
GemStone®

GemBuilder for C for GemStone/S 64 Bit

Version 3.3

February 2016



INTELLECTUAL PROPERTY OWNERSHIP

This documentation is furnished for informational use only and is subject to change without notice. GemTalk Systems LLC assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.

This documentation, or any part of it, may not be reproduced, displayed, photocopied, transmitted, or otherwise copied in any form or by any means now known or later developed, such as electronic, optical, or mechanical means, without express written authorization from GemTalk Systems.

Warning: This computer program and its documentation are protected by copyright law and international treaties. Any unauthorized copying or distribution of this program, its documentation, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted under the maximum extent possible under the law.

The software installed in accordance with this documentation is copyrighted and licensed by GemTalk Systems under separate license agreement. This software may only be used pursuant to the terms and conditions of such license agreement. Any other use may be a violation of law.

Use, duplication, or disclosure by the Government is subject to restrictions set forth in the Commercial Software - Restricted Rights clause at 52.227-19 of the Federal Acquisitions Regulations (48 CFR 52.227-19) except that the government agency shall not have the right to disclose this software to support service contractors or their subcontractors without the prior written consent of GemTalk Systems.

This software is provided by GemTalk Systems LLC and contributors "as is" and any expressed or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall GemTalk Systems LLC or any contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

COPYRIGHTS

This software product, its documentation, and its user interface © 1986-2016 GemTalk Systems LLC. All rights reserved by GemTalk Systems.

PATENTS

GemStone software is covered by U.S. Patent Number 6,256,637 "Transactional virtual machine architecture", Patent Number 6,360,219 "Object queues with concurrent updating", Patent Number 6,567,905 "Generational garbage collector with persistent object cache", and Patent Number 6,681,226 "Selective pessimistic locking for a concurrently updateable database". GemStone software may also be covered by one or more pending United States patent applications.

TRADEMARKS

GemTalk, **GemStone**, **GemBuilder**, **GemConnect**, and the GemStone and GemTalk logos are trademarks or registered trademarks of GemTalk Systems LLC, or of VMware, Inc., previously of GemStone Systems, Inc., in the United States and other countries.

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Sun, **Sun Microsystems**, and **Solaris** are trademarks or registered trademarks of Oracle and/or its affiliates. **SPARC** is a registered trademark of SPARC International, Inc.

Intel, **Pentium**, and **Itanium** are registered trademarks of Intel Corporation in the United States and other countries.

Microsoft, **MS**, **Windows**, **Windows 7**, **Windows 2008**, and **Windows 8** are registered trademarks of Microsoft Corporation in the United States and other countries.

Linux is a registered trademark of Linus Torvalds and others.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

Ubuntu is a registered trademark of Canonical Ltd., Inc., in the U.S. and other countries.

SUSE is a registered trademark of Novell, Inc. in the United States and other countries.

AIX, **POWER6**, **POWER7**, and **POWER8** are trademarks or registered trademarks of International Business Machines Corporation.

Apple, **Mac**, **Mac OS**, and **Macintosh** are trademarks of Apple Inc., in the United States and other countries.

Other company or product names mentioned herein may be trademarks or registered trademarks of their respective owners. Trademark specifications are subject to change without notice. GemTalk Systems cannot attest to the accuracy of all trademark information. Use of a term in this documentation should not be regarded as affecting the validity of any trademark or service mark.

GemTalk Systems
15220 NW Greenbrier Parkway
Suite 240
Beaverton, OR 97006



Preface

About This Manual

This manual describes GemBuilder for C — a set of C functions that provide a bridge between your application's C code and the application's database controlled by GemStone®. These functions provide your C program with complete access to a GemStone database of objects, and to a virtual machine on which to execute GemStone Smalltalk code.

This manual describes the classic GemBuilder for C interface, which is not thread-safe. A newer, thread-safe interface is provided by `gcits.hf`. This is not described in this manual; see this file, `$GEMSTONE/include/gcits.hf`, for usage details.

Prerequisites

This manual assumes you are familiar with the GemStone Smalltalk programming language, as described in the *Programming Guide for GemStone/S 64 Bit*. In addition, you must know the C programming language, as described in Kernighan and Ritchie's *The C Programming Language* (Prentice Hall, 1978). Finally, you should be familiar with your C compiler, as described in its user documentation.

You should have the GemStone system installed correctly on your host computer, as described in the *GemStone/S 64 Bit Installation Guide* for your platform.

How This Manual is Organized

- ▶ Chapter 1, "Introduction" describes the GemBuilder functions in general, and how they are used in application development with GemStone.
- ▶ Chapter 2, "Building Applications with GemBuilder for C" introduces the two versions of GemBuilder and explains how to build applications that bind to GemBuilder at run time.

- ▶ Chapter 3, "Writing C Functions To Be Called from GemStone" describes how to implement "user action" routines that can be called from GemStone Smalltalk methods.
- ▶ Chapter 4, "Compiling and Linking" describes how to compile and link your C applications and user actions, and how to install them in a GemStone environment prior to execution.
- ▶ Chapter 5, "GemBuilder for C Files and Data Structures" describes GemBuilder include files and the data structures used internally.
- ▶ Chapter 6, "GemBuilder C Functions" provides a detailed description of each GemBuilder function, including syntax, parameters, return value, a general description of what the function does, and including examples of its use.
- ▶ Appendix A, "Reserved OOPs" lists mnemonics for reserved OOPs.
- ▶ Appendix B, "GemStone C Statistics Interface" describes the GemStone C Statistics Interface (GCSI), a library of functions that allow your C application to collect GemStone statistics directly from the shared page cache.

These documents are also available on the Gemtalk customer website, as described below.

Terminology Conventions

The term "GemStone" is used to refer to the server products GemStone/S 64 Bit and GemStone/S, and the GemStone family of products; the GemStone Smalltalk programming language; and may also be used to refer to the company, now GemTalk Systems, previously GemStone Systems, Inc. and a division of VMware, Inc.

Other GemStone Documentation

You may find it useful to look at other GemStone documentation:

- ▶ *Programming Guide for GemStone/S 64 Bit* – a programmer's guide to GemStone Smalltalk, GemStone's object-oriented programming language.
- ▶ *Topaz Programming Environment* – describes Topaz, a scriptable command-line interface to GemStone Smalltalk.
- ▶ *System Administration Guide* – describes maintenance and administration of your GemStone/S system.
- ▶ *GemStone/S 64 Bit Installation Guide* for specific platforms – provides details on the supported platforms for a particular version, and includes the latest compiler version.

A description of the behavior of each GemStone kernel class is available in the class comments in the GemStone Smalltalk repository. Method comments include a description of the behavior of methods.

Technical Support

Support Website

gemtalksystems.com

GemTalk's website provides a variety of resources to help you use GemTalk products:

- ▶ **Documentation** for the current and for previous released versions of all GemTalk products, in PDF form.
- ▶ **Product download** for the current and selected recent versions of GemTalk software.
- ▶ **Bugnotes**, identifying performance issues or error conditions that you may encounter when using a GemTalk product.
- ▶ **TechTips**, providing information and instructions that are not in the documentation.
- ▶ **Compatibility matrices**, listing supported platforms for GemTalk product versions.

This material is updated regularly; we recommend checking this site on a regular basis.

Help Requests

You may need to contact Technical Support directly, if your questions are not answered in the documentation or by other material on the Technical Support site. Technical Support is available to customers with current support contracts.

Requests for technical assistance may be submitted online, by email, or by telephone. We recommend you use telephone contact only for more serious requests that require immediate evaluation, such as a production system down. The support website is the preferred way to contact Technical Support.

Website: **techsupport.gemtalksystems.com**

Email: **techsupport@gemtalksystems.com**

Telephone: **(800) 243-4772 or (503) 766-4702**

When submitting a request, please include the following information:

- ▶ Your name and company name.
- ▶ The versions of GemStone/S 64 Bit and of all related GemTalk products, and of any other related products, such as client Smalltalk products.
- ▶ The operating system and version you are using.
- ▶ A description of the problem or request.
- ▶ Exact error message(s) received, if any, including log files if appropriate.

Technical Support is available from 8am to 5pm Pacific Time, Monday through Friday, excluding GemTalk holidays.

24x7 Emergency Technical Support

GemTalk offers, at an additional charge, 24x7 emergency technical support. This support entitles customers to contact us 24 hours a day, 7 days a week, 365 days a year, for issues impacting a production system. For more details, contact GemTalk Support Renewals.

Training and Consulting

GemTalk Professional Services provide consulting to help you succeed with GemStone products. Training for GemStone/S is available at your location, and training courses are offered periodically at our offices in Beaverton, Oregon. Contact GemTalk Professional Services for more details or to obtain consulting services.



Table of Contents

<i>Chapter 1. Introduction</i>	17
1.1 GemBuilder Application Overview	17
Deciding Where to Do the Work	18
Representing GemStone Objects in C.	18
Smalltalk Access to Objects	19
Calling C Functions from Smalltalk Methods	19
The GemBuilder Functions	20
1.2 Session Control	20
Starting and Stopping GemBuilder.	20
Remote Login Setup	20
Logging In and Out	21
Transaction Management	21
Committing a Transaction	21
Aborting a Transaction.	21
Controlling Transactions Manually.	22
1.3 Representing Objects in C.	22
GemStone-Defined Object Mnemonics.	23
Converting Between Special Objects and C Values.	23
Byte-Swizzling of Binary Floating-Point Values	25
1.4 Manipulating Objects in GemStone	26
Sending Messages to GemStone Objects.	26
Executing Code in GemStone	26
Interrupting GemStone Execution	27
Modification of Classes	28
1.5 Manipulating Objects Through Structural Access	28
Direct Access to Metadata	29
Byte Objects	30
Pointer Objects	30
Nonsequenceable Collections (NSC Objects)	32

1.6 Creating Objects	33
1.7 Fetching and Storing Objects	34
Efficient Fetching and Storing with Object Traversal	34
How Object Traversal Works	34
The Object Traversal Functions	35
Efficient Fetching And Storing with Path Access	36
1.8 Nonblocking Functions.	36
1.9 Operating System Considerations	38
Signal Handling in Your GemBuilder Application	38
Executing Host File Access Methods	39
Writing Portable Code	39
1.10 Error Handling and Recovery	39
Polling for Errors	39
Error Jump Buffers	40
The Call Stack	40
GemStone System Errors	40
1.11 Garbage Collection	40
1.12 Preparing to Execute GemStone Applications	42
GemStone Environment Variables	42

Chapter 2. Building Applications with GemBuilder for C **43**

2.1 RPC and Linked applications	43
RPC for Debugging	43
Linked for Performance	44
Multiple GemStone Sessions	44
The GemBuilder for C Shared Libraries	44
2.2 GemBuilder Applications that load Shared Libraries.	44
Building the Application	44
Searching for the Library	45
How UNIX Matches Search Names with Shared Library Files	45
2.3 Building Statically Linked Applications	45

Chapter 3. Writing C Functions To Be Called from GemStone **47**

3.1 Shared User Action Libraries	47
3.2 How User Actions Work	48
3.3 Developing User Actions.	48
Write the User Action Functions.	49
Create a User Action Library.	49
The gcua.hf Header File	49
The Initialization and Shutdown Functions	50
Compiling and Linking Shared Libraries	51
Using Existing User Actions in a User Action Library.	51

Using Third-party C Code with a User Action Library	51
Loading User Actions	52
Loading User Action Libraries At Run Time	52
Specifying the User Action Library	52
Creating User Actions in Your C Application	53
Verify That Required User Actions Have Been Installed	53
Write the Code That Calls Your User Actions.	53
Remote User Actions	54
Limit on Circular Calls Among User Actions and Smalltalk.	54
Debug the User Action.	54
3.4 Executing User Actions	54
Choosing Between Session and Application User Actions.	54
Running User Actions with Applications	56
With an RPC Application	56
With a Linked Application.	56
Running User Actions with Gems	57
Running User Actions with Applications and Gems	58

Chapter 4. Compiling and Linking **61**

4.1 Development Environment and Standard Libraries	61
4.2 Compiling C Source Code for GemStone	62
The C++ Compiler	62
Listing the Version of Your Compiler	62
Compilation Options.	63
Compilation Command Lines.	63
4.3 Linking C/C++ Object Code with GemStone.	66
Risk of Database Corruption	66
Linker	67
Link Options	67
Command Line Assumptions.	67
Linking User Actions into Shared Libraries	67
Linking Applications That Bind to GemBuilder at Run Time	68

Chapter 5. GemBuilder for C Files and Data Structures **69**

5.1 GemBuilder for C Include Files.	69
5.2 GemBuilder Data Types.	70
The Structure for Representing the Date and Time.	71
The Error Report Structure	72
The Object Information Structure.	72
The Object Report Structure.	74
The Object Report Header Class	75
The User Action Information Structure	78

The Traversal Buffer Type	79
5.3 Structural Access Functions	80
5.4 environmentId	80
5.5 UNIX Signal Handling	80

Chapter 6. GemBuilder C Functions

81

6.1 Function Summary Tables	81
GciAbort	91
GciAddOopToNsc	92
GciAddOopsToNsc	93
GCI_ALIGN	94
GciAll7Bit	95
GciAllocTravBuf	96
GciAlteredObjs	97
GciAppendBytes	99
GciAppendChars	100
GciAppendOops	101
GciBegin	102
GCI_BOOL_TO_OOP	103
GciByteArrayToPointer	104
GciCallInProgress	105
GciCheckAuth	106
GCI_CHR_TO_OOP	107
GciClampedTrav	108
GciClassMethodForClass	110
GciClassNamedSize	112
GciClearStack	113
GciCommit	115
GciCompileMethod	116
GciCompress	118
GciContinue	120
GciContinueWith	121
GciCreateByteObj	122
GciCreateOopObj	124
GciCTimeToDateTime	126
GciDateTimeToCTime	127
GciDbgEstablish	128
GciDbgEstablishToFile	130
GciDbgLogString	131
GciDeclareAction	132
GciDecodeOopArray	133
GciDecSharedCounter	134
GciDirtyExportedObjs	135
GciDirtyObjsInit	137

GciDirtySaveObjs 138
GciDirtyTrackedObjs 140
Gci_doubleToSmallDouble 142
GciEnableFreeOopEncoding 143
GciEnableFullCompression 144
GciEnableSignaledErrors 145
GciEncodeOopArray 146
GciEncrypt 147
GciErr 148
GciExecute 149
GciExecute_ 149
GciExecuteFromContext 151
GciExecuteFromContext_ 151
GciExecuteStr 153
GciExecuteStr_ 153
GciExecuteStrFetchBytes 155
GciExecuteStrFromContext 157
GciExecuteStrFromContext_ 157
GciExecuteStrTrav 159
GciExecuteStrTrav_ 159
GciFetchByte 161
GciFetchBytes_ 162
GciFetchChars_ 164
GciFetchClass 165
GciFetchDateTime 167
GciFetchDynamicIv 168
GciFetchDynamicIvs 169
GciFetchNamedOop 170
GciFetchNamedOops 172
GciFetchNamedSize 174
GciFetchNameOfClass 175
GciFetchNumEncodedOops 176
GciFetchNumSharedCounters 177
GciFetchObjectInfo 178
GciFetchObjImpl 180
GciFetchOop 181
GciFetchOops 183
GciFetchPaths 185
GciFetchSharedCounterValuesNoLock 190
GciFetchSize_ 191
GciFetchUtf8Bytes_ 193
GciFetchVaryingOop 195
GciFetchVaryingOops 197
GciFetchVaryingSize_ 199
GciFindObjRep 200
GciFloatKind 202
GciFltToOop 203

GciGetFreeOop 204
GciGetFreeOops 206
GciGetFreeOopsEncoded 208
GciGetSessionId 209
GciHardBreak 210
GciHiddenSetIncludesOop 211
GCI_I64_IS_SMALL_INT 212
GciI64ToOop 213
GciIncSharedCounter 214
GciInit 215
GciInitAppName 216
GciInitAppName_ 216
GciInstallUserAction 218
GciInstallUserAction_ 218
GciInstMethodForClass 219
GciInUserAction 221
GciIsKindOf 222
GciIsKindOfClass 223
GciIsRemote 224
GciIsSubclassOf 225
GciIsSubclassOfClass 226
GciIvNameToIdx 227
GciLoadUserActionLibrary 228
GciLogin 229
GciLoginEx 231
GciLogout 232
GciLongJump 233
GciMoreTraversal 234
GciNbAbort 236
GciNbBegin 237
GciNbClampedTrav 238
GciNbCommit 239
GciNbContinue 240
GciNbContinueWith 241
GciNbEnd 242
GciNbEnd_ 242
GciNbEndPoll 244
GciNbExecute 245
GciNbExecute_ 245
GciNbExecuteStr 246
GciNbExecuteStr_ 246
GciNbExecuteStrFetchBytes 247
GciNbExecuteStrFromContext 249
GciNbExecuteStrFromContext_ 249
GciNbExecuteStrTrav 251
GciNbExecuteStrTrav_ 251
GciNbMoreTraversal 253

GciNbPerform 254
GciNbPerform_ 254
GciNbPerformNoDebug 256
GciNbPerformNoDebug_ 256
GciNbPerformTrav 258
GciNbPerformTrav_ 258
GciNbStoreTrav 260
GciNbStoreTravDo_ 261
GciNbStoreTravDoTrav_ 262
GciNbStoreTravDoTravRefs_ 263
GciNbTraverseObjs 265
GciNewByteObj 266
GciNewCharObj 267
GciNewDateTime 268
GciNewOop 269
GciNewOops 270
GciNewOopUsingObjRep 272
GciNewString 275
GciNewSymbol 276
GciNewUtf8String 277
GciNscIncludesOop 278
GciObjExists 279
GciObjInCollection 280
GciObjIsCommitted 281
GciObjRepSize_ 282
GciOldOopToNewOop 284
GCI_OOP_IS_BOOL 285
GCI_OOP_IS_SMALL_INT 286
GCI_OOP_IS_SPECIAL 287
GciOopToBool 288
GCI_OOP_TO_BOOL 289
GciOopToChar16 290
GciOopToChar32 291
GciOopToChr 292
GCI_OOP_TO_CHR 293
GciOopToFlt 294
GciOopToI32 296
GciOopToI32_ 296
GciOopToI64 297
GciOopToI64_ 297
GciPerform 298
GciPerform_ 298
GciPerformNoDebug 300
GciPerformNoDebug_ 300
GciPerformSymDbg 302
GciPerformSymDbg_ 302
GciPerformTrav 304
GciPerformTrav_ 304

GciPerformTraverse 306
GciPerformTraverse_ 306
GciPointerToByteArray 308
GciPollForSignal 309
GciPollSocketForRead 310
GciPopErrJump 311
GciProcessDeferredUpdates_ 312
GciProduct 313
GciPushErrJump 314
GciRaiseException 315
GciReadSharedCounter 316
GciReadSharedCounterNoLock 317
GciRealloc 318
GciReleaseAllGlobalOops 319
GciReleaseAllOops 320
GciReleaseAllTrackedOops 321
GciReleaseGlobalOops 322
GciReleaseOops 323
GciReleaseTrackedOops 325
GciRemoveOopFromNsc 326
GciRemoveOopsFromNsc 327
GciReplaceOops 329
GciReplaceVaryingOops 330
GciResolveSymbol 331
GciResolveSymbolObj 332
GciRtlIsLoaded 333
GciRtlLoad 334
GciRtlUnload 335
GciSaveAndTrackObjs 336
GciSaveGlobalObjs 337
GciSaveObjs 338
GciServerIsBigEndian 339
GciSessionIsRemote 340
GciSetCacheName_ 341
GciSetDynLib 342
GciSetErrJump 343
GciSetHaltOnError 345
Gci_SETJMP 346
GciSetNet 347
GciSetNetEx 349
GciSetNetEx_ 349
GciSetSessionId 351
GciSetSharedCounter 353
GciSetTraversalBufSwizzling 354
GciSetVaryingSize 355
GciShutdown 356
GciSoftBreak 357

GciStep 359
GciStep_ 359
GciStoreByte 360
GciStoreBytes 361
GciStoreBytesInstanceOf 363
GciStoreChars 365
GciStoreDynamicIv 366
GciStoreIdxOop 367
GciStoreIdxOops 369
GciStoreNamedOop 371
GciStoreNamedOops 373
GciStoreOop 375
GciStoreOops 377
GciStorePaths 379
GciStoreTrav 384
GciStoreTravDo_ 387
GciStoreTravDoTrav_ 390
GciStoreTravDoTravRefs_ 391
GciStringToInteger 394
GciStrKeyValueDictAt 395
GciStrKeyValueDictAtObj 396
GciStrKeyValueDictAtObjPut 397
GciStrKeyValueDictAtPut 398
GciSwapBytesUint 399
GciSwapBytesUshort 400
GciSymDictAt 401
GciSymDictAtObj 402
GciSymDictAtObjPut 403
GciSymDictAtPut 404
GciTrackedObjsFetchAllDirty 405
GciTrackedObjsInit 407
GciTraverseObjs 408
GciUncompress 412
GciUserActionInit 413
GciUserActionShutdown 414
GciVersion 415

Appendix A. Reserved OOPs **417**

Appendix B. GemStone C Statistics Interface **419**

B.1 Developing a GCSI Application	419
Required Header Files	419
The GCSI Shared Library	419
Compiling and Linking	420

Connecting to the Shared Page Cache	420
The Sample Program	420
B.2 GCSI Data Types	421
The Structure for Representing the GCSI Function Result	421
GcsiAllStatsForMask	423
GcsiAttachSharedCache	424
GcsiAttachSharedCacheForStone	425
GcsiDetachSharedCache	426
GcsiFetchMaxProcessesInCache	427
GcsiInit	428
GcsiStatsForGemSessionId	429
GcsiStatsForGemSessionWithName	430
GcsiStatsForPgsvrSessionId	431
GcsiStatsForProcessId	432
GcsiStatsForShrPcMon	433
GcsiStatsForStone	434
GCSI Errors	435

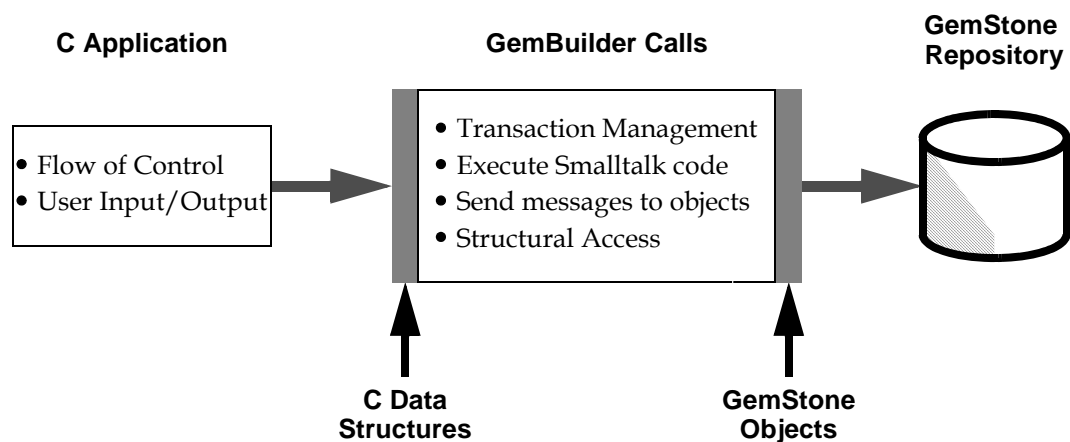
Introduction

GemBuilder for C is a set of C functions that provide your C application with complete access to a GemStone repository and its programming language, GemStone Smalltalk. The GemStone object server contains your schema (class definitions) and objects (instances of those classes), while your C program provides the user interface for your GemStone application. The GemBuilder functions allow your C program to access the GemStone repository either through structural access (the C model) or by sending messages (the Smalltalk model). Both of these approaches are discussed in detail later in this chapter.

1.1 GemBuilder Application Overview

Figure 1.1 illustrates the role of GemBuilder in developing a GemStone application. In effect, developing your GemStone application consists of two separate efforts: creating Smalltalk classes and methods, and writing C code.

Figure 1.1 The Role of GemBuilder in Application Development



We recommend the following steps for developing your hybrid application:

Step 1. Define the application's external interface.

Any GemBuilder application must manage its user interface through custom modules written in C.

Step 2. Decide where to perform the work.

Applications that are a hybrid of C functions and Smalltalk classes pose interesting problems to the designer: Where is the best place to perform the application's work? Is it better to import the representation of an object into your C program and perform the work there, or to send a message which invokes a Smalltalk method? In the next section, we'll examine this question in more detail.

Step 3. Implement and debug the application.

After you've developed a satisfactory design, you can implement and test the C-based functions using familiar techniques and tools (editor, C compiler, link editor, debugger). For information about implementing applications, see Chapter 2, "Building Applications with GemBuilder for C"

Step 4. Compile and link the application.

For instructions about compiling and linking your application, please see Chapter 4, "Compiling and Linking" For full details, see your C compiler user documentation.

Deciding Where to Do the Work

As mentioned above, you will need to decide how much of the application's work to perform in C functions and how much in Smalltalk methods. The following paragraphs discuss both approaches.

Representing GemStone Objects in C

You may choose to implement C functions that access GemStone objects for manipulation in your C program. In such cases, a representation of each object must be imported from GemStone into your C program before the C function is executed. By import, we mean that memory is allocated within your C program to contain the C equivalent of the GemStone Smalltalk object. You could also say that these values are cached in your application; rather than having a reference to the object by identity (OOP), we have the contents of its instance variables. The object in its permanent form still exists in the repository, and the cached values in your application may become obsolete if other sessions commit changes to this object. Exporting is the reverse of importing - you create a GemStone Smalltalk object that holds the equivalent to your C data, or update an existing GemStone Smalltalk object with the C data in your application.

GemBuilder provides functions for importing objects from GemStone to your C program, creating new GemStone objects, directly accessing and modifying the internal contents of objects, and exporting objects to the GemStone repository.

Of course, if you import an object to your C program and modify it, or if you create a new object within your C program, your application must export the new or modified object to GemStone before it can commit the changes to the repository.

Here are some advantages of using GemBuilder structural access functions to modify objects:

- ▶ It may be more efficient to perform a function in C than in Smalltalk.
- ▶ The function may need to be closely linked with I/O functions for the user interface.
- ▶ The function may already exist in a standard library. In this case, the data must be transported from GemStone to that function.

The section “Manipulating Objects Through Structural Access” on page 28 defines exactly how objects are represented in C as address space, and defines the GemBuilder functions for exchanging these structures between GemStone and C.

Smalltalk Access to Objects

In many cases, you will choose to perform your GemStone work directly in Smalltalk. GemBuilder provides C functions for defining and compiling Smalltalk methods for a class, and for sending a message to an object (invoking a Smalltalk method). Here are some advantages of writing a function directly in Smalltalk:

- ▶ The integrity of the data encapsulation provided by the object metaphor is preserved.
- ▶ Functions in Smalltalk are more easily shared among multiple applications.
- ▶ Functions in Smalltalk may be easier to implement. There is no need to worry about moving objects between C and Smalltalk or about space management.
- ▶ The overhead of transporting objects between C and Smalltalk is avoided.
- ▶ Classes or methods may already exist which exhibit behavior similar to the desired behavior. Thus, less effort will be required to implement a new function in Smalltalk.

The section “Manipulating Objects in GemStone” on page 26 defines the GemBuilder functions that allow C applications to send Smalltalk messages to objects and execute Smalltalk code.

Calling C Functions from Smalltalk Methods

Even though you may choose to perform your GemStone work in Smalltalk, you may find that you need to access some functions written in C. GemBuilder allows you to link your user-written C functions to a GemStone session process, and subsequently call those functions from Smalltalk. For example, operations that are computationally intensive or are external to GemStone can be written as C functions and called from within a Smalltalk method (whose high-level structure and control is written in Smalltalk). This is similar to the concept of “user-defined primitives” offered by other object-oriented systems. Here are some advantages of calling C functions from Smalltalk:

- ▶ For computationally intensive portions of a GemStone operation, C functions may execute faster than the same functions written in Smalltalk.
- ▶ Operating system services, or services of other software systems, can be accessed without the overhead of spawning a subprocess. In addition, using C functions to access such services provides greater flexibility for passing arguments and returning results.

Chapter 3, “Writing C Functions To Be Called from GemStone” describes how to implement “user action” routines that can be called from Smalltalk methods, and how to link those routines into a GemBuilder application or a Gem (GemStone session) process.

The GemBuilder Functions

The remainder of this chapter introduces you to many of the GemBuilder C functions.

- ▶ First, we’ll look at functions used in managing GemStone sessions: logging into (and out of) GemStone, switching between multiple sessions, and committing and aborting transactions.
- ▶ Next, we’ll look at functions that allow your C program to manipulate objects by sending Smalltalk messages or executing Smalltalk code fragments.
- ▶ Finally, we’ll examine those functions that perform “structural access” upon the representation of objects within your C program.

1.2 Session Control

All interactions with the GemStone repository monitor occur within the scope of a user’s GemStone session, which may encapsulate one or more individual transactions. GemBuilder provides functions for obtaining and managing GemStone repository sessions, such as logging in and logging out, committing and aborting transactions, and connecting to a different session.

Starting and Stopping GemBuilder

The functions **GciInitAppName** and **GciInit** initialize GemBuilder. When it is used, your application should call **GciInitAppName** before calling **GciInit**. Your C application must not call any other GemBuilder functions until it calls **GciInit**.

The function **GciShutdown** logs out all sessions that are connected to the Gem and deactivates GemBuilder. Your C application should call **GciShutdown** before exiting, in order to guarantee that the process deallocates its resources.

Remote Login Setup

There are two ways to prepare for remote login to a GemStone repository:

1. First, you use a NetLDI that is running in guest mode, attached to the Stone process. Guest mode provides easy access in situations where it is not considered necessary to authenticate users in the network environment before permitting them to log in.
2. Otherwise, you need to set the hostusername and host password for the NetLDI authentication, using the **GciSetNet** function. (see page 347).

You may also wish to set the `GEM_RPCGCI_TIMEOUT` configuration parameter in the GemStone configuration file you use when starting a remote Gem. This parameter sets a timeout limit for the remote Gem; if the Gem remains inactive too long, GemStone logs out the session and terminates the Gem process. See the *System Administration Guide for GemStone/S 64 Bit* for more details.

Logging In and Out

Before your C application can perform any useful repository work, it must create a session with the GemStone system by calling **GciLogin**. That function uses the network parameters initialized by **GciSetNet**.

***GciInit** must be called before the first **GciLogin** in the lifetime of a process.*

If your application calls **GciLogin** again after you are already logged in, GemBuilder will create an additional, independent, GemStone session for you. Multiple sessions can be attached to the same GemStone repository, or they can be attached to different repositories. The maximum number of sessions that may be logged in at one time depends upon your version of GemStone and the terms of your license agreement.

From the point of view of GemBuilder's classic API, only a single session is active at any one time. It is known as the *current session*. Any time you execute code that communicates with the repository, it talks to the current session only. Other sessions are unaffected.

Each session is assigned a number by GemBuilder as it is created. Your application can call **GciGetSessionId** to inquire about the number of the current session, or **GciSetSessionId** to make another session the current one. Your application is responsible for treating each session distinctly.

An application can terminate a session by calling **GciLogout**. After that call returns, the current session no longer exists.

Transaction Management

Committing a Transaction

The GemStone repository proceeds from one stable state to the next by continuously committing transactions. In Smalltalk, the message `System commitTransaction` attempts to commit changes to the repository. Similarly, when your C application calls the function **GciCommit**, GemStone will attempt to commit any changes to objects occurring within the current session.

A session within a transaction views the repository as it existed when the transaction started. By the time you are ready to commit a transaction, other sessions or users may have changed the state of the repository through intervening commit operations. Your application can call **GciAlteredObjs** to determine which objects must be reread from the repository in order to make its view current. Then, to reread those objects, use whatever kind of GemBuilder fetch or traversal functions best suits your needs.

If an attempt to commit fails, your application must call **GciAbort** to discard the transaction. If it does not do so, subsequent calls to **GciCommit** will not succeed.

As mentioned earlier, if your C code has created any new objects or has modified any objects whose representation you have imported, those objects must be exported to the GemStone repository in their new state before the transaction is committed. This ensures that the committed repository properly reflects the intended state.

Aborting a Transaction

By calling **GciAbort**, an application can discard from its current session all the changes to persistent objects that were made since the last successful commit or since the beginning of

the session (whichever is later). This has exactly the same effect as sending the Smalltalk message

```
System abortTransaction.
```

After the application aborts a transaction, it must reread any object whose state has changed.

Controlling Transactions Manually

Under automatic transaction control, a transaction is started when a user logs in to the repository. The transaction then continues until it is either committed or aborted. The call to **GciAbort** or **GciCommit** automatically starts a new transaction when it finishes processing the previous one. Thus, the user is always operating within a transaction.

Automatic transaction control is the default control mode in GemStone. However, there is some overhead associated with transactions that an application can avoid by changing the transaction mode to manual:

```
GciExecuteStr(  
    "System transactionMode: #manualBegin", OOP_NIL);
```

The transaction mode can also be returned to the automatic default:

```
GciExecuteStr(  
    "System transactionMode: #autoBegin", OOP_NIL);
```

In manual mode, the application starts a new transaction manually by calling the **GciBegin** function. The **GciAbort** and **GciCommit** functions complete the current transaction, but do not start a new transaction. Thus, they leave the user session operating outside of a transaction, without its attendant overhead. The session views the repository as it was when the last transaction was completed, or when the mode was last reset, whichever is later.

Since automatic transaction control is the default, a transaction is always started when a user logs in. To operate outside a transaction initially, an application must first set the mode to manual, and then either abort or commit the transaction.

1.3 Representing Objects in C

An important feature of the GemStone data model is its ability to preserve an object's identity distinct from its state. Within GemStone, each object is identified by a unique 64-bit object-oriented pointer, or OOP. Whenever your C program attempts to access or modify the state of a GemStone object, GemStone uses its OOP to identify it. Both the OOP and a representation of the object's state may be imported into an application's C address space.

Within your C program, object identity is represented in variables of type **OopType** (object-oriented pointer). The GemBuilder include file `gci.ht` defines type **OopType**, along with other types used by GemBuilder functions. For more information, see "GciAbort" on page 91.

GemStone-Defined Object Mnemonics

The GemBuilder include file `gcioop.ht` defines C mnemonics for all of the kernel classes in the GemStone repository, as well as the GemStone objects *nil*, *true*, and *false*, and the GemStone error dictionary.

In addition to the predefined objects mentioned above, the GemBuilder include file `gcioop.ht` also defines the C mnemonic `OOP_ILLEGAL`. That mnemonic represents a value that will never be used to represent any object in the repository. You can thus initialize the state of an OOP variable to `OOP_ILLEGAL`, and test later in your program to see if that variable contains valid information.

NOTE

Bear in mind that your C program can only use predefined OOPs, or OOPs that it has received from the GemStone. Your C program cannot create new OOPs directly – it must ask GemStone to create new OOPs for it.

Converting Between Special Objects and C Values

Some Smalltalk classes encode their objects' states directly in their OOPs:

- ▶ SmallInteger objects (for example, the number 5)
- ▶ Character (for example, the letter 'b')
- ▶ Boolean values (true and false)
- ▶ Instances of class UndefinedObject (such as *nil*)

The following GemBuilder functions and macros allow conversion between Character, Integer, or Boolean objects and the equivalent C values:

GCI_BOOL_TO_OOP – (MACRO) Convert a C Boolean value to a GemStone Boolean object.

GciByteArrayToPointer – Given a result from `GciPointerToByteArray`, return a C pointer.

GCI_CHR_TO_OOP – (MACRO) Convert a C character value to a GemStone Character object.

GciI64ToOop – Convert a C 64-bit integer value to a GemStone object.

GciOopToBool – Convert a Boolean object to a C Boolean value.

GCI_OOP_TO_BOOL – (MACRO) Convert a Boolean object to a C Boolean value.

GciOopToChar16 – Convert a Character object to a 16-bit C character value.

GciOopToChr – Convert a Character object to a C character value.

GCI_OOP_TO_CHR – (MACRO) Convert a Character object to a C character value.

GciOopToI32, **GciOopToI32_** – Convert a GemStone object to a C 32-bit integer value.

GciOopToI64, **GciOopToI64_** – Convert a GemStone object to a C 64-bit integer value.

GciPointerToByteArray – Given a C pointer, return a SmallInteger or ByteArray containing the value of the pointer.

GciStringToInteger – Convert a C string to a GemStone SmallInteger or LargeInteger object.

In addition, the following functions allow conversion between Float objects and their equivalent C values. Although a Float's OOP does not encode its state, these functions are listed here for your convenience.

GciFitToOop – Convert a C double value to a SmallDouble or Float object.

GciOopToFlt – Convert a SmallDouble, Float, or SmallFloat object to a C double.

The following macros are for testing OOPs:

GCI_OOP_IS_BOOL – (MACRO) Determine whether or not a GemStone object represents a Boolean value.

GCI_OOP_IS_SMALL_INT – (MACRO) Determine whether or not a GemStone object represents a SmallInteger.

GCI_OOP_IS_SPECIAL – (MACRO) Determine whether or not a GemStone object has a special representation.

The GemBuilder include file `gcioop.ht` uses the C mnemonics `OOP_TRUE`, `OOP_FALSE`, and `OOP_NIL` to represent the GemStone objects *true*, *false*, and *nil*, respectively.

In Example 1.1, assume that you have defined a Smalltalk class called `Address` that represents a mailing address. If the class has five instance variables, the OOPs of one instance of `Address` can be imported into a C array called *address*. Finally, assume that the fifth instance variable represents the zip code of the address.

The fifth element of *address* is the OOP of the SmallInteger object that represents the zip code, not the zip code itself. Example 1.1 imports the value of the zip code object to the C variable *zip*.

*This example and the ones that follow assume that you already have a valid session (obtained from the successful execution of **GciLogin**).*

Example 1.1

```
int64 example1_1(OopType addressId)
{
    // returns the zipcode or -1 if an error occurred,

    enum { addr_num_instVars = 5 };

    OopType instVars[addr_num_instVars];

    int numRet = GciFetchOops(addressId, 1, instVars,
addr_num_instVars);
    if (numRet != (int)addr_num_instVars)
        return -1;

    BoolType conversionError = FALSE;
    int64 zip = GciOopToI64_(instVars[4], &conversionError);
    if (! conversionError)
        return -1;

    // zip now contains an integer that has the same
    // value as the GemStone object represented by address[4]

    return zip;
}
```

Byte-Swizzling of Binary Floating-Point Values

If an application is running on a different machine than its Gem, the byte ordering of binary floating-point values may differ on the two machines. To ensure the correct interpretation of non-special floating-point objects when they are transferred between such machines, the bytes need to be reordered (*swizzled*) to match the machine to which they are transferred.

