



# **C++23**

## **Modules, std::stacktrace, std::expected**

Andi Skrgat

# Content of this deck

---

1. Modules
2. std::stacktrace
3. std::expected
4. Practical

# 01

# Modules

---



# Intro to modules

- A language feature to share declarations and definitions across translation unit
- Since C++20 but good support is fairly recent (CMake removed experimental support this year)
- Completely orthogonal to namespaces



memgraph [SSH: lima-default]

instance\_state.cppm

```
1 // Copyright 2025 Memgraph Ltd.      Andi Skrgat, 2 months ago * feat: Use num committed txns for failover and u...
2 // by the Apache License, Version 2.0, included in the file
3 // licenses/APL.txt.
4
5 module;
6
7 #include "utils/uuid.hpp"
8
9 #include <optional>
10
11 export module memgraph.coordination.instance_state;
12
13 #ifdef MG_ENTERPRISE
14
15 You, 21 hours ago | 2 authors (Andi Skrgat and one other)
16 export namespace memgraph::coordination {
17
18     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
19     struct InstanceStateV1 {
20         bool is_replica;           // MAIN or REPLICA
21         std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
22         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
23     };
24
25     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
26     struct InstanceState {
27         bool is_replica;           // MAIN or REPLICA
28         std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
29         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
30         std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
31         std::optional<std::map<std::string, std::map<std::string, int64_t>>>
32             | | replicas_num_txns; // if main, return num of committed txns for each instance
33
34         // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
35         InstanceStateV1 Downgrade() const {
36             return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
37         }
38     };
39 }
40
41 } // namespace memgraph::coordination
42
43 #endif
```

← → Q memgraph [SSH: lima-default]

instance\_state.cppm

```
1 // Copyright 2025 Memgraph Ltd.      Andi Skrgat, 2 months ago * feat: Use num committed txns for failover and u...
2 // by the Apache License, Version 2.0, included in the file
3 // licenses/APL.txt.
4
5 module;
6
7 #include "utils/uuid.hpp"
8
9 #include <optional>
10
11 export module memgraph.coordination.instance_state;
12
13 #ifdef MG_ENTERPRISE
14
15 You, 21 hours ago | 2 authors (Andi Skrgat and one other)
16 export namespace memgraph::coordination {
17
18     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
19     struct InstanceStateV1 {
20         bool is_replica;           // MAIN or REPLICA
21         std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
22         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
23     };
24
25     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
26     struct InstanceState {
27         bool is_replica;           // MAIN or REPLICA
28         std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
29         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
30         std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
31         std::optional<std::map<std::string, std::map<std::string, int64_t>>>
32             | | replicas_num_txns; // if main, return num of committed txns for each instance
33
34         // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
35         InstanceStateV1 Downgrade() const {
36             return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
37         }
38     };
39 }
40
41 } // namespace memgraph::coordination
42
43 #endif
```





← → Q memgraph [SSH: lima-default]

File instance\_state.cppm X

src > coordination > instance\_state.cppm > ...

```
1 // Copyright 2025 Memgraph Ltd.      Andi Skrgat, 2 months ago * feat: Use num committed txns for failover and u...
9 // by the Apache License, Version 2.0, included in the file
10 // licenses/APL.txt.
11
12 module;
13
14 #include "utils/uuid.hpp"
15
16 #include <optional>
17
18 export module memgraph.coordination.instance_state;
19
20 #ifdef MG_ENTERPRISE
21
22 You, 21 hours ago | 2 authors (Andi Skrgat and one other)
23 export namespace memgraph::coordination {
24
25 Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
26 struct InstanceStateV1 {
27     bool is_replica;           // MAIN or REPLICA
28     std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
29     bool is_writing_enabled;   // on replica it's never enabled. On main depends.
30 };
31
32 Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
33 struct InstanceState {
34     bool is_replica;           // MAIN or REPLICA
35     std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
36     bool is_writing_enabled;   // on replica it's never enabled. On main depends.
37     std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
38     std::optional<std::map<std::string, std::map<std::string, int64_t>>>
39     | | replicas_num_txns; // if main, return num of committed txns for each instance
40
41     // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
42     InstanceStateV1 Downgrade() const {
43         return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
44     }
45
46 } // namespace memgraph::coordination
47 #endif
```

