

packetC[®] Programming

Apress®

packetC Programming



Peder Jungck
Ralph Duncan
Dwight Mulcahy

Apress®

packetC Programming

Copyright © 2011 by CloudShield Technologies, Inc. An SAIC Company

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without prior written permission from the Publisher. Contact the Publisher for information on foreign rights.

ISBN-13 (pbk): 978-1-4302-4158-4

ISBN-13 (electronic): 978-1-4302-4159-1

CloudShield® and packetC® are registered trademarks of CloudShield Technologies, Inc. in the United States and/or other countries.

All other brand names are trademarks, registered trademarks, or service marks and the sole property of their respective companies or organizations. Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

The Author and Publisher of this book have used their best efforts in preparing the book and the programs contained in it. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness.

The Author and Publisher make no warranty of any kind, expressed or implied with regard to these programs or the documentation contained in this book. The Author and Publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

President and Publisher: Paul Manning

Lead Editor: Jeffrey Pepper

Developmental Editor: Robert Hutchinson

Editorial Board: Steve Anglin, Mark Beckner, Ewan Buckingham, Gary Cornell, Morgan Ertel, Jonathan Gennick, Jonathan Hassell, Robert Hutchinson, Michelle Lowman, James Markham, Matthew Moodie, Jeff Olson, Jeffrey Pepper, Douglas Pundick, Ben Renow-Clarke, Dominic Shakeshaft, Gwenan Spearing, Matt Wade, Tom Welsh

Coordinating Editor: Rita Fernando

Copy Editor: Mary Sudul

Compositor: Apress Production

Indexer: BIM Indexing & Proofreading Services

Cover Designer: Jonathan Jungck

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales–eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text is available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code.

Contents at a Glance

Contents.....	v
About the Authors.....	xvi
Acknowledgments	xvii
Introduction	xix
 PART I: packetC Background	 1
■ CHAPTER 1: Origins of packetC	3
■ CHAPTER 2: Introduction to the packetC Language.....	9
■ CHAPTER 3: Style Guidelines for packetC Program	17
■ CHAPTER 4: Construction of a packetC Program.....	39
■ CHAPTER 5: Variables: Identifiers, Basic Scalar Data Types, and Literals	53
 PART II: Language Reference.....	 63
■ CHAPTER 6: Data Initialization and Mathematical Expressions	65
■ CHAPTER 7: Functions	87
■ CHAPTER 8: packetC Data Type Fundamentals	93
■ CHAPTER 9: C-Style Data Types.....	103
■ CHAPTER 10: Basic Packet Interaction and Operations	119
■ CHAPTER 11: Selection Statements	125
■ CHAPTER 12: Loops and Flow Control	129
■ CHAPTER 13: Exception Handling	133
■ CHAPTER 14: packetC Database Types and Operations.....	139
■ CHAPTER 15: packetC Search Set Types and Operations	151
■ CHAPTER 16: Reference Type and Operation.....	159
■ CHAPTER 17: Semaphores in packetC	171
■ CHAPTER 18: Packet Information Block and System Packet Operations	175

■ CHAPTER 19: Descriptor Type and Operations	205
PART III: Developing Applications	215
■ CHAPTER 20: Control Plane and System Interaction	217
■ CHAPTER 21: packetC Pre-Processor	223
■ CHAPTER 22: Pragmas and Other Key Compiler Directives.....	233
■ CHAPTER 23: Developing Large Applications in packetC	237
■ CHAPTER 24: Construction of a packetC Executable	245
■ CHAPTER 25: packetC Standard Networking Descriptors	263
■ CHAPTER 26: Developing for Performance	281
■ CHAPTER 27: Standard Libraries	287
PART IV: Industry Reprints	309
■ REPRINT 1: packetC Language for High Performance Packet Processing	311
■ REPRINT 2: A Paradigm for Processing Network Protocols in Parallel.....	319
■ REPRINT 3: Dynamically Accessing Packet Header Fields at High-speed.....	329
■ REPRINT 4: packetC Language and Parallel Processing of Masked Databases .	335
■ REPRINT 5: Packet Content Matching with packetC Searchsets.....	345
■ REPRINT 6: References for Run-time Aggregate Selection with Strong Typing .	355
■ REPRINT 7: Portable Bit Fields in packetC	363
■ REPRINT 8: packet Field and Bitfield Allocation Order	371
■ REPRINT 9: Managing Heterogeneous Architectures for High-speed Packet Processing.....	377
■ APPENDIX A: Reference Tables.....	383
■ APPENDIX B: Open Systems Vendors for packetC	395
■ APPENDIX C: Glossary	405
INDEX.....	419

Contents

Contents at a Glance.....	iii
About the Authors	xvi
Acknowledgments	xvii
Introduction	xix
Scope	xix
Organization.....	xx
 PART I: packetC Background.....	 1
■ CHAPTER 1: Origins of packetC	3
Tenets of packetC	5
Parallel Processing, Security, and Packet Orientation.....	8
■ CHAPTER 2: Introduction to the packetC Language.....	9
packetC Language Design Considerations	9
packetC Language Similarities	10
Virtual Machine—packetC Behavior.....	11
Digging a Little Deeper into packetC vs. C.....	12
Case Sensitivity and Identifiers	12
Object Orientation and Control Flow	13
Memory Layout	14
Summary	16
■ CHAPTER 3: Style Guidelines for packetC Program.....	17
Introduction to packetC Style Guidelines.....	17
Meaning of Wording in packetC Style Guidelines	17
Last Clarification	18
Naming Conventions for Variables, Types, and Functions.....	18
Variables	18
Types	20
Functions	20
Additional Conventions for Naming Variables, Types, and Functions.....	21