In GemStone, a binary float is an instance of class Float (eight bytes) or SmallFloat (four bytes), or an instance of SmallDouble (a special object identifier that has no body). Instances of Float and SmallFloat have byte-format bodies whose size is fixed by GemStone and cannot be changed. The programmer must supply all the bytes, or provide a C double, for a binary floating object when creating or storing it.

Most GemBuilder functions provide automatic byte swizzling for instances of Float and SmallFloat. The following GemBuilder functions raise an error if you pass a Float or SmallFloat object to them:

GciAppendBytes – Append bytes to a byte object. (page 99)

GciStoreByte – Store one byte in a byte object. (page 360)

GciStoreBytes – (MACRO) Store multiple bytes in a byte object. (page 361)

GciStoreChars – Store multiple ASCII characters in a byte object. (page 365)

The **GciFetchBytes_** function does not raise an error if you pass an instance of Float or SmallFloat to it, but it also does not provide automatic byte swizzling. It is intended

primarily for use with other kinds of byte objects, such as strings. If you wish to use it with Floats or SmallFloats, you must perform your own byte swizzling as needed.

1.4 Manipulating Objects in GemStone

GemBuilder provides functions that allow C applications to execute Smalltalk code in the repository and to send messages directly to GemStone objects. This section describes these functions in more detail.

Sending Messages to GemStone Objects

GemBuilder provides the function **GciPerform**, which sends a message to a GemStone object. When GemStone receives a message, it invokes and executes the method associated with that message. Thus, the code execution occurs in the repository, not in the application. Example 1.2 illustrates this function.

*This example assumes that you already have a valid session (obtained from the successful execution of **GciLogin**).*

Example 1.2

```
void example_1_2(void)
{
    OopType userGlobals = GciResolveSymbol("UserGlobals", OOP_NIL);
    OopType aKey = GciNewSymbol("myNumber");
    OopType aValue = GciI32ToOop(55);

    OopType argList[2];
    argList[0] = aKey;
    argList[1] = aValue;

    /* Two statements that have the same effect when executed */
    OopType result = GciSendMsg(userGlobals, 4, "at:", aKey, "put:",
aValue);

    result = GciPerform(userGlobals, "at:put:", argList, 2);
}
```

Executing Code in GemStone

Your C application can execute Smalltalk code by calling any of the following GemBuilder functions:

GciExecute – Execute a Smalltalk expression contained in a String or Utf8 object. (page 149)

GciExecuteFromContext – Execute a Smalltalk expression contained in a String or Utf8 object as if it were a message sent to another object. (page 151)

GciExecuteStr – Execute a Smalltalk expression contained in a C string. (page 153)

GciExecuteStrFromContext – Execute a Smalltalk expression contained in a C string as if it were a message sent to an object. (page 157)

The GemBuilder function **GciExecuteStr** allows your application to send a C string containing Smalltalk code to GemStone for compilation and execution. The Smalltalk code may be a message expression, a statement, or a series of statements; in sum, any self-contained unit of code that you could execute within a Topaz **PrintIt** command.

GemStone uses the specified symbol list argument to bind any symbols contained in the Smalltalk source. If the symbol list is `OOP_NIL`, GemStone uses the symbol list associated with the currently logged-in user. Example 1.3 demonstrates the use of this GemBuilder function.

Example 1.3

```
OopType example_1_3(void)
{
    // Pass the String to GemStone for compilation and execution.
    // If it succeeds, return the result of the expression shown
    // otherwise OOP_NIL will be returned.

    OopType objSize = GciExecuteStr(" ^ myObject size ",
        OOP_NIL/*use default symbolList*/ );
    return objSize;
}
```

Your Smalltalk code has the same format as a method, and may include temporaries. In addition, although the circumflex (^) character is used in the above example to return a value after GemStone has executed Smalltalk code (`myObject size`), the circumflex is not required. GemStone returns the result of the last Smalltalk statement executed.

The other functions work similarly, with variations. Before you call **GciExecute** or **GciExecuteFromContext**, you must create or modify a GemStone String object to contain the Smalltalk text to be executed. The **GciExecuteFromContext** and **GciExecuteStrFromContext** functions execute the Smalltalk code within the context (scope) of a specified GemStone object, which implies that the code can access the object's instance variables.

Interrupting GemStone Execution

GemBuilder provides two ways for your application to handle repository interrupts:

- ▶ A *soft break* interrupts the Smalltalk virtual machine only. The only GemBuilder functions that can recognize a soft break are **GciPerform**, **GciContinue**, **GciExecute**, **GciExecuteFromContext**, **GciExecuteStr**, and **GciExecuteStrFromContext**.
- ▶ A *hard break* interrupts the Gem process itself, and is not trappable through Smalltalk exceptions.

Issuing a soft break may be desirable if, for example, your application sends a message to an object (via **GciPerform**), and for some reason the invoked Smalltalk method enters an infinite loop.

In order for GemBuilder functions in your program to recognize interrupts, your program usually needs a signal handler that can call the functions **GciSoftBreak** and **GciHardBreak**. Since GemBuilder generally does not relinquish control to an application until it has finished its processing, soft and hard breaks are then initiated from an interrupt service routine. Alternatively, if you are calling the non-blocking GemBuilder functions, you can service interrupts directly within your event loop, while awaiting the completion of a function.

If GemStone is executing when it receives the break, it replies with an error message. If it is not executing, it ignores the break.

Modification of Classes

Some class definitions are more flexible than others. With respect to modification, classes fall into three categories:

kernel classes

Predefined kernel classes cannot be modified. You can, however, create a subclass of a kernel class and redefine your subclass's behavior.

invariant classes

Once a class has been fully developed, it is normally invariant. Class invariance does not imply that it is impervious to all change. You can add or remove methods, method categories, class variables, or pool variables to any class except a predefined kernel class. You can also create instances of an invariant class.

modifiable classes

You can also create specially modifiable classes, a feature that can be useful (for example) while you are defining schema or implementing the classes. You can modify these classes in the same ways as invariant classes, but you can also add or remove named instance variables. However, you cannot create an instance of a modifiable class. To create an instance, you must first change the class to invariant.

The GemStone Behavior class provides several methods for changing the characteristics of modifiable classes. Use only these predefined methods — *do not use structural access to modify classes*.

1.5 Manipulating Objects Through Structural Access

As mentioned earlier in this chapter, GemBuilder provides a set of C functions that enable you to do the following:

- ▶ Import objects from GemStone to your C program
- ▶ Create new GemStone objects
- ▶ Directly access and modify the internal contents of objects through their C representations
- ▶ Export objects from your C program to the GemStone repository

You may need to use GemBuilder's "structural access" functions for either of two reasons:

- ▶ Speed

Because they call on GemStone's internal object manager without using the Smalltalk virtual machine, the structural access functions provide the most efficient possible access to individual objects.

► Generality

If your C application must handle GemStone objects that it did not create, using the structural access functions may be the only way you can be sure that the components of those objects will be accessible to the application. A user might, for example, define a subclass of Array in which `at:` and `at:put:` were disallowed or given new meanings. In that case, your C application could not rely on the standard GemStone kernel class methods to read and manipulate the contents of such a collection.

Despite their advantages, you should use these structural access functions *only* if you've determined that Smalltalk message-passing won't do the job at hand. GemBuilder's structural access functions violate the principles of abstract data types and encapsulation, and they bypass the consistency checks encoded in the Smalltalk kernel class methods. If your C application unwisely alters the structure of a GemStone object (by, for example, storing bytes directly into a floating-point number), the object will behave badly and your application will break.

For the same reason, do not use structural access to change the characteristics of modifiable classes. Use **GciPerform** to invoke the Smalltalk methods defined under class Behavior for this specific purpose.

For security reasons, the GemStone object AllUsers cannot be modified using structural access. If you attempt to do so, GemStone raises the `RT_ERR_OBJECT_PROTECTED` error.

Direct Access to Metadata

Your C program can use GemBuilder's structural access functions to request certain data about an object:

► Class

Each object is an instance of some class. The class defines the behavior of its instances. To find an object's class, call **GciFetchClass**.

► Format

GemStone represents the state of an object in one of four different implementations (formats): byte, pointer, NSC (non-sequenceable collection), or special. These implementations are described in greater detail in the *GemStone/S 64 Bit Programming Guide*. To find an object's implementation, call **GciFetchObjImpl**.

► Size

The function **GciFetchNamedSize** returns the number of named instance variables in an object, while **GciFetchVaryingSize_** returns the number of unnamed instance variables in an object. **GciFetchSize_** returns the object's complete size (the sum of its named and unnamed variables).

The result of **GciFetchSize_** depends on the object's implementation ("format"). For byte objects (such as instances of String or Float), **GciFetchSize_** returns the number of bytes in the object's representation. For pointer and NSC objects, this function returns the number of OOPs that represent the object. For "special" objects (such as *nil*, or instances of SmallInteger, Character, and Boolean), the size is always 0.

Byte Objects

GemStone byte objects (for example, instances of class String or Symbol) can be manipulated in C as arrays of characters. The following GemBuilder functions enable your C program to store into, or fetch from, GemStone byte objects such as Strings:

GciAppendBytes – Append bytes to a byte object. (page 99)

GciAppendChars – Append a C string to a byte object. (page 100)

GciFetchByte – Fetch one byte from an indexed byte object. (page 161)

GciFetchBytes_ – Fetch multiple bytes from an indexed byte object. (page 162)

GciFetchChars_ – Fetch multiple ASCII characters from an indexed byte object. (page 164)

GciStoreByte – Store one byte in a byte object. (page 360)

GciStoreBytes – (MACRO) Store multiple bytes in a byte object. (page 361)

GciStoreChars – Store multiple ASCII characters in a byte object. (page 365)

Although instances of Float are implemented within GemStone as byte objects, use the functions **GciOopToFlt** and **GciFltToOop** to convert between Float objects and their equivalent C values.

Assume that the C variable *suppId* contains an OOP representing an object of class String. Example 1.4 imports that String into the C variable *suppName*.

Example 1.4

```
void example_1_4(OopType suppId)
{
    char suppName[1025];
    int64 size = GciFetchBytes_(suppId, 1L, (ByteType*)suppName,
                               sizeof(suppName) - 1);
    suppName[size] = '\0';

    // suppName now contains the bytes of the GemStone object
    // referenced
    // by suppId , or the first 1024 bytes whichever is less
}
```

Pointer Objects

In your C program, a GemStone pointer object is represented as an array of OOPs. The order of the OOPs within the GemStone pointer object is preserved in the C array. GemStone represents the following kinds of objects as arrays of OOPs:

Objects with Named Instance Variables

Any object with one or more named instance variables is represented as an array of OOPs. You can determine the positional mapping of instance variables to indexes within the OOP array by calling the GemBuilder function **GciIvNameToIdx**. The following GemBuilder functions allow your C program to store into, or fetch from, GemStone pointer objects with named instance variables:

GciFetchNamedOop – Fetch the OOP of one of an object's named instance variables. (page 170)

GciFetchNamedOops – Fetch the OOPs of one or more of an object's named instance variables. (page 172)

GciStoreNamedOop – Store one OOP into an object's named instance variable. (page 371)

GciStoreNamedOops – Store one or more OOPs into an object's named instance variables. (page 373)

Indexable Objects

Any indexable object not implemented as a byte object is represented as an array of OOPs. The following GemBuilder functions allow your C program to store into, or fetch from, indexable pointer objects:

GciFetchVaryingOop – Fetch the OOP of one unnamed instance variable from an indexable pointer object or NSC. (page 195)

GciFetchVaryingOops – Fetch the OOPs of one or more unnamed instance variables from an indexable pointer object or NSC. (page 197)

GciStoreIdxOop – Store one OOP in an indexable pointer object's unnamed instance variable. (page 367)

GciStoreIdxOops – Store one or more OOPs in an indexable pointer object's unnamed instance variables. (page 369)

In each of the following functions, if the indexable object contains named instance variables, pointers to the named instance variables precede pointers to the indexable instance variables.

GciFetchOop – Fetch the OOP of one instance variable of an object. (page 181)

GciFetchOops – Fetch the OOPs of one or more instance variables of an object. (page 183)

GciStoreOop – Store one OOP into an object's instance variable. (page 375)

GciStoreOops – Store one or more OOPs into an object's instance variables. (page 377)

Assume that the C variable *currSup* contains an OOP representing an object of class Supplier (which defines seven named instance variables). Example 1.5 imports the state of the Supplier object (that is, the OOPs of its component instance variables) into the C variable *instVar*.

Example 1.5

```
void example_1_5(OopType currSup)
{
    enum { num_ivs = 7 };
    OopType instVars[num_ivs];

    int numRet = GciFetchNamedOops(currSup, 1L, instVars, num_ivs);
    if (numRet == 7) {
        // instVars now contains the OOPs of the seven instance
        // variables of the GemStone object referenced by currSup
    } else {
        // error occurred or currSup is not of expected class or size
    }
}
```

Nonsequenceable Collections (NSC Objects)

In addition to byte objects and pointer objects, GemStone exports objects implemented as nonsequenceable collections (NSCs). NSC objects (for example, instances of class IdentityBag and IdentitySet) reference other objects in a manner similar to pointer objects, except that the notion of order is not preserved when objects are added to or removed from the collection.

The following GemBuilder functions allow your C program to store into, or fetch from, GemStone NSC objects:

GciAddOopToNsc – Add an OOP to the unordered variables of a nonsequenceable collection. (page 92)

GciAddOopsToNsc – Add multiple OOPs to the unordered variables of a nonsequenceable collection. (page 93)

GciFetchOop – Fetch the OOP of one instance variable of an object. (page 181)

GciFetchOops – Fetch the OOPs of one or more instance variables of an object. (page 183)

GciRemoveOopFromNsc – Remove an OOP from an NSC. (page 326)

GciRemoveOopsFromNsc – Remove one or more OOPs from an NSC. (page 327)

GciReplaceVaryingOops – Replace all unnamed instance variables in an NSC object. (page 330)

Note that GemStone preserves the position of objects in an NSC only until the NSC is modified, or until the session is terminated (whichever comes first). Although you may use the functions **GciFetchOops** or **GciFetchOop** (defined for pointer objects) to retrieve the OOPs of an NSC's elements, you must use one of the **GciAddOopToNsc** functions to modify the unnamed instance variables of an NSC. (You can use the **GciStoreOop**, **GciStoreOops**, **GciStoreNamedOop**, and **GciStoreNamedOops** functions to modify user-defined named instance variables of an NSC. You cannot, however, use these functions to modify the named instance variables defined in class IdentityBag.)

Assume that the C variable *mySuppSet* contains an OOP representing an object of class SupplierSet (a large set of Supplier objects). Example 1.6 exports the contents of the C variable *newSupp* (a Supplier object) into that SupplierSet.

Example 1.6

```
void example_1_6(OopType mySuppSet, OopType newSupp)
{
    GciAddOopToNsc(mySuppSet, newSupp);

    GciErrSType errInfo;
    if (GciErr(&errInfo)) {
        // an error occurred
    } else {
        // The instance of SupplierSet referenced by mySuppSet now
contains
        // the OOP of the object newSupp.
    }
}
```

1.6 Creating Objects

The following GemBuilder functions allow your C program to create instances of Smalltalk classes:

GciNewOop – Create a new GemStone object. (page 269)

GciNewOops – Create multiple new GemStone objects. (page 270)

GciNewOopUsingObjRep – Create a new GemStone object from an existing object report. (page 272)

Your C application may also create a new object by executing some Smalltalk code that creates new objects as a side-effect.

Once your application has created a new object, it can export the object to the repository by performing the following steps:

Step 1. Modify a previously committed object in the repository so that it references the new object. This may be accomplished with a call to one of the **GciStore...** functions, or by sending a Smalltalk message with the new object as an argument, where the invoked method changes a committed object to reference the new object.

Step 2. Give the new object some meaningful state.

Step 3. Commit a transaction. (As mentioned earlier in this chapter, your C program must first export the object to the GemStone repository before attempting to commit the transaction.)

1.7 Fetching and Storing Objects

Efficient Fetching and Storing with Object Traversal

The functions described in the preceding sections allow your C program to import and export the components of a single GemStone object. When your application needs to obtain information about multiple objects in the repository, it can minimize the number of network calls by using GemBuilder's object traversal functions.

NOTE:

If you are using "linkable" GemBuilder, object traversal will be of little benefit to you. For details, see "RPC and Linked applications" on page 43.

Suppose, for example, that you had created a GemStone Employee class like the one in Example 1.7.

Example 1.7

```
Object subclass: 'Employee'  
  instVarNames: #( 'name' 'empNum' 'jobTitle'  
                  'department' 'address' 'favoriteTune')  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: #()  
  inDictionary: UserGlobals
```

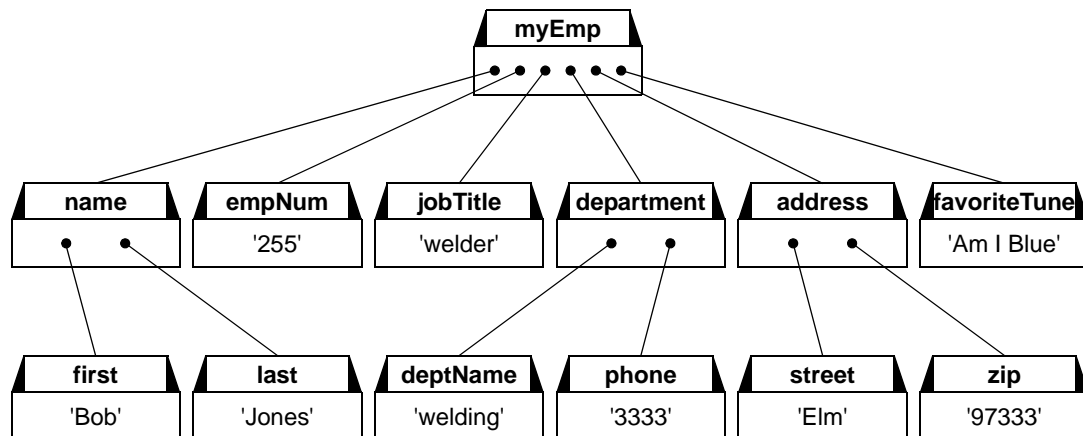
Imagine that you needed to write C code to make a two-column display of job titles and favorite tunes. By using GemBuilder's "object traversal" functions, you can minimize the number of network fetches and avoid running the Smalltalk virtual machine.

How Object Traversal Works

To understand the object traversal mechanism, think of each GemStone pointer object as the root of a tree (for now, ignore the possibility of objects containing themselves). The branches at the first level go to the object's instance variables, which in turn are connected to their own instance variables, and so on.

Figure 1.2 illustrates a piece of the tree formed by an instance of Employee.

Figure 1.2 Object Traversal and Paths



In a single call, GemStone's internal object traversal function walks such a tree post-depth-first to some specified level, building up a "traversal buffer" that is an array of "object reports" describing the classes of the objects encountered and the values of their contents. It then returns that traversal buffer to your application for selective extraction and processing of the contents.

Thus, to make your list of job titles and favorite tunes with the smallest possible amount of network traffic per employee processed, you could ask GemStone to traverse each employee to two levels (the first level is the Employee object itself and the second level is that object's instance variables). You could then pick out the object reports describing *jobTitle* and *favoriteTune*, and extract the values stored by those reports (*welder* and *Am I Blue* respectively).

This approach would minimize network traffic to a single round trip.

One further optimization is possible: instead of fetching each employee and traversing it individually to level two, you could ask GemStone to begin traversal at the *collection* of employees and to descend three levels. That way, you would get information about the whole collection of employees with just a single call over the network.

The Object Traversal Functions

The function `GciTraverseObjs` traverses object trees rooted at a collection of one or more GemStone objects, gathering object reports on the specified objects into a traversal buffer.

- ▶ *Traversal buffers* are instances of the C++ class `GciTravBufType`, which is defined in `$GEMSTONE/include/gcicmn.ht`. (For details about `GciTravBufType`, see "The Traversal Buffer Type" on page 79.)
- ▶ *Object reports* within the traversal buffer are described by the C++ classes `GciObjRepSType` and `GciObjRepHdrSType`, which are defined in `$GEMSTONE/include/gci.ht`. (For details about these classes, see "The Object Report Structure" on page 74.)

Each object report provides information about an object's identity (its OOP), class, size (the number of instance variables, named plus unnamed), object security policy id,

implementation (byte, pointer, NSC, or special), and the values stored in its instance variables.

When the amount of information obtained in a traversal exceeds the amount of available memory, your application can break the traversal into manageable amounts of information by issuing repeated calls to **GciMoreTraversal**. Generally speaking, an application can continue to call **GciMoreTraversal** until it has obtained all requested information.

Your application can call **GciFindObjRep** to scan a traversal buffer for an individual object report. Before it allocates memory for a copy of the object report, your program can call **GciObjRepSize_** to obtain the size of the report.

The function **GciStoreTrav** allows you to store values into any number of existing GemStone objects in a single network round trip. That function takes a traversal buffer of object reports as its argument.

The function **GciStoreTravDo_** is even more parsimonious of network resources. In a single network round trip, you can store values into any number of existing GemStone objects, then execute some code; the function returns a pointer to the resulting object. That function takes a structure as its argument, which defines traversal buffer of object reports and an execution string or message. After the function has completed, the structure also contains information describing the GemStone objects that have changed.

Efficient Fetching And Storing with Path Access

As you've seen, object traversal is a powerful tool for fetching information about multiple objects efficiently. But writing the code for parsing traversal buffers and object reports may not always be simple. And even if you can afford the memory for importing unwanted information, the processing time spent in parsing that information into object reports may be unacceptable.

Consider the Employee object illustrated in the Figure 1.2. If your job were to extract a list of job titles and favorite tunes from a set of such Employees, it would be reasonable to use GemBuilder's object traversal functions (as described above) to get the needed information. The time spent in building up object reports for the unwanted portions would probably be negligible. Suppose, however, that there were an additional 200 instance variables in each Employee. Then the time used in processing wasted object reports would far exceed the time spent in useful work.

Therefore, GemBuilder provides a set of path access functions that can fetch or store multiple objects at selected positions in an object tree with a single call across the network, bringing only the desired information back. The function **GciFetchPaths** lets you fetch selected components from a large set of objects with only a single network round trip. Similarly, your program can call **GciStorePaths** to store new values into disparate locations within a large number of GemStone objects.

1.8 Nonblocking Functions

Under most circumstances, when an application calls a GemBuilder function, the operation that the function specifies is completed before the function returns control to the application. That is, the GemBuilder function *blocks* the application from proceeding until the operation is finished. This effect guarantees a strict sequence of execution.

Nevertheless, in most cases a GemBuilder function calls upon GemStone (that is, the Gem) to perform some work. If the Gem and the application are running in different processes, especially on different machines, blocking implies that only one process can accomplish work at a time. GemBuilder's *nonblocking functions* were designed to take advantage of the opportunity for concurrent execution in separate Gem and application processes.

The results of performing an operation through a blocking function or through its nonblocking twin are always the same. The difference is that the nonblocking function does not wait for the operation to complete before it returns control to the session. Since the results of the operation are probably not ready when a nonblocking function returns, all nonblocking functions but one (**GciNbEnd**) return void.

While a nonblocking operation is in progress an application can do any kind of work that does not require GemBuilder. In fact, it can also call a limited set of GemBuilder functions, listed as follows:

- GciCallInProgress**
- GciErr**
- GciGetSessionId**
- GciHardBreak**
- GciNbEnd**
- GciSetSessionId**
- GciShutdown**
- GciSoftBreak**

If the application first changes sessions, and that session has no nonblocking operation in progress, then the application can call any GemBuilder function, including a nonblocking function. GemBuilder supports one repository request at a time, per session. However, nonblocking functions do not implement threads, meaning that you cannot have multiple concurrent repository requests in progress within a single session. If an application calls any GemBuilder function besides those listed here while a nonblocking operation is in progress in the current session, the error `GCI_ERR_OP_IN_PROGRESS` is generated.

Once a nonblocking operation is in progress, an application must call **GciNbEnd** at least once to determine the operation's status. Repeated calls are made if necessary, until the operation is complete. When it is complete, **GciNbEnd** hands the application a pointer to the result of the operation, the same value that the corresponding blocking call would have returned directly.

Nonblocking functions are not truly nonblocking if they are called from a linkable GemBuilder session, because the Gem and GemBuilder are part of the same process. However, those functions can still be used in linkable sessions. If they are, **GciNbEnd** must still be called at least once per nonblocking call, and it always indicates that the operation is complete.

All error-handling features are supported while nonblocking functions are used. Errors may be signalled either when the nonblocking function is called or later when **GciNbEnd** is called.

1.9 Operating System Considerations

Like your C application, GemBuilder for C is, in itself, a body of C code. Some aspects of the interface must interact with the surrounding operating system. The purpose of this section is to point out a few places where you must code with caution in order to avoid conflicts.

Signal Handling in Your GemBuilder Application

Under UNIX, it is important that signals be enabled when your code calls GemBuilder functions. Disabling signals has the effect of disabling much of the error handling within GemBuilder. Because signal handlers can execute at arbitrary points during execution of your application, your signal handling code should not call any GemBuilder functions other than **GciSoftBreak**, **GciHardBreak**, or **GciCallInProgress**.

GciInit always installs a signal handler for SIGIO. This handler chains to any previous handler.

In the linkable (GciLnk) configuration, **GciInit** also does the following:

- ▶ Installs handlers to service and ignore these signals (if no previous handler is found): SIGPIPE, SIGHUP, SIGDANGER
- ▶ Installs handlers to treat the following signals as fatal errors if they are defined by the operating system: SIGTERM, SIGXCPU, SIGABRT, SIGXFSZ, SIGXCPU, SIGEMT, SIGLOST
- ▶ Installs a handler for SIGUSR1. If you have a valid linkable session, SIGUSR1 will cause the Smalltalk interpreter to print the current Smalltalk stack to stdout or to the Topaz output file. This handler chains to any previous handler.
- ▶ Installs a handler for SIGUSR2, which is used internally by a Gemstone session. This handler chains to any previous handler.
- ▶ Installs a handler to gracefully handle SIGCHLD if no previous handler is found.
- ▶ Installs a handler to treat SIGFPE as a fatal error if no previous handler is found.
- ▶ Installs handlers for SIGILL and SIGBUS. If the program counter is found to be in libgci*.so code, or if no previous handler is available to chain to, these are fatal errors.
- ▶ Installs a handler for SIGSEGV. A Smalltalk stack overflow produces a SIGSEGV, which is translated to a Smalltalk stack overflow error. If the program counter is found to be in libgci*.so code, or if no previous handler is available to chain to, SEGV is a fatal error.

If your application installs a handler for SIGIO after calling **GciInit**, your handler must chain to the previously existing handler.

If your application uses Linkable GCI and installs any signal handlers after calling **GciInit**, you must chain to the previously existing handlers. If you install handlers for SIGSEGV, SIGILL or SIGBUS, your handler must determine if the program counter at the point of the signal is in your own C or C++ code and if not, must chain to the previously existing handler. You must only treat these signals as fatal if the program counter is in your own code.

If you are linking with other shared libraries, it is recommended that **GciInit** be called after all other libraries are loaded.

Executing Host File Access Methods

If you use **GciPerform** or any of the **GciExecute...** functions to execute a Smalltalk host file access method (as listed below), and you do not supply a full file pathname as part of the method argument, the default directory for the Smalltalk method depends on the version of GemBuilder that you are running. With **GciLnk**, the default directory is the directory in which the Gem (GemStone session) process was started. With **GciRpc**, the default directory is the home directory of the host user account, or the `#dir` specification of the network resource string. The Smalltalk methods that are affected include `System class>>performOnServer:` and the file accessing methods implemented in `GsFile`. See the file I/O information in the *GemStone/S 64 Bit Programming Guide*.

Writing Portable Code

If you want to produce code that can run in both 32-bit and 64-bit environments, observe the following guidelines:

- ▶ Don't hard-code size computations. Instead, use `sizeof` operations, so that if some structure changes, your code will still return the correct values.
- ▶ If you are using `printf` strings to print 64-bit integers, you might find it convenient to use the `FMT_*` macros in `$GEMSTONE/include/gcicmn.ht`. Those macros help you to compose a format string for a `printf` that will be portable. In particular, use of the `FMT_*` macros make the printing of 64-bit integers portable between Windows and UNIX.
- ▶ To avoid discrepancies between 32-bit and 64-bit environments, avoid the use of `long` or `unsigned long` in your code. Instead, you can use the type `intptr_t`, which makes the variable the same size as a pointer, regardless whether your application is running in 32-bit or 64-bit. Alternatively, you can use `int64` or `int` to fix the size of the variable explicitly.

1.10 Error Handling and Recovery

Your C program is responsible for processing any errors generated by GemBuilder function calls.

The GemBuilder include file `gciterr.ht` documents and defines mnemonics for all GemStone errors. Search the file for the mnemonic name or error number to locate an error in the file. The errors are divided into five groups: compiler, run-time (virtual machine), aborting, fatal, and event.

GemBuilder provides functions that allow you to poll for errors or to use error jump buffers. The following paragraphs describe both of these techniques.

Polling for Errors

Each call to GemBuilder can potentially fail for a number of reasons. Your program can call **GciErr** to determine whether the previous GemBuilder call resulted in an error. If so, **GciErr** will obtain full information about the error. If an error occurs while Smalltalk code is executing (in response to **GciPerform** or one of the **GciExecute...** functions), your program may be able to continue Smalltalk execution by calling **GciContinue**.

If you are in a `UserAction` and want to pass the error back to the Smalltalk client, you must execute `GciRaiseException`.

Error Jump Buffers

When your program makes three or more GemBuilder calls in sequence, jump buffers provide significantly faster performance than polling for errors.

When your C program calls `Gci_SETJMP`, the context of the current C environment is saved in a jump buffer designated by your program. GemBuilder maintains a stack of up to 20 error jump buffers. A buffer is pushed onto the stack when `GciPushErrJump` is called, and popped when `GciPopErrJump` is called. When an error occurs during a GemBuilder call, the GemBuilder implementation calls `GciLongJump` using the buffer currently at the top of GemBuilder's error jump stack, and pops that buffer from the stack.

For functions with local error recovery, your program can call `GciSetErrJump` to temporarily disable the `GciLongJump` mechanism (and to re-enable it afterwards).

Whenever the jump stack is empty, the application must use `GciErr` to poll for any GemBuilder errors.

The Call Stack

The Smalltalk virtual machine creates and maintains a *call stack* that provides information about the state of execution of the current Smalltalk expression or sequence of expressions. The call stack includes an ordered list of activation records related to the methods and blocks that are currently being executed. The virtual machine ordinarily clears the call stack before each new expression is executed.

If a soft break or an unexpected error occurs, the virtual machine suspends execution, creates a Process object, and raises an error. The Process object represents both the Smalltalk call stack when execution was suspended and any information that the virtual machine needs to resume execution. If there was no fatal error, your program can call `GciContinue` to resume execution. Call `GciClearStack` instead if there was a fatal error, or if you do not want your program to resume the suspended execution.

GemStone System Errors

If your application receives a GemStone system error while linked with `GciLnk`, relink your application with `GciRpc` and run it again with an uncorrupted copy of your repository. Your GemStone system administrator can refer to the repository backup and recovery procedures in the *System Administration Guide for GemStone/S 64 Bit*.

If the error can be reproduced, contact GemStone Customer Support. Otherwise, the error is in your application, and you need to debug your application before using `GciLnk` again.

1.11 Garbage Collection

GemStone performs automatic garbage collection via several mechanisms, which are discussed more fully in the chapter "GemStone Garbage Collection" in the *System Administration Guide for GemStone/S 64 Bit*.

In-memory garbage collection of non-persistent temporary objects occurs regularly, to avoid low and out of memory issues. If newly created or temporary objects are not

referenced, they run the risk of being garbage collected and disappearing prematurely during in-memory garbage collection. To avoid this problem GemStone uses several internal sets: the *PureExportSet*, the *GciTrackedObjs* set, and the user action's export set.

Before removing any objects, the GemStone in-memory garbage collector checks the *PureExportSet* and the *GciTrackedObjs* set in the user session's workspace, and if in a user action, the user action's export set. Any object in these sets is considered to be referenced. The garbage collector does not remove objects that are in these sets, or objects that are referenced by a persistent object. It also does not remove any additional objects that they refer to, or more objects that those additional objects refer to, and so on.

Some functions will automatically add the objects which they return to the export sets. Objects may also be added and removed explicitly. Objects are automatically added to an export set in these cases:

- ▶ The results of *GciNew**, *GciCreate**, *GciSend**, *GciPerform**, *GciExecute** and *GciResolve** calls are automatically added to the applicable export set - either the *PureExportSet*, or if the function is called from within a user action, to the user action's export set.
- ▶ Objects returned in the report buffer of a *GciFetchObjectInfo* or *GciClampedTrav* when *GCI_RETRIEVE_EXPORT* flag is set will be added to the *PureExportSet* or the user action's export set.
- ▶ When the function

```
GciErr(GciErrSType *errorReport);
```

returns TRUE, values of type *OopType* in the **errorReport* are added to the applicable export set.

All of these functions return their results to the C application in the form of one or more OOPs (objects), through either return values or output parameters. To protect these result objects from premature garbage collection, GemBuilder *automatically* adds all of them to the applicable export set. GemBuilder does not automatically add other objects to the export sets; the application should be careful to explicitly call the **GciSaveObjs** or **GciSaveGlobalObjs** function when it needs to be sure to retain an object that is not already in an export set.

Objects that are the contents of instance variables, such as objects returned from a call to *GciFetchOops*, are not added to the export sets. These are already referenced from the object whose instance variable references them. Note however that these objects are cached in your C code, and the values may no longer be valid if the referencing object becomes dirty due to an abort or commit.

Persistent objects may be added to any of the three sets, in which case they are protected from garbage collection on persistent objects, such as *markForCollection*.

In a user action, some of these functions behave differently. When these functions are called from within a user action, the objects are added to the user action's export set to prevent them from being garbage collected, rather than to the *PureExportSet*. When the user action comes to an end, the user action's export set ceases to exist and the objects it contained may be garbage collected. This avoids the risk of objects not being released and consuming excess memory, for example if the user action exits with an unexpected error. In order to prevent objects saved from within a user action from being released prematurely, the user action can explicitly call *GciSaveGlobalObjs*, which will save them to the *PureExportSet* regardless of the user action context.

Once the objects in the GciTrackedObjs or in the export sets are no longer needed, the application can improve performance and avoid out of memory issues by calling the **GciRelease...** functions, to reduce the size of the set and permit garbage collection of obsolete temporaries.

1.12 Preparing to Execute GemStone Applications

The following information includes the requirements and recommendations for preparing your environment to execute C applications for GemStone. Your application may have additional requirements, such as environment variables that it uses.

GemStone Environment Variables

Anyone who runs a GemStone application or process is responsible for setting the following environment variables:

GEMSTONE – A full pathname to your GemStone installation directory.

PATH – Add the GemStone bin directory to your path.

The following environment variables influence the behavior of GemStone and GemBuilder. You may wish to supply values or defaults for them when you or your users run your application or a Gem.

GEMSTONE_EXE_CONF – (not for RPC applications) A full path to a special GemStone configuration file for an executable, if any. See the *System Administration Guide for GemStone/S 64 Bit* for details.

GEMSTONE_SYS_CONF – (not for RPC applications) a full path to a special GemStone configuration file for your system, if any. See the *System Administration Guide for GemStone/S 64 Bit* for details.

GEMSTONE_NRS_ALL – A network resource string – a means for identifying certain GemStone file and process information. It can identify the name of the script to run to start an RPC Gem. See the *System Administration Guide for GemStone/S 64 Bit* for details.

Building Applications with GemBuilder for C

This chapter explains how to use GemBuilder for C to build your C application. Your application will provide the main program and API. When linking your compiled code, you will link the GemBuilder for C libraries that allow you to make GemBuilder function calls.

2.1 RPC and Linked applications

The GemBuilder for C interface provide two versions: Remote Procedure Call (RPC) and Linked.

In an RPC application, your application exists in a process separate from the Gem. The two processes communicate through remote procedure calls. This interprocess communication requires a small overhead associated with each GemBuilder call, independent of whatever object access is performed or Smalltalk code is executed.

In a linked application, your application and Gem (the GemStone session) exist as a single process. Your application is expected to provide the main entry point.

A linked application can also run RPC Gems, in addition to its linked Gem. An RPC application cannot have a linked Gem.

The function **GciIsRemote** reports whether your application was linked with the RPC version or Linked version of the GemBuilder interface.

RPC for Debugging

When debugging a new application, you must use the RPC interface. You should use **GciLnk** *only* after your application has been completely debugged and thoroughly tested, to avoid the risk of errors corrupting your repository.

When using an RPC Gem, you usually achieve the best performance by using functions such as **GciTraverseObjs**, **GciStoreTrav**, and **GciFetchPaths**. Those functions are designed to reduce the number of network round-trips by combining functions that are commonly used in sequence.

Linked for Performance

You can use a linked, single-Gem configuration to enhance performance. In a linked application, a GemBuilder function call is a machine-instruction procedure call rather than a remote call over the network to a different process.

WARNING!

*Before using GciLnk, debug your C code in a process that does not include a Gem!
For more information, see section "Risk of Database Corruption" on page 66.*

In a linked application, you usually achieve the best performance by using the simple **GciFetch...** and **GciStore...** functions instead of the complex object traversal functions. This makes the application easier to write.

However, if your application will login to GemStone multiple times, any sessions after the first will be RPC gems. The traversal functions should perform better in those sessions.

Multiple GemStone Sessions

If your application will be running multiple GemStone sessions simultaneously, or if you will need to run your application and the GemStone session on separate machines, then you will need to use either the RPC version of GemBuilder, or a non-default linked login session.

The GemBuilder for C Shared Libraries

The two versions of GemBuilder are provided as a set of shared libraries. A *shared library* is a collection of object modules that can be bound to an executable, usually at run time. The contents of a shared library are not copied into the executable. Instead, the library's main function loads all of its functions. Only one copy is loaded into memory, even if multiple client processes use the library at the same time. Thus, they "share" the library.

The GemBuilder library files `libgcilnk.*` and `libgcirpc.*` are provided in the GemStone distribution under `$GEMSTONE/lib`.

2.2 GemBuilder Applications that load Shared Libraries

Most applications will load shared libraries at run time. The binding is done by code that is part of the application. If that code is not executed, the shared library is not loaded. With this type of binding, applications can decide at run time which GemBuilder library to use. They can also unbind at run time and rebind to the same or different shared libraries. The code is free to handle a run-time-bind error however it sees fit.

Building the Application

To build an application that loads GemBuilder shared libraries:

1. Include `gci.ht` and `gcirtl.hf` in the C source code.

```
#include "gci.hf"  
#include "gcirtl.hf"
```

However, applications are free to use their own run-time-bind interface instead of `gcirtl`, which is meant to be used from C. For example, a Smalltalk application would use the mechanism provided by the Smalltalk vendor to call a shared library.

2. Call **GciRtlLoad**(*useRpc, path, errBuf, errBufSize*) to load the RPC GemBuilder.

To load the RPC version, pass **true** as the first argument; to load the linked version, pass in **false**.