Module name



← → Q memgraph [SSH: lima-default]

instance\_state.cppm

src > coordination > instance\_state.cppm > ...

```
1 // Copyright 2025 Memgraph Ltd.      Andi Skrgat, 2 months ago * feat: Use num committed txns for failover and u...
9 // by the Apache License, Version 2.0, included in the file
10 // licenses/APL.txt.
11
12 module;
13
14 #include "utils/uuid.hpp"
15
16 #include <optional>
17
18 export module memgraph.coordination.instance_state;
19
20 #ifdef MG_ENTERPRISE
21
22 You, 21 hours ago | 2 authors (Andi Skrgat and one other)
23 export namespace memgraph::coordination {
24
25     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
26     struct InstanceStateV1 {
27         bool is_replica;           // MAIN or REPLICA
28         std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
29         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
30     };
31
32     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
33     struct InstanceState {
34         bool is_replica;           // MAIN or REPLICA
35         std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
36         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
37         std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
38         std::optional<std::map<std::string, std::map<std::string, int64_t>>>
39             | | replicas_num_txns; // if main, return num of committed txns for each instance
40
41         // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
42         InstanceStateV1 Downgrade() const {
43             return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
44         }
45     };
46 } // namespace memgraph::coordination
47 #endif
```

Global module fragment

# Global module fragment

- Used to include headers when importing modules isn't supported
- It must come at the beginning of the file

← → Q memgraph [SSH: lima-default]

File instance\_state.cppm X

src > coordination > instance\_state.cppm > ...

```
1 // Copyright 2025 Memgraph Ltd.      Andi Skrgat, 2 months ago * feat: Use num committed txns for failover and u...
9 // by the Apache License, Version 2.0, included in the file
10 // licenses/APL.txt.
11
12 module;
13
14 #include "utils/uuid.hpp"
15
16 #include <optional>
17
18 export module memgraph.coordination.instance_state;
19
20 #ifdef MG_ENTERPRISE
21
22 You, 21 hours ago | 2 authors (Andi Skrgat and one other)
23 export namespace memgraph::coordination {
24
25     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
26     struct InstanceStateV1 {
27         bool is_replica;           // MAIN or REPLICA
28         std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
29         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
30     };
31
32     Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
33     struct InstanceState {
34         bool is_replica;           // MAIN or REPLICA
35         std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
36         bool is_writing_enabled;    // on replica it's never enabled. On main depends.
37         std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
38         std::optional<std::map<std::string, std::map<std::string, int64_t>>>
39             | | replicas_num_txns; // if main, return num of committed txns for each instance
40
41         // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
42         InstanceStateV1 Downgrade() const {
43             return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
44         }
45     };
46 } // namespace memgraph::coordination
47 #endif
```





```
You, 1 second ago | 2 authors (You and one other)
22  namespace memgraph::coordination {
23
24  You, 1 second ago | 2 authors (Andi Skrgat and one other)
25  export struct InstanceStateV1 {
26      bool is_replica;           // MAIN or REPLICA
27      std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
28      bool is_writing_enabled;    // on replica it's never enabled. On main depends.
29  };
30
31  export [ Not Committed Yet You, now * Uncommitted changes
32  You, 1 second ago | 1 author (You)
33  struct InstanceState {
34      bool is_replica;           // MAIN or REPLICA
35      std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
36      bool is_writing_enabled;    // on replica it's never enabled. On main depends.
37      std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
38      std::optional<std::map<std::string, std::map<std::string, int64_t>>>
39          replicas_num_txns; // if main, return num of committed txns for each instance
40
41      // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
42      InstanceStateV1 Downgrade() const {
43          return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
44      }
45  };
46 } // namespace memgraph::coordination
47
48 #endif
49
```



```
You, 1 second ago | 2 authors (You and one other)
22  namespace memgraph::coordination {
23
You, 1 second ago | 2 authors (Andi Skrgat and one other)
24  export struct InstanceStateV1 {
25      bool is_replica;           // MAIN or REPLICA
26      std::optional<utils::UUID> uuid; // MAIN's UUID or the UUID which REPLICA listens
27      bool is_writing_enabled;    // on replica it's never enabled. On main depends.
28  };
29
30  export [ Not Committed Yet You, now * Uncommitted changes
You, 1 second ago | 1 author (You)
31  struct InstanceState {
32      bool is_replica;           // MAIN or REPLICA
33      std::optional<utils::UUID> uuid;           // MAIN's UUID or the UUID which REPLICA listens
34      bool is_writing_enabled;    // on replica it's never enabled. On main depends.
35      std::optional<std::map<std::string, uint64_t>> main_num_txns; // if main, returns db->num_committed_txns
36      std::optional<std::map<std::string, std::map<std::string, int64_t>>>
37          replicas_num_txns; // if main, return num of committed txns for each instance
38
39      // Follows the logic of other RPC versioning code. For responses, we downgrade newer version to the older version
40      InstanceStateV1 Downgrade() const {
41          return {.is_replica = is_replica, .uuid = uuid, .is_writing_enabled = is_writing_enabled};
42      }
43  };
44
45 } // namespace memgraph::coordination
46
47 #endif
48
49
```



