std::is_same_v<decltype(what), ? ::S>); // POL, names invisible entity 'S'
```

Even though name 'S' is **visible at function block scope** and it is the function's return type, it is **totally invisible from outside** the function.

This is the best you can achieve in terms of name hiding in a TU. Alas, it denies foo forward-declarability from another TU despite having external linkage.

SELECTIVE VISIBILITY

```
1 // module interface TU
2 export module M;
3
4 struct S {           // POD, introduces entity 'S', not exported
5     int s_ = 1;
6 };
7
8 export S foo(int); // POD, introduces entity 'foo' and exports name 'foo'
9                 // POL, names visible entity 'S'
```


```
1 // module implementation TU
2 module M;
3
4 S foo(int x) {      // POL, names visible entities 'S' and 'foo'
5     return S{x};
6 }
```

```
1 // client TU
2 import M;           // introduces entity 'foo' by BMI, exported from module M
3
4 auto y = foo(1);   // POL, names entity 'foo'
5                 // the name of result type of 'foo' is totally invisible!
```

- without modules, **total invisibility** of entities declared within a TU is **impossible**
- moving declarations from headers into modules makes them **totally invisible** from the outside
- exporting names from a module and importing them **controls** the extent to which names become **visible** in the translation unit importing the module's interface.

REACHABILITY

of declarations



AN EXAMPLE

```
1 export module mod; // may become "necessarily reachable" if the interface of 'mod' is imported
2 import stuff;      // not exported, implementation detail, not part of module interface
3                      // creates "interface dependency" to 'stuff' which is "necessarily reachable"
4
5 struct S {          // not exported, not meant to be used elsewhere outside
6     S(const char * msg) : sth_{ msg } {}
7     auto what() const { return sth_.message(); }
8     something sth_; // there must be 'something' exported from module 'stuff'
9 };
10
11 export             // exports name 'foo'
12 S foo(const char * msg) {
13     return { msg };
14 }
```


```
1 #include <type_traits>
2 import mod;           // creates "interface dependency" to 'mod' and 'stuff'
3                      // makes 'mod' "necessarily reachable"
4
5 int main(int, char * argv[]) {
6     const auto result = foo(argv[0]); // so far, so good
7
8     const auto huh = result.what(); // why is entity 'what' nameable? 🤔
9     using mystery = decltype(huh); // it was never exported from 'mod'
10    // -> it's a reachable semantic property of 'S'!
11
12    static_assert(std::is_class_v<mystery>); // compiles
13    static_assert(sizeof(mystery::value_type) == sizeof(char)); // compiles
14    return huh.empty(); // compiles
15 }
```

SEMANTIC PROPERTIES

```
1 // header file "some.h"
2
3 extern "C++" {                                // default
4     using func = int(int)                      // mandatory
5                         noexcept(false); // default
6     struct D                                     // mandatory
7     {
8         operator int() const { return v_; }    // optional or mandatory
9         int v_ = 1;
10    };
11
12 extern "C" {                                  // optional
13     namespace N {                            // mandatory
14         func bar;                          // mandatory
15     }
16
17 [[nodiscard]]                                // optional
18 extern                               // default
19 inline                                // optional or mandatory
20 int N::bar(int                           // mandatory
21             x = D{})                     // optional or mandatory
22 {
23     return x + D{2};                      // optional or mandatory
24 }
25 }
26 }
```

LINKAGE

about relationships between TUs



NAME ISOLATION

```
1 export module mod1;
2
3 int foo(); // module linkage
4
5 export namespace A { // external
6
7 int bar() { // external linkage
8     return foo();
9 }
10
11 } // namespace A
```

↔ no clash ↔
↔ same namespace ::A ↔

```
1 export module mod2;
2
3 int foo(); // module linkage
4
5 namespace A { // external linkage
6
7 export int baz() { // external link.
8     return foo();
9 }
10
11 } // namespace A
```

name '::foo' is attached to
module 'mod1', i.e.
'::foo@mod1',
exported name '::A::bar' is also
attached to the module