Call **GciRtlLoad** before any other GemBuilder calls. To unload the GemBuilder libraries, call **GciRtlUnload**. The **GciRtlIsLoaded** function can be used to determine if an interface is loaded or not.

3. When linking, include `gcirtlobj.o`, not one of the GemBuilder libraries (`libgcirpc.*` and `libgcilnk.*`). Details of what should be included and other linker flags are provided starting on page 68.

Chapter 4, “Compiling and Linking” tells how to compile and link your application.

Searching for the Library

At run time the `gcirtl` code searches for the GemBuilder library in the following places:

1. Any directories specified by the application in the argument to **GciRtlLoad**.
2. The `$GEMSTONE/lib` directory.
3. The normal operating system search, as described in the following sections.

How UNIX Matches Search Names with Shared Library Files

The UNIX operating system loader searches the following directories for matching file names, in this order:

1. Any path specified by an environment variable:
`LD_LIBRARY_PATH` on Solaris, Linux or AIX
2. Any path recorded in the executable when it was built.
3. The global directory `/usr/lib`.

2.3 Building Statically Linked Applications

Normally, applications will load shared libraries at runtime. However, it is also possible to statically link GemBuilder applications, which in some cases may improve performance.

In a statically linked application, the executable is linked to the GemBuilder library archive file when the executable is built. When the executable runs, it automatically loads the library. Unlike when loading a shared library, your code cannot programmatically determine which library to load. For example, with run-time loading, your code can determine which version of the GemBuilder libraries to load, but this must be determined ahead of time for a statically linked application.

Static linking is only possible with the RPC shared libraries.

To build an application that build-time-binds to GemBuilder:

1. Include `gci.h` in the C source code.
2. When linking, include the RPC GemBuilder library (`libgcirpc.*`).

GemBuilder finds build-time-bound shared libraries by using the normal operating system searches described in “Searching for the Library” on page 45. If a library cannot be found, the operating system returns an error.

Writing C Functions To Be Called from GemStone

Within a GemStone Smalltalk-based application, you may choose to write a C function for certain operations, rather than to perform the work in GemStone. For example, operations that are computationally intensive or are external to GemStone can be written as C functions and called from within a Smalltalk method (whose high-level structure and control is written in Smalltalk). This approach is similar to the concept of “user-defined primitives” offered by some other object-oriented systems.

This chapter describes how to implement C user action functions that can be called from GemStone, and how to call those functions from a GemBuilder application or a Gem (GemStone session) process.

3.1 Shared User Action Libraries

Although user actions can be linked directly into an application, they are usually placed in shared libraries so they can be loaded dynamically. The contents of a library are not copied into the executable. Instead, the library’s main function loads all of its user actions. Only one copy is loaded into memory, even if multiple client processes use the library at the same time. See Chapter 2, “Building Applications with GemBuilder for C” for more information.

User action libraries are used in two ways: They can be *application user actions*, which are loaded by the application process, or *session user actions*, which are loaded by the session process. The operation that is used to load the library determines which type it is, not any quality of the library itself. Application and Gem executables can load any library.

Application user actions are the traditional GemStone user actions. They are used by the application for communication with the Gem or for an interactive interface to the user.

Session user actions add new functionality to the Gem, something like the traditional custom Gem. The difference here is that you only need one Gem, which can customize itself at run time. It loads the appropriate libraries for the code it is running. The decisions are made automatically within GemStone Smalltalk, rather than requiring the users to decide what Gem they need before they start their session.

3.2 How User Actions Work

Here's a quick overview of the sequence of events when a user action function is executed:

1. The Gem or your C application program initiates GemStone Smalltalk execution by calling one of the following functions: **GciExecute**, **GciExecuteStr**, **GciExecuteStrFromContext**, **GciPerform**, or **GciContinue**.
2. Your GemStone Smalltalk code invokes a user action function (written in C) by sending a message of the form:


```
System userAction: aSymbol with: args
```

The *args* arguments are passed to the C user action function named *aSymbol*. (You must have already initialized that function before logging in to GemStone. See Chapter , "Loading User Actions".)
3. The C user action function can call any GemBuilder functions and any C functions provided in the application or the libraries loaded by the application (for application user actions), or provided in the libraries loaded by the Gem (for session user actions).

Specifically, the C user action function can call GemBuilder's structural access functions (**GciFetch...** and **GciStore...**, etc.) to read or modify, respectively, any objects that were passed as arguments to the user action.

If a GemBuilder or other GemStone error is encountered during execution of the user action, control is returned to the Gem or your GemBuilder application as if the error had occurred during the call to **GciExecute** (or whichever GemBuilder function executed the GemStone Smalltalk code in step 1).
4. The C user action function must return an **OopType** as the function result, and must return control directly to the Smalltalk method from which it was called.

NOTE:

*Results are unpredictable if the C function uses **GCI_LONGJMP** instead of returning control to the GemStone Smalltalk virtual machine.*

3.3 Developing User Actions

For your GemStone application to take advantage of user action functions, you do the following:

- Step 1.** Determine which operations to perform in C user action functions rather than in Smalltalk. Then write the user action functions.
- Step 2.** Create a user action library to package the functions.
- Step 3.** Provide the code to load the user action library.
- Step 4.** If the application is to load the library, add the loading code to your application.
- Step 5.** If the session is to load the library, use the GemStone Smalltalk method `System class>>loadUserActionLibrary:` for loading.
- Step 6.** Write the Smalltalk code that calls your user action. Commit it to your GemStone repository.

Step 7. Debug your user action.

The following sections describe each of these steps.

Write the User Action Functions

Writing a C function to install as a user action called from Smalltalk is little different from writing other C functions. However, one important difference exists: user actions cannot reliably retain references to objects they create. The application that called the user action (whether written in C, Java, or Smalltalk) controls the export set – the set of OOPs to save after execution completes. Therefore, make sure your C application treats all argument and result objects of a user action as temporary objects. Don't save the OOPs in static C variables for use by a subsequent invocation of the user action or by another C function.

Don't rely on **GciSaveObjs** to make the objects persistent. The application that called the user action can still call **GciReleaseOops** on the object that the user action needs to retain (or **GciReleaseAllOops** to release all objects at once).

To make a newly created object a permanent part of the GemStone repository, the user action has two options:

- ▶ Store the OOP of the new object into an object known to be permanent, such as a collection created by the calling application (for example, a collection created in Smalltalk and committed to the repository).
- ▶ Return the OOP of the object as the function result.

After a user action returns, the persistence of the new object is determined by the normal semantics of the calling application.

If you are working in GemBuilder for Smalltalk, you can also explicitly save these user action objects by populating a collection in the user-definable portion of `System sessionState` using `System > sessionStateAt:put:.` Your user action can retain references to objects that you add to this collection in this way.

Create a User Action Library

Whether you have one user action or many, the way in which you prepare and package the source code for execution has significant effects upon what uses you can make of user actions at run time. It is important to visualize your intended execution configurations as you design the way in which you package your user actions.

To build a user action library:

1. Include `gciua.hf` in your C source code.
2. Define the initialization and shutdown functions.
3. Compile with shared library switches.
4. Link with `gciualib.o` and shared library switches.
5. Install the library in the `$GEMSTONE/uilib` directory.

The `gciua.hf` Header File

User action libraries must always include the `gciua.hf` file, rather than the `gci.hf` or `gcirtl.hf` file. Using the wrong file causes unpredictable results.

The Initialization and Shutdown Functions

A user action library must define the initialization function **GciUserActionInit** and the shutdown function **GciUserActionShutdown**.

Do not call **GciInit**, **GciLogin**, or **GciLogout** within a user action.

Defining the Initialization Function

Example 3.1 shows how the initialization function **GciUserActionInit** is defined, using the macro **GCIUSER_ACTION_INIT_DEF**. This macro must call **GciDeclareAction** once for each function in the set of user actions.

Example 3.1

```

static OopType doParse(void)
{
    return OOP_NIL;
}
static OopType doFetch(void)
{
    return OOP_NIL;
}

GCIUSER_ACTION_INIT_DEF()
{
    GciDeclareAction("doParse", doParse, 1, 0, TRUE);
    GciDeclareAction("doFetch", doFetch, 1, 0, TRUE);
    // ...
}

```

GciDeclareAction associates the Smalltalk name of the user action function *userActionName* (a C string) with the C address of that function, *userActionFunction*, and declares the number of arguments that the function takes. A call to **GciDeclareAction** looks similar to this:

```
GciDeclareAction("userActionName", userActionFunction, 1, 0, TRUE)
```

The function installs the user action into a table of such functions that GemBuilder maintains. Once a user action is installed, it can be called from GemStone.

The name of the user action, "*userActionName*", is a case-sensitive, null-terminated string that corresponds to the symbolic name by which the function is called from Smalltalk. The name is significant to 31 characters. It is recommended that the name of the user action be the same as the C source code name for the function, *userActionFunction*.

The third argument to **GciDeclareAction** indicates how many arguments the C function accepts. This value should correspond to the number of arguments specified in the Smalltalk message. When it is 0, the function argument is void. Similarly, a value of 1 means one argument. The maximum number of arguments is 8. Each argument is of type **OopType**.

The fourth argument to **GciDeclareAction** is rarely used. The final argument indicates whether to return an error if there is already a user action with the specified name.

Your user action library may call **GciDeclareAction** repeatedly to install multiple C functions. Each invocation of **GciDeclareAction** must specify a unique *userActionName*. However, the same *userActionFunction* argument may be used in multiple calls to **GciDeclareAction**.

Defining the Shutdown Function

The shutdown function **GciUserActionShutdown** is defined by the **GCIUSER_ACTION_SHUTDOWN_DEF** macro. **GciUserActionShutdown** is called when the user action library is unloaded. It is provided so the user action library can clean up any system resources it has allocated. Do not make GemBuilder C calls from this function, because the session may no longer exist. In fact, **GciUserActionShutdown** can be left empty. Example 3.2 shows a shutdown definition that does nothing but report that it has been called.

Example 3.2

```
#include "gciuser.hf"

GCIUSER_ACTION_SHUTDOWN_DEF()
{
    /* Nothing needs to be done. */
    fprintf(stderr, "GciUserActionShutdown called.\n");
}
```

Compiling and Linking Shared Libraries

Shared user actions are compiled for and linked into a shared library. See Chapter 4, “Compiling and Linking” for instructions.

Be sure to check the output from your link program carefully. Linking with shared libraries does not require that all entry points be resolved at link time. Those that are outside of each shared library await resolution until application execution time, or even until function invocation time. You may not find out about incorrect external references until run time.

Using Existing User Actions in a User Action Library

With slight modifications, existing user action code can be used in a user action library. You need to include `gciua.hf` instead of `gci.hf` or `gcirtl.hf`. Define a **GciUserActionShutdown**, and a **GciUserActionInit**, if it is not already present. Compile, link, and install according to the instructions for user action libraries.

Using Third-party C Code with a User Action Library

Third-party C code has to reside in the same process as the C user action code. Link the third-party code into the user action library itself, and then you can call that code. It doesn't matter where you call it from.

Loading User Actions

GemBuilder does not support the loading of any default user action library. Applications and Gems must include code that specifically loads the libraries they require.

Loading User Action Libraries At Run Time

Dynamic run-time loading of user action libraries requires some planning to avoid name conflicts. If an executable tries to load a library with the same name as a library that has already been loaded, the operation fails.

When user actions are installed in a process, they are given a name by which GemBuilder refers to them. These names must be unique. If a user action that was already loaded has the same name as one of the user actions in the library the executable is attempting to load, the load operation fails. On the other hand, if the two libraries contain functions with the same implementation but different names, the operation succeeds.

Application User Actions

If the application is to load a user action library, implement an application feature to load it. The GemStone interfaces provide a way to load user actions from your application.

- ▶ GemBuilder for C applications: the **GciLoadUserActionLibrary** call
- ▶ Topaz applications: the **loadua** command

The application must load application user actions after it initializes GemBuilder (**GciInit**) and before the user logs into GemStone (**GciLogin**). If the application attempts to install user actions after logging in, an error is returned.

Session User Actions

A linked or RPC Gem process can install and execute its own user action libraries. To cause the Gem to do this, use the `System class>>loadUserActionLibrary:` method in your GemStone Smalltalk application code. A session user action library stays loaded until the session logs out.

The session must load its user actions after the user logs into GemStone (**GciLogin**). At that time, any application user actions are already loaded. If a session tries to load a library that the application has already defined, it gets an error. The loading code can be written to handle the error appropriately. Two sessions can load the same user action library without conflict.

Specifying the User Action Library

When writing scripts or committing to the database, you can specify the user action library as a full path or a simple base name. Always use the base name when you need portability. The code that GemBuilder uses to load a user action library expands the base name *ua* to a valid shared library name for the current platform:

- ▶ Solaris: `libua.so`
- ▶ HP-UX: `libua.sl`
- ▶ AIX: `libua.so`
- ▶ Linux (x86_64): `libua.so`

► Darwin: `libua.dylib`

and searches for the file in the following places in the specified order:

1. The current directory of the application or Gem.
2. The directory the executable is in, if it can be determined.
3. The `$GEMSTONE/ua/lib` directory.
4. The normal operating system search, as described in “Searching for the Library” on page 45.

Creating User Actions in Your C Application

Loading user action libraries at run time is the preferred behavior for GemBuilder applications. For application user actions, however, you have the option to create the user actions directly in your C application, not as part of a library. When you implement user actions this way, include `gcirtl.hf` or `gci.hf` in your C source code, instead of `gciua.hf`. (Your C source code should include *exactly* one of these three include files.)

The `GciUserActionInit` and `GciUserActionShutdown` functions are not required, but the application must call `GciDeclareAction` once for each function in the set of user actions.

After your application has successfully logged in to GemStone (via `GciLogin`), it may not call `GciDeclareAction`. If your application attempts to install user actions after logging in, an error will be returned.

Verify That Required User Actions Have Been Installed

After logging in to GemStone, your application can test for the presence of specific user actions by sending the following Smalltalk message:

```
System hasUserAction: aSymbol
```

This method returns `true` if your C application has loaded the user action named `aSymbol`, `false` otherwise.

For a list of all the currently available user actions, send this message:

```
System userActionReport
```

Write the Code That Calls Your User Actions

Once your application or Gem has a way to access the user action library, your GemStone Smalltalk code invokes a user action function by sending a message to the GemStone system. The message can take one of the following forms:

```
System userAction: aSymbol
System userAction: aSymbol with:arg1 [with:arg2] ...
System userAction: aSymbol withArgs:anArrayOfUpTo8Args
```

You can use the `with` keyword from zero to seven times in a message. The `aSymbol` argument is the name of the user action function, significant to 31 characters. Each method returns the function result.

Notice that these methods allow you to pass up to eight arguments to the C user action function. If you need to pass more than eight objects to a user action, you can create a

Collection (for example, an instance of Array), store the objects into the Collection, and then pass the Collection as a single argument object to the C user action function:

```
| myArray |  
myArray := Array new: 10.
```

```
"populate myArray, then send the following message"
```

```
System userAction: #doSomething with: myArray.
```

NOTE

You can also call a user action function directly from your C code, as you would any other C function.

Remote User Actions

The user action code that is called can be remote (on a different machine) from the Gem that invokes this method.

Limit on Circular Calls Among User Actions and Smalltalk

From Smalltalk you can invoke a user action, and within the user action you can do a **GciSend**, **GciPerform**, or **GciExecute**, that may in turn invoke another user action. This kind of circular function calling is limited; no more than 2 user actions may be active at any one time on the current Smalltalk stack. If the limit is exceeded, GemStone raises an error.

Debug the User Action

Even if you intend to use your library only as session user actions, test them first as application user actions with an RPC Gem. As with applications, never debug user actions with linked versions.

CAUTION

*Debug your C code in a process that does not include a Gem.
For more information, see "Risk of Database Corruption" on page 66.*

Use the instructions for user actions in Chapter 4, "Compiling and Linking" to compile and link the user action library. Then load the user actions from the RPC version of your application or Topaz. To load from Topaz, use the **loadua** command.

3.4 Executing User Actions

User actions can be executed either in the GemBuilder application (client) process or in a Gem (server) process, or in both.

Choosing Between Session and Application User Actions

The distinction between application user actions that execute in the application and session user actions that execute in the Gem is interesting primarily when the two processes are running remotely, or when the application has more than one Gem process.

Remote Application and Gem Processes

When the application and Gem run on different machines and the Gem calls an application user action, the call is made over the network. Computation is done by the application where the application user action is running, and the result is returned across the network. Using a session user action eliminates this network traffic.

On the other hand, for overall efficiency you also need to consider which machine is more suitable for execution of the user action. For example, assume that your application acquires data from somewhere and wishes to store it in GemStone. You could write a user action to create GemStone objects from the data and then store the objects. It might make more sense to execute the user action in the application process rather than transport the raw data to the Gem.

Alternatively, assume there is a GemStone object that could require processing before the application could use it, like a matrix on which you need to perform a Fast Fourier Transform (FFT). If the Gem runs on a more powerful machine than the client, you may wish to run an FFT user action in the Gem process and send the result to your application.

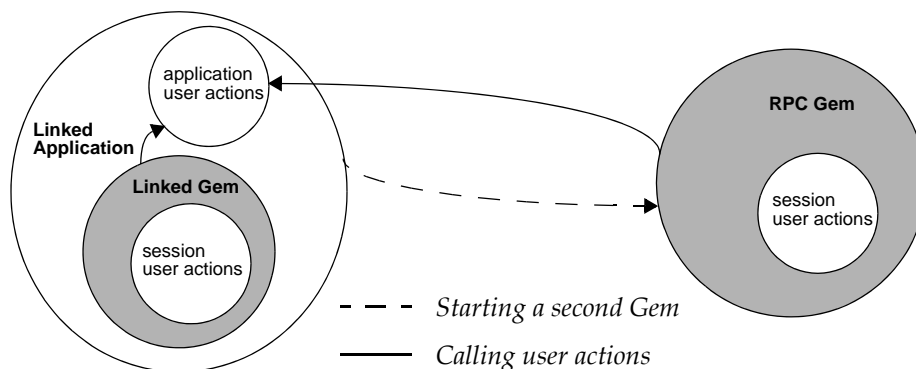
Applications With Multiple Gems

In most situations, session user actions are preferable, because the Gem does not have to make calls to the application. In the case of a linked application, however, an application user action is just as efficient for the linked Gem, because the Gem and application run as one process. Using an application user action guarantees that if any new sessions are created, they will have access to the same user action functions as the first session.

Every Gem can access its own session user actions and the application user actions loaded by its application. A Gem cannot access another Gem's session user actions, however, even when the Gems belong to the same application.

Although a linked application and its first Gem run in the same process, that process can have session and application user actions, as in Figure 3.1. Application user actions, loaded by the application's loading function, are accessible to all the Gems. Session user actions in the same process, loaded by the `System class>>loadUserActionLibrary:` method, are not accessible to the RPC Gem. Conversely, the RPC Gem's user actions are not accessible to the linked Gem.

Figure 3.1 Access to Application and Session User Actions



The following sections discuss the various possible configurations in detail.

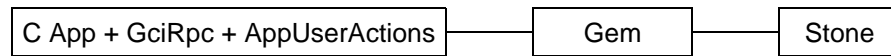
Running User Actions with Applications

User actions can be executed in the user application process under two configurations of GemStone processes. The configurations differ depending upon whether the application is linked or RPC.

With an RPC Application

Figure 3.2 illustrates how various architectural components are distributed among three GemStone processes when a set of user actions executes with an RPC application.

Figure 3.2 Application User Actions and RPC Applications in GemStone Processes



In this configuration, the application runs in a separate process from any Gem. Each time the application calls a GemBuilder C function, the function uses remote procedure calls to communicate with a Gem. The remote procedure calls are used whether the Gem is running on the same machine as the application, or on another machine across the network.

The user actions run in the same process as the application. If they call GemBuilder functions, those functions also use remote procedure calls to communicate with the Gem. In this configuration, all your code executes as a GemStone client (on the application side). It can thus execute on any GemStone client platform; it is not restricted to GemStone server platforms. Care should be taken in coding to minimize remote procedure call overhead and to avoid excessive transportation of GemStone data across the network. The following list enumerates some of the conditions in which you may find occasion to use this configuration:

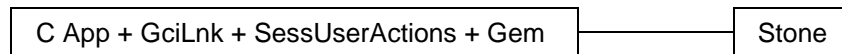
- ▶ The application and/or the user action needs to be debugged or tested.
- ▶ The user action depends on facilities or implement capabilities for the application environment. Screen management, GUI operations, and control of specialized hardware are possibilities.
- ▶ The application acquires data from somewhere and wishes to store it in GemStone. The user action creates the requisite GemStone objects from the data and then commits them to the repository.

NOTE:

*You can run RPC Topaz as the C application in this configuration for debugging to perform unit testing of user action libraries. Apply a source-level debugger to the Topaz executable, load the libraries with the Topaz **loadua** command, then call the user actions directly from GemStone Smalltalk.*

With a Linked Application

Figure 3.1 illustrates how various architectural components are distributed between two GemStone processes when a set of user actions executes with a linked application.

Figure 3.1 Session User Actions and Linked Applications in GemStone Processes

In this configuration, the application, the user actions, and one Gem all run in the same process (on the same machine). All function calls, from the application to GemBuilder and between GemBuilder and the Gem, are resolved by ordinary C-language linkage, not by remote procedure calls.

Since a Gem is required for each GemStone session, the first session uses the (linked) Gem that runs in your application process. This Gem has the advantages that it does not incur the overhead of remote procedure calls, and may not incur as much network traffic. It has the disadvantage that it must run in the same process as the Gem, so that work cannot be distributed between separate client and server processes. Since the application cannot continue processing while the Gem is at work, the non-blocking GemBuilder functions provide no benefit here.

If a linked application user logs in to GemStone more than once, GemStone creates a new RPC Gem process for each new session. (These sessions would be additions to the configuration of Figure 3.1.) If one of these sessions invokes a user action, the user action executes in the same process as the application. If the user action then calls a GemBuilder function, that call is serviced by the linked Gem, not by the Gem from which the user action was invoked.

In this configuration, your code executes only on GemStone server platforms. It cannot execute on client-only platforms because a Gem is part of the same process. The occasions for using this configuration are much the same as those for running user actions with an RPC application, except that you should not use this one for debugging.

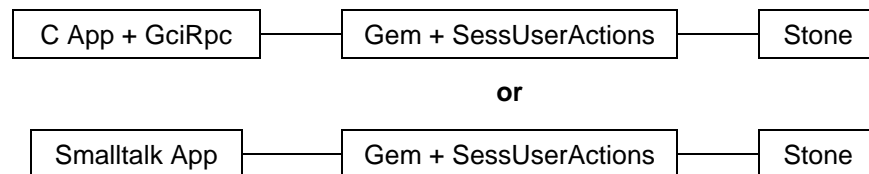
CAUTION

Debug your user actions in a process that does not include a Gem. For more information, see "Risk of Database Corruption" on page 66.

Running User Actions with Gems

Just as with applications, there are two forms of Gems: linked and RPC. The linked Gem is embedded in the `gciLnk` library and is only used with linked applications.

Figure 3.2 illustrates how various architectural components are distributed among three GemStone processes when a set of user actions executes with an RPC Gem.

Figure 3.2 Session User Actions and RPC Gems in GemStone Processes

An RPC Gem executes in a separate process that can install and execute its own user actions. The RPC Gem is RPC because it communicates by means of remote procedure calls, through an RPC GemBuilder, with an application in another process.

However, it is also a separate C program. The Gem itself also uses GemBuilder directly, to interact with the database. That is the reason why the RPC Gem is linked with the `gciLnk`

library. The user action in this configuration executes in the same process as the Gem, with the GemBuilder that does *not* use remote procedure calls.

CAUTION

Debug your user actions in a process that does not include a Gem. For more information, see “Risk of Database Corruption” on page 66.

The following list enumerates some of the conditions in which you may find occasion to use this configuration:

- ▶ You wish to execute the user action from a Smalltalk application using GemBuilder for Smalltalk. This configuration is required for that purpose.
- ▶ You wish the user action to be available to all or many other C applications.
- ▶ The user action is called frequently from GemStone. This configuration eliminates network traffic between GemBuilder and GemStone.
- ▶ The user action makes many calls to GemBuilder. This configuration avoids remote procedure call overhead.
- ▶ You have a GemStone object or objects that you wish to process first, and your application needs the result. The processing may be substantial. Your GemStone server machine may be more powerful than your client machine and could do it more quickly, or it might have specialized software the user action needs. Also, the result might be smaller and could reduce network traffic.

For example, the user action might retrieve a data matrix and a filter from GemStone, perform a Fast Fourier Transform, and send the result to the application.

Running User Actions with Applications and Gems

Figure 3.3 illustrates how various architectural components are distributed among three GemStone processes when one set of user actions executes with an RPC application and another set of user actions executes with an RPC Gem.

Figure 3.3 RPC Applications and Gems with User Actions in GemStone Processes

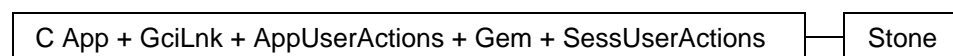


This configuration is a combination of previous configurations. The application and the Gem run in separate processes. User actions in the first set execute in the application process, and user actions in the second set execute in the Gem process.

When user actions are installed in a process, they are given a name by which GemBuilder refers to them. If a user action in the application has the same name as a user action in the Gem, then the one in the Gem is always used, and the one in the application is ignored.

The two types of user actions could also exist in one linked process, as shown in Figure 3.4.

Figure 3.4 Application and Session User Actions in GemStone Processes



In this configuration, the user actions can be loaded as either application or session user actions; it would be the same from the point of view of the linked Gem. Application user

actions would be just as efficient as session user actions, because they are part of the Gem process. If a linked application user logs in to GemStone more than once, GemStone creates a new RPC Gem process for each new session, additions to the configuration of Figure 3.4. The RPC Gems do not have access to the linked Gem's session user actions. So it is generally better to load them as application user actions, just in case.

Compiling and Linking

This chapter describes how to compile and link your C/C++ applications and user actions. The focus is directly on operations for each compiling or linking alternative on each GemStone server platform. It is assumed that you already know which alternatives you want to use, and why, and when. Those topics are part of the application design and code implementation, which are described in other chapters of this manual.

All operations are illustrated as though you are issuing commands at a command-line prompt. You may choose to take advantage of your system's programming aids, such as the UNIX **make** utility and predefined environment variables, to simplify compilation and linking. Whatever you choose, be sure that you designate options and operations that are equivalent to those shown here.

NOTE

Much of the material in this chapter is system-specific and, therefore, subject to change by compiler vendors and hardware manufacturers. Please check your GemStone/S 64 Bit Release Notes, Installation Guide, and vendor publications for possible updates.

4.1 Development Environment and Standard Libraries

For simplicity, set the GEMSTONE environment variable to your GemStone installation directory. The command lines shown in this chapter assume that this has been done. No other environment variables are required to find the GemStone C++ libraries.

GemStone requires linking with certain architecture-specific "standard" C and C++ libraries on some platforms. The order in which these libraries are specified can be significant; be sure to retain the ordering given in the command lines to follow in this section.

The environment of the supported Unix platforms is System V. On these platforms, the `/usr/bin` directory should be present in the PATH environment variable. If `/usr/ucb` is also present in PATH, then it should come after `/usr/bin`. The System V "standard" C/C++ libraries (*not* Berkeley) should be used in linking.

4.2 Compiling C Source Code for GemStone

The following information includes the requirements and recommendations for compiling C applications or user actions for GemStone. Your C code may have additional requirements, such as compile options or environment variables.

The C++ Compiler

C applications and user actions must be compiled and linked with a compiler that is compatible with GemStone libraries and object code.

The example compiler and linker command lines in this chapter assume that a compatible compiler has been installed and is in your path.

The following C++ compilers were used to produce the GemStone product, and have been tested for producing C/C++ applications and user action libraries. Other compilers, such as ANSI C++ compilers, are assumed to work, but have not been tested.

Linux:

Red Hat Linux 6.x: gcc/gcc/g++ 4.4.7

Red Hat Linux 7.x: gcc/gcc/g++ 4.8.5

Ubuntu Linux: gcc/g++ 4.6.3

SUSE Linux: gcc/g++ 4.8.5

Solaris (SPARC):

CC: Sun C++ 5.8 Patch 121017-05 2006/08/30

Solaris (x86):

Sun C++ 5.10 SunOS_i386 128229-09 2010/06/24

AIX (IBM):

AIX 6.1 and 7.1 on POWER7: IBM XL C/C++ for AIX, V11.1

AIX 7.1 on POWER8: 7.1: IBM XL C/C++ for AIX, V13.1.2

Darwin:

Apple LLVM version 6.0 (clang-600.0.56)

Windows:

Microsoft Visual Studio 2010 Version 10.0.30319.1 RTMRel

Microsoft Visual C++ 2010 01021-532-2002102-70611

Listing the Version of Your Compiler

To list the version of your compiler, execute the appropriate command line.

Linux and Darwin:

```
% g++ -v
```

Solaris (SPARC and x86):

```
% CC -V
```

AIX (IBM):

```
% /usr/vacpp/bin/xlC_r -qversion
```

Compilation Options

When you compile, specify each directory that is to be searched for include files separately by repeating the `-I` option. At a minimum, you should specify the GemStone `include` directory.

The `-c` option inhibits the “load and go” operation, so compilation ends when the compiler has produced an object file.

For more information and details on the listed compiler options and other compiler flags, please consult your compiler documentation.

Compilation Command Lines

This section presents simple example command lines for compiling C source code on each platform.

The command lines for each platform illustrates how to compile a simple application program or user action file named *userCode*, whose source contains one code file, *userCode.c*. Its result is one object file, *userCode.o*.

For simplicity in compiling code for user actions, this file is assumed to be a library containing both the source code for one set of user actions and the implementation of the function that installs them all with GemStone.

If you have multiple application or user action files, they should all be compiled under these same basic conditions.

Linux:

```
g++ -fmessage-length=0 -fcheck-new -O3 -ggdb -m64 -pipe
-D_REENTRANT -D_GNU_SOURCE -pthread -fPIC -fno-strict-aliasing
-fno-exceptions -I$GEMSTONE/include -x c++ -c userCode.c
-o userCode.o
```

The following warn flags are recommended for compilation:

```
-Wformat -Wtrigraphs -Wcomment -Wsystem-headers -Wtrigraphs
-Wno-aggregate-return -Wswitch -Wshadow -Wunused-value
-Wunused-variable -Wunused-label -Wno-unused-function
-Wchar-subscripts -Wmissing-braces -Wmultichar -Wparentheses
-Wsign-compare -Wsign-promo -Wwrite-strings -Wreturn-type
-Wuninitialized
```

If you want to stop the compilation process when any of the above warnings are encountered, use the following flag:

```
-Werror
```

To allow debugging of the resulting library, also include the optional `-g` flag and omit the optimization flag `-O3`.

Solaris (SPARC):

```
CC -xO4 -xcode=pic32 -xarch=v9 -mt -xchip=ultra2
-I$GEMSTONE/include -features=no%except -D_REENTRANT
-D_POSIX_PTHREAD_SEMANTICS -features=no%anachronisms
-c userCode.c -o userCode.o
```

To allow debugging of the resulting library, include the optional `-g` flag and omit the optimization flag `-xO4`.

Solaris (x86):

```
CC -xO4 -m64 -xarch=generic -Kpic -mt -D_REENTRANT
-D_POSIX_PTHREAD_SEMANTICS -I$GEMSTONE/include
-features=no%except -c userCode.c -o userCode.o
```

To allow debugging of the resulting library, include the optional `-g` flag and omit the optimization flag `-xO4`.

AIX (IBM):

```
xlc_r -O3 -qstrict -qalias=noansi -q64 -+ -qpic -qthreaded
-qarch=pwr6 -qtune=balanced -D_LARGEFILE64_SOURCE
-DFLG_AIX_VERSION=version -D_REENTRANT -D_THREAD_SAFE
-qminimaltoc -qlist=offset -qmaxmem=-1 -qsuppress=1500-010:1500
-029:1540-1103:1540-2907:1540-0804:1540-1281:1540-1090 -qnoeh
-I$GEMSTONE/include -c userCode.c -o userCode.o
```

Depending on your version of AIX, you need to include either `-DFLG_AIX_VERSION=61` or `-DFLG_AIX_VERSION=71`.

Note that there is no space in the `-qsuppress` arguments that are continued on the following line.

To allow debugging of the resulting library, also include the optional `-g`, `-qdbxextra` and `-qfullpath` flags, and omit the optimization flag `-O3`.

Darwin:

```
g++ -fmessage-length=0 -O3 -ggdb -m64 -pipe -fPIC
-fno-strict-aliasing -D_LARGEFILE64_SOURCE -D_XOPEN_SOURCE
-D_REENTRANT -D_GNU_SOURCE -I$GEMSTONE/include -x c++
-c userCode.c -o userCode.o
```

The following warn flags are recommended for compilation:

```
-Wformat -Wtrigraphs -Wcomment -Wsystem-headers -Wtrigraphs
-Wno-aggregate-return -Wswitch -Wshadow -Wunused-value
-Wunused-variable -Wunused-label -Wno-unused-function
-Wchar-subscripts -Wconversion -Wmissing-braces -Wmultichar
-Wparentheses -Wsign-compare -Wsign-promo -Wwrite-strings
-Wreturn-type
```

To allow debugging of the resulting library, also include the optional `-g` flag and omit the optimization flag `-O3`.

Windows:

```
> cl /W3 /Zi /MD /O2 /Oy- -DNDEBUG /TP /nologo /D_LP64 /D_AMD64_ /  
  D_CONSOLE /D_DLL /DWIN32_LEAN_AND_MEAN /  
  D_CRT_SECURE_NO_WARNINGS /DNATIVE  
  /I 'VisualStudioInstallPath\atlmfc\include'  
  /I 'VisualStudioInstallPath\VC\include'  
  /I 'C:\Program Files (x86)\Microsoft SDKs\Windows\v7.0A\Include'  
  /I '%GEMSTONE%\include' -c userCode.c -FouserCode.obj
```

4.3 Linking C/C++ Object Code with GemStone

The following information includes the requirements and recommendations for linking C/C++ applications or user actions with GemStone. Your code may have additional requirements, such as link options or libraries.

Run-time binding is done by code that is part of the application. The same application can use either the RPC or linked GemBuilder libraries with this type of binding.

Linking with shared libraries does not require that all entry points be resolved at link time. Those that are outside of each shared library await resolution until application execution time, or even until function invocation time.

NOTE

When you link a user action shared library, be aware of the dangers of incorrect unresolved external references. If you misspell a function call, you may not find out about it until run-time, when your process dies with an unresolved external reference error. Be sure to check your link program's output carefully.

Risk of Database Corruption

CAUTION

Debug your C/C++ code in a process that does not include a Gem.

Do not log into GemStone in a linked application or run a Gem with your user actions until your C/C++ code has been properly debugged.

When your C/C++ code executes in the same process as a Gem, it shares the same address space as the GemStone database buffers and object caches that are part of the Gem. If that C code has not yet been debugged, there is a danger that it might use a C pointer erroneously. Such an error could overwrite the Gem code or its data, with unpredictable and disastrous results. It is conceivable that such corruption of the Gem could lead it to perform undesired GemStone operations that might then leave your database irretrievably corrupt. The only remedy then is to restore the database from a backup.

There are three circumstances under which this risk arises:

- ▶ You are running your linked application and you have logged into GemStone.
- ▶ You are running any linked application and you are executing one of your user actions from the application.
- ▶ You are running any Gem, even a remote Gem, and you are executing one of your user actions from the Gem.

To avoid the risk, you must run your C code in some process that does not include a Gem. If the Gem is in a separate process, it has a separate address space that your C code should not be able to access. Use the RPC version of an application, and run any user actions from the application.

Linker

Use the same C++ compiler to link your GemStone C/C++ code as you use to compile it.

Link Options

The `-o` option designates the path of the executable file produced by the link operation.

Be sure to employ at the appropriate times the link option that designates symbolic debugging (often `-g`).

For information on most options, please consult your linker documentation.

Command Line Assumptions

This section presents simple example command lines for linking object code on each platform.

- ▶ How to link a user action object file named `userCode.o` with GemStone libraries to produce a user action library named `libuserAct.so`, `libuserAct.sl`, or `libuserAct.dylib`, depending on your platform.
- ▶ How to link a simple application program with one application object file, `userCode.o`. Its result is one executable file, `userAppl` or `userAppl.exe`, depending on your platform.

If you have multiple application or user action files, they should all be linked under the same basic conditions.

Use the same C++ compiler to link your GemStone C/C++ code as you used to compile it.

Linking User Actions into Shared Libraries

Linux:

```
$ g++ -shared -Wl,-Bdynamic,-hlibuserAct.so userCode.o
  $GEMSTONE/lib/gciualib.o -o libuserAct.so -m64 -lpthread -lcrypt
  -ldl -lc -lm -lrt -lpam -lpam_misc -Wl,-z,muldefs
  -Wl,--warn-unresolved-symbols
```

Solaris (SPARC):

```
$ CC -xarch=v9 -G -Bsymbolic -h libuserAct.so -i userCode.o
  $GEMSTONE/lib/gciualib.o -o libuserAct.so -Bdynamic -lc
  -lpthread -ldl -lrt -lsocket -lnsl -lm -lpam -lCrun -znodefs
```

Solaris (x86):

```
$ CC -m64 -xarch=generic -G -Bsymbolic -h libuserAct.so -i userCode.o
  $GEMSTONE/lib/gciualib.o -o libuserAct.so -Bdynamic -lc
  -lpthread -ldl -lrt -lsocket -lnsl -lm -lpam -lCrun -z nodefs
```

AIX (IBM):

```
$ xlc_r -G -Wl,-bdatapsize:64K -Wl,-btextpsize:64K
  -Wl,-bstackpsize:64K -q64 userCode.o $GEMSTONE/lib/gciualib.o
  -o libuserAct.so -e GciUserActionLibraryMain -L/usr/vacpp/lib
  -lpthreads -lc_r -lC_r -lm -ldl -lbsd -lpam -Wl,-berok
```

Darwin:

```
$ g++ -dynamiclib userCode.o $GEMSTONE/lib/gciualib.o
-o libuserAct.dylib -m64 -lpthread -ldl -lc -lm -lpam -undefined
dynamic_lookup
```

Linking Applications That Bind to GemBuilder at Run Time**Linux:**

```
g++ userCode.o $GEMSTONE/lib/gcirtlobj.o -Wl,-traditional
-Wl,--warn-unresolved-symbols -m64 -lpthread -lcrypt -ldl -lc
-lm -lrt -lpam -lpam_misc -Wl,-z,muldefs -o userAppl
```

Solaris (SPARC):

```
CC -xildoff -xarch=v9 -i userCode.o $GEMSTONE/lib/gcirtlobj.o
-z nodefs -Bdynamic -lc -lpthread -ldl -lrt -lsocket -lnsl
-lm -lpam -lCrun -o userAppl
```

Solaris (x86):

```
CC -xildoff -m64 -xarch=generic -i userCode.o
$GEMSTONE/lib/gcirtlobj.o -z nodefs -Bdynamic -lc -lpthread
-ldl -lrt -lsocket -lnsl -lm -lpam -lCrun -o userAppl
```

AIX (IBM):

```
xlC_r -Wl,-bdatapsize:64K -Wl,-btextpsize:64K
-Wl,-bstackpsize:64K -q64 userCode.o $GEMSTONE/lib/gcirtlobj.o
-Wl,-berok -L/usr/vacpp/lib -lpthreads -lc_r -lC_r -lm -ldl
-lbsd -lpam -Wl,-brtllib -o userAppl
```

Darwin:

```
g++ userCode.o $GEMSTONE/lib/gcirtlobj.o -undefined dynamic_lookup
-m64 -lpthread -ldl -lc -lm -lpam -o userAppl
```

Windows:

```
link /LIBPATH:"VisualStudioInstallPath\VC\lib\amd64"
/LIBPATH:"C:\Program Files (x86)\Microsoft SDKs\Windows\v7.0A\Lib\x64"
/OPT:REF /INCREMENTAL:NO /MAP /nologo /MANIFEST
/MANIFESTFILE:userAppl.exe.manifest
/MANIFESTUAC:"level='asInvoker'" userCode.obj
%GEMSTONE%\lib\gcirpc.lib ws2_32.lib netapi32.lib advapi32.lib
comdlg32.lib user32.lib gdi32.lib kernel32.lib winspool.lib
/out:userAppl.exe
```

GemBuilder for C Files and Data Structures

This chapter describes the GemBuilder for C include files, data structures, and other reference information that may be useful when writing your application.

