P0461R1: Proposed RCU C++ API

Doc. No.: WG21/P0461R1
Date: 2017-02-07
Reply to: Paul E. McKenney, Maged Michael, Michael Wong, Isabella Muerte, Arthur O'Dwyer, and David Hollman
Email: paulmck@linux.vnet.ibm.com, maged.michael@gmail.com, fraggamuffin@gmail.com, isabella.muerte@mnmlstc.com, arthur.j.odwyer@gmail.com, and dshollm@sandia.gov

February 7, 2017

This document is based on WG21/P0279R1 combined with feedback at the 2015 Kona and 2016 Jacksonville meetings, which most notably called for a C++-style method of handling different RCU implementations or domains within a single translation unit, and which also contains useful background material and references. Unlike WG21/P0279R1, which simply introduced RCU's C-language practice, this document presents proposals for C++-style RCU APIs. At present, it appears that these are not conflicting proposals, but rather ways of handling different C++ use cases resulting from inheritance, templates, and different levels of memory pressure. This document also incorporates content from WG21/P0232R0[4].

Note that this proposal is related to the hazard-pointer proposal in that both proposals defer destructive actions such as reclamation until all readers have completed. See P0233R3, which updates "P0233R2: Hazard Pointers: Safe Resource Reclamation for Optimistic Concurrency" at `http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0233r2.pdf`.

Note also that a redefinition of the infamous `memory_order_consume` is the subject of two separate papers:

1. P0190R3, which updates "P0190R2: Proposal for New `memory_order_consume` Definition", `http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0190r2.pdf`.
2. P0462R1, which updates "P0462R0: Marking `memory_order_consume` Dependency Chains", `http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0462r0.pdf`.

Draft wording for this proposal may be found in the new working paper "P0566R0: Proposed Wording for Concurrent Data Structures: Hazard Pointer and Read-Copy-Update (RCU)".

A detailed change log appears starting on page 12.

```
 1 void std::rcu_read_lock();
 2 void std::rcu_read_unlock();
 3 void std::synchronize_rcu();
 4 void std::call_rcu(struct std::rcu_head *rhp,
 5                    void cbf(class rcu_head *rhp));
 6 void std::rcu_barrier();
 7 void std::rcu_register_thread();
 8 void std::rcu_unregister_thread();
 9 void std::rcu_quiescent_state();
10 void std::rcu_thread_offline();
11 void std::rcu_thread_online();
```

Figure 1: Existing C-Language RCU API

1 Introduction

This document proposes C++ APIs for read-copy update (RCU). For more information on RCU, including RCU semantics, see WG21/P0462R0 ("Marking `memory_order_consume` Dependency Chains"), WG21/P0279R1 ("Read-Copy Update (RCU) for C++"), WG21/P0190R2 ("Proposal for New `memory_order_consume` Definition"), and WG21/P0098R1 ("Towards Implementation and Use of `memory_order_consume`").

Specifically, this document proposes `rcu_domain` (Figure 2), `rcu_guard` (Figure 3), and `rcu_obj_base` (Figure 4).

Section 2 presents the base (C-style) RCU API, Section 3 presents a proposal for scoped RCU readers, Section 4 presents proposals for handling of RCU callbacks (with that of Section 4.1 being the preferred implementation), Section 5 presents a table comparing reference counting, hazard pointers, and RCU, and finally Section 6 presents a summary.

2 Existing C-Language RCU API

Figure 1 shows the existing C-language RCU API as provided by implementations such as userspace RCU [1, 3]. This API is provided for compatibility with existing C-language practice as well as to provide the highest performance for fast-path code. (See Figure 2 for a proposed API that permits multiple RCU domains, as requested by several committee members.)

Lines 1 and 2 show `rcu_read_lock()` and `rcu_read_unlock()`, which mark the beginning and the end, respectively, of an *RCU read-side critical section*. These primitives may be nested, and matching `rcu_read_lock()` and `rcu_read_unlock()` calls need not be in the same scope. (That said, it is good practice to place them in the same scope in cases where the entire critical section fits comfortably into one scope.)

Line 3 shows `synchronize_rcu()`, which waits for any pre-existing RCU read-side critical sections to complete. The period of time that `synchronize_rcu()` is required to wait is called a *grace period*. Note that a given call to `synchronize_rcu()` is *not* required to wait for critical sections that start later.

