
Programming Languages — C++ Extensions for Networking

Langages de programmation — Extensions C++ pour mise en réseau



Reference number
ISO/IEC TS 19216:2018(E)

AVA Documents
Document Sharing Hub
© ISO/IEC 2018



COPYRIGHT PROTECTED DOCUMENT

© ISO/IEC 2018

All rights reserved. Unless otherwise specified, or required in the context of its implementation, no part of this publication may be reproduced or utilized otherwise in any form or by any means, electronic or mechanical, including photocopying, or posting on the internet or an intranet, without prior written permission. Permission can be requested from either ISO at the address below or ISO's member body in the country of the requester.

ISO copyright office
CP 401 • Ch. de Blandonnet 8
CH-1214 Vernier, Geneva
Phone: +41 22 749 01 11
Fax: +41 22 749 09 47
Email: copyright@iso.org
Website: www.iso.org

Published in Switzerland

Contents

Foreword	vi
1 Scope	1
2 Normative references	2
3 Terms and definitions	3
4 General Principles	4
4.1 Conformance	4
4.2 Acknowledgments	4
5 Namespaces and headers	5
6 Future plans (Informative)	6
7 Feature test macros (Informative)	7
8 Method of description (Informative)	8
8.1 Structure of each clause	8
8.2 Other conventions	8
9 Error reporting	9
9.1 Synchronous operations	9
9.2 Asynchronous operations	10
9.3 Error conditions	10
9.4 Suppression of signals	10
10 Library summary	11
11 Convenience header	13
11.1 Header <experimental/net> synopsis	13
12 Forward declarations	14
12.1 Header <experimental/netfwd> synopsis	14
13 Asynchronous model	16
13.1 Header <experimental/executor> synopsis	16
13.2 Requirements	19
13.3 Class template <code>async_result</code>	27
13.4 Class template <code>async_completion</code>	28
13.5 Class template <code>associated_allocator</code>	29
13.6 Function <code>get_associated_allocator</code>	30
13.7 Class <code>execution_context</code>	30
13.8 Class <code>execution_context::service</code>	32
13.9 Class template <code>is_executor</code>	33
13.10 Executor argument tag	33
13.11 <code>uses_executor</code>	34

13.12	Class template <code>associated_executor</code>	34
13.13	Function <code>get_associated_executor</code>	35
13.14	Class template <code>executor_binder</code>	36
13.15	Function <code>bind_executor</code>	39
13.16	Class template <code>executor_work_guard</code>	40
13.17	Function <code>make_work_guard</code>	41
13.18	Class <code>system_executor</code>	42
13.19	Class <code>system_context</code>	43
13.20	Class <code>bad_executor</code>	44
13.21	Class <code>executor</code>	45
13.22	Function <code>dispatch</code>	49
13.23	Function <code>post</code>	50
13.24	Function <code>defer</code>	51
13.25	Class template <code>strand</code>	52
13.26	Class template <code>use_future_t</code>	56
13.27	Partial specialization of <code>async_result</code> for <code>packaged_task</code>	59
14	Basic I/O services	61
14.1	Header <code><experimental/io_context></code> synopsis	61
14.2	Class <code>io_context</code>	61
14.3	Class <code>io_context::executor_type</code>	65
15	Timers	67
15.1	Header <code><experimental/timer></code> synopsis	67
15.2	Requirements	67
15.3	Class template <code>wait_traits</code>	68
15.4	Class template <code>basic_waitable_timer</code>	69
16	Buffers	73
16.1	Header <code><experimental/buffer></code> synopsis	73
16.2	Requirements	78
16.3	Error codes	82
16.4	Class <code>mutable_buffer</code>	82
16.5	Class <code>const_buffer</code>	83
16.6	Buffer type traits	84
16.7	Buffer sequence access	85
16.8	Function <code>buffer_size</code>	85
16.9	Function <code>buffer_copy</code>	85
16.10	Buffer arithmetic	86
16.11	Buffer creation functions	86
16.12	Class template <code>dynamic_vector_buffer</code>	88
16.13	Class template <code>dynamic_string_buffer</code>	89
16.14	Dynamic buffer creation functions	91
17	Buffer-oriented streams	92
17.1	Requirements	92
17.2	Class <code>transfer_all</code>	94
17.3	Class <code>transfer_at_least</code>	95
17.4	Class <code>transfer_exactly</code>	95
17.5	Synchronous read operations	96
17.6	Asynchronous read operations	98

17.7	Synchronous write operations	99
17.8	Asynchronous write operations	100
17.9	Synchronous delimited read operations	102
17.10	Asynchronous delimited read operations	102
18	Sockets	104
18.1	Header <experimental/socket> synopsis	104
18.2	Requirements	106
18.3	Error codes	115
18.4	Class <code>socket_base</code>	116
18.5	Socket options	118
18.6	Class template <code>basic_socket</code>	121
18.7	Class template <code>basic_datagram_socket</code>	131
18.8	Class template <code>basic_stream_socket</code>	139
18.9	Class template <code>basic_socket_acceptor</code>	145
19	Socket iostreams	157
19.1	Class template <code>basic_socket_streambuf</code>	157
19.2	Class template <code>basic_socket_iostream</code>	161
20	Socket algorithms	164
20.1	Synchronous connect operations	164
20.2	Asynchronous connect operations	165
21	Internet protocol	167
21.1	Header <experimental/internet> synopsis	167
21.2	Requirements	171
21.3	Error codes	173
21.4	Class <code>ip::address</code>	174
21.5	Class <code>ip::address_v4</code>	177
21.6	Class <code>ip::address_v6</code>	181
21.7	Class <code>ip::bad_address_cast</code>	186
21.8	Hash support	187
21.9	Class template <code>ip::basic_address_iterator</code> specializations	187
21.10	Class template <code>ip::basic_address_range</code> specializations	188
21.11	Class template <code>ip::network_v4</code>	190
21.12	Class template <code>ip::network_v6</code>	193
21.13	Class template <code>ip::basic_endpoint</code>	195
21.14	Class template <code>ip::basic_resolver_entry</code>	199
21.15	Class template <code>ip::basic_resolver_results</code>	201
21.16	Class <code>ip::resolver_base</code>	204
21.17	Class template <code>ip::basic_resolver</code>	205
21.18	Host name functions	211
21.19	Class <code>ip::tcp</code>	211
21.20	Class <code>ip::udp</code>	212
21.21	Internet socket options	214
Index		219
Index of library names		221
Index of implementation-defined behavior		227

Foreword

ISO (the International Organization for Standardization) and IEC (the International Electrotechnical Commission) form the specialized system for worldwide standardization. National bodies that are members of ISO or IEC participate in the development of International Standards through technical committees established by the respective organization to deal with particular fields of technical activity. ISO and IEC technical committees collaborate in fields of mutual interest. Other international organizations, governmental and non-governmental, in liaison with ISO and IEC, also take part in the work. In the field of information technology, ISO and IEC have established a joint technical committee, ISO/IEC JTC 1.

The procedures used to develop this document and those intended for its further maintenance are described in the ISO/IEC Directives, Part 1. In particular the different approval criteria needed for the different types of ISO documents should be noted. This document was drafted in accordance with the editorial rules of the ISO/IEC Directives, Part 2 (see www.iso.org/directives).

Attention is drawn to the possibility that some of the elements of this document may be the subject of patent rights. ISO shall not be held responsible for identifying any or all such patent rights. Details of any patent rights identified during the development of the document will be in the Introduction and/or on the ISO list of patent declarations received (see www.iso.org/patents).

Any trade name used in this document is information given for the convenience of users and does not constitute an endorsement.

For an explanation on the voluntary nature of standards, the meaning of ISO specific terms and expressions related to conformity assessment, as well as information about ISO's adherence to the World Trade Organization (WTO) principles in the Technical Barriers to Trade (TBT) see the following URL: www.iso.org/iso/foreword.html.

This document was prepared by Technical Committee ISO/IEC JTC 1, *Information technology*, Subcommittee SC 22, *Programming languages, their environments and system software interfaces*.

1 Scope

[scope]

- ¹ This document describes extensions to the C++ Standard Library. This document specifies requirements for implementations of an interface that computer programs written in the C++ programming language may use to perform operations related to networking, such as operations involving sockets, timers, buffer management, host name resolution and internet protocols. This document is applicable to information technology systems that can perform network operations, such as those with operating systems that conform to the POSIX interface. This document is applicable only to vendors who wish to provide the interface it describes.

2 Normative references [references]

¹ The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For undated references, the latest edition of the referenced document (including any amendments) applies.

- (1.1) — ISO/IEC 14882:2014, *Programming languages — C++*
- (1.2) — ISO/IEC TS 19568:2015, *C++ Extensions for Library Fundamentals*
- (1.3) — ISO/IEC 9945:2009, *Information Technology — Portable Operating System Interface (POSIX)*
- (1.4) — ISO/IEC 2382-1:1993, *Information technology — Vocabulary*

² The programming language and library described in ISO/IEC 14882 is herein called the C++ Standard. References to clauses within the C++ Standard are written as “C++ 2014, Clause 17”. The operating system interface described in ISO/IEC 9945 is herein called POSIX.

³ This document mentions commercially available operating systems for purposes of exposition. POSIX® is a registered trademark of The IEEE. Windows® is a registered trademark of Microsoft Corporation. This information is given for the convenience of users of this document and does not constitute an endorsement by ISO or IEC of these products.

⁴ Unless otherwise specified, the whole of the C++ Standard’s Library introduction (C++ 2014, Clause 17) is included into this document by reference.

3 Terms and definitions [defs]

- ¹ For the purposes of this document, the terms and definitions given in ISO/IEC 14882:2014, ISO/IEC 2382-1:1993, and the following apply.
- ² ISO and IEC maintain terminological databases for use in standardization at the following addresses:
 - (2.1) — IEC Electropedia: available at <http://www.electropedia.org/>
 - (2.2) — ISO Online browsing platform: available at <http://www.iso.org/obp>
- ³ Terms that are used only in a small portion of this document are defined where they are used and italicized where they are defined.