```
1 import mod1;
2 import mod2;
3
4 using namespace A;
5
6 int main(){
7     return bar() + baz();
8 }
```

name '::foo' is attached to
module 'mod2', i.e.
'::foo@mod2',
exported name '::A::baz' is also
attached to the module

namespace name '::A' is attached to the global module, as it is oblivious of module boundaries

LINKER SYMBOLS

```
1 export module mod3;
2
3 int foo();           // module linkage, attached to module 'mod3'
4
5 namespace A {        // external linkage, attached to global module
6
7 export int bar() { // external linkage, attached to module 'mod3'
8     return foo();
9 }
10
11 } // namespace A
```

msvc 16.11 name mangling:

?**foo**@mod3@_YAHXZ::<!mod3> (🤔)
?**bar**@A@_YAHXZ::<!mod3>

clang13 & gcc(modules) name mangling:

_ZW4mod3E3**foo**v
_ZN1A3**bar**Ev

The module name **will** be encoded into the linker symbol if an entity has **module linkage**, and **may** be encoded if it is **exported** and therefore has **external linkage**

OWNERSHIP

```
1 export module mod3;
2
3 namespace A {
4
5 export int bar() { // external linkage, attached to module 'mod3'
6     return foo();
7 }
8
9 } // namespace A
```

Strong ownership model, the linker symbols of exported names **contain** the **module name** they are attached to

e.g. msvc

?bar@A@@YAHXZ ::<!mod3>

Benefit:

Identical names can be exported from multiple modules and used in separate TUs without causing linker errors.

Weak ownership model, the linker symbols of exported names are **oblivious** of module attachment

e.g. clang & gcc

_ZN1A3barEv

Benefit:

Exported names can be moved freely between modules or from headers into modules.

BUT PLATFORMS...

```
1 export module awesome.v1;  
2  
3 // other stuff, not exported,  
4 // must go here because reasons  
5  
6 export namespace libawesome {  
7     // implemented elsewhere  
8     int deep_thought(...);  
9 } // namespace libawesome
```

perfectly tested™

← compatible →

```
1 export module awesome.v2;  
2  
3 // other stuff, not exported,  
4 // must go here because reasons  
5  
6 export namespace libawesome {  
7     // for compatibility  
8     int deep_thought(...);  
9     // implemented elsewhere  
10    int much_deeper_thought(...);  
11 } // namespace libawesome
```

used in implementations of OEM1 and OEM2, not exposed

```
1 // OEM 1, traditional header  
2 // implementation calls 'deep_thought'  
3  
4 namespace OEM1 {  
5     int doit();  
6 } // namespace OEM1
```

perfectly tested™,
distributed as static
libraries & header

```
1 // OEM 2, traditional header  
2 // impl. calls 'much_deeper_thought()'  
3  
4 namespace OEM2 {  
5     int makeit();  
6 } // namespace OEM2
```

Compiles on platform A 😊

Links on platform A 😊

Profit! 💰

```
1 // Poor customer's application  
2 #include "oem1/interface.h"  
3 #include "oem2/interface.h"  
4  
5 int main(){  
6     return OEM1::doit() +  
7             OEM2::makeit();  
8 }
```

Compiles on platform B 😊

Linker error on platform B 🤯

but why? 🤔(weak ownership)

DETACHING NAMES

