



EMBRACING MODERN C++ SAFELY



JOHN LAKOS | VITTORIO ROMEO | ROSTISLAV KHLBNIKOV | ALISDAIR MEREDITH

FREE SAMPLE CHAPTER

SHARE WITH OTHERS



Embracing Modern C++ *Safely*

Embracing Modern C++ *Safely*

John Lakos

Vittorio Romeo

Rostislav Khlebnikov

Alisdair Meredith

◆ Addison-Wesley

Boston • Columbus • New York • San Francisco • Amsterdam • Cape Town
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

Cover image: Unconventional/Shutterstock
Pages 109, 130: cocktail glass, Laura Humpfer/OpenMojis

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at corpsales@pearsoned.com or (800) 382-3419.

For government sales inquiries, please contact governmentsales@pearsoned.com.

For questions about sales outside the U.S., please contact intlcs@pearson.com.

Visit us on the Web: informit.com/aw

Library of Congress Control Number: 2021947542

Copyright © 2022 Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit www.pearson.com/permissions.

ISBN-13: 978-0-13-738035-0

ISBN-10: 0-13-738035-6

ScoutAutomatedPrintCode

Pearson's Commitment to Diversity, Equity, and Inclusion

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create a better life for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where:

- Everyone has an equitable and lifelong opportunity to succeed through learning.
- Our educational products and services are inclusive and represent the rich diversity of learners.
- Our educational content accurately reflects the histories and experiences of the learners we serve.
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview).

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

- Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

To my darling wife, Elyse, who I love dearly, always have, and forever will:

“‘When I use a word,’ Humpty Dumpty said in rather a scornful tone,
‘it means just what I choose it to mean—neither more nor less.’”

— Lewis Carroll, *Through the Looking-Glass*

JSL

To my aunts and my dad,
who have always supported me
in every aspect of my life.

VR

To Elena and my parents.

RK

To the late David and Mary Meredith,
loving parents who encouraged me in everything that I did
and would have been so proud to see their son finally in print.

AM

Contents

Foreword by Shawn Edwards	xiii
Foreword by Andrei Alexandrescu	xv
Acknowledgments	xix
About the Authors	xxv
Chapter 0 Introduction	1
What Makes This Book Different	1
Scope for the First Edition	2
The <i>EMC++S</i> Guiding Principles	3
What Do We Mean by <i>Safely</i> ?	4
A <i>Safe</i> Feature	5
A <i>Conditionally Safe</i> Feature	5
An <i>Unsafe</i> Feature	6
Modern C++ Feature Catalog	6
How to Use This Book	8
Chapter 1 Safe Features	11
1.1 C++11	11
Attribute Syntax	12
Consecutive <code>>s</code>	21
<code>decltype</code>	25
Defaulted Functions	33
Delegating Ctors	46
Deleted Functions	53
<code>explicit</code> Operators	61
Function <code>static</code> '11	68
Local Types '11	83
<code>long long</code>	89
<code>noreturn</code>	95
Generalized Attribute Support	12
Consecutive Right-Angle Brackets	21
Operator for Extracting Expression Types	25
Using <code>= default</code> for Special Member Functions	33
Constructors Calling Other Constructors	46
Using <code>= delete</code> for Arbitrary Functions	53
Explicit Conversion Operators	61
Thread-Safe Function-Scope <code>static</code> Variables	68
Local/Unnamed Types as Template Arguments	83
The <code>long long</code> (≥ 64 bits) Integral Type	89
The <code>[[noreturn]]</code> Attribute	95

Contents

	nullptr	The Null-Pointer-Literal Keyword	99
	override	The override Member-Function Specifier	104
	Raw String Literals	Syntax for Unprocessed String Contents	108
	static_assert	Compile-Time Assertions	115
	Trailing Return	Trailing Function Return Types	124
	Unicode Literals	Unicode String Literals	129
	using Aliases	Type/Template Aliases (Extended typedef)	133
1.2	C++14		138
	Aggregate Init '14	Aggregates Having Default Member Initializers	138
	Binary Literals	Binary Literals: The 0b Prefix	142
	deprecated	The <code>[[deprecated]]</code> Attribute	147
	Digit Separators	The Digit Separator (')	152
	Variable Templates	Templated Variable Declarations/Definitions	157
	Chapter 2 Conditionally Safe Features		167
2.1	C++11		167
	alignas	The alignas Specifier	168
	alignof	The alignof Operator	184
	auto Variables	Variables of Automatically Deduced Type	195
	Braced Init	Braced-Initialization Syntax: <code>{}</code>	215
	constexpr Functions	Compile-Time Invocable Functions	257
	constexpr Variables	Compile-Time Accessible Variables	302
	Default Member Init	Default class/union Member Initializers	318
	enum class	Strongly Typed, Scoped Enumerations	332
	extern template	Explicit-Instantiation Declarations	353
	Forwarding References	Forwarding References (T&&)	377
	Generalized PODs '11	Trivial and Standard-Layout Types	401
	Inheriting Ctors	Inheriting Base-Class Constructors	535
	initializer_list	List Initialization: <code>std::initializer_list<T></code>	553
	Lambdas	Anonymous Function Objects (Closures)	573
	noexcept Operator	Asking if an Expression Cannot throw	615
	Opaque enums	Opaque Enumeration Declarations	660
	Range for	Range-Based for Loops	679
	<i>Rvalue</i> References	Move Semantics and <i>Rvalue</i> References (&&)	710
	Underlying Type '11	Explicit Enumeration Underlying Type	829
	User-Defined Literals	User-Defined Literal Operators	835
	Variadic Templates	Variable-Argument-Count Templates	873
2.2	C++14		958
	constexpr Functions '14	Relaxed Restrictions on constexpr Functions	959
	Generic Lambdas	Lambdas Having a Templated Call Operator	968
	Lambda Captures	Lambda-Capture Expressions	986

Contents

Chapter 3 Unsafe Features		997
3.1 C++11		997
carries_dependency	The [[<code>carries_dependency</code>]] Attribute	998
final	Prohibiting Overriding and Derivation	1007
friend '11	Extended <code>friend</code> Declarations	1031
inline namespace	Transparently Nested Namespaces	1055
noexcept Specifier	The <code>noexcept</code> Function Specification	1085
Ref-Qualifiers	Reference-Qualified Member Functions	1153
union '11	Unions Having Non-Trivial Members	1174
3.2 C++14		1182
auto Return	Function (<code>auto</code>) Return-Type Deduction	1182
decltype(auto)	Deducing Types Using <code>decltype</code> Semantics	1205
Afterword: Looking Back and Looking Forward		1215
Glossary		1217
Bibliography		1281
Index		1305

Foreword by Shawn Edwards

I have been writing programs in C++ professionally for more than 25 years, even before it was standardized. The C++ language, in its mission to deliver zero overhead and maximum performance, necessarily provides few guardrails; syntax and type safety go only so far. Using C++ features in unsound ways and creating spectacular failures was always easy. But because the language was relatively stable, good developers — over time — learned how to write reliable C++ software.

The first standardized version, C++98, formalized what many already knew about the language. The second version of the Standard, C++03, included some small corrections and enhancements but did not fundamentally alter the way programs were written. What it meant to know how to program in C++, however, changed drastically with the publication of the C++11 Standard. For the first time in many years, the ISO C++ Standards Committee (WG21) added significant new functionality and removed functionality as well. For example, `noexcept` and `std::unique_ptr` were in, and the days of using dynamic exception specifications and `std::auto_ptr` were numbered.

At the same time, the Standards Committee announced its unprecedented commitment to deliver a new version of the C++ Standard every three years! For a large software organization, like Bloomberg, whose software asset lifetimes are measured in decades, relying on a language standard that is updated with such frequency is especially problematic. Bloomberg has been reliably and accurately providing indispensable information to the professional financial community for nearly 40 years, with services that span such diverse needs as financial analytics, trading solutions, and real-time market data.

To support our global business, Bloomberg has developed high-performance software systems that operate at scale and, for more than two decades now, has written them primarily in C++. As you can imagine, incorporating and validating new tool chains that underpin our company's entire code base is no simple task. Each update risks the stability of the very products upon which our customers depend.

Modern C++ has much to offer — both good and bad. Many of its newer features offer the prospect of improving performance, expressiveness, maintainability, and so on. On the other hand, many of these same features come with potential pitfalls, some of which are obvious, and others less so. With each new release of C++, now every three years, the language gets bigger, and the opportunities for misusing a feature, through lack of knowledge and experience, grow ever larger as well.

Foreword by Shawn Edwards

Using new features of an already sophisticated programming language such as C++, with which many developers might not be fully familiar, introduces its own category of risk. Less-seasoned engineers might unwittingly introduce new features into a mature code base where they could add manifestly negative value in that context. As ever, only time and experience can provide proof as to whether and under what conditions using a new C++ language feature would be prudent. We, as senior developers, team leads, and technical managers of a leading financial technology company, bear responsibility for protecting our Software Capital asset from undue risk.

We cannot justify the instability of rewriting all of our software every time a new version of the language appears, nor can we leave it in perpetual stasis and forgo the important benefits modern C++ has to offer. So we move forward but with expertise and caution, adopting features only after we fully understand them. Bloomberg is committed to extracting all of the benefit that it can from modern C++, but as a company, we must do so *safely*.

Bloomberg sponsored this book, *Embracing Modern C++ Safely*, because we felt that, despite all of the books, conferences, blogs, etc., that covered C++11/14 features, we needed to look at each feature from the point of view of how to apply it *safely* as well as effectively in the context of a large, mature corpus of production code. Therefore, this book provides detailed explanations of each C++11/14 language feature, examples of its effective use, and pitfalls to avoid. Moreover, this book could only have been written *now*, after years of gathering real-world experience. What's more, we knew that we had the right people — some of the best engineers and authors in the world — to write it.

As promised, the C++ Standards Committee has been sticking to its schedule, sometimes in the face of major world events, and two additional versions of the Standard, C++17 and C++20, have been published. As the community gains experience using the new features provided in those standards, I expect that future editions of this book will offer similar guidance and critique.

If you've been writing programs in C++ for more than a decade, you've undoubtedly noticed that being an accomplished C++ programmer is a different challenge than it used to be. This book will help you navigate the modern C++ landscape so that you too can feel confident in applying C++11/14 in ways that truly add value without undue risk to your organization's precious Software Capital investment.

— Shawn Edwards
Chief Technology Officer, Bloomberg LP
August 2021

Foreword by Andrei Alexandrescu

Do you like version control systems — Git, Perforce, Mercurial, and such? I love them! I have no idea how any of today’s complex software systems could have ever been built without using version control.

One beneficial artifact of version control software is the *diff view*, that quintessential side-by-side view of a change of a large system as a differential from the previous, known version of the system. The diff view is often the best way to review code, to assess complexity of a feature, to find a bug, and, most importantly, to get familiar with a new system. I pore over diff views almost every working day, perusing them for one or more of their advantages. The diff view is proof that we *can* actually have the proverbial nice things.

The novel concept of this book is a diff view between classic C++ — i.e. C++03, the baseline — and modern C++ — i.e., post-2011 C++, with its added features. A diff view of programming language features! Now that’s a cool idea with interesting implications.

Embracing Modern C++ Safely addresses a large category of programmers: those who work daily on complex, long-lived C++ systems and who are familiar with C++03 because said systems were written with that technology. Classes. Inheritance. Polymorphism. Templates. The STL. Yep, they know these notions well and work with them every day in complex problem domains. Rehashing those classic features is unnecessary. But some programmers are perhaps less comfortable with the cornucopia of new features standardized every three years, starting with C++11. They have no time to spend on tracking what the C++ Standards Committee is doing. Every hour spent learning new C++ features is an hour not spent on core systems functionality, so that snazzy new feature better be worth it. *Embracing Modern C++ Safely* is cleverly optimized to maximize the ratio of usefulness in production to time spent learning.

Pedagogically, this book achieves an almost impossible challenge: a *partial* diff (to allow this nerd a mathematically motivated metaphor) for each individual new feature added to C++ after 2003. What do I mean by that? When a book teaches language features, cross talk is inevitable: While discussing any one given feature, most other features interfere by necessity. As Scott Meyers once told me, “When you learn a language, all features come at you in parallel.” The authors *modularized* the teaching of each new feature, so if you want to read about, say, generic lambdas, you get to read about generic lambdas with minimal interference from any other new language feature. When necessary, the interaction between the feature being discussed and others is narrowly specified, documented, and cross-referenced. The result is a fractally self-consistent book that can be read cover to cover or chunked by themes, interconnected features, or individual topics.

Foreword by Andrei Alexandrescu

Chapters 1, 2, and 3 mimic a sort of reverse *Divine Comedy*, whereby, as you may recall, the poet Dante is led by trusted guides through Hell (*Inferno*), Purgatory (*Purgatorio*), and Heaven (*Paradiso*). The respective chapters help you navigate from *Safe* to *Conditionally Safe* to *Unsafe* features of Modern C++.

Safe features (Chapter 1) will clearly, definitely, pound-the-table improve your code wherever you use them. Acquiring and applying the teachings of Chapter 1 is the fastest way for a team to start leveraging Modern C++ in production. `override`? Enjoy. Digit separators? Have at 'em. Explicit conversion operators? Knock yourself out. Such features are recommended fully and without reservation. Chapter 2 discusses *conditionally safe* features, those that are good for you but come with caveats. Initializer lists? Let's talk. Range `for`? Couple of things to be mindful of. Rvalue references? Long discussion; grab a coffee. And last but not least, *unsafe* features are those that can be challenging and require skill and utmost attention in usage. Their use should be confined as much as possible and wrapped under interfaces. Standard-layout types? Way trickier than it may seem. The `noexcept` specifier? Careful, you're on your own. Inline namespaces? At best, don't. Extensive details, examples, and discussions are available for every single feature added after C++03.

The authors use “unsafe” in a tongue-in-cheek manner here. Nothing taught in this book is unsafe in the traditional computer science sense; instead, think of the casual meaning of “safety” when used, say, in a hardware store. What's the safety of various tools for someone just starting to use them? A screwdriver is safe; a power drill is conditionally safe; and a welding machine is unsafe.

You may be concerned, thinking, “That sounds authoritative. What is the basis of such a ranking?” In fact, *Embracing Modern C++ Safely* is emphatically not authoritative but *objective* and based on the vast community of experience that the authors collected and curated. They intentionally, sometimes painfully, withhold their opinions. The “Use Cases” and “Potential Pitfalls” sections, taken from production code, are empirical evidence as much as instructive examples to learn from.

Only the passage of time can distill the programming community's practical experience with each feature and how well it fared, which is why this book discusses features added up to C++14, even though C++20 is already out. Using features for years can replace passionate debate on language design ideas with cold, hard experience, which guides this book's remarkably clinical approach. In the words of John Lakos, “We explain the degrees of burns you could get if you put your hand on a hot stove, but we won't tell you not to do it.” The result is a refreshingly nonideological read, no more partisan than a book on experimental physics. Consistently avoiding injecting one's own ego and opinion in an analysis takes paradoxically a lot of work. *Ars est celare artem*, the Latin proverb goes in typical brief, cryptic, and slightly confusing manner. (Is Latin the APL of natural languages?) That literally translates to “the art is to conceal the art,” but the profound meaning is closer to “good art is not emphatically artsy.” Good artists don't leave fingerprints all over their work. In a very concrete way, that has been a design goal of *Embracing Modern C++ Safely*, for you won't find in it any opinion, pontification, or even gratuitously flowery language. (Fierce debates

Foreword by Andrei Alexandrescu

occurred about the perfect, most spartan choice of words in one paragraph or another.) This polished clarity will, I'm sure, shine through to any reader.

That Extra Oomph

“The only kind of writing is rewriting,” goes the famous quote. That is doubly true for technical books. The strength of a textbook stands in the willingness of its authors to redo their work and in the depth and breadth of its review team feeding the revision process. And rewriting is not easy! Have you ever written some code and then resisted reviews because you fell in love with it? Multiply that by 1024 and you'll know how book authors feel about rewriting passages they've already poured their souls into. You really need to be committed to quality to keep heart during such a trial.

The authors' insistence on quality brings to mind what I like most about this book, which is also the most difficult to explain. I call it *the extra oomph*.

I noticed something about great work — be it in engineering, art, sports, or any other challenging human endeavor. Almost always, great work is the result of talented people making an extra effort that goes beyond what one might consider reasonable. In appreciating such work, we implicitly acknowledge great capability combined with commensurately great effort in realizing it. Good work can be done glibly; great work cannot.

Through an odd turn of events I ended up getting quite involved with this book — first, for one review. And then another, and another, for a total of four thorough passes through the entire book. The quest for perfection is as contagious as the resignation to sloppiness and incomparably more fun. (“Destroy!” John Lakos pithily emailed me along with each new revision. My often caustic reviews motivated him like nothing else.) Other reviewers — C++ Standards Committee denizens, industry C++ experts, C++03 experts with no prior exposure to C++1x, software-architecture experts, multithreading experts, process experts, even LaTeX experts — have done the same, with the net result that each sentence you'll read has been critically considered dozens of times and probably rewritten a few. For my part, I got so enthused with the project and with the authors' uncompromising take on quality, that I ended up writing a full feature for the book. (Any mistake in Section 2.1. “Variadic Templates” is my fault.) This book project has been a lot of work, more than I might have reasonably expected, which is everything I'd hoped for. I thought I've gotten too old to still pull all-nighters; apparently I was wrong.

Having been thusly involved, I can tell: This book does have that extra oomph baked into it. The talk is being walked, there's no fluff, and the code examples are precise and eloquent. I think *Embracing Modern C++ Safely* is Great Work. Aside from learning from this book, I hope you derive from it inspiration to add more oomph into your own work. I know I did.

— Andrei Alexandrescu
May 2021

Acknowledgments

Embracing Modern C++ Safely is the work of the C++ community as a whole, not just the authors. This book comprises knowledge drawn from the depths of language design to the boundaries of sound software development. Those who are expert at one end of that language-design to application-development spectrum might be relatively unfamiliar with the other. Although we, the four authors named on the front of this book, are each professional senior software engineers, our combined knowledge did not initially span everything presented here, and we relied on many of our colleagues — from our fellow developers at Bloomberg to the Core Working Group of the C++ Standards Committee to Bjarne Stroustrup himself — to fill in holes in our understanding and to correct misconceptions we held.

Everyone on Bloomberg’s BDE team, founded in December 2001, contributed directly, in one way or another, to the publication of this book: Parsa Amini, Joshua Berne, Harry Bott, Steven Breitstein, Nathan Burgers, Bill Chapman, Attila Feher, Mungo Gill, Rostislav Khlebnikov, Jeffrey Mendelsohn, Alisdair Meredith, Hyman Rosen, and the BDE team’s second manager (since April 2019), Mike Verschell.

Nina Ranns, ISO C++ Standards Committee secretary and ISO C++ Foundation director, was our principal researcher and provided a window into the depths of the C++ core language standard. We relied on her to get to *the truth*: With a release coming every three years and defect reports retroactive to previous standards, *the truth* is a contextual, ephemeral, and elusive beast. Nonetheless, Nina provided us with clarity about what was in effect when and thoroughly reviewed each and every core-language-intensive feature in this book; see Section 2.1. “**constexpr** Functions” on page 257, Section 2.1. “Generalized PODs ’11” on page 401, Section 2.1. “*Rvalue* References” on page 710, and Section 3.1. “**noexcept** Specifier” on page 1085, as just a few examples.

Joshua Berne, senior software engineer on Bloomberg’s BDE team and an active member of the C++ Standards Committee’s Core Working Group (CWG) and Contracts Study Group (SG21), served multiple roles: Josh was our bridge between the core language and software development, performing structural rewrites of major features, including the features mentioned and many others. All benchmarking research conducted for this book was designed, performed, and/or reviewed by Josh. He provided the technical expertise needed to make LaTeX function to its fullest capabilities, designing and implementing the glossary, including automating the references back into the individual sections that use the terms. Importantly, Josh was the voice of reason throughout this entire project.

Lori Hughes, our project manager, frontline technical editor, and LaTeX designer and compositor, would probably tell you that herding cats is child’s play compared to what she

Acknowledgments

endured during this project. The tenacity, assertiveness, and roll-up-your-sleeves hard work she demonstrated relentlessly is arguably the only reason this book was published in 2021 (if not this decade). In short, Lori’s our rock; she is a veteran of **lakos20**, and we look forward to working with her on *all* of our planned future projects — e.g., allocators (**lakos22**), contracts (**lakos23**), Volumes II and III following **lakos20** (**lakos2a** and **lakos2b**), and anticipated future editions of this book incorporating C++17, C++20, etc.

Pablo Halpern, a former member of Bloomberg’s BDE team, an active member of the C++ Standards Committee, the creator of the `std::pmr` allocator model, and now a full-time collaborator with BDE working on language-level support for local memory allocators, served as a ghost writer for several features in this book (e.g., see Section 2.1. “User-Defined Literals” on page 835) and provided massive restructuring to many others (e.g., see Section 2.1. “Generalized PODs ’11” on page 401. Notably, out of all the nonauthors who contributed drafts in final form, only Pablo was able to write in a style approximating the authors’ voice. He also performed the research for a paper, commissioned by the authors of this book, demonstrating that **move operations**, though faster to execute initially, can have negative overall runtime implications due to **memory diffusion**; see **halpern21c**.

Dr. Andrei Alexandrescu — author of the seminal book *Modern C++ Design* (Addison-Wesley, 2001), coauthor of *C++ Coding Standards* (Addison-Wesley, 2005), and major contributor to the D language — was called upon for multiple assists in this endeavor: (1) as an expert author to provide an approachable guide to using variadic templates for those accustomed to their C++03 counterparts (see Section 2.1. “Variadic Templates” on page 873); (2) as a technical reviewer whose primary job was to reduce the tedium of John Lakos’ writing style and its numerous parenthetical phrases and footnotes; and (3) as a mascot and champion of our effort to imbue, on C++03 folk, the C++11/14 overlay of features. Andrei also generously agreed to write a foreword to this book, advocating its utility for senior developers familiar with classic C++.

Harold Bott, John’s TA in his Advanced C++ course during the 1990s at Columbia University, reconnected with John in 2019. Harry has since been a force in driving this book forward to completion. After a month of research with Nina, John entrusted Harry, a former programmer at Goldman Sachs and Executive Director at JP Morgan, with getting the flagship feature of modern C++ (see Section 2.1. “Rvalue References” on page 710) ready for review — a daunting task indeed. Once reviews were in and revisions were needed, Harry worked with John, nearly around the clock for almost three straight weeks, to incorporate reviewer feedback and to bring this important feature to the state in which it is presented here.

Mungo Gill is one of the newest full-time contributors on the BDE team and brings with him more than 30 years of professional software experience at such notable organizations as Salomon Brothers, Citigroup, Lehman Brothers, Google, and Citadel Securities. Mungo has reviewed every line of this book and has provided valuable feedback from a senior practitioner’s perspective. He also coordinated the process of assembling glossary definitions and gaining consensus among a host of eclectic domain experts.

Clay Wilson, a member of the BDE team since 2003, is another veteran of **lakos20**. Clay has, for the past 18 years, been our “closer” when it comes to reviewing software components.

Acknowledgments

His attention to detail and accuracy is, in our experience, second to none. Clay has reviewed much of this book, and we look forward to the possibility of working with him on future projects.

Steven Breitstein, a member of BDE since 2004 and an alumnus of **lakos20**, has reviewed every line and code snippet in this book and has made innumerable suggestions for manifestly improving the rendering of the material. He also stepped up and singlehandedly applied all the copy edits to our glossary.

Hyman Rosen retired from Bloomberg’s BDE team in April 2021 and was the master of pragmatic real-world use cases for some of the otherwise ostensibly *unsafe* features of modern C++, such as using extended friendship (see Section 3.1. “**friend** ’11” on page 1031) with the **curiously recurring template pattern (CRTP)**. You’ll find many others scattered throughout this book.

Stephen Dewhurst, an internationally recognized expert in C++ programming and popular repeat C++ author, conference speaker, and professional C++ trainer (including, for more than a decade, at Bloomberg), has reviewed *every* feature in this book and provided copious, practically valuable feedback, including a use case; see *Use Cases — Stateless lambdas* on page 605 within Section 2.1. “Lambdas.”

Jeffrey Olkin, who joined Bloomberg in 2011, is one of its most senior software architects, was the structural editor of **lakos20**, and has been a welcome advocate of this book from the start, reviewing many features, helping to organize the preliminary material, and providing his insightful and always valuable feedback along the way.

Steve Downey, a senior developer at Bloomberg since 2003, C++ Standards Committee member, and multidomain expert, contributed much of the advanced material found in a somewhat niche, *conditionally safe* feature of C++11; see Section 1.1. “Unicode Literals” on page 129. Mike Giroux and Oleg Subbotin fleshed out and provided benchmark material for another *conditionally safe* C++11 feature; see Section 2.1. “**extern template**” on page 353.

Sean Parent contributed a subsection assessing the strictness of current requirements on moved-from objects for standard containers; see *Annoyances — Standard Library requirements on a moved-from object are overly strict* on page 807 within Section 2.1. “Rvalue References.” Niall Douglas contributed a subsection detailing his experiences at scale with one of the *unsafe* C++ features; see *Appendix — Case study of using **inline** namespaces for versioning* on page 1083 within Section 3.1. “**inline namespace**.” Niels Dekker reviewed another *unsafe* C++11 feature (see Section 3.1. “**noexcept** Specifier” on page 1085) and provided valuable additional information as well as pointers to his own benchmark research. Kevin Klein helped organize and draft the material of yet another *unsafe* C++11 feature; see Section 3.1. “**final**” on page 1007.

Many senior C++ software engineers, instructors, and professional developers reviewed this work and provided copious feedback: Adil Al-Yasiri, Andrei Alexandrescu, Parsa Amini, Brian Bi, Frank Birbacher, Harry Bott, Steve Breitstein, Tomaz Canabrava, Bill Chapman, Marshall Clow, Stephen Dewhurst, Akshaye Dhawan, Niall Douglas, Steve Downey, Tom Eccles, Attila Feher, Kevin Fleming, J. Daniel Garcia, Mungo Gill, Mike Giroux, Kevin

Acknowledgments

Klein, Jeff Mendelsohn, Jeffrey Olkin, Nina Ranns, Hyman Rosen, Daniel Ruoso, Ben Saks, Richard Smith, Oleg Subbotin, Julian Templeman, Mike Verschell, Clay Wilson, and JC van Winkel.

In addition to reviewing the features of this book as they were being written, the BDE team fleshed out the first draft of the glossary, after which we relied again on Josh Berne, Brian Bi, Harry Bott, Mungo Gill, Pablo Halpern, and Nina Ranns to refine and consolidate it into a final draft before finally reviewing it ourselves in its totality. Jeff Mendelsohn and Nathan Burgers helped to distill the essence of this book onto its back cover. We want to thank all the Standards Committee members who provided valuable information when researching the details, history, etc., surrounding the various language features presented in this book. In particular, we would like to thank Bjarne Stroustrup for affably answering our pointed questions regarding anything related to C++. Howard Hinnant confirmed, among other things, the details of why and how *xvalues* were originally invented and how they have since morphed (a.k.a. “the delta”) from their original concept to their definition today. Michael Wong, Paul McKenney, and Maged Michael reviewed and signed off on our presentation of the `[[carries_dependency]]` attribute; see Section 3.1. “`carries_dependency`” on page 998. And we cannot thank Richard Smith enough for his thorough review and myriad suggestions on how to correct and improve our treatment of the flagship feature of modern C++ (see Section 2.1. “*Rvalue* References” on page 710). We hope that Richard will review *every* feature in subsequent editions of this book.

The team at Pearson — Greg Doench, our editor and fearless leader; Julie Nahil, our production manager; and Kim Wimpsett, our copy editor — have been very supportive of our efforts to get this book done quickly and accurately, despite its unorthodox workflow. We had originally projected that this book would contain 300–400 pages and would be complete by the end of 2020. That didn’t happen. Somehow, Greg and Julie found a way to accommodate our process and get this book printed in time for the 2021 winter holidays; thank you!

Online compilers, such as Godbolt (Compiler Explorer) and Wandbox, proved invaluable in the development of this work, allowing the team to rapidly evaluate and share code samples tested with various versions of multiple compilers accepting different dialects of the language.

We want to give a shout-out to the folks at Bloomberg involved in making sure that Bloomberg’s intellectual property and customer data were in no way compromised by anything contained herein and that appropriate attributions were made: Tom Arcidiacono, Kevin P. Fleming, and Chaim Haas.

Moreover, we want to recognize and thank our Bloomberg management for providing not just the support but the *imperative* to do this essential work for ourselves and then share it! In 2012, Vladimir Kliatchko, then and still Global Head of Engineering at Bloomberg, directed John Lakos, who collaborated with Rostislav Khlebnikov, to write a paper, **khlebnikov18**, to describe concisely the value proposition of C++11 and how best to exploit it. That short C++11 paper, 11 pages of 11-point type, was indeed well received, widely accepted, and ratified by fully 85 percent of the Standards Committee members who reviewed it. After

Acknowledgments

that, Andrei Basov, Engineering Manager, Middleware and Core Services; Akshaye Dhawan, Engineering Manager, Training, Documentation, and Work Management; and Adam Wolf, Head of Engineering, Software Infrastructure, encouraged and supported us in pursuing a more all-encompassing, practical-engineering-oriented treatment of modern C++, including C++14, in book form.

Finally, this book would not have been possible without the generous patronage of our Chief Technology Officer, Shawn Edwards. Without his support, and especially his sponsorship, the vast technical resources needed for this book to come to fruition could never have been brought to bear. Shawn, with his illustrious career as a developer, team lead, and technical manager, and now, as a senior executive, has graciously provided a foreword to this book.

About the Authors

John Lakos, author of *Large-Scale C++ Software Design* (Addison-Wesley, 1996) and *Large-Scale C++ Volume I: Process and Architecture* (Addison-Wesley, 2020), serves at Bloomberg in New York City as a senior architect and mentor for C++ software development worldwide. He is also an active voting member of the C++ Standards Committee's Evolution Working Group. From 1997 to 2001, Dr. Lakos directed the design and development of infrastructure libraries for proprietary analytic financial applications at Bear Stearns. From 1983 to 1997, Dr. Lakos was employed at Mentor Graphics, where he developed large frameworks and advanced ICCAD applications for which he holds multiple software patents. His academic credentials include a Ph.D. in Computer Science (1997) and an Sc.D. in Electrical Engineering (1989) from Columbia University. Dr. Lakos received his undergraduate degrees from MIT in Mathematics (1982) and Computer Science (1981).

Vittorio Romeo (B.Sc., Computer Science, 2016) is a senior software engineer at Bloomberg in London, where he builds mission-critical C++ middleware and delivers modern C++ training to hundreds of fellow employees. He began programming at the age of 8 and quickly fell in love with C++. Vittorio has created several open-source C++ libraries and games, has published many video courses and tutorials, and actively participates in the ISO C++ standardization process. He is an active member of the C++ community with an ardent desire to share his knowledge and learn from others: He presented more than 20 times at international C++ conferences (including CppCon, C++Now, ++it, ACCU, C++ On Sea, C++ Russia, and Meeting C++), covering topics from game development to template metaprogramming. Vittorio maintains a website (<https://vittorioromeo.info/>) with advanced C++ articles and a YouTube channel (<https://www.youtube.com/channel/UC1XihgHdkNOQd5IBHnIZWbA>) featuring well received modern C++11/14 tutorials. He is active on StackOverflow, taking great care in answering interesting C++ questions (75k+ reputation). When he is not writing code, Vittorio enjoys weightlifting and fitness-related activities as well as computer gaming and sci-fi movies.

Rostislav Khlebnikov is the lead of the BDE Solutions team that works on a variety of BDE libraries, such as the library for HTTP/2 communication, and contributes to other projects, including improving interoperability of BDE libraries with the Standard Library vocabulary types. He is an active member of the C++ Standards Committee and presented at CppCon 2019. Prior to his work at Bloomberg, Dr. Khlebnikov received his undergraduate degrees in Applied Mathematics and Computer Science from St. Petersburg State

About the Authors

Polytechnic University, Russia, and his Ph.D. in Computer Science from Graz University of Technology, Austria. He has worked professionally as a C++ software engineer for over 15 years.

Alisdair Meredith has been a member of the C++ Standards Committee since the inception of C++11 at the Oxford 2003 meeting, focusing on feature integration and actively finding and fixing language inconsistencies. Alisdair was the LWG chair when both C++11 and C++14 were published, for which he credits the hard work of the preceding chair, Howard Hinnant. Alisdair has been a perennial conference speaker for nearly 15 years, elucidating new work from the C++ Standards Committee. Alisdair joined Bloomberg's BDE team in 2009. For a decade prior, Alisdair worked as a professional C++ application programmer in F1 motor racing with the Benetton and Renault teams, winning two world championships! Between the two, Alisdair spent a year or so as a product manager at Borland, marketing their C++ products. Alisdair enjoys traveling, dining, and snorkeling.

Chapter 0

Introduction

Welcome! *Embracing Modern C++ Safely* is a reference book designed for professionals who develop and maintain large-scale, complex C++ software systems and want to leverage modern C++ features.

This book focuses on the productive value of each new language feature, starting with C++11, particularly when the systems and organizations involved are considered at scale. We deliberately left aside ideas and idioms — however clever and intellectually intriguing — that could hurt the bottom line when applied to large-scale systems. Instead, we focus on making wise economic and design decisions, with an understanding of the inevitable trade-offs that arise in any engineering discipline. In doing so, we do our best to steer clear of subjective opinions and recommendations.

Richard Feynman famously said, “If it disagrees with experiment, it’s wrong. In that simple statement is the key to science.”¹ There is no better way to experiment with a language feature than letting time do its work. We took that to heart and decided to cover only the features of modern C++ that have been part of the Standard for at least five years, which we believe provides enough perspective to properly evaluate the practical impact of new features. Thus, we are able to draw from practical experience to provide a thorough and comprehensive treatment that is worthy of your limited professional development time. If you’re looking for ways to improve your productivity by using tried and true modern C++ features, we hope this book will be the one you’ll reach for.

What’s missing from a book is as important as what’s present. *Embracing Modern C++ Safely*, known also as *EMC++S*, is not a tutorial on C++ programming or even on new features of C++. We assume you are an experienced developer, team lead, or manager; that you already have a good command of classic C++98/03; and that you are looking for clear, goal-driven ways to integrate modern C++ features into your toolbox.

What Makes This Book Different

The goal of the book you’re now reading is to be objective, empirical, and practical. We simply present features, their applicability, and their potential pitfalls as reflected by the analysis of millions of person-hours of using C++11 and C++14 in the development of

¹Richard Feynman, lecture at Cornell University, 1964. Video and commentary available at <https://fs.blog/2009/12/mental-model-scientific-method>.

varied large-scale software systems; personal preference matters have been neutralized to our best ability. We wrote down the distilled truth that remains, which should shape your understanding of what modern C++ has to offer without being skewed by our subjective opinions or domain-specific inclinations.

The final analysis and interpretation of what is appropriate for your context is left to you, the reader. This book is, by design, not a C++ style or coding-standards guide; it does, however, provide valuable input to any development organization seeking to author or enhance one.

Practicality is important to us in a real-world, economic sense. We examine modern C++ features through the lens of a large company developing and using software in a competitive environment. In addition to showing you how to best utilize a given C++ language feature in practice, our analysis takes into account the costs associated with routinely employing that feature in the ecosystem of a software development organization. Most texts omit the costs of using language features. In other words, we weigh the benefits of successfully using a feature against the hidden cost of its widespread ineffective use (or misuse) and/or the costs associated with training and code review required to reasonably ensure that such ill-conceived use does not occur. We are acutely aware that what applies to one person or a small crew of like-minded individuals is quite different from what works with a large, distributed team. The outcome of this analysis is our signature categorization of features based on how safe they are to adopt — namely, *safe*, *conditionally safe*, or *unsafe* features.

We are not aware of any similar text amid the rich offering of C++ textbooks; we wrote this book because we needed it.

Scope for the First Edition

Given the vastness of C++'s already voluminous and rapidly growing standardized libraries, we have chosen to limit this book's scope to just the language features themselves. A companion book, *Embracing Modern C++ Standard Libraries Safely*, is a separate project that we hope to tackle in the future. To be effective, this book, however, must remain focused on what expert C++ developers need to know well to be successful right now.

We chose to limit the scope of this first edition to only those features that have been included in the language standard since C++11 and widely available in practice for at least five years. This limited focus enables us to better evaluate the real-world impact of these features and to highlight any caveats that might not have been anticipated prior to standardization and sustained, active, and widespread use in industry.