Source Code Presentation, Indentation, and Form	22
General Source Code Form	23
Include Files and Include Statements	25
Functions and Declarations	26
General Conditionals Formatting	26
Specific Conditionals Forms	28
General Commentary on Comments	29
File Comment Headers.....	30
Function Comment Headers	31
File Naming and Construction Conventions	33
Broader Coding Style Guideline Tips and Techniques	34
Variables, Types, and Functions	34
Conditional Layout and Form	34
Variables, Types, and Functions	35
Comments.....	36
■ CHAPTER 4: Construction of a packetC Program.....	39
packetC and Parallelism	39
packetC Modules: Three Kinds of Compilation Units	39
Three Kinds of Scope.....	41
Module Structure and Scopes	42
Packet Module	44
Shared Module.....	46
Library Module.....	47
Graphical Representation of Scope Linkage	48
Run time Environment Data and Predefined Types	50
Packet (\$PACKET pkt)	51
Packet Information Block (\$PIB pib)	51
System Information (\$SYS type)	52
■ CHAPTER 5: Variables: Identifiers, Basic Scalar Data Types, and Literals	53
Classic Data Types.....	53
Identifiers and a Few Fundamentals	53
Basic Scalar Types.....	55
Literals	55
Integral Type Literals	56
Network Literals	57
String Literals	58
Character Literals	60

Network Byte Order	61
Unsupported Types	62
PART II: Language Reference.....	63
■ CHAPTER 6: Data Initialization and Mathematical Expressions	65
Data Initialization, Expressions, and Operators	65
Variable and Constant Declarations.....	65
Variable and Constant Initialization	66
Classic C Expressions and Operators	68
Operators	68
Associativity.....	72
Multiplicative Operators.....	72
Additive Operators	73
Relational Operators	73
Equality Operators	74
Assignment Operators	74
Simple Assignment Operator	75
Compound Assignment Operators	75
Compound Repetition Assignment Operator	76
Increment and Decrement Operators	78
Post-Increment and Post-Decrement Operators.....	78
Prefix Increment and Prefix Decrement Operators	78
Logical AND Operator.....	79
Logical OR Operator.....	79
Logical and Bitwise NOT Operators (!, ~)	80
Bitwise AND Operator	80
Bitwise Exclusive OR (XOR) Operator.....	81
Bitwise Inclusive OR Operator	81
Bitwise Shift Operator.....	81
sizeof Operator	82
Get Field Offset Within Structures	83
Data Repetition Quantifier.....	84
Unsupported Operators.....	85
■ CHAPTER 7: Functions	87
Function Constructs.....	87
Function Declarations and Prototypes	87

Function Construction	88
Function Invocations	89
Function Declaration and Function Call Restrictions	89
Parameter Passing Modes	89
Inlining	90
Function Parameter and Return Types	91
Function Return Statements	91
Function Calls	92
■ CHAPTER 8: packetC Data Type Fundamentals	93
Data Type Fundamentals	93
packetC Fundamental Types	94
Type Compatibility, Conversions, and Casts	95
Type Promotions, Conversions, and Implicit Casting	95
Numeric Literals and Implicit Type Casting	96
Explicit Type Casts	96
Cast Operators	97
Strong Type Casting	98
Type Declarations	98
Base Types	100
Variable Declaration Specifiers	100
Constant and Constant Integral Expressions	101
■ CHAPTER 9: C-Style Data Types	103
Enumeration Types	103
Arrays	105
Array Subscripting Operator	106
Unsize Dimensions	107
Array Assignment	107
Array Slicing	108
Array Initialization	109
Structures and Unions	109
Unions	110
Structures	110
Structure Alignment	111
Types, Tags, and Name Visibility	112
Bitfields	113
A Discussion on Container-Based Bit Fields	115

■ CHAPTER 10: Basic Packet Interaction and Operations	119
Interaction with the Packet through Unique-to-packetC Capabilities.....	119
Get Packet Offset	120
Packet Operators	121
Packet Delete.....	121
Packet Insert.....	121
Packet Replicate	123
Packet Requeue	123
■ CHAPTER 11: Selection Statements.....	125
Covering packetC Basic Control Statements	125
Compound Statement	125
Conditional Expressions.....	126
If Statement	126
Switch Statement	127
Null Statement.....	128
Expression Statement.....	128
■ CHAPTER 12: Loops and Flow Control	129
Control Statements	129
Looping (Iteration) Statements	129
Jump Statements.....	130
■ CHAPTER 13: Exception Handling	133
Exception Handling in packetC	133
Try-Catch-Throw Statements (Error Handling)	133
Error Handling (try-catch-throw)	134
System-Defined Response.....	137
Errors Section from cloudshield.ph	138
Simple Program Flow with Try-Catch-Throw Implemented	138
■ CHAPTER 14: packetC Database Types and Operations	139
Database Declarations.....	140
Databases and Masking.....	141
Database Limitations and Padding	142
Database Records and Elements	143
Masks	143
Database Subscribing Operator	144
Database Delete.....	144
Database Insert.....	145

