



[rkyv](#) (*archive*) is a zero-copy deserialization framework

This book covers the motivation, architecture, and how to learn and understand rkyv, but won't go as in-depth as the documentation will. Don't be afraid to consult these resources as you read through.

Resources

Learning Materials

- The [rkyv discord](#) is a great place to get help with people using rkyv
- The [rkyv github](#) hosts the source and tracks progress

Documentation

- [rkyv](#), the core library
- [rkyv_dyn](#), which adds trait object support to rkyv
- [rkyv_typename](#), a type naming library

Benchmarks

- The [rust serialization benchmark](#) is a shootout of various rust serialization solutions. It includes special benchmarks for solutions like rkyv.

Sister Crates

- [bytecheck](#), which rkyv uses for validation

- [ptr_meta](#), which rkyv uses for pointer manipul.
- [rend](#), which rkyv uses for endian-agnostic feati

Motivation

First and foremost, the motivation behind rkyv is im achieves that goal can also lead to gains in memory the way.

Familiarity with other serialization frameworks ar works will help, but isn't necessary to understand

Most serialization frameworks like [serde](#) define an i basic types such as primitives, strings, and byte arra type into two stages: the frontend and the backend. breaks it down into the serializable types of the data: data model types and writes them using some data etc. This allows a clean separation between the seri: it is written to.

Serde describes [its data model](#) in the [serde book](#) eventually boils down to some combination of th

A major downside of traditional serialization is that time to read, parse, and reconstruct types from thei

In JSON for example, strings are encoded by surr quotes and escaping invalid characters inside of t

```
{ "line": "\"All's well that ends well\""
      ^ ^                ^ ^
```

numbers are turned into characters:

```
{ "pi": 3.1415926 }
      ^^^^^^^^^^^
```

and even field names, which could be *implicit* in r

```
{ "message_size": 334 }
      ^^^^^^^^^^^^^^^^^
```

All those characters are not only taking up space,


```
Example {
  quote: str::from_utf8(&buffer[0..30]).un
  a: &buffer[30..42],
  b: u64::from_le_bytes(&buffer[42..50]),
  c: char::from_u32(u32::from_le_bytes(&bu
}
```

And we can't borrow types like `u64` or `char` that since our buffer might not be properly aligned. We store those! Even though we borrowed 42 of the the last 12 and still had to parse through the buff

Partial zero-copy deserialization can considerably in speed up some deserialiation, but with some work v

Total zero-copy

rkyv implements total zero-copy deserialization, whi during deserialization and no work is done to deser structuring its encoded representation so that it is t representation of the source type.

This is more like if our buffer *was* an Example:

```
struct Example {
  quote: String,
  a: [u8; 12],
  b: u64,
  c: char,
}
```

And our buffer looked like this:

```
I don't know, I didn't listen.__QOFFQLENAA
^-----^-----^-----^-----^-----^-----
quote bytes                                pointer a
                                           and len
                                           ^-----
                                           Example
```

In this case, the bytes are padded to the correct a Example are laid out exactly the same as they wc deserialization code can be much simpler:

```
unsafe { &*buffer.as_ptr().add(32).cast()
```

This operation is almost zero work, and more importantly, it doesn't touch any data. No matter how much or how little data we have, it's just an offset and a cast to access our data.

This opens up blazingly-fast data loading and enables us to read data more quickly than traditional serialization.

Architecture

The core of rkyv is built around [relative pointers](#) and [Serialize](#), and [Deserialize](#). Each of these traits supports unsized types: [ArchiveUnsize](#), [Serialize](#)

A good way to think about it is that sized types are built on. That's not a fluke either, rkyv is built more complex abstractions out of lower-level machinery. It's not much different from what you normally

The system is built to be flexible and can be extended. For example, the `rkyv_dyn` crate adds support for trait objects by defining how they build up to allow trait objects to be

Relative pointers

Relative pointers are the bread and butter of total zero-copy, replacing the use of normal pointers. But why can't

Consider some zero-copy data on disc. Before we can load it into memory. But we can't control *where* in memory it gets located at a different address, and therefore the data is located at a different address.

One of the major reasons for this is actually *security*. For a program, it may run in a completely different range of memory called [address space layout randomization](#) and it introduces memory corruption vulnerabilities.

At most, we can only control the *alignment* of our data within those constraints.

This means that we can't store any pointers to that data. As soon as we reload the data, it might not be at the same address, pointers dangling, and would almost definitely result in errors. Other libraries like [abomonation](#) store some extra data to take the place of deserialization, but we can do better.