We assume you are quite familiar with essentially all of the basic and important special-purpose features of classic C++98/03, so in this book we confine our attention to just the subset of C++ language features introduced in C++11 and C++14. This book is best for

you if you need to know how to safely incorporate C++11/14 language features into a predominately C++98/03 codebase, today.

We are actively planning to cover pre-C++11 material in future editions. For the time being, however, we highly recommend *Effective C++* by Scott Meyers² as a concise, practical treatment of many important and useful C++98/03 features.

The *EMC++S* Guiding Principles

Throughout the writing of *Embracing Modern C++ Safely*, we have followed a set of guiding principles, which collectively drive the style and content of this book.

Facts, Not Opinions

This book describes only beneficial uses and potential pitfalls of modern C++ features. The content presented is based on objectively verifiable facts, derived either from standards documents or from extensive practical experience; we explicitly avoid subjective opinions on the relative merits of design trade-offs (restraint that is a good exercise in humility). Although such opinions are often valuable, they are inherently biased toward the author's area of expertise.

Note that *safety* — the rating we use to segregate features by chapter — is the one exception to this objectivity guideline. Although the analysis of each feature aims at being entirely objective, each feature's chapter classification — indicating the relative safety of its quotidian use in a large software-development environment — reflects our combined decades of real-world, hands-on experience developing a variety of large-scale C++ software systems.

Elucidation, Not Prescription

We deliberately avoid prescribing any solutions to address specific feature pitfalls. Instead, we merely describe and characterize such concerns in sufficient detail to equip you to devise a solution suitable for your own development environment. In some cases, we might reference techniques or publicly available libraries that others have used to work around such speed bumps, but we do not pass judgment about which workaround should be considered a best practice.

Thorough, Not Superficial

Embracing Modern C++ Safely is neither designed nor intended to be an introduction to modern C++. This book is a handy reference for experienced C++ programmers who have

²meyers92

familiarity with earlier versions of the language (C++98/03). Our goal is to provide you with facts, detailed objective analysis, and cogent, real-world examples. By doing so, we spare you the task of wading through material that we presume you already know. If you are entirely unfamiliar with the C++ language, we suggest you start with a more elementary and language-centric text such as *The C++ Programming Language* by Bjarne Stroustrup.³

Real-World, Not Contrived, Examples

We hope you will find the examples in this book useful in multiple ways. The primary purpose of the examples is to illustrate productive use of each feature as it might occur in practice. We stay away from contrived examples that give equal importance to seldom-used aspects and to the intended, idiomatic uses of the feature. Hence, many of our examples are based on simplified code fragments extracted from real-world codebases. Though we typically change identifier names to be more appropriate to the shortened example (rather than the context and the process that led to the example), we keep the code structure of each example as close as possible to its original, real-world counterpart.

At Scale, Not Overly Simplistic, Programs

As with many aspects of software development, what works for small programs and teams often doesn't scale to larger development efforts. We attempt to simultaneously capture two distinct aspects of size: (1) the sheer product size (e.g., in bytes, source lines, separate units of release) of the programs, systems, and libraries developed and maintained by a software organization; and (2) the size of an organization itself as measured by the number of software developers, quality-assurance engineers, site-reliability engineers, operators, and so on that the organization employs.

What's more, powerful new language features in the hands of a few expert programmers working together on a prototype for their new start-up don't always fare as well when they are wantonly exercised by dozens or hundreds of developers in a large software-development organization. Hence, when we consider the relative safety of a feature, as defined in the next section, we do so with mindfulness that any given feature might be used — and occasionally misused — in large programs and by a large number of programmers having a wide range of knowledge, skill, and ability.

What Do We Mean by *Safely*?

The ISO C++ Standards Committee, of which we are members, would be remiss — and downright negligent — if it allowed any feature of the C++ language to be standardized if that feature were not reliably safe when used as intended. Still, we have chosen the word

³stroustrup13

“safely” as the moniker for the signature aspect of our book and the method by which we rank the risk-to-reward ratio for using a given feature in a large-scale development environment. By contextualizing the meaning of the term “safe,” we apply it to a real-world economy in which everything has a cost in multiple dimensions: risk of misuse, added maintenance burden borne by using a new feature in an older codebase, and training needs for developers who might not be familiar with that feature.

Several factors impact the value added by the adoption and widespread use of any new language feature, thereby reducing its intrinsic safety. By categorizing features in terms of safety, we strive to capture an appropriately weighted combination of the following factors:

- Number and severity of known deficiencies
- Difficulty in teaching consistent proper use
- Experience level required for consistent proper use
- Risks associated with widespread misuse

In this book, the degree of safety of a given feature is the relative likelihood that widespread use of that feature will have positive impact and no adverse effect on a large software company’s codebase.

A Safe Feature

Some of the new features of modern C++ add considerable value, are easy to use, and are decidedly hard to misuse unintentionally; hence, ubiquitous adoption of such features is productive, relatively unlikely to become a problem in the context of a large-scale development organization, and generally encouraged — even without training. We identify such staunchly helpful, unflappable C++ features as *safe*.

For example, we categorize the **override** contextual keyword as a safe feature because it prevents bugs, serves as documentation, cannot be easily misused, and has no serious deficiencies. If someone has heard of this feature and tried to use it and the software compiles, the codebase is likely better for it. Using **override** wherever applicable is always a sound engineering decision.

A Conditionally Safe Feature

The vast majority of new features available in modern C++ have important, frequently occurring, and valuable uses, yet how these features are used appropriately, let alone optimally, might not be obvious. What’s more, some of these features are fraught with inherent

dangers and deficiencies, requiring explicit training and extra care to circumnavigate their pitfalls.

For example, we deem default member initializers a *conditionally safe* feature because, although they are easy to use per se, the perhaps less-than-obvious unintended consequences of doing so (e.g., tight compile-time coupling) might be prohibitively costly in certain circumstances (e.g., might prevent relink-only patching in production).

An Unsafe Feature

When an expert programmer uses any C++ feature appropriately, the feature typically does no direct harm. Yet other developers — seeing the feature’s use in the codebase but failing to appreciate the highly specialized or nuanced reasoning justifying it — might attempt to use it in what they perceive to be a similar way, yet with profoundly less desirable results. Similarly, maintainers might change the use of a fragile feature, altering its semantics in subtle but damaging ways.

Features that are classified as unsafe are those that might have valid — and even important — use cases, yet our experience indicates that routine or widespread use would be counterproductive in a typical, large-scale, software-development enterprise.

For example, we deem the **final** contextual keyword an unsafe feature because the situations in which it would be misused overwhelmingly outnumber those vanishingly few isolated cases in which it is appropriate, let alone valuable. Furthermore, its widespread use would inhibit fine-grained (e.g., hierarchical) reuse, which is critically important to the success of a large organization.

Modern C++ Feature Catalog

This first edition of *Embracing Modern C++ Safely* was designed to serve as a comprehensive catalog of C++11 and C++14 language features, presenting vital information for each in a clear, consistent, and predictable format to which experienced engineers can readily refer during development or technical discourse.

Organization

This book is divided into four chapters, the last three of which form the catalog of modern C++ language features grouped by their respective safety classifications:

- Chapter 0: Introduction
- Chapter 1: *Safe* Features

- Chapter 2: *Conditionally Safe* Features
- Chapter 3: *Unsafe* Features

For this first edition, the language-feature chapters (1, 2, and 3) are divided into two sections containing, respectively, C++11 and C++14 features having the safety level (*safe*, *conditionally safe*, or *unsafe*) corresponding to that chapter. Recall, however, that Standard Library features are outside the scope of this book.

Each feature is presented in a separate section, rendered in a canonical format:

- **Description** — A brisk but comprehensive introduction of the feature’s syntax and semantics, supplemented with abundant code snippets. We do our best to avoid using other new features concurrently with the one being described, so each feature can be read independently and out of order. This might lead, on occasion, to code that is less fluent than it could otherwise be. Make sure you consult the “See Also” section (described below) to learn about crosstalk between features.
- **Use Cases** — A collection of tried-and-true use cases distilled from libraries and applications.
- **Potential Pitfalls** — Misuses of the feature that might lead to serious bugs and other problems.
- **Annoyances** — Shortcomings of the feature and unpleasant quirks that might make the feature less pleasant to use.
- **See Also** — Cross-references to other related features within this book along with a brief description of the connection.
- **Further Reading** — References to external sources discussing the feature.

Constraining our treatment of each individual feature to this canonized format facilitates rapid discovery of whatever particular aspects of a given language feature you are searching for.

Note that cross-references to subsections within a feature are in italics, and cross-references to other features are in normal text font. We refer to each feature within its relevant chapter and section: For example, Section 1.1. “Attribute Syntax” tells you that the “Attributes” feature is located in Chapter 1 (Safe) and within Section 1 (C++11). Terms that are defined within the glossary are set in a **different font**, with the first use in each feature being set in **bold**.

The commenting style is worth noting because it conveys good information in a terse format. Note that “description” or “details” provides additional descriptive information. Placeholders for irrelevant and/or unspecified code are shown with stylized comments in one of the following ways:

```

/*...*/
// ...
// ...                (<description>)

```

Code that does not compile will be marked with one of the following two comments:

```

// Error
// Error, <details>

```

Code that does not link will be marked with one of the following two comments:

```

// Link-Time Error
// Link-Time Error, <details>

```

Code that does not behave as expected at run time will be marked with one of the following two comments:

```

// Bug
// Bug, <details>

```

Code that behaves as expected will be marked with one of the following two comments:

```

// OK
// OK, <details>

```

Code that might warn but behaves as expected would be marked “OK, might warn” or similarly. For example, if a feature is deprecated until C++17 and removed in C++20, we might comment it like this:

```

// OK, deprecated4 (might warn)

```

How to Use This Book

Depending on your needs, *Embracing Modern C++ Safely* can be handy in a variety of ways.

- **Read the entire book from front to back.** If you are conversant with classic C++, consuming this book in its entirety will provide a complete and nuanced practical understanding of each of the language features introduced by C++11 and C++14.
- **Read the chapters in order but slowly over time.** An incremental, priority-driven approach is also possible and recommended, especially if you’re feeling less sure-footed. Understanding and applying first the *safe* features of Chapter 1 gets you the low-hanging fruit. In time, the *conditionally safe* features of Chapter 2 will allow you to ease into the breadth of useful modern C++ language features, prioritizing those that are least likely to prove problematic.

⁴Removed in C++20

- **Read the C++11 sections of each of the three catalog chapters first.** If you are a developer whose organization uses C++11 but not yet C++14, you can focus on learning everything that can be applied now and then circle back and learn the rest later when it becomes relevant to your evolving organization.
- **Use the book as a quick-reference guide if and as needed.** Random access is great, too, especially now that you’ve made it through Chapter 0. If you prefer not to read the book in its entirety (or simply want to refer to it periodically as a refresher), reading any arbitrary individual feature section in any order will provide timely access to all relevant details of whichever feature is of immediate interest.

We believe that you will derive value in several ways from the knowledge we imbued into *Embracing Modern C++ Safely*, irrespective of how you read it. In addition to helping you become a more knowledgeable and therefore safer developer, this book aims to clarify (whether you are a developer, a lead, or a manager) which features demand more training, attention to detail, experience, peer review, and such. The factual, objective presentation style also makes for excellent input into the preparation of coding standards and style guides that suit the particular needs of a company, project, team, or even just a single discriminating developer (which, of course, we all aim at being). Finally, any C++ software-development organization that adopts this book will be taking the first steps toward leveraging modern C++ in a way that maximizes reward while minimizing risks, i.e., by embracing modern C++ *safely*.

Last but definitely not least, this is *your* book in more than one sense of the word. It has been a collaborative effort with input from many engineers just like you, and it was “designed for maintenance” because we plan future revised editions with new features and improved treatment of the existing ones. Those future editions could greatly benefit from your contributions. Found something broken or missing? A clever use case? A hidden pitfall? An annoyance you can’t stand? We’d be happy to add it to the book. Point your browser to <http://emcpps.com>, and follow the instructions to send us feedback. Your input will be well received. You’ll find more information about the book on the website. Thank you, and happy coding!

The `override` Member-Function Specifier

Decorating a function in a derived class with the contextual keyword **override** ensures that a **virtual** function having a compatible declaration exists in one or more of its base classes.

Description

The **contextual keyword** **override** can be provided at the end of a member-function declaration to ensure that the decorated function is indeed *overriding* a corresponding **virtual** member function in a base class, as opposed to *hiding* it or otherwise inadvertently introducing a distinct function declaration:

```

struct Base
{
    virtual void f(int);
        void g(int);
    virtual void h(int) const;
    virtual void i(int) = 0;
};

struct DerivedWithoutOverride : Base
{
    void f();           // hides Base::f(int) (likely mistake)
    void f(int);       // OK, implicitly overrides Base::f(int)

    void g();           // hides Base::g(int) (likely mistake)
    void g(int);       // hides Base::g(int) (likely mistake)

    void h(int);        // hides Base::h(int) const (likely mistake)
    void h(int) const; // OK, implicitly overrides Base::h(int) const

    void i(int);       // OK, implicitly overrides Base::i(int)
};

struct DerivedWithOverride : Base
{
    void f()           override; // Error, Base::f() not found
    void f(int)       override; // OK, explicitly overrides Base::f(int)

    void g()           override; // Error, Base::g() not found
    void g(int)       override; // Error, Base::g() is not virtual.

    void h(int)       override; // Error, Base::h(int) not found
    void h(int) const override; // OK, explicitly overrides Base::h(int)
};

```

```
    void i(int)      override;    // OK, explicitly overrides Base::i(int)
};
```

Using this feature expresses design intent so that (1) human readers are aware of it and (2) compilers can validate it.

As noted, **override** is a contextual keyword. C++11 introduces keywords that have special meaning only in certain contexts. In this case, **override** is a keyword in the context of a declaration, but not otherwise using it as the identifier for a variable name, for example, is perfectly fine:

```
int override = 1; // OK
```

Use Cases

Ensuring that a member function of a base class is being overridden

Consider the following polymorphic hierarchy of error-category classes, as we might have defined them using C++03:

```
struct ErrorCategory
{
    virtual bool equivalent(const ErrorCode& code, int condition);
    virtual bool equivalent(int code, const ErrorCondition& condition);
};

struct AutomotiveErrorCategory : ErrorCategory
{
    virtual bool equivalent(const ErrorCode& code, int condition);
    virtual bool equivalent(int code, const ErrorCondition& condition);
};
```

Notice that there is a defect in the last line of the example above: `equivalent` has been misspelled. Moreover, the compiler did not catch that error. Clients calling `equivalent` on `AutomotiveErrorCategory` will incorrectly invoke the base-class function. If the function in the base class happens to be defined, the code might compile and behave unexpectedly at run time. Now, suppose that over time the interface is changed by marking the equivalence-checking function **const** to bring the interface closer to that of `std::error_category`:

```
struct ErrorCategory
{
    virtual bool equivalent(const ErrorCode& code, int condition) const;
    virtual bool equivalent(int code, const ErrorCondition& condition) const;
};
```

Without applying the corresponding modification to all classes deriving from `ErrorCategory`, the semantics of the program change due to the derived classes now hiding the base class's **virtual** member function instead of overriding it. Both errors discussed above would be detected automatically if the **virtual** functions in all derived classes were decorated with **override**:

```
struct AutomotiveErrorCategory : ErrorCategory
{
    bool equivalent(const ErrorCode& code, int condition) override;
    // Error, failed when base class changed

    bool equivalent(int code, const ErrorCondition& code) override;
    // Error, failed when first written
};
```

What's more, **override** serves as a clear indication of the derived-class author's intent to customize the behavior of `ErrorCategory`. For any given member function, using **override** necessarily renders any use of **virtual** for that function syntactically and semantically redundant. The only cosmetic reason for retaining **virtual** in the presence of **override** would be that **virtual** appears to the left of the function declaration, as it always has, instead of all the way to the right, as **override** does now.

Potential Pitfalls

Lack of consistency across a codebase

Relying on **override** as a means of ensuring that changes to base-class interfaces are propagated across a codebase can prove unreliable if this feature is used inconsistently, i.e., not applied in every circumstance where its use would be appropriate. In particular, altering the signature of a **virtual** member function in a base class and then compiling the entire code base will always flag as an error any nonmatching derived-class function where **override** was used but might fail even to warn where it is not.

Further Reading

- Various relationships among **virtual**, **override**, and **final** (see Section 3.1. “**final**” on page 1007) are presented in **boccaro20**.
- Scott Meyers advocates the use of the **override** specifier in **meyers15b**, “Item 12: Declare overriding functions **override**,” pp. 79–85.

The `[[deprecated]]` Attribute

The standard attribute `[[deprecated]]` indicates that the use of the entity to which the attribute pertains is discouraged, typically in the form of a compiler warning.

Description

The standard `[[deprecated]]` attribute is used to portably indicate that a particular **entity** is no longer recommended and to actively discourage its use. Such deprecation typically follows the introduction of alternative constructs that are superior to the original one, providing time for clients to migrate to them *asynchronously* before the deprecated one is removed in some subsequent release.

An asynchronous process for ongoing improvement of legacy codebases, sometimes referred to as **continuous refactoring**, often allows time for clients to migrate — on their own respective schedules and time frames — from existing *deprecated* constructs to newer ones, rather than having every client change in lock step. Allowing clients time to move *asynchronously* to newer alternatives is often the only viable approach unless (1) the codebase is a closed system, (2) all of the relevant code is governed by a single authority, and (3) the change can be made mechanically.

Although not strictly required, the Standard explicitly encourages¹ conforming compilers to produce a diagnostic message in case a program refers to any **entity** to which the `[[deprecated]]` attribute pertains. For instance, most popular compilers emit a warning whenever a `[[deprecated]]` function or object is used:

```
        void f();
[[deprecated]] void g();

        int a;
[[deprecated]] int b;

void h()
{
    f();
    g(); // Warning: g is deprecated.
    a;
    b; // Warning: b is deprecated.
}
```

¹The C++ Standard characterizes what constitutes a well-formed program, but compiler vendors require a great deal of leeway to facilitate the needs of their users. In case any feature induces warnings, command-line options are typically available to disable those warnings (`-wno-deprecated` in GCC), or methods are in place to suppress those warnings locally, e.g., `#pragma GCC diagnostic ignored "-wdeprecated"`.

The `[[deprecated]]` attribute can be used portably to decorate other entities: **class**, **struct**, **union**, type alias, variable, data member, function, enumeration, template specialization.²

A programmer can supply a **string literal** as an argument to the `[[deprecated]]` attribute — e.g., `[[deprecated("message")]]` — to inform human users regarding the reason for the deprecation:

```
[[deprecated("too slow, use algo1 instead")]] void algo0();
                                           void algo1();

void f()
{
    algo0(); // Warning: algo0 is deprecated; too slow, use algo1 instead.
    algo1();
}
```

An **entity** that is initially *declared* without `[[deprecated]]` can later be redeclared with the attribute and vice versa:

```
void f();
void g0() { f(); } // OK, likely no warnings

[[deprecated]] void f();
void g1() { f(); } // Warning: f is deprecated.

void f();
void g2() { f(); } // Warning: f is deprecated still.
```

As shown in `g2` in the example above, redeclaring an **entity** that was previously decorated with `[[deprecated]]` without the attribute leaves the entity still deprecated.

Use Cases

Discouraging use of an obsolete or unsafe entity

Decorating any **entity** with the `[[deprecated]]` attribute serves both to indicate a particular feature should not be used in the future and to actively encourage migration of existing uses to a better alternative. Obsolescence, lack of safety, and poor performance are common motivators for deprecation.

As an example of productive deprecation, consider the `RandomGenerator` class having a static `nextRandom` member function to generate random numbers:

²Applying `[[deprecated]]` to a specific enumerator or namespace, however, is guaranteed to be supported only since C++17; see [smith15a](#).

```

struct RandomGenerator
{
    static int nextRandom();
    // Generate a random value between 0 and 32767 (inclusive).
};

```

Although such a simple random number generator can be useful, it might become unsuitable for heavy use because good pseudorandom number generation requires more state (and the overhead of synchronizing such state for a single **static** function can be a significant performance bottleneck), while good random number generation requires potentially high overhead access to external sources of entropy. The `rand` function, inherited from C and available in C++ through the `<cstdlib>` header, has many of the same issues as our `RandomGenerator::nextRandom` function, and similarly developers are guided to use the facilities provided in the `<random>` header since C++11.

One solution is to provide an alternative random number generator that maintains more state, allows users to decide where to store that state (the random number generator objects), and overall offers more flexibility for clients. The downside of such a change is that it comes with a functionally distinct API, requiring that users update their code to move away from the inferior solution:

```

class StatefulRandomGenerator
{
    // ... (internal state of a quality pseudorandom number generator)

public:
    int nextRandom();
    // Generate a quality random value between 0 and 32767, inclusive.
};

```

Any user of the original random number generator can migrate to the new facility with little effort, but that is not a completely trivial operation, and migration will take some time before the original feature is no longer in use. The empathic maintainers of `RandomGenerator` can decide to use the `[[deprecated]]` attribute to discourage continued use of `RandomGenerator::nextRandom()` instead of removing it completely:

```

struct RandomGenerator
{
    [[deprecated("Use StatefulRandomGenerator class instead.")]
    static int nextRandom();
    // ...
};

```

By using `[[deprecated]]` as shown in the previous example, existing clients of `RandomGenerator` are informed that a superior alternative, `BetterRandomGenerator`, is available, yet they are granted time to migrate their code to the new solution rather than having their code broken by the removal of the old solution. When clients are notified of the deprecation (thanks to a compiler diagnostic), they can schedule time to rewrite their applications to consume the new interface.

Continuous refactoring is an essential responsibility of a development organization, and deciding when to go back and fix what's suboptimal instead of writing new code that will please users and contribute more immediately to the bottom line will forever be a source of tension. Allowing disparate development teams to address such improvements in their own respective time frames, perhaps subject to some reasonable overall deadline date, is a proven real-world practical way of ameliorating this tension.

Potential Pitfalls

Interaction with treating warnings as errors

In some code bases, compiler warnings are promoted to errors using compiler flags, such as `-werror` for GCC and Clang or `/WX` for MSVC, to ensure that their builds are warning-clean. For such code bases, use of the `[[deprecated]]` attribute by their dependencies as part of the API might introduce unexpected compilation failures.

Having the compilation process completely stopped due to use of a deprecated **entity** defeats the purpose of the attribute because users of such an **entity** are given no time to adapt their code to use a newer alternative. On GCC and Clang, users can selectively demote deprecation errors back to warnings by using the `-wno-error=deprecated-declarations` compiler flag. On MSVC, however, such demotion of warnings is not possible, and the available workarounds, such as entirely disabling the effects of the `/WX` flag or the deprecation diagnostics using the `-wd4996` flag, are often unsuitable.

Furthermore, this interaction between `[[deprecated]]` and treating warnings as errors makes it impossible for owners of a low-level library to deprecate a function when releasing their code requires that they do not break the ability for *any* of their higher-level clients to compile; a single client using the to-be-deprecated function in a code base that treats warnings as errors prevents the release of the code that uses the `[[deprecated]]` attribute. With the frequent advice given in practice to aggressively treat warnings as errors, the use of `[[deprecated]]` might be completely unfeasible.

Explicit-Instantiation Declarations

The **extern template** prefix can be used to suppress *implicit* generation of local object code for the definitions of particular specializations of class, function, or variable templates used within a translation unit, with the expectation that any suppressed object-code-level definitions will be provided elsewhere within the program by template definitions that are instantiated *explicitly*.

Description

Inherent in the current ecosystem for supporting template programming in C++ is the need to generate redundant definitions of fully specified function and variable templates within `.o` files. For common instantiations of popular templates, such as `std::vector`, the increased object-file size, a.k.a. **code bloat**, and potentially extended link times might become significant:

```
#include <vector>      // std::vector is a popular template.
std::vector<int> v;    // std::vector<int> is a common instantiation.

#include <string>      // std::basic_string is a popular template.
std::string s;        // std::string, an alias for std::basic_string<char>, is
                      // a common instantiation.
```

The intent of the **extern template** feature is to *suppress* the implicit generation of duplicative object code within every translation unit in which a fully specialized class template, such as `std::vector<int>` in the code snippet above, is used. Instead, **extern template** allows developers to choose a single translation unit in which to explicitly *generate* object code for all the definitions pertaining to that specific template specialization as explained next.

Explicit-instantiation definition

Creating an **explicit-instantiation definition** was possible prior to C++11.¹ The requisite syntax is to place the keyword **template** in front of the name of the fully specialized class template, function template, or, in C++14, variable template (see Section 1.2. “Variable Templates” on page 157):

¹The C++03 Standard term for the syntax used to create an **explicit-instantiation definition**, though rarely used, was **explicit-instantiation directive**. The term **explicit-instantiation directive** was clarified in C++11 and can now also refer to syntax that is used to create a *declaration* — i.e., **explicit-instantiation declaration**.

```
#include <vector> // std::vector (general template)

template class std::vector<int>;
    // Deposit all definitions for this specialization into the .o for this
    // translation unit.
```

This explicit-instantiation directive compels the compiler to instantiate *all* functions defined by the named `std::vector` class template having the specified `int` template argument; any collateral object code resulting from these instantiations will be deposited in the resulting `.o` file for the current translation unit. Importantly, even functions that are never used are still instantiated, so this solution might not be the correct one for many classes; see *Potential Pitfalls — Accidentally making matters worse* on page 373.

Explicit-instantiation declaration

C++11 introduced the explicit-instantiation declaration, a complement to the explicit-instantiation definition. The newly provided syntax allows us to place **extern template** in front of the declaration of an explicit specialization of a class template, a function template, or a variable template:

```
#include <vector> // std::vector (general template)

extern template class std::vector<int>;
    // Suppress depositing of any object code for std::vector<int> into the
    // .o file for this translation unit.
```

Using the modern **extern template** syntax above instructs the compiler to *refrain* from depositing any object code for the named specialization in the current translation unit and instead to rely on some other translation unit to provide any missing object-level definitions that might be needed at link time; see *Annoyances — No good place to put definitions for unrelated classes* on page 373.

Note, however, that declaring an explicit instantiation to be an **extern template** *in no way* affects the ability of the compiler to instantiate and to inline visible function-definition bodies for that template specialization in the translation unit:

```
// client.cpp:
#include <vector> // std::vector (general template)

extern template class std::vector<int>;

void client(std::vector<int>& inOut) // fully specialized instance of a vector
{
    if (inOut.size()) // This invocation of size can inline.
    {
        int value = inOut[0]; // This invocation of operator[] can be inlined.
    }
}
```

In the previous example, the two tiny member functions of `vector`, namely, `size` and `operator[]`, will typically be inlined — in precisely the same way they would have been had the **extern template** declaration been omitted. The *only* purpose of an **extern template** declaration is to suppress object-code generation for this particular template instantiation for the current translation unit.

Finally, note that the use of **explicit-instantiation directives** has absolutely no effect on the logical meaning of a well-formed program; in particular, when applied to specializations of function templates, they have no effect on overload resolution:

```
template <typename T> bool f(T v) { /*...*/ } // general template definition

extern template bool f(char c); // specialization of f for char
extern template bool f(int v); // specialization of f for int

bool bc = f((char) 0); // exact match: Object code is suppressed locally.
bool bs = f((short) 0); // not exact match: Object code is generated locally.
bool bi = f((int) 0); // exact match: Object code is suppressed locally.
bool bu = f((unsigned)0); // not exact match: Object code is generated locally.
```

As the example above illustrates, overload resolution and template argument deduction occur independently of any **explicit-instantiation** declarations. Only *after* the template to be instantiated is determined does the **extern template** syntax take effect; see also *Potential Pitfalls — Corresponding explicit-instantiation declarations and definitions* on page 371.

A more complete illustrative example

So far, we have seen the use of **explicit-instantiation** declarations and **explicit-instantiation** definitions applied to only a standard *class* template, `std::vector`. The same syntax shown in the previous code snippet applies also to full specializations of individual function templates and variable templates.

As a more comprehensive, albeit largely pedagogical, example, consider the overly simplistic `my::Vector` class template along with other related templates defined within a header file, `my_vector.h`:

```
// my_vector.h
#ifndef INCLUDED_MY_VECTOR // internal include guard
#define INCLUDED_MY_VECTOR

#include <cstdlib> // std::size_t
#include <utility> // std::swap

namespace my // namespace for all entities defined within this component
{

template <typename T>
class Vector
```

```

{
    static std::size_t s_count;    // track number of objects constructed
    T*                 d_data_p;  // pointer to dynamically allocated memory
    std::size_t        d_length;  // current number of elements in the vector
    std::size_t        d_capacity; // number of elements currently allocated

public:
    // ...

    std::size_t length() const { return d_length; }
        // Return the number of elements.

    // ...
};

// ...          Any partial or full specialization definitions          ...
// ...          of the class template Vector go here.                  ...

template <typename T>
void swap(Vector<T> &lhs, Vector<T> &rhs) { return std::swap(lhs, rhs); }
    // free function that operates on objects of type my::Vector via ADL

// ...          Any [full] specialization definitions          ...
// ...          of free function swap would go here.            ...

template <typename T>
const std::size_t vectorSize = sizeof(Vector<T>); // C++14 variable template
    // This nonmodifiable static variable holds the size of a my::Vector<T>.

// ...          Any [full] specialization definitions          ...
// ...          of variable vectorSize would go here.          ...

template <typename T>
std::size_t Vector<T>::s_count = 0;
    // definition of static counter in general template

// ... We might opt to add explicit-instantiation declarations here.
// ...

} // Close my namespace.

#endif // Close internal include guard.

```

In the `my_vector` component in the code snippet above, we have defined the following, in the `my` namespace.

1. A **class** template, `Vector`, parameterized on element type

2. A free-function template, `swap`, that operates on objects of corresponding specialized `Vector` type
3. A **const** C++14 variable template, `vectorSize`, that represents the number of bytes in the **footprint** of an object of the corresponding specialized `Vector` type

Any use of these templates by a client might and typically will trigger the depositing of equivalent definitions as object code in the client translation unit's resulting `.o` file, irrespective of whether the definition being used winds up getting inlined.

To eliminate object code for specializations of entities in the `my_vector` component, we must first decide where the unique definitions will go; see *Annoyances — No good place to put definitions for unrelated classes* on page 373. In this specific case, we own the component that requires specialization, and the specialization is for a ubiquitous built-in type; hence, the natural place to generate the specialized definitions is in a `.cpp` file corresponding to the component's header:

```
// my_vector.cpp:
#include <my_vector.h> // We always include the component's own header first.
    // By including this header file, we have introduced the general template
    // definitions for each of the explicit-instantiation declarations below.

namespace my // namespace for all entities defined within this component
{

template class Vector<int>;
    // Generate object code for all nontemplate member functions and definitions
    // of static data members of template my::Vector having int elements.

template std::size_t Vector<double>::length() const; // BAD IDEA
    // In addition, we could generate object code for just a particular member
    // function definition of my::Vector (e.g., length) for some other
    // argument type (e.g., double).

template void swap(Vector<int>& lhs, Vector<int>& rhs);
    // Generate object code for the full specialization of the swap free-
    // function template that operates on objects of type my::Vector<int>.

template const std::size_t vectorSize<int>; // C++14 variable template
    // Generate the object-code-level definition for the specialization of the
    // C++14 variable template instantiated for built-in type int.

template std::size_t Vector<int>::s_count;
    // Generate the object-code-level definition for the specialization of the
    // static member variable of Vector instantiated for built-in type int.

} // Close my namespace.
```

Each of the constructs introduced by the keyword **template** within the `my` namespace in the previous example represents a separate **explicit-instantiation definition**. These constructs instruct the compiler to generate object-level definitions for general templates declared in `my_vector.h` specialized on the built-in type `int`. Explicit instantiation of individual member functions, such as `length()` in the example, is, however, only rarely useful; see *Annoyances — All members of an explicitly defined template class must be valid* on page 374.

Having installed the necessary **explicit-instantiation definitions** in the component's `my_vector.cpp` file, we must now go back to its `my_vector.h` file and, without altering any of the previously existing lines of code, *add* the corresponding **explicit-instantiation declarations** to suppress redundant local code generation:

```
// my_vector.h:
#ifndef INCLUDED_MY_VECTOR // internal include guard
#define INCLUDED_MY_VECTOR

namespace my // namespace for all entities defined within this component
{

// ...
// ... everything that was in the original my namespace
// ...

// -----
// explicit-instantiation declarations
// -----

extern template class Vector<int>;
    // Suppress object code for this class template specialized for int.

extern template std::size_t Vector<double>::length() const; // BAD IDEA
    // Suppress object code for this member, only specialized for double.

extern template void swap(Vector<int>& lhs, Vector<int>& rhs);
    // Suppress object code for this free function specialized for int.

extern template std::size_t vectorSize<int>; // C++14
    // Suppress object code for this variable template specialized for int.

extern template std::size_t Vector<int>::s_count;
    // Suppress object code for this static member definition w.r.t. int.

} // Close my namespace.

#endif // Close internal include guard.
```

Each of the constructs that begins with **extern template** in the example above are **explicit-instantiation declarations**, which serve only to suppress the generation of any object code

emitted to the `.o` file of the current translation unit in which such specializations are used. These added **extern template** declarations must appear in `my_header.h` *after* the declaration of the corresponding general template and, importantly, before whatever relevant definitions are ever used.

The effect on various `.o` files

To illustrate the effect of **explicit-instantiation** declarations and **explicit-instantiation definitions** on the contents of object and executable files, we'll use a simple `lib_interval` library **component** consisting of a header file, `lib_interval.h`, and an implementation file, `lib_interval.cpp`. The latter, apart from including its corresponding header, is effectively empty:

```
// lib_interval.h:
#ifndef INCLUDED_LIB_INTERVAL // internal include guard
#define INCLUDED_LIB_INTERVAL

namespace lib // namespace for all entities defined within this component
{

template <typename T> // elided definition of a class template
class Interval
{
    T d_low; // interval's low value
    T d_high; // interval's high value

public:
    explicit Interval(const T& p) : d_low(p), d_high(p) { }
        // Construct an empty interval.

    Interval(const T& low, const T& high) : d_low(low), d_high(high) { }
        // Construct an interval having the specified boundary values.

    const T& low() const { return d_low; }
        // Return this interval's low value.

    const T& high() const { return d_high; }
        // Return this interval's high value.

    int length() const { return d_high - d_low; }
        // Return this interval's length.

    // ...
};

template <typename T> // elided definition of a function template
bool intersect(const Interval<T>& i1, const Interval<T>& i2)
    // Determine whether the specified intervals intersect.
```

```

{
    bool result = false; // nonintersecting until proven otherwise
    // ...
    return result;
}

} // Close lib namespace.

#endif // INCLUDED_LIB_INTERVAL

// lib_interval.cpp:
#include <lib_interval.h>

```

This library component above defines, in the namespace `lib`, an implementation of (1) a class template, `Interval`, and (2) a function template, `intersect`.

Let's also consider a trivial application that uses this library component:

```

// app.cpp:
#include <lib_interval.h> // Include the library component's header file.

int main(int argv, const char** argc)
{
    lib::Interval<double> a(0, 5); // instantiate with double type argument
    lib::Interval<double> b(3, 8); // instantiate with double type argument
    lib::Interval<int> c(4, 9); // instantiate with int type argument

    if (lib::intersect(a, b)) // instantiate deducing double type argument
    {
        return 0; // Return "success" as (0.0, 5.0) does intersect (3.0, 8.0).
    }

    return 1; // Return "failure" status as function apparently doesn't work.
}

```

The purpose of this application is merely to exhibit a couple of instantiations of the library *class* template, `lib::Interval`, for type arguments `int` and `double`, and of the library *function* template, `lib::intersect`, for just `double`.

Next, we compile the application and library translation units, `app.cpp` and `lib_interval.cpp`, and inspect the symbols in their respective corresponding object files, `app.o` and `lib_interval.o`:

```

$ gcc -I. -c app.cpp lib_interval.cpp
$ nm -C app.o lib_interval.o

app.o:
0000000000000000 w lib::Interval<double>::Interval(double const&, double const&)

```

```

0000000000000000 W lib::Interval<int>::Interval(int const&, int const&)
0000000000000000 W bool lib::intersect<double>(lib::Interval<double> const&,
                                                lib::Interval<double> const&)

0000000000000000 T main

lib_interval.o:

```

Looking at `app.o` in the previous example, the class and function templates used in the `main` function, which is defined in the `app.cpp` file, were instantiated *implicitly*, and the relevant code was added to the resulting object file, `app.o`, with each instantiated function definition in its own separate **section**. In the `Interval` *class* template, the generated symbols correspond to the two unique instantiations of the constructor, i.e., for **double** and **int**, respectively. The `intersect` function template, however, was implicitly instantiated for only type **double**. Note that all of the implicitly instantiated functions have the `W` symbol type, indicating that they are *weak* symbols, which are permitted to be present in multiple object files. By contrast, this file also defines the *strong* symbol `main`, marked here by a `T`. Linking `app.o` with any other object file containing such a symbol would cause the linker to report a multiply-defined-symbol error. On the other hand, the `lib_interval.o` file corresponds to the `lib_interval` library component, whose `.cpp` file served only to include its own `.h` file, and is again effectively empty.

