

Document number: P3427R2

Date: 2025-08-11

Project: Programming Language C++, LEWG

Authors: Maged M. Michael, Michael Wong, Paul McKenney

Email: maged.michael@gmail.com, fraggamuffin@gmail.com, paulmck@kernel.org

Hazard Pointer Synchronous Reclamation

Table of Contents

Introduction	2
Background: P2530R3 C++26 Hazard Pointers	2
Motivation	3
Implementation and Use Experience	3
Synchronous Reclamation	3
Global Cleanup	3
Object Cohorts	4
Proposed Interface	4
Draft Wording	5
Usage Example	6
Separating Cohort Object Retirement from Asynchronous Reclamation	7
History	8
References	8

Introduction

This paper proposes extending the working draft (N5008) C++26 hazard pointer interface to support synchronous reclamation.

The R1 revision of this paper was reviewed by SG1 in Wroclaw 2024 and forwarded to LEWG with feedback "The word 'asynchronous' doesn't belong in the name of the free function". This revision, P3427R2, revises R1 by following SG1 feedback.

Background: P2530R3 C++26 Hazard Pointers

Hazard pointer interface from P2530R3 merged into the working draft (N5008):

```
template <class T, class D = default_delete<T>>
class hazard_pointer_obj_base {
public:
    void retire(D d = D()) noexcept;
protected:
    hazard_pointer_obj_base() = default;
    hazard_pointer_obj_base(const hazard_pointer_obj_base&) = default;
    hazard_pointer_obj_base(hazard_pointer_obj_base&&) = default;
    hazard_pointer_obj_base& operator=(const hazard_pointer_obj_base&) = default;
    hazard_pointer_obj_base& operator=(hazard_pointer_obj_base&&) = default;
    ~hazard_pointer_obj_base() = default;
private:
    D deleter ; // exposition only
};

class hazard_pointer {
public:
    hazard_pointer() noexcept;
    hazard_pointer(hazard_pointer&&) noexcept;
    hazard_pointer& operator=(hazard_pointer&&) noexcept;
    ~hazard_pointer();
    [[nodiscard]] bool empty() const noexcept;
    template <class T> T* protect(const atomic<T*>& src) noexcept;
    template <class T> bool try_protect(T*& ptr, const atomic<T*>& src) noexcept;
    template <class T> void reset_protection(const T* ptr) noexcept;
    void reset_protection(nullptr_t = nullptr) noexcept;
    void swap(hazard_pointer&) noexcept;
};

hazard_pointer make_hazard_pointer();
void swap(hazard_pointer&, hazard_pointer&) noexcept;
```

Motivation

This paper proposes supporting object cohorts due to the importance of synchronous reclamation for general purpose usability. For example, a concurrent hash map that uses hazard pointers is more generally usable if it allows arbitrary key and value types rather than only types without dependence on resources with independent lifetimes.

Implementation and Use Experience

Object cohorts have been part of the Folly open-source library (under the name `hazptr_obj_cohort`) and in heavy use in production since 2018. (See [CppCon 2021 Hazard Pointer Synchronous reclamation beyond Concurrency TS2](#) for details about the evolution of support for synchronous reclamation in Folly).

Synchronous Reclamation

The P2530R3 C++26 hazard pointer interface supports only asynchronous reclamation which does not guarantee the timing of the reclamation of protectable objects that are no longer protected. As a result, hazard pointer users must guarantee separately that the deleters of such objects do not depend on resources that may become subsequently unavailable.

Support for synchronous reclamation allows users to synchronously induce and wait for the reclamation of unprotected objects.

Global Cleanup

A straightforward albeit inefficient solution to this problem is global cleanup, which guarantees the completion of deleters of all unprotected retired objects. This involves synchronously checking all retired objects against all hazard pointers.

The main drawback of the global cleanup approach is its high overhead that makes it impractical to use. For example, it may be useful to include global cleanup in the destructor of a generic container library. However, the prohibitive overhead of global cleanup makes that impractical for many use cases of such a container.

Another drawback of the global cleanup approach is that it adds overhead to the hazard pointer implementation even when users never use this feature (e.g., would require synchronization on thread local private buffers of retired objects that would otherwise be unnecessary).

We do not recommend the standardization of the global cleanup approach due to these drawbacks and the lack of production experience of the necessity of such approach in contrast to the more efficient approach discussed in the following section.