3.1 [defs.host.byte.order]

host byte order

the arrangement of bytes in any integer type when using a specific machine architecture

[SOURCE: ISO/IEC 9945:2009, 3.193]

3.2 [defs.net.byte.order]

network byte order

the way of representing any integer type such that, when transmitted over a network via a network endpoint, the **int** type is transmitted as an appropriate number of octets with the most significant octet first, followed by any other octets in descending order of significance

[SOURCE: ISO/IEC 9945:2009, 3.237]

3.3 [defs.sync.op]

synchronous operation

operation where control is not returned until the operation completes

3.4 [defs.async.op]

asynchronous operation

operation where control is returned immediately without waiting for the operation to complete

[*Note 1 to entry:* Multiple asynchronous operations may be executed concurrently. — *end note*]

4 General Principles [general]

4.1 Conformance [conformance]

- ¹ Conformance is specified in terms of behavior. Ideal behavior is not always implementable, so the conformance subclauses take that into account.

4.1.1 POSIX conformance [conformance.9945]

- ¹ Some behavior is specified by reference to POSIX. How such behavior is actually implemented is unspecified.
- ² [*Note:* This constitutes an “as if” rule allowing implementations to call native operating system or other APIs. — *end note*]
- ³ Implementations are encouraged to provide such behavior as it is defined by POSIX. Implementations shall document any behavior that differs from the behavior defined by POSIX. Implementations that do not support exact POSIX behavior are encouraged to provide behavior as close to POSIX behavior as is reasonable given the limitations of actual operating systems and file systems. If an implementation cannot provide any reasonable behavior, the implementation shall report an error as specified in Error Reporting (9).
- ⁴ [*Note:* This allows users to rely on an exception being thrown or an error code being set when an implementation cannot provide any reasonable behavior. — *end note*]
- ⁵ Implementations are not required to provide behavior that is not supported by a particular operating system.

4.1.2 Conditionally-supported features [conformance.conditional]

- ¹ This document defines conditionally-supported features, in the form of additional member functions on types that satisfy Protocol (18.2.6), Endpoint (18.2.4), SettableSocketOption (18.2.9), GettableSocketOption (18.2.8) or IoControlCommand (18.2.12) requirements.
- ² [*Note:* This is so that, when the additional member functions are available, C++ programs can extend the library to add support for other protocols and socket options. — *end note*]
- ³ For the purposes of this document, implementations that provide all of the additional member functions are known as extensible implementations.
- ⁴ [*Note:* Implementations are encouraged to provide the additional member functions, where possible. It is intended that POSIX and Windows implementations will provide them. — *end note*]

4.2 Acknowledgments [intro.ack]

- ¹ The design of this specification is based, in part, on the Asio library written by Christopher Kohlhoff.

5 Namespaces and headers [namespaces]

- ¹ The components described in this document are experimental and not part of the C++ standard library. All components described in this document are declared in namespace `std::experimental::net::v1` or a sub-namespace thereof unless otherwise specified. The headers described in this document shall import the contents of `std::experimental::net::v1` into `std::experimental::net` as if by:

```
namespace std {  
    namespace experimental {  
        namespace net {  
            inline namespace v1 {}  
        }  
    }  
}
```

- ² Unless otherwise specified, references to other entities described in this document are assumed to be qualified with `std::experimental::net::v1::`, references to entities described in the C++ standard are assumed to be qualified with `std::`, and references to entities described in C++ Extensions for Library Fundamentals are assumed to be qualified with `std::experimental::fundamentals_v2::`.

6 Future plans (Informative) [plans]

- ¹ This clause describes tentative plans for future versions of this document and plans for moving content into future versions of the C++ Standard.
- ² The C++ committee may release new versions of this document, containing networking library extensions we hope to add to a near-future version of the C++ Standard. Future versions will define their contents in `std::experimental::net::v2`, `std::experimental::net::v3`, etc., with the most recent implemented version inlined into `std::experimental::net`.
- ³ When an extension defined in this or a future version of this document represents enough existing practice, it will be moved into the next version of the C++ Standard by replacing the `experimental::net::vN` segment of its namespace with `net`, and by removing the `experimental/` prefix from its header's path.

7 Feature test macros (Informative)

[feature.test]

- ¹ These macros allow users to determine which version of this document is supported by the headers defined by the specification. All headers in this document shall define the `__cpp_lib_experimental_net` feature test macro in Table 1.
- ² If an implementation supplies all of the conditionally-supported features specified in 4.1.2, all headers in this document shall additionally define the `__cpp_lib_experimental_net_extensible` feature test macro.

Table 1 — Feature-test macro(s)

Macro name	Value
<code>__cpp_lib_experimental_net</code>	201707
<code>__cpp_lib_experimental_net_extensible</code>	201707

8 Method of description (Informative)

[description]

- ¹ This subclause describes the conventions used to specify this document, in addition to those conventions specified in C++ 2014, 17.5.

8.1 Structure of each clause [structure]

8.1.1 Detailed specifications [structure.specifications]

- ¹ In addition to the elements defined in C++ 2014, 17.5.1.4, descriptions of function semantics contain the following elements (as appropriate):
- (1.1) — *Completion signature*: if the function initiates an asynchronous operation, specifies the signature of a completion handler used to receive the result of the operation.

8.2 Other conventions [conventions]

8.2.1 Nested classes [nested.class]

- ¹ Several classes defined in this document are nested classes. For a specified nested class **A::B**, an implementation is permitted to define **A::B** as a synonym for a class with equivalent functionality to class **A::B**. [*Note*: When **A::B** is a synonym for another type **A** provides a nested type **B**, to emulate the injected class name. — *end note*]

9 Error reporting

[err.report]

9.1 Synchronous operations

[err.report.sync]

- ¹ Most synchronous network library functions provide two overloads, one that throws an exception to report system errors, and another that sets an **error_code** (C++ 2014, 19.5).

[*Note*: This supports two common use cases:

- (1.1) — Uses where system errors are truly exceptional and indicate a serious failure. Throwing an exception is the most appropriate response.
- (1.2) — Uses where system errors are routine and do not necessarily represent failure. Returning an error code is the most appropriate response. This allows application specific error handling, including simply ignoring the error.

— *end note*]

- ² Functions not having an argument of type **error_code&** report errors as follows, unless otherwise specified:

- (2.1) — When a call by the implementation to an operating system or other underlying API results in an error that prevents the function from meeting its specifications, the function exits via an exception of a type that would match a handler of type **system_error**.
- (2.2) — Destructors throw nothing.

- ³ Functions having an argument of type **error_code&** report errors as follows, unless otherwise specified:

- (3.1) — If a call by the implementation to an operating system or other underlying API results in an error that prevents the function from meeting its specifications, the **error_code&** argument **ec** is set as appropriate for the specific error. Otherwise, the **ec** argument is set such that **!ec** is **true**.

- ⁴ Where a function is specified as two overloads, with and without an argument of type **error_code&**:

```
R f(A1 a1, A2 a2, ..., AN aN);
R f(A1 a1, A2 a2, ..., AN aN, error_code& ec);
```

- ⁵ then, when **R** is non-void, the effects of the first overload are as if:

```
error_code ec;
R r(f(a1, a2, ..., aN, ec));
if (ec) throw system_error(ec, S);
return r;
```

- ⁶ otherwise, when **R** is void, the effects of the first overload are as if:

```
error_code ec;
f(a1, a2, ..., aN, ec);
if (ec) throw system_error(ec, S);
```

- ⁷ except that the type thrown may differ as specified above. **S** is an NTBS indicating where the exception was thrown. [*Note*: A possible value for **S** is **__func__**. — *end note*]

- ⁸ For both overloads, failure to allocate storage is reported by throwing an exception as described in the C++ standard (C++ 2014, 17.6.5.12).

- ⁹ In this document, when a type requirement is specified using two function call expressions *f*, with and without an argument **ec** of type **error_code**:

```
f(a1, a2, ..., aN)
f(a1, a2, ..., aN, ec)
```

then the effects of the first call expression of *f* shall be as described for the first overload above.

9.2 Asynchronous operations [err.report.async]

- ¹ Asynchronous network library functions in this document are identified by having the prefix **async_** and take a completion handler (13.2.7.2). These asynchronous operations report errors as follows:
 - (1.1) — If a call by the implementation to an operating system or other underlying API results in an error that prevents the asynchronous operation from meeting its specifications, the completion handler is invoked with an **error_code** value **ec** that is set as appropriate for the specific error. Otherwise, the **error_code** value **ec** is set such that **!ec** is **true**.
 - (1.2) — Asynchronous operations shall not fail with an error condition that indicates interruption of an operating system or underlying API by a signal. [*Note:* Such as POSIX error number **EINTR** — *end note*] Asynchronous operations shall not fail with any error condition associated with non-blocking operations. [*Note:* Such as POSIX error numbers **EWOULDBLOCK**, **EAGAIN**, or **EINPROGRESS**; Windows error numbers **WSAEWOULDBLOCK** or **WSAEINPROGRESS** — *end note*]
- ² In this document, when a type requirement is specified as a call to a function or member function having the prefix **async_**, then the function shall satisfy the error reporting requirements described above.

9.3 Error conditions [err.report.conditions]

- ¹ Unless otherwise specified, when the behavior of a synchronous or asynchronous operation is defined “as if” implemented by a POSIX function, the **error_code** produced by the function shall meet the following requirements:
 - (1.1) — If the failure condition is one that is listed by POSIX for that function, the **error_code** shall compare equal to the error’s corresponding **enum class errc** (C++ 2014, 19.5) or **enum class resolver_errc** (21.3) constant.
 - (1.2) — Otherwise, the **error_code** shall be set to an implementation-defined value that reflects the underlying operating system error.
- ² [*Example:* The POSIX specification for **shutdown** lists **EBADF** as one of its possible errors. If a function that is specified “as if” implemented by **shutdown** fails with **EBADF** then the following condition holds for the **error_code** value **ec**: **ec == errc::bad_file_descriptor** — *end example*]
- ³ When the description of a function contains the element Error conditions, this lists conditions where the operation may fail. The conditions are listed, together with a suitable explanation, as **enum class** constants. Unless otherwise specified, this list is a subset of the failure conditions associated with the function.

9.4 Suppression of signals [err.report.signal]

- ¹ Some POSIX functions referred to in this document may report errors by raising a **SIGPIPE** signal. Where a synchronous or asynchronous operation is specified in terms of these POSIX functions, the generation of **SIGPIPE** is suppressed and an error condition corresponding to POSIX **EPIPE** is produced instead.

10 Library summary

[summary]

Table 2 — Networking library summary

Clause	Header(s)
Convenience header (11)	<experimental/net>
Forward declarations (12)	<experimental/netfwd>
Asynchronous model (13)	<experimental/executor>
Basic I/O services (14)	<experimental/io_context>
Timers (15)	<experimental/timer>
Buffers (16)	<experimental/buffer>
Buffer-oriented streams (17)	
Sockets (18)	<experimental/socket>
Socket iostreams (19)	
Socket algorithms (20)	
Internet protocol (21)	<experimental/internet>

- ¹ Throughout this document, the names of the template parameters are used to express type requirements, as listed in Table 3.