Let's now link the two object files, `app.o` and `lib_interval.o`, and inspect the symbols in the resulting executable, `app2`:

```

$ gcc -o app app.o lib_interval.o
$ nm -C app
000000000040056e W lib::Interval<double>::Interval(double const&, double const&)
00000000004005a2 W lib::Interval<int>::Interval(int const&, int const&)
00000000004005ce W bool lib::intersect<double>(lib::Interval<double> const&,
                                                lib::Interval<double> const&)

00000000004004b7 T main

```

As the textual output above confirms, the final program contains exactly one copy of each weak symbol. In this tiny illustrative example, these weak symbols have been defined in only a single object file, thus not requiring the linker to select one definition out of many.

More generally, if the application comprises multiple object files, each file will potentially contain their own set of weak symbols, often leading to duplicate code **sections** for implicitly instantiated class, function, and variable templates instantiated on the same type arguments. When the linker combines object files, it will arbitrarily choose at most one of each of these respective and ideally identical weak-symbol **sections** to include in the final executable.

Imagine now that our program includes a large number of object files, many of which make use of our `lib_interval` component, particularly to operate on **double** intervals.

²We have stripped out extraneous unrelated information that the `nm` tool produces; note that the `-C` option invokes the symbol demangler, which turns encoded names like `_ZN3lib8IntervalIdEC1ERKdS3_` into something more readable like `lib::Interval<double>::Interval(double const&, double const&)`.

Suppose, for now, we decide we would like to suppress the generation of object code for templates related to just **double** type with the intent of later putting them all in one place, i.e., the currently empty `lib_interval.o`. Achieving this objective is precisely what the **extern template** syntax is designed to accomplish.

Returning to our `lib_interval.h` file, we need not change one line of code; we need only to *add* two **explicit-instantiation** declarations — one for the *class* template, `Interval<double>`, and one for the *function* template, `intersect<double>(const double&, const double&)` — to the header file anywhere *after* their respective corresponding general template declaration and definition:

```
// lib_interval.h: // No change to existing code.
#ifndef INCLUDED_LIB_INTERVAL // internal include guard
#define INCLUDED_LIB_INTERVAL

namespace lib // namespace for all entities defined within this component
{

template <typename T>
class Interval
{
    // ... (same as before)
};

template <typename T>
bool intersect(const Interval<T>& i1, const Interval<T>& i2)
{
    // ... (same as before)
}

extern template class Interval<double>; // explicit-instantiation declaration

extern template // explicit-instantiation declaration
bool intersect(const Interval<double>&, const Interval<double>&);

} // close lib namespace

#endif // INCLUDED_LIB_INTERVAL
```

Let's again compile the two `.cpp` files and inspect the corresponding `.o` files:

```
$ gcc -I. -c app.cpp lib_interval.cpp
$ nm -C app.o lib_interval.o

app.o:
        U lib::Interval<double>::Interval(double const&, double const&)
```

```

0000000000000000 W lib::Interval<int>::Interval(int const&, int const&)
                U bool lib::intersect<double>(lib::Interval<double> const&,
                lib::Interval<double> const&)

0000000000000000 T main

lib_interval.o:

```

Notice that this time some of the symbols, specifically those relating to the class and function templates instantiated for type **double**, have changed from **W**, indicating a *weak* symbol, to **U**, indicating an *undefined* one. This symbol type change means that instead of generating a weak symbol for the explicit specializations for **double**, the compiler left those symbols undefined, as if only the *declarations* of the member and free-function templates had been available when compiling `app.cpp`, yet inlining of the instantiated definitions is in no way affected. **Undefined symbols** are expected to be made available to the linker from other object files. Attempting to link this application expectedly fails because no object files being linked contain the needed definitions for those instantiations:

```

$ gcc -o app app.o lib_interval.o

app.o: In function 'main':
app.cpp:(.text+0x38): undefined reference to
`lib::Interval<double>::Interval(double const&, double const&)'
app.cpp:(.text+0x69): undefined reference to
`lib::Interval<double>::Interval(double const&, double const&)'
app.cpp:(.text+0xa1): undefined reference to
`bool lib::intersect<double>(lib::Interval<double> const&,
                lib::Interval<double> const&)'

collect2: error: ld returned 1 exit status

```

To provide the missing definitions, we will need to instantiate them explicitly. Since the type for which the class and function are being specialized is the ubiquitous built-in type, **double**, the ideal place to sequester those definitions would be within the object file of the `lib_interval` library component itself, but see *Annoyances — No good place to put definitions for unrelated classes* on page 373. To force the needed template definitions into the `lib_interval.o` file, we will need to use **explicit-instantiation definition** syntax, i.e., the **template** prefix:

```

// lib_interval.cpp:
#include <lib_interval.h>

template class lib::Interval<double>;
    // example of an explicit-instantiation definition for a class

template bool lib::intersect(const Interval<double>&, const Interval<double>&);
    // example of an explicit-instantiation definition for a function

```

We recompile once again and inspect our newly generated object files:

```
$ gcc -I. -c app.cpp lib_interval.cpp
$ nm -C app.o lib_interval.o

app.o:
                U lib::Interval<double>::Interval(double const&, double const&)
0000000000000000 W lib::Interval<int>::Interval(int const&, int const&)
                U bool lib::intersect<double>(lib::Interval<double> const&,
                                                lib::Interval<double> const&)
0000000000000000 T main

lib_interval.o:
0000000000000000 W lib::Interval<double>::Interval(double const&)
0000000000000000 W lib::Interval<double>::Interval(double const&, double const&)
0000000000000000 W lib::Interval<double>::low() const
0000000000000000 W lib::Interval<double>::high() const
0000000000000000 W lib::Interval<double>::length() const
0000000000000000 W bool lib::intersect<double>(lib::Interval<double> const&,
                                                lib::Interval<double> const&)
```

The application object file, `app.o`, naturally remained unchanged. What's new here is that the functions that were missing from the `app.o` file are now available in the `lib_interval.o` file, again as *weak* (`W`), as opposed to strong (`T`), symbols. Notice, however, that explicit instantiation forces the compiler to generate code for all of the member functions of the class template for a given specialization. These symbols might all be linked into the resulting executable unless we take explicit precautions to exclude those that aren't needed³:

```
$ gcc -o app app.o lib_interval.o -Wl,--gc-sections
$ nm -C app
00000000004005ca W lib::Interval<double>::Interval(double const&, double const&)
000000000040056e W lib::Interval<int>::Interval(int const&, int const&)
000000000040063d W bool lib::intersect<double>(lib::Interval<double> const&,
                                                lib::Interval<double> const&)
00000000004004b7 T main
```

The **extern template** feature is provided to enable software architects to reduce code bloat in individual object files for common instantiations of class, function, and, as of C++14, variable templates in large-scale C++ software systems. The practical benefit is in reducing the physical size of libraries, which *might* lead to improved link times. **Explicit-instantiation declarations** do *not* (1) affect the meaning of a program, (2) suppress inline template implicit instantiation, (3) impede the compiler's ability to **inline**, or (4) meaningfully improve

³To avoid including the explicitly generated definitions that are being used to resolve undefined symbols, we have instructed the linker to remove all unused code **sections** from the executable. The `-wl` option passes comma-separated options to the linker. The `--gc-sections` option instructs the compiler to compile and assemble and instructs the linker to omit individual unused sections, where each section contains, for example, its own instantiation of a function template.

compile time. To be clear, the *only* purpose of the **extern template** syntax is to suppress object-code generation for the current translation unit, which is then selectively overridden in the translation unit(s) of choice.

Use Cases

Reducing template code bloat in object files

The motivation for the **extern template** syntax is as a purely compile-time, not runtime, optimization, i.e., to reduce the amount of redundant code within individual object files resulting from common template instantiations in client code. As an example, consider a fixed-size-array class template, `FixedArray`, that is used widely, i.e., by many clients from separate translation units, in a large-scale game project for both integral and floating-point calculations, primarily with type arguments **int** and **double** and array sizes of either 2 or 3:

```
// game_fixedarray.h:
#ifndef INCLUDED_GAME_FIXEDARRAY // internal include guard
#define INCLUDED_GAME_FIXEDARRAY

#include <cstdint> // std::size_t

namespace game // namespace for all entities defined within this component
{

template <typename T, std::size_t N> // widely used class template
class FixedArray
{
    // ... (elided private implementation details)
public:
    FixedArray() { /*...*/ }
    FixedArray(const FixedArray<T, N>& other) { /*...*/ }
    T& operator[](std::size_t index) { /*...*/ }
    const T& operator[](std::size_t index) const { /*...*/ }
};

template <typename T, std::size_t N>
T dot(const FixedArray<T, N>& a, const FixedArray<T, N>& b) { /*...*/ }
    // Return the scalar ("dot") product of the specified 'a' and 'b'.

// Explicit-instantiation declarations for full template specializations
// commonly used by the game project are provided below.

extern template class FixedArray<int, 2>; // class template
extern template int dot(const FixedArray<int, 2>& a, // function template
                       const FixedArray<int, 2>& b); // for int and 2
```

```

extern template class FixedArray<int, 3>;           // class template
extern template int dot(const FixedArray<int, 3>& a, // function template
                       const FixedArray<int, 3>& b); // for int and 3

extern template class FixedArray<double, 2>;      // for double and 2
extern template double dot(const FixedArray<double, 2>& a,
                           const FixedArray<double, 2>& b);

extern template class FixedArray<double, 3>;      // for double and 3
extern template double dot(const FixedArray<double, 3>& a,
                           const FixedArray<double, 3>& b);

} // Close game namespace.

#endif // INCLUDED_GAME_FIXEDARRAY

```

Specializations commonly used by the `game` project are provided by the `game` library. In the component header in the example above, we have used the **extern template** syntax to suppress object-code generation for instantiations of both the class template `FixedArray` and the function template `dot` for element types `int` and `double`, each for array sizes 2 and 3. To ensure that these specialized definitions are available in every program that might need them, we use the **template** syntax counterpart to *force* object-code generation within just the one `.o` corresponding to the `game_fixedarray` library component⁴:

```

// game_fixedarray.cpp:
#include <game_fixedarray.h> // included as first substantive line of code

// Explicit-instantiation definitions for full template specializations
// commonly used by the game project are provided below.

template class game::FixedArray<int, 2>;          // class template
template int game::dot(const FixedArray<int, 2>& a, // function template
                      const FixedArray<int, 2>& b); // for int and 2

template class game::FixedArray<int, 3>;          // class template
template int game::dot(const FixedArray<int, 3>& a, // function template
                      const FixedArray<int, 3>& b); // for int and 3

template class game::FixedArray<double, 2>;      // for double and 2
template double game::dot(const FixedArray<double, 2>& a,
                          const FixedArray<double, 2>& b);

```

⁴Notice that we have chosen *not* to nest the explicit specializations — or any other definitions — of entities already declared directly within the `game` namespace, preferring instead to qualify each entity explicitly to be consistent with how we render free-function definitions to avoid self-declaration; see **lakos20**, section 2.5, “Component Source-Code Organization,” pp. 333–342, specifically Figure 2-36b, p. 340. See also *Potential Pitfalls — Corresponding explicit-instantiation declarations and definitions* on page 371.

```

template class game::FixedArray<double, 3>;           // for double and 3
template double game::dot(const FixedArray<double, 3>& a,
                          const FixedArray<double, 3>& b);

```

Compiling `game_fixedarray.cpp` and examining the resulting object file shows that the code for all explicitly instantiated classes and free functions was generated and placed into the object file, `game_fixedarray.o`, of which we show a subset of the relevant symbols:

```

$ gcc -I. -c game_fixedarray.cpp
$ nm -C game_fixedarray.o
0000000000000000 W game::FixedArray<double, 2ul>::FixedArray(
    game::FixedArray<double, 2ul> const&)
0000000000000000 W game::FixedArray<double, 2ul>::FixedArray()
0000000000000000 W game::FixedArray<double, 2ul>::operator[](unsigned long)
0000000000000000 W game::FixedArray<double, 3ul>::FixedArray(
    game::FixedArray<double, 3ul> const&)
0000000000000000 W game::FixedArray<int, 3ul>::FixedArray()
:
0000000000000000 W double game::dot<double, 2ul>(
    game::FixedArray<double, 2ul> const&, game::FixedArray<double, 2ul> const&)
0000000000000000 W double game::dot<double, 3ul>(
    game::FixedArray<double, 3ul> const&, game::FixedArray<double, 3ul> const&)
0000000000000000 W int game::dot<int, 2ul>(
    game::FixedArray<int, 2ul> const&, game::FixedArray<int, 2ul> const&)
:
0000000000000000 W game::FixedArray<int, 2ul>::operator[](unsigned long) const
0000000000000000 W game::FixedArray<int, 3ul>::operator[](unsigned long) const

```

This `FixedArray` class template is used in multiple translation units within the `game` project. The first one contains a set of geometry utilities:

```

// app_geometryutil.cpp:

#include <game_fixedarray.h> // game::FixedArray
#include <game_unit.h>      // game::Unit

using namespace game;

void translate(Unit* object, const FixedArray<double, 2>& dst)
    // Perform precise movement of the object on 2D plane.
{
    FixedArray<double, 2> objectProjection;
    // ...
}

void translate(Unit* object, const FixedArray<double, 3>& dst)
    // Perform precise movement of the object in 3D space.

```

```

{
    FixedArray<double, 3> delta;
    // ...
}

bool isOrthogonal(const FixedArray<int, 2>& a1, const FixedArray<int, 2>& a2)
    // Return true if 2d arrays are orthogonal.
{
    return dot(a1, a2) == 0;
}

bool isOrthogonal(const FixedArray<int, 3>& a1, const FixedArray<int, 3>& a2)
    // Return true if 3d arrays are orthogonal.
{
    return dot(a1, a2) == 0;
}

```

The second one deals with physics calculations:

```

// app_physics.cpp:

#include <game_fixedarray.h> // game::FixedArray
#include <game_unit.h>      // game::Unit

using namespace game;

void collide(Unit* objectA, Unit* objectB)
    // Calculate the result of object collision in 3D space.
{
    FixedArray<double, 3> centerOfMassA = objectA->centerOfMass();
    FixedArray<double, 3> centerOfMassB = objectB->centerOfMass();
    // ..
}

void accelerate(Unit* object, const FixedArray<double, 3>& force)
    // Calculate the position after applying a specified force for the
    // duration of a game tick.
{
    // ...
}

```

Note that the object files for the application components throughout the game project do not contain any of the implicitly instantiated definitions that we had chosen to uniquely sequester externally, i.e., within the `game_fixedarray.o` file:

```

$ nm -C app_geometryutil.o
000000000000003e T isOrthogonal(game::FixedArray<int, 2ul> const&,
    game::FixedArray<int, 2ul> const&)

```

```

00000000000000068 T isOrthogonal(game::FixedArray<int, 3ul> const&,
    game::FixedArray<int, 3ul> const&)
0000000000000000 T translate(game::Unit*, game::FixedArray<double, 2ul> const&)
0000000000000001f T translate(game::Unit*, game::FixedArray<double, 3ul> const&)
    U game::FixedArray<double, 2ul>::FixedArray()
    U game::FixedArray<double, 3ul>::FixedArray()
    U int game::dot<int, 2ul>(game::FixedArray<int, 2ul> const&,
game::FixedArray<int, 2ul> const&)
    U int game::dot<int, 3ul>(game::FixedArray<int, 3ul> const&,
game::FixedArray<int, 3ul> const&)

$ nm -C app_physics.o
00000000000000039 T accelerate(game::Unit*,
    game::FixedArray<double, 3ul> const&)
0000000000000000 T collide(game::Unit*, game::Unit*)
    U game::FixedArray<double, 3ul>::FixedArray()
0000000000000000 W game::Unit::centerOfMass()

```

Whether optimization involving **explicit-instantiation directives** reduces library sizes on disk has no noticeable effect or actually makes matters worse will depend on the particulars of the system at hand. Having this optimization applied to frequently used templates across a large organization has been known to decrease object file sizes, storage needs, link times, and overall build times, but see *Potential Pitfalls — Accidentally making matters worse* on page 373.

Insulating template definitions from clients

Even before the introduction of **explicit-instantiation declarations**, strategic use of **explicit-instantiation definitions** made it possible to **insulate** the *definition* of a template from client code, presenting instead just a limited set of instantiations against which clients may link. Such insulation enables the definition of the template to change without forcing clients to recompile. What's more, new explicit instantiations can be added without affecting existing clients.

As an example, suppose we have a single free-function template, `transform`, that operates on only floating-point values:

```

// transform.h:
#ifndef INCLUDED_TRANSFORM
#define INCLUDED_TRANSFORM

template <typename T> // declaration only of free-function template
T transform(const T& value);
    // Return the transform of the specified floating-point value.

#endif

```

Initially, this function template will support just two built-in types, **float** and **double**, but it is anticipated to eventually support the additional built-in type **long double** and perhaps even supplementary user-defined types (e.g., `Float128`) to be made available via separate headers (e.g., `float128.h`). By placing only the declaration of the `transform` function template in its component's header, clients will be able to link against only two supported explicit specializations provided in the `transform.cpp` file:

```
// transform.cpp:
#include <transform.h> // Ensure consistency with client-facing declaration.

template <typename T> // redeclaration/definition of free-function template
T transform(const T& value)
{
    // insulated implementation of transform function template
}

// explicit-instantiation definitions
template float transform(const float&); // Instantiate for type float.
template double transform(const double&); // Instantiate for type double.
```

Without the two `explicit-instantiation` declarations in the `transform.cpp` file above, its corresponding object file, `transform.o`, would be empty.

Note that, as of C++11, we *could* place the corresponding `explicit-instantiation` declarations in the header file for, say, documentation purposes:

```
// transform.h:
#ifndef INCLUDED_TRANSFORM
#define INCLUDED_TRANSFORM

template <typename T> // declaration only of free-function template
T transform(const T& value);
    // Return the transform of the specified floating-point value.

// explicit-instantiation declarations, available as of C++11
extern template float transform(const float&); // user documentation only;
extern template double transform(const double&); // has no effect whatsoever

#endif
```

Because no definition of the `transform` free-function template is visible in the header, no *implicit* instantiation can result from client use; hence, the two `explicit-instantiation` declarations above for **float** and **double**, respectively, do nothing.

Potential Pitfalls

Corresponding explicit-instantiation declarations and definitions

To realize a reduction in object-code size for individual translation units and yet still be able to link all valid programs successfully into a well-formed program, four moving parts have to be brought together correctly.

1. Each general template, `C<T>`, whose object code bloat is to be optimized must be declared within some designated component's header file, `c.h`.
2. The specific definition of each `C<T>` relevant to an explicit specialization being optimized — including general, partial-specialization, and full-specialization definitions — must appear in the header file prior to its corresponding **explicit-instantiation declaration**.
3. Each **explicit-instantiation declaration** for each specialization of each separate top-level — i.e., class, function, or variable — template must appear in the component's `.h` file *after* the corresponding general template declaration and the relevant general, partial-specialization, or full-specialization definition, but, in practice, always after *all* such definitions, not just the relevant one.
4. Each template specialization having an **explicit-instantiation declaration** in the header file must have a corresponding **explicit-instantiation definition** in the component's implementation file, `c.cpp`.

Absent items (1) and (2), clients would have no way to safely separate out the usability and inlineability of the template definitions yet consolidate the otherwise redundantly generated object-level definitions within just a single translation unit. Moreover, failing to provide the relevant definition would mean that any clients using one of these specializations would either fail to compile or, arguably worse, pick up the general definitions when a more specialized definition was intended, likely resulting in an ill-formed program.

Failing item (3), the object code for that particular specialization of that template will be generated locally in the client's translation unit as usual, negating any benefits with respect to local object-code size, irrespective of what is specified in the `c.cpp` file.

Finally, unless we provide a matching **explicit-instantiation definition** in the `c.cpp` file for each and every corresponding **explicit-instantiation declaration** in the `c.h` file as in item (4), our optimization attempts might well result in a library component that compiles, links, and even passes some unit tests but, when released to our clients, fails to link. Additionally, any **explicit-instantiation definition** in the `c.cpp` file that is not accompanied by a corresponding

explicit-instantiation declaration in the `c.h` file will inflate the size of the `c.o` file with no possibility of reducing code bloat in client code:

```
// c.h:
#ifndef INCLUDED_C // internal include guard
#define INCLUDED_C

template <typename T> void f(T v) { /*...*/ } // general template definition

extern template void f<int>(int v); // OK, matched in c.cpp
extern template void f<char>(char c); // Error, unmatched in .cpp file

#endif

// c.cpp:
#include <c.h> // incorporate own header first

template void f<int>(int v); // OK, matched in c.h
template void f<double>(double v); // Bug, unmatched in c.h file

// client.cpp:
#include <c.h>

void client()
{
    int i = 1;
    char c = 'a';
    double d = 2.0;

    f(i); // OK, matching explicit-instantiation directives
    f(c); // Link-Time Error, no matching explicit-instantiation definition
    f(d); // Bug, size increased due to no matching explicit-instantiation
        // declaration.
}

```

In the example above, `f(i)` works as expected, with the linker finding the definition of `f<int>` in `c.o`; `f(c)` fails to link because no definition of `f<char>` is guaranteed to be found anywhere; and `f(d)` accidentally works by silently generating a *redundant* local copy of `f<double>` in `client.o`, while another, identical definition is generated explicitly in `c.o`. These extra instantiations do not result in multiply-defined symbols because they still reside in their own **sections** and are marked as *weak* symbols. Importantly, note that **extern template** has *absolutely no effect* on overload resolution because the call to `f(c)` did *not* resolve to `f<int>`.

Accidentally making matters worse

When making the decision to explicitly instantiate common specializations of popular templates within some designated object file, it is important to consider that not all programs necessarily need every (or even any) such instantiation. Classes that have many member functions but typically use only a few require special attention.

For such classes, it might be beneficial to explicitly instantiate individual member functions instead of the entire class template. However, selecting *which* member functions to explicitly instantiate and with *which* template arguments they should be instantiated without carefully measuring the effect on the overall object size might result in not only overall pessimization, but also to an unnecessary maintenance burden. Finally, remember that one might need to explicitly tell the linker to strip unused sections resulting, for example, from forced instantiation of common template specializations, to avoid inadvertently bloating executables, which could adversely affect load times.

Annoyances

No good place to put definitions for unrelated classes

When we consider the implications of physical dependency,^{5,6} determining in which component to deposit the specialized definitions can be problematic. For example, consider a codebase implementing a core library that provides both a nontemplated `String` class and a `Vector` container class template. These fundamentally unrelated entities would ideally live in separate physical **components** (i.e., `.h/.cpp` pairs), neither of which depends physically on the other. That is, an application using just one of these components could be compiled, linked, tested, and deployed entirely independently of the other. Now, consider a large codebase that makes heavy use of `Vector<String>`: In what component should the object-code-level definitions for the `Vector<String>` specialization reside?⁷ There are two obvious alternatives.

1. `vector` — In this case, `vector.h` would hold **extern template class** `Vector<String>`; — the **explicit-instantiation** declaration. `vector.cpp` would hold **template class** `Vector<String>`; — the **explicit-instantiation** definition. With this approach, we would create a physical dependency of the `vector` component on `string`. Any client program wanting to use a `Vector` would also depend on `string` regardless of whether it was needed.

⁵See [lakos96](#).

⁶See [lakos20](#).

⁷Note that the problem of determining in which component to instantiate the object-level implementation of a template for a user-defined type is similar to that of specializing an arbitrary user-defined trait for a user-defined type.

2. `string` — In this case, `string.h` and `string.cpp` would instead be modified so as to depend on `vector`. Clients wanting to use a `string` would also be forced to depend physically on `vector` *at compile time*.

Another possibility might be to create a third component, called `stringvector`, that itself depends on both `vector` and `string`. By **escalating**⁸ the mutual dependency to a higher level in the physical hierarchy, we avoid forcing any client to depend on more than what is actually needed. The practical drawback to this approach is that only those clients that proactively include the composite `stringvector.h` header would realize any benefit; fortunately, in this case, there is no **one-definition rule (ODR)** violation if they don't.

Finally, complex machinery could be added to both `string.h` and `vector.h` to conditionally include `stringvector.h` whenever both of the other headers are included; such heroic efforts would, nonetheless, involve a **cyclic physical dependency** among all three of these components. Circular intercomponent collaborations are best avoided.⁹

All members of an explicitly defined template class must be valid

In general, when using a class template, only those members that are actually used get implicitly instantiated. This hallmark allows class templates to provide functionality for parameter types having certain capabilities, e.g., default constructible, while also providing partial support for types lacking those same capabilities. When providing an **explicit-instantiation definition**, however, *all* members of a class template are instantiated.

Consider a simple class template having a data member that can be either default-initialized via the template's default constructor or initialized with an instance of the member's type supplied at construction:

```
template <typename T>
class W
{
    T d_t; // a data member of type T

public:
    W() : d_t() {}
        // Create an instance of W with a default-constructed T member.

    W(const T& t) : d_t(t) {}
        // Create an instance of W with a copy of the specified t.

    void doStuff() { /* do stuff */ }
};
```

This class template can be used successfully with a type, such as `U` in the following code snippet, that is not default constructible:

⁸Iakos20, section 3.5.2, “Escalation,” pp. 604–614

⁹Iakos20, section 3.4, “Avoiding Cyclic Link-Time Dependencies,” pp. 592–601

```

struct U
{
    U(int i) { /* construct from i */ }
    // ...
};

void useWU()
{
    W<U> wu1(U(17)); // OK, using copy constructor for U
    wu1.doStuff();
}

```

As it stands, the code above is well formed even though `W<U>::W()` would fail to compile if instantiated. Consequently, although providing an `explicit-instantiation` declaration for `W<U>` is valid, a corresponding `explicit-instantiation` definition for `W<U>` fails to compile, as would an implicit instantiation of `W<U>::W()`:

```

extern template class W<U>; // Valid: Suppress implicit instantiation of W<U>.

template class W<U>;      // Error, U::U() not available for W<U>::W()

void useWU0()
{
    W<U> wu0;              // Error, U::U() not available for W<U>::W()
}

```

Unfortunately, the only workaround to achieve a comparable reduction in code bloat is to provide `explicit-instantiation` directives for each valid member function of `W<U>`, an approach that would likely carry a significantly greater maintenance burden:

```

extern template W<U>::W(const U& u); // suppress individual member
extern template void W<U>::doStuff(); // " " "
// ... Repeat for all other functions in W except W<U>::W().

template W<U>::W(const U& u); // instantiate individual member
template void W<U>::doStuff(); // " " "
// ... Repeat for all other functions in W except W<U>::W().

```

The power and flexibility to make it all work — albeit annoyingly — are there nonetheless.

See Also

- “Variable Templates” (§1.2, p. 157) covers an extension of the template syntax for defining a family of like-named variables or static data members that can be instantiated explicitly.

Further Reading

- For a different perspective on this feature, see **lakos20**, section 1.3.16, “extern Templates,” pp. 183–185.
- For a more complete discussion of how compilers and linkers work with respect to C++, see **lakos20**, Chapter 1, “Compilers, Linkers, and Components,” pp. 123–268.

Transparently Nested Namespaces

An **inline namespace** is a nested namespace whose member entities closely behave as if they were declared directly within the enclosing namespace.

Description

To a first approximation, an **inline namespace** (e.g., `v2` in the code snippet below) acts a lot like a conventional nested namespace (e.g., `v1`) followed by a **using** directive for that namespace in its enclosing namespace¹:

```
// example.cpp:
namespace n
{
    namespace v1 // conventional nested namespace followed by using directive
    {
        struct T { }; // nested type declaration (identified as ::n::v1::T)
        int d; // ::n::v1::d at, e.g., 0x01a64e90
    }

    using namespace v1; // Import names T and d into namespace n.
}

namespace n
{
    inline namespace v2 // similar to being followed by using namespace v2
    {
        struct T { }; // nested type declaration (identified as ::n::v2::T)
        int d; // ::n::v2::d at, e.g., 0x01a64e94
    }

    // using namespace v2; // redundant when used with an inline namespace
}

```

¹C++17 allows developers to concisely declare nested namespaces with shorthand notation:

```
namespace a::b { /*...*/ }
// is the same as
namespace a { namespace b { /*...*/ } }
```

C++20 expands on the above syntax by allowing the insertion of the **inline** keyword in front of any of the namespaces, except the first one:

```
namespace a::inline b::inline c { /*...*/ }
// is the same as
namespace a { inline namespace b { inline namespace c { /*...*/ } } }
```

```
inline namespace a::b { } // Error, cannot start with inline for compound namespace names
namespace inline a::b { } // Error, inline at front of sequence explicitly disallowed
```

Four subtle details distinguish these approaches.

1. Name collisions with existing names behave differently due to differing name-lookup rules.
2. **Argument-dependent lookup (ADL)** gives special treatment to **inline** namespaces.
3. Template specializations can refer to the primary template in an **inline** namespace even if written in the enclosing namespace.
4. Reopening namespaces might reopen an **inline** namespace.

One important aspect that all forms of namespaces share, however, is that (1) nested symbolic names (e.g., `n::v1::T`) at the **API** level, (2) **mangled names** (e.g., `_ZN1n2v11dE`, `_ZN1n2v21dE`), and (3) assigned relocatable addresses (e.g., `0x01a64e90`, `0x01a64e94`) at the **ABI** level remain unaffected by the use of either **inline** or **using** or both. To be precise, source files containing, alternately, `namespace n { inline namespace v { int d; } }` and `namespace n { namespace v { int d; } using namespace v; }`, will produce identical assembly.² Note that a **using** directive immediately following an **inline** namespace is superfluous; name lookup will always consider names in **inline** namespaces before those imported by a **using** directive. Such a directive can, however, be used to import the contents of an **inline** namespace to some other namespace, albeit only in the conventional, **using directive** sense; see *Annoyances — Only one namespace can contain any given inline namespace* on page 1082.

More generally, each namespace has what is called its **inline namespace set**, which is the transitive closure of all **inline** namespaces within the namespace. All names in the **inline namespace set** are roughly intended to behave as if they are defined in the enclosing namespace. Conversely, each **inline** namespace has an *enclosing namespace set* that comprises all enclosing namespaces up to and including the first non**inline** namespace.

Loss of access to duplicate names in enclosing namespace

When both a type and a variable are declared with the same name in the same scope, the variable name hides the type name — such behavior can be demonstrated by using the form of **sizeof** that accepts a nonparenthesized *expression* (recall that the form of **sizeof** that accepts a *type* as its argument requires parentheses):

```
struct A { double d; }; static_assert(sizeof( A) == 8, ""); // type
                          // static_assert(sizeof A == 8, ""); // Error

int A;                    static_assert(sizeof( A) == 4, ""); // data
                          static_assert(sizeof A == 4, ""); // OK
```

²These mangled names can be seen with GCC by running `g++ -S <file>.cpp` and viewing the contents of the generated `<file>.s`. Note that Compiler Explorer is another valuable tool for learning about what comes out the other end of a C++ compiler: see <https://godbolt.org/>.

Unless both type and variable entities are declared within the same scope, no preference is given to variable names; the name of an entity in an inner scope hides a like-named entity in an enclosing scope:

```
void f()
{
    double B;          static_assert(sizeof(B) == 8, ""); // variable
    {
        struct B { int d; }; static_assert(sizeof(B) == 8, ""); // variable
        static_assert(sizeof(B) == 4, ""); // type
    }
    static_assert(sizeof(B) == 8, ""); // variable
}
```

When an entity is declared in an enclosing **namespace** and another entity having the same name hides it in a *lexically* nested scope, then (apart from **inline** namespaces) access to a hidden element can generally be recovered by using scope resolution:

```
struct C { double d; }; static_assert(sizeof( C) == 8, "");

void g()
{
    static_assert(sizeof( C) == 8, ""); // type
    int C;          static_assert(sizeof( C) == 4, ""); // variable
    static_assert(sizeof(::C) == 8, ""); // type
}
static_assert(sizeof( C) == 8, ""); // type
```

A conventional nested namespace behaves as one might expect:

```
namespace outer
{
    struct D { double d; }; static_assert(sizeof( D) == 8, ""); // type

    namespace inner
    {
        static_assert(sizeof( D) == 8, ""); // type
        int D;          static_assert(sizeof( D) == 4, ""); // var
    }
    static_assert(sizeof( D) == 8, ""); // type
    static_assert(sizeof(inner::D) == 4, ""); // var
    static_assert(sizeof(outer::D) == 8, ""); // type
    using namespace inner; //static_assert(sizeof( D) == 0, ""); // Error
    static_assert(sizeof(inner::D) == 4, ""); // var
    static_assert(sizeof(outer::D) == 8, ""); // type
}
static_assert(sizeof(outer::D) == 8, ""); // type
```

In the example above, the inner variable name, `D`, hides the outer type with the same name, starting from the point of `D`'s declaration in `inner` until `inner` is closed, after which the unqualified name `D` reverts to the type in the outer namespace. Then, right after the subsequent `using namespace inner;` directive, the meaning of the unqualified name `D` in `outer` becomes ambiguous, shown here with a `static_assert` that is commented out; any attempt to refer to an unqualified `D` from here to the end of the scope of `outer` will fail to compile. The type entity declared as `D` in the outer namespace can, however, still be

accessed — from inside or outside of the `outer` namespace, as shown in the example — via its qualified name, `outer::D`.

If an **inline** namespace were used instead of a nested namespace followed by a **using** directive, however, the ability to recover by name the hidden entity in the enclosing namespace is lost. Unqualified name lookup considers the inline namespace set and the used namespace set simultaneously. Qualified name lookup first considers the **inline** namespace set and *then* goes on to look into used namespaces. These lookup rules mean we can still refer to `outer::D` in the example above, but doing so would still be ambiguous if `inner` were an inline namespace. This subtle difference in behavior is a byproduct of the highly specific use case that motivated this feature and for which it was explicitly designed; see *Use Cases — Link-safe ABI versioning* on page 1067.

Argument-dependent-lookup interoperability across inline namespace boundaries

Another important aspect of **inline** namespaces is that they allow ADL to work seamlessly across **inline** namespace boundaries. Whenever unqualified function names are being resolved, a list of *associated namespaces* is built for each argument of the function. This list of associated namespaces comprises the namespace of the argument, its enclosing namespace set, plus the **inline** namespace set.

Consider the case of a type, `U`, defined in an `outer` namespace, and a function, `f(U)`, declared in an `inner` namespace nested within `outer`. A second type, `V`, is defined in the `inner` namespace, and a function, `g`, is declared, after the close of `inner`, in the `outer` namespace:

```
namespace outer
{
    struct U { };

    // inline // Uncommenting this line fixes the problem.
    namespace inner
    {
        void f(U) { }
        struct V { };
    }

    using namespace inner; // If we inline inner, we don't need this line.

    void g(V) { }
}

void client()
{
    f(outer::U()); // Error, f is not declared in this scope.
    g(outer::inner::V()); // Error, g is not declared in this scope.
}
```

In the example above, a `client` invoking `f` with an object of type `outer::U` fails to compile because `f(outer::U)` is declared in the nested `inner` namespace, which is not the same as declaring it in `outer`. Because ADL does not look into namespaces added with the `using` directive, ADL does not find the needed `outer::inner::f` function. Similarly, the type `V`, defined in namespace `outer::inner`, is not declared in the same namespace as the function `g` that operates on it. Hence, when `g` is invoked from within `client` on an object of type `outer::inner::V`, ADL again does not find the needed function `outer::g(outer::V)`.

Simply making the `inner` namespace **inline** solves both of these ADL-related problems. All transitively nested **inline** namespaces — up to and including the most proximate non-**inline** enclosing namespace — are treated as one with respect to ADL.