5.1 GemBuilder for C Include Files

The following include files are provided for use with GemBuilder for C functions. These files are in the `$GEMSTONE/include` directory.

Your C source code should include *exactly* one of these include files:

<code>gciua.hf</code>	Used instead of <code>gcirtl.hf</code> in modules that define user actions.
<code>gci.hf</code>	Forward references to the GemBuilder functions; indirectly included by <code>gcirtl.hf</code> and <code>gciua.hf</code> . Used for a C application that will call <code>GciInit</code> and <code>GciLogin</code> , on platforms that allow shared libraries to be built containing references to unresolved symbols (which are defined in <code>gci.hf</code> and resolved at run time).

If you want to directly call functions to load the `libgci*.so` files, rather than rely on implicit loading, you can include the following:

<code>gcirtl.hf</code>	Forward references to the GemBuilder functions, to be included in code that will bind to GemBuilder at run time. For a discussion of how to load a library that was compiled against <code>gcirtl.hf</code> , see “GemBuilder Applications that load Shared Libraries” on page 44. For modules that define user actions, use <code>gciua.hf</code> instead of this file.
------------------------	--

In addition, your code can include these files:

<code>gcifloat.hf</code>	Macros, constants and functions for accessing the bodies of instances of GemStone classes <code>Float</code> and <code>SmallFloat</code> . Optional for code that includes <code>gci.hf</code> and <code>gciua.hf</code> , not used with <code>gcirtl.hf</code> .
--------------------------	---

<code>gcisend.hf</code>	Inline implementations of deprecated <code>GciSendMsg</code> , which no longer uses variable arguments. New code should use <code>GciPerform</code> . For convenience, if you are using <code>GciSendMsg</code> in your GemBuilder application, include this file after the include of <code>gcirtl.hf</code> .
You do not include the following files explicitly; they are listed here for your information.	
<code>flag.ht</code>	Contains host-specific C definitions for compilation.
<code>gci.ht</code>	Defines C types for use by GemBuilder functions. See “GemBuilder Data Types” on page 70.
<code>gcicmn.ht</code>	Defines common C types and macros used by <code>gcirtl.hf</code> , <code>gci.hf</code> , and <code>gciua.hf</code> .
<code>gcierr.ht</code>	Defines mnemonics for all GemStone errors.
<code>gcilegacy.ht</code>	Defines types used to <code>gci.hf</code> that are not used by the thread-safe GCI.
<code>gcioct.ht</code>	Defines C mnemonics for sizes and offsets into objects.
<code>gcioop.ht</code>	Defines C mnemonics for predefined GemStone objects. See Appendix A, “Reserved OOPs” for a list of constants defined in this file.
<code>gciuser.hf</code>	Defines a macro to be used to install user actions. Include <code>gciua.hf</code> instead of this file.
<code>version.ht</code>	Defines C mnemonics for version-dependent strings.

5.2 GemBuilder Data Types

The following C types are used by GemBuilder functions. The file `gci.ht` defines each of the GemBuilder types (shown in capital letters below). That file is in the `$GEMSTONE/include` directory.

<code>BoolType</code>	An int.
<code>ByteType</code>	An unsigned 8-bit integer.
<code>OopType</code>	Object-oriented pointer, an unsigned 32-bit integer.
<code>FloatKindEType</code>	Enumerated type that defines the possible kinds of an IEEE binary float.
<code>GciClampedTravArgsSType</code>	A C++ class for clamped traversal arguments.
<code>GciDateTimeSType</code>	A structure for representing GemStone dates and times.
<code>GciDbgFuncType</code>	The type of C function called by <code>GciDbgEstablish</code> .
<code>GciErrSType</code>	A GemStone error report (see “The Error Report Structure” on page 72).

GciJumpBufSType	Jump buffer, defined in the setjmp.h file.
GciObjInfoSType	A C++ class for a GemStone object information report (see “The Object Information Structure” on page 72).
GciObjRepHdrSType	A C++ class for an object report header (see “The Object Report Header Class” on page 75).
GciObjRepSType	A C++ class for an object report (see “The Object Report Structure” on page 74).
GciSessionIdType	A signed 32-bit integer.
GciStoreTravDoArgsSType	A C++ class for store traversal arguments.
GciTravBufType	A traversal buffer. See “The Traversal Buffer Type” on page 79.
GciUserActionSType	A structure for describing a user action (see “The User Action Information Structure” on page 78).

The Structure for Representing the Date and Time

GemBuilder includes some functions to facilitate access to objects of type `DateTime`. (These functions also make use of the C representation for time, `time_t`.)

The structured type `GciDateTimeSType`, which provides a C representation of an instance of class `DateTime`, contains the following fields:

```
#if !defined(GCICMN_HT)

typedef struct {
    int year;
    int dayOfYear;
    int milliseconds;
    OopType timeZone;
} GciDateTimeSType;

#endif
```

The year value must be less than 1,000,000.

In addition, a C mnemonic supports representation of `DateTime` objects.

```
#define GCI_SECONDS_PER_DAY      86400
/* conversion constant */
```

NOTE:

The OOP of the Smalltalk `DateTime` class is `OOP_CLASS_DATE_TIME`.

The Error Report Structure

An error report is a C structured type named `GciErrSType`. This structure contains the following fields:

<code>OopType</code>	<i>category</i> Deprecated. The value is always <code>OOP_GEMSTONE_ERROR_CAT</code> .
<code>OopType</code>	<i>context</i> The OOP of a <code>GsProcess</code> that provides the state of the virtual machine for use in debugging. This <code>GsProcess</code> can be used as the argument to <code>GciContinue</code> or <code>GciClearStack</code> . If the virtual machine was not running, then <i>context</i> is <code>OOP_NIL</code> . If you are not interested in debugging or in continuing from an error, your program can ignore this value.
<code>OopType</code>	<i>exceptionObj</i> Either an instance of <code>Exception</code> or <code>nil</code> (if the error was not signaled from Smalltalk execution).
<code>OopType</code>	<i>args</i> [<code>GCI_MAX_ERR_ARGS</code>] An optional array of error arguments. In this release, <code>GCI_MAX_ERR_ARGS</code> is defined to be 10.
<code>int</code>	<i>number</i> The GemStone error number (a positive integer).
<code>int</code>	<i>argCount</i> The number of arguments in the <i>args</i> array.
<code>unsigned char</code>	<i>fatal</i> Nonzero if this error is fatal.
<code>char</code>	<i>message</i> [<code>GCI_ERR_STR_SIZE + 1</code>] The null-terminated string which contains the text of the error message. In this release, <code>GCI_ERR_STR_SIZE</code> is defined to be 300.

The arguments (*args*) are specific to the error encountered. In the case of a compiler error, this is a single argument – the OOP of an array of error identifiers. Each identifier is an Array with three elements: (1) the error number (a `SmallInteger`); (2) the offset into the source string at which the error occurred (also a `SmallInteger`); and (3) the text of the error message (a `String`). See the `gcierr.ht` file for a full list of errors and their arguments.

In the case of a fatal error, *fatal* is set to nonzero (`TRUE`). Your connection to GemStone is lost, and the current session ID (from `GciGetSessionId`) is reset to `GCI_INVALID_SESSION_ID`.

The Object Information Structure

Object information is placed in a C++ class named `GciObjInfoSType`. Object information access functions provide information about objects in the database. These functions offer C-style access to much information that would otherwise be available only through calls to GemStone. For more information about the `GciObjInfoSType` structured type, refer to `GciFetchObjImpl` (page 180).

<code>OopType</code>	<i>objId</i> OOP of the object.
----------------------	------------------------------------

OopType	<i>objClass</i> Class of the object; see GciFetchClass (page 165).
int64	<i>objSize</i> Object's total size in bytes or OOPs; see GciFetchSize_ (page 191).
int	<i>namedSize</i> Number of named instance variables in the object.
unsigned short	<i>objectSecurityPolicyId</i> The ID of the object's security policy.

Functions

The object information class GciObjInfoSType provides the following functions:

enum {	<pre> implem_mask = 0x03, indexable_mask = 0x04, invariant_mask = 0x08, partial_mask = 0x10, overlay_mask = 0x20 }; </pre> Defines bits to use in evaluating whether this instance is invariable, indexable, partial, or overlaid.
unsigned char isInvariant();	Returns non-zero if this object is invariant. Returns zero otherwise.
unsigned char isIndexable();	Returns non-zero if this object is indexable. Returns zero otherwise.
unsigned char isPartial();	Returns non-zero if the value buffer does not contain the entire object; that is, the operation truncated the object's instance variables. Returns zero otherwise.
unsigned char isOverlaid();	Returns non-zero if overlay semantics were used on this operation. Returns zero otherwise. When the traversal is overlaid, you can use OOP_ILLEGAL to mask out instance variables that you don't want to modify, and then store into the remaining instance variables.
unsigned char objImpl();	Returns an unsigned char in the range 0..3 that corresponds to the object's implementation format. See the description on page 74.
void clearBits();	Sets the invariant, indexable, partial, and overlaid bits to FALSE.
void setBits(unsigned char <i>bits</i>);	Sets the invariant, indexable, partial, and overlaid bits.
void setObjImpl(unsigned char <i>impl</i>);	Defines the object's implementation format. The argument must be an integer in the range 0..3 corresponding to the implementation format. See the description on page 74.

```
void setInvariant(unsigned char val);
    If val is non-zero, make this object invariant.

void setIndexable(unsigned char val);
    If val is non-zero, make this object indexable.

void setPartial(unsigned char val);
    This function has no practical effect.

void setOverlaid(unsigned char val);
    If val is non-zero, use overlay semantics on this store traversal.
```

Description

The `gcioc.ht` include file defines four mnemonics that can be of assistance when you are handling the object implementation field: `GC_FORMAT_OOP`, `GC_FORMAT_BYTE`, `GC_FORMAT_NSC`, and `GC_FORMAT_SPECIAL`. These mnemonics, and no other values, should be used to supply values for the `objImpl` field, or to test its contents.

The Object Report Structure

Each object report has two parts: a fixed-size header (as defined in the C++ class `GciObjRepHdrSType`) and a variable-size value buffer (an array of the values of the object's instance variables):

```
#if !defined(GCI_HT)
class GciObjRepSType { /* object report struct */
public:
    GciObjRepHdrSType hdr;          /* object report header */
    union {
        ByteType      bytes[1];    /* Byte obj impl. obj's value buff */
        OopType       oops[1];    /* Pointer obj impl. obj's value buff*/
    } u;

    inline int64 usedBytes() const {
        return this->hdr.usedBytes();
    }

    inline GciObjRepSType* nextReport() const {
        return (GciObjRepSType*) this->hdr.nextReport();
    }

    inline ByteType* valueBufferBytes() const {
        return (ByteType*)this->u.bytes;
    }

    inline OopType* valueBufferOops() const {
        return (OopType*)this->u.oops;
    }
};
#endif
```

Functions

The object report class `GciObjRepSType` provides these functions:

`int64 usedBytes()` ;

When constructing an object report buffer, returns the size of the object report, including any alignment considerations.

`GciObjRepHdrSType * nextReport()` ;

Given a pointer to an object report in a traversal buffer, this function increments the pointer by *usedBytes* (the size of the object report).

`ByteType* valueBufferBytes()` ;

Returns a pointer to the start of the body, as bytes.

`OopType* valueBufferOops()`

Returns a pointer to the start of the body, as OOPs.

The Object Report Header Class

An object report header is a C++ class named `GciObjRepHdrSType`. This class holds a general description of an object, and contains the following fields:

<code>int</code>	<i>valueBuffSize</i>	Size (in bytes) of the object's value buffer.
<code>short</code>	<i>namedSize</i>	Number of named instance variables in the object.
<code>unsigned short</code>	<i>objectSecurityPolicyId</i>	The ID of the object's security policy.
<code>OopType</code>	<i>objId</i>	OOP of the object.
<code>OopType</code>	<i>oclass</i>	Class of the object; see <code>GciFetchClass</code> (page 165).
<code>int64</code>	<i>firstOffset</i>	Offset of first value to fetch or store.

Functions

The object report header class `GciObjRepHdrSType` provides the following functions:

```
enum {
    implem_mask = 0x03,
    indexable_mask = 0x04,
    invariant_mask = 0x08,
    partial_mask = 0x10,
    overlay_mask = 0x20,
    no_read_auth_mask = 0x40,
    clamped_mask = 0x80,
    unused_mask = 0xFF00,
    all_bits_mask = 0xFFFF
};
```

Defines bits to use in evaluating this instance's implementation format, and whether this instance is indexable, invariable, partial, overlaid, readable, or clamped.

- `int64 idxSize();`
Returns the number of indexable or varying instance variables.
- `void setIdxSize(int64 size);`
Sets the number of indexable or varying instance variables.
- `void setIdxSizeBits(int64 size, unsigned char bits);`
Sets both the indexable size and the eight bits defined by the enum of the mask values. Intended for GemStone use only.
- `int objImpl();`
Returns an integer in the range 0..3 that corresponds to the object's implementation format. See the description on page 77.
- `int setObjImpl(int impl);`
Defines the object's implementation format. The argument must be an integer in the range 0..3 corresponding to the implementation format. See the description on page 77.
- `int64 objSize();`
Returns the total number of instance variables in the object (both indexable and named). See **GciFetchSize_** (page 191).
- `void clearBits();`
Sets indexable, invariable, partial, overlaid, non-readable, and clamped to FALSE.
- `unsigned char isClamped();`
Returns non-zero if this object report is clamped. Returns zero otherwise.
- `unsigned char noReadAuthorization();`
Returns non-zero if this object report is not readable. Returns zero otherwise.
- `unsigned char isInvariant();`
Returns non-zero if this object report is invariant. Returns zero otherwise.
- `unsigned char isIndexable();`
Returns non-zero if this object report is indexable. Returns zero otherwise.
- `unsigned char isPartial();`
Returns non-zero if the value buffer does not contain the entire object; that is, the traversal operation truncated the object's instance variables. Returns zero otherwise.
- `unsigned char isOverlaid();`
Returns non-zero if overlay semantics were used on this store traversal operation. Returns zero otherwise.
- When the traversal is overlaid, you can use `OOP_ILLEGAL` to mask out instance variables that you don't want to modify, and then store into the remaining instance variables.
- `void setIsClamped(unsigned char val);`
If *val* is non-zero, make this object report clamped.

```
void setNoReadAuth(unsigned char val);
    If val is non-zero, make this object report non-readable.

void setInvariant(unsigned char val);
    If val is non-zero, make this object report invariant.

void setIndexable(unsigned char val);
    If val is non-zero, make this object report indexable.

void setPartial(unsigned char val);
    This function has no practical effect.

void clearPartial(unsigned char val);
    This function has no practical effect.

void setOverlaid(unsigned char val);
    If val is non-zero, use overlay semantics on this store traversal.

ByteType* valueBufferBytes() ;
    Returns a pointer to the start of the body, as bytes.

OopType* valueBufferOops()
    Returns a pointer to the start of the body, as OOPs.

int64 usedBytes() ;
    Returns the size (in bytes) of this object report, including the size of the
    header, value buffer, and any padding bytes needed at the end of the
    report so that the next report in the buffer begins on an address that is
    a multiple of 8.

GciObjRepHdrSType * nextReport() ;
    Given a pointer to an object report in a traversal buffer, this function
    increments the pointer by usedBytes (the size of the object report).
```

Description

During a store traversal operation, if the specified *idxSize* is inadequate to accommodate the contents of the value buffer (the values in *u.bytes* or *u.oops*), the store operation will automatically increase *idxSize* (the number of the object's indexed variables) as needed. If the specified *objClass* is not indexable, then the *idxSize* is ignored; in addition, if there are more OOPs in the value buffer than there are named instance variables, and the object is not an NSC, an error will be generated.

During a store traversal operation, the *firstOffset* indicates where to begin storing values into the object's array of instance variables. In that array, the object's named instance variables are followed by its unnamed variables. If *firstOffset* is not 1, all instance variables (named or indexed) up to the *firstOffset* will be initialized to nil or 0. The *firstOffset* must be in the range (1, *objSize*+1).

The `gcio.c.ht` include file defines four mnemonics that can be of assistance when you are handling the object implementation field (*objImpl*): `GC_FORMAT_OOP`, `GC_FORMAT_BYTE`, `GC_FORMAT_NSC`, and `GC_FORMAT_SPECIAL`. These mnemonics, and no other values, should be used to supply values for *objImpl*, or to test its contents. However, the `gcio.c.ht` file also defines other mnemonics that can be used in other contexts related to object implementations, indexability, and invariance.

An object's implementation may restrict the number of its named instance variables (*namedSize*) and its indexed instance variables (*idxSize*), as contained in the object report header.

- ▶ If the object implementation is `GC_FORMAT_OOP`, the object can have both named and unnamed instance variables.
- ▶ If the object implementation is `GC_FORMAT_BYTE`, the object can only have indexed instance variables, and its *namedSize* is always zero.
- ▶ If the object implementation is `GC_FORMAT_NSC`, the object can have both named and unnamed instance variables. (The NSC's *idxSize* reports the number of unnamed instance variables, even though they are unordered, not indexed.)
- ▶ If the object implementation is `GC_FORMAT_SPECIAL`, the object cannot have any instance variables, and the number of both its named and unnamed variables is always zero.

The `isInvariant()` value is true if the object itself is invariant. This can happen in one of three ways:

- ▶ The application program sends the message `immediateInvariant` to the object.
- ▶ The application program explicitly executes `setInvariant()` in the report header and then uses that report header in a call to `GciStoreTrav`.
- ▶ The object's class was created with `instancesInvariant: true` and the object has been committed.

Table 5.1 Object Implementation Restrictions on Instance Variables

Object Implementation	Named Instance Variables OK?	Unnamed Instance Variables OK?
0=Pointer	YES	YES (always indexed)
1=Byte	NO	YES (always indexed)
2=NSC	YES	YES (always unordered)
3=Special	NO	NO

For more information about object implementation types, see "Manipulating Objects Through Structural Access" on page 28.

The User Action Information Structure

The structured type `GciUserActionSType` describes a user action function. It defines the following fields:

```
char          userActionName[GCI_MAX_ACTION_NAME+1]
              The user action name (a case-insensitive, null-terminated string). In
              this release, GCI_MAX_ACTION_NAME is defined to be 31.

int          userActionNumArgs
              The number of arguments in the C function.
```

GciUserActionFType	<i>userAction</i> A pointer to the C user action function.
unsigned int	<i>userActionFlags</i> Mainly for internal use. If you use it, set it to 0 before passing a pointer to it.

The Traversal Buffer Type

The C++ class **GciTravBufType** describes a traversal buffer, and defines the following fields:

int64	<i>allocatedBytes</i> The allocated size of the body.
int64	<i>usedBytes</i> The used bytes of the body.
ByteType	<i>body[8]</i> The actual body size is variable, with a minimum of GCI_MIN_TRAV_BUFF_SIZE.

Functions

The following function call is used to create an instance of GciTravBufType:

```
static GciTravBufType* malloc(size_t allocationSize);
```

Returns an instance obtained from ::malloc initialized with *allocatedBytes* equal to *allocationSize* and *usedBytes*== 0. (If *allocationSize* is not a multiple of 8 bytes, *allocatedBytes* is rounded up to the next 8-byte multiple.) Returns NULL if malloc fails.

The traversal buffer class GciTravBufType provides these functions:

```
GciObjRepSType* firstReport() ;
```

Returns a pointer to the first object report in the buffer.

```
GciObjRepSType* readLimit() ;
```

Used when reading object reports out of a buffer. Returns a pointer past the end of last object report in the buffer.

If `readLimit() == firstReport()`, the buffer is empty.

```
GciObjRepSType* writeLimit() ;
```

Used when composing a buffer. Returns a pointer one byte past the end of the allocated buffer.

```
GciObjRepHdrSType* firstReportHdr() ;
```

Returns a pointer to the first object report in the buffer.

```
GciObjRepHdrSType* readLimitHdr() ;
```

Used when reading object reports out of a buffer. Returns a pointer past the end of last object report in the buffer.

```
GciObjRepHdrSType* writeLimitHdr() ;
```

Used when composing a buffer. Returns a pointer one byte past the end of the allocated buffer.

5.3 Structural Access Functions

A number of functions access Smalltalk objects structurally, rather than via executing message sends. A list of these functions is in Table 6.8 on page 88.

Exercise caution when using structural access functions. Although they can improve the speed of GemStone database operations, these functions bypass GemStone's message-sending metaphor. That is, structural access functions may bypass any checking that might be coded into your application's methods.

Structural access functions do not bypass authorization checks or other checks that are not done in Smalltalk code.

5.4 environmentId

Many GCI functions come in two variants; the standard one, and one with the same name followed by a trailing underscore, that takes an additional environmentId argument.

The environmentId argument allows a GCI function to specify one of up to 256 execution environments, for use in Ruby applications.

Smalltalk applications do not need to know anything about environmentId. With Smalltalk applications, it is preferable to use the existing GCI function (without the trailing underscore).

For an example of this, see **GciExecute** (page 149). The syntax section on that page shows both variants: **GciExecute** (used with Smalltalk applications) and **GciExecute_** (used with Ruby applications).

5.5 UNIX Signal Handling

Both versions of GemBuilder (GciLnk and GciRpc) use the SIGIO signal handler. GciLnk also uses the signals SIGSEGV and SIGVTALRM. SIGVTALRM is used by the ProfMonitor class.

If you must install your own signal handler (using signal or sigvec) for any of these signals, be sure that your application signal handler chains to the previous handler when done. Similar chaining is required for SIGVTALRM, if you intend to use ProfMonitor.

SIGSEGV occurs normally when a Smalltalk stack overflow occurs, and is translated to a Smalltalk stack overflow error by the GemStone SIGSEGV handler. If you use GciLnk and install handlers for this signal after calling GciLogin, your own SIGSEGV handler must determine whether the SIGSEGV was produced by your own C code, and if not, chain to the GemStone handler.

CAUTION

Do not, under any circumstances, turn off SIGIO.

GemBuilder C Functions

This chapter describes the GemBuilder functions that may be called by your C application program.

6.1 Function Summary Tables

Tables 6.1 through 6.9 summarize the GemBuilder C functions and the services that they provide to your application.

Table 6.1 Functions for Controlling Sessions and Transactions

GciAbort, GciNbAbort	Abort the current transaction.
GciAlteredObjs	Find all exported or dirty objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.
GciBegin, GciNbBegin	Begin a new transaction.
GciCommit, GciNbCommit	Write the current transaction to the database.
GciDeclareAction	An alternative way to associate a C function with a Smalltalk user action.
GciDirtyExportedObjs	Find all objects in the ExportedDirtyObjs set.
GciDirtyObjsInit	Begin tracking which objects in the session workspace change.
GciDirtySaveObjs	Find all exported or tracked objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.
GciDirtyTrackedObjs	Find all tracked objects that have changed and are therefore in the TrackedDirtyObjs set.
GciGetSessionId	Find the ID number of the current user session.
GciHardBreak	Interrupt GemStone and abort the current transaction.
GciInit	Initialize GemBuilder.

Table 6.1 Functions for Controlling Sessions and Transactions (Continued)

GciInitAppName	Override the default application configuration file name.
GciInstallUserAction	Associate a C function with a Smalltalk user action.
GciIsRemote	Determine whether the application is running linked or remotely.
GciLoadUserActionLibrary	Load an application user action library.
GciLogin, GciLoginEx	Start a user session.
GciLogout	End the current user session.
GciNbEnd, GciNbEndPoll	Test the status of nonblocking call in progress for completion.
GciProcessDeferredUpdates_	Process deferred updates to objects that do not allow direct structural update.
GciReleaseAllGlobalOops	Remove all OOPS from the PureExportSet, making these objects eligible for garbage collection.
GciReleaseAllOops	Remove all OOPS from the PureExportSet, or if in a user action, from the user action's export set, making these objects eligible for garbage collection.
GciReleaseAllTrackedOops	Clear the GciTrackedObjs set, making all tracked OOPs eligible for garbage collection.
GciReleaseGlobalOops	Remove an array of GemStone OOPs from the PureExportSet, making them eligible for garbage collection.
GciReleaseOops	Remove an array of GemStone OOPs from the PureExportSet, or if in a user action, remove them from the user action's export set, making them eligible for garbage collection.
GciReleaseTrackedOops	Remove an array of OOPs from the GciTrackedObjs set, making them eligible for garbage collection.
GciRtlIsLoaded	Report whether a GemBuilder library is loaded.
GciRtlLoad	Load a GemBuilder library.
GciRtlUnload	Unload a GemBuilder library.
GciSaveAndTrackObjs	Add objects to GemStone's internal GciTrackedObjs set to prevent them from being garbage collected.
GciSaveGlobalObjs	Add an array of OOPs to the PureExportSet, making them ineligible for garbage collection.
GciSaveObjs	Add an array of OOPs to the PureExportSet, or if in a user action to the user action's export set, making them ineligible for garbage collection.
GciServerIsBigEndian	Determine whether or not the server process is big-endian.
GciSessionIsRemote	Determine whether or not the current session is using a Gem on another machine.
GciSetCacheName_	Set the name that a linked application will be known by in the shared cache.
GciSetNet, GciSetNetEx	Set network parameters for connecting the user to the Gem and Stone processes.

Table 6.1 Functions for Controlling Sessions and Transactions (Continued)

GciSetSessionId	Set an active session to be the current one.
GciShutdown	Logout from all sessions and deactivate GemBuilder.
GciStep	Continue code execution in GemStone with specified single-step semantics.
GciTrackedObjsFetchAllDirty	Find all exported or tracked objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.
GciTrackedObjsInit	Reinitialize the set of tracked objects maintained by GemStone.
GciUserActionInit	Declare user actions for GemStone.
GciUserActionShutdown	Enable user-defined clean-up for user actions.

Table 6.2 Functions for Handling Errors and Interrupts and for Debugging

GciCallInProgress	Determine if a GemBuilder call is currently in progress.
GciClearStack	Clear the Smalltalk call stack.
GciContinue, GciNbContinue	Continue code execution in GemStone after an error.
GciContinueWith GciNbContinueWith	Continue code execution in GemStone after an error.
GciDbgEstablish	Specify the debugging function for GemBuilder to execute before most calls to GemBuilder functions.
GciDbgEstablishToFile	Write trace information for most GemBuilder functions to a file.
GciDbgLogString	Pass a message to a trace function.
GciEnableSignaledErrors	Establish or remove GemBuilder visibility to signaled errors from GemStone.
GciErr	Prepare a report describing the most recent GemBuilder error.
GciInUserAction	Determine whether or not the current process is executing a user action.
GciLongJump	Provides equivalent functionality to the corresponding longjmp() or _longjmp() function.
GciPollForSignal	Poll GemStone for signal errors without executing any Smalltalk methods.
GciPopErrJump	Discard a previously saved error jump buffer.
GciPushErrJump	Associate GemBuilder error handling with a jump buffer by pushing a jump buffer onto the stack.
GciRaiseException	Signal an error, synchronously, within a user action.
GciSetErrJump	Enable or disable the current error handler.
GciSetHaltOnError	Halt the current session when a specified error occurs.
Gci_SETJMP	(MACRO) Save a jump buffer in GemBuilder's error jump stack.
GciSoftBreak	Interrupt the execution of Smalltalk code, but permit it to be restarted.

Table 6.3 Functions for Compiling and Executing Smalltalk Code in the Database

GciClassMethodForClass	Compile a class method for a class.
GciCompileMethod	Compile a method.
GciExecute, GciNbExecute	Execute a Smalltalk expression contained in a String or Utf8 object.
GciExecuteFromContext	Execute a Smalltalk expression contained in a String or Utf8 object as if it were a message sent to another object.
GciExecuteStr, GciNbExecuteStr	Execute a Smalltalk expression contained in a C string.
GciExecuteStrFetchBytes, GciNbExecuteStrFetchBytes	Execute a Smalltalk expression contained in a C string, returning byte-format results.
GciExecuteStrFromContext, GciNbExecuteStrFromContext	Execute a Smalltalk expression contained in a C string as if it were a message sent to an object.
GciInstMethodForClass	Compile an instance method for a class.
GciPerform, GciNbPerform	Send a message to a GemStone object.
GciPerformNoDebug, GciNbPerformNoDebug	Send a message to a GemStone object, and temporarily disable debugging.
GciPerformSymDbg	Send a message to a GemStone object, using a String object as a selector.
GciPerformTraverse	First send a message to a GemStone object, then traverse the result of the message.

Table 6.4 Functions for Accessing Symbol Dictionaries

GciResolveSymbol	Find the OOP of the object to which a symbol name refers, in the context of the current session's user profile.
GciResolveSymbolObj	Find the OOP of the Symbol to which a String object refers, in the context of the current session's user profile.
GciStrKeyValueDictAt	Find the value in a symbol KeyValue dictionary at the corresponding string key.
GciStrKeyValueDictAtObj	Find the value in a symbol KeyValue dictionary at the corresponding object key.
GciStrKeyValueDictAtObj Put	Store a value into a symbol KeyValue dictionary at the corresponding object key.
GciStrKeyValueDictAtPut	Store a value into a symbol KeyValue dictionary at the corresponding string key.
GciSymDictAt	Find the value in a symbol dictionary at the corresponding string key.
GciSymDictAtObj	Find the value in a symbol dictionary corresponding to the key object.
GciSymDictAtObjPut	Store a value into a symbol dictionary at the corresponding object key.

Table 6.4 Functions for Accessing Symbol Dictionaries (Continued)

GciSymDictAtPut	Store a value into a symbol dictionary at the corresponding string key.
GciTraverseObjs	Traverse an array of GemStone objects.

Table 6.5 Functions for Creating and Initializing Objects

GciCreateByteObj	Create a new byte-format object.
GciCreateOopObj	Create a new pointer-format object.
GciGetFreeOop	Allocate an OOP.
GciGetFreeOops	Allocate multiple OOPs.
GciGetFreeOopsEncoded	Allocate multiple OOPs.
GciNewByteObj	Create and initialize a new byte object.
GciNewCharObj	Create and initialize a new character object.
GciNewDateTime	Create and initialize a new date-time object.
GciNewOop	Create a new GemStone object.
GciNewOops	Create multiple new GemStone objects.
GciNewOopUsingObjRep	Create a new GemStone object from an existing object report.
GciNewString	Create a new String object from a C character string.
GciNewSymbol	Create a new Symbol object from a C character string.
GciNewUtf8String	Create a new Unicode string object from a UTF-8 encoded C character string.

Table 6.6 Functions and Macros for Converting Objects and Values

GCI_BOOL_TO_OOP	(MACRO) Convert a C Boolean value to a GemStone Boolean object.
GciByteArrayToPointer	Given a result from GciPointerToByteArray, return a C pointer.
GCI_CHR_TO_OOP	(MACRO) Convert a C character value to a GemStone Character object.
GciCTimeToDateTime	Convert a C date-time representation to the equivalent GemStone representation.
GciDateTimeToCTime	Convert a GemStone date-time representation to the equivalent C representation.
Gci_doubleToSmallDouble	Convert a C double to a SmallDouble object.
GciFetchDateTime	Convert the contents of a DateTime object and place the results in a C structure.
GciFetchUtf8Bytes_	Encode a String, MultiByteString, or Utf8 as UTF-8, and fetch the bytes of the encoded result.
GciFloatKind	Obtain the float kind corresponding to a C double value.
GciFltToOop	Convert a C double value to a SmallDouble or Float object.

Table 6.6 Functions and Macros for Converting Objects and Values (Continued)

GCI_I64_IS_SMALL_INT	Determine whether or not a C 64-bit integer value can be translated into a SmallInteger object.
GciI64ToOop	Convert a C 64-bit integer value to a GemStone object.
GCI_OOP_IS_BOOL	(MACRO) Determine whether or not a GemStone object represents a Boolean value.
GCI_OOP_IS_SMALL_INT	(MACRO) Determine whether or not a GemStone object represents a SmallInteger.
GCI_OOP_IS_SPECIAL	(MACRO) Determine whether or not a GemStone object has a special representation.
GciOopToBool	Convert a Boolean object to a C Boolean value.
GCI_OOP_TO_BOOL	(MACRO) Convert a Boolean object to a C Boolean value.
GciOopToChar16	Convert a Character object to a 16-bit C character value.
GciOopToChar32	Convert a Character object to a 32-bit C character value.
GciOopToChr	Convert a Character object to a C character value.
GCI_OOP_TO_CHR	(MACRO) Convert a Character object to a C character value.
GciOopToFlt	Convert a SmallDouble, Float, or SmallFloat object to a C double.
GciOopToI32	Convert a GemStone object to a C 32-bit integer value.
GciOopToI64	Convert a GemStone object to a C 64-bit integer value.
GciPointerToByteArray	Given a C pointer, return a SmallInteger or ByteArray containing the value of the pointer.
GciStringToInteger	Convert a C string to a GemStone SmallInteger or LargeInteger object.

Table 6.7 Object Traversal and Path Functions and Macros

GCI_ALIGN	(MACRO) Align an address to a word boundary.
GciClampedTrav, GciNbClampedTrav	Traverse an array of objects, subject to clamps.
GciExecuteStrTrav, GciNbExecuteStrTrav	Execute a string and traverse the result of the execution.
GciFetchPaths	Fetch selected multiple OOPs from an object tree.
GciFindObjRep	Fetch an object report in a traversal buffer.
GciMoreTraversal, GciNbMoreTraversal	Continue object traversal, reusing a given buffer.
GciObjRepSize_	Find the number of bytes in an object report.
GciPerformTrav GciNbPerformTrav	First send a message to a GemStone object, then traverse the result of the message.
GciPerformTraverse	First send a message to a GemStone object, then traverse the result of the message.
GciSetTraversalBufSwizzling	Control swizzling of the traversal buffers.
GciStorePaths	Store selected multiple OOPs into an object tree.
GciStoreTrav GciNbStoreTrav	Store multiple traversal buffer values in objects.
GciStoreTravDo_ GciNbStoreTravDo_	Store multiple traversal buffer values in objects, execute the specified code, and return the resulting object.
GciStoreTravDoTrav_ GciNbStoreTravDoTrav_	Combine in a single function the calls to GciStoreTravDo_ and GciClampedTrav, to store multiple traversal buffer values in objects, execute the specified code, and traverse the result object.
GciStoreTravDoTravRefs_ GciNbStoreTravDoTravRefs_	Combine in a single function modifications to session sets, traversal of objects to the server, optional Smalltalk execution, and traversal to the client of changed objects and (optionally) the result object.
GciTraverseObjs, GciNbTraverseObjs	Traverse an array of GemStone objects.

CAUTION

Exercise caution when using the following structural access functions. Although they can improve the speed of GemStone database operations, these functions bypass GemStone's message-sending metaphor. That is, structural access functions may bypass any checking that might be coded into your application's methods. In using structural access functions, you implicitly assume full responsibility for safeguarding the integrity of your system.

Note, however, that structural access functions do not bypass checks on authorization violations or concurrency conflicts.

Table 6.8 Structural Access Functions and Macros

GciAddOopToNsc	Add an OOP to the unordered variables of a nonsequenceable collection.
GciAddOopsToNsc	Add multiple OOPs to the unordered variables of a nonsequenceable collection.
GciAppendBytes	Append bytes to a byte object.
GciAppendChars	Append a C string to a byte object.
GciAppendOops	Append OOPs to the unnamed variables of a collection.
GciClassNamedSize	Find the number of named instance variables in a class.
GciFetchByte	Fetch one byte from an indexed byte object.
GciFetchBytes_	Fetch multiple bytes from an indexed byte object.
GciFetchChars_	Fetch multiple ASCII characters from an indexed byte object.
GciFetchClass	Fetch the class of an object.
GciFetchNamedOop	Fetch the OOP of one of an object's named instance variables.
GciFetchNamedOops	Fetch the OOPs of one or more of an object's named instance variables.
GciFetchNamedSize	Fetch the number of named instance variables in an object.
GciFetchNameOfClass	Fetch the class name object for a given class.
GciFetchObjImpl	Fetch the implementation of an object.
GciFetchOop	Fetch the OOP of one instance variable of an object.
GciFetchOops	Fetch the OOPs of one or more instance variables of an object.
GciFetchSize_	Fetch the size of an object.
GciFetchVaryingOop	Fetch the OOP of one unnamed instance variable from an indexable pointer object or NSC.
GciFetchVaryingOops	Fetch the OOPs of one or more unnamed instance variables from an indexable pointer object or NSC.
GciFetchVaryingSize_	Fetch the number of unnamed instance variables in a pointer object or NSC.
GciHiddenSetIncludesOop	Determines whether the given OOP is present in the specified hidden set.
GciIsKindOf	Determine whether or not an object is some kind of a given class or class history.
GciIsKindOfClass	Determine whether or not an object is some kind of a given class.
GciIsSubclassOf	Determine whether or not a class is a subclass of a given class or class history.
GciIsSubclassOfClass	Determine whether or not a class is a subclass of a given class.
GciIvNameToIdx	Fetch the index of an instance variable name.
GciNscIncludesOop	Determines whether the given OOP is present in the specified unordered collection.