Table 3 — Template parameters and type requirements

template parameter name	type requirements
AcceptableProtocol	acceptable protocol (18.2.7)
Allocator	C++ 2014, 17.6.3.5
AsyncReadStream	buffer-oriented asynchronous read stream (17.1.2)
AsyncWriteStream	buffer-oriented asynchronous write stream (17.1.4)
Clock	C++ 2014, 20.12.3
CompletionCondition	completion condition (17.1.5)
CompletionToken	completion token (13.2.7.2)
ConnectCondition	connect condition (18.2.13)
ConstBufferSequence	constant buffer sequence (16.2.2)
DynamicBuffer	dynamic buffer (16.2.4)
EndpointSequence	endpoint sequence (18.2.5)
ExecutionContext	execution context (13.2.3)
Executor	executor (13.2.2)
GettableSocketOption	gettable socket option (18.2.8)
InternetProtocol	Internet protocol (21.2.1)
IoControlCommand	I/O control command (18.2.12)
MutableBufferSequence	mutable buffer sequence (16.2.1)
ProtoAllocator	proto-allocator (13.2.1)
Protocol	protocol (18.2.6)
Service	service (13.2.4)
SettableSocketOption	settable socket option (18.2.9)
Signature	signature (13.2.5)

Table 3 — Template parameters and type requirements (continued)

template parameter name	type requirements
<code>SyncReadStream</code>	buffer-oriented synchronous read stream (17.1.1)
<code>SyncWriteStream</code>	buffer-oriented synchronous write stream (17.1.3)
<code>WaitTraits</code>	wait traits (15.2.1)

11 Convenience header [convenience.hdr]

11.1 Header <experimental/net> synopsis [convenience.hdr.synop]

```
#include <experimental/executor>
#include <experimental/io_context>
#include <experimental/timer>
#include <experimental/buffer>
#include <experimental/socket>
#include <experimental/internet>
```

- ¹ [*Note:* This header is provided as a convenience for programs so that they may access all networking facilities via a single, self-contained `#include`. — *end note*]

12 Forward declarations

[fwd.decl]

12.1 Header <experimental/netfwd> synopsis

[fwd.decl.synop]

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    class execution_context;
    template<class T, class Executor>
        class executor_binder;
    template<class Executor>
        class executor_work_guard;
    class system_executor;
    class executor;
    template<class Executor>
        class strand;

    class io_context;

    template<class Clock> struct wait_traits;
    template<class Clock, class WaitTraits = wait_traits<Clock>>
        class basic_waitable_timer;
    using system_timer = basic_waitable_timer<chrono::system_clock>;
    using steady_timer = basic_waitable_timer<chrono::steady_clock>;
    using high_resolution_timer = basic_waitable_timer<chrono::high_resolution_clock>;

    template<class Protocol>
        class basic_socket;
    template<class Protocol>
        class basic_datagram_socket;
    template<class Protocol>
        class basic_stream_socket;
    template<class Protocol>
        class basic_socket_acceptor;
    template<class Protocol, class Clock = chrono::steady_clock,
        class WaitTraits = wait_traits<Clock>>
        class basic_socket_streambuf;
    template<class Protocol, class Clock = chrono::steady_clock,
        class WaitTraits = wait_traits<Clock>>
        class basic_socket_iostream;

namespace ip {

    class address;
    class address_v4;
    class address_v6;
    template<class Address>
        class basic_address_iterator;
    using address_v4_iterator = basic_address_iterator<address_v4>;
    using address_v6_iterator = basic_address_iterator<address_v6>;

```

```

template<class Address>
    class basic_address_range;
using address_v4_range = basic_address_range<address_v4>;
using address_v6_range = basic_address_range<address_v6>;
class network_v4;
class network_v6;
template<class InternetProtocol>
    class basic_endpoint;
template<class InternetProtocol>
    class basic_resolver_entry;
template<class InternetProtocol>
    class basic_resolver_results;
template<class InternetProtocol>
    class basic_resolver;
class tcp;
class udp;

} // namespace ip
} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

- ¹ Default template arguments are described as appearing both in `<netfwd>` and in the synopsis of other headers but it is well-formed to include both `<netfwd>` and one or more of the other headers. [*Note: It is the implementation's responsibility to implement headers so that including `<netfwd>` and other headers does not violate the rules about multiple occurrences of default arguments. — end note*]

13 Asynchronous model

[async]

13.1 Header <experimental/executor> synopsis

[async.synop]

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class CompletionToken, class Signature>
        class async_result;

    template<class CompletionToken, class Signature>
        struct async_completion;

    template<class T, class ProtoAllocator = allocator<void>>
        struct associated_allocator;

    template<class T, class ProtoAllocator = allocator<void>>
        using associated_allocator_t = typename associated_allocator<T, ProtoAllocator>::type;

    // 13.6, get_associated_allocator:

    template<class T>
        associated_allocator_t<T> get_associated_allocator(const T& t) noexcept;
    template<class T, class ProtoAllocator>
        associated_allocator_t<T, ProtoAllocator>
        get_associated_allocator(const T& t, const ProtoAllocator& a) noexcept;

    enum class fork_event {
        prepare,
        parent,
        child
    };

    class execution_context;

    class service_already_exists;

    template<class Service> Service& use_service(execution_context& ctx);
    template<class Service, class... Args> Service&
        make_service(execution_context& ctx, Args&&... args);
    template<class Service> bool has_service(execution_context& ctx) noexcept;

    template<class T> struct is_executor;

    template<class T>
        constexpr bool is_executor_v = is_executor<T>::value;

    struct executor_arg_t { };
    constexpr executor_arg_t executor_arg = executor_arg_t();

```

```

template<class T, class Executor> struct uses_executor;

template<class T, class Executor>
    constexpr bool uses_executor_v = uses_executor<T, Executor>::value;

template<class T, class Executor = system_executor>
    struct associated_executor;

template<class T, class Executor = system_executor>
    using associated_executor_t = typename associated_executor<T, Executor>::type;

// 13.13, get_associated_executor:

template<class T>
    associated_executor_t<T> get_associated_executor(const T& t) noexcept;
template<class T, class Executor>
    associated_executor_t<T, Executor>
        get_associated_executor(const T& t, const Executor& ex) noexcept;
template<class T, class ExecutionContext>
    associated_executor_t<T, typename ExecutionContext::executor_type>
        get_associated_executor(const T& t, ExecutionContext& ctx) noexcept;

template<class T, class Executor>
    class executor_binder;

template<class T, class Executor, class Signature>
    class async_result<executor_binder<T, Executor>, Signature>;

template<class T, class Executor, class ProtoAllocator>
    struct associated_allocator<executor_binder<T, Executor>, ProtoAllocator>;

template<class T, class Executor, class Executor1>
    struct associated_executor<executor_binder<T, Executor>, Executor1>;

// 13.15, bind_executor:

template<class Executor, class T>
    executor_binder<decay_t<T>, Executor>
        bind_executor(const Executor& ex, T&& t);
template<class ExecutionContext, class T>
    executor_binder<decay_t<T>, typename ExecutionContext::executor_type>
        bind_executor(ExecutionContext& ctx, T&& t);

template<class Executor>
    class executor_work_guard;

// 13.17, make_work_guard:

template<class Executor>
    executor_work_guard<Executor>
        make_work_guard(const Executor& ex);
template<class ExecutionContext>
    executor_work_guard<typename ExecutionContext::executor_type>
        make_work_guard(ExecutionContext& ctx);
template<class T>

```

```

    executor_work_guard<associated_executor_t<T>>
        make_work_guard(const T& t);
template<class T, class U>
    auto make_work_guard(const T& t, U&& u)
        -> decltype(make_work_guard(get_associated_executor(t, forward<U>(u))));

class system_executor;
class system_context;

bool operator==(const system_executor&, const system_executor&);
bool operator!=(const system_executor&, const system_executor&);

class bad_executor;

class executor;

bool operator==(const executor& a, const executor& b) noexcept;
bool operator==(const executor& e, nullptr_t) noexcept;
bool operator==(nullptr_t, const executor& e) noexcept;
bool operator!=(const executor& a, const executor& b) noexcept;
bool operator!=(const executor& e, nullptr_t) noexcept;
bool operator!=(nullptr_t, const executor& e) noexcept;

// 13.22, dispatch:

template<class CompletionToken>
    DEDUCED dispatch(CompletionToken&& token);
template<class Executor, class CompletionToken>
    DEDUCED dispatch(const Executor& ex, CompletionToken&& token);
template<class ExecutionContext, class CompletionToken>
    DEDUCED dispatch(ExecutionContext& ctx, CompletionToken&& token);

// 13.23, post:

template<class CompletionToken>
    DEDUCED post(CompletionToken&& token);
template<class Executor, class CompletionToken>
    DEDUCED post(const Executor& ex, CompletionToken&& token);
template<class ExecutionContext, class CompletionToken>
    DEDUCED post(ExecutionContext& ctx, CompletionToken&& token);

// 13.24, defer:

template<class CompletionToken>
    DEDUCED defer(CompletionToken&& token);
template<class Executor, class CompletionToken>
    DEDUCED defer(const Executor& ex, CompletionToken&& token);
template<class ExecutionContext, class CompletionToken>
    DEDUCED defer(ExecutionContext& ctx, CompletionToken&& token);

template<class Executor>
    class strand;

template<class Executor>
    bool operator==(const strand<Executor>& a, const strand<Executor>& b);

```

```

template<class Executor>
    bool operator!=(const strand<Executor>& a, const strand<Executor>& b);

template<class ProtoAllocator = allocator<void>>
    class use_future_t;

constexpr use_future_t<> use_future = use_future_t<>();

template<class ProtoAllocator, class Result, class... Args>
    class async_result<use_future_t<ProtoAllocator>, Result(Args...)>;

template<class Result, class... Args, class Signature>
    class async_result<packaged_task<Result(Args...)>, Signature>;

} // inline namespace v1
} // namespace net
} // namespace experimental

template<class Allocator>
    struct uses_allocator<experimental::net::v1::executor, Allocator>
        : true_type {};