The ability to specialize templates declared in a nested inline namespace

The third property that distinguishes **inline** namespaces from conventional ones, even when followed by a `using` directive, is the ability to specialize a class template defined within an **inline** namespace from within an enclosing one; this ability holds transitively up to and including the most proximate non**inline** namespace:

```
namespace out                                // proximate noninline outer namespace
{
    inline namespace in1                     // first-level nested inline namespace
    {
        inline namespace in2               // second-level nested inline namespace
        {
            template <typename T>          // primary class template general definition
            struct S { };

            template <>                     // class template full specialization
            struct S<char> { };
        }

        template <>                         // class template full specialization
        struct S<short> { };
    }

    template <>                             // class template full specialization
    struct S<int> { };
}

using namespace out;                         // conventional using directive

template <>
struct S<int> { };                           // Error, cannot specialize from this scope
```

Note that the conventional nested namespace `out` followed by a `using` directive in the enclosing namespace does not admit specialization from that outermost namespace, whereas

all of the **inline** namespaces do. Function templates behave similarly except that — unlike class templates, whose definitions must reside entirely within the namespace in which they are declared — a function template can be *declared* within a nested namespace and then be *defined* from anywhere via a **qualified name**:

```
namespace out // proximate noninline outer namespace
{
    inline namespace in1 // first-level nested inline namespace
    {
        template <typename T> // function template declaration
        void f();

        template <> // function template (full) specialization
        void f<short>() { }
    }

    template <> // function template (full) specialization
    void f<int>() { }
}

template <typename T> // function template general definition
void out::in1::f() { }
```

An important takeaway from the examples above is that every template entity — be it class or function — *must* be declared in *exactly* one place within the collection of namespaces that comprise the **inline** namespace set. In particular, declaring a class template in a nested **inline** namespace and then subsequently defining it in a containing namespace is not possible because, unlike a function definition, a type definition cannot be placed into a namespace via name qualification alone:

```
namespace outer
{
    inline namespace inner
    {
        template <typename T> // class template declaration
        struct Z; // (if defined, must be within same namespace)

        template <> // class template full specialization
        struct Z<float> { };
    }

    template <typename T> // inconsistent declaration (and definition)
    struct Z { }; // Z is now ambiguous in namespace outer.

    const int i = sizeof(Z<int>); // Error, reference to Z is ambiguous.

    template <> // attempted class template full specialization
```

```

    struct Z<double> { };           // Error, outer::Z or outer::inner::Z?
}

```

Reopening namespaces can reopen nested inline ones

Another subtlety specific to **inline** namespaces is related to reopening namespaces. Consider a namespace `outer` that declares a nested namespace `outer::m` and an **inline** namespace `inner` that, in turn, declares a nested namespace `outer::inner::m`. In this case, subsequent attempts to reopen namespace `m` cause an ambiguity error:

```

namespace outer
{
    namespace m { }           // opens and closes ::outer::m

    inline namespace inner
    {
        namespace n { }     // opens and closes ::outer::inner::n
        namespace m { }     // opens and closes ::outer::inner::m
    }

    namespace n             // OK, reopens ::outer::inner::n
    {
        struct S { };      // defines ::outer::inner::n::S
    }

    namespace m             // Error, namespace m is ambiguous.
    {
        struct T { };      // with clang defines ::outer::m::T
    }
}

static_assert(std::is_same<outer::n::S, outer::inner::n::S>::value, "");

```

In the code snippet above, no issue occurs with reopening `outer::inner::n` and no issue would have occurred with reopening `outer::m` but for the `inner` namespaces having been declared **inline**. When a new namespace declaration is encountered, a lookup determines if a matching namespace having that name appears anywhere in the *inline namespace set* of the current namespace. If the namespace is ambiguous, as is the case with `m` in the example above, one can get the surprising error shown.³ If a matching namespace is found

³Note that reopening already declared namespaces, such as `m` and `n` in the `inner` and `outer` example, is handled incorrectly on several popular platforms. Clang, for example, performs a name lookup when encountering a new namespace declaration and give preference to the outermost namespace found, causing the last declaration of `m` to reopen `::outer::m` instead of being ambiguous. GCC, prior to 8.1 (c. 2018), does not perform name lookup and will place *any* nested namespace declarations directly within their enclosing namespace. This defect causes the last declaration of `m` to reopen `::outer::m` instead of `::outer::inner::m` and the last declaration of `n` to open a new namespace, `::outer::n`, instead of reopening `::outer::inner::n`.

unambiguously inside an **inline** namespace, `n` in this case, then it is that nested namespace that is reopened — here, `::outer::inner::n`. The inner namespace is reopened even though the last declaration of `n` is not lexically scoped within `inner`. Notice that the definition of `S` is perhaps surprisingly defining `::outer::inner::n::S`, not `::outer::n::S`. For more on what is *not* supported by this feature, see *Annoyances — Inability to redeclare across namespaces impedes code factoring* on page 1079.

Use Cases

Facilitating API migration

Getting a large codebase to *promptly* upgrade to a new version of a library in any sort of timely fashion can be challenging. As a simplistic illustration, imagine that we have just developed a new library, `parselib`, comprising a class template, `Parser`, and a function template, `analyze`, that takes a `Parser` object as its only argument:

```
namespace parselib
{
    template <typename T>
    class Parser
    {
        // ...

    public:
        Parser();
        int parse(T* result, const char* input);
        // Load result from null-terminated input; return 0 (on
        // success) or nonzero (with no effect on result).
    };

    template <typename T>
    double analyze(const Parser<T>& parser);
}

```

To use our library, clients will need to specialize our `Parser` class directly within the `parselib` namespace:

```
struct MyClass { /*...*/ }; // end-user-defined type

namespace parselib // necessary to specialize Parser
{
    template <> // Create full specialization of class
    class Parser<MyClass> // Parser for user-type MyClass.
    {
        // ...
    }
}

```

```

    public:
        Parser();
        int parse(MyClass* result, const char* input);
            // The contract for a specialization typically remains the same.
};

double analyze(const Parser<MyClass>& parser);
}

```

Typical client code will also look for the `Parser` class directly within the `parselib` namespace:

```

void client()
{
    MyClass result;
    parselib::Parser<MyClass> parser;

    int status = parser.parse(&result, "...( MyClass value )...");
    if (status != 0)
    {
        return;
    }

    double value = analyze(parser);
    // ...
}

```

Note that invoking `analyze` on objects of some instantiated type of the `Parser` class template will rely on ADL to find the corresponding overload.

We anticipate that our library's API will evolve over time, so we want to enhance the design of `parselib` accordingly. One of our goals is to somehow encourage clients to move essentially all at once, yet also to accommodate both the early adopters and the inevitable stragglers that make up a typical adoption curve. Our approach will be to create, within our outer `parselib` namespace, a nested **inline** namespace, `v1`, which will hold the current implementation of our library software:

```

namespace parselib
{
    inline namespace v1 // Note our use of inline namespace here.
    {
        template <typename T>
        class Parser
        {
            // ...

```

```

public:
    Parser();
    int parse(T* result, const char* input);
        // Load result from null-terminated input; return 0 (on
        // success) or nonzero (with no effect on result).
};

template <typename T>
double analyze(const Parser<T>& parser);
}
}

```

As suggested by the name `v1`, this namespace serves primarily as a mechanism to support library evolution through API and ABI versioning (see *Link-safe ABI versioning* on page 1067 and *Build modes and ABI link safety* on page 1071). The need to specialize `class Parser` and, independently, the reliance on ADL to find the free function template `analyze` require the use of `inline` namespaces, as opposed to a conventional namespace followed by a `using` directive.

Note that, whenever a subsystem starts out directly in a first-level namespace and is subsequently moved to a second-level nested namespace for the purpose of versioning, declaring the inner namespace `inline` is the most reliable way to avoid inadvertently destabilizing existing clients; see also *Enabling selective using directives for short-named entities* on page 1074.

Now suppose we decide to enhance `parselib` in a non-backwards-compatible manner, such that the signature of `parse` takes a second argument `size` of type `std::size_t` to allow parsing of non-null-terminated strings and to reduce the risk of buffer overruns. Instead of unilaterally removing all support for the previous version in the new release, we can create a second namespace, `v2`, containing the new implementation and then, at some point, make `v2` the `inline` namespace instead of `v1`:

```

#include <cstdlib> // std::size_t

namespace parselib
{
    namespace v1 // Notice that v1 is now just a nested namespace.
    {
        template <typename T>
        class Parser
        {
            // ...

        public:
            Parser();
            int parse(T* result, const char* input);
        };
    };
};

```

```

        // Load result from null-terminated input; return 0 (on
        // success) or nonzero (with no effect on result).
    };

    template <typename T>
    double analyze(const Parser<T>& parser);
}

inline namespace v2 // Notice that use of inline keyword has moved here.
{
    template <typename T>
    class Parser
    {
        // ...

    public: // Note incompatible change to Parser's essential API.
        Parser();
        int parse(T* result, const char* input, std::size_t size);
        // Load result from input of specified size; return 0
        // on success) or nonzero (with no effect on result).
    };

    template <typename T>
    double analyze(const Parser<T>& parser);
}
}

```

When we release this new version with `v2` made **inline**, all existing clients that rely on the version supported directly in `parselib` will, by design, break when they recompile. At that point, each client will have two options. The first one is to upgrade the code immediately by passing in the size of the input string (e.g., 23) along with the address of its first character:

```

void client()
{
    // ...
    int status = parser.parse(&result, "...( MyClass value )...", 23);
    // ...
}

```

^^^ Look here!

The second option is to change all references to `parselib` to refer to the original version in `v1` explicitly:

```

namespace parselib
{
    namespace v1 // specializations moved to nested namespace
    {

```

```

    template <>
    class Parser<MyClass>
    {
        // ...

    public:
        Parser();
        int parse(MyClass* result, const char* input);
    };

    double analyze(const Parser<MyClass>& parser);
}

void client1()
{
    MyClass result;
    parselib::v1::Parser<MyClass> parser; // reference nested namespace v1

    int status = parser.parse(&result, "...( MyClass value )...");
    if (status != 0)
    {
        return;
    }

    double value = analyze(parser);
    // ...
}

```

Providing the updated version in a new **inline** namespace **v2** provides a more flexible migration path — especially for a large population of independent client programs — compared to manual targeted changes in client code.

Although new users would pick up the latest version automatically either way, existing users of `parselib` will have the option of converting immediately by making a few small syntactic changes or opting to remain with the original version for a while longer by making all references to the library namespace refer explicitly to the desired version. If the library is released before the **inline** keyword is moved, early adopters will have the option of opting in by referring to **v2** explicitly until it becomes the default. Those who have no need for enhancements can achieve stability by referring to a particular version in perpetuity or until it is physically removed from the library source.

Although this same functionality can sometimes be realized without using **inline** namespaces (i.e., by adding a **using namespace** directive at the end of the `parselib` namespace), any benefit of ADL and the ability to specialize templates from within the enclosing `parselib` namespace itself would be lost. Note that, because specialization doesn't kick in until overload resolution is completed, specializing overloaded functions is dubious at

best; see *Potential Pitfalls — Relying on **inline** namespaces to solve library evolution* on page 1077.

Providing separate namespaces for each successive version has an additional advantage in an entirely separate dimension: avoiding inadvertent, difficult-to-diagnose, latent linkage defects. Though not demonstrated by this specific example, cases do arise where simply changing which of the version namespaces is declared **inline** might lead to an **ill formed, no-diagnostic required (IFNDR)** program. This issue might ensue when one or more of its translation units that use the library are not recompiled before the program is relinked to the new static or dynamic library containing the updated version of the library software; see *Link-safe ABI versioning* below.

For distinct nested namespaces to guard effectively against accidental link-time errors, the symbols involved have to (1) reside in object code (e.g., a **header-only library** would fail this requirement) and (2) have the same **name mangling** (i.e., linker symbol) in both versions. In this particular instance, however, the signature of the `parse` member function of `parser` did change, and its mangled name will consequently change as well; hence the same **undefined symbol link error** would result either way.

Link-safe ABI versioning

inline namespaces are not intended as a mechanism for source-code versioning; instead, they prevent programs from being **ill formed** due to linking some version of a library with client code compiled using some other, typically older version of the same library. Below, we present two examples: a simple pedagogical example to illustrate the principle followed by a more real-world example. Suppose we have a library component `my_thing` that implements an example type, `Thing`, which wraps an **int** and initializes it with some value in its default constructor defined out-of-line in the `cpp` file:

```
struct Thing // version 1 of class Thing
{
    int i;    // integer data member (size is 4)
    Thing(); // original noninline constructor (defined in .cpp file)
};
```

Compiling a source file with this version of the header included might produce an object file that can be incompatible yet linkable with an object file resulting from compiling a different source file with a different version of this header included:

```
struct Thing // version 2 of class Thing
{
    double d; // double-precision floating-point data member (size is 8)
    Thing(); // updated noninline constructor (defined in .cpp file)
};
```

To make the problem that we are illustrating concrete, let's represent the client as a `main` program that does nothing but create a `Thing` and print the value of its only data member, `i`.

```
// main.cpp:
#include <my_thing.h> // my::Thing (version 1)
#include <iostream>   // std::cout

int main()
{
    my::Thing t;
    std::cout << t.i << '\n';
}
```

If we compile this program, a reference to a locally undefined linker symbol, such as `_ZN2my7impl_v15ThingC1Ev`,⁴ which represents the `my::Thing::Thing` constructor, will be generated in the `main.o` file:

```
$ g++ -c main.cpp
```

Without explicit intervention, the spelling of this linker symbol would be unaffected by any subsequent changes made to the implementation of `my::Thing`, such as its data members or implementation of its default constructor, even after recompiling. The same, of course, applies to its definition in a separate translation unit.

We now turn to the translation unit implementing type `my::Thing`. The `my_thing` **component** consists of a `.h/.cpp` pair: `my_thing.h` and `my_thing.cpp`. The header file `my_thing.h` provides the physical interface, such as the definition of the principal type, `Thing`, its member and associated free function declarations, plus definitions for inline functions and function templates, if any:

```
// my_thing.h:
#ifndef INCLUDED_MY_THING
#define INCLUDED_MY_THING

namespace my // outer namespace (used directly by clients)
{
    inline namespace impl_v1 // inner namespace (for implementer use only)
    {
        struct Thing
        {
            int i; // original data member, size = 4
            Thing(); // default constructor (defined in my_thing.cpp)
        };
    }
}
```

⁴On a Unix machine, typing `nm main.o` reveals the symbols used in the specified object file. A symbol prefaced with a capital `U` represents an undefined symbol that must be resolved by the linker. Note that the linker symbol shown here incorporates an intervening `inline` namespace, `impl_v1`, as will be explained shortly.

```

}

#endif

```

The implementation file `my_thing.cpp` contains all of the non**inline** function bodies that will be translated separately into the `my_thing.o` file:

```

// my_thing.cpp:
#include <my_thing.h>

namespace my                // outer namespace (used directly by clients)
{
    inline namespace impl_v1 // inner namespace (for implementer use only)
    {
        Thing::Thing() : i(0) // Load a 4-byte value into Thing's data member.
        {
        }
    }
}

```

Observing common good practice, we include the header file of the component as the first substantive line of code to ensure that — irrespective of anything else — the header always compiles in isolation, thereby avoiding insidious include-order dependencies.⁵ When we compile the source file `my_thing.cpp`, we produce an object file `my_thing.o` containing the definition of the same linker symbol, such as `_ZN2my7impl_v15ThingC1Ev`, for the default constructor of `my::Thing` needed by the client:

```
$ g++ -c my_thing.cpp
```

We can then link `main.o` and `my_thing.o` into an executable and run it:

```
$ g++ -o prog main.o my_thing.o
$ ./prog
```

```
0
```

Now, suppose we were to change the definition of `my::Thing` to hold a **double** instead of an **int**, recompile `my_thing.cpp`, and then relink with the original `main.o` without recompiling `main.cpp` first. None of the relevant linker symbols would change, and the code would recompile and link just fine, but the resulting binary `prog` would be IFNDR: the client would be trying to print a 4-byte, **int** data member, `i`, in `main.o` that was loaded by the library component as an 8-byte, **double** into `d` in `my_thing.o`. We can resolve this problem by changing — or, if we didn't think of it in advance, by adding — a new **inline** namespace and making that change there:

⁵See [lakos20](#), section 1.6.1, “Component Property 1,” pp. 210–212.

```

// my_thing.cpp:
#include <my_thing.h>

namespace my // outer namespace (used directly by clients)
{
    inline namespace impl_v2 // inner namespace (for implementer use only)
    {
        Thing::Thing() : d(0.0) // Load 8-byte value into Thing's data member.
        {
        }
    }
}

```

Now clients that attempt to link against the new library will not find the linker symbol, such as `_Z...impl_v1...v`, and the link stage will fail. Once clients recompile, however, the undefined linker symbol will match the one available in the new `my_thing.o`, such as `_Z...impl_v2...v`, the link stage will succeed, and the program will again work as expected. What's more, we have the option of keeping the original implementation. In that case, existing clients that have not as yet recompiled will continue to link against the old version until it is eventually removed after some suitable deprecation period.

As a more realistic second example of using **inline** namespaces to guard against linking incompatible versions, suppose we have two versions of a `Key` class in a security library in the enclosing namespace, `auth` — the original version in a regular nested namespace `v1`, and the new current version in an **inline** nested namespace `v2`:

```

#include <cstdint> // std::uint32_t, std::uint64_t

namespace auth // outer namespace (used directly by clients)
{
    namespace v1 // inner namespace (optionally used by clients)
    {
        class Key
        {
        private:
            std::uint32_t d_key;
            // sizeof(Key) is 4 bytes.

        public:
            std::uint32_t key() const; // stable interface function

            // ...
        };
    }

    inline namespace v2 // inner namespace (default current version)
    {
        class Key

```

```

    {
    private:
        std::uint64_t d_securityHash;
        std::uint32_t d_key;
        // sizeof(Key) is 16 bytes.

    public:
        std::uint32_t key() const; // stable interface function

        // ...
    };
}

```

Attempting to link together older binary artifacts built against version 1 with binary artifacts built against version 2 will result in a link-time error rather than allowing an ill formed program to be created. Note, however, that this approach works only if functionality essential to typical use is defined out of line in a `.cpp` file. For example, it would add absolutely no value for libraries that are shipped entirely as header files, since the versioning offered here occurs strictly at the binary level (i.e., between object files) during the link stage.

Build modes and ABI link safety

In certain scenarios, a class might have two different memory layouts depending on compilation flags. For instance, consider a low-level `ManualBuffer` class template in which an additional data member is added for debugging purposes:

```

template <typename T>
struct ManualBuffer
{
private:
    alignas(T) char d_data[sizeof(T)]; // aligned and big enough to hold a T

#ifdef NDEBUG
    bool d_engaged; // tracks whether buffer is full (debug builds only)
#endif

public:
    void construct(const T& obj);
        // Emplace obj. (Engage the buffer.) The behavior is undefined unless
        // the buffer was not previously engaged.

    void destroy();
        // Destroy the current obj. (Disengage the buffer.) The behavior is
        // undefined unless the buffer was previously engaged.

    // ...
};

```

Note that we have employed the C++11 **alignas** attribute (see Section 2.1. “**alignas**” on page 168) here because it is exactly what’s needed for this usage example.

The `d_engaged` flag in the example above serves as a way to detect misuse of the `ManualBuffer` class but only in debug builds. The extra space and run time required to maintain this Boolean flag is undesirable in a release build because `ManualBuffer` is intended to be an efficient, lightweight abstraction over the direct use of **placement new** and explicit destruction.

The linker symbol names generated for the methods of `ManualBuffer` are the same irrespective of the chosen build mode. If the same program links together two object files where `ManualBuffer` is used — one built in debug mode and one built in release mode — the **one-definition rule (ODR)** will be violated, and the program will again be IFNDR.

Prior to **inline namespaces**, it was possible to control the ABI-level name of linked symbols by creating separate template instantiations on a per-build-mode basis:

```
#ifndef NDEBUG
enum { is_debug_build = 1 };
#else
enum { is_debug_build = 0 };
#endif

template <typename T, bool Debug = is_debug_build>
struct ManualBuffer { /*...*/ };
```

While the code above changes the interface of `ManualBuffer` to accept an additional template parameter, it also allows debug and release versions of the same class to coexist in the same program, which might prove useful, e.g., for testing.

Another way of avoiding incompatibilities at link time is to introduce two **inline namespaces**, the entire purpose of which is to change the ABI-level names of the linker symbols associated with `ManualBuffer` depending on the build mode:

```
#ifndef NDEBUG // perhaps a BAD IDEA
inline namespace release
#else
inline namespace debug
#endif
{
    template <typename T>
    struct ManualBuffer
    {
        // ... (same as above)
    };
}
```

The approach demonstrated in this example tries to ensure that a linker error will occur if any attempt is made to link objects built with a build mode different from that of

manualbuffer.o. Tying it to the NDEBUG flag, however, might have unintended consequences; we might introduce unwanted restrictions in what we call **mixed-mode builds**. Most modern platforms support the notion of linking a collection of object files irrespective of their optimization levels. The same is certainly true for whether or not C-style `assert` is enabled. In other words, we might want to have a mixed-mode build where we link object files that differ in their optimization and assertion options, as long as they are binary compatible — i.e., in this case, they all must be uniform with respect to the implementation of `ManualBuffer`. Hence, a more general, albeit more complicated and manual, approach would be to tie the noninteroperable behavior associated with this “safe” or “defensive” build mode to a different switch entirely. Another consideration would be to avoid ever inlining a namespace into the global namespace since no method is available to recover a symbol when there is a collision:

```
namespace buflib // GOOD IDEA: enclosing namespace for nested inline namespace
{
#ifdef SAFE_MODE // GOOD IDEA: separate control of non-interoperable versions
    inline namespace safe_build_mode
#else
    inline namespace normal_build_mode
#endif
    {
        template <typename T>
        struct ManualBuffer
        {
        private:
            alignas(T) char d_data[sizeof(T)]; // aligned/sized to hold a T

#ifdef SAFE_MODE
            bool d_engaged; // tracks whether buffer is full (safe mode only)
#endif

        public:
            void construct(const T& obj); // sets d_engaged (safe mode only)
            void destroy(); // sets d_engaged (safe mode only)
            // ...
        };
    }
}
```

And, of course, the appropriate conditional compilation within the function bodies would need to be in the corresponding `.cpp` file.

Finally, if we have two implementations of a particular entity that are sufficiently distinct, we might choose to represent them in their entirety, controlled by their own bespoke conditional-compilation switches, as illustrated here using the `my::VersionedThing` type (see *Link-safe ABI versioning* on page 1067):

```
// my_versionedthing.h:
#ifndef INCLUDED_MY_VERSIONEDTHING
#define INCLUDED_MY_VERSIONEDTHING

namespace my
{
#ifdef MY_THING_VERSION_1 // bespoke switch for this component version
    inline
#endif
    namespace v1
    {
        struct VersionedThing
        {
            int d_i;
            VersionedThing();
        };
    }

#ifdef MY_THING_VERSION_2 // bespoke switch for this component version
    inline
#endif
    namespace v2
    {
        struct VersionedThing
        {
            double d_i;
            VersionedThing();
        };
    }
}
#endif
```

However, see *Potential Pitfalls—**inline**-namespace-based versioning doesn't scale* on page 1076.

Enabling selective using directives for short-named entities

Introducing a large number of small names into client code that doesn't follow rigorous nomenclature can be problematic. Hoisting these names into one or more nested namespaces so that they are easier to identify as a unit and can be used more selectively by clients, such as through explicit qualification or using directives, can sometimes be an effective way of organizing shared codebases. For example, `std::literals` and its nested namespaces, such as `chrono_literals`, were introduced as **inline** namespaces in C++14. As it turns out, clients of these nested namespaces have no need to specialize any templates defined in these namespaces nor do they define types that must be found through ADL, but one can at least imagine special circumstances in which such tiny-named entities are either templates that

require specialization or operator-like functions, such as `swap`, defined for local types within those nested namespaces. In those cases, **inline** namespaces would be required to preserve the desired “as if” properties.

Even without either of these two needs, another property of an **inline** namespace differentiates it from a non**inline** one followed by a **using** directive. Recall from *Description — Loss of access to duplicate names in enclosing namespace* on page 1056 that a name in an outer namespace will hide a duplicate name imported via a **using** directive, whereas any access to that duplicate name within the enclosing namespace would be ambiguous when that symbol is installed by way of an **inline** namespace. To see why this more forceful clobbering behavior might be preferred over hiding, suppose we have a communal namespace `abc` that is shared across multiple disparate headers. The first header, `abc_header1.h`, represents a collection of logically related small functions declared directly in `abc`:

```
// abc_header1.h:
namespace abc
{
    int i();
    int am();
    int smart();
}
```

A second header, `abc_header2.h`, creates a suite of many functions having tiny function names. In a perhaps misguided effort to avoid clobbering other symbols within the `abc` namespace having the same name, all of these tiny functions are sequestered within a nested namespace:

```
// abc_header2.h:
namespace abc
{
    namespace nested // Should this namespace have been inline instead?
    {
        int a(); // lots of functions with tiny names
        int b();
        int c();
        // ...
        int h();
        int i(); // might collide with another name declared in abc
        // ...
        int z();
    }

    using namespace nested; // becomes superfluous if nested is made inline
}
```

Now suppose that a client application includes both of these headers to accomplish some task:

```

// client.cpp:
#include <abc_header1.h>
#include <abc_header2.h>

int function()
{
    if (abc::smart() < 0) { return -1; } // uses smart() from abc_header1.h
    return abc::z() + abc::i() + abc::a() + abc::h() + abc::c(); // Oops!
    // Bug, silently uses the abc::i() defined in abc_header1.h
}

```

In trying to cede control to the client as to whether the declared or imported `abc::i()` function is to be used, we have, in effect, invited the defect illustrated in the above example whereby the client was expecting the `abc::i()` from `abc_header2.h` and yet picked up the one from `abc_header1.h` by default. Had the nested namespace in `abc_header2.h` been declared **inline**, the qualified name `abc::i()` would have automatically been rendered *ambiguous* in namespace `abc`, the translation would have failed *safely*, and the defect would have been exposed at compile time. The downside, however, is that no method would be available to recover nominal access to the `abc::i()` defined in `abc_header1.h` once `abc_header2.h` is included, even though the two functions (e.g., including their mangled names at the ABI level) remain distinct.

Potential Pitfalls

inline-namespace-based versioning doesn't scale

The problem with using **inline** namespaces for ABI link safety is that the protection they offer is only partial; in a few major places, critical problems can linger until run time instead of being caught at compile time.

Controlling which namespace is **inline** using macros, such as was done in the `my::VersionedThing` example in *Use Cases — Link-safe ABI versioning* on page 1067, will result in code that directly uses the unversioned name, `my::VersionedThing` being bound directly to the versioned name `my::v1::VersionedThing` or `my::v2::VersionedThing`, along with the class layout of that particular entity. Sometimes details of using the **inline** namespace member are not resolved by the linker, such as the object layout when we use types from that namespace as member variables in other objects:

```

// my_thingaggregate.h:

// ...
#include <my_versionedthing.h>
// ...

namespace my
{

```

```

struct ThingAggregate
{
    // ...
    VersionedThing d_thing;
    // ...
};
}

```

This new `ThingAggregate` type does not have the versioned **inline** namespace as part of its mangled name; it does, however, have a completely different layout if built with `MY_THING_VERSION_1` defined versus `MY_THING_VERSION_2` defined. Linking a program with mixed versions of these flags will result in runtime failures that are decidedly difficult to diagnose.

This same sort of problem will arise for functions taking arguments of such types; calling a function from code that is wrong about the layout of a particular type will result in stack corruption and other undefined and unpredictable behavior. This macro-induced problem will also arise in cases where an old object file is linked against new code that changes which namespace is **inlined** but still provides the definitions for the old version namespace. The old object file for the client can still link, but new object files using the headers for the old objects might attempt to manipulate those objects using the new namespace.

The only viable workaround for this approach is to propagate the **inline** namespace hierarchy through the entire software stack. Every object or function that uses `my::VersionedThing` needs to also be in a namespace that differs based on the same control macro. In the case of `ThingAggregate`, one could just use the same `my::v1` and `my::v2` namespaces, but higher-level libraries would need their own `my`-specific nested namespaces. Even worse, for higher-level libraries, every lower-level library having a versioning scheme of this nature would need to be considered, resulting in having to provide the full cross-product of nested namespaces to get link-time protection against mixed-mode builds.

This need for layers above a library to be aware of and to integrate into their own structure the same namespaces the library has removes all or most of the benefits of using **inline** namespaces for versioning. For an authentic real-world case study of heroic industrial use — and eventual disuse — of **inline**-namespaces for versioning, see *Appendix — Case study of using inline namespaces for versioning* on page 1083.

Relying on inline namespaces to solve library evolution

Inline namespaces might be misperceived as a complete solution for the owner of a library to evolve its API. As an especially relevant example, consider the C++ Standard Library, which itself does not use inline namespaces for versioning. Instead, to allow for its anticipated essential evolution, the Standard Library imposes certain special restrictions on what is permitted to occur within its own `std` namespace by dint of deeming certain problematic uses as either ill formed or otherwise engendering **undefined behavior**.

Since C++11, several restrictions related to the Standard Library were put in place.

- Users may not add any new declarations within namespace `std`, meaning that users cannot add new *functions*, *overloads*, *types*, or *templates* to `std`. This restriction gives the Standard Library freedom to add new *names* in future versions of the Standard.
- Users may not specialize member functions, member function templates, or member class templates. Specializing any of those entities might significantly inhibit a Standard Library vendor's ability to maintain its otherwise encapsulated implementation details.
- Users may add specializations of top-level Standard Library templates only if the declaration depends on the name of a nonstandard user-defined type and only if that user-defined type meets all requirements of the original template. Specialization of function templates is allowed but generally discouraged because this practice doesn't scale since function templates cannot be partially specialized. Specializing of standard class templates when the specialization names a nonstandard user-defined type, such as `std::vector<MyType*>`, is allowed but also problematic when not explicitly supported. While certain specific types, such as `std::hash`, are designed for user specialization, steering clear of the practice for any other type helps to avoid surprises.

Several other good practices facilitate smooth evolution for the Standard Library.⁶

- Avoid specializing variable templates, even if dependent on user-defined types, except for those variable templates where specialization is explicitly allowed.⁷
- Other than a few specific exceptions, avoiding the forming of pointers to Standard Library functions — either explicitly or implicitly — allows the library to add overloads, either as part of the Standard or as an implementation detail for a particular Standard Library, without breaking user code.⁸
- Overloads of Standard Library functions that depend on user-defined types are permitted, but, as with specializing Standard Library templates, users must still meet the requirements of the Standard Library function. Some functions, such as `std::swap`, are designed to be customization points via overloading, but leaving functions not specifically designed for this purpose to vendor implementations only helps to avoid surprises.

Finally, upon reading about this **inline** namespace feature, one might think that all names in namespace `std` could be made available at a global scope simply by inserting an

⁶These restrictions are normative in C++20, having finally formalized what were long identified as best practices. Though these restrictions might not be codified in the Standard for pre-C++20 software, they have been recognized best practices for as long as the Standard Library has existed and adherence to them will materially improve the ability of software to migrate to future language standards irrespective of what version of the language standard is being targeted.

⁷C++20 limits the specialization of variable templates to only those instances where specialization is explicitly allowed and does so only for the mathematical constants in `<numbers>`.

⁸C++20 identifies these functions as `addressable` and gives that property to only `iostream` manipulators since those are the only functions in the Standard Library for which taking their address is part of normal usage.

`inline namespace std {}` before including any standard headers. This practice is, however, explicitly called out as ill-formed within the C++11 Standard. Although not uniformly diagnosed as an error by all compilers, attempting this forbidden practice is apt to lead to surprising problems even if not diagnosed as an error immediately.

Inconsistent use of inline keyword is ill formed, no diagnostic required

It is an ODR violation, IFNDR, for a nested namespace to be `inline` in one translation unit and `noninline` in another. And yet, the motivating use case of this feature relies on the linker to actively complain whenever different, incompatible versions — nested within different, possibly `inline`-inconsistent, namespaces of an ABI — are used within a single executable. Because declaring a nested namespace `inline` does not, by design, affect linker-level symbols, developers must take appropriate care, such as effective use of header files, to defend against such preventable inconsistencies.

Annoyances

Inability to redeclare across namespaces impedes code factoring

An essential feature of an `inline` namespace is the ability to declare a template within a nested `inline` namespace and then specialize it within its enclosing namespace. For example, we can declare

- a type template, `S0`
- a couple of function templates, `f0` and `g0`
- and a member function template `h0`, which is similar to `f0`

in an `inline` namespace, `inner`, and specialize each of them, such as for `int`, in the enclosing namespace, `outer`:

```
namespace outer                                     // enclosing namespace
{
    inline namespace inner                          // nested namespace
    {
        template<typename T> struct S0;             // declarations of
        template<typename T> void f0();             // various class
        template<typename T> void g0(T v);         // and function
        struct A0 { template <typename T> void h0(); }; // templates
    }

    template<> struct S0<int> { };                  // specializations
    template<> void f0<int>() { }                   // of the various
    void g0(int) { } /* overload not specialization */ // class and function
    template<> void A0::h0<int>() { }               // declarations above
}                                                  // in outer namespace
```

Note that, in the case of `g0` in this example, the “specialization” `void g0(int)` is a non-template *overload* of the function template `g0` rather than a specialization of it. We *cannot*, however, portably⁹ declare these templates within the **outer** namespace and then specialize them within the inner one, even though the inner namespace is **inline**:

```
namespace outer                                // enclosing namespace
{
    template<typename T> struct S1;             // class template
    template<typename T> void f1();            // function template
    template<typename T> void g1(T v);         // function template

    struct A1 { template <typename T> void h1(); }; // member function template

    inline namespace inner                      // nested namespace
    {
        template<> struct S1<int> { };          // BAD IDEA
        template<> void f1<int>() { };          // Error, S1 not a template
        void g1(int) { }                       // OK, overloaded function
        template<> void A1::h1<int>() { }       // Error, h1 not a template
    }
}
```

Attempting to declare a template in the **outer** namespace and then define, effectively redeclaring, it in an **inline** inner one causes the name to be inaccessible within the **outer** namespace:

```
namespace outer                                // enclosing namespace
{
    template<typename T> struct S2;             // BAD IDEA
    template<typename T> void f2();            // declarations of
    template<typename T> void g2(T v);         // various class and
                                                // function templates

    inline namespace inner                      // nested namespace
    {
        template<typename T> struct S2 { };    // definitions of
        template<typename T> void f2() { };    // unrelated class and
        template<typename T> void g2(T v) { }  // function templates
    }

    template<> struct S2<int> { };              // Error, S2 is ambiguous in outer.
    template<> void f2<int>() { };              // Error, f2 is ambiguous in outer.
    void g2(int) { }                           // OK, g2 is an overload definition.
}
```

⁹GCC provides the `-fpermissive` flag, which allows the example containing specializations within the inner namespace to compile with warnings. Note again that `g1(int)`, being an *overload* and not a *specialization*, wasn't an error and, therefore, isn't a warning either.

Finally, declaring a template in the nested **inline** namespace **inner** in the example above and then subsequently defining it in the enclosing **outer** namespace has the same effect of making declared symbols ambiguous in the **outer** namespace:

```

namespace outer                                // enclosing namespace
{
    inline namespace inner                      // BAD IDEA
    {
        template<typename T> struct S3;         // declarations of
        template<typename T> void f3();         // various class
        template<typename T> void g3(T v);     // and function
        struct A3 { template <typename T> void h3(); }; // templates
    }

    template<typename T> struct S3 { };         // definitions of
    template<typename T> void f3() { }          // unrelated class
    template<typename T> void g3(T v) { }      // and function
    template<typename T> void A3::h3() { }     // templates

    template<> struct S3<int> { };              // Error, S3 is ambiguous in outer.
    template<> void f3<int>() { }               // Error, f3 is ambiguous in outer.
    void g3(int) { }                           // OK, g3 is an overload definition.
    template<> void A3::h3<int>() { }           // Error, h2 is ambiguous in outer.
}

```

Note that, although the definition for a member function template must be located directly within the namespace in which it is declared, a class or function template, once declared, may instead be defined in a different scope by using an appropriate name qualification:

```

template <typename T> struct outer::S3 { };     // OK, enclosing namespace
template <typename T> void outer::inner::f3() { } // OK, nested namespace
template <typename T> void outer::g3(T v) { }   // OK, enclosing namespace
template <typename T> void outer::A3::h3<T>() { } // Error, ill-formed

namespace outer
{
    inline namespace inner
    {
        template <typename T> void A3::h3() { } // OK, within same namespace
    }
}

```

Also note that, as ever, the corresponding definition of the declared template must have been seen before it can be used in a context requiring a complete type. The importance of ensuring that all specializations of a template have been seen before it is used substantively (i.e., **ODR-used**) cannot be overstated, giving rise to the only limerick, which is actually part of the normative text, in the C++ Language Standard¹⁰:

¹⁰See **iso11a**, section 14.7.3, “Explicit specialization,” paragraph 7, pp. 375–376, specifically p. 376.