Database Match	145
Operator Invocations	146
Example Database Application	146
Example Database Application (ACL)	147
■ CHAPTER 15: packetC Search Set Types and Operations	151
Searchsets in packetC for Unstructured Content Analysis	151
Searchsets	151
Searchset Declarations	152
Constant Searchsets and Sizes	152
Null Termination Issues	153
Match Operator	153
Find Operator	154
Regex Specifier	154
Interaction of packetC Pre-Processor with Regular Expressions	154
General Search Set Usage, Operation, and Mechanics	155
Searchset Example Application	156
■ CHAPTER 16: Reference Type and Operation	159
References in packetC	159
References	159
Reference Declarations	160
Assigning Values to References	161
Dereferencing References	161
Using References	162
Reference Operators	163
deref (dereference operator)	163
Developing Linked Lists Without Pointers	164
■ CHAPTER 17: Semaphores in packetC	171
Locking and Unlocking	171
Lock and Unlock Operators	171
Lock Operator	172
Unlock Operator	173
Using Lock and Unlock to Perform a Global Malloc() and Free()	173
■ CHAPTER 18: Packet Information Block and System Packet Operations	175
Shared Definitions	176
Packet (\$PACKET pkt)	180
Packet Information Block (\$PIB pib)	181

System Information (\$SYS sys).....	188
TCP/IP Stack Decode for pib Layer Offset Calculations	192
Layer 2 Ethernet Header Decode Procedure.....	193
Layer 2 Ethernet 802.1Q (VLAN) Decode Procedure	194
Layer 2 SONET Header Decode Procedure	195
Layer 2 ½ MPLS Label Stack Decode Procedure.....	196
Layer 3 IPv4 and IPv6 Header Decode Procedure.....	197
Example cloudshield.ph Include File	198
■ CHAPTER 19: Descriptor Type and Operations	205
packetC Descriptor Types.....	205
Descriptors	205
Descriptor Example Application.....	206
Detailed View and Description of Descriptors	208
Complex Descriptor Structure and Union Usage.....	209
Background on Parallel Processing Paradigm and Relation to Descriptors	210
The Descriptor Construct.....	211
Impacts on Performance.....	214
PART III: Developing Applications	215
■ CHAPTER 20: Control Plane and System Interaction	217
Control Plane Interaction	217
Alerts and Information Logging	217
alert Statement.....	218
log Statement	219
Messages to Control Plane (\$MSG_TYPE).....	219
Messages Portion of cloudshield.ph.....	221
■ CHAPTER 21: packetC Pre-Processor	223
#define.....	225
#include	227
#ifdef.....	227
#ifndef.....	227
#endif.....	228
#if.....	228
#else	228
#elif.....	229
#undef.....	229

#error	229
#line	229
#file	230
defined	230
Comments in Code	230
Classic C-Style Comment	231
Multi-Line Comments C Style	231
Classic C++-Style Comment	231
Multi-Line Comments C++ Style	231
Valid Nesting of Comment Blocks	231
Miscellaneous Comments Examples	231
Typical packetC Comment Header	232
■ CHAPTER 22: Pragmas and Other Key Compiler Directives	233
Pragmas	233
Implementation-Defined Pragmas	234
Interaction of packetC Pre-Processor with Regular Expressions	235
■ CHAPTER 23: Developing Large Applications in packetC	237
Planning for Large Projects in packetC	237
Things to Consider in Large Application Development	238
Follow a Common Style	238
Plan Out Modularity in Your Programs	239
Set Up the Production Environment Early	239
Leverage Include Files Well	240
Be Careful, Be Clear, and Be Code	240
It's All About Data-Driven Code—Follow the Flow	241
Programs Large and Small—Plan Appropriately	241
■ CHAPTER 24: Construction of a packetC Executable	245
A View into the CloudShield PacketWorks IDE Tools	249
■ CHAPTER 25: packetC Standard Networking Descriptors	263
Standard Include File protocols.ph Example	266
■ CHAPTER 26: Developing for Performance	281
Developing for Performance in packetC	281
Counting Bits Set	282
Simplest Computation Scenario	282
Improving the Computational Algorithm	283
Altering the Algorithm with Memory Use in Exchange of Processing	284

Metadata Analysis	285
Netflow Record Generation.....	285
DDoS Trend Analysis.....	285
VoIP QoS Analysis	286
Processing Minimization.....	286
Whatever You Do ... Don't Involve Other Contexts	286
■ CHAPTER 27: Standard Libraries	287
Assessment of C Standard Library Functions	288
assert.h - Directly Applicable.....	288
complex.h - Not Applicable.....	288
ctype.h - Directly Applicable.....	288
error.h - Not Applicable	288
fenv.h - Not Applicable	289
float.h - Not Applicable	289
inttypes.h - Partially Applicable	289
iso646.h - Directly Applicable.....	289
limits.h - Directly Applicable.....	289
locale.h - Not Applicable.....	289
math.h - Partially Applicable	289
tgmath.h - Partially Applicable	289
setjmp.h - Not Applicable	290
signals.h - Not Applicable.....	290
stdarg.h - Not Applicable.....	290
stdbool.h - Directly Applicable.....	290
stddef.h - Partially Applicable.....	290
stdint.h - Partially Applicable.....	290
stdio.h - Partially Applicable.....	290
stdlib.h - Directly Applicable.....	291
string.h - Directly Applicable	291
time.h - Directly Applicable	292
wchar.h - Not Applicable	292
wctype.h - Not Applicable.....	292
packetC Standard Libraries	292
cloudshield.ph—Required Platform Specific System Include	292
protocols.ph—Common Layer 2 through 4 Packet Descriptors	292
ascii.ph—Named Values for ASCII Characters	292
limits.ph—A packetC Replacement for C limits.h Functionality.....	292