} // namespace std

```

13.2 Requirements

[async.reqmts]

13.2.1 Proto-allocator requirements

[async.reqmts.proto.allocator]

- ¹ A type *A* meets the proto-allocator requirements if *A* is `CopyConstructible` (C++ 2014, 17.6.3.1), `Destructible` (C++ 2014, 17.6.3.1), and `allocator_traits<A>::rebind_alloc<U>` meets the allocator requirements (C++ 2014, 17.6.3.5), where *U* is an object type. [*Note*: For example, `allocator<void>` meets the proto-allocator requirements but not the allocator requirements. — *end note*] No comparison operator, copy operation, move operation, or swap operation on these types shall exit via an exception.

13.2.2 Executor requirements

[async.reqmts.executor]

- ¹ The library describes a standard set of requirements for executors. A type meeting the `Executor` requirements embodies a set of rules for determining how submitted function objects are to be executed.
- ² A type *X* meets the `Executor` requirements if it satisfies the requirements of `CopyConstructible` (C++ 2014, 17.6.3.1) and `Destructible` (C++ 2014, 17.6.3.1), as well as the additional requirements listed below.
- ³ No comparison operator, copy operation, move operation, swap operation, or member functions `context`, `on_work_started`, and `on_work_finished` on these types shall exit via an exception.
- ⁴ The executor copy constructor, comparison operators, and other member functions defined in these requirements shall not introduce data races as a result of concurrent calls to those functions from different threads. The member function `dispatch` may be recursively reentered.
- ⁵ Let `ctx` be the execution context returned by the executor's `context()` member function. An executor becomes invalid when the first call to `ctx.shutdown()` returns. The effect of calling `on_work_started`, `on_work_finished`, `dispatch`, `post`, or `defer` on an invalid executor is undefined. [*Note*: The copy constructor, comparison operators, and `context()` member function continue to remain valid until `ctx` is destroyed. — *end note*]
- ⁶ In Table 4, `x1` and `x2` denote (possibly const) values of type *X*, `mx1` denotes an xvalue of type *X*, `f` denotes a function object of `MoveConstructible` (C++ 2014, 17.6.3.1) type `Func` such that `f()` is a valid expression, `a` denotes a (possibly const) value of type *A* where *A* is a type meeting the `ProtoAllocator` requirements

(13.2.1), and *u* denotes an identifier.

Table 4 — Executor requirements

expression	type	assertion/note pre/post-conditions
<code>X u(x1);</code>		Shall not exit via an exception. <i>post</i> : <code>u == x1</code> and <code>std::addressof(u.context()) ==</code> <code>std::addressof(x1.context())</code> .
<code>X u(mx1);</code>		Shall not exit via an exception. <i>post</i> : <code>u</code> equals the prior value of <code>mx1</code> and <code>std::addressof(u.context())</code> equals the prior value of <code>std::addressof(mx1.context())</code> .
<code>x1 == x2</code>	<code>bool</code>	Returns <code>true</code> only if <code>x1</code> and <code>x2</code> can be interchanged with identical effects in any of the expressions defined in these type requirements. [<i>Note</i> : Returning <code>false</code> does not necessarily imply that the effects are not identical. — <i>end note</i>] <code>operator==</code> shall be reflexive, symmetric, and transitive, and shall not exit via an exception.
<code>x1 != x2</code>	<code>bool</code>	Same as <code>!(x1 == x2)</code> .
<code>x1.context()</code>	<code>execution_</code> <code>context&</code> , or <code>E&</code> where <code>E</code> is a type that satisfies the <code>ExecutionContext</code> (13.2.3) requirements.	Shall not exit via an exception. The comparison operators and member functions defined in these requirements shall not alter the reference returned by this function.
<code>x1.on_work_started()</code>		Shall not exit via an exception.
<code>x1.on_work_finished()</code>		Shall not exit via an exception. <i>Precondition</i> : A preceding call <code>x2.on_work_started()</code> where <code>x1 == x2</code> .
<code>x1.dispatch(std::move(f), a)</code>		<i>Effects</i> : Creates an object <code>f1</code> initialized with <code>DECAY_COPY(std::forward<Func>(f))</code> (C++ 2014, 30.2.6) in the current thread of execution. Calls <code>f1()</code> at most once. The executor may block forward progress of the caller until <code>f1()</code> finishes execution. Executor implementations should use the supplied allocator to allocate any memory required to store the function object. Prior to invoking the function object, the executor shall deallocate any memory allocated. [<i>Note</i> : Executors defined in this document always use the supplied allocator unless otherwise specified. — <i>end note</i>] <i>Synchronization</i> : The invocation of <code>dispatch</code> synchronizes with (C++ 2014, 1.10) the invocation of <code>f1</code> .

Table 4 — Executor requirements (continued)

expression	type	assertion/note pre/post-conditions
<code>x1.post(std::move(f), a)</code>		<i>Effects:</i> Creates an object <code>f1</code> initialized with <code>DECAY_COPY(std::forward<Func>(f))</code> in the current thread of execution. Calls <code>f1()</code> at most once. The executor shall not block forward progress of the caller pending completion of <code>f1()</code>. The executor may begin <code>f1</code>'s progress before the call to <code>post</code> completes. Executor implementations should use the supplied allocator to allocate any memory required to store the function object. Prior to invoking the function object, the executor shall deallocate any memory allocated. [<i>Note:</i> Executors defined in this document always use the supplied allocator unless otherwise specified. — <i>end note</i>] <i>Synchronization:</i> The invocation of <code>post</code> synchronizes with (C++ 2014, 1.10) the invocation of <code>f1</code>.
<code>x1.defer(std::move(f), a)</code>		<i>Effects:</i> Creates an object <code>f1</code> initialized with <code>DECAY_COPY(std::forward<Func>(f))</code> in the current thread of execution. Calls <code>f1()</code> at most once. The executor shall not block forward progress of the caller pending completion of <code>f1()</code>. The executor should not begin <code>f1</code>'s progress before the call to <code>defer</code> completes. [<i>Note:</i> One use of <code>defer</code> is to convey the intention of the caller that <code>f1</code> is a continuation of the current call context. The executor can use this information to optimize or otherwise adjust the way in which <code>f1</code> is invoked. — <i>end note</i>] Executor implementations should use the supplied allocator to allocate any memory required to store the function object. Prior to invoking the function object, the executor shall deallocate any memory allocated. [<i>Note:</i> Executors defined in this document always use the supplied allocator unless otherwise specified. — <i>end note</i>] <i>Synchronization:</i> The invocation of <code>defer</code> synchronizes with (C++ 2014, 1.10) the invocation of <code>f1</code>.

13.2.3 Execution context requirements**[`async.reqmts.executioncontext`]**

- ¹ A type `X` meets the `ExecutionContext` requirements if it is publicly and unambiguously derived from `execution_context`, and satisfies the additional requirements listed below.
- ² In Table 5, `x` denotes a value of type `X`.

Table 5 — ExecutionContext requirements

expression	return type	assertion/note pre/post-condition
<code>X::executor_type</code>	type meeting <code>Executor</code> (13.2.2) requirements	
<code>x.~X()</code>		Destroys all unexecuted function objects that were submitted via an executor object that is associated with the execution context.
<code>x.get_executor()</code>	<code>X::executor_type</code>	Returns an executor object that is associated with the execution context.

13.2.4 Service requirements

[async.reqmts.service]

- ¹ A class is a service if it is publicly and unambiguously derived from `execution_context::service`, or if it is publicly and unambiguously derived from another service. For a service `S`, `S::key_type` shall be valid and denote a type (C++ 2014, 14.8.2), `is_base_of_v<typename S::key_type, S>` shall be true, and `S` shall satisfy the `Destructible` requirements (C++ 2014, 17.6.3.1).
- ² The first parameter of all service constructors shall be an lvalue reference to `execution_context`. This parameter denotes the `execution_context` object that represents a set of services, of which the service object will be a member. [*Note*: These constructors can be called by the `make_service` function. — *end note*]
- ³ A service shall provide an explicit constructor with a single parameter of lvalue reference to `execution_context`. [*Note*: This constructor can be called by the `use_service` function. — *end note*]
- ⁴ [*Example*:

```
class my_service : public execution_context::service
{
public:
    using key_type = my_service;
    explicit my_service(execution_context& ctx);
    my_service(execution_context& ctx, int some_value);
private:
    virtual void shutdown() noexcept override;
    ...
};
```

— *end example*]

- ⁵ A service's `shutdown` member function shall destroy all copies of function objects that are held by the service.

13.2.5 Signature requirements

[async.reqmts.signature]

- ¹ A type satisfies the signature requirements if it is a call signature (C++ 2014, 20.9.1).

13.2.6 Associator requirements

[async.reqmts.associator]

- ¹ An associator defines a relationship between different types and objects where, given:
 - (1.1) — a source object `s` of type `S`,
 - (1.2) — type requirements `R`, and
 - (1.3) — a candidate object `c` of type `C` meeting the type requirements `R`,

an associated type **A** meeting the type requirements **R** may be computed, and an associated object **a** of type **A** may be obtained.

- ² An associator shall be a class template that takes two template type arguments. The first template argument is the source type **S**. The second template argument is the candidate type **C**. The second template argument shall be defaulted to some default candidate type **D** that satisfies the type requirements **R**.
- ³ An associator shall additionally satisfy the requirements in Table 6. In this table, **X** is a class template that meets the associator requirements, **S** is the source type, **s** is a value of type **S** or **const S**, **C** is the candidate type, **c** is a (possibly **const**) value of type **C**, **D** is the default candidate type, and **d** is a (possibly **const**) value of type **D** that is the default candidate object.

Table 6 — Associator requirements

expression	return type	assertion/note pre/post-conditions
X<S>::type	X<S, D>::type	
X<S, C>::type		The associated type.
X<S>::get(s)	X<S>::type	Returns X<S>::get(S, d) .
X<S, C>::get(s, c)	X<S, C>::type	Returns the associated object.

- ⁴ The associator's primary template shall be defined. A program may partially specialize the associator class template for some user-defined type **S**.
- ⁵ Finally, the associator shall provide the following type alias and function template in the enclosing namespace:

```
template<class S, class C = D> using X_t = typename X<S, C>::type;

template<class S, class C = D>
typename X<S, C>::type get_X(const S& s, const C& c = d)
{
    return X<S, C>::get(s, c);
}
```

where **X** is replaced with the name of the associator class template. [*Note*: This function template is provided as a convenience, to automatically deduce the source and candidate types. — *end note*]

13.2.7 Requirements on asynchronous operations [async.reqmts.async]

- ¹ This subclause uses the names **Alloc1**, **Alloc2**, **alloc1**, **alloc2**, **Args**, **CompletionHandler**, **completion_handler**, **Executor1**, **Executor2**, **ex1**, **ex2**, **f**, **i**, **N**, **Signature**, **token**, **T_i**, **t_i**, **work1**, and **work2** as placeholders for specifying the requirements below.

13.2.7.1 General asynchronous operation concepts [async.reqmts.async.concepts]

- ¹ An initiating function is a function which may be called to start an asynchronous operation. A completion handler is a function object that will be invoked, at most once, with the result of the asynchronous operation.
- ² The life cycle of an asynchronous operation is comprised of the following events and phases:
 - (2.1) — Event 1: The asynchronous operation is started by a call to the initiating function.
 - (2.2) — Phase 1: The asynchronous operation is now outstanding.
 - (2.3) — Event 2: The externally observable side effects of the asynchronous operation, if any, are fully established. The completion handler is submitted to an executor.
 - (2.4) — Phase 2: The asynchronous operation is now completed.

(2.5) — Event 3: The completion handler is called with the result of the asynchronous operation.

³ In this document, all functions with the prefix `async_` are initiating functions.

13.2.7.2 Completion tokens and handlers [async.reqmts.async.token]

¹ Initiating functions:

(1.1) — are function templates with template parameter `CompletionToken`;

(1.2) — accept, as the final parameter, a completion token object `token` of type `CompletionToken`;

(1.3) — specify a completion signature, which is a call signature (C++ 2014, 20.9.1) `Signature` that determines the arguments to the completion handler.

² An initiating function determines the type `CompletionHandler` of its completion handler function object by performing `typename async_result<decay_t<CompletionToken>, Signature>::completion_handler_type`. The completion handler object `completion_handler` is initialized with `std::forward<CompletionToken>(token)`. [*Note*: No other requirements are placed on the type `CompletionToken`. — *end note*]

³ The type `CompletionHandler` shall satisfy the requirements of `Destructible` (C++ 2014, 17.6.3.1) and `MoveConstructible` (C++ 2014, 17.6.3.1), and be callable with the specified call signature.

⁴ In this document, all initiating functions specify a *Completion signature*: element that defines the call signature `Signature`. The *Completion signature*: elements in this document have named parameters, and the results of an asynchronous operation are specified in terms of these names.

13.2.7.3 Deduction of initiating function return type [async.reqmts.async.return.type]

¹ The return type of an initiating function is `typename async_result<decay_t<CompletionToken>, Signature>::return_type`.

² For the sake of exposition, this document sometimes annotates functions with a return type *DEDUCED*. For every function declaration that returns *DEDUCED*, the meaning is equivalent to specifying the return type as `typename async_result<decay_t<CompletionToken>, Signature>::return_type`.

13.2.7.4 Production of initiating function return value [async.reqmts.async.return.value]

¹ An initiating function produces its return type as follows:

(1.1) — constructing an object `result` of type `async_result<decay_t<CompletionToken>, Signature>`, initialized as `result(completion_handler)`; and

(1.2) — using `result.get()` as the operand of the return statement.

² [*Example*: Given an asynchronous operation with Completion signature `void(R1 r1, R2 r2)`, an initiating function meeting these requirements may be implemented as follows:

```
template<class CompletionToken>
auto async_xyz(T1 t1, T2 t2, CompletionToken&& token)
{
    typename async_result<decay_t<CompletionToken>, void(R1, R2)>::completion_handler_type
        completion_handler(forward<CompletionToken>(token));

    async_result<decay_t<CompletionToken>, void(R1, R2)> result(completion_handler);

    // initiate the operation and cause completion_handler to be invoked with
    // the result

    return result.get();
}
```

- ³ For convenience, initiating functions may be implemented using the `async_completion` template:

```
template<class CompletionToken>
auto async_xyz(T1 t1, T2 t2, CompletionToken&& token)
{
    async_completion<CompletionToken, void(R1, R2)> init(token);

    // initiate the operation and cause init.completion_handler to be invoked
    // with the result

    return init.result.get();
}
```

— end example]

13.2.7.5 Lifetime of initiating function arguments [async.reqmts.async.lifetime]

- ¹ Unless otherwise specified, the lifetime of arguments to initiating functions shall be treated as follows:
- (1.1) — If the parameter has a pointer type or has a type of lvalue reference to non-const, the implementation may assume the validity of the pointee or referent, respectively, until the completion handler is invoked. [Note: In other words, the program is responsible for guaranteeing the validity of the argument until the completion handler is invoked. — end note]
 - (1.2) — Otherwise, the implementation does not assume the validity of the argument after the initiating function completes. [Note: In other words, the program is not required to guarantee the validity of the argument after the initiating function completes. — end note] The implementation may make copies of the argument, and all copies shall be destroyed no later than immediately after invocation of the completion handler.

13.2.7.6 Non-blocking requirements on initiating functions [async.reqmts.async.non.blocking]

- ¹ An initiating function shall not block (C++ 2014, 17.3.2) the calling thread pending completion of the outstanding operation.
- ² [Note: Initiating functions can still block the calling thread for other reasons. For example, if an initiating function locks a mutex in order to synchronize access to shared data. — end note]

13.2.7.7 Associated executor [async.reqmts.async.assoc.exec]

- ¹ Certain objects that participate in asynchronous operations have an associated executor. These are obtained as specified below.

13.2.7.8 I/O executor [async.reqmts.async.io.exec]

- ¹ An asynchronous operation has an associated executor satisfying the **Executor** (13.2.2) requirements. If not otherwise specified by the asynchronous operation, this associated executor is an object of type `system_executor`.
- ² All asynchronous operations in this document have an associated executor object that is determined as follows:
 - (2.1) — If the initiating function is a member function, the associated executor is that returned by the `get_executor` member function on the same object.
 - (2.2) — If the initiating function is not a member function, the associated executor is that returned by the `get_executor` member function of the first argument to the initiating function.
- ³ Let `Executor1` be the type of the associated executor. Let `ex1` be a value of type `Executor1`, representing the associated executor object obtained as described above.

13.2.7.9 Completion handler executor**[async.reqmts.async.handler.exec]**

- ¹ A completion handler object of type `CompletionHandler` has an associated executor satisfying the Executor requirements (13.2.2). The type of this associated executor is `associated_executor_t<CompletionHandler, Executor1>`. Let `Executor2` be the type `associated_executor_t<CompletionHandler, Executor1>`. Let `ex2` be a value of type `Executor2` obtained by performing `get_associated_executor(completion_handler, ex1)`.

13.2.7.10 Outstanding work**[async.reqmts.async.work]**

- ¹ Until the asynchronous operation has completed, the asynchronous operation shall maintain:
- (1.1) — an object `work1` of type `executor_work_guard<Executor1>`, initialized as `work1(ex1)`, and where `work1.owns_work() == true`; and
 - (1.2) — an object `work2` of type `executor_work_guard<Executor2>`, initialized as `work2(ex2)`, and where `work2.owns_work() == true`.

13.2.7.11 Allocation of intermediate storage**[async.reqmts.async.alloc]**

- ¹ Asynchronous operations may allocate memory. [*Note: Such as a data structure to store copies of the `completion_handler` object and the initiating function's arguments. — end note*]
- ² Let `Alloc1` be a type, satisfying the `ProtoAllocator` (13.2.1) requirements, that represents the asynchronous operation's default allocation strategy. [*Note: Typically `allocator<void>`. — end note*] Let `alloc1` be a value of type `Alloc1`.
- ³ A completion handler object of type `CompletionHandler` has an associated allocator object `alloc2` of type `Alloc2` satisfying the `ProtoAllocator` (13.2.1) requirements. The type `Alloc2` is `associated_allocator_t<CompletionHandler, Alloc1>`. Let `alloc2` be a value of type `Alloc2` obtained by performing `get_associated_allocator(completion_handler, alloc1)`.
- ⁴ The asynchronous operations defined in this document:
- (4.1) — If required, allocate memory using only the completion handler's associated allocator.
 - (4.2) — Prior to completion handler execution, deallocate any memory allocated using the completion handler's associated allocator.
- ⁵ [*Note: The implementation can perform operating system or underlying API calls that perform memory allocations not using the associated allocator. Invocations of the allocator functions do not introduce data races (see C++ 2014, 17.6.5.9). — end note*]

13.2.7.12 Execution of completion handler on completion of asynchronous operation**[async.reqmts.async.completion]**

- ¹ Let `Args...` be the argument types of the completion signature `Signature` and let `N` be `sizeof...(Args)`. Let `i` be in the range `[0, N)`. Let `Ti` be the *i*th type in `Args...` and let `ti` be the *i*th completion handler argument associated with `Ti`.
- ² Let `f` be a function object, callable as `f()`, that invokes `completion_handler` as if by `completion_handler(forward<T00), ..., forward<TN-1N-1))`.
- ³ If an asynchronous operation completes immediately (that is, within the thread of execution calling the initiating function, and before the initiating function returns), the completion handler shall be submitted for execution as if by performing `ex2.post(std::move(f), alloc2)`. Otherwise, the completion handler shall be submitted for execution as if by performing `ex2.dispatch(std::move(f), alloc2)`.

13.2.7.13 Completion handlers and exceptions**[async.reqmts.async.exceptions]**

- ¹ Completion handlers are permitted to throw exceptions. The effect of any exception propagated from the execution of a completion handler is determined by the executor which is executing the completion handler.

13.2.7.14 Composed asynchronous operations**[async.reqmts.async.composed]**

- ¹ In this document, a *composed asynchronous operation* is an asynchronous operation that is implemented in terms of zero or more intermediate calls to other asynchronous operations. The intermediate asynchronous operations are performed sequentially. [*Note:* That is, the completion handler of an intermediate operation initiates the next operation in the sequence. — *end note*]

An intermediate operation's completion handler shall have an associated executor that is either:

- (1.1) — the type `Executor2` and object `ex2` obtained from the completion handler type `CompletionHandler` and object `completion_handler`; or
- (1.2) — an object of an unspecified type satisfying the Executor requirements (13.2.2), that delegates executor operations to the type `Executor2` and object `ex2`.

An intermediate operation's completion handler shall have an associated allocator that is either:

- (1.3) — the type `Alloc2` and object `alloc2` obtained from the completion handler type `CompletionHandler` and object `completion_handler`; or
- (1.4) — an object of an unspecified type satisfying the ProtoAllocator requirements (13.2.1), that delegates allocator operations to the type `Alloc2` and object `alloc2`.

13.3 Class template `async_result`**[async.async.result]**

- ¹ The `async_result` class template is a customization point for asynchronous operations. Template parameter `CompletionToken` specifies the model used to obtain the result of the asynchronous operation. Template parameter `Signature` is the call signature (C++ 2014, 20.9.1) for the completion handler type invoked on completion of the asynchronous operation. The `async_result` template:

- (1.1) — transforms a `CompletionToken` into a completion handler type that is based on a `Signature`; and
- (1.2) — determines the return type and return value of an asynchronous operation's initiating function.

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class CompletionToken, class Signature>
    class async_result
    {
    public:
        using completion_handler_type = CompletionToken;
        using return_type = void;

        explicit async_result(completion_handler_type&) {}
        async_result(const async_result&) = delete;
        async_result& operator=(const async_result&) = delete;

        return_type get() {}
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

- ² The template parameter `CompletionToken` shall be an object type. The template parameter `Signature` shall be a call signature (C++ 2014, 20.9.1).
- ³ Specializations of `async_result` shall satisfy the `Destructible` requirements (C++ 2014, 17.6.3.1) in addition to the requirements in Table 7. In this table, `R` is a specialization of `async_result`; `r` is a modifiable lvalue of type `R`; and `h` is a modifiable lvalue of type `R::completion_handler_type`.

Table 7 — `async_result` specialization requirements

Expression	Return type	Requirement
<code>R::completion_handler_type</code>		A type satisfying <code>MoveConstructible</code> requirements (C++ 2014, 17.6.3.1). An object of type <code>completion_handler_type</code> shall be a function object with call signature <code>Signature</code> , and <code>completion_handler_type</code> shall be constructible with an rvalue of type <code>CompletionToken</code> .
<code>R::return_type</code>		<code>void</code> ; or a type satisfying <code>MoveConstructible</code> requirements (C++ 2014, 17.6.3.1)
<code>R r(h);</code> <code>r.get()</code>	<code>R::return_type</code>	[<i>Note:</i> An asynchronous operation's initiating function uses the <code>get()</code> member function as the sole operand of a return statement. — <i>end note</i>]

13.4 Class template `async_completion`

[`async.async.completion`]

- ¹ Class template `async_completion` is provided as a convenience, to simplify the implementation of asynchronous operations that use `async_result`.

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class CompletionToken, class Signature>
    struct async_completion
    {
        using completion_handler_type = async_result<decay_t<CompletionToken>,
            Signature>::completion_handler_type;

        explicit async_completion(CompletionToken& t);
        async_completion(const async_completion&) = delete;
        async_completion& operator=(const async_completion&) = delete;

        see below completion_handler;
        async_result<decay_t<CompletionToken>, Signature> result;
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

- ² The template parameter **Signature** shall be a call signature (C++ 2014, 20.9.1).