Table 6.8 Structural Access Functions and Macros (Continued)

GciObjExists	Determine whether or not a GemStone object exists.
GciObjInCollection	Determine whether or not a GemStone object is in a Collection.
GciObjIsCommitted	Determine whether or not an object is committed.
GciRemoveOopFromNsc	Remove an OOP from an NSC.
GciRemoveOopsFromNsc	Remove one or more OOPs from an NSC.
GciReplaceOops	Replace all instance variables in a GemStone object.
GciReplaceVaryingOops	Replace all unnamed instance variables in an NSC object.
GciSetVaryingSize	Set the size of a collection.
GciStoreByte	Store one byte in a byte object.
GciStoreBytes	(MACRO) Store multiple bytes in a byte object.
GciStoreBytesInstanceOf	Store multiple bytes in a byte object.
GciStoreChars	Store multiple ASCII characters in a byte object.
GciStoreIdxOop	Store one OOP in an indexable pointer object's unnamed instance variable.
GciStoreIdxOops	Store one or more OOPs in an indexable pointer object's unnamed instance variables.
GciStoreNamedOop	Store one OOP into an object's named instance variable.
GciStoreNamedOops	Store one or more OOPs into an object's named instance variables.
GciStoreOop	Store one OOP into an object's instance variable.
GciStoreOops	Store one or more OOPs into an object's instance variables.

Table 6.9 Utility Functions

GciAll7Bit	Determine if a String contains only 7-bit ASCII characters.
GciCompress	Compress the supplied data, which can be uncompressed with GciUncompress.
GciDecodeOopArray	Decode an OOP array that was previously run-length encoded.
GciDecSharedCounter	Decrement the value of a shared counter.
GciEnableFreeOopEncoding	Enable run-length encoding of free OOPs.
GciEnableFullCompression	Enable full compression between the client and the RPC version of GemBuilder.
GciEncodeOopArray	Encode an array of OOPs, using run-length encoding.
GciEncrypt	Encrypt a password string.
GciFetchNumEncodedOops	Obtain the size of an encoded OOP array.
GciFetchNumSharedCounters	Obtain the number of shared counters available on the shared page cache used by this session.
GciFetchSharedCounterValuesNoLock	Fetch the value of multiple shared counters without locking them.
GciIncSharedCounter	Increment the value of a shared counter.

Table 6.9 Utility Functions (Continued)

GciProduct	Return an 8-bit unsigned integer that indicates the GemStone/S product.
GciReadSharedCounter	Lock and fetch the value of a shared counter.
GciReadSharedCounterNoLock	Fetch the value of a shared counter without locking it.
GciSetSharedCounter	Set the value of a shared counter.
GciUncompress	Uncompress the supplied data, assumed to have been compressed with GciCompress.
GciVersion	Return a string that describes the GemBuilder version.

GciAbort

Abort the current transaction.

Syntax

```
void GciAbort()
```

Description

This function causes the GemStone system to abort the current transaction. All changes to persistent objects that were made since the last committed transaction are lost, and the application is connected to the most recent version of the database. Your application must fetch again from GemStone any changed persistent objects, to refresh the copies of these objects in your C program. Use the **GciDirtySaveObjs** function to determine which of the fetched objects were also changed.

This function has the same effect as issuing a hard break, or the function call `GciExecuteStr("System abortTransaction", OOP_NIL)`. For more information, see "Interrupting GemStone Execution" on page 27.

See Also

[GCL_CHR_TO_OOP](#), page 107
[GciCommit](#), page 115
[GciNbAbort](#), page 236
[GciNbCommit](#), page 239

GciAddOopToNsc

Add an OOP to the unordered variables of a nonsequenceable collection.

Syntax

```
void GciAddOopToNsc(
    OopType      theNsc,
    OopType      theOop );
```

Input Arguments

<i>theNsc</i>	The OOP of the NSC.
<i>theOop</i>	The OOP to be added.

Description

This function adds an OOP to the unordered variables of an NSC, using structural access.

Example

```
OopType GciAddOopToNsc_example(void)
{
    // return an IdentityBag containing the SmallIntegers with value
    0..99

    OopType oNsc = GciNewOop(OOP_CLASS_IDENTITY_BAG);
    for (int i = 0; i < 100; i ++) {
        OopType oNum = GciI32ToOop(i);
        GciAddOopToNsc(oNsc, oNum);
    }
    return oNsc;
}
```

See Also

- GciAddOopsToNsc, page 93
- GciNscIncludesOop, page 278
- GciRemoveOopFromNsc, page 326
- GciRemoveOopsFromNsc, page 327

GciAddOopsToNsc

Add multiple OOPs to the unordered variables of a nonsequenceable collection.

Syntax

```
void GciAddOopsToNsc(  
    OopType          theNsc,  
    const OopType    theOops[],  
    int              numOops);
```

Input Arguments

<i>theNsc</i>	The OOP of the NSC.
<i>theOops</i>	An array of OOPs to be added.
<i>numOops</i>	The number of OOPs to add.

Description

This function adds multiple OOPs to the unordered variables of an NSC, using structural access.

Example

```
OopType GciAddOopsToNsc_example(void)  
{  
    // return an IdentityBag containing the SmallIntegers with value  
    0..99  
  
    enum { AddOopsToNsc_SIZE = 100 };  
  
    OopType oNsc = GciNewOop(OOP_CLASS_IDENTITY_BAG);  
  
    OopType values[AddOopsToNsc_SIZE];  
    for (int i = 0; i < AddOopsToNsc_SIZE; i++) {  
        values[i] = GciI32ToOop(i);  
    }  
    GciAddOopsToNsc(oNsc, values, AddOopsToNsc_SIZE);  
    return oNsc;  
}
```

See Also

[GciAddOopToNsc](#), page 92
[GciNscIncludesOop](#), page 278
[GciRemoveOopFromNsc](#), page 326
[GciRemoveOopsFromNsc](#), page 327

GCI_ALIGN

(MACRO) Align an address to a word boundary.

Syntax

```
uintptr_t * GCI_ALIGN(argument)
```

Input Arguments

argument The pointer or integer to be aligned.

Return Value

The first multiple of 8 that is greater than or equal to the input *argument*.

Description

This macro can be used to round up a pointer or size to be a multiple of `sizeof(OopType)`.

Provided for compatibility. New code should use the accessor functions in `GciObjRepHdrSType` (page 72).

See Also

`GciMoreTraversal`, page 234

`GciNewOopUsingObjRep`, page 272

`GciTraverseObjs`, page 408

GciAll7Bit

Determine if a String contains only 7-bit ASCII characters.

Syntax

```
BoolType GciAll7Bit(  
    const char *      aString,  
    size_t *          stringLength  
);
```

Input Arguments

aString A null-terminated string.

Result Arguments

stringLength The length of *aString*, computed by `strlen(aString)`.

Return Value

Returns TRUE if each character in *aString* is in the 7-bit ASCII range, that is, a value ≤ 127 ; FALSE if any characters have values of 128 or higher.

Description

The function **GciAll7Bit** is used to test if a String is in the ASCII range, which will not involve encoding to UTF-8 for storage.

See Also

[GciFetchUtf8Bytes_](#), page 193

GciAllocTravBuf

Allocate and initialize a new GciTravBufType structure.

Syntax

```
(GciTravBufType *) GciAllocTravBuf(  
    size_t          allocationSize);
```

Input Arguments

allocationSize The size of the traversal buffer.

Description

This function allocates and initializes a new GciTravBufType structure.

See Also

“The Traversal Buffer Type” on page 79

GciAlteredObjs

Find all exported or dirty objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.

Syntax

```
BoolType GciAlteredObjs(
    OopType      theOops[ ],
    int *        numOops );
```

Input Arguments

<i>theOops</i>	An array that will contain the of OOPs of the objects in the ExportedDirtyObjs or TrackedDirtyObjs sets.
<i>numOops</i>	Pointer to the maximum number of OOPs that can be returned in this call, that is, the size (in OOPs) of the buffer specified by <i>theOops</i> .

Result Arguments

<i>theOops</i>	The resulting array of OOPs of objects that are in either the ExportedDirtyObjs or TrackedDirtyObjs sets.
* <i>numOops</i>	The number of actual OOPs in the result array <i>theOops</i> .

Return Value

The function result indicates whether all dirty objects have been returned. If the operation is not complete, **GciAlteredObjs** returns FALSE, and it is expected that the application will make repeated calls to this function until it returns TRUE, indicating that all of the dirty objects have been returned. If repeated calls are not made, then the unreturned objects persist in the list until the next time **GciAlteredObjs**, or another call that destructively accesses the ExportedDirtyObjs or TrackedDirtyObjs sets, is called.

Description

Typically, a GemStone C application program caches some database objects in its local object space, generally in the PureExportSet or if in a user action, in the user action's export set (see **GciSaveObjs**). It may also track them by storing them in the GciTrackedObjs set (see **GciSaveAndTrackObjs**). After an abort or a successful commit, the user's session is resynchronized with the most recent version of the database. The values of instance variables cached in your C program may no longer accurately represent the corresponding GemStone objects. In such cases, your C program must update its representation of those objects. The function **GciAlteredObjs** permits you to determine which objects your application needs to reread from the database.

This function returns a list of all objects that are in the PureExportSet *and* are “dirty”. An object is considered dirty (changed) under one or more of the following conditions:

- The object was changed by Smalltalk execution from this session.
- The object was changed by a call from this session to any GemBuilder function from within a user action.
- The object was changed by a call from this session to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.
- The object was read by this session, and after this session did a commit, begin, or abort transaction, the session now has visibility to changes to the object committed by another session.
- The object is persistent, and this session aborted its changes to the object, thus rolling back the Smalltalk in-memory state to the previously committed state.

Calls to **GciStore...** (other than **GciStorePaths**), **GciAppend...**, **GciReplace...**, and **GciCreate...** do not put the modified object into the set of dirty objects (unless the call is from within a user action). The assumption is that the client does not want the dirty set to include modifications that the client has explicitly made.

You must call **GciDirtyObjsInit** once after **GciLogin** before you can use **GciAlteredObjs**.

Note that **GciAlteredObjs** removes OOPs from the ExportedDirtyObjs set and TrackedDirtyObjs sets as it enumerates.

See Also

GciAbort, page 91

GciCommit, page 115

GciDirtyObjsInit, page 137

GciReleaseAllOops, page 320

GciReleaseOops, page 323

GciSaveAndTrackObjs, page 336

GciSaveGlobalObjs, page 337

GciSaveObjs, page 338

GciAppendBytes

Append bytes to a byte object.

Syntax

```
void GciAppendBytes(  
    OopType          theObject,  
    int64            numBytes,  
    const ByteType * theBytes );
```

Input Arguments

<i>theObject</i>	A byte object to which bytes are to be appended.
<i>numBytes</i>	The number of bytes to be appended.
<i>theBytes</i>	A pointer to the bytes to be appended.

Result Arguments

<i>theObject</i>	The resulting byte object, with the appended bytes.
------------------	---

Description

This function appends *numBytes* bytes to byte object *theObject*. Its effect is equivalent to `GciStoreBytes(x, GciFetchSize_(x)+1, theBytes, numBytes)`.

GciAppendBytes raises an error if *theObject* is a Float or SmallFloat. Float and SmallFloat objects are of a fixed and unchangeable size.

See Also

[GciAppendChars](#), page 100

GciAppendChars

Append a C string to a byte object.

Syntax

```
void GciAppendChars(  
    OopType          theObject,  
    const char *     aString );
```

Input Arguments

<i>theObject</i>	A byte object to which the string is to be appended.
<i>aString</i>	A pointer to the string to be appended.

Result Arguments

<i>theObject</i>	The resulting byte object, with the appended string.
------------------	--

Description

This function appends the characters of *aString* to byte object *theObject*.

See Also

[GciAppendBytes](#), page 99

GciAppendOops

Append OOPs to the unnamed variables of a collection.

Syntax

```
void GciAppendOops(  
    OopType          theObject,  
    int              numOops,  
    const OopType*   theOops);
```

Input Arguments

<i>theObject</i>	A collection to which additional OOPs are to be added.
<i>numOops</i>	The number of OOPs to be added.
<i>theOops</i>	A pointer to the OOPs to be added.

Result Arguments

<i>theObject</i>	The resulting collection, with the added OOPs.
------------------	--

Description

Appends *numOops* OOPs to the unnamed variables of the collection *theObject*. If the collection is indexable, this is equivalent to:

```
GciStoreOops(theObject, GciFetchSize_(theObject)+1, theOops, numOops);
```

If the collection is an NSC, this is equivalent to:

```
GciAddOopsToNsc(theObject, theOops, numOops);
```

If the object is neither indexable nor an NSC, an error is generated.

GciBegin

Begin a new transaction.

Syntax

```
void GciBegin()
```

Description

This function begins a new transaction. If there is a transaction currently in progress, it aborts that transaction. Calling **GciBegin** is equivalent to the function call `GciExecuteStr("System beginTransaction", OOP_NIL)`.

See Also

GciAbort, page 91

GciExecuteStr, page 153

GciNbAbort, page 236

GciNbBegin, page 237

GciNbExecuteStr, page 246

GCI_BOOL_TO_OOP

(MACRO) Convert a C Boolean value to a GemStone Boolean object.

Syntax

```
OopType GCI_BOOL_TO_OOP(aBoolean)
```

Input Arguments

aBoolean The C Boolean value to be translated into a GemStone object.

Result Value

The OOP of the GemStone Boolean object that is equivalent to *aBoolean*.

Description

This macro translates a C Boolean value into the equivalent GemStone Boolean object. A C value of 0 translates to the GemStone Boolean object *false* (represented in your C program as OOP_FALSE). Any other C value translates to the GemStone Boolean object *true* (represented as OOP_TRUE). For more information, see Appendix A, "Reserved OOPs"

Example

```
int GCI_BOOL_TO_OOP_example(void)
{
    int z = 0;
    int nonZ = 99;

    OopType Fa = GCI_BOOL_TO_OOP(z);

    // any non-zero argument will produce a result of OOP_TRUE
    OopType Tr = GCI_BOOL_TO_OOP(nonZ);

    // the following will always be true
    return Fa == OOP_FALSE && Tr == OOP_TRUE;
}
```

See Also

[GciOopToBool](#), page 288

GciByteArrayToPointer

Given a result from **GciPointerToByteArray**, return a C pointer.

Syntax

```
void * GciByteArrayToPointer(  
    OopType      arg );
```

Input Arguments

arg A GemStone SmallInteger or ByteArray that was returned by **GciPointerToByteArray**.

Description

Given an argument that was the result of **GciPointerToByteArray**, this function returns the corresponding C pointer.

See Also

GciPointerToByteArray, page 308

GciCallInProgress

Determine if a GemBuilder call is currently in progress.

Syntax

```
BoolType GciCallInProgress()
```

Return Value

This function returns TRUE if a GemBuilder call is in progress, and FALSE otherwise.

Description

This function is intended for use within signal handlers. It can be called any time after **GciInit**.

GciCallInProgress returns FALSE if the process is currently executing within a user action and the user action's code is not within a GemBuilder call. It considers the highest (most recent) call context only.

See Also

GciInUserAction, page 221

GciCheckAuth

Gather the current authorizations for an array of database objects.

Syntax

```
void GciCheckAuth(  
    const OopType      oopArray[ ];  
    ArraySizeType     arraySize;  
    unsigned char      authCodeArray[ ] );
```

Input Arguments

<i>oopArray</i>	An array of OOPs of objects for which the user's authorization level. is to be ascertained. The caller must provide these values.
<i>arraySize</i>	The number of OOPs in <i>oopArray</i> .

Result Arguments

<i>authCodeArray</i>	The resulting array, having at least <i>arraySize</i> elements, in which the authorization values of the objects in <i>oopArray</i> are returned as 1-byte integer values.
----------------------	--

Description

GciCheckAuth checks the current user's authorization for each object in *oopArray* up to *arraySize*, returning each authorization code in the corresponding element of *authCodeArray*. The calling context is responsible for allocating enough space to hold the results.

Authorization levels are:

1. No authorization
2. Read authorization
3. Write authorization

Special objects, such as instances of `SmallInteger`, are reported as having read authorization.

Authorization values returned are those that have been committed to the database; they do not reflect changes you might have made in your local workspace. To query the local workspace, send an authorization query message to a particular object security policy using **GciPerform**.

If any member of *oopArray* is not a legal OOP, **GciCheckAuth** generates the error `OBJ_ERR_DOES_NOT_EXIST`. In that case, the contents of *authCodeArray* are undefined.

GCI_CHR_TO_OOP

(MACRO) Convert a C character value to a GemStone Character object.

Syntax

```
OopType GCI_CHR_TO_OOP(aChar)
```

Input Arguments

aChar The C character value to be translated into a GemStone object.

Result Value

The OOP of the GemStone Character object that is equivalent to *aChar*.

Description

This macro translates a C character value into the equivalent GemStone Character object. For more information, see Appendix A, “Reserved OOPs”.

Example

```
OopType GCI_CHR_TO_OOP_example(void)
{
    // return the OOP for the ASCII character 'a'
    OopType theOop = GCI_CHR_TO_OOP('a');
    return theOop;
}
```

See Also

[GciOopToChr](#), page 292

GciClampedTrav

Traverse an array of objects, subject to clamps.

Syntax

```
BoolType GciClampedTrav(
    const OopType *    theOops,
    int                numOops,
    GciClampedTravArgsSType *travArgs );
```

Input Arguments

<i>theOops</i>	An array of OOPs representing the objects to traverse.
<i>numOops</i>	The number of elements in <i>theOops</i> .
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType containing the following input argument fields:
OopType	<i>clampSpec</i> The OOP of the Smalltalk ClampSpecification to be used, or OOP_NIL, if the traversal is to operate without clamping.
int	<i>level</i> Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in <i>theOops</i> . When the level is 2, an object report is also obtained for the instance variables of each level-1 object. When the level is 0, the number of levels in the traversal is not restricted.
GciTravBufType *	<i>travBuff</i> A pointer to the traversal buffer.
int	<i>retrievalFlags</i> If (<i>retrievalFlags</i> & GCI_RETRIEVE_EXPORT != 0) then OOPs of non-special objects for which an object report header is returned in the traversal buffer are automatically added to the PureExportSet or the user action's export set (see GciSaveObjs). The value of <i>retrievalFlags</i> should be given by using the following GemBuilder mnemonics: GCI_RETRIEVE_DEFAULT GCI_RETRIEVE_EXPORT GCI_CLEAR_EXPORT causes the traversal to clear the PureExportSet or the user action's export set before it adds any OOPs to the traverse buffer.

Result Arguments

<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType containing the following result argument field:
-----------------	--

GciTravBufType * *travBuff*

The buffer for the results of the traversal. The first element placed in the buffer is the *actualBufferSize*, an integer that indicates how many bytes were actually stored in the buffer by this function. The remainder of the traversal buffer consists of a series of object reports, each of which is of type GciObjRepSType.

Return Value

Returns FALSE if the traversal is not yet completed. Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal** (that is, an object report was constructed for each object, minus the special objects).

Description

GciClampedTrav initiates a traversal of the specified objects, subject to the clamps in the specified ClampSpecification. In order to guarantee that the root object of the traversal will always have an entry in the traversal buffer, the root object is not subject to the specified clamps. Refer to "GciTraverseObjs" on page 408 for a detailed discussion of object traversal.

GemBuilder clamped traversal functions are used by the GemBuilder for Smalltalk implementation of object replication and are intended for similar sophisticated client applications.

See Also

GciMoreTraversal, page 234

GciSaveObjs, page 338

GciClassMethodForClass

Compile a class method for a class.

Syntax

```

OopType GciClassMethodForClass(
    OopType          source,
    OopType          oclass,
    OopType          category,
    OopType          symbolList );

```

Input Arguments

<i>source</i>	The OOP of a Smalltalk string to be compiled as a class method.
<i>oclass</i>	The OOP of the class with which the method is to be associated.
<i>category</i>	The OOP of a Smalltalk string, which contains the name of the category to which the method is added. If the category is nil (OOP_NIL), the compiler adds this method to the category "(as yet unclassified)".
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). Smalltalk resolves symbolic references in source code by using symbols that are available from <i>symbolList</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).

Return Value

Returns OOP_NIL, unless there were compiler warnings (such as variables declared but not used, etc.), in which case the return will be the OOP of a string containing the warning messages.

Description

This function compiles a class method for the given class. You may not compile any method whose selector begins with an underscore (_) character. Such selectors are reserved for use by the GemStone development team as private methods.

In addition, the Smalltalk virtual machine optimizes a small number of selectors. You may not compile any methods with any of those selectors. See the *Programming Guide* for a list of the optimized selectors.

To *remove* a class method, use **GciExecuteStr** instead.

Example

```

void GciClassMethodForClass_example(void)
{
    // Assumes the topaz code for GciFetchVaryingOop example
    // has been executed.

    OopType theClass = GciResolveSymbol("Component", OOP_NIL);

```

```
OopType oCateg = GciNewString("instance creation");
// method to create a new instance with a specified part number
OopType oMethodSrc = GciNewString(
"newWithNumber: aNum . | o | o := self new . o partNumber: aNum. ^
o");

GciClassMethodForClass(oMethodSrc, theClass, oCateg, OOP_NIL);
GciErrSType errInfo;
if (GciErr(&errInfo)) {
    printf("error category "FMT_OID" number %d, %s\n",
        errInfo.category, errInfo.number, errInfo.message);
}
}
```

See Also

[GciInstMethodForClass](#), page 219

GciClassNamedSize

Find the number of named instance variables in a class.

Syntax

```
int GciClassNamedSize(  
    OopType          oclass);
```

Input Arguments

oclass The OOP of the class from which to obtain information about instance variables. Appendix A, "Reserved OOPs" lists the OOP of each Smalltalk kernel class.

Return Value

Returns the number of named instance variables in the class, or zero if an error occurs.

Description

This function returns the number of named instance variables for the specified class, including those inherited from superclasses.

Example

```
int namedSizeExample(void)  
{  
    // find the class named Employee in the current symbolList  
    OopType empClass = GciResolveSymbol("Employee", OOP_NIL);  
    if (empClass == OOP_NIL) {  
        return -1; // class not found or other error.  
    }  
  
    int numIvs = GciClassNamedSize(empClass);  
    GciErrSType errInfo;  
    if (GciErr(&errInfo)) {  
        return -1; // error occurred  
    }  
  
    // return the number of named instance variables which will  
    // be >= 0  
    return numIvs;  
}
```

See Also

[GciIvNameToIdx](#), page 227

GciClearStack

Clear the Smalltalk call stack.

Syntax

```
void GciClearStack(
    OopType          process);
```

Input Arguments

process The OOP of a GsProcess object (obtained as the value of the *context* field of an error report returned by **GciErr**).

Description

Whenever a session executes a Smalltalk expression or sequence of expressions, the virtual machine creates and maintains a call stack that provides information about its state of execution. The call stack includes an ordered list of activation records related to the methods and blocks that are currently being executed.

If a soft break or an unexpected error occurs, the virtual machine suspends execution, creates a GsProcess object, and raises an error. The GsProcess object represents both the call stack when execution was suspended and any information that the virtual machine needs to resume execution. If there was no fatal error, your program can call **GciContinue** to resume execution. Call **GciClearStack** instead if there was a fatal error, or if you do not want your program to resume the suspended execution.

Example

The following example shows how an application can handle an error and either continue or terminate Smalltalk execution.

```
void clearStackExample(void)
{
    OopType result = GciExecuteStr(
        "| a | a := 10 + 10. nil halt . ^ a + 100",
        OOP_NIL/*use default symbolList for execution*/);

    // halt method is expected to generate error number
    RT_ERR_GENERIC_ERROR
    GciErrSType errInfo;
    if (! GciErr(&errInfo)) {
        printf("expected an error but none found\n");
        return;
    }
    if (errInfo.number == ERR_Halt) {
        // now continue the execution to finish the computation
        result = GciContinue(errInfo.context);
    } else {
        // FMT_OID format string is defined in gci.ht
```

```
    printf("unexpected error category \"FMT_OID\" number %d, %s\n",
errInfo.category, errInfo.number, errInfo.message);
    // terminate the execution
    GciClearStack(errInfo.context);
    return;
}
int val = GciOopToI32(result);
if (GciErr(&errInfo)) {
    printf("unexpected error category \"FMT_OID\" number %d, %s\n",
errInfo.category, errInfo.number, errInfo.message);
} else {
    if (val != 120) {
        printf("Wrong answer = %d\n", val);
    } else {
        printf("result = %d\n", val);
    }
}
}
```

See Also

[GciContinue](#), page 120

[GciSoftBreak](#), page 357

GciCommit

Write the current transaction to the database.

Syntax

```
BoolType GciCommit()
```

Return Value

Returns TRUE if the transaction committed successfully. Returns FALSE if the transaction fails to commit due to a concurrency conflict or in case of error.

Description

The **GciCommit** function attempts to commit the current transaction to the GemStone database.

GciCommit ignores any commit pending action that may be defined in the current GemStone session state.

Example

```
void GciCommit_example(void)
{
    // Call GciCommit and see if there was an error
    if ( ! GciCommit() ) {
        GciErrSType errInfo;
        if (GciErr(&errInfo)) {
            printf("commit failed with error %d , %s \n",
                errInfo.number, errInfo.message );
        } else {
            printf("commit failed due to transaction conflicts\n");
        }
    }
}
```

See Also

GciAbort, page 91
GCI_CHR_TO_OOP, page 107
GciNbAbort, page 236
GciNbCommit, page 239

GciCompileMethod

Compile a method.

Syntax

```

OopType GciCompileMethod(
    OopType      source,
    OopType      oclass,
    OopType      category,
    OopType      symbolList,
    OopType      overrideSelector,
    int          compileFlags,
    ushort      environmentId );

```

Input Arguments

<i>source</i>	The OOP of a Smalltalk string to be compiled as a class method.
<i>oclass</i>	The OOP of the class with which the method is to be associated.
<i>category</i>	The OOP of a Smalltalk string, which contains the name of the category to which the method is added. If the category is nil (OOP_NIL), the compiler adds this method to the category "(as yet unclassified)".
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). Smalltalk resolves symbolic references in source code by using symbols that are available from <i>symbolList</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>overrideSelector</i>	If not nil, this is a string that is converted to a symbol and used in precedence to the selector pattern in the method source when installing the method in the method dictionary. Sending 'selector' to the resulting method will also reflect the overrideSelector argument.
<i>compileFlags</i>	Compiler flags used for bootstrapping the Ruby environment.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see "environmentId" on page 80.

Return Value

Returns OOP_NIL, unless there were compiler warnings (such as variables declared but not used, etc.), in which case the return will be the OOP of a string containing the warning messages.

Description

This function is used for compiling a method. Replaces both GciInstMethodForClass and GciClassMethodForClass, and adds the environmentId argument.

This function compiles a method for the given class. You may not compile any method whose selector begins with an underscore (_) character. Such selectors are reserved for use by the GemStone development team as private methods.

In addition, the Smalltalk virtual machine optimizes a small number of selectors. You may not compile any methods with any of those selectors. See the *Programming Guide* for a list of the optimized selectors.

To *remove* a method, use **GciExecuteStr** instead.

See Also

GciClassMethodForClass, page 110

GciInstMethodForClass, page 219

GciCompress

Compress the supplied data, which can be uncompressed with **GciUncompress**.

Syntax

```
int GciCompress(
    char *          dest,
    uint *         destLen,
    const char *   source,
    uint           sourceLen );
```

Input Arguments

<i>dest</i>	Pointer to the buffer intended to hold the resulting compressed data.
<i>destLen</i>	Length, in bytes, of the buffer intended to hold the compressed data.
<i>source</i>	Pointer to the source data to compress.
<i>sourceLen</i>	Length, in bytes, of the source data.

Result Arguments

<i>dest</i>	The resulting compressed data.
-------------	--------------------------------

Return Value

GciCompress returns `Z_OK` (equal to 0) if the compression succeeded, or various error values if it failed; see the documentation for the `compress` function in the GNU `zlib` library at <http://www.gzip.org>.

Description

GciCompress passes the supplied inputs unchanged to the `compress` function in the GNU `zlib` library Version 1.2.3, and returns the result exactly as the GNU `compress` function returns it.

Example

```
#include <limits.h>

OopType compressByteArray(OopType byteArray)
{
    // given an input ByteArray , return a new ByteArray with
    // the contents of the input compressed .

    if (!GciIsKindOfClass(byteArray, OOP_CLASS_BYTE_ARRAY) )
        return OOP_NIL; /* error: input arg is not a ByteArray */

    int64 inputSize = GciFetchSize_(byteArray);
    if (inputSize > INT_MAX) {
        return OOP_NIL; // GciCompress supports max 2G bytes input
    }
}
```

```
}

int64 outputSize = inputSize;

ByteType *inputBuffer = (ByteType*)malloc( inputSize);
if (inputBuffer == NULL) {
    return OOP_NIL; // malloc failure
}
ByteType *outputBuffer = (ByteType*)malloc( outputSize);
if (outputBuffer == NULL) {
    free(inputBuffer);
    return OOP_NIL; // malloc failure
}

OopType resultOop = OOP_NIL;

int64 numRet = GciFetchBytes_(byteArray, 1/* start at first element */,
    inputBuffer, inputSize /* max bytes to fetch */ );
if (numRet == inputSize) {
    uint compressedSize;
    int status = GciCompress( (char *)outputBuffer,
        &compressedSize,
        (char *) inputBuffer, inputSize);
    if (status == 0) {
        // compress ok
        resultOop = GciNewByteObj(OOP_CLASS_BYTE_ARRAY,
            outputBuffer, (int64)compressedSize );
    } else {
        // compress failed
    }
} else {
    // error during FetchBytes
}
free(inputBuffer);
free(outputBuffer);
return resultOop;
}
```

See Also

[GciUncompress](#), page 412

GciContinue

Continue code execution in GemStone after an error.

Syntax

```
OopsType GciContinue(  
    OopsType  
    process );
```

Input Arguments

process The OOP of a GsProcess object (obtained as the value of the *context* field of an error report returned by **GciErr**).

Return Value

Returns the OOP of the result of the Smalltalk code that was executed. Returns OOP_NIL in case of error.

Description

The **GciContinue** function attempts to continue Smalltalk execution sometime after it was suspended. It is most useful for proceeding after GemStone encounters a pause message, a soft break (**GciSoftBreak**), or an application-defined error, since continuation is always possible after these events. Because **GciContinue** calls the virtual machine, the application user can also issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27.

It may also be possible to continue Smalltalk execution if the virtual machine detects a nonfatal error during a **GciExecute...** or **GciPerform** call. You may then want to use structural access functions to investigate (or modify) the state of the database before you call **GciContinue**.

Example

See the example for **GciClearStack** on page 113.

See Also

GciClearStack, page 113
GciErr, page 148
GciExecute, page 149
GciNbContinue, page 240
GciNbExecute, page 245

GciContinueWith

Continue code execution in GemStone after an error.

Syntax

```
OopsType GciContinueWith (  
    OopsType      process,  
    OopsType      replaceTopOfStack,  
    int           flags,  
    GciErrSType * error );
```

Input Arguments

<i>process</i>	The OOP of a GsProcess object (obtained as the value of the <i>context</i> field of an error report returned by GciErr).
<i>replaceTopOfStack</i>	If not OOP_ILLEGAL, replace the top of the Smalltalk evaluation stack with this value before continuing. If OOP_ILLEGAL, the evaluation stack is not changed.
<i>flags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk, as defined for GciPerformNoDebug .
<i>error</i>	If not NULL, continue with an error. This argument takes precedence over <i>replaceTopOfStack</i> .

Return Value

Returns the OOP of the result of the Smalltalk code that was executed. In case of error, this function returns OOP_NIL.

Description

This function is a variant of the **GciContinue** function, except that it allows you to modify the call stack and the state of the database before attempting to continue the suspended Smalltalk execution. This feature is typically used while implementing a Smalltalk debugger.

See Also

GciContinue, page 120
GciErr, page 148
GciExecute, page 149
GciNbContinueWith, page 241
GciNbExecute, page 245
GciPerformNoDebug, page 300

GciCreateByteObj

Create a new byte-format object.

Syntax

```

OopType GciCreateByteObj(
    OopType          oclass,
    OopType          objId,
    const ByteType * values,
    int64            numValues,
    int              clusterId,
    BoolType         makePermanent );

```

Input Arguments

<i>oclass</i>	The OOP of the class of the new object.
<i>objId</i>	The new object's OOP (obtained from GciGetFreeOop), or OOP_ILLEGAL. If you are trying to create a Symbol or DoubleByteSymbol, <i>objId</i> must be OOP_ILLEGAL. You cannot use the result of GciGetFreeOop to create a type of Symbol object.
<i>values</i>	Array of instance variable values.
<i>numValues</i>	Number of elements in values.
<i>clusterId</i>	ID of the cluster bucket in which to place the object. If <i>clusterId</i> is 0, use the cluster bucket (System currentClusterId). Otherwise, <i>clusterId</i> must be a positive integer <= GciFetchSize_(OOP_ALL_CLUSTER_BUCKETS).
<i>makePermanent</i>	Has no effect.

Return Value

GciCreateByteObj returns the OOP of the object it creates. The return value is the same as *objId* unless that value is OOP_ILLEGAL, in which case **GciCreateByteObj** assigns and returns a new OOP itself.

Description

Creates a new object using an object identifier (*objId*) previously obtained from **GciGetFreeOop** or **GciGetFreeOops**. For more about the semantics of such object identifiers, see **GciGetFreeOop** on page 204.

The object is created in temporary object space, and the garbage collector makes it permanent if the object is referenced, or becomes referenced, by another permanent object.

Values are stored into the object starting at the first named instance variable (if any) and continuing to the indexable (or NSC) instance variables if *oclass* is indexable or NSC. The caller must initialize any unused elements of **values* to OOP_NIL.

If *oclass* is an indexable or NSC class, then *numValues* may be as large or as small as desired. If *oclass* is neither indexable nor NSC, *numValues* must not exceed the number of named instance variables in the class. If *numValues* is less than number of named instance variables, then the size of the

newly-created object is the number of named instance variables and any instance variables beyond *numValues* are initialized to OOP_NIL.

For certain classes of byte format, namely DateTime, Float, and BigInteger, additional size restrictions apply.

For an indexable object, if *numValues* is greater than zero and *values* is NULL, then the object is created of size *numValues*, and is initialized to logical size *numValues*. (This is equivalent to new : *aSize* for classes Array or String.)

If **GciCreateByteObj** is being used to create an instance of OOP_CLASS_FLOAT or OOP_CLASS_SMALL_FLOAT, then the correct number of *value* bytes must be supplied at the time of creation.

If you are trying to create a Symbol or DoubleByteSymbol, *objId* **must** be OOP_ILLEGAL.

See Also

GciCreateOopObj, page 124

GciGetFreeOop, page 204

GciGetFreeOops, page 206

GciCreateOopObj

Create a new pointer-format object.

Syntax

```
OopType GciCreateOopObj(
    OopType          oclass,
    OopType          objId,
    const OopType *  values,
    int              numValues,
    int              clusterId,
    BoolType         makePermanent );
```

Input Arguments

<i>oclass</i>	The OOP of the class of the new object.
<i>objId</i>	The new object's OOP (obtained from GciGetFreeOop), or OOP_ILLEGAL.
<i>values</i>	Array of instance variable values.
<i>numValues</i>	Number of elements in values.
<i>clusterId</i>	ID of the cluster bucket in which to place the object. If <i>clusterId</i> is 0, use the cluster bucket (System currentClusterId). Otherwise, <i>clusterId</i> must be a positive integer <= GciFetchSize_(OOP_ALL_CLUSTER_BUCKETS).
<i>makePermanent</i>	Has no effect.

Return Value

GciCreateOopObj returns the OOP of the object it creates. The return value is the same as *objId* unless that value is OOP_ILLEGAL, in which case **GciCreateOopObj** assigns and returns a new OOP itself.

Description

Creates a new object using an object identifier (*objId*) previously obtained from **GciGetFreeOop** or **GciGetFreeOops**. For more about the semantics of such object identifiers, see **GciGetFreeOop** on page 204.

The object is created in temporary object space, and the garbage collector makes it permanent if the object is referenced, or becomes referenced, by another permanent object.

Values are stored into the object starting at the first named instance variable (if any) and continuing to the indexable (or NSC) instance variables if *oclass* is indexable or NSC. Values may be forward references to objects whose identifier has been allocated with **GciGetFreeOop**, but for which the object has not yet been created with **GciCreate...**. The caller must initialize any unused elements of **values* to OOP_NIL.

Because it is illegal to create a forward reference to a Symbol, any **GciCreate...** call that creates a Symbol will fail if the client's *objId* of the created object was already used as a forward reference.

If *oclass* is an indexable or NSC class, then *numValues* may be as large or as small as desired. If *oclass* is neither indexable nor NSC, *numValues* must not exceed the number of named instance variables

in the class. If *numValues* is less than number of named instance variables, then the size of the newly-created object is the number of named instance variables and any instance variables beyond *numValues* are initialized to OOP_NIL.

For an indexable object, if *numValues* is greater than zero and *values* is NULL, then the object is created of size *numValues*, and is initialized to logical size *numValues*. (This is equivalent to `new : aSize` for classes Array or String.)

See Also

GciCreateByteObj, page 122

GciGetFreeOop, page 204

GciGetFreeOops, page 206

GciCTimeToDateTime

Convert a C date-time representation to the equivalent GemStone representation.

Syntax

```
BoolType GciCTimeToDateTime(  
    time_t          arg,  
    GciDateTimeSType * result );
```

Input Arguments

arg The C time value to be converted.

Result Arguments

result A pointer to the C struct in which to place the converted value.

Return Value

Returns TRUE if the conversion succeeds; otherwise returns FALSE.

Description

Converts a **time_t** value to **GciDateTimeSType**. On systems where **time_t** is a signed value, **GciCTimeToDateTime** generates an error if *arg* is negative.

See Also

GciDateTimeToCTime, page 127

GciDateTimeToCTime

Convert a GemStone date-time representation to the equivalent C representation.

Syntax

```
time_t GciDateTimeToCTime(  
    const GciDateTimeSType *arg );
```

Input Arguments

arg An instance of **GciDateTimeSType** to be converted.

Return Value

A C time value of type **time_t**.

Description

Converts an instance of **GciDateTimeSType** to the equivalent **time_t** value.

See Also

GciCTimeToDateTime, page 126

GciDbgEstablish

Specify the debugging function for GemBuilder to execute before most calls to GemBuilder functions.

Syntax

```
GciDbgFuncType * GciDbgEstablish(  
    GciDbgFuncType * newDebugFunc );
```

Input Arguments

newDebugFunc A pointer to a C function that will be called before each subsequent GemBuilder call. Note that this function will not be called before any of the following GemBuilder functions or macros: `GCI_ALIGN`, `GCI_BOOL_TO_OOP`, `GCI_CHR_TO_OOP`, `GciErr`, or `GciDbgEstablish` itself.

The `newDebugFunc` function is passed a single null-terminated string argument, (of type `const char []`), the name of the GemBuilder function about to be called.

Return Value

Returns a pointer to the *newDebugFunc* specified in the previous `GciDbgEstablish` call (if any).

Description

This function establishes the name of a C function (most likely a debugging routine) to be called before your program calls any GemBuilder function or macro (except those named above). Before each GemBuilder call, a single argument, a null-terminated string that names the GemBuilder function about to be executed, is passed to the specified *newDebugFunc*.