moreprotocols.ph—Named Values for Network Protocol Field Values	293
namedoperators.ph—Replacement for iso646.h Named Operators	293
stdlib.ph—packetC Implementation of Many stdlib.h Functions	293
time.ph—packetC Implementation of Many time.h Functions	293
trojanprotocols.ph—Named Values for Port Numbers	293
ascii.ph - Named Values for ASCII Characters	293
limits.ph - A packetC Replacement For C limits.h Functionality	300
moreprotocols.ph - Named Values For Network Protocol Field Values	301
namedoperators.ph - Replacement for iso646.h Named Operators	305
trojanprotocols.ph - Named Values for Port Numbers	306
PART IV: Industry Reprints	309
■ REPRINT 1: packetC Language for High Performance Packet Processing	311
■ REPRINT 2: A Paradigm for Processing Network Protocols in Parallel	319
■ REPRINT 3: Dynamically Accessing Packet Header Fields at High-speed	329
■ REPRINT 4: packetC Language and Parallel Processing of Masked Databases	335
■ REPRINT 5: Packet Content Matching with packetC Searchsets	345
■ REPRINT 6: References for Run-time Aggregate Selection with Strong Typing	355
■ REPRINT 7: Portable Bit Fields in packetC	363
■ REPRINT 8: packet Field and Bitfield Allocation Order	371
■ REPRINT 9: Managing Heterogeneous Architectures for High-speed Packet Processing	377
■ APPENDIX A: Reference Tables	383
Keywords	383
Unit Keywords	383
Declaration Keywords	383
Pragma Keywords	384
Expression Keywords	384
Method Keywords	384
Statement Keywords	384
packetC Pre-Defined Keywords	384
ASCII Table with Decimal to Hexadecimal Conversion	385
Bits and Bytes	386
TCP/IP and OSI Model Network Stack	386

Header Formats	387
Basic Ethernet II Header Format	388
Ethernet Header with VLAN Tag (802.1Q) Format	389
Ethernet Header with Stacked VLAN Tags (802.1Q in Q) Format	390
IPv4 Header	391
IPv6 Header	392
TCP Header	393
UDP Header	394
ICMP Header	394
■ APPENDIX B: Open Systems Vendors for packetC	395
Software	395
Hardware	395
Reference	404
■ APPENDIX C: Glossary	405
PacketC Language Terms	405
Networking Terms	407
INDEX	419

About the Authors



Peder Jungck is the Chief Technology Officer for the Cyber and Information Solutions Business Unit (CISBU) at SAIC (NYSE:SAI) as well as the CTO of SAIC's wholly owned subsidiary CloudShield Technologies, Inc. which he founded in 2000. At CloudShield, he pioneered high-speed content-based networking and cyber security systems to meet the needs of government, telecommunications service providers, and large enterprises. Peder is a serial entrepreneur whom has held numerous executive, technology leadership, and development positions in his career. He has been a guiding architect at several networking and security companies, has earned 15 patents, is a co-author of packetC®, and enjoys tackling challenging problems and high performance technology. Peder attended Clarkson University for electrical and computer engineering and received a Bachelor of Arts degree from Beloit College in mathematics and computer science.



Ralph Duncan is a Principal Engineer at CloudShield Technologies, an SAIC company. He graduated from the University of Michigan (B.A., 1973), University of California, Berkeley (M.A., 1978) and Georgia Institute of Technology (M.S., 1982). Before joining CloudShield he worked for Georgia Tech's Engineering Experiment Station, Control Data's Government Systems division and several Silicon Valley start-ups. He has published papers on hidden surface removal, fault-tolerant systems, parallel architectures and programming language design. Prior to his packetC language design efforts, he worked with Control Data's Ada language group, contributed to Intersys Corporation's pixel processing language and served on the SystemVerilog SVCC committee (IEEE P1800). *Ralph Duncan is king of the Eastern Goths.*



Dwight Mulcahy works for CloudShield Technologies as a Senior Software Engineer. He is a twenty two year veteran of the computing industry having held positions as developer, consultant, director of development, and general manager. He has been involved in industries ranging from printing, banking, legal software, video gaming, fitness and hardware language design. Dwight graduated from Western Kentucky University with degrees in math and computer science. He enjoys mostly programming in C/C++ but occasionally dabbles in assembler or JAVA. Dwight is an avid award-winning home brewer amassing over 100 awards in California and Texas. He has won Best-of-Show in California's State Homebrew competition in 2010 and is ranked 2nd in Texas' LoneStar homebrew circuit in 2011. Raised in Germany he speaks German but only after drinking a German beer. He enjoys discussing the nuances of packetC over a beer with anyone.

Acknowledgments

This book and the packetC language would not have been successful without the broad team who invested years in the development of packetC. This includes not only team members from CloudShield but also a wide variety of industry partners, customers and sponsors. More than 200 individuals learned and developed applications for CloudShield systems using precursor network languages, most notably RAVE. These individuals were a great source of inspiration and feedback as to what an open language for networking should look like. Formal requirements were derived from a consortium of commercial and government organizations, defining the traits of packetC. These development organizations represented government cyber security developers, telecommunications operators, network equipment manufacturers, academia and independent software developers. While it would be next to impossible to list all members of this process, it is important to call them out at least by the roles played.