```
explicit async_completion(CompletionToken& t);
```

- ³ *Effects:* If `CompletionToken` and `completion_handler_type` are the same type, binds `completion_handler` to `t`; otherwise, initializes `completion_handler` with `std::forward<CompletionToken>(t)`. Initializes `result` with `completion_handler`.

see below completion_handler;

- ⁴ *Type:* `completion_handler_type&` if `CompletionToken` and `completion_handler_type` are the same type; otherwise, `completion_handler_type`.

13.5 Class template `associated_allocator` [async.assoc.alloc]

- ¹ Class template `associated_allocator` is an associator (13.2.6) for the `ProtoAllocator` (13.2.1) type requirements, with default candidate type `allocator<void>` and default candidate object `allocator<void>()`.

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class ProtoAllocator = allocator<void>>
    struct associated_allocator
    {
        using type = see below;

        static type get(const T& t, const ProtoAllocator& a = ProtoAllocator()) noexcept;
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

- ² Specializations of `associated_allocator` shall satisfy the requirements in Table 8. In this table, `X` is a specialization of `associated_allocator` for the template parameters `T` and `ProtoAllocator`; `t` is a (possibly `const`) value of type `T`; and `a` is an object of type `ProtoAllocator`.

Table 8 — `associated_allocator` specialization requirements

Expression	Return type	Note
<code>typename X::type</code>	A type meeting the proto-allocator (13.2.1) requirements.	
<code>X::get(t)</code>	<code>X::type</code>	Shall not exit via an exception. Equivalent to <code>X::get(t, ProtoAllocator())</code> .
<code>X::get(t, a)</code>	<code>X::type</code>	Shall not exit via an exception.

13.5.1 `associated_allocator` members [async.assoc.alloc.members]

```
using type = see below;
```

- ¹ *Type:* `T::allocator_type` if the *qualified-id* `T::allocator_type` is valid and denotes a type (C++ 2014, 14.8.2). Otherwise `ProtoAllocator`.

```
type get(const T& t, const ProtoAllocator& a = ProtoAllocator()) noexcept;
```

- ² *Returns:* `t.get_allocator()` if the *qualified-id* `T::allocator_type` is valid and denotes a type (C++ 2014, 14.8.2). Otherwise `a`.

13.6 Function `get_associated_allocator` [async.assoc.alloc.get]

```
template<class T>
    associated_allocator_t<T> get_associated_allocator(const T& t) noexcept;
```

- ¹ *Returns:* `associated_allocator<T>::get(t)`.

```
template<class T, class ProtoAllocator>
    associated_allocator_t<T, ProtoAllocator>
    get_associated_allocator(const T& t, const ProtoAllocator& a) noexcept;
```

- ² *Returns:* `associated_allocator<T, ProtoAllocator>::get(t, a)`.

13.7 Class `execution_context` [async.exec.ctx]

- ¹ Class `execution_context` implements an extensible, type-safe, polymorphic set of services, indexed by service type.

```
namespace std {
    namespace experimental {
        namespace net {
            inline namespace v1 {

                class execution_context
                {
                public:
                    class service;

                    // 13.7.1, construct / copy / destroy:

                    execution_context();
                    execution_context(const execution_context&) = delete;
                    execution_context& operator=(const execution_context&) = delete;
                    virtual ~execution_context();

                    // 13.7.3, execution context operations:

                    void notify_fork(fork_event e);

                protected:

                    // 13.7.4, execution context protected operations:

                    void shutdown() noexcept;
                    void destroy() noexcept;
                };

                // service access:
                template<class Service> typename Service::key_type&
                    use_service(execution_context& ctx);
            }
        }
    }
}
```

```

template<class Service, class... Args> Service&
    make_service(execution_context& ctx, Args&&... args);
template<class Service> bool has_service(const execution_context& ctx) noexcept;
class service_already_exists : public logic_error { };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

- ² Access to the services of an `execution_context` is via three function templates, `use_service`, `make_service` and `has_service`.
- ³ In a call to `use_service<Service>`, the type argument chooses a service, identified by `Service::key_type`, from a set of services in an `execution_context`. If the service is not present in the `execution_context`, an object of type `Service` is created and added to the `execution_context`. A program can check if an `execution_context` implements a particular service with the function template `has_service<Service>`.
- ⁴ Service objects may be explicitly added to an `execution_context` using the function template `make_service<Service>`. If the service is already present, `make_service` exits via an exception of type `service_already_exists`.
- ⁵ Once a service reference is obtained from an `execution_context` object by calling `use_service`, that reference remains usable until a call to `destroy()`.
- ⁶ If a call to a specialization of `use_service` or `make_service` recursively calls another specialization of `use_service` or `make_service` which would choose the same service (identified by `key_type`) from the same `execution_context`, then the behavior is undefined. [*Note*: Nested calls to specializations for different service types are well-defined. — *end note*]

13.7.1 execution_context constructor

[async.exec.ctx.cons]

```
execution_context();
```

- ¹ *Effects*: Creates an object of class `execution_context` which contains no services. [*Note*: An implementation can preload services of internal service types for its own use. — *end note*]

13.7.2 execution_context destructor

[async.exec.ctx.dtor]

```
~execution_context();
```

- ¹ *Effects*: Destroys an object of class `execution_context`. Performs `shutdown()` followed by `destroy()`.

13.7.3 execution_context operations

[async.exec.ctx.ops]

```
void notify_fork(fork_event e);
```

- ¹ *Effects*: For each service object `svc` in the set:
- (1.1) — If `e == fork_event::prepare`, performs `svc->notify_fork(e)` in reverse order of addition to the set.
- (1.2) — Otherwise, performs `svc->notify_fork(e)` in order of addition to the set.

13.7.4 execution_context protected operations

[async.exec.ctx.protected]

```
void shutdown() noexcept;
```

- ¹ *Effects*: For each service object `svc` in the `execution_context` set, in reverse order of addition to the set, performs `svc->shutdown()`. For each service in the set, `svc->shutdown()` is called only once irrespective of the number of calls to `shutdown` on the `execution_context`.

```
void destroy() noexcept;
```

- 2 *Effects:* Destroys each service object in the `execution_context` set, and removes it from the set, in reverse order of addition to the set.

13.7.5 `execution_context` globals [async.exec.ctx.globals]

- 1 The functions `use_service`, `make_service`, and `has_service` do not introduce data races as a result of concurrent calls to those functions from different threads.

```
template<class Service> typename Service::key_type&
  use_service(execution_context& ctx);
```

- 2 *Effects:* If an object of type `Service::key_type` does not already exist in the `execution_context` set identified by `ctx`, creates an object of type `Service`, initialized as `Service(ctx)`, and adds it to the set.
- 3 *Returns:* A reference to the corresponding service of `ctx`.
- 4 *Remarks:* The reference returned remains valid until a call to `destroy`.

```
template<class Service, class... Args> Service&
  make_service(execution_context& ctx, Args&&... args);
```

- 5 *Requires:* A service object of type `Service::key_type` does not already exist in the `execution_context` set identified by `ctx`.
- 6 *Effects:* Creates an object of type `Service`, initialized as `Service(ctx, forward<Args>(args)...)` , and adds it to the `execution_context` set identified by `ctx`.
- 7 *Throws:* `service_already_exists` if a corresponding service object of type `Service::key_type` is already present in the set.
- 8 *Remarks:* The reference returned remains valid until a call to `destroy`.

```
template<class Service> bool has_service(const execution_context& ctx) noexcept;
```

- 9 *Returns:* `true` if an object of type `Service::key_type` is present in `ctx`, otherwise `false`.