```
1 export module mod4;
2
3 int foo();           // module linkage, attached to module 'mod4'
4 extern "C" int bar(); // external linkage, attached to global module
5
6 extern "C++" int baz() { // external linkage, attached to global module
7     return foo() + bar();
8 }
```

msvc 16.11 name mangling:

?**foo**@A@mod4@@YAHXZ ::::<!mod4> (😐)
_bar
?**baz**@@YAHXZ


gcc(modules) name mangling:

_ZW4mod4E3**foo**v
bar
_ZW4mod4E3**baz**v (wat? 🤔)

Giving explicit language **linkage** specifications **reattaches** the names to the **global module** and mangles them accordingly into linker symbols

TRANSITIONING TO MODULES

The road forward



TRANSITIONING TO MODULES

Options available and advice on how to proceed into the modules world:

- If the interface of a library is (mostly) separate from the implementation
 - consider a **named module** by turning the interface headers into a module interface unit with optionally some interface partitions (see slide 9.1)
 - consider **refactoring the interface** to make this happen
 - think about **macros in the interface** and how to **get rid of them**
- If macros are indispensable
 - consider a named module like above plus a **header** file which imports the module and **augments it with the macros**
- Otherwise, consider using the existing headers as **header units** (discouraged)
- If a library must still be usable as a non-modular, traditional library consider a **dual-mode library** which can - by default - be #included as before, or alternatively be provided through a module interface unit.

DUAL-MODE LIBRARY

CASE STUDY: THE {FMT} LIBRARY

For the most part, the code is located in 12 headers defining the API

- core.h
- format.h
- compile.h
- printf.h

...

Plus 2 source files that can be precompiled into a static or shared library

- format.cc (+ format-inl.h)
- os.cc

THE {FMT} LIBRARY

Question: how can this traditional library become a full-fledged module of the same name, i.e. become a ***dual-mode*** library?

Requirements:

- a **lot of macros** are used in the implementation to support as many platforms, compilers and language standards as possible. This must still work.
- there are even two macros as API features in the interface.
- the **unrestricted usability** as a traditional library as before must be maintained.
- all implementation details must be **hidden** when the "Named Module" option is selected in the configuration.

THE {FMT} LIBRARY

Answer: this set of requirements is unattainable!

Neither a header module nor a named module has all necessary properties:

- header modules *can't hide* the implementation details
- named modules *can't export* macros

C++20 to the rescue: screw macros and support only the modern alternatives (i.e. UDLs) provided by the latest version of {fmt} !

THE {FMT} LIBRARY

Question: which implementation strategy? (refer to slides 9.x)

There is a lot of coupling between most headers because of

- the *vast* amount of *macros used internally*
- the *liberal* use of *implementation details* from *other* headers

And this applies to the compilable sources just as well.

This is not bad per se if the library is seen as a whole and therefore it is not untypical in traditional libraries. If a clean, layered module structure is the primary goal, untangling that 'mess' becomes necessary.

THE {FMT} LIBRARY

Answer: restructuring a dual-mode library is probably not worth the effort as long as the details are clearly separatable from the API!

Strategy:

- Wrap the existing headers and source files into a *single-file module*
- Mark the *exported* API with some *opt-in* syntax
- Make everything in the *source* files *strictly invisible* and unreachable

In other words:

- Apply preprocessor gymnastics to separate the API from details
- Move the contents of the source files into the *private module fragment*

THE {FMT} MODULE INTERFACE UNIT

```
1 module;           // start of the 'global module fragment' (GMF)
2
3 // put *all* non-library headers (STL, platform headers, etc.) here
4 // to prevent further inclusion into the module purview!
5 #include <algorithm>
6 #include <sys/stat.h>
7 ...
8           // end of external code attached to the 'global module'
9 export module fmt; // start of the 'module purview'
10
11 #define macros to differentiate between interface and details
12
13 // put *all* library headers here to become
14 // * the exported interface (API)
15 // * the non-exported, but reachable implementation details
16 #include "format.h"
17 #include "chrono.h"
18 ...
19           // end of declarations affecting other TUs
20 module : private; // start of the 'private module fragment' (PMF)
21
22 // put *all* library sources here to become part of the compiled object file
23 // all required macros are available, yay!
24 #include "format.cc"
25 #include "os.cc"
```

THE {FMT} LIBRARY

This single module interface TU compiles into **two artefacts**:

- the compiled interface (a.k.a. **BMI**)
- the compiled implementation (a single **object file**)

This is basically a **unity** build of the whole library that provides the **full API**.

The object file may then optionally be wrapped into a static or shared library.

THE {FMT} LIBRARY

Benefits of a dual-mode library:

- usable as both a traditional library and a named module from identical sources
- has the same properties as a named library, i.e.
 - total isolation from other sources changing the compile environment
 - (hopefully) cleaner interface free of implementation details being visible
- enables gradual transitioning into the modules world depending on the maturity of compilers
- doesn't require changes to existing tests
- doesn't require re-architecting the inner dependencies

LESSONS LEARNED

THE POTENTIAL PITFALLS

On the journey to making `{fmt}` a full-fledged named module, I've encountered a couple of stumbling blocks that needed to be addressed.

The properties of module interfaces require special care when implementing them. Stuff that never had to be taken into consideration with headers becomes really important now!

Most of them are due to the **separation** of visibility of names when

- compiling the **interface** TU (unrestricted visibility)
- compiling TUs that **import** the module (restricted visibility)

This applies to **both named modules and header units**!

THE POTENTIAL PITFALLS

Instantiations of templates in user code perform **name lookup** of dependent entities from **outside** of the module. Non-exported names are invisible now and may cause compile failures.