# Key points

- For clang preferred extension is .cppm
- Modules are replacement for header files
- Global module fragment for includes
- Export means I am dealing with module interface units



```
coordination > C+ coordinator_cluster_state.cppm > ...
// Copyright 2025 Memgraph Ltd.
// by the Apache License, Version 2.0, included in the file
// licenses/APL.txt.

module;

#include "replication_coordination_glue/role.hpp"
#include "utils/resource_lock.hpp"
#include "utils/uuid.hpp"

#include <libnurraft/nurraft.hpp>
#include <nlohmann/json_fwd.hpp>

#include <string>

export module memgraph.coordination.coordinator_cluster_state;

#ifndef MG_ENTERPRISE

import memgraph.coordination.coordinator_instance_context;      feat: Modularize cluster state, as51340 (18 hours ago)
import memgraph.coordination.data_instance_context;
import memgraph.coordination.coordinator_communication_config;

You, 18 hours ago | 4 authors (Andi Skrgat and others)
export namespace memgraph::coordination {

    using nurraft::buffer;
    // NOLINTNEXTLINE
    using nurraft::buffer_serializer;
    using nurraft::ptr;
    // NOLINTNEXTLINE
    using replication_coordination_glue::ReplicationRole;

Andi Skrgat, 2 months ago | 1 author (Andi Skrgat)
struct CoordinatorClusterStateDelta {
    std::optional<std::vector<DataInstanceContext>> data_instances_;
    std::optional<std::vector<CoordinatorInstanceContext>> coordinator_instances_;
    std::optional<utils::UUID> current_main_uuid_;
    std::optional<bool> enabled_reads_on_main_;
    std::optional<bool> sync_failover_only_;
    std::optional<uint64_t> max_failover_replica_lag_;
    std::optional<uint64_t> max_replica_read_lag_;

    bool operator==(const CoordinatorClusterStateDelta &other) const = default;
};

    
```



```
module;

#include <mutex>
#include <nlohmann/json.hpp>
#include <shared_mutex>

#include "utils/uuid.hpp"

module memgraph.coordination.coordinator_cluster_state;

#ifndef MG_ENTERPRISE

import memgraph.coordination.constants;

Andi Skrgat, 2 months ago | 3 authors (Andi Skrgat and others)
namespace memgraph::coordination {      Fix client-side routing in K8s HA deployment scenario, Andi Skrgat

CoordinatorClusterState::CoordinatorClusterState(CoordinatorClusterState const &other) {
    auto lock: std::lock_guard<ResourceLock> = std::lock_guard{&other.app_lock_};
    // NOLINTBEGIN
    data_instances_ = other.data_instances_;
    coordinator_instances_ = other.coordinator_instances_;
    current_main_uuid_ = other.current_main_uuid_;
    enabled_reads_on_main_ = other.enabled_reads_on_main_;
    sync_failover_only_ = other.sync_failover_only_;
    max_failover_replica_lag_ = other.max_failover_replica_lag_;
    max_replica_read_lag_ = other.max_replica_read_lag_;
    // NOLINTEND
}

CoordinatorClusterState &CoordinatorClusterState::operator=(CoordinatorClusterState const &other) {
    if (this == &other) {
        return *this;
    }
    std::scoped_lock const lock{app_lock_, other.app_lock_};

    data_instances_ = other.data_instances_;
    coordinator_instances_ = other.coordinator_instances_;
```



```
module;