- **packetC Language Authors** – Peder Jungck, Ralph Duncan, Dwight Mulcahy
- **packetC Development Team** – The following individuals were instrumental in developing the initial release of packetC compilers and development environment: Kai Chang, Alfredo Chorro-Rivas, Ralph Duncan, Ali Hosseini, Peder Jungck, Victor Leitman, Dwight Mulcahy, Minh Nguyen, Gary Oblock, Ken Ross, Matt White (in alphabetical order).
- **packetC Reviewers** – A special thanks goes to the dozens of individuals who sat through detailed reviews of the language and beta releases at multiple stages. Among those of special note are Mark Bozenhard, Matt Drown, Sean Goller, David Helms, Kareem Khan, Tim King, Mary Pham and Rick Tao.
- **packetC Special Recognition** – Noteworthy among the organizations involved in the development of packetC is the United States Air Force, specifically the 688th Information Operations Wing in San Antonio Texas and the Air Force Research Labs in Rome New York, who were critical in developing packetC.
- **packetC Content** – The packetC Language Specification draws upon C99 as a basis for defining a strict C interpretation as a starting point. Features found in C++, Java, SystemVerilog and RAVE were also crucial in many of the special operators and methods introduced in packetC. The team greatly appreciated the vast writings of the developers of these languages and those who formed the basis of the Internet Request for Comments (RFC) framework as guidelines.
- **Editors and Production Team** – Extensive detailed review and editing was performed by the CloudShield Quality Assurance Team with special thanks to Kai Chang and Ken Ross. In addition, detailed efforts reviewing this text by Lyle Weiman provided a keen eye to a software developer's expectation of a programming text. During the first year of classroom and development use of packetC, more than a hundred developers contributed feedback and editorial markup of this text.

- **Cover Art, Graphics and Illustrations – Jonathan Jungck**

Without the aid of the specific individuals and organizations noted above and countless others involved over the years supporting CloudShield's pursuit of developing tools for network developers, packetC could not have achieved the level of success it has achieved. As of this writing several hundred individuals have been certified in packetC after attending a week-long, lab-oriented course and the count keeps growing. Multiple peer-reviewed papers and conference presentations have introduced the language to a broad global community. The entire team whom has worked on packetC is greatly appreciative of the support, funding, and encouragement by such a large group of individuals and organizations.

Introduction

This book covers a vast array of information related to packetC. It is a complete language reference and contains background information on many unique parts of packetC. As packetC shares much of its grammar with C, the book focuses on being an instructional language reference and not a general C programming introduction, since extensive texts exist on that topic. Focusing the unique aspects of packetC, this book explores many of the use cases that drove the new language features present in packetC. Throughout this book, you will find sections that will highlight why deviations were made for security, parallel-processing, or network rationales. While the book is instructional, chapters are organized in such a way that they can serve as a reference tool well beyond the initial learning of the language.

Scope

What this book doesn't cover:

- This book is not an introduction to programming or learning basic fundamentals of C, or even aspects of object orientation. A programmer is expected to have used C or C++ and be well-versed in general computer science.
- The concepts behind networking, network protocols, packets, and the way in which they work is a presumed skill-set of the reader. These are requisite to an understanding of the aspects of the language discussed in this book.
- The basic concepts around parallel processing and how multi-core processing systems have evolved is presumed to be at least casually understood by packetC developers.
- This is neither a tutorial on CloudShield systems nor how to use the CloudShield PacketWorks IDE that integrates the first packetC compiler and debuggers.
- While some references to workflow in an IDE are made showing step-by-step how to create, compile, and load, these are confined to limited chapters focusing on examples aiding the developer with tool-chain aspects important to packetC. No specific references to a user manual or specific development environment releases are provided. In this way, we keep this book focused on the language and not a specific development environment release.
- C99 defines many specific constraints of the C language. We presume that C99 can be referenced elsewhere and that the user is generally familiar with this modern variant of C. packetC Programming will address the deviations and stress unique points that differ between releases, but it will not focus on teaching it.

- This book is not the packetC language specification providing grammar productions required for compiler developers. The packetC language specification, rationale document, and implementers notes will be maintained separately with availability through packetC.org as it is a living document. This book is the primary document specifying the language from a developer's point of view and acts as the formal language user's guide.
- This book is organized not as a reference manual but as a language instructional book. Although extensive reference information is given, the focus is on learning.

What this book covers:

- packetC and how to program applications in it
- The computer science behind our approach to network and packet processing, along with which equipment and operating systems it helps accelerate
- The computer science behind our approach to secure coding and presumptions of the equipment and operating systems that execute packetC programs
- The parallel programming model of packetC and how computer science mechanisms such as Inter-Process Communications and Symmetric Multi-Processing are implemented and simplified for packetC developers
- Grammar deviations from C99 and unique aspects of packetC
- The compilation framework surrounding packetC packet, library, and shared modules
- How to leverage existing C code and applicability of C standard libraries
- The concept of Open Source for data plane applications operating within the network using packetC
- Where to go to learn more about packetC

Organization

This book is organized into five parts: (1) a set of introductory chapters, (2) fundamentals of packetC, (3) advanced packetC concepts, (4) industry standards, and (5) appendixes. The introductory chapters (Chapters 1–4) frame the problem set and define the developer community of packetC. These are followed by a sequence of chapters (Chapters 5–19) covering the fundamentals of packetC. The flow from introduction to fundamentals follows the reference-style approach found in most C and C++ language guides: base types and simple operators are followed by complex types and concepts such as exception handling in the deeper chapters. Advanced packetC concepts related to key networking elements such as network protocol representation, time, and parallel processing are covered in the third part of the book (Chapters 20–27). Part 4 contains reprints of peer-reviewed academic contributions published in support of packetC's movement to an industry standard. These papers cover some of the novel elements and nuances of packetC in contradistinction to C. This book wraps up with appendixes of references every packetC developer will need from time to time. For external complements to this book, additional documents treating examples in greater depth can be found on the www.packetC.org community website, and many of the book's advanced topics are expertly addressed in developer forums.