```
1 export module M;
2
3 namespace detail {
4
5 template <typename T>
6 void baz(T) {}
7
8 template <typename T>
9 void bar(T t) {
10     baz(t);           // ok while compiling the module interface
11 }                     // fails to find 'baz' when 'bar' is implicitly instantiated
12
13 } // namespace detail
14
15 export template <typename T>
16 void foo(T t) {
17     detail::bar(t);    // ok, fully qualified name lookup is done at module compilation time
18 }
```

THE POTENTIAL PITFALLS

Two potential solutions:

```
1 export module M;
2
3 namespace detail {
4     template <typename T> void baz(T) {}
5
6     template <typename T>
7     void bar(T t) {
8         detail::baz(t);           // do fully-qualified name lookup if you *really* mean
9     }                           // to call 'detail::baz' only (i.e. disable ADL)
10 } // namespace detail
11 ...
```

```
1 export module M;
2
3 namespace detail {
4
5     void baz(int) {}
6
7     template <typename T>
8     void bar(T t) {
9         using detail::baz;       // "symbolic link" 'detail::baz' (looked up at module compilation time)
10        baz(t);              // into the function body (thereby available at template instantiation time)
11    }                      // if you want to make the call to 'baz' a customization point (i.e. enable ADL)
12
13 } // namespace detail
14 ...
```

THE POTENTIAL PITFALLS

beware of entities with **internal linkage** at namespace-scope when importing headers as **header-units**.

This is quite common when defining constants.

```
1 // header file 'some.h'  
2  
3 static const int the_answer = 42; // internal linkage -> not exported from header unit  
4  
5 namespace {  
6     constexpr int the_beast = 666; // internal linkage -> not exported from header unit  
7 }
```

```
→ 1 import "some.h"                                // import rather than #include!  
2  
3 int main() {  
4     return the_answer;                          // name 'the_answer' is unknown because  
5 }                                              // it wasn't eligible to be exported from 'some.h'
```

THE POTENTIAL PITFALLS

Solution:

- make them entities with external linkage
- or wrap them into other entities

```
1 // header file 'some.h'
2
3 inline const int the_answer = 42; // define variable with external linkage
4
5 enum int_constants : int {           // enum definition has external linkage
6     the_beast = 666;
7 };
8
9 struct constants {                 // struct definition has external linkage
10    static constexpr int      no_answer = 0;
11    static constexpr unsigned pi = 4;
12 };
```



```
1 import "some.h"                      // import rather than #include!
2
3 int main() {
4     return the_answer;                // name 'the_answer' is known now
5 }
```

THE POTENTIAL PITFALLS

Within the **purview** of a module, the 'inline' specifier gets its original meaning back!

Member bodies with function **definitions** at class scope are **no longer implicitly 'inline'**. You may give 'inline' hints if you mean it.