#include <mutex>
#include <nlohmann/json.hpp>
#include <shared_mutex>

#include "utils/uuid.hpp"

module memgraph.coordination.coordinator_cluster_state;

#ifndef MG_ENTERPRISE

import memgraph.coordination.constants;

Andi Skrgat, 2 months ago | 3 authors (Andi Skrgat and others)
namespace memgraph::coordination {      Fix client-side routing in K8s HA deployment scenario, Andi Skrgat

CoordinatorClusterState::CoordinatorClusterState(CoordinatorClusterState const &other) {
    auto lock: std::lock_guard<ResourceLock> = std::lock_guard{&other.app_lock_};
    // NOLINTBEGIN
    data_instances_ = other.data_instances_;
    coordinator_instances_ = other.coordinator_instances_;
    current_main_uuid_ = other.current_main_uuid_;
    enabled_reads_on_main_ = other.enabled_reads_on_main_;
    sync_failover_only_ = other.sync_failover_only_;
    max_failover_replica_lag_ = other.max_failover_replica_lag_;
    max_replica_read_lag_ = other.max_replica_read_lag_;
    // NOLINTEND
}

CoordinatorClusterState &CoordinatorClusterState::operator=(CoordinatorClusterState const &other) {
    if (this == &other) {
        return *this;
    }
    std::scoped_lock const lock{app_lock_, other.app_lock_};

    data_instances_ = other.data_instances_;
    coordinator_instances_ = other.coordinator_instances_;
```

# Named modules

- A collection of module units with the same module name
- If there is an export ⇒ module interface units
- Everything else ⇒ module implementation units
- There can be more than one impl. units but only a single interface unit



memgraph [SSH: lima-default]

src > coordination > coordinator\_cluster\_state\_1.cpp U

```
1 // Copyright 2025 Memgraph Ltd.
2 //
3 // Use of this software is governed by the Business Source License
4 // included in the file licenses/BSL.txt; by using this file, you agree to be bound by the terms of the Business Source
5 // License, and you may not use this file except in compliance with the Business Source License.
6 //
7 // As of the Change Date specified in that file, in accordance with
8 // the Business Source License, use of this software will be governed
9 // by the Apache License, Version 2.0, included in the file
10 // licenses/APL.txt.
11
12 module;
13
14 #include <mutex>
15
16 module memgraph.coordination.coordinator_cluster_state;
17
18 #ifdef MG_ENTERPRISE
19
20 import memgraph.coordination.constants;
21
22 namespace memgraph::coordination {
23
24 CoordinatorClusterState::CoordinatorClusterState(CoordinatorClusterState const &other) {
25     auto lock: std::lock_guard<ResourceLock> = std::lock_guard{&other.app_lock_};
26     // NOLINTBEGIN
27     data_instances_ = other.data_instances_;
28     coordinator_instances_ = other.coordinator_instances_;
29     current_main_uuid_ = other.current_main_uuid_;
30     enabled_reads_on_main_ = other.enabled_reads_on_main_;
31     sync_failover_only_ = other.sync_failover_only_;
32     max_failover_replica_lag_ = other.max_failover_replica_lag_;
33     max_replica_read_lag_ = other.max_replica_read_lag_;
34     // NOLINTEND
35 }
36
37 } // namespace memgraph::coordination
38#endif
39
```