Lines 4 and 5 show `call_rcu()`, which, after a subsequent grace period elapses, causes the `cbf(rhp)` *RCU callback function* to be invoked. Thus, `call_rcu()` is the asynchronous counterpart to `synchronize_rcu()`. In most cases, `synchronize_rcu()` is easier to use, however, `call_rcu()` has the benefit of moving the grace-period delay off of the updater's critical path. Use of `call_rcu()` is thus critically important for good performance of update-heavy workloads, as has been repeatedly discovered by any number of people new to RCU [2].

Note that although `call_rcu()`'s callbacks are guaranteed not to be invoked too early, there is no guarantee that their execution won't be deferred for a considerable time. This can be a problem if a given program requires that all outstanding RCU callbacks be invoked before that program terminates. The `rcu_barrier()` function shown on line 6 is intended for this situation. This function blocks until all callbacks corresponding to previous `call_rcu()` invocations have been invoked and also until after those invocations have returned. Therefore, taking the following steps just before terminating a program will guarantee that all callbacks have completed:

1. Take whatever steps are required to ensure that there are no further invocations of `call_rcu()`.

2. Invoke `rcu_barrier()`.

Carrying out this procedure just prior to program termination can be very helpful for avoiding false positives when using tools such as valgrind.

Many RCU implementations require that every thread announce itself to RCU prior to entering the first RCU read-side critical section, and to announce its departure after exiting the last RCU read-side critical section. These tasks are carried out via the `rcu_register_thread()` and `rcu_unregister_thread()`, respectively.

The implementations of RCU that feature the most aggressive implementations of `rcu_read_lock()` and `rcu_read_unlock()` require that each thread periodically pass through a *quiescent state*, which is announced to RCU using `rcu_quiescent_state()`. A thread in a quiescent state is guaranteed not to be in an RCU read-side critical section. Threads can also announce entry into and exit from *extended quiescent states*, for example, before and after blocking system calls, using `rcu_thread_offline()` and `rcu_thread_online()`.

2.1 RCU Domains

The userspace RCU library features several RCU implementations, each optimized for different use cases.

The quiescent-state based reclamation (QSBR) implementation is intended for standalone applications where the developers have full control over the entire application, and where extreme read-side performance and scalability is required. Applications use `#include "urcu-qsbr.hpp"` to select QSBR and

```
 1 class rcu_domain {
 2 public:
 3   constexpr explicit rcu_domain() noexcept { };
 4   rcu_domain(const rcu_domain&) = delete;
 5   rcu_domain(rcu_domain&&) = delete;
 6   rcu_domain& operator=(const rcu_domain&) = delete;
 7   rcu_domain& operator=(rcu_domain&&) = delete;
 8   virtual void register_thread() = 0;
 9   virtual void unregister_thread() = 0;
10   static constexpr bool register_thread_needed() { return true; }
11   virtual void quiescent_state() noexcept = 0;
12   virtual void thread_offline() noexcept = 0;
13   virtual void thread_online() noexcept = 0;
14   static constexpr bool quiescent_state_needed() { return false; }
15   virtual void read_lock() noexcept = 0;
16   virtual void read_unlock() noexcept = 0;
17   virtual void synchronize() noexcept = 0;
18   virtual void retire(rcu_head *rhp, void (*cbf)(rcu_head *rhp)) = 0;
19   virtual void barrier() noexcept = 0;
20 };
```

Figure 2: RCU Domain Base Class

`-lurcu -lurcu-qsbr` to link to it. These applications must use `rcu_register_thread()` and `rcu_unregister_thread()` to announce the coming and going of each thread that is to execute `rcu_read_lock()` and `rcu_read_unlock()`. They must also use `rcu_quiescent_state()`, `rcu_thread_offline()`, and `rcu_thread_online()` to announce quiescent states to RCU.

The memory-barrier implementation is intended for applications that can announce threads (again using `rcu_register_thread()` and `rcu_unregister_thread()`), but for which announcing quiescent states is impractical. Such applications use `#include "urcu-mb.hpp"` and `-lurcu-mb` to select the memory-barrier implementation. Such applications will incur the overhead of a full memory barrier in each call to `rcu_read_lock()` and `rcu_read_unlock()`.

The signal-based implementation represents a midpoint between the QSBR and memory-barrier implementations. Like the memory-barrier implementation, applications must announce threads, but need not announce quiescent states. On the one hand, readers are almost as fast as in the QSBR implementation, but on the other applications must give up a signal to RCU, by default `SIGUSR1`. Such applications use `#include "urcu-signal.hpp"` and `-lurcu-signal` to select signal-based RCU.