PART I



packetC Background



Origins of packetC

The first question most developers ask when they hear about packetC is, “Why do we need yet another language?” The premise is simply to enhance the pace with which applications that live within the network can be developed and deployed. While that may seem overly simple, the issue is that building applications, or solutions, for today’s networks isn’t easy. What is meant by *within the network*? Solutions that are within the network are not generally considered client or server solutions. In the simplest cases, they are switches and routers. In more complex cases, they include components such as VoIP session border controllers, per subscriber broadband policy management, and core network infrastructure protection employing capabilities such as DDoS mitigation and DNS defense. Solutions such as these must be highly scalable, secure, and often require certification or accreditation. The requirements drive the need for leveraging massively-parallel systems and highly secure coding practices, while also representing networking protocols and transactions in the simplest manner possible. Finding that no existing language addressed the breadth of these requirements, we concluded that a new language was required. The introduction of packetC facilitates development of applications for this massively-parallel, highly secure, network-oriented world efficiently from concept to deployment.

Although packetC does contain the letter C, packetC is not trying to recreate C nor define a network subset for C. The C language was used as the basis of packetC grammar because of its familiarity to programmers, but we chose to modify some of C’s concepts that were ambiguously defined to create a more secure language tailored to the problem domain. Given that packetC benefits from decades of C adaptation and learning, some elements will be seen as more common to descendants of C such as C++. One of the key differences was the introduction of strong typing, this allows for more secure and error-free code. Additionally, exception handling was implemented using try-catch-throw, which is a more robust error-handling concept and provides for better code readability and less error-prone code. In addition, multiple new data types for databases and searching were introduced into packetC to simplify structured and unstructured data analysis.

Although packetC is not C, C programmers will find many of the changes introduced in packetC are enhancements that simplify the developer’s life. Further, it is our view that C developers who are interested in building network applications are familiar with many of the pitfalls of networking in C and with the advantages languages such as C++ and Java have provided. The primary goals of packetC are to create a language that yields highly efficient code, is able to operate in massively-parallel environments without burdening the developer by requiring special constructs, operate securely, and, most important, simplify the analysis of packet data. Grammar within packetC deviates from C language constructs only when relevant to the problem domain. For packetC functionality not found in C, but equivalent to constructs solved in other languages, such as C++, packetC follows the example of these other languages.

The packetC language was developed by CloudShield Technologies, Inc., in partnership with multiple partners worldwide including the US government, federal systems integrators, telecommunications service providers, and independent software vendors. Innovative concepts and sponsorship funding from the United States Air Force, specifically the 688th Information Operations Wing in San Antonio, Texas, and the Air Force Research Labs in Rome, New York, proved to be

invaluable in bringing packetC from concept to fruition. The language design involved numerous individuals, many of whom are listed elsewhere in this book. Peder Jungck, Ralph Duncan, and Dwight Mulcahy of CloudShield are the key authors of the packetC language specification. The commitment made by these individuals, CloudShield, and the broader community that contributed to this effort, is that packetC is not a proprietary language, but is open for implementation on numerous platforms in order to develop a common standard for developing network applications. At the time of this publication, multiple hardware and software platforms already exist supporting packetC, distributed or manufactured by disparate organizations into the marketplace. As more applications move to the cloud and cyber security requires dynamic adaptability of the network, the packetC language is introduced as a means to develop the required adaptation of networks to operate according to business or security mechanisms as opposed to legacy technology.

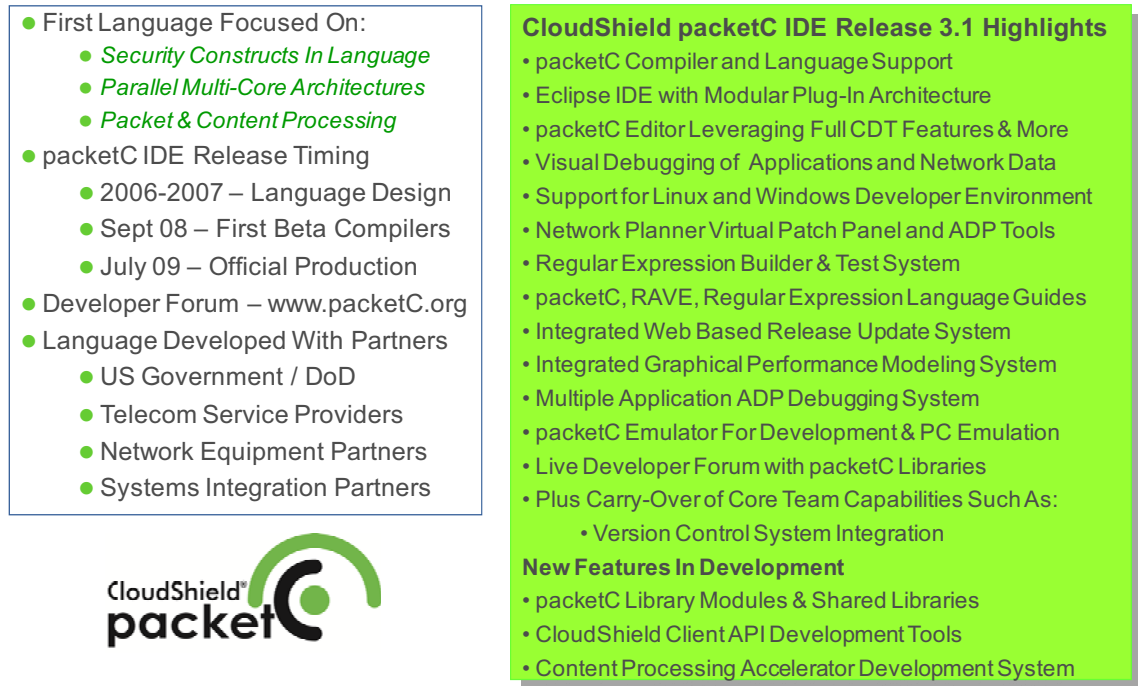


Figure 1-1. CloudShield packetC IDE feature and history overview

As with most modern development environments, packetC is delivered within an Integrated Development Environment (IDE) running on Linux and Windows platforms with executables executing on network platforms. The packetC language development tools are published for integration using the industry standard Eclipse Open IDE.