Module impl unit 1



←

→

memgraph [SSH: lima-default]



src > coordination > coordinator\_cluster\_state\_2.cpp > ...

```
1 // Copyright 2025 Memgraph Ltd.
2 //
3 // Use of this software is governed by the Business Source License
4 // included in the file licenses/BSL.txt; by using this file, you agree to be bound by the terms of the Business Source
5 // License, and you may not use this file except in compliance with the Business Source License.
6 //
7 // As of the Change Date specified in that file, in accordance with
8 // the Business Source License, use of this software will be governed
9 // by the Apache License, Version 2.0, included in the file
10 // licenses/APL.txt.
11
12 module;
13
14 #include <mutex>
15 #include <nlohmann/json.hpp>
16 #include <shared_mutex>
17
18 #include "utils/uuid.hpp"
19
20 module memgraph.coordination.coordinator_cluster_state
21
22 #ifdef MG_ENTERPRISE
23
24 import memgraph.coordination.constants;
25
26 namespace memgraph::coordination {
27
28 CoordinatorClusterState &CoordinatorClusterState::operator=(CoordinatorClusterState const &other) {
29     if (this == &other) {
30         return *this;
31     }
32     std::scoped_lock const lock{app_lock_, other.app_lock_};
33
34     data_instances_ = other.data_instances_;
35     coordinator_instances_ = other.coordinator_instances_;
36     current_main_uuid_ = other.current_main_uuid_;
37     enabled_reads_on_main_ = other.enabled_reads_on_main_;
38     sync_failover_only_ = other.sync_failover_only_;
39     max_failover_replica_lag_ = other.max_failover_replica_lag_;
40     max_replica_read_lag_ = other.max_replica_read_lag_;
41     return *this;
42 }
43
44 CoordinatorClusterState::CoordinatorClusterState(CoordinatorClusterState &&other) noexcept
45     : data_instances_{std::move(other.data_instances_)},
46       coordinator_instances_{std::move(other.coordinator_instances_)},
```

## Module impl unit 2



```
class CoordinatorInstanceManagementServer {
public:
    explicit CoordinatorInstanceManagementServer(const ManagementServerConfig &config);
    CoordinatorInstanceManagementServer(const CoordinatorInstanceManagementServer &) = delete;
    CoordinatorInstanceManagementServer(CoordinatorInstanceManagementServer &&) = delete;
    CoordinatorInstanceManagementServer &operator=(const CoordinatorInstanceManagementServer &) = delete;
    CoordinatorInstanceManagementServer &operator=(CoordinatorInstanceManagementServer &&) = delete;

    ~CoordinatorInstanceManagementServer();

    bool Start();

    template <typename TRequestResponse, typename F>
    void Register(F &&callback) {
        rpc_server_.Register<TRequestResponse>(callback: std::forward<F>(callback));
    }

private:
    communication::ServerContext rpc_server_context_;
    rpc::Server rpc_server_;
};

} // namespace memgraph::coordination

module :private;
You, 22 hours ago | 1 author (You)
namespace memgraph::coordination {

You, 22 hours ago | 1 author (You)
namespace {

// NOTE: The coordinator server doesn't need more than 1 processing thread - it's not a bottleneck
constexpr auto kCoordInstanceManagementServerThreads: const int = 1;

} // namespace

CoordinatorInstanceManagementServer::CoordinatorInstanceManagementServer(const ManagementServerConfig &config)
    : rpc_server_context_(communication::ServerContext{)},
      rpc_server_{endpoint: config.endpoint, context: &rpc_server_context_, workers_count: kCoordInstanceManagementServerThreads} {}

CoordinatorInstanceManagementServer::~CoordinatorInstanceManagementServer() {
    if (rpc_server_.IsRunning()) {
        rpc_server_.Shutdown();
    }
    rpc_server_.AwaitShutdown();
}

bool CoordinatorInstanceManagementServer::Start() { return rpc_server_.Start(); }

} // namespace memgraph::coordination

#endif
```