In order to perform that fixup step, [abomonation](#) uses *mutable backing*. This is okay for many use cases, but you won't be able to mutate our buffer. One example is [files](#).

While normal pointers hold an absolute address in memory, relative pointers hold an offset to an address. This changes how the pointer behaves.

Pointer	Self is moved	
Absolute	✓ Target is still at address	✗
Relative	✗ Relative distance has changed	✓ di

This is exactly the property we need to build data structures for deserialization. By using relative pointers, we can load data and still have valid pointers inside of it. Relative pointers work in memory either, so we can memory map entire files

data in a structured manner.

rkyv's implementation of relative pointers is the [Re1](#)

Archive

Types that implement `Archive` have an alternate `resolve` method for deserialization. The construction of archived types follows these rules:

1. Any dependencies of the type are serialized. For strings, the resolver would be the position of the string, for boxes it would be the boxed value and the position of the contained elements. Any bookkeeping from the resolver is held onto for later. This is the *serialize* method.
2. The resolver and original value are used to construct the original value. For strings the resolver would be the position of the archived elements. With the original value and the archived version can be constructed. This is the *resolve* method.

Resolvers

A good example of why resolvers are necessary is with a tuple of two strings:

```
let value = ("hello".to_string(), "world".to_string());
```

The archived tuple needs to have both of the strings:

```
0x0000      AA AA AA AA BB BB BB BB
0x0008      CC CC CC CC DD DD DD DD
```

A and B might be the length and pointer for the first string, C and D might be the length and pointer for the second string.

When archiving, we might be tempted to serialize and resolve the second one. But this might result in the second string ("world") being written before the first string ("hello"). Instead, we need to write the first string, then finish archiving both of them. The tuple doesn't need to finish archiving themselves, so they have to implement `Resolver`.

This way, the tuple can:

1. Archive the first string (save the resolver)
2. Archive the second string (save the resolver)
3. Resolve the first string with its resolver

4. Resolve the second string with its resolver

And we're guaranteed that the two strings are place need.

Serialize

Types implement `Serialize` separately from `Archive`. Some object, then `Archive` turns the value and that. Having a separate `Serialize` trait is necessary because one archived representation, you may have options in order to create one.

The `Serialize` trait is parameterized over the `ScratchSpace` mutable object that helps the type serialize itself. `char` don't *bound* their serializer type because they can use any kind of serializer. More complex types like `Vec` that implements `Serializer`, and even more complex types require a serializer that additionally implements `ScratchSpace`.

Unlike `Serialize`, `Archive` doesn't parameterize `ScratchSpace`. It shouldn't matter what serializer a resolver was made with.

Serializer

rkyv provides serializers that provide all the functionality of the library types, as well as serializers that combine other serializers. All of the components' capabilities.

The [provided serializers](#) offer a wide range of strategies. Some cases will be best suited by `AllocSerializer`.

Many types require *scratch space* to serialize. This space they can use temporarily and return when they're done. They request scratch space to store the resolvers for it. Requesting scratch space from the serializer is reused many times, which reduces the number of allocations performed while serializing.

Deserialize

Similarly to `Serialize`, `Deserialize` parameterize converts a type from its archived form back to its original form. `Deserialize` occurs in a single step and doesn't have any intermediate steps.

`Deserialize` also parameterizes over the type that the archived data is being deserialized into. This allows the same archived type to deserialize into different types depending on what's being asked for. This helps with generic abstractions, but might require you to annotate the type.

This provides a more or less a traditional deserialization that is sped up somewhat by having very compatible representations in memory and performance penalties of traditional deserialization. You need to load what you need before you use it. Deserialization is required as long as you can do so through the archived version.

Even the highest-performance serialization framework has a speed limit because of the amount of memory allocated and performed.

A good use for `Deserialize` is deserializing portions of the archived data to locate some subobject, then deserializing the archive as a whole. This granular approach provides fine-grained deserialization as well as traditional deserialization.

Deserializer

Deserializers, like serializers, provide capabilities to deserialize types. Some types don't bound their deserializers, but some like `Vec` in order to deserialize memory properly.