```
1 export module M;
2
3 struct S {
4     int foo() { return bar(); }          // no longer implicitly 'inline',
5                               // the function call might be
6                               // invalid in module context!
7
8     inline int bar() { return 42; }      // safe to inline
9 }
```

THE POTENTIAL PITFALLS

Beware of entities that are **local to the TU**.


Do not expose them e.g. by naming them in non-TU-local inline functions!

Learn more at [\[basic.link\]#14](#) in the standard

```
1 export module M;
2
3 static void foo();           // TU-local
4
5 static inline void bar() { foo(); } // ok, TU-local
6
7 inline void baz() { bar(); }      // error, 'baz()' has module linkage
8                                // must not "expose" TU-local 'bar()'!
```

FROM HEADER TO MODULE

A reality check




USAGE SCENARIOS

1. Use {fmt} in traditional way by #including the required {fmt} headers
2. As before, but with a modularized standard library and #include translation for all standard library headers included by {fmt}
3. Use existing {fmt} headers as header units and import them
4. Use {fmt} as named module

Let's examine them in detail.

USAGE SCENARIOS

TEST SCENARIO



```
1 // #include <...>, import <...>, import fmt; go here
2 // The fictitious code requires at least the API from fmt/format.h and fmt/xchar.h
3
4 int main {
5     /* empty main to zoom in on making the API available to TU
6     fictitious call:
7
8     fmt::print(L"The answer is {}", 42);
9     */
10 }
```

Baseline, no `#include` or `import`:

compile time 31 ms

all configurations taken on an AMD Ryzen 9 5900X, compiled with msvc 16.11.3, release mode

USAGE SCENARIOS

#INCLUDE THE HEADERS

```
→ 1 #include <fmt/format.h>
→ 2 #include <fmt/xchar.h>
→ 3
→ 4 int main {
→ 5     /* empty main to zoom in on making the API available to TU
→ 6     fictitious call:
→ 7
→ 8     fmt::print(L"The answer is {}", 42);
→ 9     */
→ 10 }
```

Two configurations

header-only: compile time 944 ms (baseline + 913 ms),
6896 non-blank {fmt} code lines, 59'430 lines after preprocessing

static library: compile time 562 ms (baseline + 531 ms),
4685 non-blank {fmt} code lines, 42'735 lines after preprocessing

USAGE SCENARIOS

MODULARIZED STANDARD LIBRARY + #INCLUDE TRANSLATION

[CPP.INCLUDE]#7

```
====> 1 #include <fmt/format.h>
====> 2 #include <fmt/xchar.h>
3
4 int main {
5     /* empty main to zoom in on making the API available to TU
6     fictitious call:
7
8     fmt::print(L"The answer is {}", 42);
9     */
10 }
```

Two configurations

header-only: compile time 511 ms (baseline + 480 ms)

static library: compile time 304 ms (baseline + 273 ms)

Total std lib BMI size 41 MB (461 MB if std lib user-compiled from original headers)

USAGE SCENARIOS

IMPORT THE HEADERS

```
====> 1 import <fmt/format.h>;
====> 2 import <fmt/xchar.h>;
3
4 int main {
5     /* empty main to zoom in on making the API available to TU
6     fictitious call:
7
8     fmt::print(L"The answer is {}", 42);
9     */
10 }
```

Two configurations

header-only: compile time 64 ms (baseline + 33 ms),
BMI size 22 MB

static library: compile time 64 ms (baseline + 33 ms),
BMI size 16 MB

USAGE SCENARIOS

MAKE THE COMPARISON FAIR!