13.8 `Class execution_context::service` [async.exec.ctx.svc]

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

class execution_context::service
{
protected:
    // construct / copy / destroy:

    explicit service(execution_context& owner);
    service(const service&) = delete;
    service& operator=(const service&) = delete;
    virtual ~service();

    // service observers:

    execution_context& context() noexcept;
```

```

private:
    // service operations:

    virtual void shutdown() noexcept = 0;
    virtual void notify_fork(fork_event e) {}

    execution_context& context_; // exposition only
};

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

```
explicit service(execution_context& owner);
```

¹ *Postconditions:* `std::addressof(context_) == std::addressof(owner)`.

```
execution_context& context() noexcept;
```

² *Returns:* `context_`.

13.9 Class template `is_executor`

[async.is.exec]

¹ The class template `is_executor` can be used to detect executor types satisfying the Executor (13.2.2) type requirements.

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T> struct is_executor;

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

² T shall be a complete type.

³ Class template `is_executor` is a `UnaryTypeTrait` (C++ 2014, 20.10.1) with a `BaseCharacteristic` of `true_type` if the type T meets the syntactic requirements for Executor (13.2.2), otherwise `false_type`.

13.10 Executor argument tag

[async.executor.arg]

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    struct executor_arg_t { };
    constexpr executor_arg_t executor_arg = executor_arg_t();

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

- ¹ The `executor_arg_t` struct is an empty structure type used as a unique type to disambiguate constructor and function overloading. Specifically, types may have constructors with `executor_arg_t` as the first argument, immediately followed by an argument of a type that satisfies the Executor requirements (13.2.2).

13.11 uses_executor [async.uses.executor]

13.11.1 uses_executor trait [async.uses.executor.trait]

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor> struct uses_executor;

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

- ¹ Remark: Detects whether `T` has a nested `executor_type` that is convertible from `Executor`. Meets the `BinaryTypeTrait` requirements (C++ 2014, 20.10.1). The implementation provides a definition that is derived from `true_type` if a type `T::executor_type` exists and `is_convertible<Executor, T::executor_type>::value != false`, otherwise it is derived from `false_type`. A program may specialize this template to derive from `true_type` for a user-defined type `T` that does not have a nested `executor_type` but nonetheless can be constructed with an executor if the first argument of a constructor has type `executor_arg_t` and the second argument has type `Executor`.

13.11.2 uses-executor construction [async.uses.executor.cons]

- ¹ Uses-executor construction with executor `Executor` refers to the construction of an object `obj` of type `T`, using constructor arguments `v1`, `v2`, ..., `vN` of types `V1`, `V2`, ..., `VN`, respectively, and an executor `ex` of type `Executor`, according to the following rules:
- (1.1) — if `uses_executor_v<T, Executor>` is `true` and `is_constructible<T, executor_arg_t, Executor, V1, V2, ..., VN>::value` is `true`, then `obj` is initialized as `obj(executor_arg, ex, v1, v2, ..., vN)`;
 - (1.2) — otherwise, `obj` is initialized as `obj(v1, v2, ..., vN)`.

13.12 Class template associated_executor [async.assoc.exec]

- ¹ Class template `associated_executor` is an associator (13.2.6) for the `Executor` (13.2.2) type requirements, with default candidate type `system_executor` and default candidate object `system_executor()`.

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor = system_executor>
    struct associated_executor
    {
        using type = see below;

        static type get(const T& t, const Executor& e = Executor()) noexcept;
    };
}
```

```

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

- ² Specializations of `associated_executor` shall satisfy the requirements in Table 9. In this table, **X** is a specialization of `associated_executor` for the template parameters **T** and **Executor**; **t** is a (possible const) value of **T**; and **e** is an object of type **Executor**.

Table 9 — `associated_executor` specialization requirements

Expression	Return type	Note
<code>typename X::type</code>	A type meeting Executor requirements (13.2.2).	
<code>X::get(t)</code>	<code>X::type</code>	Shall not exit via an exception. Equivalent to <code>X::get(t, Executor())</code> .
<code>X::get(t, e)</code>	<code>X::type</code>	Shall not exit via an exception.

13.12.1 `associated_executor` members

[`async.assoc.exec.members`]

using type = *see below*;

- ¹ *Type:* `T::executor_type` if the *qualified-id* `T::executor_type` is valid and denotes a type (C++ 2014, 14.8.2). Otherwise `Executor`.

type `get(const T& t, const Executor& e = Executor()) noexcept;`

- ² *Returns:* `t.get_executor()` if the *qualified-id* `T::executor_type` is valid and denotes a type (C++ 2014, 14.8.2). Otherwise `e`.

13.13 Function `get_associated_executor`

[`async.assoc.exec.get`]

```

template<class T>
associated_executor_t<T> get_associated_executor(const T& t) noexcept;

```

- ¹ *Returns:* `associated_executor<T>::get(t)`.

```

template<class T, class Executor>
associated_executor_t<T, Executor>
get_associated_executor(const T& t, const Executor& ex) noexcept;

```

- ² *Returns:* `associated_executor<T, Executor>::get(t, ex)`.

- ³ *Remarks:* This function shall not participate in overload resolution unless `is_executor_v<Executor>` is `true`.

```

template<class T, class ExecutionContext>
associated_executor_t<T, typename ExecutionContext::executor_type>
get_associated_executor(const T& t, ExecutionContext& ctx) noexcept;

```

- ⁴ *Returns:* `get_associated_executor(t, ctx.get_executor())`.

- ⁵ *Remarks:* This function shall not participate in overload resolution unless `is_convertible<ExecutionContext&, execution_context&>::value` is `true`.

13.14 Class template executor_binder**[async.exec.binder]**

- ¹ The class template `executor_binder` binds executors to objects. A specialization `executor_binder<T, Executor>` binds an executor of type `Executor` satisfying the `Executor` requirements (13.2.2) to an object or function of type `T`.

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor>
    class executor_binder
    {
    public:
        // types:

        using target_type = T;
        using executor_type = Executor;

        // 13.14.1, construct / copy / destroy:

        executor_binder(T t, const Executor& ex);
        executor_binder(const executor_binder& other) = default;
        executor_binder(executor_binder&& other) = default;
        template<class U, class OtherExecutor>
            executor_binder(const executor_binder<U, OtherExecutor>& other);
        template<class U, class OtherExecutor>
            executor_binder(executor_binder<U, OtherExecutor>&& other);
        template<class U, class OtherExecutor>
            executor_binder(executor_arg_t, const Executor& ex,
                const executor_binder<U, OtherExecutor>& other);
        template<class U, class OtherExecutor>
            executor_binder(executor_arg_t, const Executor& ex,
                executor_binder<U, OtherExecutor>&& other);

        ~executor_binder();

        // 13.14.2, executor binder access:

        T& get() noexcept;
        const T& get() const noexcept;
        executor_type get_executor() const noexcept;

        // 13.14.3, executor binder invocation:

        template<class... Args>
            result_of_t<T&(Args&&...)> operator()(Args&&... args);
        template<class... Args>
            result_of_t<const T&(Args&&...)> operator()(Args&&... args) const;

    private:
        Executor ex_; // exposition only
        T target_; // exposition only
    };

```

```

template<class T, class Executor, class Signature>
    class async_result<executor_binder<T, Executor>, Signature>;

template<class T, class Executor, class ProtoAllocator>
    struct associated_allocator<executor_binder<T, Executor>, ProtoAllocator>;

template<class T, class Executor, class Executor1>
    struct associated_executor<executor_binder<T, Executor>, Executor1>;

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

13.14.1 executor_binder constructors

[async.exec.binder.cons]

```
executor_binder(T t, const Executor& ex);
```

- 1 *Effects:* Initializes `ex_` with `ex`. Initializes `target_` by performing uses-executor construction, using the constructor argument `std::move(t)` and the executor `ex_`.

```

template<class U, class OtherExecutor>
    executor_binder(const executor_binder<U, OtherExecutor>& other);

```

- 2 *Requires:* If `U` is not convertible to `T`, or if `OtherExecutor` is not convertible to `Executor`, the program is ill-formed.
- 3 *Effects:* Initializes `ex_` with `other.get_executor()`. Initializes `target_` by performing uses-executor construction, using the constructor argument `other.get()` and the executor `ex_`.

```

template<class U, class OtherExecutor>
    executor_binder(executor_binder<U, OtherExecutor>&& other);

```

- 4 *Requires:* If `U` is not convertible to `T`, or if `OtherExecutor` is not convertible to `Executor`, the program is ill-formed.
- 5 *Effects:* Initializes `ex_` with `other.get_executor()`. Initializes `target_` by performing uses-executor construction, using the constructor argument `std::move(other.get())` and the executor `ex_`.

```

template<class U, class OtherExecutor>
    executor_binder(executor_arg_t, const Executor& ex,
        const executor_binder<U, OtherExecutor>& other);

```

- 6 *Requires:* If `U` is not convertible to `T` the program is ill-formed.
- 7 *Effects:* Initializes `ex_` with `ex`. Initializes `target_` by performing uses-executor construction, using the constructor argument `other.get()` and the executor `ex_`.

```

template<class U, class OtherExecutor>
    executor_binder(executor_arg_t, const Executor& ex,
        executor_binder<U, OtherExecutor>&& other);

```

- 8 *Requires:* `U` is `T` or convertible to `T`.
- 9 *Effects:* Initializes `ex_` with `ex`. Initializes `target_` by performing uses-executor construction, using the constructor argument `std::move(other.get())` and the executor `ex_`.

13.14.2 executor_binder access

[async.exec.binder.access]

```
T& get() noexcept;
const T& get() const noexcept;
```

¹ *Returns:* target_.

```
executor_type get_executor() const noexcept;
```

² *Returns:* executor_.

13.14.3 executor_binder invocation

[async.exec.binder.invocation]

```
template<class... Args>
    result_of_t<T&(Args&&...)> operator()(Args&&... args);
template<class... Args>
    result_of_t<const T&(Args&&...)> operator()(Args&&... args) const;
```

¹ *Returns:* INVOKE(get(), forward<Args>(args)...) (C++ 2014, 20.9.2).

13.14.4 Class template partial specialization async_result

[async.exec.binder.async.result]

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor, class Signature>
    class async_result<executor_binder<T, Executor>, Signature>
    {
    public:
        using completion_handler_type = executor_binder<
            typename async_result<T, Signature>::completion_handler_type,
            Executor>;
        using return_type = typename async_result<T, Signature>::return_type;

        explicit async_result(completion_handler_type& h);
        async_result(const async_result&) = delete;
        async_result& operator=(const async_result&) = delete;

        return_type get();

    private:
        async_result<T, Signature> target_; // exposition only
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

```
explicit async_result(completion_handler_type& h);
```

¹ *Effects:* Initializes target_ as target_(h.get()).

```
return_type get();
```

² *Returns:* target_.get().