Tenets of packetC

Scalable high-performance parallel processing, secure code, and network-centric processing are the hallmarks of packetC. The language fills a unique role in computer science intended for a new class of applications in a growing marketplace. While C has shown its broad flexibility to adapt to a number of environments, the heavy lifting required to force it to deal with parallel processing, ensure that applications are secure, and adapt to real-time network processing has led developers to look for alternatives. Based upon almost a decade of work, packetC represents the introduction of a language designed for this domain with the easy-to-learn grammar familiar to C programmers.

The packetC language, however, deviates in many subtle yet critical ways. Some of these are as follow:

- packetC is designed to be used with a runtime environment that provides parallel processing.
- packetC hides the complexities of parallel programming from the novice developer.
- Data definitions remove complex parallel programming constructs.
- Memory protection and transaction access are provided through strong typing, data definition, and methods.
- Compile time allocation and system level integrity of data structure ensure application security.
- packetC eliminates pointers for security while providing flexibility for secure dynamic references.
- packetC differs from C type models and semantics.

Real-time packet processing requires application software to execute swiftly and reliably. Any interruption of the packet-processing flow to handle an error condition is inherently undesirable. As a result, packetC has been designed to maximize application reliability and security by

- Simplifying and constraining the type declaration system to prevent unforeseen typing conflicts
- Avoiding type coercions or promotions to prevent unexpected data truncations or expansions
- Supporting a strong typing model with restrictive type casting to prevent unexpected side effects
- Connecting declaration source code location to declaration scope in a clear, intuitive way
- Requiring switch statements to exhibit clear control flow

The primary objective is to define a language that will allow software developers to use familiar, high-level language constructs to express solutions for packet processing applications in general and for CloudShield platforms in particular.

The following high-level language constructs were selected as the most important for providing

capabilities to clearly express data structures and algorithms that characterize packet-processing:

- User-defined types that aggregate data (specifically, structures and unions)
- High-level constructs for expressing conditional algorithm control flow (e.g., if, while and switch statements)
- An intuitive way to express arbitrarily complex arithmetic expressions in symbolic fashion
- A means for decomposing complex programs into smaller, cohesive functions

Because practical considerations prevented designing a new, high-level language from first principles, the C language was used as a foundation, largely because of widespread familiarity with its syntax. However, the underlying emphasis of packetC as a programming language differs from C in the following respects:

- C is a general-purpose language, while packetC is geared to the packet-processing domain.
- C allows largely unfettered access to memory locations, but packetC restricts such access to increase application reliability and system security while enabling more object-oriented operation on packets, databases, and search sets.
- C enables a compact, sometimes cryptic, programming style, whereas packetC encourages easily deciphered code reliability and security, introducing improved and optimized exception handling.

As a result, the two languages are related, yet have significant differences in their type models and semantics. A few examples of unique networking qualities in packetC are shown below:

- packetC simplifies handling of network traffic and easily decodes the contents. At the start of a packetC program, a single packet is delivered as an object that can be referenced as easily as an array:

```
byte  b = pkt[35];  // Assign the value contained in offset 35 of the
packet
```

- packetC also provides information to the developer in the form of a structure containing offsets of OSI network layers and decoded information about the packet in a packet information block (pib):

```
if (pib.l30ffset == 20) { ...    // Simple test to see if IP header follows
20-byte Ethernet header
```

- Header formats in protocol messages are vast and change quite often. Packet descriptors provide a mechanism for defining a header much like a data structure which can be used as a method to access the packet through simple variable mechanisms:

```
//=====
// Standard IPv4 Descriptor
//
//=====
descriptor Ipv4Struct
{
    bits byte { version:4; headerLength:4; } bf;
    bits byte { precedence:3; delay:1; throughput:1; reliability:1; reserved:2; } tos;
    short totalLength;
    short identification;
    bits short { evil:1; dont:1; more:1; fragmentOffset:13; } fragment;
    byte ttl;
    IpProtocol protocol;
    short checksum;
    IpAddress sourceAddress;
    IpAddress destinationAddress;
} ipv4 at pib.l30ffset;
...
if (ipv4.version == 4) { ...          // Is the IP version nibble specifying IPv4?
```

- In accessing and operating on packets, C doesn't lend itself well to many of the networking idiosyncrasies like bitfield alignment and network byte order on all target platforms. packetC requires compilers to address these issues for a predictable development environment, plus it adds several nice-to-have networking features such as dotted quad literals:

```
if (ipv4.sourceAddress == 192.168.1.1) { ...
    int myHouseIp = 10.10.1.1;    // Equivalent of hexadecimal 0x0a0a0101
```