To disable previous debugging routines, your program can use the following statement:

```
GciDbgEstablish(NULL);
```

Example

```
void traceGciFunc(const char* gciFname)  
{  
    printf("trace gci call %s \n", gciFname);  
}  
  
void debugEstablishExample(void)  
{  
    GciDbgEstablish(traceGciFunc); // enable tracing  
  
    GciFetchSize_(OOP_CLASS_STRING); // this call will be traced  
  
    GciDbgEstablish(NULL); // shut off tracing  
}
```


See Also

`GciDbgEstablishToFile`, page 130

`GciErr`, page 148

GciDbgEstablishToFile

Write trace information for most GemBuilder functions to a file.

Syntax

```
BoolType GciDbgEstablishToFile(  
    const char *      fileName );
```

Input Arguments

fileName The file to which trace information is to be written.

Return Value

Returns TRUE if the file operation was successful.

Description

This function causes trace information for most GemBuilder functions to be written to a file. If the file already exists, it is opened in append mode. If *fileName* is NULL and tracing to a file is not currently active, trace information will be written to stdout.

Calling **GciDbgEstablishToFile** supersedes the effect of any previous calls to **GciDbgEstablish** or **GciDbgEstablishToFile**.

To terminate tracing to an active file, your program can use the following statement:

```
GciDbgEstablishToFile(NULL);
```

Alternatively, your program can call **GciShutdown**.

For details about the trace information generated, see **GciDbgEstablish**.

See Also

GciDbgEstablish, page 128

GciErr, page 148

GciDbgLogString

Pass a message to a trace function.

Syntax

```
void GciDbgLogString(  
    const char *      message);
```

Input Arguments

message A message to be passed to **GciDbgEstablish** or **GciDbgEstablishToFile**.

Description

If either **GciDbgEstablish** or **GciDbgEstablishToFile** has been called to activate tracing of GemBuilder calls, this function passes the argument to the trace function.

If tracing is not active, this function has no effect.

See Also

GciDbgEstablish, page 128

GciDbgEstablishToFile, page 130

GciDeclareAction

An alternative way to associate a C function with a Smalltalk user action.

Syntax

```
void GciDeclareAction(  
    const char*      name,  
    void*            func,  
    int              nargs,  
    uint             flags,  
    BoolType         errorIfDuplicate );
```

Input Arguments

<i>name</i>	The user action name (a case-insensitive, null-terminated string).
<i>func</i>	A pointer to the C user action function.
<i>nargs</i>	The number of arguments in the C function.
<i>flags</i>	Flags that are rarely used. Mainly for internal use.
<i>errorIfDuplicate</i>	If True, return an error if there is already a user action with the specified name. If False, leave the existing user action in place and ignore the current call.

Description

This function associates a user action name (declared in Smalltalk) with a user-written C function. **GciDeclareAction** allows you to declare a user action by passing each field of the user action structure to the function as a separate argument. Because the user action structure is encapsulated within the function itself, there's no need to explicitly allocate and free memory, as is required with **GciInstallUserAction** (which uses the data structure defined by `GciUserActionSType`).

See Also

Chapter 3, "Writing C Functions To Be Called from GemStone", starting on page 47
"The User Action Information Structure" on page 78
GciInstallUserAction, page 218
GciUserActionInit, page 413
GciUserActionShutdown, page 414

GciDecodeOopArray

Decode an OOP array that was previously run-length encoded.

Syntax

```
int GciDecodeOopArray(  
    OopType *          encodedOopArray,  
    const int          numEncodedOops,  
    OopType *          decodedOopArray,  
    const int          decodedOopArraySize);
```

Input Arguments

encodedOopArray An OOP array that was encoded by a call to **GciEncodeOopArray**.
numEncodedOops The number of OOPs in *encodedOopArray*.
decodedOopArraySize The maximum number of OOPs in *decodedOopArray*.

Result Arguments

decodedOopArray The decoded OOP array that had been run-length encoded.

Return Value

Returns the number of OOPs placed in *decodedOopArray*.

Description

This function decodes the OOPs in *encodedOopArray* that were run-length encoded using **GciEncodeOopArray** and places the result in *decodedOopArray*.

The *decodedOopArraySize* must be large enough to hold all decoded OOPs. If it is not, no decode is performed and **decodedOopArraySize* is set to -1.

See Also

GciFetchNumEncodedOops, page 176
GciEnableFreeOopEncoding, page 143
GciEncodeOopArray, page 146
GciGetFreeOopsEncoded, page 208

GciDecSharedCounter

Decrement the value of a shared counter.

Syntax

```
BoolType GciDecSharedCounter(  
    int64_t          counterIdx,  
    int64_t *       value,  
    int64_t *       floor);
```

Input Arguments

<i>counterIdx</i>	The offset into the shared counters array of the value to decrement.
<i>value</i>	Pointer to a value that indicates how much to decrement the shared counter by.
<i>floor</i>	The minimum possible value for the shared counter. The counter cannot be decremented below this value. If <i>floor</i> is NULL, then a <i>floor</i> value of INT_MIN (-2147483647) will be used.

Result Arguments

<i>value</i>	Pointer to a value that indicates the new value of the shared counter, after the decrement.
--------------	---

Return Value

Returns a C Boolean value indicating if the shared counter was successfully decremented by the given amount. Returns TRUE if successful, FALSE if an error occurred.

Description

This function decrements the value of a particular shared counter by a specified amount. The shared counter is specified by index. The value of this shared counter cannot be decremented to a value lower than *floor*.

This function is not supported for remote GCI interfaces, and will always return FALSE.

See Also

[GciFetchNumSharedCounters](#), page 177
[GciIncSharedCounter](#), page 214
[GciSetSharedCounter](#), page 353
[GciReadSharedCounter](#), page 316
[GciReadSharedCounterNoLock](#), page 317
[GciFetchSharedCounterValuesNoLock](#), page 190

GciDirtyExportedObjs

Find all objects in the ExportedDirtyObjs set.

Syntax

```
BoolType GciDirtyExportedObjs(
    OopType      theOops[ ],
    int *        numOops );
```

Input Arguments

numOops The maximum number of objects that can be put into *theOops* buffer.

Result Arguments

theOops An array of the dirty exported objects found.
numOops The number of dirty exported objects found.

Return Value

This function returns a C Boolean value indicating whether or not the complete set of dirty objects has been returned in *theOops* in one or more calls. TRUE indicates that the complete set has been returned, and FALSE indicates that it has not.

Description

This function returns a list of all objects that are in the ExportedDirtyObjs set, which includes all objects in the PureExportSet that have been made “dirty” since the ExportedDirtyObjs set was last initialized or retrieved using **GciDirtyAlteredObjs**, **GciDirtyExportedObjs**, **GciDirtyObjsInit**, **GciDirtySaveObjs**, or **GciTrackedObjsFetchAllDirty**. Object are added to the PureExportSet using **GciSaveObjs** or by other functions that invoke **GciSaveObjs**. An object is considered dirty (changed) under one or more of the following conditions:

- The object was changed by Smalltalk execution from this session.
- The object was changed by a call from this session to any GemBuilder function from within a user action.
- The object was changed by a call from this session to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.
- The object was read by this session, and after this session did a commit, begin, or abort transaction, the session now has visibility to changes to the object committed by another session.
- The object is persistent, and this session aborted its changes to the object, thus rolling back the Smalltalk in-memory state to the previously committed state.

Calls to **GciStore...** (other than **GciStorePaths**), **GciAppend...**, **GciReplace...**, and **GciCreate...** do not put the modified object into the set of dirty objects (unless the call is from within a user action).

The assumption is that the client does not want the dirty set to include modifications that the client has explicitly made.

The function **GciDirtyObjsInit** must be executed once after **GciLogin** before this function can be called, because it depends upon GemStone's set of dirty objects.

The user is expected to call this function repeatedly while it returns FALSE, until it finally returns TRUE. When this function returns TRUE, it first clears the set of dirty objects.

Note that **GciDirtyExportedObjs** removes OOPs from the ExportedDirtyObjs set as they are enumerated.

See Also

"Garbage Collection" on page 40

GciDirtyObjsInit, page 137

GciDirtySaveObjs, page 138

GciDirtyTrackedObjs, page 140

GciTrackedObjsFetchAllDirty, page 405

GciHiddenSetIncludesOop, page 211

GciReleaseAllGlobalOops, page 319

GciSaveGlobalObjs, page 337

GciSaveObjs, page 338

GciDirtyObjsInit

Begin tracking which objects in the session workspace change.

Syntax

```
void GciDirtyObjsInit()
```

Description

GemStone can track which objects in a session change, but doing so has a measurable cost. By default, GemStone does not do it. The **GciDirtyObjsInit** function permits an application to request GemStone to maintain that set of dirty objects, the `ExportedDirtyObjects`, when it is needed. Once initialized, GemStone tracks dirty objects until **GciLogout** is executed.

GciDirtyObjsInit must be called once after **GciLogin** before **GciDirtyExportedObjs**, **GciDirtySaveObjs**, or **GciTrackedObjsFetchAllDirty** in order for those functions to operate properly, because they depend upon GemStone's set of dirty objects.

An object is considered dirty (changed) under one or more of the following conditions:

- The object was changed by Smalltalk execution from this session.
- The object was changed by a call from this session to any GemBuilder function from within a user action.
- The object was changed by a call from this session to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.
- The object was read by this session, and after this session did a commit, begin, or abort transaction, the session now has visibility to changes to the object committed by another session.
- The object is persistent, and this session aborted its changes to the object, thus rolling back the Smalltalk in-memory state to the previously committed state.

See Also

GciDirtyExportedObjs, page 135
GciDirtySaveObjs, page 138
GciTrackedObjsFetchAllDirty, page 405
GciHiddenSetIncludesOop, page 211

GciDirtySaveObjs

Find all exported or tracked objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.

Syntax

```
BoolType GciDirtySaveObjs(  
    OopType      theOops[ ],  
    int *        numOops );
```

Input Arguments

numOops The number of objects that can be put into *theOops* buffer.

Result Arguments

theOops An array of the dirty cached objects found.
numOops The number of dirty cached objects found.

Return Value

This function returns a C Boolean value indicating whether or not the complete set of dirty objects has been returned in *theOops* in one or more calls. TRUE indicates that the complete set has been returned, and FALSE indicates that it has not.

Description

GciDirtySaveObjs finds all objects that are in the ExportedDirtyObjs or TrackedDirtyObjs sets. The ExportedDirtyObjs set includes all objects in PureExportSet that have been made “dirty” since the ExportedDirtyObjs set was last reset, and the TrackedDirtyObjs set includes all objects in the GciTrackedObjs set that have been made “dirty” since the TrackedDirtyObjs set was last reset.

The ExportedDirtyObjs set is initialized by **GciDirtyObjsInit**; it is cleared by calls to **GciDirtyAlteredObjs**, **GciDirtyExportedObjs**, **GciDirtySaveObjs** (this function), or **GciTrackedObjsFetchAllDirty**. The TrackedDirtyObjs set is initialized by **GciTrackedObjsInit** and cleared by calls to **GciDirtyAlteredObjs**, **GciDirtySaveObjs** (this function), **GciDirtyTrackedObjs**, or **GciTrackedObjsFetchAllDirty**.

An object is considered dirty (changed) under one or more of the following conditions:

- The object was changed by Smalltalk execution from this session.
- The object was changed by a call from this session to any GemBuilder function from within a user action.
- The object was changed by a call from this session to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.
- The object was read by this session, and after this session did a commit, begin, or abort transaction, the session now has visibility to changes to the object committed by another session.
- The object is persistent, and this session aborted its changes to the object, thus rolling back the Smalltalk in-memory state to the previously committed state.

Calls to **GciStore...** (other than **GciStorePaths**), **GciAppend...**, **GciReplace...**, and **GciCreate...** do not put the modified object into the set of dirty objects (unless the call is from within a user action). The assumption is that the client does not want the dirty set to include modifications that the client has explicitly made.

GciDirtyObjsInit must be called once after **GciLogin** before **GciDirtySaveObjs** can be executed, because it depends upon GemStone's set of dirty objects.

The user is expected to call **GciDirtySaveObjs** repeatedly while it returns FALSE, until it finally returns TRUE. When **GciDirtySaveObjs** returns TRUE, it first clears the set of dirty objects.

For details about the PureExportSet, see **GciSaveObjs**. For details about the GciTrackedObjs set, see **GciSaveAndTrackObjs**.

Note that **GciDirtySaveObjs** removes OOPs from the ExportedDirtyObjs and TrackedDirtyObjs sets.

See Also

"Garbage Collection" on page 40

GciDirtyExportedObjs, page 135

GciDirtyObjsInit, page 137

GciDirtyTrackedObjs, page 140

GciTrackedObjsFetchAllDirty, page 405

GciSaveObjs, page 338

GciDirtyTrackedObjs

Find all tracked objects that have changed and are therefore in the TrackedDirtyObjs set.

Syntax

```
BoolType GciDirtyTrackedObjs(
    OopType          theOops[ ],
    int *            numOops );
```

Input Arguments

numOops The maximum number of objects that can be put into *theOops* buffer.

Result Arguments

theOops An array of the dirty tracked objects found.
numOops The number of dirty tracked objects found.

Return Value

This function returns a C Boolean value indicating whether or not the complete set of dirty tracked objects has been returned in *theOops* in one or more calls. TRUE indicates that the complete set has been returned, and FALSE indicates that it has not.

Description

This function returns a list of all objects that are in the TrackedDirtyObjs set, which includes all objects that are in the GciTrackedObjs set and have been made “dirty” since the GciTrackedObjs set was initialized or cleared. Functions that initialize or remove objects from the TrackedDirtyObjs set are **GciDirtyAlteredObjs**, **GciDirtySaveObjs**, **GciDirtyTrackedObjs** (this function), **GciTrackedObjsFetchAllDirty** and **GciTrackedObjsInit**.

An object is considered dirty (changed) under one or more of the following conditions:

- The object was changed by Smalltalk execution from this session.
- The object was changed by a call from this session to any GemBuilder function from within a user action.
- The object was changed by a call from this session to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.
- The object was read by this session, and after this session did a commit, begin, or abort transaction, the session now has visibility to changes to the object committed by another session.
- The object is persistent, and this session aborted its changes to the object, thus rolling back the Smalltalk in-memory state to the previously committed state.

Calls to **GciStore...** (other than **GciStorePaths**), **GciAppend...**, **GciReplace...**, and **GciCreate...** do not put the modified object into the set of dirty objects (unless the call is from within a user action).

The assumption is that the client does not want the dirty set to include modifications that the client has explicitly made.

This function may only be called after **GciTrackedObjsInit** has been executed, because it depends upon GemStone's set of tracked objects. The user is expected to call this function repeatedly while it returns FALSE, until it finally returns TRUE. When this function returns TRUE, it first clears the set of dirty objects.

Note that **GciDirtyTrackedObjs** removes OOPs from the TrackedDirtyObjs set.

See Also

"Garbage Collection" on page 40

GciDirtySaveObjs, page 138

GciHiddenSetIncludesOop, page 211

GciReleaseAllTrackedOops, page 321

GciSaveAndTrackObjs, page 336

GciTrackedObjsFetchAllDirty, page 405

GciTrackedObjsInit, page 407

Gci_doubleToSmallDouble

Convert a C double to a SmallDouble object.

Syntax

```
OopType Gci_doubleToSmallDouble(  
    double          aFloat );
```

Return Value

Returns the OOP of the GemStone SmallDouble object that corresponds to the C value. If the C value is not representable as a GemStone SmallDouble, return OOP_ILLEGAL.

Description

This function translates a C double into the equivalent GemStone SmallDouble object.

See Also

[GciFltToOop](#), page 203

GciEnableFreeOopEncoding

Enable run-length encoding of free OOPs.

Syntax

```
void GciEnableFreeOopEncoding( )
```

Description

This function enables run-length encoding of free OOPs sent between the Gem and the GemBuilder client. This function increases CPU consumption on both the client and the Gem, and decreases the number of bytes passed on the network.

See Also

[GciDecodeOopArray](#), page 133
[GciEncodeOopArray](#), page 146
[GciFetchNumEncodedOops](#), page 176
[GciGetFreeOopsEncoded](#), page 208

GciEnableFullCompression

Enable full compression between the client and the RPC version of GemBuilder.

Syntax

```
void GciEnableFullCompression()
```

Description

This function enables full compression (in both directions) between the client and GciRpc (the “remote procedure call” version of GemBuilder). This function has no effect for linked sessions.

See Also

[GciIsRemote](#), page 224

GciEnableSignaledErrors

Establish or remove GemBuilder visibility to signaled errors from GemStone.

Syntax

```
BoolType GciEnableSignaledErrors(  
    BoolType          newState );
```

Input Arguments

newState The new state of signaled error visibility: TRUE for visible.

Return Value

This function returns TRUE if signaled errors are already visible when it is called.

Description

GemStone permits selective response to signal errors: `RT_ERR_SIGNAL_ABORT`, `RT_ERR_SIGNAL_COMMIT`, and `RT_ERR_SIGNAL_GEMSTONE_SESSION`. The default condition is to leave them all invisible. GemStone responds to each single kind of signal error only after an associated method of class System has been executed: `enableSignaledAbortError`, `enableSignaledObjectsError`, and `enableSignaledGemStoneSessionError` respectively.

After **GciInit** executes successfully, the GemBuilder default condition also leaves all signal errors invisible. The **GciEnableSignaledErrors** function permits GemBuilder to respond automatically to signal errors. However, GemStone must respond to each kind of error in order for GemBuilder to respond to it. Thus, if an application calls **GciEnableSignaledErrors** with *newState* equal to TRUE, then GemBuilder responds automatically to exactly the same kinds of signal errors as GemStone. If GemStone has not executed any of the appropriate System methods, then this call has no effect until it does.

When enabled, GemBuilder checks for signal errors at the start of each function that accesses the database. It treats any that it finds just like any other errors, through **GciErr** or the **GciLongJump** mechanism, as appropriate.

Automatic checking for signalled errors incurs no extra runtime cost. The check is optimized into the check for a valid session. However, instead of checking automatically, these errors can be polled by calling the **GciPollForSignal** function.

GciEnableSignaledErrors may be called before calling **GciLogin**.

See Also

GciErr, page 148
GciPollForSignal, page 309

GciEncodeOopArray

Encode an array of OOPs, using run-length encoding.

Syntax

```
int GciEncodeOopArray(  
    OopType *          oopArray,  
    const int         numOops,  
    OopType *         encodedOopArray,  
    BoolType          needsSorting );
```

Input Arguments

<i>oopArray</i>	An OOP array to be encoded.
<i>numOops</i>	The number of OOPs in <i>oopArray</i> .
<i>needsSorting</i>	If <i>oopArray</i> is known to be in ascending order, set this to FALSE; otherwise set it to TRUE.

Result Arguments

<i>encodedOopArray</i>	The encoded OOP array.
------------------------	------------------------

Return Value

Returns the number of elements in the encoded array. Returns -1 indicating an error if the input array was found to be out of sequence and *needsSorting* was set to FALSE.

Description

This function encodes the OOPs in *oopArray* using run-length encoding and places the result in *encodedOopArray*. Both *oopArray* and *encodedOopArray* must have the size *numOops*.

See Also

[GciDecodeOopArray](#), page 133
[GciEnableFreeOopEncoding](#), page 143
[GciFetchNumEncodedOops](#), page 176
[GciGetFreeOopsEncoded](#), page 208

GciEncrypt

Encrypt a password string.

Syntax

```
char * GciEncrypt(  
    const char*      password,  
    char            outBuff[],  
    unsigned int     outBuffSize);
```

Input Arguments

<i>password</i>	String containing a password.
<i>outBuffSize</i>	The maximum number of bytes to place in <i>outBuff</i> .

Result Arguments

<i>outBuff</i>	Array in which to put the encrypted password.
----------------	---

Return Value

Returns a pointer to the first character of the encrypted password in *outBuff*[], or NULL if the encrypted password is larger than *outBuffSize*.

Description

This function encrypts a host or GemStone password for use in **GciSetNetEx** or **GciLogin** with password encryption enabled.

See Also

GciLoginEx, page 231
GciSetNetEx, page 349

GciErr

Prepare a report describing the most recent GemBuilder error.

Syntax

```
BoolType GciErr(  
    GciErrSType *    errorReport );
```

Result Arguments

errorReport Address of a GemBuilder error report structure.

Return Value

TRUE indicates that an error has occurred. The *errorReport* parameter has been modified to contain the latest error information, and the internal error buffer in GemBuilder has been cleared. You can only call **GciErr** once for a given error. If **GciErr** is called a second time, the function returns FALSE.

If the result is TRUE, all objects referenced from *errorReport* have been added to the PureExportSet, unless the error occurred during a **GciStoreTravDoTravRefs_**, in which case all objects referenced from *errorReport* have been added to the ReferencedSet rather than the PureExportSet.

FALSE indicates no error occurred, and the contents of *errorReport* are unchanged.

Description

Your application program can call **GciErr** to determine whether or not the previous GemBuilder function call resulted in an error. If an error has occurred, this function provides information about the error and about the state of the GemStone system. In the case of a fatal error, your connection to GemStone is lost, and the current session ID (from **GciGetSessionId**) is reset to GCI_INVALID_SESSION_ID.

The **GciErr** function is especially useful when error traps are disabled or are not present. See “GciPopErrJump” on page 311 for information about using general-purpose error traps in GemBuilder. The section “The Error Report Structure” on page 72 describes the C structure for error reports.

See Also

GciClearStack, page 113
GciContinue, page 120
GciExecute, page 149
GciPopErrJump, page 311
GciRaiseException, page 315

GciExecute

GciExecute_

Execute a Smalltalk expression contained in a String or Utf8 object.

Syntax

```

OopType GciExecute(
    OopType          source,
    OopType          symbolList );

OopType GciExecute_(
    OopType          source,
    OopType          symbolList,
    ushort          environmentId);

```

Input Arguments

<i>source</i>	The OOP of a String or Utf8 containing a sequence of one or more statements to be executed.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the execution result. In case of error, this function returns OOP_NIL.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution. This is roughly equivalent to executing the body of a nameless procedure (method).

In most cases, you may find it more efficient to use **GciExecuteStr**. That function takes a C string as its argument, thus reducing the number of network round-trips required to execute the code. With **GciExecute**, you must first convert the source to a String or Utf8 object (see the following example.) If the source is already one of these kinds of object, however, **GciExecute** will be more efficient.

Because **GciExecute** calls the virtual machine, the user can issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27.

Example

```
void executeExample(void)
```

```
{
  OopType oString = GciNewString(" ^ 3 + 4 ");

  OopType result = GciExecute(oString, OOP_NIL);
  if (result == OOP_NIL) {
    printf("error from execution\n");
  } else {
    BoolType conversionErr = FALSE;
    int val = GciOopToI32_(result, &conversionErr);
    if (conversionErr) {
      printf("Error converting result to C int\n");
    } else {
      printf("result = %d\n", val);
    }
  }
}
```

See Also

GciContinue, page 120
GciErr, page 148
GciExecuteFromContext, page 151
GciExecuteStr, page 153
GciExecuteStrFetchBytes, page 155
GciExecuteStrFromContext, page 157
GciNbExecute, page 245

GciExecuteFromContext

GciExecuteFromContext_

Execute a Smalltalk expression contained in a String or Utf8 object as if it were a message sent to another object.

Syntax

```

OopType GciExecuteFromContext(
    OopType          source,
    OopType          contextObject,
    OopType          symbolList );

OopType GciExecuteFromContext_(
    OopType          source,
    OopType          contextObject,
    OopType          symbolList,
    ushort           environmentId );

```

Input Arguments

<i>source</i>	The OOP of a String or Utf8 containing a sequence of one or more statements to be executed.
<i>contextObject</i>	The OOP of any GemStone object.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the execution result. In case of error, this function returns OOP_NIL.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution. The source is executed as though *contextObject* were the receiver. That is, the pseudo-variable *self* will have the value *contextObject* during the execution. Messages in the source are executed as defined for *contextObject*.

For example, if *contextObject* is an instance of Association, the *source* can reference the pseudo-variables *key* and *value* (referring to the instance variables of the Association *contextObject*). If any pool dictionaries were available to Association, the *source* could reference them too.

In most cases, you may find it more efficient to use **GciExecuteStrFromContext**. That function takes a C string as its argument, thus reducing the number of network round-trips required to execute

the code. With **GciExecuteFromContext**, you must first convert the source to a String or Utf8 object. If the source is already one of these kinds of object, however, **GciExecuteFromContext** will be more efficient.

Because **GciExecuteFromContext** calls the virtual machine, the user can issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27.

See Also

GciContinue, page 120

GciErr, page 148

GciExecute, page 149

GciExecuteStr, page 153

GciExecuteStrFetchBytes, page 155

GciExecuteStrFromContext, page 157

GciNbExecuteStrFromContext, page 249

GciExecuteStr

GciExecuteStr_

Execute a Smalltalk expression contained in a C string.

Syntax

```

OopType GciExecuteStr(
    const char    source[ ],
    OopType      symbolList );

OopType GciExecuteStr_(
    const char    source[ ],
    OopType      symbolList,
    ushort       environmentId );

```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the execution result. In case of error, this function returns OOP_NIL.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution.

If the source is already a String object, you may find it more efficient to use **GciExecute**. That function takes the OOP of a String as its argument.

Because **GciExecuteStr** calls the virtual machine, the user can issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27

Example

```
void executeStrExample(void)
{
    // get the symbolList for UserProfile named 'romeo'
    OopType symbolList = GciExecuteStr(
"(AllUsers userWithId: 'Newton') symbolList", OOP_NIL);

    // get the value associated with key "#GciStructsMd5" in that
    // symbolList ; expected to be a kind of String
    OopType keysum = GciExecuteStr("GciStructsMd5", symbolList);

    // fetch characters of the String
    char buf[1024];
    GciFetchChars_(keysum, 1, buf, sizeof(buf));

    GciErrSType errInfo;
    if ( GciErr(&errInfo)) {
        // FMT_OID format string is defined in gci.ht
        printf("unexpected error category "FMT_OID" number %d, %s\n",
            errInfo.category, errInfo.number, errInfo.message);
    } else {
        printf("#GciStructsMd5 is %s \n", buf);
    }
}
```

See Also

[GciContinue](#), page 120
[GciErr](#), page 148
[GciExecute](#), page 149
[GciExecuteFromContext](#), page 151
[GciExecuteStrFetchBytes](#), page 155
[GciExecuteStrFromContext](#), page 157
[GciNbExecuteStr](#), page 246

GciExecuteStrFetchBytes

Execute a Smalltalk expression contained in a C string, returning byte-format results.

Syntax

```
int64 GciExecuteStrFetchBytes(
    const char *      source,
    int64             sourceSize,
    OopType           sourceClass,
    OopType           contextObject,
    OopType           symbolList,
    ByteType *       result,
    int64             maxResultSize );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>sourceSize</i>	The number of bytes in the <i>source</i> , or -1. If <i>sourceSize</i> is -1, <code>strlen(source)</code> is used.
<i>sourceClass</i>	The OOP of the class that <i>source</i> should be converted to. Examples are <code>OOP_CLASS_STRING</code> , <code>OOP_CLASS_Utf8</code> , <code>OOP_CLASS_Unicode7</code> .
<i>contextObject</i>	The OOP of any GemStone object. The code to be executed is compiled as if it were an instance method in the class of <i>contextObject</i> . A value of <code>OOP_NO_CONTEXT</code> means no context.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of <code>SymbolListDictionary</code>). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of <code>OOP_NIL</code> means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>maxResultSize</i>	Maximum size of the resulting string.

Result Arguments

<i>result</i>	Array in which to put the results of executing the <i>source</i> .
---------------	--

Return Value

Returns the the number of bytes returned in the result buffer, or -1 if an error occurred.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution. The execution result, which should be a byte format object is returned in the **result* buffer.

Execution is in environment 0 using `GCI_PERFORM_FLAG_ENABLE_DEBUG`.

See Also

GciContinue, page 120
GciErr, page 148
GciExecute, page 149
GciExecuteFromContext, page 151
GciExecuteStr, page 153
GciExecuteStrFromContext, page 157
GciNbExecuteStrFetchBytes, page 247

GciExecuteStrFromContext

GciExecuteStrFromContext_

Execute a Smalltalk expression contained in a C string as if it were a message sent to an object.

Syntax

```

OopType GciExecuteStrFromContext(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList );

```

```

OopType GciExecuteStrFromContext_(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    ushort             environmentId );

```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>contextObject</i>	The OOP of any GemStone object.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the execution result. In case of error, this function returns OOP_NIL.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution. The source is executed as though *contextObject* were the receiver. That is, the pseudo-variable `self` will have the value *contextObject* during the execution. Messages in the source are executed as defined for *contextObject*.

For example, if *contextObject* is an instance of Association, the source can reference the pseudo-variables *key* and *value* (referring to the instance variables of the Association *contextObject*). If any pool dictionaries were available to Association, the source could reference them too.

Because **GciExecuteStrFromContext** calls the virtual machine, the user can issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27.

Example

```
void executeFromContextExample(void)
{
    // get the Association with key UserProfileSet in dictionary
    Globals
    OopType oAssoc = GciExecuteStr("Globals associationAt:
#UserProfileSet",
                                OOP_NIL);

    OopType oResult = GciExecuteStrFromContext(" ^ value ", oAssoc,
OOP_NIL);

    if (oResult != OOP_CLASS_USERPROFILE_SET) {
        printf("unexpected result" FMT_OID " \n", oResult);
    }
}
```

See Also

[GciContinue](#), page 120
[GciErr](#), page 148
[GciExecute](#), page 149
[GciExecuteFromContext](#), page 151
[GciExecuteStr](#), page 153
[GciExecuteStrFetchBytes](#), page 155
[GciNbExecuteStrFromContext](#), page 249

GciExecuteStrTrav

GciExecuteStrTrav_

First execute a Smalltalk expression contained in a C string as if it were a message sent to an object, then traverse the result of the execution.

Syntax

```
BoolType GciExecuteStrTrav(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    GciClampedTravArgsSType *travArgs );
```

```
BoolType GciExecuteStrTrav_(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    GciClampedTravArgsSType *travArgs,
    ushort             environmentId );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.						
<i>contextObject</i>	The OOP of any GemStone object. A value of OOP_ILLEGAL means no context.						
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).						
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType containing the following input argument fields: <table> <tr> <td>OopType</td> <td><i>clampSpec</i> The OOP of the Smalltalk ClampSpecification to be used, or OOP_NIL, if the traversal is to operate without clamping.</td> </tr> <tr> <td>int</td> <td><i>level</i> Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in the array of OOPs representing the objects to traverse. When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.</td> </tr> <tr> <td>int</td> <td><i>retrievalFlags</i> Flags to control object retrieval. The value of <i>retrievalFlags</i></td> </tr> </table>	OopType	<i>clampSpec</i> The OOP of the Smalltalk ClampSpecification to be used, or OOP_NIL, if the traversal is to operate without clamping.	int	<i>level</i> Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in the array of OOPs representing the objects to traverse. When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.	int	<i>retrievalFlags</i> Flags to control object retrieval. The value of <i>retrievalFlags</i>
OopType	<i>clampSpec</i> The OOP of the Smalltalk ClampSpecification to be used, or OOP_NIL, if the traversal is to operate without clamping.						
int	<i>level</i> Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in the array of OOPs representing the objects to traverse. When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.						
int	<i>retrievalFlags</i> Flags to control object retrieval. The value of <i>retrievalFlags</i>						

should be given by using the following GemBuilder mnemonics:
GCI_RETRIEVE_DEFAULT
GCI_RETRIEVE_EXPORT
GCI_CLEAR_EXPORT causes the traversal to clear the PureExportSet or the user action's export set before it adds any OOPs to the traverse buffer.

environmentId The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see "environmentId" on page 80.

Result Arguments

travArgs Pointer to an instance of **GciClampedTravArgsSType** containing the following result argument field:

ByteType * *travBuff*
 The buffer for the results of the traversal. The first element placed in the buffer is the *actualBufferSize*, an integer that indicates how many bytes were actually stored in the buffer by this function. The remainder of the traversal buffer consists of a series of object reports, each of which is of type **GciObjRepSType**.

Return Value

Returns FALSE if the traversal is not yet completed. You can then call **GciMoreTraversal** to proceed. Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal**.

Description

This function is like **GciPerformTrav**, except that it first does a **GciExecuteStr** instead of a **GciPerform**.

See Also

GciExecuteStr, page 153
GciMoreTraversal, page 234
GciPerformTrav, page 304
GciPerformTraverse, page 306

GciFetchByte

Fetch one byte from an indexed byte object.

Syntax

```
ByteType GciFetchByte(  
    OopType      theObject,  
    int64        atIndex);
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone byte object.
<i>atIndex</i>	The index into <i>theObject</i> of the element to be fetched. The index of the first element is 1.

Return Value

Returns the byte value at the specified index. In case of error, this function returns zero.

Description

This function fetches a single element from a byte object at the specified index, using structural access.

Example

```
void fetchByteExample(void)  
{  
    OopType oString = GciNewString("abc");  
  
    ByteType theChar = GciFetchByte(oString, 2);  
    if (theChar != 'b') {  
        printf("unexpected result %d \n", theChar);  
    }  
}
```

See Also

[GciFetchBytes_](#), page 162
[GciStoreByte](#), page 360
[GciStoreBytes](#), page 361

GciFetchBytes_

Fetch multiple bytes from an indexed byte object.

Syntax

```
int64 GciFetchBytes_(
    OopType      theObject,
    int64        startIndex,
    ByteType     theBytes[ ],
    int64        numBytes);
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone byte object.
<i>startIndex</i>	The index into <i>theObject</i> at which to begin fetching bytes. (The index of the first element is 1.) Note that if <i>startIndex</i> is 1 greater than the size of the object, this function returns a byte array of size 0, but no error is generated.
<i>numBytes</i>	The maximum number of bytes to return.

Result Arguments

<i>theBytes</i>	The array of fetched bytes
-----------------	----------------------------

Return Value

Returns the number of bytes fetched. (This may be less than *numBytes*, depending upon the size of *theObject*.) In case of error, this function returns zero.

Description

This function fetches multiple elements from a byte object starting at the specified index, using structural access. A common application of **GciFetchBytes_** would be to fetch a text string.

GciFetchBytes_ permits *theObject* to be a byte object with multiple bytes per character or digit, such as DoubleByteString, Float and LargeInteger. In this case, **GciFetchBytes_** provides automatic byte swizzling to client native byte order. For more about byte swizzling, see page 29. For MultiByteStrings, *startIndex* must be aligned on character boundaries and *numbytes* must be a multiple of the number of bytes per character; for numeric objects *startindex* must be one and *numBytes* the size of the numeric class.

Example

This example illustrates a C function that incrementally processes a GemStone String of arbitrary size, while using a limited amount of C memory space.

```
void displayByteObject(OopType oObject)
{
    enum { BUF_SIZE = 5000 };
    char displayBuff[BUF_SIZE];
```

```
    BoolType done = FALSE;
    int idx = 1;
    while (! done) {
        int64 numRet = GciFetchBytes_(oObject, idx,
        (ByteType*)displayBuff,
            BUF_SIZE - 1);
        if (numRet == 0) {
            done = TRUE; // hit end of object or error
            GciErrSType errInfo;
            if (GciErr(&errInfo)) {
                printf("unexpected error category "FMT_OID" number %d, %s\n",
                errInfo.category, errInfo.number, errInfo.message);
            }
        } else {
            displayBuff[numRet] = '\0';
            printf("%s\n", displayBuff);
            idx += numRet;
        }
    }
}
```

See Also

[GciFetchByte](#), page 161
[GciFetchOop](#), page 181
[GciStoreByte](#), page 360
[GciStoreBytes](#), page 361

GciFetchChars_

Fetch multiple ASCII characters from an indexed byte object.

Syntax

```
int64 GciFetchChars_(
    OopType      theObject,
    int64        startIndex,
    char *       cString,
    int64        maxSize );
```

Input Arguments

<i>theObject</i>	The OOP of a text object.
<i>startIndex</i>	The index of the first character to retrieve.
<i>maxSize</i>	Maximum number of characters to fetch.

Result Arguments

<i>cString</i>	Pointer to the location in which to store the returned string.
----------------	--

Return Value

Returns the number of characters fetched.

Description

Equivalent to **GciFetchBytes_**, except that it is assumed that *theObject* contains ASCII text. The bytes fetched are stored in memory starting at *cString*. At most *maxSize* - 1 bytes will be fetched from the object, and a `\0` character will be stored in memory following the bytes fetched. The function returns the number of characters fetched, excluding the null terminator character, which is equivalent to `strlen(cString)` if the object does not contain any null characters. If an error occurs, the function result is 0, and the contents of *cString* are undefined.

See Also

GciFetchBytes_, page 162

GciFetchClass

Fetch the class of an object.

Syntax

```

    OopType GciFetchClass(
        OopType          theObject );

```

Input Arguments

theObject The OOP of the specified object.

Return Value

Returns the OOP of the object's class. In case of error, this function returns OOP_NIL.

The GemBuilder include file `gcioop.ht` defines a C constant for each of the Smalltalk kernel classes. Those C constants are listed in Appendix A, "Reserved OOPs".

Description

The **GciFetchClass** function obtains the class of an object from GemStone. The GemBuilder session must be valid when **GciFetchClass** is called, unless theObject is an instance of one of the following classes: Boolean, Character, JisCharacter, SmallInteger, SmallDouble, or UndefinedObject.

Example

```

#include <stdlib.h>

void fetchClassExample(void)
{
    // random double to Oop conversion producing a Float or SmallFloat
    double rand = drand48() * 1.0e38 ;
    OopType oFltObj = GciFltToOop(rand);

    OopType oClass = GciFetchClass(oFltObj);
    const char* kind;
    if (oClass == OOP_CLASS_SMALL_DOUBLE) {
        kind = "SmallDouble";
    } else if (oClass == OOP_CLASS_FLOAT) {
        kind = "Float";
    } else {
        kind = "Unexpected";
    }
    printf("result is a %s, class oop = "FMT_OID"\n", kind, oClass);
}

```

See Also

[GciFetchNamedSize](#), page 174

[GciFetchObjImpl](#), page 180

[GciFetchSize_](#), page 191

[GciFetchVaryingSize_](#), page 199

GciFetchDateTime

Convert the contents of a DateTime object and place the results in a C structure.

Syntax

```
void GciFetchDateTime(  
    OopType          datetimeObj,  
    GciDateTimeSType * result );
```

Input Arguments

datetimeObj OOP of the object to fetch.

Result Arguments

result C pointer to the structure for the returned object.

Description

Fetches the contents of a DateTime object into the specified C result. Generates an error if *datetimeObj* is not an instance of DateTime. The value that *result* points to is undefined if an error occurs.

GciFetchDynamicIv

Fetch the OOP of one of an object's dynamic instance variables.