Alignment

The *alignment* of a type restricts where it can be loaded and stores. Because rkyv creates references to bytes, it has to ensure that the references it creates

In order to perform arithmetic and logical operations, the CPU has to *load* that data from memory into its registers. However, there is a limitation on how the CPU can access that data: it has to access it at *word boundaries*. These words are the natural size of the CPU. The word size is 4 bytes for 32-bit machines and 8 bytes for 64-bit machines. If we had some data laid out like this:

```
0  4  8  C
AAAABBBBCCCCDDDD
```

On a 32-bit CPU, accesses could occur at any address. For example, one could access `A` by loading 4 bytes from address 0, `B` by loading 4 bytes from address 4, and so on. This works great for aligned data. *Unaligned* data can throw a wrench into this.

```
0  4  8  C
..AAAABBBBCCCC
```

Now if we want to load `A` into memory, we have to:

1. Load 4 bytes from address 0
2. Throw away the first two bytes
3. Load 4 bytes from address 4
4. Throw away the last two bytes
5. Combine our four bytes together

That forces us to do twice as many loads *and* per byte we can have a real impact on our performance across the board. Our data has to be properly aligned.

rkyv provides two main utilities for aligning byte buffers:

- [AlignedVec](#) is a drop-in replacement for `Vec<T>`
- [AlignedBytes](#) is a wrapper around `[u8; N]`

Both of these types align the bytes inside to 16-byte boundaries. For almost all use cases, but if your particular situation

then you may need to manually align your bytes.

In practice

rkyv has a very basic unaligned data check built in that also [validate](#) your data, then it will always make sure

Common pitfalls

In some cases, your archived data may be prefixed with extra data in the buffer. If this extra data misaligns the following data, you can have the prefixing data removed before accessing it.

In other cases, your archived data may not be tight to the [archived_root](#) if you rely on the end of the buffer being tight. You may miscalculate the positions of the contained values if

Format

Types which derive `Archive` generate an archived version of themselves.

- Member types are replaced with their archived counterparts.
- Enums have `#[repr(N)]` where N is `u8`, `u16`, `u32`, or the smallest possible type that can represent all of the enum variants.

For example, a struct like:

```
struct Example {  
    a: u32,  
    b: String,  
    c: Box<(u32, String)>,  
}
```

Would have the archived counterpart:

```
struct ArchivedExample {  
    a: u32,  
    b: ArchivedString,  
    c: ArchivedBox<(u32, ArchivedString)>,  
}
```


With the `strict` feature, these structs are additionally guaranteed portability and stability.

In most cases, the `strict` feature will not be necessary for the efficiency of archived types. Make sure you understand the implications of this feature by reading the crate documentation for details on the `strict` feature.

rkyv provides `Archive` implementations for common types. In general they follow the same format as derived implementations. For example, `ArchivedString` performs a smart allocation to reduce memory use.

Object order

rkyv lays out subobjects in depth-first order from the root object. The root object is stored at the end of the buffer, not at the beginning of the tree:



would be laid out like this in the buffer:

```
b d e c a
```

from this serialization order:

```
a -> b
a -> c -> d
a -> c -> e
a -> c
a
```

This deterministic layout means that you don't need an object in most cases. As long as your buffer ends right, you can use `archived_root` with your buffer.

Wrapper types

Wrapper types make it easy to customize the way that rkyv makes it easier to adapt rkyv to existing data models by making deserializing idiomatic for even complicated types.

Annotating a field with `#[with(...)]` will *wrap* that struct is serialized or deserialized. There's no performance cost but doing more or less work during serialization and deserialization. This excerpt is from the documentation:

```
#[derive(Archive, Deserialize, Serialize)]
struct Example {
    #[with(Incremented)]
    a: i32,
    // Another i32 field, but not incremented
    b: i32,
}
```

The `Incremented` wrapper is wrapping `a`, and the `Incremented` wrapper increments `a` in its archived form.

With

The core type behind wrappers is `With`. This struct wraps a type and provides another name for the type inside of it. rkyv uses `With` for wrapping types during serializing and deserializing, and when you write your own wrappers you use `With` as well.

See `ArchiveWith` for an example of how to write your own wrapper.

Shared Pointers

The implementation details of shared pointers may vary. Specifically, the rules surrounding how and when shared and pooled may affect how you choose to use them.

Serialization

Shared pointers (`Rc` and `Arc`) are serialized when they are dropped, and the data address is reused when subsequent data is serialized. This means that you can expect shared pointers to point to the same data when archived, even if they are unsized to different types.

Weak pointers (`Rc::Weak` and `Arc::Weak`) have separate serialization rules. They're encountered. The serialization process upgrades them to shared pointers. Otherwise, it serializes them like shared pointers. Otherwise, it serializes them like shared pointers.

Deserialization

Similarly, shared pointers are deserialized on the first use. Weak pointers do a similar upgrade attempt when they are first used.