When writing a specialization,
be careful about its location;
or to make it compile
will be such a trial
as to kindle its self-immolation.

Only one namespace can contain any given inline namespace

Unlike conventional **using** directives, which can be used to generate arbitrary many-to-many relationships between different namespaces, **inline** namespaces can be used only to contribute names to the sequence of enclosing namespaces up to the first **noninline** one. In cases in which the names from a namespace are desired in multiple other namespaces, the classical **using** directive must be used, with the subtle differences between the two modes properly addressed.

As an example, the C++14 Standard Library provides a hierarchy of nested **inline** namespaces for literals of different sorts within namespace `std`.

- `std::literals::complex_literals`
- `std::literals::chrono_literals`
- `std::literals::string_literals`
- `std::literals::string_view_literals`

These namespaces can be imported to a local scope in one shot via a **using** `std::literals` or instead, more selectively, by **using** the nested namespaces directly. This separation of the types used with user-defined literals, which are all in namespace `std`, from the user-defined literals that can be used to create those types led to some frustration; those who had a **using namespace** `std`; could reasonably have expected to get the user-defined literals associated with their `std` types. However, the types in the nested namespace `std::chrono` did *not* meet this expectation.¹¹

Eventually *both* solutions for incorporating literal namespaces, **inline** from `std::literals` and **noninline** from `std::chrono`, were pressed into service when, in C++17, a **using namespace** `literals::chrono_literals`; was added to the `std::chrono` namespace. The Standard does not, however, benefit in any objective way from any of these namespaces being **inline** since the artifacts in the `literals` namespace neither depend on ADL nor are templates in need of user-defined specializations; hence, having all **noninline** namespaces with appropriate **using** declarations would have been functionally indistinguishable from the bifurcated approach taken.

¹¹CWG issue 2278; [hinnant17](#)

See Also

- “**alignas**” (§2.1, p. 168) provides properly aligned storage for an object of arbitrary type `T` in the example in *Use Cases — Build modes and ABI link safety* on page 1071.

Further Reading

- **sutter14a** uses inline namespaces as part of a proposal for a portable ABI across compilers.
- **lopez-gomez20** uses inline namespaces as part of a solution to avoid ODR violation in an interpreter.

Appendix

Case study of using inline namespaces for versioning

By Niall Douglas

Let me tell you what I (don't) use them for. It is not a conventional opinion.

At a previous well-regarded company, they were shipping no less than forty-three copies of Boost in their application. Boost was not on the approved libraries list, but the great thing about header-only libraries is that they don't obviously appear in final binaries, unless you look for them. So each individual team was including bits of Boost quietly and without telling their legal department. Why? Because it saved time. (This was C++98, and `boost::shared_ptr` and `boost::function` are both extremely attractive facilities.)

Here's the really interesting part: Most of these copies of Boost were not the same version. They were varying over a five-year release period. And, unfortunately, Boost makes no API or ABI guarantees. So, theoretically, you could get two different incompatible versions of Boost appearing in the same program binary, and BOOM! there goes memory corruption.

I advocated to Boost that a simple solution would be for Boost to wrap up their implementation into an internal inline namespace. That inline namespace ought to mean something.

- `lib::v1` is the *stable*, version-1 ABI, which is guaranteed to be compatible with all past and future `lib::v1` ABIs, forever, as determined by the ABI-compliance-check tool that runs on **CI**. The same goes for `v2`, `v3`, and so on.
- `lib::v2_a7fe42d` is the *unstable*, version-2 ABI, which may be incompatible with any other `lib::*` ABI; hence, the seven hex chars after the underscore are the git short **SHA**, permuted by every commit to the git repository but, in practice, per CMake configure, because nobody wants to rebuild everything per commit. This ensures that no symbols from any revision of `lib` will *ever* silently collide or otherwise interfere with any other revision of `lib`, when combined into a single binary by a dumb linker.

I have been steadily making progress on getting Boost to avoid putting anything in the global namespace, so a straightforward find-and-replace can let you “fix” on a particular version of Boost.

That’s all the same as the pitch for **inline** namespaces. You’ll see the same technique used in `libstdc++` and many other major modern C++ codebases.

But I’ll tell you now, I don’t use **inline** namespaces anymore. Now what I do is use a macro defined to a uniquely named namespace. My build system uses the git SHA to synthesize namespace macros for my namespace name, beginning the namespace and ending the namespace. Finally, in the documentation, I teach people to always use a namespace alias to a macro to denote the namespace:

```
namespace output = OUTCOME_V2_NAMESPACE;
```

That macro expands to something like `::outcome_v2_ee9abc2`; that is, I don’t use **inline** namespaces anymore.

Why?

Well, for *existing* libraries that don’t want to break backward source compatibility, I think **inline** namespaces serve a need. For *new* libraries, I think a macro-defined namespace is clearer.

- It causes users to publicly commit to “I know what you’re doing here, what it means, and what its consequences are.”
- It declares to *other* users that something unusual (i.e., go read the documentation) is happening here, instead of silent magic behind the scenes.
- It prevents accidents that interfere with ADL and other customization points, which induce surprise, such as accidentally injecting a customization point into `lib`, not into `lib::v2`.
- Using macros to denote namespace lets us reuse the preprocessor machinery to generate C++ modules using the exact same codebase; C++ modules are used if the compiler supports them, else we fall back to inclusion.

Finally, and here’s the real rub, because we now have namespace aliases, if I were tempted to use an **inline** namespace, nowadays I probably would instead use a uniquely named namespace instead, and, in the `include` file, I’d alias a user-friendly name to that uniquely named namespace. I think that approach is less likely to induce surprise in the typical developer’s likely use cases than **inline** namespaces, such as injecting customization points into the wrong namespace.

So now I hope you’ve got a good handle on **inline** namespaces: I was once keen on them, but after some years of experience, I’ve gone off them in favor of better-in-my-opinion alternatives. Unfortunately, if your type `x::S` has members of type `a::T` and macros decide if that is `a::v1::T` or `a::v2::T`, then no linker protects the higher-level types from ODR bugs, unless you also version `x`.

Index

Symbols

- { } (braced-initialization syntax)
 - annoyances, 247–255
 - allowing narrowing conversions, 247–248
 - auto** deduction, 253–254
 - broken macro-invocation syntax, 248–249
 - copy list initialization in member initializer lists, lack of, 249–250
 - explicit** constructors passed multiple arguments, 250–252
 - narrowing aggregate initialization, 247
 - obfuscation with opaque usage, 252–253
 - operator acceptance, 254–255
 - C++03 aggregate initialization, 219–222
 - C++03 initialization syntax, 215–219
 - C++11 aggregate initialization, 224–225
 - copy initialization and scalars, 235–236
 - copy list initialization, 226–228
 - default member initialization and, 233
 - description of, 215
 - direct list initialization, 228–231
 - further reading for, 256
 - list initialization, 233–234
 - potential pitfalls, 242–247
 - aggregates with deleted constructors, 247
 - inadvertently calling initializer-list constructors, 242–244
 - NRVO and implicit moves disabled, 244–246
 - restrictions on narrowing conversions, 222–224
 - support for, 562–564
 - type name omissions, 234
 - use cases, 236–242
 - avoiding the most vexing parse, 237–238
 - defining value-initialized variables, 236–237
 - uniform initialization in factory functions, 239–241
 - uniform initialization in generic code, 238–239
 - uniform member initialization in generic code, 241–242
 - variables in conditional expressions, 235
- , (comma) operator, 268, 928, 955n25
 - constexpr** functions, 265
 - rvalue references, 816
 - >> (consecutive right-angle brackets)
 - description of, 21
 - further reading for, 24
 - potential pitfalls with, 22–24
 - use cases, 22
 - =**default** syntax. *See also* deleted functions; rvalue references; **static_assert**
 - annoyances, 42–43
 - description of, 33–36
 - first declaration of special member function, 34–35
 - further reading for, 44
 - implementation of user-provided special member function, 35–36
 - implicit generation of special member functions, 44–45
 - potential pitfalls, 41–42
 - use cases, 36–41
 - making explicit class APIs with no runtime cost, 38–39
 - physically decoupling interface from implementation, 40–41
 - preserving type triviality, 39–40
 - restoring generation of suppressed special member function, 36–37
 - =**delete** syntax. *See also* defaulted functions; rvalue references
 - annoyances, 58–59
 - description of, 53
 - further reading for, 60
 - use cases, 53–57
 - hiding structural base class member functions, 56–57
 - preventing implicit conversion, 55–56
 - suppressing special member function generation, 53–55
 - ' (digit separator). *See also* binary literals
 - description of, 152–153
 - further reading for, 154
 - loss of precision in floating-point literals, 154–156
 - use cases, 153
 - (()) (double parentheses) notation, 25, 30
 - > (greater-than) operator, 21–22

Index

`||` (logical or) operator, 265
() (parentheses), with **decltype** operands, 25, 30
&& (rvalue references)
 annoyances, 804–812
 destructive move, lack of, 811–812
 evolution of value categories, 807
 moved-from object requirements overly strict, 807–811
 RVO and NRVO require declared copy/move constructors, 804–805
 std::move does not move, 805–806
 visual similarity to forwarding references, 806–807
decltype results as, 26
description of, 710–741
in expressions, 730–731
extended value categories in C++11/14, 716–723
further reading for, 813
lvalue references, comparison, 710–711
modifiable, 820–821
motivation for, 715–716
move operations, 714–715
moves in **return** statements, 734–740
necessity of, 824
overload resolution, 713
overloading on reference types, 727–730
potential pitfalls, 782–804
 disabling NRVO, 783–784
 failure to std::move named rvalue references, 784–785
 implementing move-only types without std::unique_ptr, 791–794
 inconsistent expectations on moved-from objects, 794–803
 making noncopyable type movable without just cause, 788–791
 move operations that throw, 787
 repeatedly calling std::move on named rvalue references, 785–786
 requiring owned resources to be valid, 803–804
 returning **const rvalues** pessimizes performance, 786–787
 sink arguments require copying, 782–783
 some moves equivalent to copies, 788
special member functions, 732–733
std::move, 731–732
use cases, 741–781
 identifying value categories, 779–781
 move operations as optimizations of copying, 741–767
 move-only types, 768–771

 passing around resource-owning objects by value, 771–775
 sink arguments, 775–779
 value category evolution, 813–828
 xvalues, 712–713
[[]] (square brackets), 12
0b prefix. *See also* digit separator (')
description of, 142–143
further reading for, 146
use cases, 144–146
 bit masking and bitwise operations, 144–145
 replicating constant binary data, 145–146

A

ABI link safety, build modes and, 1071–1074
ABI versioning
 exception specification incompatibility, 1148–1149
 link-safe, 1067–1071
abstract classes, 1008, 1009
 extracting, 1018
 final contextual keyword, 1008–1009
 VShape as, 440
abstract interfaces, 1048, 1200
 deduced return types and, 1200
 pure, 540, 1020, 1021
abstract machine, 1118
abstract-syntax-tree (AST), 1054
access levels, 421, 489, 549–551, 550n5
access specifiers, 34, 35, 439, 537, 550n5, 884
accessible (from a context), 410
accessible copy constructor, 641, 644
accidental terminate, 1124–1128
acquire/release memory barrier, 80n7
active members, 406
adapter requirements in range-based **for** loops, 706
aggregate class, 415
aggregate initialization. *See also* default member initializers
 annoyances, 140–141
 in C++03, 219–222
 in C++11, 224–225, 241
 constexpr functions, 273–274
 default member initializers, 330
 with deleted constructors, 247
 description of, 138–139
 narrowing, 247
 potential pitfalls, 140
 rvalue references, 752
 of scalar members, 463
 use cases, 139
aggregate types, 138, 1087
 direct list initialization, 230

- generalized PODs, 410
 - as literal types, 279–280
- aggregates, 877
 - default member initializers, 330
 - POD types, 402
- algebra, 961
- algorithm selection, 947
- algorithms. *See also* functions
 - configuring via lambda expressions, 86–87
 - conflating optimization with code-size reduction, 1143–1144
 - constexpr** functions in, 294n19
 - divide and conquer, 297
 - nonrecursive **constexpr** algorithms, 961–962
 - optimized C++11 example algorithms, 965–967
- optimized metaprogramming algorithms, 963–964
- alias templates, 887
- aliases. *See also* inheriting constructors; trailing return
 - creating with **using** declarations, 133–137
 - description of, 133–134
 - use cases, 134–137
 - binding arguments to template parameters, 135–136
 - simplified **typedef** declarations, 134–135
 - type trait notation, 136–137
- aliasing, 638–639
- alignas** specifier
 - description of, 168–172
 - memory allocation, 181–183
 - natural alignment, 179–181
 - pack expansion, 921–922
 - potential pitfalls, 176–179
 - ill-formed, no-diagnostic required (IFNDR), 176–177
 - misleading applications to user-defined types (UDTs), 177–178
 - overlooking alternatives to avoiding false sharing, 178–179
 - underspecifying alignment, 176
 - strengthening alignment
 - of data members, 170–171
 - of particular objects, 169–170
 - of user-defined types (UDTs), 171
 - supported alignments, 168–169
 - type identifier as argument, 172
 - use cases, 172–176
 - false sharing, avoiding, 174–176
 - proper alignment for architecture-specific instructions, 173–174
 - sufficiently aligned object buffers, 172–173
- alignment. *See also* **alignas** specifier; **alignof** operator
 - for architecture-specific instructions, 173–174
 - incompatibly specified, 176–177
 - maximal fundamental, 193
 - natural, 179–181, 193
 - strengthening, 168
 - of data members, 170–171
 - of particular objects, 169–170
 - of user-defined types (UDTs), 171
 - supported, 168–169
 - underspecifying, 176
- alignment requirements, 168, 184
- alignof** operator. *See also* **alignas** specifier; **decltype**
 - annoyances, 193–194
 - description of, 184
 - fundamental types, 184
 - use cases, 186–193
 - monotonic memory allocation, 190–193
 - probing alignment of type during development, 186–187
 - sufficiently aligned buffers, 187–190
 - user-defined types, 185–186
- allocating objects, 634
- almost trivially destructible, 464–470
- amortized constant time (of a repeated operation), 636
- annotations. *See* attribute support; attributes
- anonymous function objects. *See* closure objects; lambda expressions
- API migration, facilitating, 1062–1067
- Application Binary Interface (ABI)
 - build modes and link safety, 1071–1074
 - changes in future C++ versions, 1089n5, 1114, 1114n24, 1148–1149
 - inline namespaces, 1056, 1064, 1083
 - link-safe ABI versioning, 1067–1071
 - POD types, 402
- Application Programming Interface (API), 402, 445
- arbitrary values, conflating with indeterminate values, 493–497
- architecture-specific instructions, alignment for, 173–174
- argument-dependent lookup (ADL), 472, 1056, 1058
 - inline namespaces, 1056, 1058–1059
 - range-based **for** loops, 681, 707–709
 - user-defined literals (UDLs), 841
- arguments
 - passing multiple to explicit constructors, 250–252
 - of same type, 564–565
 - template, local/unnamed types, 83–88

Index

- value initialization, avoiding the most vexing parse, 237–238
- arithmetic operators, braced lists and, 254–255
- arithmetic types
 - enum** class, 334, 337–339
 - implicit conversion, avoiding, 337–339
- array types
 - alignof** operator, 184
 - as literal types, 280
 - as standard-layout types, 417
 - as trivial types, 425
- arrays
 - built-in, deducing, 211–212
 - initialization with `std::initializer_list`
 - annoyances, 567–571
 - description of, 553–561
 - further reading for, 571
 - potential pitfalls, 566–567
 - range-based **for** loops, 571–572
 - use cases, 561–566
 - size deduction, lack of, 330
 - traversing with range-based **for** loops, 683–684
- array-to-pointer decay, 220, 222
- ASCII
 - basic source character set, 130
 - Unicode string literals, 129
- as if, 307
- assert, 656
- assert statements in dependency chain, 1002
- assignable (type), 486
- assignment operator, 521–522
 - braced lists and, 254–255
 - lvalue references, 816
- atomic (operation), 80–82
- attribute lists, 922
- attribute support. *See also* attributes
 - description of
 - attribute placement, 13
 - attribute syntax, 12–13
 - standardized compiler-specific attributes, 13–14
 - potential pitfalls
 - undefined behavior, 19
 - unrecognized attributes, 18–19
 - use cases
 - control of external static analysis, 17–18
 - hints for additional optimization opportunities, 15–16
 - prompting of compiler diagnostics, 14–15
 - statement of explicit assumptions, 16–17
 - statements of semantic properties, 18
- attributes. *See also* attribute support
 - `[[carries_dependency]]`
 - description of, 998–1000
 - further reading for, 1006
 - potential pitfalls, 1005
 - use cases, 1000–1005
 - `[[clang::no_sanitize]]`, 14
 - definition of, 12
 - `[[deprecated]]`, 14
 - description of, 147–148
 - potential pitfalls, 150
 - use cases, 148–150
 - `[[gnu::cold]]`, 15
 - `[[gnu::const]]`, 16–17, 19
 - `[[gnu::pure]]`, 14, 16
 - `[[gnu::warn_unused_result]]`, 14–15, 15n7
 - `[[gsl::suppress]]`, 17–18
 - `[[noreturn]]`, 13
 - description of, 95
 - further reading for, 98
 - potential pitfalls, 97–98
 - use cases, 95–97
- auto** variables
 - annoyances, 212–213
 - nonstatic data members, not allowed, 212
 - template argument deductions, not all allowed, 212–213
 - braced initialization and, 253–254
 - decltype(auto)** placeholder
 - annoyances, 1213
 - description of, 1205–1210
 - potential pitfalls, 1212–1213
 - use cases, 1210–1212
 - description of, 195–199
 - further reading for, 214
 - idioms for, 1213
 - potential pitfalls, 204–212
 - compromised readability, 204
 - deducing built-in arrays, 211–212
 - deduction for list initialization, 210–211
 - hidden properties of fundamental types, 209–210
 - interface restrictions, lack of, 208–209
 - unexpected conversions, 206–208
 - unintentional copies, 204–206
 - return-type deduction
 - annoyances, 1201–1203
 - description of, 1182–1194
 - potential pitfalls, 1200
 - use cases, 1194–1200
 - use cases, 200–203
 - deeply nested variable types, 202–203
 - ensuring variable initialization, 200
 - implementation-defined or compiler-synthesized variable types, 202
 - preventing unexpected implicit conversions, 201

- redundant type name repetition, avoiding, 200–201
 - resilience to library code changes, 203
- auto&&**, 383–384
- automatic objects, 69
- automatic storage duration, 68, 195, 582, 731, 735, 740n16
- automatic variables, 526–527, 735–737
- auxiliary variables, creating with **decltype**, 28
- B**
- backward compatibility, 1113
- barriers, 80n7
- base classes, hiding member functions, 56–57
- base name (of a component), 667n5
- base specifier list, 883, 884, 915–917, 925
- base-class constructors. *See* constructors
- basic exception-safety guarantee, 644, 651
- basic source character set, 110, 130
- behavior, undefined. *See* undefined behavior
- benchmark tests, 114
- big-endian **float** layouts, 531–534
- binary literals. *See also* digit separator ('')
 - description of, 142–143
 - further reading for, 146
 - use cases, 144–146
 - bit masking and bitwise operations, 144–145
 - replicating constant binary data, 145–146
- binary searches, 292, 294, 944–946
- binary trees, 1050–1054
- bind function, 14
- bit fields, 329n4, 526
- bit flags, 347–348
- bit masking, 144–145
- bit representation of PODs, 530–534
- bitwise copies of PODs, exporting, 479–480
- bitwise operations, 21–24, 144–145
- block scope, 587
- boilerplate code
 - aliases, 136–137
 - default member initializers, 322
 - enum** class, 333
 - implementation inheritance, avoiding with, 540–541
 - repetition, avoiding, 323–325
 - structural inheritance, avoiding with, 540
 - variable templates, 161
- Boost, 1083–1084
- boost::variant**, 1180n2
- brace elision, 140, 220
- braced initialization
 - annoyances, 247–255
 - allowing narrowing conversions, 247–248
 - auto** deduction, 253–254
 - broken macro-invocation syntax, 248–249
 - copy list initialization in member initializer lists, lack of, 249–250
 - explicit** constructors passed multiple arguments, 250–252
 - narrowing aggregate initialization, 247
 - obfuscation with opaque usage, 252–253
 - operator acceptance, 254–255
 - C++03 aggregate initialization, 219–222
 - C++03 initialization syntax, 215–219
 - C++11 aggregate initialization, 224–225
 - copy initialization and scalars, 235–236
 - copy list initialization, 226–228
 - default member initialization and, 233
 - description of, 215
 - direct list initialization, 228–231
 - further reading for, 256
 - list initialization, 233–234
 - potential pitfalls, 242–247
 - aggregates with deleted constructors, 247
 - inadvertently calling initializer-list constructors, 242–244
 - NRVO and implicit moves disabled, 244–246
 - range-based **for** loops, 684
 - restrictions on narrowing conversions, 222–224
 - support for, 562–564
 - type name omissions, 234
 - use cases, 236–242
 - avoiding the most vexing parse, 237–238
 - defining value-initialized variables, 236–237
 - uniform initialization in factory functions, 239–241
 - uniform initialization in generic code, 238–239
 - uniform member initialization in generic code, 241–242
 - variables in conditional expressions, 235
 - variadic templates, 926
- braced initialized, 752
- braced initializer lists, 554–557, 912–914
- brackets ([]), 12
- buffers
 - creating with sufficient alignment, 172–173
 - sufficiently aligned, 187–190
- build modes, ABI link safety and, 1071–1074
- builder classes, optimizing, 1167–1170
- built-in arrays, deducing, 211–212
- bytes, 153, 286–287, 503–505, 533–534, 748, 1107

Index

C

C linkage, 403–404

C Standard Library, **noexcept** operator and, 631–632

C++ Language Standard, limerick in, 1081–1082
The C++ Programming Language (Stroustrup), 4

C++03

- aggregate initialization, 219–222, 247
- double-checked-lock pattern, 81–82
- dynamic exception specifications, 1089
- explicit-instantiation directives, 353n1
- initialization syntax, 215–219
- nested templated types, 22
- passing multiple arguments to **explicit** constructors, 250–252
- POD types, 412–415
- right-shift operator (**>>**), 22–24
- unscoped enumerators, workarounds for, 332–333
- user-declared, 413n6
- weakly typed enumerators, drawbacks to, 333–335

C++11

- aggregate initialization, 224–225, 241
- conditionally safe features
 - alignas** specifier, 168–183
 - alignof** operator, 184–194
 - auto** variables, 195–214
 - constexpr** functions, 257–301
 - constexpr** variables, 302–317
 - default member initializers, 318–331
 - enum** class, 332–352
 - extern template**, 353–376
 - forwarding references, 377–400
 - generalized plain old data types (PODs), 401–534
 - inheriting constructors, 535–552
 - lambda expressions, 573–614
 - noexcept** operator, 615–659
 - opaque enumerations, 660–678
 - range-based **for** loops, 679–709
 - rvalue references, 26, 710–828
 - `std::initializer_list`, 553–572
 - underlying types (UTs), 829–834
 - user-defined literals (UDLs), 835–872
 - variadic templates, 873–958
- interface test, 275
- lvalue references, 717–720
- lvalue-reference declarations prior to, 815–818
- memory allocation, 763n25
- new keywords in, 1023
- optimized example algorithms, 965–967
- POD types. *See* POD types

prvalues, 720–721

safe features

- attribute support, 12–20
- consecutive right-angle brackets (**>>**), 21–24

decltype, 25–32

- defaulted functions, 33–45
- delegating constructors, 46–52
- deleted functions, 53–60
- explicit conversion operators, 61–67
- local/unnamed types, 83–88
- long long** integral type, 89–94
- `[[noreturn]]` attribute, 13
- `[[noreturn]]` attribute in, 95–98
- nullptr** keyword, 99–103
- override** member-function specifier, 5, 104–107
- raw string literals, 108–114
- static_assert**, 115–123
- thread-safe function-scope static variables, 68–82
- trailing return, 28. *See also* **decltype**;
- deduced return type
- trailing return types, 124–128
- type/template aliases, 133–137
- Unicode string literals, 129–132

scoped enumerations, 335–336

`std::unique_ptr<T>`, 42n3

unsafe features

- `[[carries_dependency]]` attribute, 998–1006
- final** contextual keyword, 1007–1030
- friend** declarations, 1031–1054
- inline** namespaces, 1055–1084
- noexcept** exception specification, 1085–1152
- ref-qualifiers, 1153–1173
- union type, 1174–1181

user-provided, 413n6

value categories prior to, 814–815

xvalues, 721–723

C++14

- capturing ***this** by copy, 612
- conditionally safe features
 - constexpr** functions, 959–967
 - generic lambdas, 968–985
 - lambda-capture expressions, 986–995
- lvalue references, 717–720
- new keywords in, 1023n7
- prvalues*, 720–721
- safe features
 - aggregate initialization, 138–141
 - binary literals, 142–146
 - `[[deprecated]]` attribute, 14, 147–151

- digit separator (`'`), 152–156
 - templated variable declarations, 157–166
 - `std::index_sequence`, 293
 - `<type_traits>` header, 1014
 - unsafe features
 - auto** return-type deduction, 1182–1204
 - decltype(auto)** placeholder, 1205–1214
 - user-defined literals (UDLs) in, 852–853
 - xvalues*, 721–723
- C++17
- capturing ***this** by copy, 611n7
 - conditionally supported, 425n7
 - dynamic exception specifications, 1085n1
 - exception specifications and type system, 1089n5
 - false sharing, avoiding, 175n6
 - fold expressions, 955n25
 - guaranteed copy elision, 216n1, 648n11, 805n30, 827n54
 - if **constexpr** language feature, 641n10
 - nested namespaces, 1055n1
 - new keywords in, 1023n7
 - pmr allocators, 763n25
 - polymorphic memory resources, 190n3
 - range-based **for** loops, 681n2
 - sentinels, 707n12
 - `std::any`, 187n2
 - `std::pmr::monotonic_resource`, 468n27
 - `std::pmr::unsynchronized_pool_resource`, 468n27
 - `std::string_view`, 874n1
 - `std::variant`, 452n19, 1180n2
 - structured binding, 201n2, 685n3
 - trivial types, 425n7
 - type traits, 651n12
- C++20
- bit field initialization, 329n4
 - char**-like object, 479n29
 - concepts, 122n5, 208n3, 480n30, 1201n5
 - constexpr** functions as destructors, 463n25
 - constexpr** functions in algorithms, 294n19
 - constexpr** keyword, 75n5, 304n1, 316n8
 - contracts, 467n26
 - deleted constructors, 247n8
 - designated initializers, 139n1
 - destructors, 407n3
 - encapsulation of helper types, 85n3
 - enumeration comparisons, 335n1
 - floating-point non-type template parameters, 903n7
 - generic lambdas, explicit parameter types, 193–194
 - implicit conversion, 223n3
 - implicitly movable entities, 735n13
 - manifestly constant evaluated, 258n1
 - moves in **return** statements, 740n16
 - nested namespaces, 1055n1
 - new keywords in, 1023n7
 - `[[no_unique_address]]` attribute, 1029n15
 - Ranges Library, 686n4, 687n5
 - ranges library, 391–393
 - relaxed restrictions on **constexpr** functions, 960n1
 - reordering data members, 178n10
 - requires clause, 486n31
 - sentinels, 707n12
 - Standard Library–related restrictions, 1078n6
 - `std::bit_cast`, 514n41, 516n42
 - `std::is_constant_evaluated()`, 297n20
 - `std::is_pod`, 438n14
 - `std::remove_cvref<T>`, 399n6
 - terse concept notation syntax, 398n5
 - trivially destructible types, 430n9
 - typename** disambiguator, 382n1
 - unscoped enumerated types, 833n2
 - user-declared constructors, 274n7
- C++23
- guaranteed copy elision, 805n30
 - reordering data members, 178n10
- C++-only types, translating to C, 452–456
- C99, flexible array members, 404n1
- cache associativity, 182n11
- cache hit, 181
- cache lines, 174–175, 181–183, 459, 1142
- cache miss, 182
- call operators in functor classes, 574–575
- callable objects, 70, 994
- callback functions, 669. *See also* lambda expressions
- callbacks, event-driven, 603–604
- capture default, 582–583, 600, 608
- captured by copy, 582, 611–612, 990–992
- captured by reference, 582
- captured by value. *See* captured by copy
- captured variables, 582–585, 590–591, 602, 609–610, 990
- `[[carries_dependency]]` attribute
- description of, 998–1000
 - further reading for, 1006
 - potential pitfalls, 1005
 - use cases, 1000–1005
- Carruth, Chandler, 1134
- carry dependency, 999
- cast, 345
- character literals, 837, 844n1
- char**-like object, 479, 479n29
- `checkBalance` function, 15
- `checksumLength` function, 27, 28n1

Index

Clang

- acquire/release memory barrier, 80n7
- attribute support
 - [[gnu::cold]] attribute, 15
 - [[gnu::const]] attribute, 16–17, 19
 - [[gnu::warn_unused_result]] attribute, 14–15, 15n7
- standardized compiler-specific attributes, 14
- compiler warnings, 150
- delegating constructors, 50n2
- explicit expression of type-consistency, 28n1
- incompatibly specified alignment, 177
- indirect calls, 947n22
- inline** namespaces, 1061n3
- namespace-qualified name support, 13n2
- nonrecursive **constexpr** algorithms, 961n2, 962n3
- pair mismatches, 699n8
- reducing code size, 1104n16, 1110–1111
- stack unwinding, 621n4
- template instantiation with deduced return type, 1192n3
- trivial copy/move constructors, 528n62
- underspecifying alignment, 176
- [[clang::no_sanitiz]] attribute, 14
- class APIs, making explicit, 38–39
- class keys, 414–415
- class member functions. *See* functions
- class template specialization, 1059–1061
- class templates, 892. *See also* variadic class templates
 - preventing misuse of, 118–119
 - std::initializer_list usage, 555–558
- class types, 405
- classes. *See also* constructors; templates
 - constraints in hierarchies, 655–658
 - enum**
 - annoyances, 351
 - description of, 332–337
 - further reading for, 352
 - potential pitfalls, 344–350
 - use cases, 337–344
 - Packet, 27–28
 - variadic class templates, 878–880
 - member functions, 892–894
 - non-type template parameter packs, 901–903
 - specialization of, 884–887
 - type template parameter packs, 880–884
- class-specific memory management, 1088n4
- clients
 - API migration, facilitating, 1063
 - inline namespaces, 1059

- closure objects, 974, 978, 986. *See also* lambda expressions; lambda-capture expressions
 - deduced return types and, 1197–1198
 - forwarding variables into, 992–993
 - identity, 968
 - moving objects into, 988–989
 - mutable state, 989–990
- closure types, 87, 578–581, 968–970, 1197
- code bloat, 1054
 - extern template**, 353, 364–365, 371–372, 375
 - reducing in object files, 365–369
- code duplication, 48–50, 1144–1147
- code elision, 1136, 1139
- code factoring, impeding with inline namespaces, 1079–1082
- code motion, 1136, 1137
- code point, 129, 131
- code size, reducing, 1101–1111, 1143–1144
- code units, 131, 476
- cold path, 1103, 1104n16, 1134–1135
- Collatz conjecture, 313
- Collatz function, 313
- Collatz length, 313
- Collatz sequence, 313
- collisions, 109–111
- comma (,) operator, 268, 928, 955n25
 - constexpr** functions, 265
 - rvalue references, 816
- common initial member sequence (CIMS), 406, 421–423, 447
- common type, 1186
- compilation errors, forcing with =**delete** syntax. *See* defaulted functions; rvalue references
- compiler diagnostics, prompting, 14–15
- Compiler Explorer, 1110
- compiler warnings, 150
- compiler-generated special member functions, 621–626
- compiler-synthesized types, 202
- compile-time accessible variables. *See* **constexpr** variables
- compile-time assertions with **static_assert**
 - annoyances, 123
 - description of, 115–118
 - evaluation in templates, 116–118
 - further reading for, 123
 - potential pitfalls, 120–122
 - misuse to restrict overload sets, 121–122
 - unintended compilation failures, 120–121
 - syntax and semantics, 115–116

- use cases, 118–119
 - preventing misuse of class and function templates, 118–119
 - verifying assumptions about target platform, 118
- compile-time constants
 - as conditional expression, 615
 - in constant expressions, 838
 - constexpr** functions, 260–262
 - enum** class, 346–347
 - enumerators as, 164
 - pi, 160
 - POD types, 463
 - user-defined literals (UDLs), 862
- compile-time constructible, literal types, 462–464
- compile-time coupling, 6, 40–41, 663, 677n11, 1200
- compile-time dispatch, 121
- compile-time evaluation
 - of data tables, 291–295
 - diagnosing undefined behavior, 312–314
 - low compiler limits, 295
 - overhead costs of **constexpr** functions, 298–299
 - penalizing run time for, 299–300
 - string traversal, 287–291
- compile-time introspection, 947
- compile-time invocable functions. *See* **constexpr** functions
- compile-time polymorphism, 1046–1050
- compile-time visitation, 1050–1054
- complete types, 184, 316, 832, 891, 1014, 1081
 - alignof** operator, 184
 - default member initializers, 319
 - enum** class, 350
 - opaque enumerations, 661
 - rvalues*, 720
 - rvalue references, 720, 807
- complete-class context, 319, 1086n2
- component local, 664
- components
 - extern template**, 359–360
 - friend** declarations, 1035–1036, 1041
 - link-safe ABI versioning, 1068
 - opaque enumerations, 665, 667
- composite patterns, 1020
- compound assignment operators, braced lists and, 254–255
- compound expressions, **noexcept** operator and, 626–627
- concepts, 122n5, 208n3, 480n30, 571, 1201n5
- concrete classes, 56, 540
 - final** contextual keyword, 1008–1009
 - mocking, 1017–1020
 - performance, 1015–1017
- concrete monotonic allocators, 1022
- concurrent initialization, 68–69
- conditional compilation, 403, 469, 540, 1024n10, 1073
- conditional exception specifications, 1091–1092
- conditional expressions
 - compile-time constants as, 615
 - noexcept** operator, 615
 - variables in, initialization, 235
- conditional instantiation, 979–981
- conditional **noexcept** specifications, 639, 644
- conditionally compile, 469–470
- conditionally safe features
 - alignas** specifier
 - description of, 168–172
 - memory allocation, 181–183
 - natural alignment, 179–181
 - potential pitfalls, 176–179
 - use cases, 172–176
 - alignof** operator
 - annoyances, 193–194
 - description of, 184
 - fundamental types, 184
 - use cases, 186–193
 - user-defined types, 185–186
 - auto** variables
 - annoyances, 212–213
 - description of, 195–199
 - further reading for, 214
 - potential pitfalls, 204–212
 - use cases, 200–203
- braced initialization
 - annoyances, 247–255
 - C++03 aggregate initialization, 219–222
 - C++03 initialization syntax, 215–219
 - C++11 aggregate initialization, 224–225
 - copy initialization and scalars, 235–236
 - copy list initialization, 226–228
 - default member initialization and, 233
 - description of, 215
 - direct list initialization, 228–231
 - further reading for, 256
 - list initialization, 233–234
 - potential pitfalls, 242–247
 - restrictions on narrowing conversions, 222–224
 - type name omissions, 234
 - use cases, 236–242
 - variables in conditional expressions, 235
- constexpr** functions
 - annoyances, 299–300
 - compile-time evaluation, 284–286
 - constructor constraints, 269–276
 - description of, 257–261, 959–960
 - further reading for, 301, 965
 - inlining and definition visibility, 262–265