So-called "bullet-proof RCU" avoids the need to announce either threads or quiescent states, and is therefore the best choice for use by libraries that might well be linked with RCU-oblivious applications. The penalty is that `rcu_read_lock()` incurs both a memory barrier and a test and `rcu_read_unlock()` incurs a memory barrier. Such applications or libraries use `#include urcu-bp.hpp` and `-lurcu-bp`.

```
 1  class rcu_guard {
 2  public:
 3    rcu_guard() noexcept
 4    {
 5      this->rd = nullptr;
 6      rcu_read_lock();
 7    }
 8
 9    explicit rcu_guard(rcu_domain *rd)
10    {
11      this->rd = rd;
12      rd->read_lock();
13    }
14
15    rcu_guard(const rcu_guard &) = delete;
16
17    rcu_guard&operator=(const rcu_guard &) = delete;
18
19    ~rcu_guard() noexcept
20    {
21      if (this->rd)
22        this->rd->read_unlock();
23      else
24        rcu_read_unlock();
25    }
26
27  private:
28    rcu_domain *rd;
29  };
```

Figure 3: RCU Guarded Readers

2.2 Run-Time Domain Selection

Figure 2 shows the abstract base class for runtime selection of RCU domains. Each domain creates a concrete subclass that implements its RCU APIs:

- Bullet-proof RCU: `class rcu_bp`
- Memory-barrier RCU: `class rcu_mb`
- QSBR RCU: `class rcu_qsbr`
- Signal-based RCU: `class rcu_signal`

Of course, additional implementations of RCU may be constructed by deriving from `rcu_domain` and/or by implementing the API shown in Figure 1.

3 RCU Guarded Readers

In some cases, it might be convenient to use a guard style for RCU readers, especially if the read-side critical section might be exited via exception. The `rcu_guard` class shown in Figure 3 may be used for this purpose. An argumentless constructor uses the API, or an `rcu_domain` class may be passed to the constructor to use the specified RCU implementation.

This is intended to be used in a manner similar to `std::lock_guard`.

4 RCU Callback Handling

The traditional C-language RCU callback uses address arithmetic to map from the `rcu_head` structure to the enclosing struct, for example, via the `container_of()` macro. Of course, this approach also works for C++, but this section first looks at some approaches that leverage C++ overloading and inheritance, which has the benefit of avoiding macros and providing better type safety. This will not be an either-or situation: Several of these approaches are likely to be generally useful.

However, the approach discussed in Section 4.1 is the preferred approach, and is compatible with the proposal for hazard pointers. The approaches in the other sections are presented for informational purposes: Section 4.2 uses a pointer to an enclosing type and Section 4.3 uses address arithmetic (illustrating the C-language approach used in the Linux kernel).

4.1 Derived-Type Approach

The `rcu_obj_base` class provides overloaded `retire()` methods, as shown in Figure 4. These methods take a deleter and an optional `rcu_domain` class instance. The deleter's operator() is invoked after a grace period. The deleter type defaults to `std::default_delete<T>`, but one could also use a custom functor class with an `operator()` that carries out teardown actions before freeing the object, or a raw function pointer type such as `void(*)(T*)`, or a lambda type. We recommend avoiding deleter types such as `std::function<void(T*)>` (and also any other type requiring memory allocation) because allocating memory on the free path can result in out-of-memory deadlocks.

If an `rcu_domain` is supplied, its `retire()` member function is used, otherwise the `call_rcu()` free function is used.

The next section provides a specialization that only permits `delete`, which allows omitting the deleter, thus saving a bit of memory.

4.2 Pointer To Enclosing Class

If complex inheritance networks make inheriting from an `rcu_head` derived type impractical, one alternative is to maintain a pointer to the enclosing class as shown in Figure 5. This `rcu_head_ptr` class is included as a member of the RCU-protected class. The `rcu_head_ptr` class's pointer must be initialized, for example, in the RCU-protected class's constructor.