Syntax

```
OopType GciFetchDynamicIv(  
    OopType      theObject,  
    OopType      aSymbol );
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone object.
<i>aSymbol</i>	Specifies the dynamic instance variable to fetch.

Return Value

Returns the OOP of the specified dynamic instance variable. If no such dynamic instance variable exists in the object, this function returns OOP_NIL.

Description

This function fetches the contents of an object's dynamic instance variable, as specified by *aSymbol*.

See Also

[GciFetchDynamicIvs](#), page 169

[GciStoreDynamicIv](#), page 366

GciFetchDynamicIvs

Fetches the OOPs of one or more of an object's dynamic instance variables.

Syntax

```
int GciFetchDynamicIvs(  
    OopType          theObject,  
    OopType *        buf,  
    int              numOops);
```

Input Arguments

<i>theObject</i>	The OOP of the source GemStone object.
<i>numOops</i>	The maximum number of elements to return.

Result Arguments

<i>buf</i>	C pointer to the buffer that will contain the object's dynamic instance variables.
------------	--

Return Value

Returns the number of OOPs fetched. (This may be less than *numOops*, depending upon the size of *theObject*.)

Description

The number of dynamic instance variable pairs returned is (function result / 2). To obtain all dynamic instance variables in one call, use a buffer.

See Also

[GciFetchDynamicIv](#), page 168
[GciStoreDynamicIv](#), page 366

GciFetchNamedOop

Fetch the OOP of one of an object's named instance variables.

Syntax

```
OopType GciFetchNamedOop(  
    OopType      theObject,  
    int          atIndex);
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone object.
<i>atIndex</i>	The index into <i>theObject</i> 's named instance variables of the element to be fetched. The index of the first named instance variable is 1.

Return Value

Returns the OOP of the specified named instance variable. In case of error, this function returns OOP_NIL.

Description

This function fetches the contents of an object's named instance variable at the specified index, using structural access.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void fetchNamedOopExample(void)
{
    // C constants to match Smalltalk class definition
    enum { COMPONENT_OFF_PARTNUMBER = 1,
          COMPONENT_OFF_NAME       = 2,
          COMPONENT_OFF_COST       = 3 };

    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // error during execution or detect found nothing
        return;
    }

    // fetch the name instance variable of aComponent
    OopType oName = GciFetchNamedOop(aComponent, COMPONENT_OFF_NAME);
```

```
// fetch nameinstance variable without fixing its offset at C
compile time
int ivOffset = GciIvNameToIdx(GciFetchClass(aComponent), "name");
oName = GciFetchNamedOop(aComponent, ivOffset);
}
```

See Also

GciFetchNamedOops, page 172
GciFetchVaryingOop, page 195
GciFetchVaryingOops, page 197
GciIvNameToIdx, page 227
GciStoreIdxOop, page 367
GciStoreIdxOops, page 369
GciStoreNamedOop, page 371
GciStoreNamedOops, page 373

GciFetchNamedOops

Fetch the OOPs of one or more of an object's named instance variables.

Syntax

```
int GciFetchNamedOops(
    OopsType      theObject,
    int           startIndex,
    OopsType      theOops[],
    int           numOops);
```

Input Arguments

<i>theObject</i>	The OOP of the source GemStone object.
<i>startIndex</i>	The index into <i>theObject</i> 's named instance variables at which to begin fetching. (The index of the first named instance variable is 1.) Note that if <i>startIndex</i> is 1 greater than the number of the object's named instance variables, this function returns an array of size 0, but no error is generated.
<i>numOops</i>	The maximum number of elements to return.

Result Arguments

<i>theOops</i>	The array of fetched OOPs.
----------------	----------------------------

Return Value

Returns the number of OOPs fetched. (This may be less than *numOops*, depending upon the size of *theObject*.) In case of error, this function returns zero.

Description

This function uses structural access to fetch multiple values from an object's named instance variables, starting at the specified index.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void fetchNamedOops_example(void)
{
    // retrieve a random instance of class Component
    OopsType aComponent = GciExecuteStr(
    "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // execution error, or detect: found nothing
        return;
    }
}
```

```
// fetch name instance variables without knowing offset at C
compile time
int namedSize = GciFetchNamedSize(aComponent);
if (namedSize == 0) {
    // error during fetch
    return;
}
OopType *oBuffer = (OopType*) malloc( sizeof(OopType) * namedSize
);
if (oBuffer != NULL) {
    int numRet = GciFetchNamedOops(aComponent, 1, oBuffer,
namedSize);
    if (numRet != namedSize) {
        // error during fetch
    } else {
        // do something with contents of oBuffer
    }
    free(oBuffer);
} else {
    // malloc failure
}
}
```

See Also

GciFetchNamedOop, page 170
GciFetchVaryingOop, page 195
GciIvNameToIdx, page 227
GciStoreIdxOop, page 367
GciStoreNamedOop, page 371

GciFetchNamedSize

Fetch the number of named instance variables in an object.

Syntax

```
int GciFetchNamedSize(  
    OopType          theObject );
```

Input Arguments

theObject The OOP of the specified object.

Return Value

Returns the number of named instance variables in *theObject*. In case of error, this function returns zero.

Description

This function returns the number of named instance variables in a GemStone object. See the example for **GciFetchNamedOops** on page 172.

GciFetchNameOfClass

Fetch the class name object for a given class.

Syntax

```
oopType GciFetchNameOfClass(  
    oopType aClass );
```

Input Arguments

aClass The OOP of a class.

Return Value

The OOP of the class's name, or OOP_NIL if an error occurred.

Description

Given the OOP of a class, this function returns the object identifier of the String object that is the name of the class.

GciFetchNumEncodedOops

Obtain the size of an encoded OOP array.

Syntax

```
int GciFetchNumEncodedOops(  
    OopType *          encodedOopArray,  
    const int         numEncodedOops );
```

Input Arguments

encodedOopArray An OOP array that was encoded by a call to **GciEncodeOopArray**.

Result Arguments

numEncodedOops The number of OOPs in *encodedOopArray*.

Return Value

Returns the number of OOPs that will be decoded by a call to **GciDecodeOopArray**.

Description

This function returns the total number of OOPs in an OOP array that was encoded by a call to **GciEncodeOopArray**.

See Also

GciDecodeOopArray, page 133
GciEnableFreeOopEncoding, page 143
GciEncodeOopArray, page 146
GciGetFreeOopsEncoded, page 208

GciFetchNumSharedCounters

Obtain the number of shared counters available on the shared page cache used by this session.

Syntax

```
int GciFetchNumSharedCounters();
```

Return Value

Returns the number of shared counters available on the shared page cache used by this session, or -1 if the session is not logged in.

Description

This function returns the total number of shared counters available on the shared page cache used by this session.

Not supported for remote GCI interfaces.

See Also

[GciDecSharedCounter](#), page 134
[GciIncSharedCounter](#), page 214
[GciSetSharedCounter](#), page 353
[GciReadSharedCounter](#), page 316
[GciReadSharedCounterNoLock](#), page 317
[GciFetchSharedCounterValuesNoLock](#), page 190

GciFetchObjectInfo

Fetch information and values from an object.

Syntax

```
BoolType GciFetchObjectInfo(
    OopType          theObject,
    GciFetchObjInfoArgsSType *args );
```

Input Arguments

<i>theObject</i>	OOP of any object with byte, pointer, or NSC format.
<i>args</i>	Pointer to an instance of GciFetchObjInfoArgsSType with the following input argument fields:
int64	<i>startIndex</i> The offset in the object at which to start fetching, using GciFetchOops or GciFetchBytes_ semantics. <i>startIndex</i> is ignored if <i>bufSize</i> == 0 or <i>buffer</i> == NULL.
int64	<i>bufSize</i> The size in bytes of the buffer, maximum number of elements fetched for a byte object. For an OOP object, the maximum number of elements fetched for an OOP object will be <i>bufSize</i> /8. If greater than zero, and if a Float or BinaryFloat is being fetched, it must be large enough to fetch the complete object.
int	<i>retrievalFlags</i> If (<i>retrievalFlags</i> & GCI_RETRIEVE_EXPORT) != 0 then if <i>theObject</i> is non-special, <i>theObject</i> is automatically added to the PureExportSet or the user action's export set (see GciSaveObjs on page 338).

Result Arguments

<i>args</i>	Pointer to an instance of GciFetchObjInfoArgsSType with the following result argument fields:
GciObjInfoSType * <i>info</i>	Pointer to an instance of GciObjInfoSType ; may be NULL.
ByteType * <i>buffer</i>	Pointer to an area where byte or OOP values will be returned; may be NULL.
int64	<i>numReturned</i> Number of logical elements (bytes or OOPs) returned in <i>buffer</i> . Remember that the size of (OopType) is 8 bytes.

If either *info* or *buffer* is NULL, that portion of the result is not filled in.

Return Value

TRUE if successful, FALSE if an error occurs.

Description

GciFetchObjectInfo fetches information and values from an object starting at the specified index using structural access. If either *info* or *buffer* is NULL, then that part of the result is not filled in. If *numReturned* is NULL, then *buffer* will not be filled in.

If *theObject* is a byte object with multiple bytes per character or digit, such as `DoubleByteString`, `LargeInteger`, or `Float`, the results in *args->buffer* will be automatically byte swizzled to client native byte order. For more about byte swizzling, see page 25.

For `MultiByteStrings`, *args->startIndex* must be aligned on a character boundary and *args->bufSize* must be a multiple of the number of bytes per character in the string. For numeric objects, *args->startIndex* must be one and *args->bufSize* must be the size of the numeric class.

See Also

GciFetchOops, page 183
GciFetchBytes_, page 162
GciFetchOop, page 181
GciSaveObjs, page 338

GciFetchObjImpl

Fetch the implementation of an object.

Syntax

```
int GciFetchObjImpl(  
    OopType          theObject );
```

Input Arguments

theObject The OOP of the specified object.

Return Value

Returns an integer representing the implementation type of *theObject* (0=pointer, 1=byte, 2=NSC, or 3=special). In case of error, the return value is undefined.

Description

This function obtains the implementation of an object (pointer, byte, NSC, special) from GemStone. For more information about implementation types, see "Direct Access to Metadata" on page 29.

See Also

[GciFetchClass](#), page 165
[GciFetchNamedSize](#), page 174
[GciFetchSize_](#), page 191
[GciFetchVaryingSize_](#), page 199

GciFetchOop

Fetch the OOP of one instance variable of an object.

Syntax

```
OopType GciFetchOop(
    OopType      theObject,
    int64        atIndex);
```

Input Arguments

<i>theObject</i>	The OOP of the source object.
<i>atIndex</i>	The index into <i>theObject</i> of the OOP to be fetched. The index of the first OOP is 1.

Return Value

Returns the OOP at the specified index of the source object. In case of error, this function returns OOP_NIL.

Description

This function fetches the OOP of a single instance variable from any object at the specified index, using structural access. It does not distinguish between named and unnamed instance variables. Indices are based at the beginning of the object's array of instance variables. In that array, any existing named instance variables are followed by any existing unnamed instance variables.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void fetchOop_example(void)
{
    // C constant to match Smalltalk class definition
    enum { COMPONENT_OFF_NAME = 2 };

    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // error during execution, select: found nothing
        return ;
    }

    // Two ways to fetch the name instance variable of aComponent */
    OopType oName = GciFetchOop(aComponent, COMPONENT_OFF_NAME);
    oName = GciFetchNamedOop(aComponent, COMPONENT_OFF_NAME);
```

```
// Fetch the 3rd element of aComponent's partsList,  
// without knowing exactly how many named instance variables exist.  
int namedSize = GciFetchNamedSize(aComponent);  
if (namedSize == 0) {  
    // error during fetch  
    return ;  
}  
OopType aSubComponent = GciFetchOop(aComponent, namedSize + 3);  
  
// alternate way to Fetch the 3rd element of aComponent's partsList  
aSubComponent = GciFetchVaryingOop(aComponent, 3);  
}
```

GciFetchOops, page 183

GciStoreOop, page 375

GciStoreOops, page 377

GciFetchOops

Fetch the OOPs of one or more instance variables of an object.

Syntax

```
int GciFetchOops(
    OopType          theObject,
    int64            startIndex,
    OopType          theOops[],
    int              numOops);
```

Input Arguments

<i>theObject</i>	The OOP of the source object.
<i>startIndex</i>	The index into <i>theObject</i> at which to begin fetching OOPs. The index of the first OOP is 1. If <i>startIndex</i> is 1 greater than the size of the object, this function returns an array of size 0, but no error is generated.
<i>numOops</i>	The maximum number of OOPs to return.

Result Arguments

<i>theOops</i>	The array of fetched OOPs.
----------------	----------------------------

Return Value

Returns the number of OOPs fetched. (This may be less than *numOops*, depending upon the size of *theObject*.) In case of error, this function returns zero.

Description

This function fetches the OOPs of multiple instance variables from any object starting at the specified index, using structural access. It does not distinguish between named and unnamed instance variables. Indices are based at the beginning of the object's array of instance variables. In that array, any existing named instance variables are followed by any existing unnamed instance variables.

Example

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void fetchOops_example(void)
{
    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // error during execution, or detect: found nothing
        return ;
    }
}
```

```
    }

    enum { BUF_SIZE = 60 };
    OopType oBuf[BUF_SIZE];

    int namedSize = GciFetchNamedSize(aComponent);
    if (namedSize == 0 || namedSize > 50) {
        // error during fetch, or too many named instVars for buffer
        return;
    }

    // Two ways to fetch first 5 elements of aComponent's partsList
    GciFetchOops(aComponent, namedSize + 1, oBuf, 5);
    GciFetchVaryingOops(aComponent, 1, oBuf, 5);

    // Fetch the named instance variables PLUS
    // the first 5 elements of partsList
    GciFetchOops(aComponent, 1, oBuf, namedSize + 5);
    // oBuf[0..namedSize-1] are named instVar values,
    // oBuf[namedSize] is first varying instVar value
}
```

See Also

GciFetchOop, page 181
GciFetchVaryingOop, page 195
GciStoreOop, page 375
GciStoreOops, page 377

GciFetchPaths

Fetch selected multiple OOPs from an object tree.

Syntax

```
BoolType GciFetchPaths(
    const OopType    theOops [ ],
    int              numOops,
    const int        paths [ ],
    const int        pathSizes [ ],
    int              numPaths,
    OopType          results [ ] );
```

Input Arguments

<i>theOops</i>	A collection of OOPs from which you want to fetch.
<i>numOops</i>	The size of <i>theOops</i> .
<i>paths</i>	An array of integers. This one-dimensional array contains the elements of all constituent paths, laid end to end.
<i>pathSizes</i>	An array of integers. Each element of this array is the length of the corresponding path in the <i>paths</i> array (that is, the number of elements in each constituent path).
<i>numPaths</i>	The number of paths in the <i>paths</i> array. This should be the same as the number of integers in the <i>pathSizes</i> array.

Result Arguments

<i>results</i>	An array containing the OOPs that were fetched.
----------------	---

Return Value

Returns TRUE if all desired objects were successfully fetched. Returns FALSE if the fetch on any path fails for any reason.

Description

This function allows you to fetch multiple OOPs from selected positions in an object tree with a single GemBuilder call, importing only the desired information from the database.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

Each path in the *paths* array is itself an array of integers. Those integers are offsets that specify a path from which to fetch objects. In each path, a positive integer *x* refers to an offset within an

object's named instance variables (see **GciFetchNamedOop**), while a negative integer *-x* refers to an offset within an object's indexed instance variables (see **GciFetchVaryingOop**).

From each object in *theOops*, this function fetches the object pointed to by each element of the paths array, and stores the fetched object into the results array. The *results* array contains (*numOops* * *numPaths*) elements, stored in the following order:

```
[ 0, 0 ] .. [ 0, numPaths-1 ] ..
[ 1, 0 ] .. [ 1, numPaths-1 ] ..
[ numOops-1, 0 ] .. [ numOops-1, numPaths-1 ]
```

That is, all paths are first applied in order to the first element of *theOops*. This step is repeated for each subsequent object, until all paths have been applied to all elements of *theOops*. The result for object *i* and path *j* is represented as:

```
results[ ((i-1) * numPaths) + (j-1) ]
```

If the fetch on any path fails for any reason, the result of that fetch is reported in the results array as OOP_ILLEGAL. Because some path-fetching errors do not necessarily invalidate the remainder of the information fetched, the system will then attempt to continue its fetching with the remaining paths and objects.

This ability to complete a fetching sequence despite errors means that your application won't be slowed by a round-trip to GemStone on each fetch to check for errors. Instead, after a fetch is complete, you can cycle through the result and deal selectively at that time with any errors you find.

The appropriate response to an error in path fetching depends both upon the error itself and on your application. Here are some of the reasons why a fetch might not succeed:

- The user had no read authorization for some object in the path. The seriousness of this depends on your application. In some applications, you may simply wish to ignore the inaccessible data.
- The path was invalid for the object to which it was applied. This can happen if the object from which you're fetching is not of the correct class, or if the path itself is faulty for the class of the object.
- The path was valid but simply not filled out for the object being processed. This would be the case, for example, if you attempted to access *address.zip* when an Employee's Address instance variable contained only *nil*. This is probably the most common path fetching error, and may require only that the application program detect the condition and display some suitable indication to the user that a field is not yet filled in with meaningful data.

Example

Example 1: Calling sequence for a single object and a single path

```
void fetchPath1(void)
{
    enum { path_size = 5 };
    int    aPath[path_size]; /* the path itself */
    int    aSize = path_size; /* the size of the path */

    for (int j = 0; j < path_size; j++) {
        aPath[j] = j + 1; // arbitrary offsets
    }
}
```

```

    OopType anOop; // the OOP to use as the root of the path
    anOop = GciExecuteStr("AllComponents detect[:i|i partNumber =
1234]", OOP_NIL);
    if (anOop == OOP_NIL) {
        return; // error during resolve
    }

    OopType result;
    GciFetchPaths(&anOop, 1, aPath, &aSize, 1, &result);
}

```

Example 2: Calling sequence for multiple objects with a single path

```

void fetchPath2(void)
{
    OopType coll = GciResolveSymbol("AllComponents", OOP_NIL);
    if (coll == OOP_NIL) {
        return ; // error during resolve
    }
    enum { num_roots = 3 ,
          path_size = 5 };
    OopType oops[num_roots];
    int numRet = GciFetchVaryingOops(coll, 1, oops, num_roots);
    if (numRet != num_roots) {
        return; // error during fetch or collection too small
    }

    int aPath[path_size];
    int aSize = path_size;
    for (int j = 0; j < path_size; j++) {
        aPath[j] = 1; // arbitrary offsets
    }
    OopType results[num_roots];
    GciFetchPaths(oops, num_roots, aPath, &aSize, 1, results);
}

```

Example 3: Calling sequence for a single object with multiple paths

```

void fetchPath3(void)
{
    OopType anOop; // the OOP to use as the root of the path
    anOop = GciExecuteStr("AllComponents detect[:i|i partNumber =
1234]", OOP_NIL);
    if (anOop == OOP_NIL) {
        return; // error during execution
    }

    enum { num_paths = 10,
          path_size = 5 };
}

```

```
int pathSizes[num_paths];
int paths[path_size * num_paths ];
int idx = 0;
for (int j = 0; j < num_paths; j++) {
    for (int k = 0; k < path_size; k++) {
        paths[idx++] = k + 1; // arbitrary offset
    }
}
OopType results[num_paths];
GciFetchPaths (&anOop, 1, paths, pathSizes, num_paths, results);
}
```

Example 4: Calling sequence for multiple objects with multiple paths

```
void fetchPath4(void)
{
    OopType coll = GciResolveSymbol("AllComponents", OOP_NIL);
    if (coll == OOP_NIL) {
        return ; // error during resolve
    }

    enum { num_roots = 10,
          num_paths = 3,
          path_size = 5 };

    OopType oops[num_roots];
    int numRet = GciFetchVaryingOops(coll, 1, oops, num_roots);
    if (numRet != num_roots) {
        return; // error during fetch or collection too small
    }

    int pathSizes[num_paths];
    int paths[path_size * num_paths ];
    int idx = 0;
    for (int j = 0; j < num_paths; j++) {
        for (int k = 0; k < path_size; k++) {
            paths[idx++] = k + 1; // arbitrary offset
        }
    }

    OopType results[num_roots * num_paths];
    GciFetchPaths(oops, num_roots, paths, pathSizes, num_paths,
results);
}
```

Example 5: Integrated Code

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for `GciFetchVaryingOop` on page 195.

```
void fetchPath5(void)
{
    // retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
    "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        return; // error in execute, or detect: found nothing
    }

    // fetch name instVar of 5th element of aComponent's partsList */
    enum { path_size = 2 };
    int path[path_size];
    path[0] = -5; // 5th varying instVar
    path[1] = GciIvNameToIdx(GciFetchClass(aComponent), "name");
    int pathSizes = path_size;

    OopType oName;
    GciFetchPaths(&aComponent, 1, path, &pathSizes, 1, &oName);
};
```

See Also

`GciStorePaths`, page 379

GciFetchSharedCounterValuesNoLock

Fetch the value of multiple shared counters without locking them.

Syntax

```
int GciFetchSharedCounterValuesNoLock(  
    int                startIndex,  
    int64_t            buffer[],  
    size_t*            maxReturn);
```

Input Arguments

<i>startIndex</i>	The offset into the shared counters array of the first shared counter value to fetch.
<i>maxReturn</i>	Pointer to a value that indicates the maximum number of shared counters to fetch.

Result Arguments

<i>buffer</i>	Pointer to a buffer where the shared counter values will be stored. The buffer must be at least $8 * \textit{maxReturn}$ bytes and the address must be aligned on an 8-byte boundary.
---------------	---

Return Value

Returns an int indicating the number of shared counter values successfully stored in the buffer. Returns -1 if a bad argument is detected.

Description

Fetch the values of multiple shared counters in a single call, without locking any of them. The values of the *maxReturn* count of shared counters starting at the offset indicated by *counterIdx* (0-based) are put into the buffer *buffer*. *buffer* must be large enough to accommodate *maxReturn* 8-byte values, and be aligned on an 8-byte boundary.

Not supported for remote GCI interfaces.

See Also

- [GciFetchNumSharedCounters](#), page 177
- [GciDecSharedCounter](#), page 134
- [GciIncSharedCounter](#), page 214
- [GciSetSharedCounter](#), page 353
- [GciReadSharedCounter](#), page 316
- [GciReadSharedCounterNoLock](#), page 317

GciFetchSize_

Fetch the size of an object.

Syntax

```
int64 GciFetchSize_(
    OopType          theObject );
```

Input Arguments

theObject The OOP of the specified object.

Return Value

Returns the size of *theObject*. In case of error, this function returns zero.

Description

This function obtains the size of an object from GemStone.

The result of this function depends on the object's implementation (see **GciFetchObjImpl**). For byte objects, this function returns the number of bytes in the object. (For Strings, this is the number of Characters in the String; for Floats, the size is 23.) For pointer objects, this function returns the number of named instance variables (**GciFetchNamedSize**) plus the number of indexed instance variables, if any (**GciFetchVaryingSize_**). For NSC objects, this function returns the cardinality of the collection. For special objects, the size is always zero.

This differs somewhat from the result of executing the Smalltalk method `Object>>size`, as shown in Table 6.10:

Table 6.10 Differences in Reported Object Size

Implementation	Object>>size (Smalltalk)	GciFetchSize_
0=Pointer	Number of indexed elements in the object (0 if not indexed)	Number of indexed elements PLUS number of named instance variables
1=Byte	Number of indexed elements in the object	Same as Smalltalk message "size"
2=NSC	Number of elements in the object	Same as Smalltalk message "size"
3=Special	0	0

Example

```
void fetchSize_example(void)
{
    const char* str = "abcdef";
    OopType oString = GciNewString(str);
```

```
int64 itsSize = GciFetchSize_(oString);
if (itsSize != (int64)strlen(str)) {
    printf("error during fetch size\n");
}
}
```

See Also

[GciFetchClass](#), page 165
[GciFetchNamedSize](#), page 174
[GciFetchObjImpl](#), page 180
[GciFetchVaryingOop](#), page 195

GciFetchUtf8Bytes_

Encode a String, MultiByteString, or Utf8 as UTF-8, and fetch the bytes of the encoded result.

Syntax

```
int64 GciFetchUtf8Bytes_(
    OopType          oopOfAString,
    int64            startIndex,
    ByteType *       theBytes,
    int64            numBytes,
    OopType *        utf8String,
    int              flags);
```

Input Arguments

<i>oopOfAString</i>	The OOP of the GemStone String, MultiByteString or Utf8.
<i>startIndex</i>	The index into <i>aString</i> at which to begin encoding bytes. The index of the first element is 1. Note that if <i>startIndex</i> is 1 greater than the size of the object, this function returns a byte array of size 0, but no error is generated.
<i>numBytes</i>	The maximum number of bytes to return.
<i>flags</i>	If <i>flags</i> = 0, generate an error on illegal codePoints in the input. If <i>flags</i> = 1, substitute '.' for any illegal codePoints in input string while performing the encoding. If <i>flags</i> = 2, put a generated description in <i>*theBytes</i> instead of signalling an error.

Result Arguments

<i>theBytes</i>	The array of encoded bytes
<i>utf8String</i>	Pointer to the OOP of the encoded string.

Return Value

Returns the number of bytes fetched. (This may be less than *numBytes*, depending upon the size of *theObject*.) In case of error, this function returns zero.

Description

This function encodes a kind of String, MultiByteString, or Utf8 into UTF-8. The encoded bytes are placed in the buffer *theBytes*, and the OOP of the encoded object is placed at *utf8String*.

The OOP of *utf8String* is also placed in the ExportSet. The caller must pass the OOP of *utf8String* to **GciReleaseOops** after fetching all bytes.

If all characters in *aString* are < 128, or if the class of *aString* is Utf8, then the the behavior is the same as **GciFetchBytes_**. No encoding is done; the bytes of *aString* are place in *theBytes*. *utf8String* is the same as *aString*, and is not added to the ExportSet.

See Also

[GciFetchBytes_](#), page 162
[GciNewUtf8String](#), page 277

[GciReleaseOops](#), page 323

GciFetchVaryingOop

Fetch the OOP of one unnamed instance variable from an indexable pointer object or NSC.

Syntax

```
OopType GciFetchVaryingOop(
    OopType      theObject,
    int64        atIndex);
```

Input Arguments

<i>theObject</i>	The OOP of the pointer object or NSC.
<i>atIndex</i>	The index of the OOP to be fetched. The index of the first unnamed instance variable's OOP is 1.

Return Value

Returns the OOP of the unnamed instance variable at index *atIndex*. In case of error, this function returns OOP_NIL.

Description

This function fetches the OOP of a single unnamed instance variable at the specified index, using structural access. The numerical index of any unordered variable of an NSC can change whenever the NSC is modified.

Example

In the following example, assume that you've executed the following Smalltalk code to define the class Component and to populate the set AllComponents:

```
run
" define the class Component and compile accessor methods"
| cls |
cls := Array subclass: #Component
    instVarNames: (#partNumber #name #cost
                  "varying instVars form the partsList")
    classVars: #()
    classInstVars: #()
    poolDictionaries: #()
    inDictionary: UserGlobals.
cls compileAccessingMethodsFor: cls instVarNames .
^ cls
%
run
"create and populate the set of all Components"
| allC |
allC := IdentitySet new .
UserGlobals at: #AllComponents put: allC .
```

```

1 to: 100 do:[:j || aComp |
  aComp := Component new .
  aComp partNumber: 1200 + j .
  aComp name: 'part' + j asString .
  aComp cost: j asFloat .
  allC add: aComp .
] .
^ allC size
%
run
"build a parts list for each part."
| allC idx |
allC := Array withAll: AllComponents .
idx := 1 .
AllComponents do:[ :aComp | | list |
  list := Array new: (idx \\ 10 ) . "list size varies from 0 to 9"
  idx > 75 ifTrue:[ idx := 1 ].
  1 to: list size do:[:k |
    list at: k put: (allC at: idx + (k * 2)).
  ].
  aComp addAll: list .
  idx := idx + 1.
].
%

```

Now execute this C code:

```

OopType fetchVaryingOopExample(void)
{
  /* retrieve a random instance of class Component */
  OopType aComponent = GciExecuteStr(
    "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
  if (aComponent == OOP_NIL) {
    return OOP_NIL; /* error in execute, or detect: found nothing
  }

  /* fetch 3rd element of aComponent's parts list */
  OopType aSubComponent = GciFetchVaryingOop(aComponent, 3);
  return aSubComponent;
}

```

See Also

[GciFetchNamedOop](#), page 170
[GciFetchNamedOops](#), page 172
[GciFetchVaryingOops](#), page 197
[GciStoreIdxOop](#), page 367
[GciStoreIdxOops](#), page 369
[GciStoreNamedOop](#), page 371
[GciStoreNamedOops](#), page 373

GciFetchVaryingOops

Fetch the OOPs of one or more unnamed instance variables from an indexable pointer object or NSC.

Syntax

```
int GciFetchVaryingOops(
    OopType      theObject,
    int64        startIndex,
    OopType      theOops[],
    int          numOops);
```

Input Arguments

<i>theObject</i>	The OOP of the pointer object or NSC.
<i>startIndex</i>	The index of the first OOP to be fetched. The index of the first unnamed instance variable's OOP is 1. Note that if <i>startIndex</i> is 1 greater than the number of <i>theObject</i> 's unnamed instance variables, this function returns an array of size 0, but no error is generated.
<i>numOops</i>	Maximum number of elements to return.

Result Arguments

<i>theOops</i>	The array of fetched OOPs.
----------------	----------------------------

Return Value

Returns the number of OOPs fetched. (This may be less than *numOops*, depending upon the size of *theObject*.) In case of error, this function returns zero.

Description

This function fetches the OOPs of multiple unnamed instance variables beginning at the specified index, using structural access. The numerical index of any unordered variable of an NSC can change whenever the NSC is modified.

Example

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for `GciFetchVaryingOop` on page 195.

```
int fetchVaryingOopsExample(void)
{
    // retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        return -1; // error in execute, or detect: found nothing
    }
}
```

```
    /* fetch the up to the first 5 elements of aComponent's parts list
    */
    enum { num_oops = 5 };
    OopType oBuf[num_oops];

    int numRet = GciFetchVaryingOops(aComponent, 1, oBuf, num_oops);
    // at this point we have    0 <= numRet <= 5
    return numRet;
}
```

See Also

GciFetchNamedOop, page 170
GciFetchNamedOops, page 172
GciFetchVaryingOop, page 195
GciStoreIdxOop, page 367
GciStoreIdxOops, page 369
GciStoreNamedOop, page 371
GciStoreNamedOops, page 373

GciFetchVaryingSize_

Fetch the number of unnamed instance variables in a pointer object or NSC.

Syntax

```
int64 GciFetchVaryingSize_(
    OopType      theObject );
```

Input Arguments

theObject The OOP of the specified object.

Return Value

Returns the number of unnamed instance variables in *theObject*. In case of error, this function returns zero.

Description

This function obtains from GemStone the number of indexed variables in an indexable object or the number of unordered variables in an NSC.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
int64 fetchVaryingSizeExample(void)
{
    // retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        return -1; // error in execute, or detect: found nothing
    }

    /* fetch the size of aComponent's partsList */
    int64 theSize = GciFetchVaryingSize_(aComponent);
    return theSize;
}
```

See Also

GciFetchClass, page 165

GciFetchNamedSize, page 174

GciFetchObjImpl, page 180

GciFetchSize_, page 191

GciSetVaryingSize, page 355

GciFindObjRep

Fetch an object report in a traversal buffer.

Syntax

```
GciObjRepHdrSType * GciFindObjRep(
    GciTravBufType *   travBuff,
    OopType            theObject );
```

Input Arguments

<i>travBuff</i>	A traversal buffer returned by a call to GciTraverseObjs .
<i>theObject</i>	The OOP of the object to find.

Return Value

Returns a pointer to an object report within the traversal buffer. In case of error, this function returns NULL.

Description

This function locates an object report within a traversal buffer that was previously returned by **GciTraverseObjs**. If the report is not found within the buffer, this function generates the error `GCL_ERR_TRAV_OBJ_NOT_FOUND`.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

```
GciObjRepHdrSType* findObjRepExample(GciTravBufType *buf, OopType
objId)
{
    GciObjRepHdrSType *theReport = GciFindObjRep(buf, objId);
    if (theReport == NULL) {
        GciErrSType errInfo;
        if (GciErr(&errInfo)) {
            printf("error category "FMT_OID" number %d, %s\n",
                errInfo.category, errInfo.number, errInfo.message);
        }
    }
    return theReport;
}
```


See Also

GciMoreTraversal, page 234
GciObjRepSize_, page 282

GciTraverseObjs, page 408

GciFloatKind

Obtain the float kind corresponding to a C double value.

Syntax

```
GciFloatKindEType GciFloatKind(  
    double          aReal );
```

Input Arguments

aReal A floating point value.

Return Value

Returns the type of GemStone Float object that corresponds to the C value.

Description

This function obtains the kind of GemStone Float object that corresponds to the C floating point value *aReal*.

See Also

[GciFltToOop](#), page 203

[GciOopToFlt](#), page 294

GciFltToOop

Convert a C double value to a SmallDouble or Float object.

Syntax

```
OopType GciFltToOop(  
    double          aReal);
```

Input Arguments

aReal The floating point value to be translated into an object.

Return Value

Returns the OOP of the GemStone SmallDouble or Float object that corresponds to the C value. In case of error, this function returns OOP_NIL.

Description

This function translates a C double precision value into the equivalent GemStone Float object.

Example

```
#include <stdlib.h>  
  
void fltToOopExample(void)  
{  
    // random double to Oop conversion producing a Float or SmallFloat  
    double rand = drand48() * 1.0e38 ;  
    OopType oFltObj = GciFltToOop(rand);  
  
    OopType oClass = GciFetchClass(oFltObj);  
    const char* kind;  
    if (oClass == OOP_CLASS_SMALL_DOUBLE) {  
        kind = "SmallDouble";  
    } else if (oClass == OOP_CLASS_FLOAT) {  
        kind = "Float";  
    } else {  
        kind = "Unexpected";  
    }  
    printf("result is a %s, class oop = "FMT_OID"\n", kind, oClass);  
}
```

See Also

[GciOopToFlt](#), page 294

[GciGetFreeOopsEncoded](#), page 208

GciGetFreeOop

Allocate an OOP.

Syntax

```
OopType GciGetFreeOop( )
```

Return Value

Returns an unused object identifier (OOP).

You cannot use the result of **GciGetFreeOop** to create a Symbol object.

Description

Allocates an object identifier without creating an object.

The object identifier returned from this function remains allocated to the Gci session until the session calls **GciLogout** or until the identifier is used as an argument to a function call.

If an object identifier returned from **GciGetFreeOop** is used as a value in a **GciStore...** call before it is used as the *objld* argument of a **GciCreate...** call, then an unresolved forward reference is created in object memory. This is a reference to an object that does not yet exist. This forward reference must be satisfied by using the identifier as the *objld* argument to a **GciCreate...** call before a **GciCommit** can be successfully executed.

If **GciCommit** is attempted prior to satisfying all unresolved forward references, an error is generated and **GciCommit** returns FALSE. In this case, **GciCreate** can be used to satisfy the forward references and **GciCommit** can be attempted again. **GciAbort** removes all unsatisfied forward references from the session's object space, just as it removes any other uncommitted modifications.

As long as it remains an unresolved forward reference, the identifier returned by **GciGetFreeOop** can be used only as a parameter to the following function calls, under the given restrictions:

- As the *objID* of the object to be created
 - GciCreateByteObj
- As the *objID* of the object to be created, or as an element of the value buffer
 - GciCreateOopObj
- As an element of the value buffer only
 - GciStoreOop
 - GciStoreOops
 - GciStoreIdxOop
 - GciStoreIdxOops
 - GciStoreNamedOop
 - GciStoreNamedOops
 - GciStoreTrav
 - GciAppendOops
 - GciAddOopToNsc
 - GciAddOopsToNsc
 - GciNewOopUsingObjRep
- As an element of *newValues* only
 - GciStorePaths

See Also

GciCreateByteObj, page 122
GciCreateOopObj, page 124
GciGetFreeOops, page 206
GciGetFreeOopsEncoded, page 208

GciGetFreeOops

Allocate multiple OOPs.

Syntax

```
void GciGetFreeOops(  
    int count,  
    OopsType * resultOops );
```

Input Arguments

count The number of OOPs to allocate.

Result Arguments

resultOops An array to hold the returned OOPs.

Return Value

Returns an unused object identifier (OOP).

Description

Allocates object identifiers without creating objects.

If an object identifier returned from **GciGetFreeOops** is used as a value in a **GciStore...** call before it is used as the *objId* argument of a **GciCreate...** call, then an unresolved forward reference is created in object memory. This is a reference to an object that does not yet exist. This forward reference must be satisfied by using the identifier as the *objId* argument to a **GciCreate...** call before a **GciCommit** can be successfully executed.

If **GciCommit** is attempted prior to satisfying all unresolved forward references, an error is generated and **GciCommit** returns false. In this case, **GciCreate** can be used to satisfy the forward references and **GciCommit** can be attempted again. **GciAbort** removes all unsatisfied forward references from the session's object space, just as it removes any other uncommitted modifications.

As long as it remains an unresolved forward reference, the identifier returned by **GciGetFreeOops** can be used only as a parameter to the following function calls, under the given restrictions:

- As the *objID* of the object to be created
 - `GciCreateByteObj`
- As the *objID* of the object to be created, or as an element of the value buffer
 - `GciCreateOopObj`
- As an element of the value buffer, only
 - `GciStoreOop`
 - `GciStoreOops`
 - `GciStoreIdxOop`
 - `GciStoreIdxOops`
 - `GciStoreNamedOop`
 - `GciStoreNamedOops`
 - `GciStoreTrav`
 - `GciAppendOops`
 - `GciAddOopToNsc`
 - `GciAddOopsToNsc`
 - `GciNewOopUsingObjRep`
- As an element of *newValues*, only
 - `GciStorePaths`

See Also

GciCreateByteObj, page 122
GciCreateOopObj, page 124
GciGetFreeOop, page 204
GciGetFreeOopsEncoded, page 208

GciGetFreeOopsEncoded

Allocate multiple OOPs.

Syntax

```
void GciGetFreeOopsEncoded(  
    int *                count,  
    OopType *           encodedOops );
```

Input Arguments

<i>count</i>	The number of OOPs to allocate.
<i>encodedOops</i>	A pointer to memory for holding encoded oops. Must be large enough to hold at least the input value of <i>count</i> .

Result Arguments

<i>count</i>	The number of OOPs returned in the encoded OOP array.
<i>encodedOops</i>	An array to hold the returned encoded oops. Must be large enough to hold at least the input value of <i>count</i> .

Description

This function is identical to **GciGetFreeOops**, except that it returns OOPs in an encoded array that is more compact for less network I/O. Before the OOPs can be used, the encoded array must be decoded by calling **GciDecodeOopArray()**.

See Also

- GciGetFreeOop**, page 204
- GciGetFreeOops**, page 206
- GciFetchNumEncodedOops**, page 176
- GciEnableFreeOopEncoding**, page 143
- GciEncodeOopArray**, page 146
- GciDecodeOopArray**, page 133

GciGetSessionId

Find the ID number of the current user session.

Syntax