Index

- conditionally safe features (cont.)
- constexpr** functions (cont.)
 - literal types, 278–284
 - member functions, 266–268
 - optimized C++11 example algorithms, 965–967
 - parameter and return types, 277–278
 - as part of contract, 261–262
 - potential pitfalls, 295–299, 314
 - relaxed restrictions in C++14, 959–967
 - restrictions on, 268–269
 - templates, 276–277
 - type system and function pointers, 265–266
 - use cases, 286–295, 961–964
- constexpr** variables
 - annoyances, 314–316
 - description of, 302–307
 - potential pitfalls, 314
 - use cases, 307–314
- default member initializers, 6
 - annoyances, 328–330
 - description of, 318–321
 - potential pitfalls, 326–328
 - use cases, 322–325
- definition of, 5–6
- enum** class
 - annoyances, 351
 - description of, 332–337
 - further reading for, 352
 - potential pitfalls, 344–350
 - use cases, 337–344
- extern** template
 - annoyances, 373–375
 - description of, 353–365
 - further reading for, 376
 - potential pitfalls, 371–373
 - use cases, 365–370
- forwarding references
 - annoyances, 397–400
 - description of, 377–385
 - further reading for, 400
 - potential pitfalls, 394–397
 - use cases, 386–393
- generalized plain old data types (PODs)
 - annoyances, 521–529
 - bit representation, 530–534
 - C++03 POD types, 412–415
 - C++11 POD types, 415–417
 - description of, 401–402
 - further reading for, 530
 - future direction, 438–439
 - potential pitfalls, 479–521
 - privileges, 402–412
 - standard-layout class special properties, 420–425
 - standard-layout types, 417–420
 - trivial subcategories, 429–436
 - trivial types, 425–429
 - type traits, 436–438
 - use cases, 439–479
- generic lambdas
 - annoyances, 981–984
 - description of, 968–975
 - further reading for, 985
 - potential pitfalls, 981
 - use cases, 975–981
- inheriting constructors
 - annoyances, 549–552
 - description of, 535–539
 - potential pitfalls, 546–549
 - use cases, 539–545
- lambda expressions
 - annoyances, 611–614
 - description of, 573–597
 - further reading for, 614
 - potential pitfalls, 607–611
 - use cases, 597–607
- lambda-capture expressions
 - annoyances, 993–994
 - description of, 986–988
 - further reading for, 995
 - potential pitfalls, 992–993
 - use cases, 988–992
- noexcept** operator
 - annoyances, 650–658
 - description of, 615–634
 - further reading for, 658
 - move operations, 658–659
 - potential pitfalls, 647–650
 - use cases, 634–647
- opaque enumerations
 - annoyances, 677–678
 - description of, 660–663
 - further reading for, 678
 - potential pitfalls, 675–677
 - use cases, 663–675
- range-based **for** loops
 - annoyances, 703–709
 - description of, 679–684
 - further reading for, 709
 - potential pitfalls, 691–703
 - use cases, 684–691
- rvalue references
 - annoyances, 804–812
 - decltype** results as, 26
 - description of, 710–741
 - further reading for, 813
 - potential pitfalls, 782–804

- use cases, 741–781
 - value category evolution, 813–828
 - `std::initializer_list`
 - annoyances, 567–571
 - description of, 553–561
 - further reading for, 571
 - potential pitfalls, 566–567
 - range-based **for** loops, 571–572
 - use cases, 561–566
 - underlying types (UTs)
 - description of, 829–830
 - further reading for, 834
 - potential pitfalls, 832–833
 - use cases, 830–832
 - user-defined literals (UDLs)
 - annoyances, 869–871
 - description of, 835–853
 - further reading for, 872
 - potential pitfalls, 867–869
 - use cases, 853–867
 - variadic templates
 - annoyances, 953–957
 - description of, 873–925
 - further reading for, 958
 - potential pitfalls, 952–953
 - use cases, 925–951
 - conditionally supported, 425n7
 - conditionally supported behavior, 13n2
 - configuration structs, 139
 - conforming implementations, 14n4, 171n2
 - consecutive right-angle brackets (>>)
 - description of, 21
 - further reading for, 24
 - potential pitfalls with, 22–24
 - use cases, 22
 - const** data members
 - difficulty of synthesizing, 993–994
 - memcpy usage on, 489–493
 - returning as *rvalues* pessimizes performance, 786–787
 - const** default constructible, 218
 - const** objects, representation by initializer lists, 570
 - const** variables, capturing modifiable copy of, 990–992
 - constant binary data, replicating, 145–146
 - constant expressions, 115, 224, 257. *See also*
 - constexpr** functions; **constexpr** variables
 - compile-time constants in, 838
 - conditional exception specifications, 1091
 - noexcept** exception specifications, 1129
 - trivially destructible types, 431
 - user-defined literals (UDLs), 836
- constant initialization, 75
- constant time, 636, 823
- constants, named, 346–347
- const**-default-constructible, 218
 - constexpr** data structures, storing, 311–312
- constexpr** functions
- in algorithms, 294n19
 - annoyances, 299–300
 - implicit **const**-qualification, 300
 - penalizing run time to enable compile time, 299–300
 - compile-time evaluation, 284–286
 - constructor constraints, 269–276
 - description of, 257–261
 - destructors, 463n25
 - further reading for, 301
 - inlining and definition visibility, 262–265
 - literal types, 278–284
 - member functions, 266–268
 - noexcept** operator and, 654–655
 - parameter and return types, 277–278
 - as part of contract, 261–262
 - potential pitfalls, 295–299, 314
 - implementation difficulties, 296–297
 - low compiler limits, 295
 - overhead costs, 298–299
 - overzealous usage, 298
 - premature commitment, 297–298
 - relaxed restrictions in C++14, 959–967
 - description of, 959–960
 - further reading for, 965
 - optimized C++11 example algorithms, 965–967
 - use cases, 961–964
 - relaxed restrictions on. *See* **constexpr** variables; variadic templates
 - restrictions on, 268–269
 - templates, 276–277
 - type system and function pointers, 265–266
 - use cases, 286–295
 - alternative to function-like macros, 286–287
 - compile-time data table evaluation, 291–295
 - compile-time string traversal, 287–291
 - user-defined literals (UDLs), 838–839
- constexpr** specifier, 28n1
- constexpr** variables
- annoyances, 314–316
 - static **constexpr** member variables not defined in own class, 316
 - static member variables require external definitions, 314–315
 - description of, 302–307
 - initializer undefined behavior, 306–307
 - internal linkage, 307

Index

- constexpr** variables (cont.)
 - potential pitfalls, 314
 - use cases, 307–314
 - alternative to enumerated compile-time integral constants, 307–310
 - diagnosing undefined behavior at compile time, 312–314
 - nonintegral symbolic numeric constants, 310–311
 - storing **constexpr** data structures, 311–312
- constexpr** keyword, 75n5, 304n1, 316n8
- const**-qualified member functions, 300
- constraining
 - deduced parameters, 970–973
 - multiple arguments, 983–984
- constructors. *See also* copy constructors; move constructors
 - boilerplate code repetition, avoiding, 323–325
 - code duplication, avoiding, 48–50
 - delegating
 - description of, 46–48
 - potential pitfalls, 50–51
 - use cases, 48–50
 - deleted in aggregates, 247
 - explicit, passing multiple arguments, 250–252
 - inheriting
 - annoyances, 549–552
 - description of, 535–539
 - potential pitfalls, 546–549
 - use cases, 539–545
 - restrictions on, 269–276
 - for `std::initializer_list`, inadvertently calling, 242–244
 - as trivial, 437
 - user-declared, 274n7, 1087
 - value initializing arguments, avoiding the most vexing parse, 237–238
- containers
 - initialization, 561–562
 - iterating all elements, 684–685
 - nested, 22
- contextual keywords, 1007. *See also* keywords
 - override**
 - description of, 104–105
 - further reading for, 107
 - potential pitfalls, 106
 - use cases, 105–106
 - potential pitfalls, 1023
- contextually convertible to **bool**, 63–65, 1129
- continuous refactoring, 147
- contract guarantees
 - nofail functions, 1117–1122
 - overly strong, 1112–1116
- contract violations, 485
- contracts, 467n26, 485
 - constexpr** functions as part of, 261–262
 - new operator, 616
 - overly restrictive, 480–482
 - rvalue references, 714
- control constructs
 - emulating, 599–600
 - in lambda expressions, 600–601
- controlling constant expressions, 285
- conventional string literals, 113
- conversion operators
 - explicit
 - description of, 61–63
 - potential pitfalls, 66–67
 - use cases, 63–65
 - as placeholders, 1193–1194
- converting constructors, 61
- cooked UDL operators, 841, 843–845, 870
- cookies, 669–675
- copy assignable, 485–486
- copy assignment, 485, 758
- copy assignment operator
 - deleted functions, 54
 - rvalue references, 714
 - user-provided, 759
 - vertical encoding, 451
- copy constructible, 455
- copy construction, 489–492
- copy constructors
 - declaring special member functions, 34
 - deleted functions, 54
 - hijacking with perfect-forwarding constructor, 395–397
 - literal types and, 281
 - rvalue references, 714
 - RVO and NRVO requirements, 804–805
 - as trivial, 437
 - user-provided, 758–759
 - vertical encoding, 450
- copy elision, 390
- copy initialization, 215–216
 - in aggregate initialization, 221
 - in generic code, 239
 - for nonstatic data members, 318
 - scalar type, 235–236
 - unions, 506
- copy list initialization, 226–228
 - direct list initialization, compared, 231–232
 - in factory functions, 240
 - in generic code, 239
 - in member initializer lists, 249–250

- for nonstatic data members, 318
 - `std::initializer_list`, 555
 - copy operations, 522, 627
 - deleted functions, 53
 - move operations as optimization of, 741–767
 - rvalue references, 715
 - sink arguments, 782–783
 - some equivalent to moves, 788
 - copy semantics, 54, 627, 742, 852
 - copy/direct, 215
 - copy/swap idiom, 636, 1097
 - core constant expression, 960n1
 - core language specification, 482
 - core traits, 482
 - .cpp files, 41n2
 - critical section, 71
 - C-style ellipsis, 952
 - C-style functions, 158
 - curiously recurring template pattern (CRTP), 1042–1054
 - compile-time polymorphism, 1046–1050
 - compile-time visitation, 1050–1054
 - refactoring, 1042–1044
 - synthesizing equality, 1045–1046
 - currying, 597–598
 - cv-qualifiers, 1153–1154, 1157–1158, 1207–1208
 - forwarding references, 379
 - as literal types, 280
 - rvalue references, 724
 - as standard-layout types, 417
 - as trivial types, 425
 - cyclic physical dependency, 374
 - cyclically dependent, 75
- D**
- `d_engaged` flag, 1072
 - dangling references, 566–567, 607–608, 1171, 1212
 - data dependency, 999
 - data dependency chains, 998–999, 1002
 - data members
 - const**
 - difficulty of synthesizing, 993–994
 - memcpy usage on, 489–493
 - returning as *rvalues* pessimizes performance, 786–787
 - reordering, 178n10
 - strengthening alignment, 168, 170–171
 - data races, 68–69
 - data structures, **constexpr**, 311–312
 - data tables, compile-time evaluation, 291–295
 - death tests, 656
 - debug build, 468
 - debugging lambda expressions, 611
 - decay (of a type), 815
 - decimal floating-point (DFP), 862
 - declarations, 121, 315, 879
 - friend**
 - curiously recurring template pattern (CRTP) use cases, 1042–1054
 - description of, 1031–1033
 - further reading for, 1042
 - potential pitfalls, 1041
 - use cases, 1033–1041
 - prior to C++11, 815–818
 - user-provided destructors, 1105
 - declarator operators, 889
 - declared interface, 987
 - declared type (of an object), 25
 - declaring
 - deleted functions, 58–59
 - function pointers, 127–128
 - decltype**. *See also* **auto** variables; **decltype(auto)** placeholder; rvalue references
 - annoyances, 31
 - description of
 - use with entities, 25
 - use with expressions, 25–26
 - potential pitfalls, 30
 - use cases
 - avoidance of explicit typenames, 26–27
 - creation of auxiliary variable of generic type, 28
 - explicit expression of type-consistency, 27–28, 28n1
 - validation of generic expressions, 28–30
 - decltype(auto)** placeholder
 - annoyances, 1213
 - description of, 1205–1210
 - in new expressions, 1210
 - potential pitfalls, 1212–1213
 - specification, 1206–1208
 - syntactic restrictions, 1208–1209
 - use cases, 1210–1212
 - `declval` function, 31
 - deduced parameters, constraints on, 970–973
 - deduced return types, 593–594, 1146
 - annoyances, 1201–1203
 - description of, 1182–1194
 - for lambda expressions, 1189–1190, 1197–1198
 - potential pitfalls, 1200
 - use cases, 1194–1200
 - compiler-applied rules, 1197
 - complicated return types, 1194–1196
 - delaying return-type deduction, 1199–1120
 - perfect returning wrapped functions, 1198
 - returning lambda expressions, 1197–1198

Index

- deducing
 - built-in arrays, 211–212
 - list initialization, 210–211
 - pointer types, 197–198
 - reference types, 198
- deeply nested variable types, 202–203
- default constructed, 478, 752
- default constructors, 754, 1136
 - declaring special member functions, 33–34
 - suppressed by `std::initializer_list`, 568–570
 - as trivial, 437
 - user-provided, 755
- default initialization, 216–219, 765
 - in aggregate initialization, 221
 - constexpr** functions, 273
 - for nonstatic data members, 322–323
- default initialized, 493, 752
- default member initialization, 233
- default member initializers
 - aggregate initialization with, 138–141
 - annoyances, 328–330
 - applicability limitations, 329
 - array size deduction, lack of, 330
 - loss of aggregate status, 330
 - loss of triviality, 329–330
 - parenthesized direct-initialization syntax, lack of, 328–329
 - constexpr** functions, 270
 - description of, 318–321
 - potential pitfalls, 326–328
 - inconsistent subobject initialization, 326–328
 - loss of insulation, 326
 - safety of, 6
 - trivial types, 426
 - union interactions, 320–321
 - use cases, 322–325
 - boilerplate repetition, avoiding, 323–325
 - documentation of default values, 325
 - nonstatic data member initialization, 322–323
 - simple struct initialization, 322
- default values, documentation of, 325
- defaulted default constructors, exception specifications and, 1087
- defaulted functions, 522, 649. *See also* deleted functions; rvalue references; **static_assert**
 - annoyances, 42–43
 - description of, 33–36
 - exception specifications and, 1086
 - first declaration of special member function, 34–35
 - further reading for, 44
 - implementation of user-provided special member function, 35–36
 - implicit generation of special member functions, 44–45
 - potential pitfalls, 41–42
 - use cases, 36–41
 - making explicit class APIs with no runtime cost, 38–39
 - physically decoupling interface from implementation, 40–41
 - preserving type triviality, 39–40
 - restoring generation of suppressed special member function, 36–37
- defaulted special member functions. *See* defaulted functions
- defaulted template parameters, 31
- default/value, 215
- defect reports (DR), 432n10, 615n2, 722n8, 1086n2
- defensive checks, 468, 744
- defensive programming, 1024
- defined behavior, 1112–1113
- defining declarations, 729
- definition (of objects), 68
- definitions, 315, 879
- delaying return-type deduction, 1199–1200
- delegating constructors
 - description of, 46–48
 - potential pitfalls, 50–51
 - delegation cycles, 50–51
 - suboptimal factoring, 51
 - use cases, 48–50
- delegation cycles, 50–51
- deleted functions, 33–34, 757, 1086n2. *See also* defaulted functions; rvalue references
 - annoyances, 58–59
 - description of, 53
 - further reading for, 60
 - as trivial, 523
 - use cases, 53–57
 - hiding structural base class member functions, 56–57
 - preventing implicit conversion, 55–56
 - suppressing special member function generation, 53–55
- dependency. *See* data dependency
- dependent base classes. *See* inheriting constructors
- dependent types
 - generic lambdas, 981
 - inheriting constructors, 538
- [[deprecated]] attribute, 14
 - description of, 147–148
 - potential pitfalls, 150
 - use cases, 148–150

- derived classes
 - compile-time visitation, 1050–1054
 - preventing with **final** contextual keyword, 1007, 1014–1015
 - design patterns, 669
 - designated initializers, 139n1
 - destructive move, lack of, 811–812
 - destructors
 - in C++20, 407n3
 - as **constexpr** functions, 463n25
 - declaring special member functions, 34
 - exception specifications and, 1086
 - final** contextual keyword, 1008
 - noexcept** by default, 653–654
 - rvalue references, 752
 - skippable, 464–470
 - user-provided, 755–757
 - vertical encoding, 450
 - devirtualize, 1011
 - diagnostics, compiler, 14–15
 - diffusion, 183n14
 - digit separator ('). *See also* binary literals
 - description of, 152–153
 - further reading for, 154
 - loss of precision in floating-point literals, 154–156
 - use cases, 153
 - dimensional unit types, 863–865
 - direct aggregate initialization, 493
 - direct braced initialized, 455
 - direct initialization, 215, 754
 - explicit conversion operators, 62
 - in factory functions, 240
 - in generic code, 239
 - of members, 241–242
 - for nonstatic data members, 318
 - syntax, 328–329
 - direct initialized, 230
 - direct list initialization, 228–231
 - copy list initialization, compared, 231–232
 - in factory functions, 240
 - in generic code, 239
 - of members, 241–242
 - for nonstatic data members, 318
 - `std::initializer_list`, 555
 - direct mapped, 182n11
 - disabling
 - implicit moves, 244–246
 - named return-value optimization (NRVO), 783–784
 - NRVO, 244–246
 - disambiguators, 28–30
 - discriminated unions, 937–948, 1177–1180
 - divide and conquer, 297
 - documentation of default values, 325
 - double-checked-lock pattern (C++03), 81–82
 - duck typing, 1052
 - dumb data, 668n7
 - duplicate names, loss of access in namespaces, 1056–1058
 - dynamic binding, 1015
 - dynamic dispatch, 1011
 - dynamic exception specifications, 618–619, 1085, 1089, 1090
 - compatibility with **noexcept** specifications, 621
 - noexcept** exception specification, compared, 1101–1102
 - violating, 1093
 - dynamic types, 416
- ## E
- EBCDIC, 129n1
 - Effective C++* (Meyers), 3
 - elaborated type specifiers, 1031–1032
 - embedded development, 145
 - embedded systems, 1101
 - embedding code in C++ programs, 111–112
 - emplacement, 390–391
 - empty-base optimization (EBO), 185, 499, 607, 933, 1028–1030
 - encapsulation
 - of helper types, 85n3
 - of implementation details, 343–344
 - opaque enumerations, 663
 - types within functions, 84–85
 - encoding prefixes, 844
 - entities
 - decltype** use with, 25–26
 - [[deprecated]] attribute, 147–150
 - enum** class
 - annoyances, 351
 - description of, 332–337
 - further reading for, 352
 - potential pitfalls, 344–350
 - bit flags, 347–348
 - collections of named constants, 346–347
 - external use of opaque enumerators, 350
 - iteration, 348–350
 - strong typing can be counterproductive, 344–346
 - scoped enumerations, 335–336
 - underlying types (UTs) and, 337
 - unscoped C++03 enumerations, work-arounds for, 332–333
 - use cases, 337–344
 - encapsulating implementation details, 343–344
 - implicit conversion to arithmetic types, avoiding, 337–339

Index

- enum** class (cont.)
 - use cases (cont.)
 - namespace pollution, avoiding, 339–340
 - overloading disambiguation, 340–343
 - weakly typed C++03 enumerators, drawbacks to, 333–335
- enumerations. *See also* opaque enumerations
 - comparisons in C++20, 335n1
 - underlying types (UTs)
 - description of, 829–830
 - further reading for, 834
 - potential pitfalls, 832–833
 - use cases, 830–832
- enumerators, as compile-time constants, 164
- errors
 - compiler warnings as, 150
 - compile-time, 22
- escalation, 374
- essential behavior, 102
- event-driven callbacks, 603–604
- exception agnostic, 644, 1126
- exception free path, 1134, 1136, 1143
- exception safe, 644, 1126
- exception specifications, 593. *See also* **noexcept**
 - exception specifications
 - conditional, 1091–1092
 - constraints in class hierarchies, 655–658
 - dynamic, 618–619
 - function types and, 1147–1148
 - text-segment size comparison, 1108
 - type system and, 1089n5
 - unconditional, 1085–1089
 - violating, 1093
- exceptions, 615–618, 1104
- excess *N* notation, 155
- executable images, 1135
- execution character sets, 844
- expansion. *See* pack expansion
- expiring objects, 741–742, 749
- expiring value
 - rvalue references, 712–713
 - xvalues*, 721–723
- explicit class APIs, 38–39
- explicit constructors, passing multiple arguments, 250–252
- explicit conversion operators
 - description of, 61–63
 - potential pitfalls, 66–67
 - use cases, 63–65
- explicit instantiation declarations
 - annoyances, 373–375
 - member validity, 374–375
 - unrelated class definitions, 373–374
 - description of, 353–365
 - further reading for, 376
 - illustrative example, 355–359
 - .o files, effect on, 359–365
 - potential pitfalls, 371–373
 - corresponding explicit-instantiation declarations and definitions, 371–372
 - pessimization over optimization, 373
 - use cases, 365–370
 - insulation from client code, 369–370
 - reducing code bloat in object files, 365–369
- explicit instantiation definitions, 353–355, 358–359, 363, 370–375
- explicit instantiation directives, 353n1, 354–355, 369, 375
- explicit template argument specifications, 895
- explicit typenames, 26–27
- explicitly captured, 582–583
- explicitly copied, 583
- explicitly declaring special member functions, 33–34
- exporting bitwise copies of PODs, 479–480
- expression alias, 1146–1147
- expression SFINAE, 29n3, 122, 126
- expression templates, 202–203
- expressions. *See also* lambda expressions
 - compound, **noexcept** operator and, 626–627
 - decltype** use with, 25–26
 - decomposing complex, 391–393
 - rvalue references in, 730–731
 - validation of, 28–30
- extended alignment, 168–170
- extended **friend** declarations. *See* **friend** declarations
- extended **typedef**. *See* aliases
- extern template**
 - annoyances, 373–375
 - member validity, 374–375
 - unrelated class definitions, 373–374
 - description of, 353–365
 - further reading for, 376
 - illustrative example, 355–359
 - .o files, effect on, 359–365
 - potential pitfalls, 371–373
 - corresponding explicit-instantiation declarations and definitions, 371–372
 - pessimization over optimization, 373
 - use cases, 365–370
 - insulation from client code, 369–370
 - reducing code bloat in object files, 365–369
- external definitions for static member variables, 314–315
- external linkages, 307
- external static analysis, control of, 17–18

F

- factory functions, 239–241, 929–930
 - perfect forwarding, 388–389
 - rvalue references, 790, 804–805
 - sink arguments, 778–779
 - uniform initialization in, 239–241
 - user-defined literals (UDLs), 836–838, 851
 - wrapping initialization in, 389–390
- factory operator, 837
- fallible, 1118
- fallible implementation, 1120
- false sharing, 174–179, 182
- fault tolerant, 1123
- fault-tolerant nofail guarantee, 1123
- fences, 82
- Feynman, Richard, 1
- file extensions, 667n5
- final** contextual keyword
 - annoyances, 1028–1030
 - description of, 1007–1014
 - further reading for, 1030
 - override** member-function specifier, interactions with, 1007, 1009–1011
 - potential pitfalls, 1023–1027
 - as contextual keyword, 1023
 - hiding nonvirtual functions, 1026–1027
 - systemic lost reusability, 1023–1026
 - with pure virtual functions, 1008–1009
 - as unsafe, 6
 - use cases, 1014–1023
 - performance of concrete classes, 1015–1017
 - protocol hierarchy performance improvements, 1020–1023
 - restoration of performance lost to mocking, 1017–1020
 - suppressed derivation for portability, 1014–1015
 - with user-defined types (UDTs), 1011–1014
 - with virtual destructors, 1008
 - virtual keyword, interactions with, 1009–1011
 - with virtual member functions, 1007–1008
- fixed-capacity strings, 470–479
- flexible array members, 404n1
- floating-point literals, 154–156, 837, 869–870
- floating-point non-type template parameters, 903n7
- floating-point types, 223
 - big-endian and little-endian layouts, 531–534
 - IEEE 754, 530–534
 - precision of, 155
 - user-defined literals (UDLs), 862
- floating-point-to-integer conversion, 843
- flow of control, 68
- fold expressions, 955n25
- footprint, 1114
 - extern template**, 357
 - POD types, 452, 475
 - rvalue references, 734, 747
- for** loops, range-based, 571–572
 - annoyances, 703–709
 - description of, 679–684
 - further reading for, 709
 - potential pitfalls, 691–703
 - use cases, 684–691
- for** range declaration, 681–682
- forward class declarations, 675
- forward declarations, 662
- forward declared, 664–665
- forwarding references
 - annoyances, 397–400
 - metafunction requirements in constraints, 398–400
 - similarity to rvalue references, 397–398
 - auto** return-type deduction, 1184
 - auto&&**, 383–384
 - description of, 377–385
 - function template argument type deduction, 379–380
 - further reading for, 400
 - generic lambdas, 971
 - identifying, 382–383
 - lambda-capture expressions, 992
 - not forwarding, 384–385
 - potential pitfalls, 394–397
 - hijacking copy constructor, 395–397
 - `std::forward<T>`, enabling move operations, 395
 - template instantiations with string literals, 394–395
 - range-based **for** loops, 680
 - reference collapsing, 380–382
 - rvalue references, 732, 806–807
 - `std::forward<T>`, 385
 - use cases, 386–393
 - decomposing complex expressions, 391–393
 - emplacement, 390–391
 - forwarding expressions to downstream consumers, 386
 - multiple parameter handling, 386–388
 - perfect forwarding for generic factory functions, 388–389
 - wrapping initialization in generic factory functions, 389–390
 - fpermissive flag, 1080n9
- fragmentation, 183n14

Index

- free functions, 442
 - declaring, 58
 - overloading, 570–571
 - range-based **for** loops, 571–572, 707–709
 - `std::initializer_list`, 558
- free operators, 839
- freestanding, 570
- friend** declarations
 - description of, 1031–1033
 - further reading for, 1042
 - potential pitfalls, 1041
 - use cases, 1033–1041
 - curiously recurring template pattern (CRTP), 1041–1054
 - declaring previously declared type as friend, 1033–1034
 - enforcing initialization with PassKey idiom, 1036–1038
 - special type access, 1038–1041
 - type aliases as customization point, 1034–1036
- full expressions, 693
- full specialization, 355–357, 1059–1062
- fully associative, 182n11
- fully constructed, 47
- function call argument list, 912–914
- function calls, hooking, 930–931
- function declarations. *See also* functions
 - `[[carries_dependency]]` attribute, 1000
 - `[[noreturn]]` attribute
 - description of, 95
 - further reading for, 98
 - potential pitfalls, 97–98
 - use cases, 95–97
 - override** keyword in
 - description of, 104–105
 - further reading for, 107
 - potential pitfalls, 106
 - use cases, 105–106
 - trailing return types
 - description of, 124–126
 - further reading for, 128
 - use cases, 126–128
 - virtual**, **override**, **final** keywords, 1009–1011
- function definitions, reducing code size, 1105
- function designators, 815
- function objects, 328, 574, 990
- function parameter packs, 879, 888–892
 - pack expansion, 911–912
 - Rule of Fair Matching, 898–899
- function pointers
 - calls through, 574
 - generic lambda conversion to, 974–975
 - noexcept** and, 1089–1091
 - `[[noreturn]]` attribute misuse, 98
 - readability of declarations, 127–128
 - type system and, 265–266
- function prototypes, 733
- function references, **noexcept** and, 1089–1091
- function template argument matching, 900–901
- function template argument type deduction, 195, 379–380
- function templates
 - instantiation and specialization, 1190–1192
 - preventing misuse of, 118–119
 - return type dependent on parameter type, 126
- functions. *See also* **constexpr** functions; constructors; defaulted functions; deleted functions; destructors; factory functions; special member functions
 - arguments of same type, 564–565
 - auto** return-type deduction
 - annoyances, 1201–1203
 - description of, 1182–1194
 - potential pitfalls, 1200
 - use cases, 1194–1200
 - `bind`, 14
 - `checkBalance`, 15
 - `checksumLength`, 27, 28n1
 - `declval`, 31
 - dynamic exception specifications, 618–619
 - encapsulating types within, 84–85
 - generic variadic functions, 925–926
 - `hereticalFunction`, 17–18
 - `linearInterpolation`, 16–17
 - `loggedSum`, 28, 31
 - `myRandom`, 19
 - noexcept** exception specifications, 619–621
 - overloading, 1089n6
 - preconditions, 18
 - `pure`, 16
 - ref-qualifiers
 - annoyances, 1171–1172
 - description of, 1153–1160
 - further reading for, 1173
 - potential pitfalls, 1170–1171
 - use cases, 1160–1170
 - `reportError`, 15
 - `sortRange`, 28–30
 - `sortRangeImpl`, 28–30
 - `start`, 14
 - `std::kill_dependency`, 999–1000
 - unusable in variadic templates, 953–954
 - variadic function templates, 888
 - function parameter packs, 888–892
 - function template argument matching, 900–901
 - template argument deductions, 894–896

- variadic member functions, 892–894
 - `vectorLerp`, 16–17
 - function-scope static variables
 - annoyances, 80
 - C++03 double-checked-lock pattern, 81–82
 - concurrent initialization, 68–69
 - description of, 68–71
 - destruction, 69
 - further reading for, 81
 - logger example, 69–70
 - multithreaded contexts, 70–71
 - potential pitfalls, 75–80
 - dangerous recursive initialization, 77
 - dependence on order-of-destruction of local objects, 78–80
 - initialization not guaranteed, 75–77
 - recursion subtleties, 77–78
 - use cases, 71–75
 - `function-try-block`, 268
 - functor classes, 574–575
 - functor types, 573
 - functors, 573
 - fundamental alignment, 168
 - fundamental integral types
 - historical perspective on, 93–94
 - long long**
 - description of, 89
 - further reading for, 92
 - potential pitfalls, 91–92
 - use cases, 89–91
 - fundamental types, 803, 1014
 - alignof** operator, 184
 - hidden properties of, 209–210
 - union membership and, 1174
- G**
- GCC
- acquire/release memory barrier, 80n7
 - ambiguity errors, 340n2
 - attribute support
 - [[gnu::cold]] attribute, 15
 - [[gnu::const]] attribute, 16–17, 19
 - [[gnu::warn_unused_result]] attribute, 14–15, 15n7
 - [[gsl::suppress]] attribute, 17–18
 - standardized compiler-specific attributes, 14
 - auto** redeclaration, 1209
 - binary literals, 142n1
 - compiler warnings, 150
 - deduced parameters, 972n1
 - default initialization, 218
 - delegating constructors, 50n2
 - explicit expression of type-consistency in, 28n1
 - fpermissive flag, 1080n9
 - incompatibly specified alignment, 177
 - indirect calls, 947n22
 - inline namespaces, 1061n3
 - namespace-qualified name support, 13n2
 - nonrecursive **constexpr** algorithms, 962n3
 - pointer compatibility, 1090n7
 - reducing code size, 1104n16, 1110
 - stack unwinding, 621n4
 - template instantiation with deduced return type, 1192n3
 - trivial copy/move constructors, 528n62
 - underspecifying alignment, 176
 - generalized attribute support. *See* attribute support
 - generalized plain old data types (PODs)
 - annoyances, 521–529
 - C++ Standard not stabilized, 521–527
 - standard type traits unreliable, 527–528
 - `std::pair` and `std::tuple` of PODs are not PODs, 528–529
 - bit representation, 530–534
 - C++03 POD types, 412–415
 - C++11 POD types, 415–417
 - description of, 401–402
 - further reading for, 530
 - future direction, 438–439
 - potential pitfalls, 479–521
 - abuse of **reinterpret_cast**, 506–519
 - aggressive use of `offsetof`, 520–521
 - conflating arbitrary and indeterminate values, 493–497
 - exporting bitwise copies of PODs, 479–480
 - ineligible use of `std::memcpy`, 497–501
 - `memcpy` usage on **const** or reference sub-objects, 489–493
 - misuse of unions, 505–506
 - naive copying other than `std::memcpy`, 501–505
 - requiring PODs or trivial types, 480–482
 - sloppy terminology, 488–489
 - wrong type traits, 482–488
 - privileges, 402–412
 - bitwise copyability, 409–410
 - contiguous storage, 405
 - object lifetime begins at allocation, 407–409
 - `offsetof` macro usage, 410–412
 - predictable layout, 405–407
 - standard-layout class special properties, 420–425
 - standard-layout types, 417–420
 - trivial subcategories, 429–436
 - trivial types, 425–429

Index

generalized plain old data types (PODs) (cont.)
 type traits, 436–438
 use cases, 439–479
 compile-time constructible, literal types, 462–464
 fixed-capacity string elements, 476–479
 fixed-capacity strings, 470–475
 navigating compound objects with
 offsetof, 456–460
 secure buffers, 460–462
 skippable destructors, 464–470
 translating C++-only types to C, 452–456
 vertical encoding for non-trivial types, 448–452
 vertical encoding within a union, 439–448
general-purpose machines, 93
generic code
 member initialization in, 241–242
 uniform initialization in, 238–239
generic expressions, validating with `decltype`, 28–30
generic factory functions. *See* factory functions
generic lambdas
 annoyances, 981–984
 cannot use full power of template-argument deduction, 981–982
 difficulty constraining multiple arguments, 983–984
 constraints on deduced parameters, 970–973
 conversion to function pointer, 974–975
 description of, 968–975
 explicit parameter types, 193–194
 further reading for, 985
 lambda captures, 969–970
 mutable closures, 969–970
 potential pitfalls, 981
 use cases, 975–981
 applying lambdas to tuple elements, 975–976
 conditional instantiation, 979–981
 recursive lambdas, 977–979
 reusable lambda expressions, 975
 terse, robust lambdas, 976–977
 variadic, 973–974
generic programming, 615
generic types, 28, 878
generic value-semantic types (VSTs), creating, 762–767
generic variadic functions, 925–926
glvalues, 717
GNU, nonstandard primitives, 956n27
[[gnu::cold]] attribute, 15
[[gnu::const]] attribute, 16–17, 19

[[gnu::pure]] attribute, 14, 16
[[gnu::warn_unused_result]] attribute, 14–15, 15n7
golden files, 114
greater-than operator (>), 21–22
grouping macros, 520
gsl::span type, 17
[[gsl::suppress]] attribute, 17–18
guaranteed copy elision, 216n1, 567n2, 648n11, 717n4, 790–791, 805n30, 807n31, 827n54, 1163n1
Guidelines Support Library, 17

H

handle types, 792
hard UB. *See* language undefined behavior
header files, 41n2, 663–665
header-only library, 1067
helper functions. *See* functions
helper types, encapsulation of, 85n3
hereticalFunction, 17–18
hidden friend idiom, 472
hidden properties of fundamental types, 209–210
hiding
 member functions, 56–57
 nonvirtual functions, 1026–1027
hierarchical reuse, 1012, 1023–1026
higher-order functions, 125
high-level value semantic types (VSTs), creating, 751–762
hooking function calls, 930–931
horizontal microcode, 445n17
hot paths, 1134–1136, 1139, 1142
Hyrum’s law, 85, 1012, 1014, 1036

I

ICC
 incompatibly specified alignment, 177
 underspecifying alignment, 176
identity closure objects, 968
identityInt, 968
id-expression, 25, 780
IEEE 754 floating-point types, 530–534
if constexpr language feature, 641n10
ill formed, 120, 1067, 1071, 1077, 1203
ill formed, no diagnostic required (IFNDR), 1000
 constexpr functions, 262–263
 delegating constructors, 50
 enum class, 350
 incompatibly specified alignment, 176–177
 inline namespaces, 1067, 1072, 1079
 [[noreturn]] attribute, 97
 opaque enumerations, 666, 675–676, 832
 static assertions in templates, 116–118
 variadic templates, 900