If the RCU-protected class is `foo` and the name of the `rcu_head_ptr` member function is `rh`, then `foo1.rh.retire(my_cb)` would cause the function `my_cb()` to be invoked after the end of a subsequent grace period. As with the previous classes, omitting the deleter results in the object being passed to `delete` and an `rcu_domain` object may be specified.

```
template<typename T,
         typename D = default_delete<T>,
         bool E = is_empty<D>::value>
class rcu_obj_base: private rcu_head {
  D deleter;
public:
  static void trampoline(rcu_head *rhp)
  {
    auto rhdp = static_cast<rcu_obj_base *>(rhp);
    auto obj = static_cast<T *>(rhdp);
    rhdp->deleter(obj);
  }

  void retire(D d = {})
  {
    deleter = d;
    call_rcu(static_cast<rcu_head *>(this), trampoline);
  }

  void retire(rcu_domain &rd, D d = {})
  {
    deleter = d;
    rd.retire(static_cast<rcu_head *>(this), trampoline);
  }
};

template<typename T, typename D>
class rcu_obj_base<T,D,true>: private rcu_head {
public:
  static void trampoline(rcu_head *rhp)
  {
    auto rhdp = static_cast<rcu_obj_base *>(rhp);
    auto obj = static_cast<T *>(rhdp);
    D()(obj);
  }

  void retire(D d = {})
  {
    call_rcu(static_cast<rcu_head *>(this), trampoline);
  }

  void retire(rcu_domain &rd, D d = {})
  {
    rd.retire(static_cast<rcu_head *>(this), trampoline);
  }
};
```

Figure 4: RCU Callbacks: Derived-Type Approach

```
1   template<typename T>
2   class rcu_head_ptr: public rcu_head {
3   public:
4     rcu_head_ptr()
5     {
6       this->container_ptr = nullptr;
7     }
8
9     rcu_head_ptr(T *containing_class)
10    {
11      this->container_ptr = containing_class;
12    }
13
14    static void trampoline(rcu_head *rhp)
15    {
16      T *obj;
17      rcu_head_ptr<T> *rhdp;
18
19      rhdp = static_cast<rcu_head_ptr<T> *>(rhp);
20      obj = rhdp->container_ptr;
21      if (rhdp->callback_func)
22        rhdp->callback_func(obj);
23      else
24        delete obj;
25    }
26
27    void retire(void callback_func(T *obj) = nullptr)
28    {
29      this->callback_func = callback_func;
30      call_rcu(static_cast<rcu_head *>(this), trampoline);
31    }
32
33    void retire(class rcu_domain &rd,
34        void callback_func(T *obj) = nullptr)
35    {
36      this->callback_func = callback_func;
37      rd.retire(static_cast<rcu_head *>(this), trampoline);
38    }
39
40  private:
41    void (*callback_func)(T *obj);
42    T *container_ptr;
43  };
```

Figure 5: RCU Callbacks: Pointer

```
 1   template<typename T>
 2   class rcu_head_container_of {
 3   public:
 4     static void set_field(const struct rcu_head T::*rh_field)
 5     {
 6       T t;
 7       T *p = &t;
 8
 9       rh_offset = ((char *)&(p->*rh_field)) - (char *)p;
10     }
11
12     static T *enclosing_class(struct rcu_head *rhp)
13     {
14       return (T *)((char *)rhp - rh_offset);
15     }
16
17   private:
18     static inline size_t rh_offset;
19   };
20
21   template<typename T>
22   size_t rcu_head_container_of<T>::rh_offset;
```

Figure 6: RCU Callbacks: Address Arithmetic

```
1 void my_cb(struct std::rcu_head *rhp)
2 {
3   struct foo *fp;
4
5   fp = std::rcu_head_container_of<struct foo>::enclosing_class(rhp);
6   std::cout << "Callback fp->a: " << fp->a << "\n";
7 }
```

Figure 7: RCU Callbacks: Address Arithmetic in Callback

4.3 Address Arithmetic

Figure 6 shows an approach that can be used if memory is at a premium and the inheritance techniques cannot be used. The `set_field()` method sets the offset of the `rcu_head_container_of` member within the enclosing RCU-protected structure, and the `enclosing_class()` member function applies that offset to translate a pointer to the `rcu_head_container_of` member to the enclosing RCU-protected structure.

This address arithmetic must be carried out in the callback function, as shown in Figure 7.

5 Hazard Pointers and RCU: Which to Use?

Table 1 provides a rough summary of the relative advantages of reference counting, RCU, and hazard pointers. Advantages are marked in bold with green background, or with a blue background for strong advantages.