# Private module fragment

- Both declaration and implementation in the .cppm
- Implementation changes don't require whole module recompilation and its dependencies
- You can have less files with the same functionality

coordinator\_cluster\_state.cppm

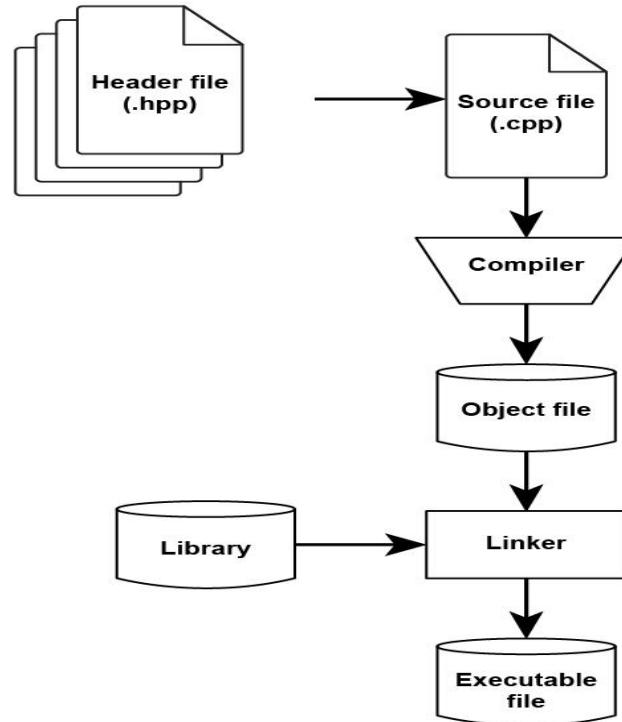
CMakeLists.txt

memgraph [SSH: lima-default]



```
src > coordination > CMakeLists.txt
  You, 18 minutes ago | 5 authors (You and others)
1 add_library(mg-coordination STATIC)
2 add_library(mg::coordination ALIAS mg-coordination)
3 target_sources(mg-coordination
4   PUBLIC
5   FILE_SET CXX_MODULES
6   FILES      feat: Start modularizing mg-coordination lib, a551340 (2 days ago)      You, 3 days ago • feat: Start modularizing mg-coordination lib
7   coordinator_cluster_state.cppm
8   coordinator_communication_config.cppm
9   coordinator_exceptions.cppm
10  coordinator_instance_aux.cppm
11  coordinator_instance_context.cppm
12  coordinator_log_store.cppm
13  coordinator_ops_status.cppm
14  constants.cppm
15  data_instance_context.cppm
16  instance_state.cppm
17  instance_status.cppm
18  logger.cppm
19  log_level.cppm
20  logger_wrapper.cppm
21  replication_lag_info.cppm
22  utils.cppm
23 FILE_SET HEADERS
24 FILES
25 include/coordination/coordinator_instance.hpp
26 include/coordination/coordinator_instance_client.hpp
27 include/coordination/coordinator_instance_connector.hpp
28 include/coordination/coordinator_instance_management_server.hpp
29 include/coordination/coordinator_instance_management_server_handlers.hpp
30 include/coordination/coordinator_observer.hpp
31 include/coordination/coordinator_rpc.hpp
32 include/coordination/coordinator_slk.hpp
33 include/coordination/coordinator_state.hpp
34 include/coordination/coordinator_state_machine.hpp
35 include/coordination/coordinator_state_manager.hpp
36 include/coordination/data_instance_management_server.hpp
37 include/coordination/data_instance_management_server_handlers.hpp
38 include/coordination/raft_state.hpp
39 include/coordination/replication_instance_client.hpp
40 include/coordination/replication_instance_connector.hpp
41
42
43 PRIVATE
44 coordinator_cluster_state_1.cpp
45 coordinator_cluster_state_2.cpp
```