Serializers and Deserializers

The serializers for shared pointers hold the location of the data. It's safe to serialize shared pointers to an archive across different processes if you use the same serializer for each one. Using a different serializer but may end up duplicating the shared data.

The deserializers for shared pointers hold a shared pointer to the data and will hold them in memory until the deserializer is dropped. They serialize only weak pointers to some shared data, then they will point to nothing as soon as the shared data is dropped.

Unsize Types

rkyv supports unsize types out of the box and ship common unsize types (`str s` and slices). Trait objects `rkyv_dyn`, see "[Trait Objects](#)" for more details.

Metadata

The core concept that enables unsize types is metadata. Metadata allows for different sizes, in contrast with languages like C and C++ where the size is fixed. This is important for the concept of sizing, which is implemented in rust's [Sized](#) trait.

Pointers are composed of two pieces: a data address and metadata. The data address is what most people think of when they think of a pointer. The metadata for a pointer is some extra information stored safely with the data at the pointed location. It can be used to determine the size of `Sized` types. Pointers with no extra metadata are called "narrow" pointers *with* metadata are sometimes called "wide" pointers.

rkyv uses the [ptr_meta](#) crate to perform these operations. These operations may be incorporated as [part of the standard library](#).

Fundamentally, the metadata of a pointer exists to provide the necessary information to safely access, drop, and deallocate slices. For slices, the metadata carries the length of the slice, for function pointers (vtable) the metadata carries the function pointer, and for custom unsize types the metadata carries the single trailing unsize member.

Archived Metadata

For unsize types, the metadata for a type is archived into the data. This mirrors how rust works internally for slices and other exotic use cases. This does complicate things, but the metadata archiving process will end up as just `returning ()`.

This is definitely one of the more complicated parts of the system, and it is difficult to wrap your head around. Reading the code may help you understand how the system works.

Trait Objects

Trait object serialization is supported through the `rkyv_dyn` crate, which is maintained as part of `rkyv`, but is separate from the `rkyv` crate. This section will discuss the implementations to be used instead. This section will discuss the use of `rkyv_dyn` and how to use it effectively.

`rkyv_dyn` may not work in some exotic environments, such as those used to register trait objects. If you want these capabilities in your environment, feel free to file an issue or drop a pull request through.

Core traits

The new traits introduced by `rkyv_dyn` are `SerializeDyn` and `DeserializeDyn`. These are effectively type-erased versions of `Serialize` and `Deserialize` that the traits are object-safe. Likewise, it introduces `DynSerializer` and `DynDeserializer`, which are basic functionality required to serialize most types, but some custom types require.

`DynSerializer` implements the `Serializer` and `Deserializer` traits, but may not be suitable for all use cases. If you need more functionality, you can talk it through in the discord.

Architecture

It is highly recommended to use the provided `archive_dyn` trait and set everything up correctly.

Using `archive_dyn` on a trait definition creates an `archive_dyn` trait and `SerializeDyn`. This "shim" trait is blank, so you should implement your trait and `SerializeDyn`, so you should use the `archive_dyn` trait to use it.

The shim trait should be used everywhere that you want to serialize. By default, it will be named "SerializeDyn".

approach that similar libraries take is directly adding your trait. While more ergonomic, this approach does not work for the trait on types that cannot or should not implement the trait. The `archive_dyn` approach was favored for `rkyv_dyn`.

When the shim trait is serialized, it stores the type hash and metadata so it can get the correct vtable for it when deserializing. For vtables for implementing types must be known ahead of time. `archive_dyn` for the second time.

Using `archive_dyn` on a trait implementation registers the implementation with a global lookup, allowing it to be found. The process can be slow, the `vtable_cache` feature allows only the first time, then cached locally for future lookups. Alternate implementations may take a different approach for benefits and tradeoffs.

Validation

Validation can be enabled with the `validation` feature of the `bytecheck` crate to perform archive validation, and detect and malicious data.

To validate an archive, you first have to derive `CheckBytes` and `Serialize`:

```
use rkyv::{Archive, Deserialize, Serialize};

#[derive(Archive, Deserialize, Serialize)]
#[archive(check_bytes)]
pub struct Example {
    a: i32,
    b: String,
    c: Vec<bool>,
}
```

The `#[archive(check_bytes)]` attribute derives `CheckBytes` for the struct. Finally, you can use `check_archived_root` to check an archived value if it was successful:

```
use rkyv::check_archived_root;

let archived_example = check_archived_root::<Example>(<...>);
```

More examples of how to enable and perform validation are available in the crate's `validation` module.