Although reference counting has normally had quite limited capabilities and been quite tricky to apply for general linked data-structure traversal, given a

	Reference Counting	Reference Counting with DCAS	RCU	Hazard Pointers
Unreclaimed objects	**Bounded**	**Bounded**	Unbounded	**Bounded**
Contention among readers	Can be very high	Can be very high	**No contention**	**No contention**
Traversal forward progress	Either blocking or lock-free with limited reclamation	**Lock free**	**Bounded population oblivious wait-free**	**Lock-free**
Reclamation forward progress *	Either blocking or lock-free with limited reclamation	**Lock free**	Blocking	**Bounded wait-free**
Traversal speed	Atomic read-modify-write updates	Atomic read-modify-write updates	**No or low overhead**	Store-load fence
Reference acquisition	**Unconditional**	**Unconditional**	**Unconditional**	Conditional
Automatic reclamation	**Yes**	**Yes**	No	No
Purpose of domains	N/A	N/A	Isolate long-latency readers	Limit contention, reduce space bounds, etc.

Table 1: Comparison of Deferred-Reclamation Mechanisms

* Does not include memory allocator, just the reclamation itself.

double-pointer-width compare-and-swap instruction, it can work quite well, as shown in the "Reference Counting with DCAS" column.

As a rough rule of thumb, for best performance and scalability, you should use RCU for read-intensive workloads and hazard pointers for workloads that have significant update rates. As another rough rule of thumb, a significant update rate has updates as part of more than 10% of its operations. Reference counting with DCAS is well-suited for small systems and/or low read-side contention, and particularly on systems that have limited thread-local-storage capabilities. Both RCU and reference counting with DCAS allow unconditional reference acquisition.

Specialized workloads will have other considerations. For example, small-memory multiprocessor systems might be best-served by hazard pointers, while the read-mostly data structures in real-time systems might be best-served by RCU.

The relationship between the Hazard Pointers proposal and this RCU proposal is as follows:

1. The `hazptr_obj_base` class is analogous to `rcu_obj_base`.
2. The `hazptr_domain` class is analogous to `rcu_domain`.
3. The private `hazptr_obj` class is analogous to the pre-existing `rcu_head` struct. Because this is a private hazard-pointers class, there is no need to have compatible names.
4. There is no RCU class analogous to `hazptr_rec` because RCU does not track (or need to track) references to individual RCU-protected objects.
5. There is no hazard pointers counterpart to the `rcu_guard` class. This is because hazard pointers does not have (or need) a counterpart to `rcu_read_lock()` and `rcu_read_unlock()`.

6 Summary

This paper demonstrates a way of creating C++ bindings for a C-language RCU implementation, which has been tested against the userspace RCU library. Specifically, this document proposes `rcu_domain` (Figure 2), `rcu_guard` (Figure 3), and `rcu_obj_base` (Figure 4). We believe that these bindings are also appropriate for the type-oblivious C++ RCU implementations that information-hiding considerations are likely to favor.

Acknowledgments

We owe thanks to Pedro Ramalhete for his review and comments. We are grateful to Jim Wasko for his support of this effort.

References

[1] Desnoyers, M. [RFC git tree] userspace RCU (urcu) for Linux. `http://liburcu.org`, February 2009.

[2] McKenney, P. E. Recent read-mostly research in 2015. `http://lwn.net/Articles/667593/`, December 2015.

[3] McKenney, P. E., Desnoyers, M., and Jiangshan, L. User-space RCU. `https://lwn.net/Articles/573424/`, November 2013.

[4] McKenney, P. E., Wong, M., and Michael, M. P0232r0: A concurrency toolkit for structured deferral or optimistic speculation. `http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2016/p0232r0.pdf`, February 2016.

Change Log

This paper first appeared as **P0461R0** in October of 2016. Revisions to this document are as follows:

- Convert to single-column mode. (November 16, 2016.)
- Change `call()` to `retire()` for hazard-pointer compatibility. (January 4, 2017.)
- Change `rcu_scoped_reader` to `rcu_guard` for compatibility with existing RAII mechanisms. (January 19, 2017.)
- Change `rcu_head_delete` to `rcu_obj_base` for compatibility with hazard pointers. (January 19, 2017.)
- Update to indicate preferred C++ RCU approach. (January 19, 2017.)
- Call out relationships between classes for RCU and for hazard pointers. (January 19, 2017.)
- Add constructors to `rcu_domain` to match those of `hazptr_domain`. (February 1, 2017.)
- Add `quiescent_state_needed()` member function to `rcu_domain` to allow code using RCU to complain if its requirements are not met, based on discussions with Geoffrey Romer and Andrew Hunter. (February 3, 2017.)
- Added references to related papers. (February 5, 2017.)

At this point, the paper was published as **P0461R1**.