Object Cohorts

An alternative approach is the use of object cohorts, which are sets of protectable objects. Object cohorts support synchronous reclamation by guaranteeing that

all the deleters of the object cohort members are completed before the completion of the cohorts destructor.

The main advantage of the object cohort approach is its performance. While its guarantees are weaker and less flexible than global cleanup, its efficiency enables users to use it in performance sensitive cases where the cost of global cleanup may be impractical. It strikes a better balance between practicality and performance. Therefore, we recommend object cohorts for standardization.

Proposed Interface

```
class hazard_pointer_cohort {
    hazard_pointer_cohort() noexcept;
    hazard_pointer_cohort(const hazard_pointer_cohort&) = delete;
    hazard_pointer_cohort(hazard_pointer_cohort&&) = delete;
    hazard_pointer_cohort& operator=(const hazard_pointer_cohort&) = delete;
    hazard_pointer_cohort& operator=(hazard_pointer_cohort&&) = delete;
    ~hazard_pointer_cohort();
};

template <class T, class D = default_delete<T>>
class hazard_pointer_obj_base {
public:
    void retire_to_cohort(hazard_pointer_cohort& c, D d = D()) noexcept;
};

void hazard_pointer_try_reclamation() noexcept;
```

Draft Wording

A *hazard-cohort* is a set of hazard-protectable objects. An object of type **hazard_pointer_cohort** is a hazard-cohort.

When the member function **hazard_pointer_obj_base<T,D>::retire_to_cohort** is invoked on an object **x**, cohort **c**, and a deleter of type **D**, **x** is retired, **x** becomes a member of **c**. Any given object **x** shall be a member of at most one hazard-cohort.

hazard_pointer_cohort();

Postconditions: ***this**. is empty.

~hazard_pointer_cohort();

Effects: Reclaims all members of ***this**.

Synchronization: The completion of the deleter for each reclaimed object synchronizes with the return from this destructor.

void retire_to_cohort(hazard_pointer_cohort& c, D d = D()) noexcept;

Preconditions: ***this** is a base class subobject of an object **x** of type **T**. **x** is not retired. Move-assigning **d** to **deleter** does not exit via an exception.

Effects: Move-assigns **d** to **deleter**, thereby setting it as the deleter of **x**, then retires **x**. May reclaim possibly-reclaimable members of **c**.

Postcondition: ***this** is a member of **c**.

void hazard_pointer_try_reclamation() noexcept;

Effects: May reclaim possibly-reclaimable objects.

Usage Example

The following table shows two code snippets one using the P2530R3 C++26 hazard pointer interface and one using object cohorts, respectively. The latter supports synchronous reclamation.

P2530R3 C++26 (Asynchronous Reclamation Only)	Cohort-Based Synchronous Reclamation
<pre> /// Library template <class T> class Container { class Obj : hazard_pointer_obj_base<Obj> { T data; <i>/* etc */</i> }; void insert(T data) { Obj* obj = new Obj(data); <i>/* Insert obj in container */</i> } void erase(Args args) { Obj* obj = find(args); <i>/* Remove obj from container */</i> executor_.add([] { obj->retire(); }); } }; /// User class A { <i>// Deleter cannot depend on resources</i> <i>// with independent lifetime.</i> ~A(); }; { Container<A> container; container.insert(a); container.erase(a); } <i>// Obj containing 'a' may be not deleted</i> <i>// yet.</i> </pre>	<pre> /// Library template <class T> class Container { class Obj : hazard_pointer_obj_base<Obj> { T data; <i>/* etc */</i> }; hazard_pointer_cohort cohort_; void insert(T data) { Obj* obj = new Obj(data); <i>/* Insert obj in container */</i> } void erase(Args args) { Obj* obj = find(args); <i>/* Remove obj from container */</i> obj->retire_to_cohort(cohort_); executor_.add([] { hazard_pointer_try_reclamation(); }); } }; /// User class B { <i>// Deleter may depend on resources</i> <i>// with independent lifetime.</i> ~B() { use_resource_XYZ(); } }; make_resource_XYZ(); { Container container; container.insert(b); container.erase(b); } <i>// Obj containing 'b' was deleted.</i> destroy_resource_XYZ(); </pre>

In the above example:

- In the code snippet on the left side, the removed object is retired asynchronously by submitting the retirement code to a dedicated thread pool (details not shown) to be executed asynchronously. Doing so avoids burdening the current thread with potentially performing amortized reclamation of tens of thousands of possibly unrelated retired objects, because calling **retire** may trigger amortized asynchronous reclamation.
- In contrast, the code snippet on the right side, retirement to the cohort must be executed synchronously in order for the object to be included as intended in synchronous reclamation of the associated cohort. Therefore, the call to **retire_to_cohort** must be synchronous. But in order not to burden worker threads with amortized asynchronous reclamation, **retire_to_cohort** does not try to perform asynchronous reclamation. Instead, if desired, the current thread may submit an asynchronous task to a dedicated thread pool to try to perform asynchronous reclamation (which attempts to reclaim unprotected retired objects, both cohort-associated and not).

Separating Cohort Object Retirement from Asynchronous Reclamation

Asynchronous Reclamation of Cohort Objects

If a cohort is long-lived, large numbers (e.g., billions) of objects may be retired to it. If such objects are not included in asynchronous reclamation, they would remain not reclaimed. Therefore, cohort objects need to be included in asynchronous reclamation.

Cohort Object Retirement Must to Be Synchronous

In order for an object to be included in synchronous reclamation in association with a cohort, the object must be retired to the cohort. Therefore, cohort object retirement needs to be synchronous.

Why Doesn't **retire_to_cohort** Try to Invoke Asynchronous Reclamation Implicitly?

As indicated above cohort-object retirement must be synchronous. However, it is often desirable to invoke asynchronous reclamation asynchronously in order to avoid burdening the thread retiring the object (typically a worker thread) with reclaiming a large number (e.g., tens of thousands) of (possibly unrelated) retired objects.

In contrast, the retirement of non-cohort objects (using **retire**) by convention may implicitly invoke asynchronous reclamation. If desired, a user may retire a non-cohort object **obj** asynchronously to avoid inline asynchronous reclamation as follows (assuming an executor associated with a separate execution resource):

```
executor_.add([obj] { obj->retire(); });
```

However, the retirement of a cohort object needs to be synchronous, as mentioned above. Therefore a free function **hazard_pointer_try_reclamation()** is added, so that users can write the cohort equivalent to the above non-cohort code snippet as follows:

```
obj->retire_to_cohort(cohort_);  
executor_.add([] { hazard_pointer_try_reclamation(); });
```

History

The Varna 2023 plenary voted in favor of including hazard pointers in the C++26 standard library ([2023-06 LWG Motion 7] P2530R3 Hazard Pointers for C++26).

P3135R1 was presented to SG1 in Tokyo 2024, reviewing potential extensions of the P2530R3 C++26 interface, and proposing two of those for inclusion in the standard library. The proposal for extending the P2530R3 C++26 interface to support synchronous reclamation was voted on by SG1 with unanimous consent:

```
we want to continue work on hazard pointer cohorts (synchronous reclamation) for C++26,
with association to the cohort registered at the time of retire()
```

R0

P3427R0 was a follow up on P3135R1, focusing on extending the P2530R3 C++26 hazard pointer interface to support synchronous reclamation, revised to take into account the feedback from SG1:

- Associate a hazard-protectable object with a cohort at the time of its call to retire().

R1

R1 added draft wording to R0.

D3427R1 was reviewed by SG1 in Wroclaw 2024 and SG1 voted with unanimous consent to forward (D)P3427R1 to LEWG for C++26 with feedback:

```
Forward (D)P3427R1 to LEWG for C++26 with notes:
```

```
* The word 'asynchronous' doesn't belong in the name of the free function
```

R2

R2 revises R1 to follow the feedback from SG1 and for review by LEWG. It makes the following changes from R1:

- Change **hazard_pointer_asynchronous_reclamation** to **hazard_pointer_try_reclamation**.
- Added "May reclaim possibly-reclaimable members of c." to the wording of `reclaim to cohort`.

References

- [P2530R3](#): Hazard Pointers for C++26 (2023-03-02).
- [P3135R1](#): Hazard Pointer Extensions (2024-04-12).
- [Folly](#): Facebook Open-source Library.
- [CppCon 2021: Hazard Pointer Synchronous reclamation beyond Concurrency TS2](#)