- immutable types, optimizing, 1167–1170
- imperative programming, 959
- implementation defined, 1093
 - alignments, 168–169
 - NULL macro, 100
 - opaque enumerations, 660
- implementation inheritance, avoiding boilerplate code with, 540–541
- implementation-defined behavior
 - enum** class, 335
 - limits on, 295
 - unrecognized attributes, 12, 18–19
- implementation-defined types, 202, 501
- implicit **const**-qualification, 300
- implicit constructors, inheriting, 546–549
- implicit conversion, 66, 223n3
 - to arithmetic types, avoiding, 337–339
 - preventing, 55–56, 201
- implicit generation of special member functions, 44–45
- implicit moves
 - disabling, 244–246
 - in **return** statements, 735–737
- implicitly captured, 582–583
- implicitly declared, 522
- implicitly declared default constructors, 568–570
- implicitly movable entities, 735n13
- in contract, 1122–1123
- in place, 734
- indentation of string literals, 112–113
- indeterminate values, 435, 493–497
- infallible, 1118
- infallible implementation, 1118–1123
- inheritance
 - improving concrete class performance, 1015–1017
 - preventing with **final** contextual keyword, 1008, 1012
- inheriting constructors. *See also* default member initializers; defaulted functions; delegating constructors; deleted functions; forwarding references; **override** member-function specifier; variadic templates
 - annoyances, 549–552
 - access levels same as in base class, 549–551
 - cannot select individually, 549
 - flawed initial specification, 551–552
 - description of, 535–539
 - potential pitfalls, 546–549
 - implicit constructors, 546–549
 - new constructors in base class alters behavior, 546
 - use cases, 539–545
 - implementation inheritance, avoiding boilerplate code, 540–541
 - reusable functionality through mix-ins, 545
 - strong **typedef** implementation, 541–544
 - structural inheritance, avoiding boilerplate code, 540
- init capture, 986
- initialization. *See also* aggregate initialization; braced initialization; copy initialization; copy list initialization; default initialization; default member initializers; direct initialization; direct list initialization; list initialization; **std::initializer_list**; uniform initialization; value initialization
 - of bit fields, 329n4
 - concurrent, 68–69
 - constant, 75
 - enforcing with PassKey idiom, 1036–1038
 - recursive, 77–78, 163–165
 - of simple structs, 322
 - subobject, inconsistency in, 326–328
 - of subobjects, inconsistency in, 326–328
 - thread-safe function-scope static variables, 68–69
 - trivial default, 1087
 - of variables, 200
 - wrapping in factory functions, 389–390
- initializer lists. *See* **std::initializer_list**
- initializer_list**. *See* **std::initializer_list**
- initializers, undefined behavior with **constexpr** variables, 306–307
- inline namespace sets, 1056
- inline namespaces. *See also* **alignas** specifier
 - annoyances, 1079–1082
 - code factoring, impeding, 1079–1082
 - one-to-one relationship with namespaces, 1082
 - argument-dependent lookup (ADL) interoperability, 1058–1059
 - class template specialization, 1059–1061
 - description of, 1055–1062
 - duplicate names, loss of access to, 1056–1058
 - further reading for, 1083
 - potential pitfalls, 1076–1079
 - inconsistent use of inline keyword, 1079
 - lack of scalability, 1076–1077
 - library evolution, 1077–1079
 - reopening, 1061–1062

Index

- inline namespaces (cont.)
 - use cases, 1062–1076
 - ABI link safety and build modes, 1071–1074
 - API migration, facilitating, 1062–1067
 - link-safe ABI versioning, 1067–1071
 - selective **using** directives for short-named entities, 1074–1076
 - versioning case study, 1083–1084
 - inline** specifier, 262–265
 - in-process value-semantic type (VST), 1034
 - instantiation, conditional, 979–981
 - instantiation time, 120
 - instruction selection, 1136
 - insulate, 96, 299, 369, 665
 - insulation, 299, 1200
 - from client code, 369–370
 - loss of, 326
 - opaque enumerations, 663, 665
 - int** type, relative size of, 91–92
 - integer literals, 837, 869–870. *See also* binary literals
 - integer types. *See* integral types
 - integer-to-floating-point conversion, 843
 - integral constant expressions. *See also* **alignof** operator
 - alignas** specifier, 169
 - alignof** operator, 184
 - constexpr** variables as alternative, 307–310
 - requirements, 303
 - integral constants, 223
 - integral promotion, 334, 726, 832–833
 - integral types. *See also* fundamental integral types
 - enumerations, 829
 - reinterpret_cast** keyword, 510–512
 - interface inheritance, 541
 - interface test (C++11), 275
 - interface traits, 482–489
 - interface widening, 1021
 - interfaces. *See also* inheriting constructors
 - adaptation with lambda expressions, 597–598
 - gs1::span in, 17n10
 - physically decoupling from implementation, 40–41
 - internal linkage, 307
 - intra-thread dependencies, 998
 - invocable, 482, 526, 986
 - ISO C++ Standards Committee, 4
 - iteration
 - enum** class and, 348–350
 - lack of access to state, 703–706
 - over all container elements, 684–685
 - over fixed number of objects, 565–566
 - over simple values, 690–691
 - sentinels, lack of support, 706–707
 - iterators, vectors of, 26–27
- ## K
- keywords. *See also* functions; **using** declarations
 - adding new, 1023
 - auto**
 - annoyances, 212–213
 - description of, 195–199
 - further reading for, 214
 - potential pitfalls, 204–212
 - use cases, 200–203
 - constexpr**, 75n5, 304n1, 316n8
 - decltype**
 - annoyances, 31
 - description of, 25–26
 - potential pitfalls, 30
 - use cases, 26–30, 28n1
 - final**
 - annoyances, 1028–1030
 - description of, 1007–1014
 - further reading for, 1030
 - potential pitfalls, 1023–1027
 - as unsafe, 6
 - use cases, 1014–1023
 - nullptr**
 - description of, 99–100
 - further reading for, 103
 - use cases, 100–103
 - override**
 - description of, 104–105
 - further reading for, 107
 - potential pitfalls, 106
 - safety of, 5
 - use cases, 105–106
 - register**, 195n1
 - reinterpret_cast**, 506–519
 - static_assert**
 - annoyances, 123
 - description of, 115–118
 - further reading for, 123
 - potential pitfalls, 120–122
 - use cases, 118–119
- ## L
- L1 cache, 181–183
 - L2 cache, 181–183
 - L3 cache, 181–183
 - Lakos Rule, 1116
 - lambda body, 581, 595–597
 - lambda captures, 577, 581–591, 919, 969–970
 - lambda closure, 584
 - lambda declarators, 591–595

- lambda expressions
 - annoyances, 611–614
 - capturing ***this** by copy, 611–612
 - debugging, 611
 - mixing immediate and deferred-execution code, 612–613
 - trailing punctuation, 613–614
 - configuring algorithms via, 86–87
 - decltype(auto)** placeholders and, 1206
 - deduced return types for, 1189–1190, 1197–1198
 - description of, 573–597
 - further reading for, 614
 - generic lambdas
 - annoyances, 981–984
 - description of, 968–975
 - further reading for, 985
 - potential pitfalls, 981
 - use cases, 975–981
 - local/unnamed types, 83–84
 - parts of, 577–578
 - closures, 578–581
 - lambda body, 595–597
 - lambda captures, 581–591
 - lambda declarators, 591–595
 - lambda introducers, 581–591
 - potential pitfalls, 607–611
 - dangling references, 607–608
 - local variables in unevaluated contexts, 610–611
 - mixing captured and noncaptured variables, 609
 - overuse, 609
 - use cases, 597–607
 - emulating local functions, 598–599
 - emulating user-defined control constructs, 599–600
 - event-driven callbacks, 603–604
 - interface adaptation, partial application, currying, 597–598
 - recursion, 604–605
 - stateless lambdas, 605–607
 - with `std::function`, 601–603
 - variables and control constructs in expressions, 600–601
- lambda introducers, 581–591, 986
- lambda-capture expressions. *See also* **auto** variables; braced initialization; forwarding references; lambda expressions; rvalue references
 - annoyances, 993–994
 - difficulty of synthesizing **const** data members, 993–994
 - `std::function` supports only copyable callable objects, 994
 - description of, 986–988
 - further reading for, 995
 - potential pitfalls, 992–993
 - use cases, 988–992
 - capturing modifiable copy of **const** variable, 990–992
 - moving objects into closure, 988–989
 - providing mutable state for closure, 989–990
- lambda-capture list, 919–921
- language undefined behavior, 1115
- libraries
 - Guidelines Support Library*, 17
 - Ranges Library, 391–393, 686n4, 687n5
 - resilience to code changes, 203
- library undefined behaviors, 1115
- lifetime extensions, 1162, 1213
 - prvalues*, 720
 - range-based **for** loops, 680, 691–696
 - temporary objects, 819–820
- limerick in C++ Language Standard, 1081–1082
- linear search in variadic templates, 957
- `linearInterpolation` function, 16–17
- linkage, 83
- link-safe ABI versioning, 1067–1071
- link-time optimization, 1094, 1143
- Liskov, Barbara, 1026, 1030
- Liskov Substitution Principal (LSP), 1030
- list initialization
 - braced initialization and, 215, 233–234
 - deducing, 210–211
- list initialized literal types, 260
- literal types, 278–284
 - aggregate types as, 279–280
 - array types as, 280
 - compile-time constructible, 462–464
 - in constant expressions, 260–261, 273, 277–278
 - constexpr** constructors and, 281
 - cv-qualifiers as, 280
 - identifying, 282–284
 - pointers as, 281
 - reference types as, 279
 - scalar types as, 278
 - `std::initializer_list`, 556
 - `std::is_literal_type`, 283n14
 - trivially destructible types as, 431
 - user-defined, 280
 - variable templates of, 302
 - void** return type as, 280
- literals
 - binary
 - description of, 142–143
 - further reading for, 146
 - use cases, 144–146

Index

literals (cont.)

- digit separators (') in
 - description of, 152–153
 - further reading for, 154
 - loss of precision in floating-point literals, 154–156
 - use cases, 153
- floating-point, 154–156, 837, 869–870
- integer, 837, 869–870
- raw string
 - description of, 108–111
 - potential pitfalls, 112–114
 - use cases, 111–112
- Unicode
 - description of, 129–130
 - potential pitfalls, 130–132
 - use cases, 130
- user-defined
 - annoyances, 869–871
 - description of, 835–853
 - further reading for, 872
 - potential pitfalls, 867–869
 - use cases, 853–867

little-endian **float** layouts, 531–534

local declarations, 662, 675–677

local functions, emulating, 598–599. *See also* lambda expressions

local scope. *See* block scope

local variables in unevaluated contexts, 610–611

locality of reference, 181, 742, 773n26

local/unnamed types. *See also* **decltype**; lambda expressions

- description of, 83–84
- use cases, 84–87
 - configuring algorithms via lambda expressions, 86–87
 - encapsulating types within functions, 84–85
 - instantiating templates with local function objects as type arguments, 85–86

loggedSum function, 28, 31

logical optimization, 365

logical or (||) operator, 265

long long integral type

- description of, 89
- further reading for, 92
- potential pitfalls, 91–92
- use cases, 89–91

long type, relative size of, 91–92

long-distance friendship, 1035–1036, 1041

loops. *See* range-based **for** loops

lossy conversions, restrictions on, 222–224

low-level value-semantic types (VSTs), creating, 742–751

lvalue references, 26, 716, 1118, 1133

- in C++11/14, 717–720
- declarations prior to C++11, 815–818
- evolution of, 807, 813–828
- forbidding operations on, 1165–1167
- implicit moves in **return** statements, 735–737
- range-based **for** loops, 703
- rvalue references, introduction to, 710–711

lvalue-to-rvalue conversion, 501

M

macro-defined namespaces, 1083–1084

macro-invocation syntax, 248–249

macros. *See also* functions

- alternatives to, 286–287
- offsetof
 - aggressive usage, 520–521
 - navigating compound objects, 456–460
 - POD type usage, 410–412
 - support for, 423–425

magic constants, 308

managed allocators, 1021–1022

mandatory RVO, 807n31

mangled names, 1056, 1114n24

manifestly constant evaluated, 258n1

mantissa, 155

materialization, 717

materialize, 1163n1

maximal fundamental alignment, 193

mebibyte conversion, 286–287

mechanisms, 51

member functions

- constexpr** as implicitly **const**-qualified, 300
- hiding, 56–57
- overriding, 105–106
- variadic member functions, 892–894

member initialization lists, 230

member initializer lists

- copy list initialization in, 249–250
- delegating constructors, 46
- nonstatic data member initialization, 318
- pack expansion, 917–918

member initializers, default

- annoyances, 328–330
 - applicability limitations, 329
 - array size deduction, lack of, 330
 - loss of aggregate status, 330
 - loss of triviality, 329–330
 - parenthesized direct-initialization syntax, lack of, 328–329
- description of, 318–321

- potential pitfalls, 326–328
 - inconsistent subobject initialization, 326–328
 - loss of insulation, 326
- safety of, 6
- union interactions, 320–321
- use cases, 322–325
 - boilerplate repetition, avoiding, 323–325
 - documentation of default values, 325
 - nonstatic data member initialization, 322–323
 - simple **struct** initialization, 322
- `memcpy`. *See* `std::memcpy`
- memory allocation, 75n4, 181–183
 - in C++11, 763n25
 - monotonic, 190–193, 1021–1022
 - secure buffers, 460–462
- memory barriers, 80n7
- memory diffusion, 628, 788
- memory leak, 74
- memory models, synchronization paradigms for, 998
- `memory_order_acquire`, 1005n2
- `memory_order_consume`, 1005n2
- memory-fence instructions, 999–1000
- metafunctions, 469, 963
 - forwarding references, 381
 - requirements in constraints, 398–400
 - `std::remove_cvref<T>`, 399n6
- metaparameters, 948
- metaprogramming, 876, 963–964
- metaprograms, 257
- Meyers, Scott, 3
- Meyers singleton, 71–75
- microbenchmarks, 1137–1141
- mixed-mode builds, 1073
- mix-ins, reusable functionality through, 545
- mocking, 1017–1020
- mocks, 1017–1020
- modifiable *rvalues*, 820–821
- modules, 85n3, 1041
- monotonic allocators, 1021–1022
- monotonic memory allocation, 190–193
- Moore’s law, 93n5
- most vexing parse, avoiding, 237–238
- move assignable, 524
- move assignment, 750, 756
- move construction, 750
- move constructors
 - literal types and, 281
 - noexcept** operator and, 653–654
 - rvalue references, 710, 714, 732–733
 - RVO and NRVO requirements, 804–805
 - `std::list`, 1114
 - as trivial, 437
 - user-provided, 760
- move operations
 - avoiding, 183n14
 - deleted functions, 53
 - destructive move, lack of, 811–812
 - enabling with `std::forward<T>`, 395
 - noexcept** operator, 627–631, 658–659
 - on noncopyable types, 788–791
 - nonthrowing, 1094–1097
 - objects into closure, 988–989
 - as optimization of copying, 741–767
 - rvalue references, 710, 714–715
 - some equivalent to copies, 788
 - throwing in, 787
 - wrappers for **noexcept**, 1099–1101
- move semantics
 - necessity of, 821–823
 - rvalue references, 710, 715–716
- move-assignment operator
 - rvalue references, 710, 714, 733
 - user-provided, 760–761
- moved-from objects
 - inconsistent expectations, 794–803
 - overly strict requirements, 807–811
 - rvalue references, 714–715, 788, 807–812
- moved-from state, 789, 791–803
- move-only types, 570, 641, 644
 - implementing without `std::unique_ptr`, 791–794
 - rvalue references, 716, 768–771, 790
- moving iterators, return types of, 1211–1212
- MSVC
 - auto** redeclaration, 1209
 - compiler warnings, 150
 - deduced parameters, 972n1
 - incompatibly specified alignment, 177
 - reducing code size, 1104n16, 1111
 - stack unwinding, 621n4
 - standardized compiler-specific attributes, 14
 - trivial copy/move constructors, 528n62
 - underspecifying alignment, 176
- multiple arguments
 - constraining, 983–984
 - passing to explicit constructors, 250–252
- multiple parameters, handling, 386–388
- multiple **return** statements, 1185–1187
- multithreaded programs, avoiding false sharing, 174–175
- multithreading context, 68, 70–71
- mutable closures, 969–970
- mutable state, providing for closure, 989–990
- `myRandom` function, 19

Index

N

- naked literals, 839–846, 849, 851, 861
- name collisions, 870
- name mangling, 1067, 1089n5, 1149
- named constants, **enum** class for collections of, 346–347
- named functions, 66–67
- named return-value optimization (NRVO), 805n30
 - disabling, 244–245, 783–784
 - requires declared copy/move constructors, 804–805
 - rvalue references, 734, 736–739, 790, 804
- namespace-qualified names, 13n2
- namespaces
 - inline
 - annoyances, 1079–1082
 - description of, 1055–1062
 - further reading for, 1083
 - potential pitfalls, 1076–1079
 - use cases, 1062–1076
 - versioning case study, 1083–1084
 - pollution, avoiding, 339–340
- NaN (Not a Number) representations, 530–534
- narrow contracts, 715, 1021, 1112–1118, 1122
- narrowing aggregate initialization, 247
- narrowing conversions, 1091
 - allowing, 247–248
 - restrictions on, 222–224
- narrowing the contract, 793
- natural alignment, 179–181, 193, 831
- negative testing, 794
- nested containers, 22
- nested namespaces. *See* inline namespaces
- new** expressions, **decltype(auto)** in, 1210
- new** handler, 193
- new line encoding, 113–114
- nibbles, 153–154
- [[no_unique_address]] attribute, 1029n15
- noexcept** exception specifications, 619–621
 - annoyances, 1143–1150
 - ABI changes in future C++ versions, 1148–1149
 - code duplication, 1144–1147
 - exception specifications not part of function’s type, 1147–1148
 - optimization conflated with reducing code size, 1143–1144
 - SFINAE triggering, 1149–1150
 - compatibility with dynamic specifications, 621
 - conditional exception specifications, 1091–1092
 - constraints for virtual functions, 632–634
 - description of, 1085–1094
 - efficiencies with, 1093–1094
 - function pointers and references, 1089–1091
 - further reading for, 1151–1152
 - potential pitfalls, 1112–1143
 - accidental terminate, 1124–1128
 - conflating with nofail, 1116–1123
 - forgotten **noexcept** operator, 1129–1130
 - imprecise expressions, 1130–1134
 - overly strong contract guarantees, 1112–1116
 - theoretical opportunities for performance improvement, 1136–1143
 - unrealizable runtime performance benefits, 1134–1136
 - unconditional exception specifications, 1085–1089
 - use cases, 1094–1111
 - noexcept** swap definition, 1097–1099
 - nonthrowing move operations, 1094–1097
 - reduction of object-code size, 1101–1111
 - wrappers for **noexcept** move operations, 1099–1101
 - violating, 1093
- noexcept** operator
 - annoyances, 650–658
 - change in unspecified behavior when `std::vector` grows, 652–653
 - destructors, not move constructors, **noexcept** by default, 653–654
 - exception specification constraints in class hierarchies, 655–658
 - older compilers invade **constexpr** function bodies, 654–655
 - sensitivity for direct usage, 650–651
 - strong exception-safety guarantee, 651–652
- C Standard Library functions and, 631–632
- compatibility of dynamic and **noexcept** exception specifications, 621
- compiler-generated special member functions, 621–626
- compound expressions and, 626–627
- constraints for virtual functions, 632–634
- description of, 615–634
- dynamic exception specifications for functions, 618–619
- exception specifications with, 1092–1093
- forgetting in **noexcept** exception specifications, 1129–1130
- further reading for, 658
- move operations, 627–631, 658–659
- noexcept** exception specifications for functions, 619–621
- operator-produced exceptions, 615–618

- potential pitfalls, 647–650
 - direct usage, 647–649
 - function bodies, lack of consideration, 649–650
 - use cases, 634–647
 - appending elements to `std::vector`, 634–639
 - enforcing `noexcept` contract using `static_assert`, 639–640
 - `std::move_if_noexcept`, 640–644
 - `std::vector::push_back(T&&)`, 644–647
 - noexcept** swap, defining, 1097–1099
 - nofail functions, 1116–1123
 - nofail guarantee, 1117, 1122–1123
 - noncaptured variables, mixing with captured, 609
 - noncopyable types, making movable, 788–791
 - nondefining declarations, 729
 - nonintegral symbolic numeric constants, 310–311
 - nonprimitive functionality, 67
 - nonrecursive **constexpr** algorithms, 961–962
 - nonreporting contracts, 1120–1122
 - nonreporting functions, 1119, 1122
 - nonstatic data members
 - auto** not allowed, 212
 - constexpr** variables, 305
 - initialization, 318, 322–323
 - nonthrowing move operations, 1094–1097
 - non-trivial, 1011
 - non-trivial constructors, union membership and, 1174
 - non-trivial destructors, 1101–1104, 1118, 1136
 - non-trivial special member functions, union type and, 1174–1181
 - non-trivial types, vertical encoding for, 448–452
 - non-trivially destructible, 1102–1109, 1137
 - non-type parameters, 902
 - non-type template parameter packs, 901–903
 - non-type template parameters, 901–903, 903n7
 - nonvirtual functions, hiding, 1026–1027
 - [[noreturn]] attribute, 13. *See also* attribute support
 - description of, 95
 - further reading for, 98
 - potential pitfalls, 97–98
 - inadvertently break working programs, 97
 - misuse on function pointers, 98
 - use cases, 95–97
 - compiler diagnostics, 95–96
 - runtime performance, 96–97
 - normative wording, 808
 - NRVO. *See* named return-value optimization (NRVO)
 - null address, 99–102
 - NULL macro, 100
 - null pointer value, 743
 - null statements, 268
 - null terminated strings, 743
 - null-pointer-literal. *See* **nullptr** keyword
 - nullptr** keyword
 - description of, 99–100
 - further reading for, 103
 - use cases, 100–103
 - overload resolution, 101–102
 - overloading literal null pointer, 102–103
 - type safety, 100–101
 - numeric literals
 - digit separators (') in
 - description of, 152–153
 - further reading for, 154
 - loss of precision in floating-point literals, 154–156
 - use cases, 153
 - user-defined, 858–862
- O**
- .o files
 - extern template**, effect on, 359–365
 - reducing code bloat, 365–369
 - object factories, 929–930
 - object files
 - extern template**, effect on, 359–365
 - reducing code bloat, 365–369
 - object invariants, 539, 742
 - object orientation, 1015
 - object representation
 - POD types, 405
 - reinterpret_cast** keyword, 510, 515–516
 - object-oriented design, vertical encoding comparison, 440–441
 - object-oriented programming, 1015
 - objects
 - creating, 516n42
 - iterating over fixed number, 565–566
 - moving into closure, 988–989
 - reducing code size, 1101–1111, 1143–1144
 - resource-owning, passing around, 771–775
 - `std::initializer_list<E>` initialization, 559
 - strengthening alignment, 169–170
 - obsolete entities, [[deprecated]] attribute for
 - description of, 147–148
 - potential pitfalls, 150
 - use cases, 148–150
 - ODR-used, 581–582, 590, 988n2, 1081
 - offsetof macro
 - aggressive usage, 520–521
 - navigating compound objects, 456–460
 - POD type usage, 410–412
 - support for, 423–425

Index

- one-definition rule (ODR), 263, 1072, 1079
 - constexpr** functions, 263
 - extern template**, 374
 - violating, 1189
 - opaque declarations, 660–662
 - opaque enumerations
 - annoyances, 677–678
 - description of, 660–663
 - external usage, 350, 832
 - further reading for, 678
 - potential pitfalls, 675–677
 - inciting local enumeration declarations, 677
 - redeclaring externally defined enumeration locally, 675–677
 - use cases, 663–675
 - cookies, 669–675
 - insulating some external clients from enumerator list, 665–668
 - within header files, 663–665
 - operands, for **decltype**
 - () versus (()) notation for, 30
 - entities, 25
 - expressions, 25–26
 - operators. *See also* keywords
 - || (logical or), 265
 - alignof**
 - annoyances, 193–194
 - description of, 184
 - fundamental types, 184
 - use cases, 186–193
 - user-defined types, 185–186
 - bitwise right-shift, 21–24
 - braced lists and, 254–255
 - decltype**
 - annoyances, 31
 - description of, 25–26
 - potential pitfalls, 30
 - use cases, 26–30, 28n1
 - explicit
 - description of, 61–63
 - potential pitfalls, 66–67
 - use cases, 63–65
 - greater-than (>), 21–22
 - noexcept**
 - annoyances, 650–658
 - description of, 615–634
 - further reading for, 658
 - move operations, 658–659
 - potential pitfalls, 647–650
 - use cases, 634–647
 - sequencing, 265
 - UDL operators, 840–842
 - cooked, 843–845
 - raw, 845–849
 - templates, 849–851
 - optimization
 - attributes for
 - hints for additional optimization opportunities, 15–16
 - statement of explicit assumptions, 16–17
 - builder classes, 1167–1170
 - conflating with reducing code size, 1143–1144
 - immutable types, 1167–1170
 - optimized metaprogramming algorithms, 963–964
 - ordered after, 998
 - ordinary character types, 501–505
 - out clause, 1118
 - out of contract, 744, 1117
 - outermost expressions, 820
 - over-aligned, 185
 - overhead costs, single-threaded applications, 80
 - overload resolution
 - deleted functions, 53
 - nullptr** keyword, 101–102
 - priorities, 730
 - rvalue references, 713
 - `std::initializer_list`, 561
 - user-defined literals (UDLs), 841
 - overloading, 741
 - free functions, 570–571
 - functions, 1089n6
 - improving disambiguation, 340–343
 - null pointer, 102–103
 - reference types, 727–730
 - overloads, ref-qualified, 1171–1172
 - overly strong contract guarantees, 1112–1116
 - override** member-function specifier
 - as contextual keyword, 1023
 - description of, 104–105
 - final** contextual keyword, interactions
 - with, 1007, 1009–1011
 - further reading for, 107
 - potential pitfalls, 106
 - safety of, 5
 - use cases, 105–106
 - overriding, 539
 - member functions, 105–106
 - preventing with **final** contextual keyword, 1007
 - owned resources, 741, 803–804
- ## P
- pack expansion, 882, 908–911, 925, 964
 - alignas** specifier, 921–922
 - attribute lists, 922
 - base specifier list, 915–917
 - braced initializer lists, 912–914

- cannot use unexpanded, 956
- disallowed, 924
- expansion is rigid and requires verbose support code, 957
- function call argument list, 912–914
- function parameter packs, 911–912
- lambda-capture list, 919–921
- limitations on contexts, 954–955
- member initializer list, 917–918
- sizeof**... expressions, 923
- template argument list, 914
- template parameter list, 923–924
- pack expansion context, 883, 929
- Packet class, 27–28
- Packet::checksumLength, 27, 28n1
- padding bytes, 475
- pages, 181–183
- pair mismatches, 699n8
- parameter count, 597
- parameter declarations, 888, 1000
- parameter pack expansion, 590
- parameter packs, 879, 964
 - function parameter packs, 888–892
 - non-type template parameter packs, 901–903
- pack expansion, 908–911, 925
 - alignas** specifier, 921–922
 - attribute lists, 922
 - base specifier list, 915–917
 - braced initializer lists, 912–914
 - cannot use unexpanded, 956
 - disallowed, 924
 - expansion is rigid and requires verbose support code, 957
 - function call argument list, 912–914
 - function parameter packs, 911–912
 - lambda-capture list, 919–921
 - limitations on contexts, 954–955
 - member initializer list, 917–918
 - sizeof**... expressions, 923
 - template argument list, 914
 - template parameter list, 923–924
- Rule of Fair Matching, 898–899
- Rule of Greedy Matching, 896–898
- template template parameter packs, 903–908
- type template parameter packs, 880–884
- variable templates, 159
- parameter types, return types dependent on, 126
- parameterized constants, 160–161
- parameters
 - constexpr** functions, 277–278
 - handling multiple, 386–388
- parentheses with **decltype** operands, 25, 30
- partial application, 597–598
- partial class template specialization, 963
- partial implementation, 1021
- partial implementation classes, 540
- partial ordering of class template specialization, 886
- partial specialization, 529, 884–887
- partially constructed, 47
- passing resource-owning objects, 771–775
- PassKey idiom
 - enforcing initialization with, 1036–1038
 - special type access with, 1039–1041
- perfect forwarding, 807, 942, 1131, 1198
 - expressions to downstream consumers, 386
 - in factory functions, 240
 - hijacking copy constructors, 395–397
 - lambda-capture expressions, 992
- perfectly forwarded, 992
- performance
 - of concrete classes, 1015–1017
 - of protocol hierarchy, 1020–1023
 - theoretical opportunities for improvement, 1136–1143
 - unrealizable runtime benefits, 1134–1136
- pessimization, returning **const rvalues**, 786–787
- physical dependency, 374
- physical design, 663
- physical memory, 1135
- physical optimization, 365
- pi, 160
- pipelined, 1137
- placeholder types, 195
- placeholders, 1182
 - conversion functions, 1193–1194
- decltype(auto)**
 - annoyances, 1213
 - description of, 1205–1210
 - potential pitfalls, 1212–1213
 - use cases, 1210–1212
 - in trailing return types, 1189
- placement **new**, 452, 638, 940, 1175, 1180
- placement of attributes, 13
- plain old data (POD). *See* POD types
- platonic values, 742
- pmr allocators in C++17, 763n25
- POD types
 - annoyances, 521–529
 - C++ Standard not stabilized, 521–527
 - standard type traits unreliable, 527–528
 - std::pair and std::tuple of PODs are not PODs, 528–529
 - bit representation, 530–534
 - C++03 POD types, 412–415
 - C++11 POD types, 415–417
 - description of, 401–402
 - further reading for, 530

Index

POD types (cont.)

- future direction, 438–439
- potential pitfalls, 479–521
 - abuse of **reinterpret_cast**, 506–519
 - aggressive use of **offsetof**, 520–521
 - conflating arbitrary and indeterminate values, 493–497
 - exporting bitwise copies of PODs, 479–480
 - ineligible use of **std::memcpy**, 497–501
 - memcpy** usage on **const** or reference sub-objects, 489–493
 - misuse of unions, 505–506
 - naive copying other than **std::memcpy**, 501–505
 - requiring PODs or trivial types, 480–482
 - sloppy terminology, 488–489
 - wrong type traits, 482–488
- privileges, 402–412
 - bitwise copyability, 409–410
 - contiguous storage, 405
 - object lifetime begins at allocation, 407–409
 - offsetof** macro usage, 410–412
 - predictable layout, 405–407
- standard-layout class special properties, 420–425
- standard-layout types, 417–420
- trivial subcategories, 429–436
- trivial types, 425–429
- type traits, 436–438
- use cases, 439–479
 - compile-time constructible, literal types, 462–464
 - fixed-capacity string elements, 476–479
 - fixed-capacity strings, 470–475
 - navigating compound objects with **offsetof**, 456–460
 - secure buffers, 460–462
 - skippable destructors, 464–470
 - translating C++-only types to C, 452–456
 - vertical encoding for non-trivial types, 448–452
 - vertical encoding within a union, 439–448
- POD-struct, 405–407, 412–415
- POD-union, 412–415
- pointer semantics, 558–559
- pointer types, deducing, 197–198
- pointers. *See also* function pointers
 - as literal types, 281
 - noexcept** and, 1089–1091
 - nullptr** keyword
 - description of, 99–100
 - further reading for, 103
 - use cases, 100–103
 - reinterpret_cast** keyword, 506–519
 - semantics, 558–559
 - smart pointers, 948–951
 - pointers to members, 456, 459, 509
 - polymorphic classes, 617
 - polymorphic memory resources, 190n3
 - polymorphic types, 616, 1011
 - polymorphism, compile-time, 1046–1050
 - portability with **final** contextual keyword, 1014–1015
 - positive **semidefinite**, 655
 - POSIX epoch, 291
 - postconditions, 807–811
 - potentially evaluated, 615
 - preconditions, 18, 472
 - predicate functions, 86
 - predicate functors, 575
 - predicates, 575
 - preprocessor macros. *See* macros
 - primary declarations, 881
 - primary-class-template declarations, 881
 - private functions, 1038–1041
 - private inheritance, 1029
 - proctor classes, 646
 - proctors, 646, 1139
 - producer-consumer programming pattern, 1000–1005
 - production build, 469
 - programmatically accessible, 1085, 1144
 - protocol hierarchy, performance of, 1020–1023
 - protocols, 440, 540, 1018, 1020
 - proxy iterators, return types of, 1211–1212
 - rvalues**, 513, 716
 - in C++11/14, 720–721
 - evolution of, 807
 - passing to **decltype**, 25
 - publicly accessible, 489
 - pure abstract classes, extracting, 1018
 - pure abstract interfaces, 540, 1020, 1021
 - pure functions, 16
 - pure interfaces, 1020
 - pure virtual functions
 - final** contextual keyword, 1008–1009
 - in protocol hierarchy, 1020

Q

- qualified ids. *See* id-expression
- qualified names, 127, 1060
- qualifiers, 889
- quality of implementation (QoI), 277, 529
- quiet NaN (qNaN), 531

R

- range expressions
 - lifetime of temporary objects, 691–696
 - range-based **for** loops, 680
- range generators, 687–690
- range-based **for** loops, 571–572
 - annoyances, 703–709
 - adapter requirements, 706
 - argument-dependent lookup (ADL), 707–709
 - sentinel iterator types, lack of support, 706–707
 - state of iteration, lack of access, 703–706
 - description of, 679–684
 - further reading for, 709
 - potential pitfalls, 691–703
 - differences in simple and reference-proxy behaviors, 700–703
 - inadvertent element copying, 696–700
 - lifetime of temporary objects, 691–696
 - specification, 680–683
 - traversing arrays and initializer lists, 683–684
 - use cases, 684–691
 - iterating all container elements, 684–685
 - iterating simple values, 690–691
 - range generators, 687–690
 - subranges, 686–687
- Ranges Library, 391–393, 686n4, 687n5
- raw string literals
 - collisions, 109–111
 - description of, 108–111
 - potential pitfalls, 112–114
 - encoding new lines and whitespace, 113–114
 - unexpected indentation, 112–113
 - use cases, 111–112
- raw UDL operators, 841, 845–849, 870
- reachable, 712
- reaching scope, 587–588
- read-copy-update (RCU) synchronization mechanism, 999
- recursion, 604–605, 875
- recursive initialization, 77–78, 163–165
- recursive lambdas, 977–979
- reducing code size, 1101–1111, 1143–1144
- redundant check, 115
- refactoring with curiously recurring template pattern (CRTP), 1042–1044
- reference collapsing, 380–382
- reference related, 726
- reference types
 - alignof** operator, 184
 - deducing, 198
 - `gs1::span`, 17
 - as literal types, 279
 - `memcpy` usage on, 489–493
 - overloading, 727–730
 - union membership and, 1174
- references, **noexcept** and, 1089–1091
- reflection, 520n46
- ref-qualified, 1154
- ref-qualified overloads, 1171–1172
- ref-qualifiers
 - annoyances, 1171–1172
 - description of, 1153–1160
 - forwarding references, 380
 - further reading for, 1173
 - potential pitfalls, 1170–1171
 - syntax and restrictions, 1157–1160
 - use cases, 1160–1170
 - forbidding *lvalue* operations, 1165–1167
 - forbidding *rvalue*-modifying operations, 1163–1165
 - optimizing immutable types and builder classes, 1167–1170
 - returning *rvalue* subobjects, 1160–1163
- register** keyword, 195n1
- regular types, 187n2, 751. *See also* types
- reinterpret_cast** keyword, 506–519
- relaxed restrictions on **constexpr** functions, 959–967. *See also* **constexpr** variables; variadic templates
 - description of, 959–960
 - further reading for, 965
 - optimized C++11 example algorithms, 965–967
 - use cases, 961–964
 - nonrecursive **constexpr** algorithms, 961–962
 - optimized metaprogramming algorithms, 963–964
- release-acquire synchronization paradigm, 998, 1000–1002, 1005
- release-consume synchronization paradigm, 998–999, 1002–1003, 1005
- reopening inline namespaces, 1061–1062
- reordering data members, 178n10
- `reportError` function, 15
- reporting contracts, 1120
- representation, 480, 570
- requires clause in C++20, 486n31
- reserved identifiers, 840
- Resource Acquisition is Initialization (RAII), 388
- resource-owning objects, passing around, 771–775
- return** statements
 - disabling NRVO and implicit move, 244–246
 - moves in, 734–740
 - multiple, 1185–1187

Index

return types

- auto** deduction
 - annoyances, 1201–1203
 - description of, 1182–1194
 - potential pitfalls, 1200
 - use cases, 1194–1200
- constexpr** functions, 277–278
- dependent on parameter type, 126
- of moving iterators, 1211–1212
- of proxy iterators, 1211–1212
- qualified names in, 127
- return by value, 774–775
- trailing return
 - description of, 124–126
 - further reading for, 128
 - inferring type of, 28
 - use cases, 126–128
- return value optimization (RVO), 390
 - requires declared copy/move constructors, 804–805
 - rvalue references, 734
- return values, `[[carries_dependency]]` attribute, 1000
- return-type deduction, delaying, 1199–1200
- reusable lambda expressions, 975
- reuse, lost with **final** contextual keyword, 1023–1026
- right-angle brackets (`>>`)
 - description of, 21
 - further reading for, 24
 - potential pitfalls with, 22–24
 - use cases, 22
- risk-to-reward ratio. *See* safety of adoption
- Rule of Fair Matching, 898–899
- Rule of Greedy Matching, 896–898
- rule of zero, 631, 788
- runtime performance
 - overhead costs of **constexpr** functions, 298–299
 - penalizing to enable compile time, 299–300
- runtime type identification (RTTI), 617
- rvalue references, 1133
 - annoyances, 804–812
 - destructive move, lack of, 811–812
 - evolution of value categories, 807
 - moved-from object requirements overly strict, 807–811
 - RVO and NRVO require declared copy/move constructors, 804–805
 - `std::move` does not move, 805–806
 - visual similarity to forwarding references, 806–807
- decltype** results as, 26
- description of, 710–741
- in expressions, 730–731