13.14.5 Class template partial specialization `associated_allocator` [`async.exec.binder.assoc.alloc`]

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor, class ProtoAllocator>
    struct associated_allocator<executor_binder<T, Executor>, ProtoAllocator>
    {
        using type = associated_allocator_t<T, ProtoAllocator>;

        static type get(const executor_binder<T, Executor>& b,
                        const ProtoAllocator& a = ProtoAllocator()) noexcept;
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

```
static type get(const executor_binder<T, Executor>& b,
               const ProtoAllocator& a = ProtoAllocator()) noexcept;
```

¹ *Returns:* `associated_allocator<T, ProtoAllocator>::get(b.get(), a)`.

13.14.6 Class template partial specialization `associated_executor` [`async.exec.binder.assoc.exec`]

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class T, class Executor, class Executor1>
    struct associated_executor<executor_binder<T, Executor>, Executor1>
    {
        using type = Executor;

        static type get(const executor_binder<T, Executor>& b,
                        const Executor1& e = Executor1()) noexcept;
    };

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std
```

```
static type get(const executor_binder<T, Executor>& b,
               const Executor1& e = Executor1()) noexcept;
```

¹ *Returns:* `b.get_executor()`.

13.15 Function `bind_executor`

[`async.bind.executor`]

```
template<class Executor, class T>
```

```

executor_binder<decay_t<T>, Executor>
bind_executor(const Executor& ex, T&& t);

```

1 *Returns:* executor_binder<decay_t<T>, Executor>(forward<T>(t), ex).

2 *Remarks:* This function shall not participate in overload resolution unless is_executor_v<Executor> is true.

```

template<class ExecutionContext, class CompletionToken>
executor_binder<decay_t<T>, typename ExecutionContext::executor_type>
bind_executor(ExecutionContext& ctx, T&& t);

```

3 *Returns:* bind_executor(ctx.get_executor(), forward<T>(t)).

4 *Remarks:* This function shall not participate in overload resolution unless is_convertible<ExecutionContext&, execution_context_t&>::value is true.

13.16 Class template executor_work_guard

[async.exec.work.guard]

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    template<class Executor>
    class executor_work_guard
    {
    public:
        // types:

        using executor_type = Executor;

        // construct / copy / destroy:

        explicit executor_work_guard(const executor_type& ex) noexcept;
        executor_work_guard(const executor_work_guard& other) noexcept;
        executor_work_guard(executor_work_guard&& other) noexcept;

        executor_work_guard& operator=(const executor_work_guard&) = delete;

        ~executor_work_guard();

        // executor work guard observers:

        executor_type get_executor() const noexcept;
        bool owns_work() const noexcept;

        // executor work guard modifiers:

        void reset() noexcept;

    private:
        Executor ex_; // exposition only
        bool owns_; // exposition only
    };

} // inline namespace v1
} // namespace net

```

```

} // namespace experimental
} // namespace std

```

13.16.1 executor_work_guard members [async.exec.work.guard.members]

```
explicit executor_work_guard(const executor_type& ex) noexcept;
```

1 *Effects:* Initializes `ex_` with `ex`, and then performs `ex_.on_work_started()`.

2 *Postconditions:* `ex_ == ex` and `owns_ == true`.

```
executor_work_guard(const executor_work_guard& other) noexcept;
```

3 *Effects:* Initializes `ex_` with `other.ex_`. If `other.owns_ == true`, performs `ex_.on_work_started()`.

4 *Postconditions:* `ex_ == other.ex_` and `owns_ == other.owns_`.

```
executor_work_guard(executor_work_guard&& other) noexcept;
```

5 *Effects:* Initializes `ex_` with `std::move(other.ex_)` and initializes `owns_` with `other.owns_`, and sets `other.owns_` to false.

```
~executor_work_guard();
```

6 *Effects:* If `owns_` is true, performs `ex_.on_work_finished()`.

```
executor_type get_executor() const noexcept;
```

7 *Returns:* `ex_`.

```
bool owns_work() const noexcept;
```

8 *Returns:* `owns_`.

```
void reset() noexcept;
```

9 *Effects:* If `owns_` is true, performs `ex_.on_work_finished()`.

10 *Postconditions:* `owns_ == false`.

13.17 Function make_work_guard [async.make.work.guard]

```
template<class Executor>
executor_work_guard<Executor>
make_work_guard(const Executor& ex);
```

1 *Returns:* `executor_work_guard<Executor>(ex)`.

2 *Remarks:* This function shall not participate in overload resolution unless `is_executor_v<Executor>` is true.

```
template<class ExecutionContext>
executor_work_guard<typename ExecutionContext::executor_type>
make_work_guard(ExecutionContext& ctx);
```

3 *Returns:* `make_work_guard(ctx.get_executor())`.

4 *Remarks:* This function shall not participate in overload resolution unless `is_convertible<ExecutionContext&, execution_context_t::value>` is true.

```
template<class T>
executor_work_guard<associated_executor_t<T>>
make_work_guard(const T& t);
```

- 5 *Returns:* `make_work_guard(get_associated_executor(t))`.
- 6 *Remarks:* This function shall not participate in overload resolution unless `is_executor_v<T>` is false and `is_convertible<T&, execution_context&>::value` is false.

```
template<class T, class U>
    auto make_work_guard(const T& t, U&& u)
        -> decltype(make_work_guard(get_associated_executor(t, forward<U>(u))));
```

7 *Returns:* `make_work_guard(get_associated_executor(t, forward<U>(u)))`.

13.18 Class `system_executor` [async.system.exec]

- 1 Class `system_executor` represents a set of rules where function objects are permitted to execute on any thread.

```
namespace std {
    namespace experimental {
        namespace net {
            inline namespace v1 {

                class system_executor
                {
                public:
                    // constructors:

                    system_executor() {}

                    // 13.18.1, executor operations:

                    system_context& context() const noexcept;

                    void on_work_started() const noexcept {}
                    void on_work_finished() const noexcept {}

                    template<class Func, class ProtoAllocator>
                        void dispatch(Func&& f, const ProtoAllocator& a) const;
                    template<class Func, class ProtoAllocator>
                        void post(Func&& f, const ProtoAllocator& a) const;
                    template<class Func, class ProtoAllocator>
                        void defer(Func&& f, const ProtoAllocator& a) const;
                };

                bool operator==(const system_executor&, const system_executor&) noexcept;
                bool operator!=(const system_executor&, const system_executor&) noexcept;

            } // inline namespace v1
        } // namespace net
    } // namespace experimental
} // namespace std
```

- 2 Class `system_executor` satisfies the `Destructible` (C++ 2014, 17.6.3.1), `DefaultConstructible` (C++ 2014, 17.6.3.1), and `Executor` (13.2.2) type requirements.
- 3 To satisfy the `Executor` requirements for the `post` and `defer` member functions, the system executor may create `thread` objects to run the submitted function objects. These `thread` objects are collectively referred to as system threads.

13.18.1 system_executor operations**[async.system.exec.ops]**

```
system_context& context() const noexcept;
```

- 1 *Returns:* A reference to an object with static storage duration. All calls to this function return references to the same object.

```
template<class Func, class ProtoAllocator>
void dispatch(Func&& f, const ProtoAllocator& a) const;
```

- 2 *Effects:* Equivalent to `DECAY_COPY(forward<Func>(f))()` (C++ 2014, 30.2.6).

```
template<class Func, class ProtoAllocator>
void post(Func&& f, const ProtoAllocator& a) const;
template<class Func, class ProtoAllocator>
void defer(Func&& f, const ProtoAllocator& a) const;
```

- 3 *Effects:* If `context().stopped()` is false, creates an object `f1` initialized with `DECAY_COPY(forward<Func>(f))`, and calls `f1()` as if in a thread of execution represented by a `thread` object. Any exception propagated from the execution of `DECAY_COPY(forward<Func>(f))()` results in a call to `std::terminate`.

13.18.2 system_executor comparisons**[async.system.exec.comparisons]**

```
bool operator==(const system_executor&, const system_executor&) noexcept;
```

- 1 *Returns:* true.

```
bool operator!=(const system_executor&, const system_executor&) noexcept;
```

- 2 *Returns:* false.

13.19 Class system_context**[async.system.context]**

- 1 Class `system_context` implements the execution context associated with `system_executor` objects.

```
namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

class system_context : public execution_context
{
public:
    // types:

    using executor_type = system_executor;

    // construct / copy / destroy:

    system_context() = delete;
    system_context(const system_context&) = delete;
    system_context& operator=(const system_context&) = delete;
    ~system_context();

    // system_context operations:

    executor_type get_executor() noexcept;
```

```

    void stop();
    bool stopped() const noexcept;
    void join();
};

```

```

} // inline namespace v1
} // namespace net
} // namespace experimental
} // namespace std

```

² The class `system_context` satisfies the `ExecutionContext` (13.2.3) type requirements.

³ The `system_context` member functions `get_executor`, `stop`, and `stopped`, and the `system_executor` copy constructors, member functions and comparison operators, do not introduce data races as a result of concurrent calls to those functions from different threads of execution.

```
~system_context();
```

⁴ *Effects:* Performs `stop()` followed by `join()`.

```
executor_type get_executor() noexcept;
```

⁵ *Returns:* `system_executor()`.

```
void stop();
```

⁶ *Effects:* Signals all system threads to exit as soon as possible. If a system thread is currently executing a function object, the thread will exit only after completion of that function object. Returns without waiting for the system threads to complete.

⁷ *Postconditions:* `stopped() == true`.

```
bool stopped() const noexcept;
```

⁸ *Returns:* `true` if the `system_context` has been stopped by a prior call to `stop`.

```
void join();
```

⁹ *Effects:* Blocks the calling thread (C++ 2014, 17.3.2) until all system threads have completed.

¹⁰ *Synchronization:* The completion of each system thread synchronizes with (C++ 2014, 1.10) the corresponding successful `join()` return.

13.20 Class `bad_executor`

[`async.bad.exec`]

¹ An exception of type `bad_executor` is thrown by `executor` member functions `dispatch`, `post`, and `defer` when the executor object has no target.

```

namespace std {
namespace experimental {
namespace net {
inline namespace v1 {

    class bad_executor : public exception
    {
    public:
        // constructor:
        bad_executor() noexcept;
    };
}
}
}
}

```