The validation context

When checking an archive, a validation context is created. The `DefaultValidator` is the default context that will work for most archived types. If you need custom validation logic, you may need to augment the capabilities of the `DefaultValidator` by implementing `check_archived_root_with_context` and use `check_archived_root_with_context`.

The `DefaultValidator` supports all builtin rkyv types. It also supports `alloc` types, whether you have the `alloc` feature enabled or not.

Bounds checking and subtree ra

All pointers are checked to make sure that they:

- point inside the archive
- are properly aligned
- and have enough space afterward to hold the

However, this alone is not enough to secure against sharing violations, so rkyv uses a system to verify the ownership model.

Archive validation uses a memory model where all subobjects are kept in memory. This is called a *subtree range*. When validating, rkyv keeps track of where subobjects are allowed to be located. A range from the beginning with `push_prefix_subtree_range` and `push_suffix_subtree_range`. After pushing a subobject, its pointer can be checked by calling their `CheckBytes` implementation. After checked, `pop_prefix_subtree_range` and `pop_suffix_subtree_range` restore the original range with the checked section.

Validation and Shared Pointers

While validating shared pointers is supported, some use cases can prevent malicious data from validating:

Shared pointers that point to the same object will fail validation for different pointer types. This can cause issues if you have a shared pointer that is both an array pointer and a slice pointer. Since shared pointers are not allowed to point to the same value as a concrete type (e.g. `dyn Any`).

rkyv still supports these use cases, but it's not possible to validate them with these use cases. Alternative validation solutions like pointer hashes may be a better approach in these cases.

Feature Comparison

This is a best-effort feature comparison between rkyv by no means completely comprehensive, and pull requests are welcomed.

Feature matrix

Feature	rkyv
Open type system	yes
Scalars	yes
Tables	no*
Schema evolution	no*
Zero-copy	yes
Random-access reads	yes
Validation	upfront*
Reflection	no*
Object order	bottom-up
Schema language	derive
Usable as mutable state	yes
Padding takes space on wire?	yes*
Unset fields take space on wire?	yes
Pointers take space on wire?	yes
Cross-language	no
Hash maps and B-trees	yes
Shared pointers	yes

* rkyv's open type system allows extension types that provide

Open type system

One of rkyv's primary features is that its type system allows users to write custom types and control their properties very easily. This provides a solid foundation to build many other features on top of, and is already a fundamental part of how rkyv works.

Unsize types

Even though they're part of the main library, unsize serialization functionality. Types like `Box` and `Rc/Ar` entry points for unsize types into the size system.

Trait objects

Trait objects are further built on top of unsize type objects easy and safe.

FAQ

Because it's so different from traditional serialization, we have a lot of questions about rkyv. This is meant to serve as a collection of answers.

How is rkyv zero-copy? It definitely doesn't copy data into memory.

Traditional serialization works in two steps:

1. Read the data from disk into a buffer (maybe in memory).
2. Process the data in the buffer into the deserialized form.

The copy happens when the data in the buffer ends up being written to disk. Zero-copy deserialization doesn't deserialize the buffer into memory, so it avoids this copy.

You can actually even avoid reading the data from disk in some environments by using memory mapping.

How does rkyv handle endianness?

rkyv supports three endiannesses: native, little, and big. It doesn't support little or big, but removes the abstraction layer to make it easier to use with native types.

You can enable specific endiannesses with the `little` or `big` macros.

Is rkyv cross-platform?

Yes, but rkyv has been tested mostly on x86 machines. There are some things that need to get fixed for other architectures.

Can I use this in embedded and microcontroller environments?

Yes, disable the `std` feature for `no_std`. You can also disable all memory allocation capabilities.

Safety

Isn't this very unsafe if you accept

Yes, *but* you can still access untrusted data if you validate it. It's an extra step, but it's usually still less than the cost of the `std` format. `rkyv` has proven to round-trip faster than `bincode`.

Doesn't that mean I always have

No. There are many other ways you can verify your data and signed buffers.

Isn't it kind of deceptive to say `unsafe` require validation?

The fastest path to access archived data is marked as `unsafe`, it means that it's only safe to call if you can

The value must be archived at the given position

As long as you can (reasonably) guarantee that, then every archive needs to be validated, and you can use `rkyv` to guarantee data integrity and security.

Even if you do need to always validate your data before using it, it's still faster than deserializing with other high-performance formats, even though it's not by the same margins.

Contributors

Thanks to all the contributors who have helped docu

- David Koloski ([djkoloski](#))

If you feel you're missing from this list, feel free to a