# Before modules



# Standard translation unit

- Highly parallel build
- A lot of code gets propagated even if not used through transitive inclusion to the translation unit
- Relationship between a header file and translation unit is more a convention (cpp file can be used independently)

# Modules

- Need an interface in order to import them
- When a module is compiled we get two files:
  - An object file for linking
  - Binary Module Interface (BMI)

# Binary Module Interface

- From OpenStd: The artifact created by a compiler to represent a module unit. The format is implementation specific.
  - All entities that are exported by a module in a format suitable for fast imports and lookups
- Reduced code bloating and better physical isolation but strong dependencies between modules ⇒ hard to parallelize, we are searching for DAG between modules
- clang-scan-deps tool used to find deps between modules



Alacritty

```
(env) *[feat/update-ha-config]~/Memgraph/code/memgraph]$ clang++ -std=c++20 -fmodules -fmodule-output=coordinator_ops.pcm -c src/coordination/coordinator_ops_status.cpp
(env) *[feat/update-ha-config]~/Memgraph/code/memgraph]$ ls
ADRs           CODE_OF_CONDUCT.md  README.md  cmake-build-debug  config           env      init      mg_data      src
CHANGELOG.md    CONTRIBUTING.md   build      conan.lock       coordinator_ops.pcm  environment  init-test  pyproject.toml  tests
CMakeLists.txt  Doxyfile        build.sh   conan_config    coordinator_ops_status.o  import      libs      query_modules  tools
CMakeUserPresets.json  LICENSE      cmake      conanfile.py   docs           include    licenses  release      tsan.sup
(env) *[feat/update-ha-config]~/Memgraph/code/memgraph$
```

# Modules != pre-compiled headers

- Pre-compiled headers still need to comply with the header inclusion model (linear dependency chain)
- Best to use PCH when having a big-sized header included at multiple places

# Modularizing mg-coordination

- Experiment: Modularize mg-coordination and measure build times
  - Build time of memgraph
  - Build time when mg-coordination.a is missing
  - Size of mg-coordination.a
- <https://github.com/memgraph/memgraph/pull/3491>

# Modules are space and time efficient

- Module contents aren't transitively propagated if not requested explicitly
- For this you have **export import**
- Instead of including each header file into a cpp file and then compiling it, you are compiling it just once ⇒ smaller translation units
- You use BMI to only import some of the symbols from the module, not whole module code

# Results

	Build memgraph	Build libmg-coordination.a	Lib size
No modules	3:20.05	19.878	37MB
Modules	4:39.11	1:04.38	38MB



# Modules lead to better abstractions

- Very fine-grained control over what's visible and what's not ⇒ you can be very precise in specifying what are your needs
- Comparison with microservices
- Boost example with `property_value.hpp`



Changes from all commits ▾ File filter ▾

replication\_instance\_con...

replication\_lag\_info.cppm

utils.cpp

utils.cppm

dbms

coordinator\_handler.cpp

coordinator\_handler.hpp

dbms\_handler.hpp

inmemory

replication\_handlers.cpp

replication\_handlers.hpp

replication\_handlers.cpp

memgraph.cpp

query

dump.cpp

interpreter.cpp

interpreter.hpp

interpreter\_context.hpp

replication\_coordination\_glue

common.hpp

mode.hpp

replication\_handler

include/replication\_handler

replication\_handler.hpp

replication\_handler.cpp

MQL

RPC

Viewed

...

...