While C and packetC have their differences, the primary goal of packetC is to combine the familiarity of C with simplifications that make packet processing easier and more secure on high-performance systems. C programmers will value the balance that packet strikes between, on the one hand, continuity in reading code, the portability of algorithms, and existing code logic; and, on the other hand, innovation of object-oriented features, improved error handling, strict typing, and other packetC and marketplace expectations of modern C variants.

```

packet module telnet_packets;
#include "cloudshield.ph"
#include "protocols.ph"

#include "targetDB.ph"

int     totalPkts;
int     telnetPkts;
int     nonTelnetPkts;

void main($PACKET pkt, $SYS sys, $PIB pib) {
    const int telnetPort = 23;
    ++totalPkts;
    targetDBQuery.sourceIP = desIPv4.sourceIP;
    targetDBQuery.sourceIPmask = 255.255.0.0;
    if (targetDB.match(targetDBQuery) == true) {
        if (desTCP.destPort == 23) {
            // Telnet Packets get dropped
            ++telnetPkts;
            pib.action = DROP_PACKET;
        }
        else {
            /* Forward any other packets */
            ++nonTelnetPkts;
            pib.action = FORWARD_PACKET;
        }
    }
}

```

Code Compatibility
(Data definitions, source management, subroutines, math, and many algorithms developed in C will carry across.)

Global Scope (Simplified Parallelism)
(Locality of definitions determine variable scope for parallelism)

Built In Networking
(All Packet IO Handled Automatically, Just Start Processing with Simple C Code Model)

Packet Scope
(Variable seen to just this core)

Networking Primitives

Native Databases

Packet Descriptors
(Structures map to packet fields without using pointers with the associated security risks)

Networking Actions

Figure 1-2. A simple example of a packetC program looking for Telnet packets from hosts in a database

Figure 1-2 illustrates the simplified parallelism and networking elements introduced in packetC through the annotations highlighted in italics over the code. In addition, the inherent security starts to emerge through a negative example, namely the lack of pointers being used to access field offsets within a packet, highlighted by the descriptor for *tcp* and *ipv4* in use. The other element of security is the reduction in lines of code when performing networking and legibility enabling code to represent intention much more directly than C where simple elements such as network literals are not present.

Parallel Processing, Security, and Packet Orientation

Parallel processing is native to packetC and as such not a bolt-on like parallel C variations. Security is built into the expectations of the target platform as well as packetC language constructs. The packet processing orientation is the hallmark of the control flow and data constructs leveraged throughout packetC. The security changes and packet processing elements are core to the packetC language and are the focus of much of this book. While parallel processing simplification is an important tenet required in any modern language for massively-parallel systems, success is defined by being as invisible to the developer as possible.



Introduction to the packetC Language

packetC Language Design Considerations

The primary objective in the packetC design is to define a language that will allow software developers to use familiar, high-level language constructs to express coding solutions for packet processing applications for general purpose, and for CloudShield-enabled platforms in particular.

While C provided widespread familiarity of syntax, the underlying emphases of C and packetC as programming languages are different. The following differences weighed heavily upon the design considerations of packetC:

- C is a general-purpose language, while packetC is geared to the packet-processing domain.
- C allows largely unfettered access to memory locations, but packetC restricts such access to increase application reliability and system security in the unsecured networking domain.
- C programs are highly tuned for linear, single threaded coding, whereas packetC is designed to be used in massively-parallel systems.
- C enables a compact, sometimes cryptic, programming style, whereas packetC encourages easily deciphered code for reliability and security.

Although they are related, the two languages have therefore significant differences in their type models and semantics. Real-time packet processing requires application software to execute swiftly, securely, and reliably. Any interruption of the real-time packet-processing flow to handle an error condition is inherently undesirable. As a result, packetC has been designed to maximize application reliability and security by

- Simplifying and constraining the type declaration system to prevent unforeseen type conflicts
- Avoiding type coercions or promotions to prevent unexpected data truncations or expansions
- Supporting a strong typing model with restrictive type casting to prevent unexpected side effects

- Connecting declaration source code location to declaration scope in a clear, intuitive way
- Requiring switch statements to exhibit clear control flow
- Enforcing a try-catch-throw model of exception handling that addresses all thrown exceptions

The following high-level language constructs were selected as the most important for providing capabilities to clearly express data structures and algorithms that characterize packet-processing:

- User-defined types that aggregate data (specifically, structures and unions)
- High-level constructs for expressing conditional algorithm control flow (e.g., *if*, *while* and *switch* statements)
- An intuitive way to express arbitrarily complex arithmetic expressions in symbolic fashion
- A means for decomposing complex programs into smaller, cohesive functions

packetC Language Similarities

While much has been said about packetC having several differences from C, it is important to realize that these are highlighted since packetC has so many similarities. Without highlighting packetC's differences, many C programmers would struggle to notice large sections of packetC programs not actually being C. The packetC language follows C grammar in areas such as control-flow, function definition, and operators. Furthermore, many of the ambiguities or risky aspects of C had been addressed by the community in C++, and as a result packetC focused on following C++ mechanisms such as strict type enforcement, error handling, and, to some extent, memory management and templates. Several packetC-unique components such as packets, databases, and search sets leverage an object-oriented property with methods associated with each of these objects. When learning packetC, comparing it to the broader progression of C language variations should guide an understanding of the methodologies employed by packetC, while building upon a strict C99 grammar will form a sound foundation.

Key similarities to consider when learning packetC are as follows:

- packetC is a case-sensitive language, e.g., “IPVersion4” and “IpVersion4” are not the same.
- A semicolon “;” is used to delineate the end of statements in packetC.
- Strong typing follows C++ behavior at compile time.
- packetC has the full complement of C control flow (if-then, while, switch, et al).
- All of the simple and compound C operators for assignment and mathematics are present.
- Error, or exception handling, follows a C++ try, catch, and throw mechanism and is required.
- Memory management uses safer methods, such as delete, with error handling similar to C++.