- extended value categories in C++11/14, 716–723
- forbidding modifying operations, 1163–1165, 1170–1171
- further reading for, 813
- lvalue references, comparison, 710–711
- modifiable, 820–821
- motivation for, 715–716
- move operations, 714–715
- moves in **return** statements, 734–740
- necessity of, 824
- overload resolution, 713
- overloading on reference types, 727–730
- potential pitfalls, 782–804
 - disabling NRVO, 783–784
 - failure to `std::move` named rvalue references, 784–785
 - implementing move-only types without `std::unique_ptr`, 791–794
 - inconsistent expectations on moved-from objects, 794–803
 - making noncopyable type movable without just cause, 788–791
 - move operations that throw, 787
 - repeatedly calling `std::move` on named rvalue references, 785–786
 - requiring owned resources to be valid, 803–804
 - returning **const rvalues** pessimizes performance, 786–787
 - sink arguments require copying, 782–783
 - some moves equivalent to copies, 788
- range-based **for** loops, 703
 - returning subobjects of, 1160–1163
 - similarity to forwarding references, 397–398
- special member functions, 732–733
- `std::move`, 731–732
- use cases, 741–781
 - identifying value categories, 779–781
 - move operations as optimizations of copying, 741–767
 - move-only types, 768–771
 - passing around resource-owning objects by value, 771–775
 - sink arguments, 775–779
- value category evolution, 813–828
- xvalues*, 712–713

S

- safe features
 - aggregate initialization
 - annoyances, 140–141
 - description of, 138–139
 - potential pitfalls, 140
 - use cases, 139

- attribute support
 - description of, 12–14
 - potential pitfalls with, 18–19
 - use cases, 14–18
- binary literals
 - description of, 142–143
 - further reading for, 146
 - use cases, 144–146
- consecutive right-angle brackets (>>)
 - description of, 21
 - further reading for, 24
 - potential pitfalls with, 22–24
 - use cases, 22
- decltype**
 - description of, 25–26
 - potential pitfalls, 30
 - use cases, 26–30
- defaulted functions
 - annoyances, 42–43
 - description of, 33–36
 - further reading for, 44
 - implicit generation of special member
 - functions, 44–45
 - potential pitfalls, 41–42
 - use cases, 36–41
- definition of, 5
- delegating constructors
 - description of, 46–48
 - potential pitfalls, 50–51
 - use cases, 48–50
- deleted functions
 - annoyances, 58–59
 - description of, 53
 - further reading for, 60
 - use cases, 53–57
- [[*deprecated*]] attribute, 14
 - description of, 147–148
 - potential pitfalls, 150
 - use cases, 148–150
- digit separator (')
 - description of, 152–153
 - further reading for, 154
 - loss of precision in floating-point literals,
 - 154–156
 - use cases, 153
- explicit conversion operators
 - description of, 61–63
 - potential pitfalls, 66–67
 - use cases, 63–65
- local/unnamed types
 - description of, 83–84
 - use cases, 84–87
- long long** integral type
 - description of, 89
 - further reading for, 92
 - potential pitfalls, 91–92
 - use cases, 89–91
- [[*noreturn*]] attribute, 13
 - description of, 95
 - further reading for, 98
 - potential pitfalls, 97–98
 - use cases, 95–97
- nullptr** keyword
 - description of, 99–100
 - further reading for, 103
 - use cases, 100–103
- override** member-function specifier, 5
 - description of, 104–105
 - further reading for, 107
 - potential pitfalls, 106
 - use cases, 105–106
- raw string literals
 - description of, 108–111
 - potential pitfalls, 112–114
 - use cases, 111–112
- static_assert**
 - annoyances, 123
 - description of, 115–118
 - further reading for, 123
 - potential pitfalls, 120–122
 - use cases, 118–119
- thread-safe function-scope static variables
 - annoyances, 80
 - C++03 double-checked-lock pattern, 81–82
 - description of, 68–71
 - further reading for, 81
 - potential pitfalls, 75–80
 - use cases, 71–75
- trailing return, 28
 - description of, 124–126
 - further reading for, 128
 - inferring type of, 28
 - use cases, 126–128
- type/template aliases, creating with **using**
 - declarations, 133–137
- variable templates
 - annoyances, 165
 - description of, 157–160
 - potential pitfalls, 163–165
 - use cases, 160–163
- safe-bool** idiom, 64
- safety of adoption, 2, 4–5. *See also* conditionally
 - safe features; safe features; unsafe features
- salient values, 634, 830–832
- sanitizers, 802
- scalar types, 1207
 - aggregate initialization, 222
 - braced lists and, 254–255

Index

- aggregate initialization (cont.)
 - in C++03, 414
 - copy initialization, 235–236
 - initialization, 217
 - as literal types, 278
 - as standard-layout types, 417
 - as trivial types, 425
- scope
 - duplicate names, loss of access to, 1056–1058
 - function-scope static variables
 - annoyances, 80
 - C++03 double-checked-lock pattern, 81–82
 - description of, 68–71
 - further reading for, 81
 - potential pitfalls, 75–80
 - use cases, 71–75
- scoped allocator model, 328n3
- scoped enumerations, 335–336, 660
- scoped guard, 645–646
- sections, **extern template**, 361
- secure buffers, 460–462
- Secure Hash Algorithms (SHA), 1083–1084
- selective **using** directives for short-named entities, 1074–1076
- semantics, 12, 18, 558–559
- sentinels, 1114
 - iterator types, lack of support, 706–707
 - rvalue references, 743
- sequencing operator, 265. *See also* comma (,) operator
- serial dates, 453
- set associative, 182n11
- SFINAE (substitution failure is not an error), 400, 1089
 - deduced return types and, 1201–1203
 - exception specifications and, 1149–1150
 - perfect forwarding, 397
 - template instantiation and specialization, 1190
- SFINAE evaluation context
 - decltype** with, 28–30
 - expression SFINAE, 29n3
- shadowed, 987
- short-named entities, **using** directives for, 1074–1076
- side effects, 16
- signaling NaN (sNaN), 531
- signals/signaling, 1120, 1213
- signatures, 1052
 - inheriting constructors, 536
 - overloading functions, 1089
 - rvalue references, 729
- signed integer overflow, 90
- simple structs, initialization, 322
- simple type specifiers, 1032
- single-thread-aware objects, avoiding false sharing, 175–176
- single-threaded applications, overhead costs, 80
- sink arguments, 775–779, 782–783
- size constructors, 764
- sizeof**... expressions, 923
- skippable destructors, 464–470
- slicing, 539, 1025
- SmallObjectBuffer, 118n4
- smart pointers, 948–951
- soft UB. *See* library undefined behaviors
- sortRange function, 28–30
- sortRangeImpl function, 28–30
- space. *See* whitespace
- special member functions. *See also* defaulted functions; deleted functions; functions; user-provided special member functions
 - compiler-generated, 621–626
 - constexpr**, 266–268
 - creating high-level value-semantic types (VSTs), 751–762
 - declaring explicitly, 33–34
 - defaulting first declaration of, 34–35
 - exception specifications and, 1086
 - implicit generation of, 44–45
 - initializer lists, 553
 - non-trivial, union type and, 1174–1181
 - restoring suppressed, 36–37
 - rvalue references, 710, 714, 732–733
 - standard-layout types, 421
 - suppressing generation of, 53–55
 - as trivial, 1012
 - user-declared versus user-provided, 413n6
- specialization of variadic class templates, 884–887
- specifiers and arguments. *See also* exception specifications; keywords
 - alignas**
 - description of, 168–172
 - memory allocation, 181–183
 - natural alignment, 179–181
 - potential pitfalls, 176–179
 - use cases, 172–176
 - inline**, 262–265
- square brackets ([[]]), 12
- stable reuse, 1012
- stack frame, 1101
- stack unwinding, 621n4, 1135
- standard conversion, 509
 - enum** class, 334
 - user-defined literals (UDLs), 835
- Standard Library–related restrictions, 1078
- standardized compiler-specific attributes, 13–14

- standard-layout class types, special properties, 420–425
- standard-layout classes, 422
- standard-layout types, 178
 - accessing subobjects via `reinterpret_cast`, 517–519
 - `alignof` operator, 186
 - generalized PODs, 401, 416, 417–420
 - translating C++-only types to C, 452–456
 - vertical encoding for, 448–452
- start function, 14
- stateless lambdas, 605–607
- static assertion declarations, 115
- static data space, 165
- static member variables
 - external definitions, 314–315
 - not defined in own class, 316
- static storage duration, 68, 478
- static variables, function-scope
 - annoyances, 80
 - C++03 double-checked-lock pattern, 81–82
 - description of, 68–71
 - further reading for, 81
 - potential pitfalls, 75–80
 - use cases, 71–75
- static_assert**. *See also* trailing return
 - annoyances, 123
 - description of, 115–118
 - enforcing `noexcept` contract, 639–640
 - evaluation in templates, 116–118
 - further reading for, 123
 - potential pitfalls, 120–122
 - misuse to restrict overload sets, 121–122
 - unintended compilation failures, 120–121
 - syntax and semantics, 115–116
 - use cases, 118–119
 - preventing misuse of class and function templates, 118–119
 - verifying assumptions about target platform, 118
- static-analysis tools, control of external, 17–18
- `std::any`, 187n2
- `std::bit_cast`, 514n41, 516n42
- `std::declval`, 31
- `std::enable_if`, 486n31
- `std::forward`. *See* forwarding references
- `std::forward<T>`, 385, 395
- `std::function`
 - lambda expressions with, 601–603
 - limitations, 994
- `std::index_sequence`, 293
- `std::initializer_list`, 233
 - annoyances, 567–571
 - constructor suppresses implicitly declared default, 568–570
 - homogeneous initializer lists, 567
 - overloaded free function templates, 570–571
 - representation of `const` objects, 570
 - class template usage, 555–558
 - description of, 553–561
 - further reading for, 571
 - inadvertently calling constructors, 242–244
 - overload resolution, 561
 - pointer semantics and temporary lifetimes, 558–559
 - potential pitfalls, 566–567
 - range-based `for` loops, 571–572
 - `std::initializer_list<E>` object initialization, 559
 - traversing with range-based `for` loops, 683–684
 - type deduction, 559–561
 - use cases, 561–566
 - function arguments of same type, 564–565
 - iterating over fixed number of objects, 565–566
 - population of standard containers, 561–562
 - support for braced lists, 562–564
- `std::initializer_list<E>` object initialization, 559
- `std::is_constant_evaluated()`, 297n20
- `std::is_final`, 1014
- `std::is_literal_type`, 283n14
- `std::is_lvalue_reference`, 378
- `std::is_pod`, 438n14
- `std::kill_dependency` function, 999–1000
- `std::list`, move constructors, 1114
- `std::literals`, 1082
- `std::memcpy`, 484–485
 - `const` and reference subobject usage, 489–493
 - ineligible usage, 497–501
- `std::move`, 731–732
 - failure to use with named rvalue references, 784–785
 - lack of movement with, 805–806
 - repeatedly calling on named rvalue references, 785–786
- `std::move_if_noexcept`, 640–644
- `std::pair`, 528–529
- `std::pmr`, 190n3
- `std::pmr::monotonic_resource`, 468n27
- `std::pmr::unsynchronized_pool_resource`, 468n27
- `std::remove_cvref<T>`, 399n6
- `std::set_terminate`, 1104
- `std::string_view`, 874n1

Index

- `std::terminate`, 1104, 1109, 1124–1128
 - `std::thread`, 70
 - `std::tr2::_bases`, 956n27
 - `std::tr2::_direct_bases`, 956n27
 - `std::tuple`, 528–529
 - `std::uint8_t` value, 27
 - `std::uint16_t`, 27, 28n1
 - `std::unique_ptr`, implementing move-only types
 - without, 791–794
 - `std::unique_ptr<T>`, 42n3
 - `std::unordered_map`, 135n1
 - `std::upper_bound`, 294n19
 - `std::variant`, 452n19, 1180n2
 - `std::vector`, 1024–1026
 - appending elements, 634–639
 - change in unspecified behavior, 652–653
 - `std::vector::push_back(T&&)`, 644–647
 - storage class specifiers, 195
 - storing **constexpr** data structures, 311–312
 - Streaming SIMD Extensions (SSE), 173–174
 - strengthening alignment, 168
 - of data members, 170–171
 - of particular objects, 169–170
 - of user-defined types (UDTs), 171
 - strict aliasing, 401
 - string literals, 837, 862–863, 870
 - compile-time traversal, 287–291
 - [[deprecated]] attribute, 148
 - raw
 - description of, 108–111
 - potential pitfalls, 112–114
 - use cases, 111–112
 - static_assert**, 123
 - template instantiations with, 394–395
 - Unicode
 - description of, 129–130
 - potential pitfalls, 130–132
 - use cases, 130
 - strong exception-safety guarantee
 - noexcept** operator, 634–639, 651–652, 658–659, 1097
 - rvalue references, 750, 751, 762, 787
 - strong guarantee, 634, 746
 - strong **typedef** idiom, 73–74
 - strong **typedef** implementation, 541–544
 - strongly typed enumerations. *See* **enum** class
 - Stroustrup, Bjarne, 4
 - undefined behavior, avoiding, 1024n10
 - “unnecessary nannyism,” 1024n8
 - structs, initialization, 322
 - structural base classes, hiding member functions, 56–57
 - structural inheritance, 57, 180, 1025–1026
 - boilerplate code with, avoiding, 540
 - with mix-ins, 545
 - natural alignment, 180
 - structured binding, 201n2, 685n3
 - subobjects
 - initialization, inconsistency in, 326–328
 - of *rvalues*, returning, 1160–1163
 - value categories of, 722n8
 - subranges, 686–687
 - substitution failure is not an error. *See* SFINAE
 - sum type, 1177–1180
 - suppressed
 - constructors
 - by `std::initializer_list`, 568–570
 - special member functions, restoring, 36–37
 - symbol demangler, 361n2
 - symbolic numeric constants, nonintegral, 310–311
 - synchronization paradigms, 998–999
 - syntax of direct initialization, 328–329
 - synthesizing equality with curiously recurring template pattern (CRTP), 1045–1046
- ## T
- T** (forwarding references)
 - annoyances, 397–400
 - metafunction requirements in constraints, 398–400
 - similarity to rvalue references, 397–398
 - auto&&**, 383–384
 - description of, 377–385
 - function template argument type deduction, 379–380
 - further reading for, 400
 - identifying, 382–383
 - not forwarding, 384–385
 - potential pitfalls, 394–397
 - hijacking copy constructor, 395–397
 - `std::forward<T>`, enabling move operations, 395
 - template instantiations with string literals, 394–395
 - reference collapsing, 380–382
 - `std::forward<T>`, 385
 - use cases, 386–393
 - decomposing complex expressions, 391–393
 - emplacement, 390–391
 - forwarding expressions to downstream consumers, 386
 - multiple parameter handling, 386–388
 - perfect forwarding for generic factory functions, 388–389
 - wrapping initialization in generic factory functions, 389–390

- template aliases. *See also* inheriting constructors; trailing return
 - creating with **using** declarations, 133–137
 - description of, 133–134
 - use cases, 134–137
 - binding arguments to template parameters, 135–136
 - simplified **typedef** declarations, 134–135
 - type trait notation, 136–137
- template argument deductions, 212–213, 894–896
- template argument list, 882, 914
- template arguments, 899
- template head, 157
- template instantiation
 - forwarding references, 382
 - with string literals, 394–395
- template instantiation time, 116, 120
- template parameter list, 888, 923–924
- template parameter packs, 437, 879–884, 896–898
- template parameters, 135–136, 896
- template template class parameters, 165
- template template parameter packs, 903–908
- template template parameters, 165, 902
- template-argument expressions, 21–22
- templated call operator. *See* generic lambdas
- templated variable declarations. *See also*
 - constexpr** variables
 - annoyances, 165
 - description of, 157–160
 - potential pitfalls, 163–165
 - use cases, 160–163
 - parameterized constants, 160–161
 - reducing verbosity of type traits, 161–163
- templates
 - constexpr** functions, 276–277
 - evaluation of static assertions in. *See*
 - static_assert**
 - extern**
 - annoyances, 373–375
 - description of, 353–365
 - further reading for, 376
 - potential pitfalls, 371–373
 - use cases, 365–370
 - instantiation and specialization, 1190–1192
 - local/unnamed types as arguments to
 - description of, 83–84
 - use cases, 84–87
 - preventing misuse of, 118–119
 - static assertion evaluation in, 116–118
 - `std::initializer_list` usage, 555–558
 - UDL operator templates, 841, 849–851
 - variable
 - annoyances, 165
 - description of, 157–160
 - potential pitfalls, 163–165
 - use cases, 160–163
 - variadic
 - annoyances, 953–957
 - description of, 873–925
 - further reading for, 958
 - potential pitfalls, 952–953
 - use cases, 925–951
- temporary materialization, 717
- temporary objects, 818–819
 - arrays, 555
 - lifetime extensions, 819–820
 - lifetime in range expressions, 691–696
 - modifiable *rvalues*, 820–821
- temporary rvalue references, 724
- ternary operator, 268, 615, 1186–1187
- test drivers, 114, 866–867
- ***this**, captured by copy, 611–612
- thrashing, 183n14
- thread pool, 989
- thread-safe function-scope static variables
 - annoyances, 80
 - C++03 double-checked-lock pattern, 81–82
 - concurrent initialization, 68–69
 - description of, 68–71
 - destruction, 69
 - further reading for, 81
 - logger example, 69–70
 - multithreaded contexts, 70–71
 - potential pitfalls, 75–80
 - dangerous recursive initialization, 77
 - dependence on order-of-destruction of local objects, 78–80
 - initialization not guaranteed, 75–77
 - recursion subtleties, 77–78
 - use cases, 71–75
- top level **const**, 729
- trailing punctuation in lambda expressions, 613–614
- trailing return. *See also* **decltype**; deduced return type
 - description of, 124–126
 - further reading for, 128
 - inferring type of, 28
 - use cases, 126–128
 - function template whose return type depends on parameter type, 126
 - qualifying names, avoiding redundantly in return types, 127
 - readability of declarations with function pointers, 127–128
- trailing return types, 593–594, 1189

Index

- translation unit (TU)
 - opaque enumerations, 660
 - thread-safe function-scope static variables, 71
- translation-lookaside buffer (TLB), 182n13
- transparently nested namespaces. *See* inline namespaces
- trivial, 38
- trivial classes, 521
- trivial constructors, 273–274, 408n4
- trivial copy constructors, 470, 528n62
- trivial copy operation, 483, 733, 812
- trivial copy-assignment operator, 470
- trivial default constructors, 461
- trivial default initialization, 1087
- trivial destructibility, 468–469
- trivial destructors, 408n4
- trivial move constructors, 484, 528n62
- trivial move operation, 733
- trivial operations, 33
- trivial types
 - in C++17, 425n7
 - fixed-capacity string elements, 476–479
 - future direction of PODs, 438–439
 - generalized PODs, 401, 416–417, 425–429
 - preserving, 39–40
 - requiring, 480–482
 - special member functions and, 1012
 - subcategories, 429–436
 - union membership and, 1174
- triviality, loss of, 329–330
- trivially constructible, 80n7
 - POD types, 431–432
 - secure buffers, 460–462
- trivially copy assignable, 486–487
- trivially copy constructible, 488
- trivially copyable, 39, 41–42
 - C++ Standard not stabilized, 521–527
 - fixed-capacity strings, 470–475
 - ineligible use of `std::memcpy`, 497–501
 - `memcpy` usage on **const** or reference subobjects, 489–493
 - naive copying other than `std::memcpy`, 501–505
 - POD types, 401, 434–436
 - sloppy terminology, 488–489
 - wrong usage of type traits, 482–488
- trivially copyable class, 521
- trivially copyable types, 468
- trivially default constructible, 401, 430–436
- trivially destructible
 - compile-time constructible, literal types, 462–464
 - constexpr** variables, 305
 - POD types, 402, 430–434
 - reducing code size, 1104
 - sloppy terminology, 488–489
- trivially destructible types in C++20, 430n9
- true sharing, 183n15
- tuples, 932–937, 975–976
- type aliases. *See also* inheriting constructors; trailing return
 - befriending as customization point, 1034–1036
 - creating with **using** declarations, 133–137
 - description of, 133–134
 - exception specifications and, 1090, 1147
 - use cases, 134–137
 - binding arguments to template parameters, 135–136
 - simplified **typedef** declarations, 134–135
 - type trait notation, 136–137
- type categories, 837, 843
- type deduction
 - forwarding references, 379–380
 - of `std::initializer_list`, 559–561
- type erasure, 602
- type identifiers as **alignas** specifier argument, 172
- type inference, 193
- type lists, 963
- type parameter packs, 903
- type punning, 401
- type safety, 100–101
- type suffix, 837
- type template parameter packs, 880–884
- type template parameters, 902
- type traits, 436–438
 - in C++17, 651n12
 - notation, 136–137
 - reducing verbosity, 161–163
 - static_assert**, 119
 - `std::is_lvalue_reference`, 378
 - as unreliable, 527–528
 - wrong usage, 482–488
- <type_traits> header, 1014
- type-consistency, explicit expression of, 27–28, 28n1
- typedef**. *See also* aliases
 - capturing results of **decltype** expressions in, 31
 - in <cstdint>, 92
 - strong implementation, 541–544
- typename disambiguator, 382n1
- typename specifiers, 1032
- typenamees
 - explicit, 26–27
 - in **friend** declarations, 1033n1

- types. *See also* POD types; trivial types; type aliases; type safety; type traits; user-defined types (UDTs); value-semantic types (VSTs)
 - as **alignof** argument, 193–194
 - function pointers and, 265–266
 - historical perspective on, 93–94
 - literal, 278–284
 - local/unnamed
 - description of, 83–84
 - use cases, 84–87
 - long long**
 - description of, 89
 - further reading for, 92
 - potential pitfalls, 91–92
 - use cases, 89–91
 - redundant repetition, avoiding, 200–201
 - relative sizes of, 91–92
 - scalar
 - aggregate initialization, 222
 - copy initialization, 235–236
 - initialization, 217
 - trailing return. *See also* **decltype**; deduced return type
 - description of, 124–126
 - further reading for, 128
 - inferring type of, 28
 - use cases, 126–128
 - underlying types (UTs)
 - description of, 829–830
 - further reading for, 834
 - potential pitfalls, 832–833
 - use cases, 830–832
 - union
 - description of, 1174–1177
 - further reading for, 1181
 - potential pitfalls, 1180
 - use cases, 1177–1180
 - variant, 937–948
- U**
- UDL operator templates, 841, 849–851, 870
 - UDL operators, 840–842
 - cooked, 843–845
 - raw, 845–849
 - templates, 849–851
 - UDL suffix, 837
 - UDL type categories, 843
 - UDTs. *See* user-defined types (UDTs)
 - unconditional exception specifications, 1085–1089
 - undefined behavior (UB), 1024n10, 1077, 1104, 1175
 - attributes and, 18–19
 - auto** return-type deduction, 1187
 - constexpr** variable initializers, 306–307
 - contract guarantees, 1115
 - delegating constructors, 50n2
 - diagnosing at compile time, 312–314
 - friend** declarations, 1049
 - generalized PODs, 401
 - long long** integral type, 90
 - [[noreturn]] attribute, 97
 - range-based **for** loops, 692
 - rvalue references, 715
 - thread-safe function-scope static variables, 70
 - uninitialized values, 218
 - union type and, 1180
 - undefined symbol links, 1068n4
 - undefined symbols, 363
 - underlying types (UTs)
 - constexpr** variables, 308–309
 - description of, 829–830
 - enum** class, 337
 - enumerations, 333–334
 - further reading for, 834
 - opaque enumerations, 660
 - potential pitfalls, 832–833
 - Unicode string literals, 131
 - use cases, 830–832
 - underspecifying alignment, 176
 - unevaluated contexts, `std::declval` used in, 31, 1132
 - unevaluated operands, 615
 - Unicode string literals
 - description of, 129–130
 - potential pitfalls, 130–132
 - embedding Unicode graphemes, 130–131
 - library support, lack of, 131
 - UTF-8, problematic treatment of, 131–132
 - use cases, 130
 - unification, 901
 - uniform initialization, 215
 - in factory functions, 239–241
 - in generic code, 238–239
 - member initialization in generic code, 241–242
 - union type
 - description of, 1174–1177
 - discriminated unions, 937–948
 - further reading for, 1181
 - misuse of, 505–506
 - potential pitfalls, 1180
 - use cases, 1177–1180
 - vertical encoding within, 439–448
 - unions
 - default member initializers and, 320–321
 - final** contextual keyword in, 1013
 - unique object address, 418

Index

- unique ownership, 768
- unique-object-address requirement, 418
- unit conversions, 863–865
- universally unique identifier (UUID), 862–863
- unnamed namespaces, 77
- unprocessed string contents, syntax for. *See* raw string literals
- unqualified name lookup, 841
- unreachable rvalue references, 712
- unrecognized attributes, implementation-defined behavior of, 18–19
- unrelated types, 507
- unsafe features
 - auto** return-type deduction
 - annoyances, 1201–1203
 - description of, 1182–1194
 - potential pitfalls, 1200
 - use cases, 1194–1200
 - [[*carries_dependency*]] attribute
 - description of, 998–1000
 - further reading for, 1006
 - potential pitfalls, 1005
 - use cases, 1000–1005
 - decltype(auto)** placeholder
 - annoyances, 1213
 - description of, 1205–1210
 - potential pitfalls, 1212–1213
 - use cases, 1210–1212
- definition of, 6
- final** contextual keyword, 6
 - annoyances, 1028–1030
 - description of, 1007–1014
 - further reading for, 1030
 - potential pitfalls, 1023–1027
 - use cases, 1014–1023
- friend** declarations
 - curiously recurring template pattern (CRTP) use cases, 1042–1054
 - description of, 1031–1033
 - further reading for, 1042
 - potential pitfalls, 1041
 - use cases, 1033–1041
- inline namespaces
 - annoyances, 1079–1082
 - description of, 1055–1062
 - further reading for, 1083
 - potential pitfalls, 1076–1079
 - use cases, 1062–1076
 - versioning case study, 1083–1084
- noexcept** exception specification
 - annoyances, 1143–1150
 - description of, 1085–1094
 - further reading for, 1151–1152
 - potential pitfalls, 1112–1143
 - use cases, 1094–1111
- ref-qualifiers
 - annoyances, 1171–1172
 - description of, 1153–1160
- ref-qualifiers (cont.)
 - further reading for, 1173
 - potential pitfalls, 1170–1171
 - use cases, 1160–1170
- union type
 - description of, 1174–1177
 - further reading for, 1181
 - potential pitfalls, 1180
 - use cases, 1177–1180
- unscoped C++03 enumerations, workarounds for, 332–333
- unsigned long long** type
 - description of, 89
 - further reading for, 92
 - potential pitfalls, 91–92
 - use cases, 89–91
- unsigned ordinary character types, 515
- unspecified rvalue references, 715
- unwinding logic, 1103
- usable literal types, 282–284
- user declared, 413n6, 1105
- user provided
 - defaulted functions, 33–36
 - generalized PODs, 466–472, 477–478
 - replacement for user declared, 413n6
 - rvalue references, 742, 794
- user-declared constructors, 274n7
- user-declared default constructors, 1087
- user-declared special member functions, 1086
- user-defined control constructs, 599–600
- user-defined conversion, 61, 580
- user-defined literals (UDLs), 462
 - annoyances, 869–871
 - confusing raw and string operators, 870
 - floating-point to integer, lack of conversion, 869–870
 - parsing problems, 870–871
 - potential suffix-name collisions, 870
 - UDL operator templates for string literals, lack of, 870
- in C++14 Standard Library, 852–853
- description of, 835–853
- further reading for, 872
- operators, 840–842
 - cooked, 843–845
 - raw, 845–849
 - templates, 849–851
- potential pitfalls, 867–869
 - overuse, 868–869
 - preprocessor surprises, 869
 - unexpected characters yield bad values, 867–868

- restrictions on, 839–840
 - use cases, 853–867
 - test drivers, 866–867
 - unit conversions and dimensional units, 863–865
 - user-defined numeric types, 858–862
 - user-defined types with string representations, 862–863
 - wrappers, 853–857
 - user-defined types (UDTs), 835
 - alignas** specifier, misleading application of, 177–178
 - alignof** operator, 185–186
 - compile-time constructible, literal types, 462
 - creating high-level value-semantic types (VSTs), 751–762
 - default initialization, 322
 - delegating constructors, 46
 - final** contextual keyword in, 1007, 1011–1015
 - friend** declarations
 - curiously recurring template pattern (CRTP) use cases, 1042–1054
 - description of, 1031–1033
 - further reading for, 1042
 - potential pitfalls, 1041
 - use cases, 1033–1041
 - initializer lists, 553
 - as literal types, 280
 - noexcept** operator, 622
 - numeric literals, 858–862
 - strengthening alignment, 168, 171
 - with string representations, 862–863
 - user-provided copy assignment operator, 759
 - user-provided copy constructors, 758–759
 - user-provided default constructors, 80, 217–219, 755, 1087
 - exception specifications and, 1087
 - initialization, 217–218
 - user-provided destructors, 755–756
 - declaration of, 1105
 - exception specifications and, 1088
 - user-provided functions, exception specifications and, 1088
 - user-provided move constructors, 760
 - user-provided move-assignment operator, 760–761
 - user-provided special member functions, 33, 751–753, 1012, 1088
 - defaulting implementation of, 35–36
 - exception specifications and, 1088
 - rvalue references, 753, 755
 - user-provided value constructors, 753–755
 - using** declarations, 535
 - alias creation with, 133–137
 - constexpr** functions, 268
 - with inline namespaces, 1055–1056
 - using** directives, 842
 - constexpr** functions, 268
 - for short-named entities, 1074–1076
 - using-namespace** directives, 1066
 - UTF-8, 129–131, 844
 - UTF-16, 129–131, 844
 - UTF-32, 129–131, 844
- V**
- valid but unspecified, 715, 801
 - value categories, 25, 26, 30, 590, 710, 1145. *See also*
 - lvalue references; prvalues; rvalue references; xvalues
 - auto** return-type deduction, 1184, 1186
 - evolution of, 807, 813–828
 - exact capture with **decltype(auto)**, 1210–1211
 - extended categories in C++11/14, 716–723
 - forwarding references, 377
 - generic lambdas, 972
 - identifying, 779–781
 - lambda-capture expressions, 992
 - prior to C++11, 814–815
 - range-based **for** loops, 680
 - ref-qualifiers, 1153, 1155, 1159–1160, 1172
 - of subobjects, 722n8
 - value constructors, 37, 942
 - user-defined literals (UDLs), 836
 - user-provided, 753–755
 - vertical encoding, 450
 - value initialization, 216–219, 764
 - constexpr** functions, 273–274
 - of constructor arguments, avoiding the most vexing parse, 237–238
 - value initialize, 493
 - value representation, 405, 409, 452, 500, 503, 517n43
 - value semantics, 627, 811
 - value-initialized variables, defining, 236–237
 - values, 51, 741
 - value-semantic classes, 36, 48n1, 187–188, 743
 - value-semantic mechanisms, 663
 - value-semantic types (VSTs), 51, 1034
 - forwarding references, 386
 - generic, creating, 762–767
 - high-level, creating, 751–762
 - lambda-capture expressions, 992
 - low-level, creating, 742–751
 - POD types, 452
 - in-process, 1034
 - rvalue references, 742, 751–752, 761

Index

- variable templates. *See also* **constexpr** variables
 - annoyances, 165
 - description of, 157–160
 - of literal type, 302
 - potential pitfalls, 163–165
 - specialization of, 1078n7
 - use cases, 160–163
 - parameterized constants, 160–161
 - reducing verbosity of type traits, 161–163
- variables. *See also* **auto** variables; **constexpr** variables; function-scope static variables
 - auxiliary, 28
 - in conditional expressions, initialization, 235
 - const**, capturing modifiable copy of, 990–992
 - forwarding into closure, 992–993
 - initialization, 200
 - in lambda expressions, 600–601
 - local, in unevaluated contexts, 610–611
 - mixing captured and noncaptured, 609
 - strengthening alignment, 168
 - value-initialized, defining, 236–237
- variadic alias templates, 887
- variadic class templates, 875, 878–880
 - member functions, 892–894
 - non-type template parameter packs, 901–903
 - specialization of, 884–887
 - type template parameter packs, 880–884
- variadic function templates, 878, 912, 926
 - function parameter packs, 888–892
 - function template argument matching, 900–901
 - generic lambdas, 978
 - lambda expressions, 590
 - template argument deductions, 894–896
- variadic generic lambdas, 973–974
- variadic macros, 249, 781
- variadic member function templates, 892
- variadic member functions, 892–894
- variadic templates. *See also* variadic class templates; variadic function templates
 - annoyances, 953–957
 - expansion is rigid and requires verbose support code, 957
 - limitations on expansion contexts, 954–955
 - linear search, 957
 - parameter packs cannot be used unexpanded, 956
 - unusable functions, 953–954
 - description of, 873–925
 - further reading for, 958
 - pack expansion, 908–911, 925
 - alignas** specifier, 921–922
 - attribute lists, 922
 - base specifier list, 915–917
 - braced initializer lists, 912–914
 - disallowed, 924
 - function call argument list, 912–914
 - function parameter packs, 911–912
 - lambda-capture list, 919–921
 - member initializer list, 917–918
 - sizeof**... expressions, 923
 - template argument list, 914
 - template parameter list, 923–924
 - potential pitfalls, 952–953
 - accidental use of C-style ellipsis, 952
 - compiler limits on number of arguments, 953
 - undiagnosed errors, 952–953
 - Rule of Fair Matching, 898–899
 - Rule of Greedy Matching, 896–898
 - template template parameter packs, 903–908
 - use cases, 925–951
 - advanced traits, 948–951
 - generic variadic functions, 925–926
 - hooking function calls, 930–931
 - object factories, 929–930
 - processing variadic arguments in order, 926–929
 - tuples, 932–937
 - variant types, 937–948
 - variadic alias templates, 887
 - variadic member functions, 892–894
- variant types, 937–948
- vectorization, 1137
- vectorLerp function, 16–17
- vectors of iterators, 26–27. *See also* `std::vector`
- verbosity of type traits, reducing, 161–163
- versioning
 - with inline namespaces, 1083–1084
 - lack of scalability, 1076–1077
- vertical encoding, 439–452
 - for non-trivial types, 448–452
 - within a union, 439–448
 - Xlib library, 445–448
- vertical microcode, 445n17
- virtual base pointers, 409, 416–417, 426
- virtual destructors, 1008
- virtual dispatch, 202, 1015–1017, 1023
- virtual functions, 632–634
- virtual** keyword, with **final** contextual keyword, 1009–1011
- virtual member functions, 1007–1008
 - overriding
 - description of, 104–105

- further reading for, 107
 - potential pitfalls, 106
 - use cases, 105–106
- virtual memory, 181–183
- virtual-function tables (vtables), 441, 617
- vocabulary types, 91, 94, 131
- void** return type
 - deducing, 1184–1185
 - as literal types, 280
- vtable pointers, 409, 416–417, 426, 441–442, 1011

W

- Wall (GCC), 28n1
- weakly typed C++03 enumerators, drawbacks to, 333–335
- well-formed programs, 147n1, 169, 276n8, 355, 371
- Wextra (GCC), 28n1
- whitespace, 22, 113–114
- wide contracts
 - final** contextual keyword, 1021
 - noexcept** operator, 1112–1113
 - rvalue* references, 750
- widgetIterators, 26–27

- Wing, Jeanette, 1030
- witness arguments, 283–284
- witnesses, 284
- working sets, 182–183, 183n14, 628, 1139
- Wpedantic (GCC), 28n1
- wrappers, 853–857
 - for **noexcept** move operations, 1099–1101
 - perfect returning, 1198

X

- Xlib library, 445–448
- xvalues, 712–713, 717
 - in C++11/14, 721–723
 - evolution of, 807, 825–828

Y

- y combinators, 605, 978, 979n4

Z

- zero cost, 1101n11, 1136
- zero initialized, 75, 77n6, 218, 222, 493
- zero-cost exception model, 1134–1136
- zero-overhead exception model, 1101