3 CMakeLists.txt

@@ -206,6 +206,9 @@ find\_package(spdlog REQUIRED)

```
206 find_package(fmt REQUIRED)
207 find_package(strong_type REQUIRED)
208 find_package(Boost REQUIRED)

209 find_package(BZip2 REQUIRED)
210 + find_package(antlr4-runtime REQUIRED)
211 find_package(cppitertools REQUIRED)
```

206 find\_package(fmt REQUIRED)
207 find\_package(strong\_type REQUIRED)
208 find\_package(Boost REQUIRED)

209 + # Ensure Conan boost headers take precedence over toolchain
210 + get\_target\_property(BOOST\_INCLUDE\_DIRS Boost::headers
211 + INTERFACE\_INCLUDE\_DIRECTORIES)
212 + include\_directories(BEFORE SYSTEM \${BOOST\_INCLUDE\_DIRS})
213 find\_package(BZip2 REQUIRED)
214 find\_package(antlr4-runtime REQUIRED)
215 find\_package(cppitertools REQUIRED)

# Caveats

- Not every .hpp+.cpp pair can be translated into a .cppm with a private module fragment (example is the issue with nlohmann/json.hpp)
- Forward declarations don't work across modulem, either use module partitions or refactor
- Hard to mix .hpp and .cppm files ⇒ couldn't modularize "coordinator\_rpc.hpp" independently from "rpc/utils.hpp"
- Build order matters, changes in the mg-coordination caused a failure to compile a property\_value.hpp
- import std ⇒ not working in our case



# Resources

- [https://cmake.org/cmake/help/latest/command/target\\_sources.html#file-sets](https://cmake.org/cmake/help/latest/command/target_sources.html#file-sets)
- <https://www.youtube.com/watch?v=7WK42YSfE9s&t=1s>
- <https://cryos.net/2024/01/c-modules-and-cmake/>
- [https://www.youtube.com/watch?v=x9K9\\_q2ZXE&t=3209s](https://www.youtube.com/watch?v=x9K9_q2ZXE&t=3209s)
- [https://www.youtube.com/watch?v=l\\_83IyxWGtE&t=3040s](https://www.youtube.com/watch?v=l_83IyxWGtE&t=3040s)

# 02

## std::stacktrace

---





<https://godbolt.org/z/dPb75GMs>

# 03

## **std::expected**





Run this code

```
#include <cmath>
#include <expected>
#include <iomanip>
#include <iostream>
#include <string_view>

enum class parse_error
{
    invalid_input,
    overflow
};

auto parse_number(std::string_view str) -> std::expected<double, parse_error>
{
    const char* begin = str.data();
    char* end;
    double retval = std::strtod(begin, &end);

    if (begin == end)
        return std::unexpected(parse_error::invalid_input);
    else if (std::isinf(retval))
        return std::unexpected(parse_error::overflow);

    str.remove_prefix(end - begin);
    return retval;
}

int main()
{
    auto process = [] (std::string_view str)
    {
        std::cout << "str: " << std::quoted(str) << ", ";
        if (const auto num = parse_number(str); num.has_value())
            std::cout << "value: " << *num << '\n';
        // If num did not have a value, dereferencing num
        // would cause an undefined behavior, and
        // num.value() would throw std::bad_expected_access.
        // num.value_or(123) uses specified default value 123.
        else if (num.error() == parse_error::invalid_input)
            std::cout << "error: invalid input\n";
        else if (num.error() == parse_error::overflow)
            std::cout << "error: overflow\n";
        else
            std::cout << "unexpected!\n"; // or invoke std::unreachable();
    };

    for (auto src : {"42", "42abc", "meow", "inf"})
        process(src);
}
```

# 04

# Practical

---





# Task recommendations

- Figure out why modularizing mg-coordination slowed down the build ⇒ RESEARCH
- Modularize some other library (partially) and measure the effect ⇒ IMPL
- Make use of std::stacktrace in our exceptions
- Replace utils::BasicResult with std::expected
- Remove range-v3 dependency from Conan
- Std::move\_only\_function
-



# Thank you for your time!



[www.memgraph.com](http://www.memgraph.com)