```
→ 1 #include <fmt/args.h>           // provide the *full* API
→ 2 #include 8 more {fmt} headers here just as the named module does!
→ 3 #include <fmt/xchar.h>
→ 4
→ 5 int main {
→ 6 }
```

| | | |
|----------------------------|---------|---|
| #include (header-only): | 1599 ms | (baseline + 1568 ms), 90'431 lines preprocessed |
| #include (static library): | 1422 ms | (baseline + 1391 ms), 88'576 lines preprocessed |
| Mod. STL (header-only): | 658 ms | (baseline + 627 ms), 10'249 code lines |
| Mod. STL (static library): | 430 ms | (baseline + 399 ms), 8'038 code lines |
| import (header-only): | 160 ms | (baseline + 129 ms), BMI size 117 MB |
| import (static library): | 155 ms | (baseline + 124 ms), BMI size 91 MB |

USAGE SCENARIOS

IMPORT NAMED MODULE

```
→ 1 import fmt;
2
3 int main {
4     /* empty main to zoom in on making the API available to TU
5     fictitious call:
6
7     fmt::print(L"The answer is {}", 42);
8     */
9 }
```

Sorry, only one configuration with **everything** that {fmt} can provide!

Module interface unit: 10'672 non-blank code lines from {fmt}, 128'431 lines after preprocessing

compile time about 31 ms

There is **no measurable difference** to baseline!

USAGE SCENARIOS


THE FINAL COMPARISON RESULT

```
1 #include <fmt/*.h>                                // provide the *full* API
2 import <fmt/*.h>                                  // provide the *full* API
3 import fmt;                                         // provide the *full* API
4
5 int main {
6 }
```

| | | | |
|------------------|----------------------------|-----------------------------------|---|
| #include | #include (header-only): | 1599 ms | (baseline + 1568 ms), 90'431 lines preprocessed |
| | #include (static library): | 1422 ms | (baseline + 1391 ms), 88'576 lines preprocessed |
| | Mod. STL (header-only): | 658 ms | (baseline + 627 ms), 10'249 code lines |
| | Mod. STL (static library): | 430 ms | (baseline + 399 ms), 8'038 code lines |
| import | import (header-only): | 160 ms | (baseline + 129 ms), BMI size 117 MB |
| | import (static library): | 155 ms | (baseline + 124 ms), BMI size 91 MB |
| named module: | 31 ms | (baseline + <1 ms), BMI size 8 MB | 128'431 lines preprocessed |
| This is the way! | | | |

IMPLEMENTATION STATUS

bumpy roads ahead ...



LANGUAGE / LIBRARY FEATURES

| | gcc (trunk) | clang | msvc |
|-------------------------|------------------|--|---|
| Syntax specification | C++20 | <= 8.0: Modules TS >= 9.0: TS and C++20 | <= 19.22: Modules TS >= 19.23: C++20 |
| Named modules | ✓ | ✓ | ✓ |
| Module partitions | ✓ | ✗ | ✓ |
| Header modules | █ (undocumented) | █ (undocumented) | ✓ |
| Private mod. fragment | ✗ | ✓ | ✓ |
| Name attachment | ✓ weak model | ✓ weak model | ✓ strong model |
| #include → import | ✗ | ✗ | ✓ |
| _cpp_modules | ⚠ 201810L | ✗ | ✓ 201907L |
| Modularized std library | ✗ | ✗ | █ (experimental) |

BUILD SYSTEMS

Build systems with support for C++ modules are rare

- **build2** (by Boris Kolpackov, [build2.org](#))
 - supports clang, gcc, and msvc
- **Evoke** (by Peter Bindels, [GitHub](#))
 - clang only
- **MSBuild** (by Microsoft, since msvc16.8, [Visual Studio](#))
 - msvc only
- **make**
 - bring your own build rules, f.e. like Bloomberg's P2473
- more ?

RESOURCES

Papers

- Modules in C++, 2004, Daveed Vandevoorde
- Modules, 2012, Doug Gregor
- A Module System for C++, 2014, Gabriel Dos Reis, Mark Hall, Gor Nishanov
- C++ Modules TS, 2018, Gabriel Dos Reis
- Another take on Modules, 2018, Richard Smith
- Merging Modules, 2019, Richard Smith
- C++23 Draft

Contact

- dani@ngrt.de
- [@DanielaKEngert](https://twitter.com/@DanielaKEngert) on Twitter



Images: Bayeux Tapestry, 11th century, world heritage

QUESTIONS?



Ceterum censeo ABI esse frangendam