```
GciSessionIdType GciGetSessionId()
```

Return Value

Returns the session ID currently being used for communication with GemStone. Returns `GCI_INVALID_SESSION_ID` if there is no session ID (that is, if the application is not logged in).

Description

This function obtains the unique session ID number that identifies the current user session to GemStone. An application can have more than one active session, but only one current session.

The ID numbers assigned to your application's sessions are unique within your application, but bear no meaningful relationship to the session IDs assigned to other GemStone applications that may be executing at the same time or accessing the same database.

Example

```
void getSessionExample(const char* userId, const char* password)
{
    if (GciLogin(userId, password)) {
        GciSessionIdType sessId = GciGetSessionId();
        printf("sessionId is %d \n", sessId);
    }
    GciLogout();
    GciSessionIdType sessId = GciGetSessionId();
    if (sessId != GCI_INVALID_SESSION_ID) {
        printf("unexpected sessionId %d after logout \n", sessId);
    }
}
```

See Also

[GciLogin](#), page 229

[GciSetSessionId](#), page 351

GciHardBreak

Interrupt GemStone and abort the current transaction.

Syntax

```
void GciHardBreak()
```

Description

GciHardBreak sends a hard break to the current user session (set by the last **GciLogin** or **GciSetSessionId**), which interrupts Smalltalk execution.

All GemBuilder functions can recognize a hard break, so the users of your application can terminate Smalltalk execution. For example, if your application sends a message to an object (via **GciPerform**), and for some reason the invoked Smalltalk method enters an infinite loop, the user can interrupt the application. **GciHardBreak** has no effect if called from within a User Action.

In order for GemBuilder functions in your program to recognize interrupts, your program will need a signal handler that can call the functions **GciSoftBreak** and **GciHardBreak**. Since GemBuilder does not relinquish control to an application until it has finished its processing, soft and hard breaks must be initiated from a signal handler.

If GemStone is executing when it receives the break, it replies with the error message `RT_ERR_HARD_BREAK`. Otherwise, it ignores the break.

If GemStone is executing any of the following methods of class Repository, then a hard break terminates the entire session, not just Smalltalk execution:

```
fullBackupTo:  
restoreFromBackup(s):  
markForCollection  
objectAudit  
auditWithLimit:  
repairWithLimit:  
pagesWithPercentFree
```

See Also

GciSoftBreak, page 357

GciHiddenSetIncludesOop

Determines whether the given OOP is present in the specified hidden set.

Syntax

```
BoolType GciHiddenSetIncludesOop(  
    OopType      theOop,  
    int          hiddenSetId );
```

Input Arguments

<i>theOop</i>	The OOP to search for.
<i>hiddenSetId</i>	The index to the hidden set to search.

Return Value

True if the OOP was found; false otherwise.

Description

The Gem holds objects in a number of sets ordinarily hidden from the user, including the *PureExportSet* and the *GciTrackedObjs* (among others). **GciHiddenSetIncludesOop** allows you to pass in an index to a specified hidden set to determine if the set includes a specific object. For indexes of available hidden sets, see the GemStone Smalltalk method `System Class >> HiddenSetSpecifiers`.

Example

```
OopType TrackedSetContainsOop(OopType anOop)  
{  
    if (GciHiddenSetIncludesOop(anOop, 40/* GciTrackedObjs */))  
        return OOP_TRUE;  
    else  
        return OOP_FALSE;  
}
```

GCI_I64_IS_SMALL_INT

Determine whether or not a C 64-bit integer value can be translated into a SmallInteger object.

Syntax

```
static inline BoolType GCI_I64_IS_SMALL_INT(anInt)
```

Input Arguments

anInt A C 64-bit signed integer.

Result Value

A C Boolean value. Returns TRUE if *anInt* is within SmallInteger range, FALSE otherwise. A SmallInteger has a 61-bit two's-complement integer and three tag bits.

For a positive argument to be within the range of the GemStone SmallInteger class, its top four bits must be 2r0000. For a negative argument, the top four bits must be 2r1111.

Description

This macro tests to see if *anInt* can be represented as a SmallInteger.

See Also

GCI_OOP_IS_SMALL_INT, page 286

GciI64ToOop

Convert a C 64-bit integer value to a GemStone object.

Syntax

```
OopType GciI64ToOop(  
    int64          anInt );
```

Input Arguments

anInt A C 64-bit signed integer.

Return Value

The **GciI64ToOop** function returns the OOP of a GemStone object whose value is equivalent to *anInt*.

Description

The **GciI64ToOop** function translates a C 64-bit integer (`int64_t`) value into the equivalent GemStone object. If the result is not a `SmallInteger`, the result is automatically saved by a `GciSaveObjs()` call.

See Also

GciOopToI64, page 297
GciSaveObjs, page 338

GciIncSharedCounter

Increment the value of a shared counter.

Syntax

```
BoolType GciIncSharedCounter(  
    int64_t          counterIdx,  
    int64_t*        value);
```

Input Arguments

<i>counterIdx</i>	The offset into the shared counters array of the value to increment.
<i>value</i>	Pointer to a value that indicates how much to increment the shared counter by. Shared counters cannot be incremented to a value greater than INT_MAX (2147483647). Attempt to do so will not cause an error, but will set the counter to a value of INT_MAX.

Result Arguments

<i>value</i>	Pointer to a value that indicates the new value of the shared counter, after incrementing.
--------------	--

Return Value

Returns a C Boolean value indicating if the shared counter was successfully incremented. Returns TRUE if successful, FALSE if an error occurred.

Description

This function increments the value of a particular shared counter by a specified amount. The shared counter is specified by index. The maximum value of this shared counter is INT_MAX (2147483647), attempts to increase a shared counter to higher values is not an error, but does not cause the value to increase further.

This function is not supported for remote GCI interfaces, and will always return FALSE.

See Also

- [GciFetchNumSharedCounters](#), page 177
- [GciDecSharedCounter](#), page 134
- [GciSetSharedCounter](#), page 353
- [GciReadSharedCounter](#), page 316
- [GciReadSharedCounterNoLock](#), page 317
- [GciFetchSharedCounterValuesNoLock](#), page 190

Gcilnit

Initialize GemBuilder.

Syntax

```
BoolType Gcilnit();
```

Return Value

The function **Gcilnit** returns TRUE or FALSE to indicate successful or unsuccessful initialization of GemBuilder.

Description

This function initializes GemBuilder. Among other things, it establishes the default GemStone login parameters.

If your C application program is linkable, you may wish to call **GcilnitAppName**, which you must do before you call **Gcilnit**. After **GcilnitAppName**, you *must* call **Gcilnit** before calling any other GemBuilder functions. Otherwise, GemBuilder behavior will be unpredictable.

(Note that when doing run-time binding, you would call **GcilRtlLoad** before calling **Gcilnit**. For details, see “Building the Application” on page 44.)

When GemBuilder is initialized on UNIX platforms, it establishes its own handler for SIGIO interrupts. See “Signal Handling in Your GemBuilder Application” on page 38 for information on **Gcilnit**'s handling of interrupts and pointers on making GemBuilder, application, and third-party handlers work together.

See Also

GcilnitAppName, page 216

GciInitAppName

GciInitAppName_

Override the default application configuration file name.

Syntax

```
void GciInitAppName(  
    const char *      applicationName,  
    BoolType          logWarnings );  
  
void GciInitAppName_(  
    const char *      applicationName,  
    BoolType          logWarnings,  
    unsigned int      gemTempObjCacheOverrideKB,  
    int               gemNativeCodeOverride );
```

Input Arguments

<i>applicationName</i>	The application's name, as a character string.
<i>logWarnings</i>	If TRUE, causes the configuration file parser to print any warnings to standard output at executable startup.
<i>gemTempObjCacheOverrideKB</i>	If non-zero, defines the maximum size (in KB) of temporary object memory for this application. This value overrides any GEM_TEMPOBJ_CACHE_SIZE settings in configuration files read by GciInit .
<i>gemNativeCodeOverride</i>	If non-zero, value overrides the GEM_NATIVE_CODE_ENABLED config file settings in configuration files read by GciInit .

Description

This function affects only linkable applications. It has no effect on RPC applications. If you do not call this function *before* you call **GciInit**, it will have no effect.

A linkable GemBuilder application reads a configuration file called *applicationName.conf* when **GciInit** is called. This file can alter the behavior of the underlying GemStone session. For complete information, please see the *System Administration Guide for GemStone/S 64 Bit*.

A linkable GemBuilder application uses defaults until it calls this function (if it does) and reads the configuration file (which it always does). For linkable GemBuilder applications, the default application name is *gci*, so the default executable configuration file is *gci.conf*. The *applicationName* argument overrides the default application name with one of your choice, which causes your linkable GemBuilder application to read its own executable configuration file.

The *logWarnings* argument determines whether or not warnings that are generated while reading the configuration file are written to standard output. If your application does not call **GciInitAppName**, the default log warnings setting is FALSE. The *logWarnings* argument resets the

default for your application, which is used in the absence of a LOG_WARNINGS entry in the configuration file, or until that entry is read.

The **GciInitAppName_** variant allows you to specify additional arguments that are used by the linkable GemBuilder application: the maximum temporary object cache size, and an override for the native code configuration.

Your application may not call both **GciInitAppName** and **GciInitAppName_**.

See Also

GciInit, page 215

GciInstallUserAction

GciInstallUserAction_

Associate a C function with a Smalltalk user action.

Syntax

```
void GciInstallUserAction(  
    GciUserActionSType * userAction );  
  
void GciInstallUserAction_(  
    GciUserActionSType * userAction,  
    BoolType errorIfDuplicate );
```

Input Arguments

<i>userAction</i>	A pointer to a C structure that describes the user-written C function.
<i>errorIfDuplicate</i>	If True, return an error if there is already a user action with the specified name. If False, leave the existing user action in place and ignore the current call.

Description

This function associates a user action name (declared in Smalltalk) with a user-written C function. Your application must call **GciInstallUserAction** before beginning any GemStone sessions with **GciLogin**. This function is typically called from **GciUserActionInit**. For more information, see Chapter 3, “Writing C Functions To Be Called from GemStone”

See Also

Chapter 3, “Writing C Functions To Be Called from GemStone”, starting on page 47
“The User Action Information Structure” on page 78
GciDeclareAction, page 132
GciUserActionInit, page 413
GciUserActionShutdown, page 414

GciInstMethodForClass

Compile an instance method for a class.

Syntax

```

OopType GciInstMethodForClass(
    OopType      source,
    OopType      oclass,
    OopType      category,
    OopType      symbolList );

```

Input Arguments

<i>source</i>	The OOP of a Smalltalk string to be compiled as an instance method.
<i>oclass</i>	The OOP of the class with which the method is to be associated.
<i>category</i>	The OOP of a Smalltalk string which contains the name of the category to which the method is added. If the category is nil (OOP_NIL), the compiler will add this method to the category "as yet unclassified".
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). Smalltalk resolves symbolic references in source code using symbols that are available from <i>symbolList</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).

Return Value

Returns OOP_NIL, unless there were compiler warnings (such as variables declared but not used, etc.), in which case the return will be the OOP of a string containing the warning messages.

Description

This function compiles an instance method for the given class.

In addition, the Smalltalk virtual machine optimizes a small number of selectors. You may not compile any methods with any of those selectors. See the *Programming Guide* for a list of the optimized selectors.

To *remove* a class method, use **GciExecuteStr** instead.

Example

```

void instanceMethodExample(void)
{
    // Assumes the topaz code for GciFetchVaryingOop example
    // has been executed.

    OopType theClass = GciResolveSymbol("Component", OOP_NIL);
    OopType oCateg = GciNewString("printing");
    // method to return the part number as a String

```

```
OopType oMethodSrc = GciNewString("partNumString ^ partNumber
asString ") ;

GciInstMethodForClass(oMethodSrc, theClass, oCateg, OOP_NIL);
GciErrSType errInfo;
if (GciErr(&errInfo)) {
    printf("error category "FMT_OID" number %d, %s\n",
        errInfo.category, errInfo.number, errInfo.message);
}
}
```

See Also

[GciClassMethodForClass](#), page 110

GciInUserAction

Determine whether or not the current process is executing a user action.

Syntax

```
BoolType GciInUserAction()
```

Return Value

This function returns TRUE if it is called from within a user action, and FALSE otherwise.

Description

This function is intended for use within signal handlers. It can be called any time after **GciInit**.

GciInUserAction returns FALSE if the process is currently executing within a GemBuilder call that was made from a user action. It considers the highest (most recent) call context only.

See Also

GciCallInProgress, page 105

GciIsKindOf

Determine whether or not an object is some kind of a given class or class history.

Syntax

```
BoolType GciIsKindOf(  
    OopType      anObj,  
    OopType      givenClass );
```

Input Arguments

<i>anObj</i>	The object whose kind is to be checked.
<i>givenClass</i>	A class or class history to compare with the object's kind.

Return Value

GciIsKindOf returns TRUE when the class of *anObj* or any of its superclasses is in the class history of *givenClass*. It returns FALSE otherwise.

Description

GciIsKindOf performs structural access that is equivalent to the `isKindOf:` method of the Smalltalk class `Object`. It compares *anObj*'s class and superclasses to see if any of them are in a given class history. When *givenClass* is simply a class (which is typical), **GciIsKindOf** uses *givenClass*'s class history. When *givenClass* is itself a class history, **GciIsKindOf** uses *givenClass* directly.

Since **GciIsKindOf** does consider class histories, it *can* help to support schema modification by simplifying checks on the relationship of types when they can change over time. To accomplish a similar operation without seeing the effects of class histories, use the **GciIsKindOfClass** function.

See Also

- GciIsKindOfClass**, page 223
- GciIsSubclassOf**, page 225
- GciIsSubclassOfClass**, page 226

GciIsKindOfClass

Determine whether or not an object is some kind of a given class.

Syntax

```
BoolType GciIsKindOfClass(  
    OopType      anObj,  
    OopType      givenClass );
```

Input Arguments

<i>anObj</i>	The object whose kind is to be checked.
<i>givenClass</i>	A class to compare with the object's kind.

Return Value

GciIsKindOfClass returns TRUE when the class of *anObj* or any of its superclasses is *givenClass*. It returns FALSE otherwise.

Description

GciIsKindOfClass performs structural access that is equivalent to the `isKindOf:` method of the Smalltalk class `Object`. It compares *anObj*'s class and superclasses to see if any of them are the *givenClass*.

Since **GciIsKindOfClass** does *not* consider class histories, it *cannot* help to support schema modification. To accomplish a similar operation when the relationship of types can change over time, use the **GciIsKindOf** function.

See Also

GciIsKindOf, page 222
GciIsSubclassOf, page 225
GciIsSubclassOfClass, page 226

GciIsRemote

Determine whether the application is running linked or remotely.

Syntax

```
BoolType GciIsRemote( )
```

Return Value

Returns TRUE if this application is running with GciRpc (the remote procedure call version of GemBuilder). Returns FALSE if this application is running with GciLnk (that is, if GemBuilder is linked with your GemStone session).

Description

This function reports whether the current application is using the GciRpc (remote procedure call) or GciLnk (linkable) version of GemBuilder.

GciIsSubclassOf

Determine whether or not a class is a subclass of a given class or class history.

Syntax

```
BoolType GciIsSubclassOf(  
    OopType      aClass,  
    OopType      givenClass );
```

Input Arguments

<i>aClass</i>	The class that is to be checked.
<i>givenClass</i>	A class or class history to compare with the first class.

Return Value

`GciIsSubclassOf` returns TRUE when *aClass* or any of its superclasses is in the class history of *givenClass*. It returns FALSE otherwise.

Description

`GciIsSubclassOf` performs structural access that is equivalent to the `isSubclassOf:` method of the Smalltalk class Behavior. It compares *aClass* and *aClass*'s superclasses to see if any of them are in a given class history. When *givenClass* is simply a class (which is typical), `GciIsSubclassOf` uses *givenClass*'s class history. When *givenClass* is itself a class history, `GciIsSubclassOf` uses *givenClass* directly.

Since `GciIsSubclassOf` does consider class histories, it *can* help to support schema modification by simplifying checks on the relationship of types when they can change over time. To accomplish a similar operation without seeing the effects of class histories, use the `GciIsSubclassOfClass` function.

See Also

`GciIsKindOf`, page 222
`GciIsKindOfClass`, page 223
`GciIsSubclassOfClass`, page 226

GciIsSubclassOfClass

Determine whether or not a class is a subclass of a given class.

Syntax

```
BoolType GciIsSubclassOf(  
    OopType      aClass,  
    OopType      givenClass );
```

Input Arguments

<i>aClass</i>	The class that is to be checked.
<i>givenClass</i>	A class to compare with the first class.

Return Value

GciIsSubclassOf returns TRUE when *aClass* or any of its superclasses is *givenClass*. It returns FALSE otherwise.

Description

GciIsSubclassOfClass performs structural access that is equivalent to the `isSubclassOf:` method of the Smalltalk class Behavior. It compares *aClass* and *aClass*'s superclasses to see if any of them are the *givenClass*.

Since **GciIsSubclassOfClass** does *not* consider class histories, it *cannot* help to support schema modification. To accomplish a similar operation when the relationship of types can change over time, use the **GciIsSubclassOf** function.

See Also

GciIsKindOf, page 222
GciIsKindOfClass, page 223
GciIsSubclassOf, page 225

GciIvNameToIdx

Fetch the index of an instance variable name.

Syntax

```
int GciIvNameToIdx(  
    OopType          oclass,  
    const char       instVarName[ ]);
```

Input Arguments

<i>oclass</i>	The OOP of the class from which to obtain information about instance variables.
<i>instVarName</i>	The instance variable name to search for.

Return Value

Returns the index of *instVarName* into the array of named instance variables for the specified class. Returns 0 if the name is not found or if an error is encountered.

Description

This function searches the array of instance variable names for the specified class (including those inherited from superclasses), and returns the index of the specified instance variable name. This index could then be used as the *atIndex* parameter in the **GciFetchNamedOop** or **GciStoreNamedOop** function call.

Example

```
int nameToIdx_example(void)  
{  
    // Assumes topaz code for GciFetchVaryingOop example has run  
  
    OopType theClass = GciResolveSymbol("Component", OOP_NIL);  
    int idx = GciIvNameToIdx(theClass, "cost");  
    if (idx < 1) {  
        printf("error during GciIvNameToIdx\n");  
    }  
    return idx;  
}
```

See Also

GciClassNamedSize, page 112
GciFetchNamedOop, page 170
GciFetchNamedOops, page 172

GciStoreNamedOop, page 371
GciStoreNamedOops, page 373

GciLoadUserActionLibrary

Load an application user action library.

Syntax

```
BoolType GciLoadUserActionLibrary(  
    const char *          uaLibraryName[ ],  
    BoolType             mustExist,  
    void **              libHandlePtr,  
    char                 infoBuf[ ],  
    int64                infoBufSize );
```

Input Arguments

<i>uaLibraryName</i>	The name and location of the user action library file (a null-terminated string).
<i>mustExist</i>	A flag to make the library required or optional.
<i>libHandlePtr</i>	A variable to store the status of the loading operation.
<i>infoBuf</i>	A buffer to store the name of the user action library or an error message.
<i>infoBufSize</i>	The size of <i>infoBuf</i> .

Return Value

A C Boolean value. If an error occurs, the return value is FALSE, and the error message is stored in *infoBuf*, unless *infoBuf* is NULL. Otherwise, the return value is TRUE, and the name of the user action library is stored in *infoBuf*.

Description

This function loads a user action shared library at run time. If *uaLibraryName* does not contain a path, then a standard user action library search is done. The proper prefix and suffix for the current platform are added to the basename if necessary. For more information, see Chapter 3, "Writing C Functions To Be Called from GemStone"

If a library is loaded, *libHandlePtr* is set to a value that represents the loaded library, if *libHandlePtr* is not NULL. If *mustExist* is TRUE, then an error is generated if the library can not be found. If *mustExist* is FALSE, then the library does not need to exist. In this case, TRUE is returned and *libHandlePtr* is NULL if the library does not exist and non-NULL if it exists.

See Also

GciInstallUserAction, page 218
GciInUserAction, page 221
GciUserActionShutdown, page 414

GciLogin

Start a user session.

Syntax

```
BoolType GciLogin(
    const char      gemstoneUsername[],
    const char      gemstonePassword []);
```

Input Arguments

gemstoneUsername The user's GemStone user name (a null-terminated string).
gemstonePassword The user's GemStone password (a null-terminated string).

Description

The GemStone system is much like a time-shared computer system in that the user must log in before any work may be performed. This function creates a user session and its corresponding transaction workspace.

This function uses the current network parameters (as specified by **GciSetNet**) to establish the user's GemStone session.

Example

```
BoolType login_example(void)
{
    // assume the netldi on machine lichen been started with -a -g
    // so that host userId and host password are not required.
    const char* StoneName      = "!tcp@lichen!gs64stone";
    const char* HostUserId     = "";
    const char* HostPassword   = "";
    const char* GemService     = "!tcp@lichen!gemnetobject";
    const char* gsUserName     = "isaacNewton";
    const char* gsPassword     = "pomme";

    // GciInit required before first login
    if (!GciInit()) {
        printf("GciInit failed\n");
        return FALSE;
    }
    GciSetNet(StoneName, HostUserId, HostPassword, GemService);
    BoolType success = GciLogin(gsUserName, gsPassword);
    if (!success) {
        GciErrSType errInfo;
        if (GciErr(&errInfo)) {
            printf("error category \"FMT_OID\" number %d, %s\n",
                errInfo.category, errInfo.number, errInfo.message);
        }
    }
}
```

```
    }  
    return success;  
}
```

See Also

GciGetSessionId, page 209

GciLogout, page 232

GciSetNet, page 347

GciSetSessionId, page 351

GciLoginEx

Start a user session with session configuration

Syntax

```
BoolType GciLoginEx(
    const char          gemstoneUsername[ ],
    const char          gemstonePassword[ ],
    unsigned int        loginFlags,
    int                 haltOnErrNum);
```

Input Arguments

<i>gemstoneUsername</i>	The user's GemStone user name (a null-terminated string).
<i>gemstonePassword</i>	The user's GemStone password (a null-terminated string).
<i>loginFlags</i>	Login flags, as described below.
<i>haltOnErrNum</i>	A legacy error number.

Description

This function creates a user session and its corresponding transaction workspace. The current network parameters, as specified by **GciSetNet**, are used to establish the user's GemStone session.

Other login attributes are provided by *loginFlags*. The value of *loginFlags* should be given using the following GemBuilder mnemonics:

GCI_LOGIN_PW_ENCRYPTED indicates the *gemstonePassword* is encrypted. When specifying this flag, you must encrypt the password using **GciEncrypt**.

GCI_LOGIN_IS_SUBORDINATE Creates the new session as a child of the current session. It will be terminated if the current session terminates. The primary use of this is for **GciInterface**, and is only valid for RPC sessions.

GCI_LOGIN_FULL_COMPRESSION_ENABLED the login will use full compression.

GCI_LOGIN_ERRS_USE_REF_SET sets the OOPs of objects related to errors to be put into the ReferencedSet, rather than the PureExportSet.

GCI_LOGIN_QUIET suppresses printing of banner information during login.

GCI_CLIENT_DOES_SESSION_INIT after login, the session will manually execute **GciPerform(OOP_CLASS_GSCURRENT_SESSION, "initialize", NULL, 0)**; otherwise, the VM will automatically execute that code as part of login.

GCI_TS_CLIENT the client will use the thread-safe API. Using this flag, **GciLoginEx** is equivalent to **GciTsLogin**, from the *gcits.hf* (not included in this documentation).

If *haltOnErrNum* is set to a GemStone error number, and the session encounters the error associated with that error number, the session will halt.

See Also

GciEncrypt, page 147
GciGetSessionId, page 209
GciLogout, page 232

GciSetNet, page 347
GciSetSessionId, page 351

GciLogout

End the current user session.

Syntax

```
void GciLogout()
```

Description

This function terminates the current user session (set by the last **GciLogin** or **GciSetSessionId**), and allows GemStone to release all uncommitted objects created by the application program in the corresponding transaction workspace. The current session ID is reset to `GCI_INVALID_SESSION_ID`, indicating that the application is no longer logged in. (See “GciGetSessionId” on page 209 for more information.)

See Also

GciGetSessionId, page 209

GciLogin, page 229

GciSetSessionId, page 351

GciLongJump

Provides equivalent functionality to the corresponding `longjmp()` or `_longjmp()` function.

Syntax

```
void GciLongJump(  
    GciJumpBufSType *   jumpBuffer,  
    int                 val );
```

Input Arguments

jumpBuffer A pointer to a jump buffer.

Description

Except for the difference in the first argument type, the semantics of this function are the same as for `longjmp()` on Solaris and `_longjmp()` on HP-UX.

See Also

[GciErr](#), page 148
[GciPopErrJump](#), page 311
[GciPushErrJump](#), page 314
[GciSetErrJump](#), page 343
[Gci_SETJMP](#), page 346

GciMoreTraversal

Continue object traversal, reusing a given buffer.

Syntax

```
BoolType GciMoreTraversal(  
    GciTravBufType *    travBuff);
```

Result Arguments

travBuff A buffer in which the results of the traversal will be placed.

Return Value

Returns FALSE if the traversal is not yet completed, but further traversal would cause the *travBuffSize* to be exceeded. If the *travBuffSize* is reached before the traversal is complete, you can continue to call **GciMoreTraversal** to proceed from the point where *travBuffSize* was exceeded.

Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal**.

Description

When the amount of information obtained in a traversal exceeds the amount of memory available to the buffer (as specified with *travBuffSize*), your application can call **GciMoreTraversal** repeatedly to break the traversal into manageable amounts of information. The information returned by this function begins with the object report following where the previous unfinished traversal left off. The level value is retained from the initial **GciTraverseObjs** call.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

Generally speaking, an application can continue to call **GciMoreTraversal** until it has obtained all requested information.

Naturally, GemStone will not continue an incomplete traversal if there is any chance that changes to the database in the intervening period might have invalidated the previous report or changed the connectivity of the objects in the path of the traversal. Specifically, GemStone will refuse to continue a traversal if, in the interval before attempting to continue, you:

- Modify the objects in the database directly by calling any of the **GciStore...** or **GciAdd...** functions;
- Call one of the Smalltalk message-sending functions **GciPerform**, **GciContinue**, or any of the **GciExecute...** functions.
- Abort your transaction, thus invalidating any subsequent information from that traversal.

Any attempt to call **GciMoreTraversal** after one of these events will generate an error.

Note that this holds true across multiple GemBuilder applications sharing the same GemStone session. Suppose, for example, that you were holding on to an incomplete traversal buffer and the user moved from the current application to another, did some work that required executing Smalltalk code, and then returned to the original application. You would be unable to continue the interrupted traversal.

If you attempt to call **GciMoreTraversal** when no traversal is underway, this function will generate the error `GCI_ERR_TRAV_COMPLETED`.

During the entire sequence of **GciTraverseObjs** and **GciMoreTraversal** calls that constitute a traversal, any single object report will be returned exactly once. Regardless of the connectivity of objects in the GemStone database, only one report will be generated for any non-special object.

The section "Organization of the Traversal Buffer" on page 409 describes the organization of traversal buffers in detail.

GciMoreTraversal provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see "Byte-Swizzling of Binary Floating-Point Values" on page 25.

```
void moreTraversalExample(void)
{
    // Assumes topaz code for GciFetchVaryingOops example has run

    OopType rootObj = GciResolveSymbol("AllComponents", OOP_NIL);
    GciTravBufType *buf = GciTravBufType::malloc(8000);

    int totalCount = 0;
    // traverse the AllComponents collection to 10 levels deep
    BoolType done = GciTraverseObjs(&rootObj, 1, buf, 10);
    while (! done) {
        int objCount = 0;
        GciObjRepHdrSType *rpt = buf->firstReportHdr();
        GciObjRepHdrSType *limit = buf->readLimitHdr();
        while (rpt < limit) {
            objCount++ ;
            rpt = rpt->nextReport();
        }
        totalCount += objCount;
        done = GciMoreTraversal(buf);
    }
    buf->free();
    printf("traversal returned %d total objects\n", totalCount);
}
```

See Also

GCI_ALIGN, page 94

GciFindObjectRep, page 200

GciNbMoreTraversal, page 253

GciNbTraverseObjs, page 265

GciNbTraverseObjs, page 265

GciObjRepSize_, page 282

GciTraverseObjs, page 408

GciNbAbort

Abort the current transaction (nonblocking).

Syntax

```
void GciNbAbort()
```

Description

The **GciNbAbort** function is equivalent in effect to **GciAbort**. However, **GciNbAbort** permits the application to proceed with non-GemStone tasks while the transaction is aborted, and **GciAbort** does not.

See Also

GciAbort, page 91

GCI_CHR_TO_OOP, page 107

GciCommit, page 115

GciNbCommit, page 239

GciNbBegin

Begin a new transaction (nonblocking).

Syntax

```
void GciNbBegin()
```

Description

The **GciNbBegin** function is equivalent in effect to **GciBegin**. However, **GciNbBegin** permits the application to proceed with non-GemStone tasks while a new transaction is started, and **GciBegin** does not.

See Also

GciAbort, page 91

GciBegin, page 102

GciExecuteStr, page 153

GciNbAbort, page 236

GciNbExecuteStr, page 246

GciNbClampedTrav

Traverse an array of objects, subject to clamps (nonblocking).

Syntax

```
void GciNbClampedTrav(  
    const OopType *    theOops,  
    int                numOops,  
    GciClampedTravArgsSType *travArgs );
```

Input Arguments

<i>theOops</i>	An array of OOPs representing the objects to traverse.
<i>numOops</i>	The number of elements in <i>theOops</i> .
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType . See GciClampedTrav (page 108) for documentation on the fields in <i>travArgs</i> .

Result Arguments

<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType containing the result argument field <i>travBuff</i> .
-----------------	---

Return Value

The **GciNbClampedTrav** function, unlike **GciClampedTrav**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciClampedTrav** by using the argument to **GciNbEnd**.

Description

This function is equivalent in effect to **GciClampedTrav**. However, **GciClampedTrav** permits the application to proceed with non-GemStone tasks while a traversal is carried out, and **GciClampedTrav** does not.

See Also

GciClampedTrav, page 108

GciNbCommit

Write the current transaction to the database (nonblocking).

Syntax

```
void GciNbCommit()
```

Return Value

The **GciNbCommit** function, unlike **GciCommit**, does not have a return value. However, when the commit operation is complete, you can access a value identical in meaning to the return value of **GciCommit** by using the argument to **GciNbEnd**.

Description

This function is equivalent in effect to **GciCommit**. However, **GciNbCommit** permits the application to proceed with non-GemStone tasks while the transaction is committed, and **GciCommit** does not.

See Also

GciAbort, page 91
GCI_CHR_TO_OOP, page 107
GciCommit, page 115
GciNbAbort, page 236

GciNbContinue

Continue code execution in GemStone after an error (nonblocking).

Syntax

```
void GciNbContinue(  
    OopType          process );
```

Input Arguments

process The OOP of a GsProcess object (obtained as the value of the *context* field of an error report returned by **GciErr**).

Return Value

The **GciNbContinue** function, unlike **GciContinue**, does not have a return value. However, when the continued operation is complete, you can access a value identical in meaning to the return value of **GciContinue** by using the argument to **GciNbEnd**.

Description

The **GciNbContinue** function is equivalent in effect to **GciContinue**. However, **GciNbContinue** permits the application to proceed with non-GemStone tasks while the operation continues, and **GciContinue** does not.

See Also

GciClearStack, page 113
GciContinue, page 120
GciErr, page 148
GciExecute, page 149
GciNbExecute, page 245

GciNbContinueWith

Continue code execution in GemStone after an error (nonblocking).

Syntax

```
void GciNbContinueWith (  
    OopType          process,  
    OopType          replaceTopOfStack,  
    int              flags,  
    GciErrSType *    error );
```

Input Arguments

<i>process</i>	The OOP of a GsProcess object (obtained as the value of the <i>context</i> field of an error report returned by GciErr).
<i>replaceTopOfStack</i>	If not OOP_ILLEGAL, replace the top of the Smalltalk evaluation stack with this value before continuing. If OOP_ILLEGAL, the evaluation stack is not changed.
<i>flags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk, as defined for GciPerformNoDebug .
<i>error</i>	If not NULL, continue with an error. This argument takes precedence over <i>replaceTopOfStack</i> .

Description

The **GciNbContinueWith** function is equivalent in effect to **GciContinueWith**. However, **GciNbContinueWith** permits the application to proceed with non-GemStone tasks while the operation continues, and **GciContinueWith** does not.

See Also

GciContinue, page 120
GciContinueWith, page 121
GciErr, page 148
GciExecute, page 149
GciNbExecute, page 245
GciPerformNoDebug, page 300

GciNbEnd

GciNbEnd_

Test the status of nonblocking call in progress for completion.

Syntax

```
GciNbProgressEType GciNbEnd(  
    void **  
    result );
```

```
GciNbProgressEType GciNbEnd_  
    int64*  
    result,  
    );
```

Input Arguments

result The address at which **GciNbEnd** or **GciNbEnd_** should place a pointer to the result of the nonblocking call when it is complete. The actual result is that of the corresponding equivalent blocking GCI call.

With **GciNbEnd**, the result may be either 4 byte or 8 byte, and the caller must be aware of whether this is 4 byte or 8 byte, and dereference result accordingly on big endian machines.

GciNbEnd_ produces 8 byte results that are correct on big endian machines.

Return Value

The **GciNbEnd** function returns an enumerated type. Its value is `GCI_RESULT_READY` if the outstanding nonblocking call has completed execution and its result is ready, `GCI_RESULT_NOT_READY` if the call is not complete and there has been no change since the last inquiry, and `GCI_RESULT_PROGRESSED` if the call is not complete but progress has been made towards its completion.

Description

Once an application calls a nonblocking function, it must call **GciNbEnd**, **GciNbEnd_**, or **GciNbEndPoll** at least once, and must continue to do so until that nonblocking function has completed execution. The intent of the return values is to give the scheduler a hint about whether it is calling one of these functions too often or not often enough.

Once an operation is complete, you are permitted to call **GciNbEnd** or a related function repeatedly. It returns `GCI_RESULT_READY` and a pointer to the same result each time, until you call a nonblocking function again. It is an error to call **GciNbEnd** or a related function before you call any nonblocking functions at all. Use the **GciCallInProgress** function to determine whether or not there is a GemBuilder call currently in progress.

Example

```
void nbEnd_example(void)
```

```
{
  void *resultPtr;
  GciNbExecuteStr("Globals size", OOP_NIL);
  do {
    // wait for non-blocking result
    GciHostMilliSleep(1);
  } while (GciNbEnd(&resultPtr) != GCI_RESULT_READY);

  OopType result = *(OopType*)resultPtr;
  BoolType conversionErr = FALSE;
  int gSize = GciOopToI32_(result, &conversionErr);
  if (conversionErr) {
    printf("error in execution\n");
  } else {
    printf("Globals size = %d \n", gSize);
  }
}
```

See Also

[GciCallInProgress](#), page 105

[GciNbEndPoll](#), page 244

GciNbEndPoll

Test the status of nonblocking call in progress for completion, with timeout.

Syntax

```
GciNbProgressEType GciNbEndPoll(  
    int64*          result,  
    int            timeoutMs  
);
```

Input Arguments .

result The address at which **GciNbEndPoll** should place a pointer to the result of the nonblocking call when it is complete. The actual result is that of the corresponding equivalent blocking GCI call. The result will be 8 bytes, and correct on big endian machines.

Return Value

The **GciNbEndPoll** function returns an enumerated type. Its value is `GCI_RESULT_READY` if the outstanding nonblocking call has completed execution and its result is ready, `GCI_RESULT_NOT_READY` if the call is not complete and there has been no change since the last inquiry, and `GCI_RESULT_PROGRESSED` if the call is not complete but progress has been made towards its completion.

Description

Once an application calls a nonblocking function, it must call **GciNbEndPoll**, **GciNbEnd**, or **GciNbEnd_** at least once, and must continue to do so until that nonblocking function has completed execution. The intent of the return values is to give the scheduler a hint about whether it is calling **GciNbEndPoll** too often or not often enough.

GciNbEndPoll allows you to specify a timeout, and will wait for *timeoutMs* milliseconds before returning status. This avoids too-frequent polling for long-running code.

Once an operation is complete, you are permitted to call **GciNbEndPoll** repeatedly. It returns `GCI_RESULT_READY` and a pointer to the same result each time, until you call a nonblocking function again. It is an error to call **GciNbEndPoll** before you call any nonblocking functions at all. Use the **GciCallInProgress** function to determine whether or not there is a GemBuilder call currently in progress.

See Also

GciNbEnd, page 242
GciCallInProgress, page 105

GciNbExecute

GciNbExecute_

Execute a Smalltalk expression contained in a String object (nonblocking).

Syntax

```
void GciNbExecute(
    OopType          source,
    OopType          symbolList );

void GciNbExecute_(
    OopType          source,
    OopType          symbolList,
    ushort          environmentId );
```

Input Arguments

<i>source</i>	The OOP of a String containing a sequence of one or more statements to be executed.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbExecute** function, unlike **GciExecute**, does not have a return value. However, when the executed operation is complete, you can access a value identical in meaning to the return value of **GciExecute** by using the argument to **GciNbEnd**.

Description

The **GciNbExecute** function is equivalent in effect to **GciExecute**. However, **GciNbExecute** permits the application to proceed with non-GemStone tasks while the Smalltalk expression is executed, and **GciExecute** does not.

See Also

- GciErr**, page 148
- GciExecute**, page 149
- GciNbContinue**, page 240
- GciNbExecuteStr**, page 246
- GciNbExecuteStrFromContext**, page 249

GciNbExecuteStr

GciNbExecuteStr_

Execute a Smalltalk expression contained in a C string (nonblocking).

Syntax

```
void GciNbExecuteStr(
    const char      source[ ],
    OopType         symbolList );

void GciNbExecuteStr_(
    const char      source[ ],
    OopType         symbolList,
    ushort         environmentId );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbExecuteStr** function, unlike **GciExecuteStr**, does not have a return value. However, when the executed operation is complete, you can access a value identical in meaning to the return value of **GciExecuteStr** by using the argument to **GciNbEnd**.

Description

The **GciNbExecuteStr** function is equivalent in effect to **GciExecuteStr**. However, **GciNbExecuteStr** permits the application to proceed with non-GemStone tasks while the Smalltalk expression is executed, and **GciExecuteStr** does not.

See Also

GciErr, page 148
GciExecuteStr, page 153
GciNbEnd, page 242
GciNbExecute, page 245
GciNbExecuteStrFromContext, page 249

GciNbExecuteStrFetchBytes

Execute a Smalltalk expression contained in a C string, returning byte-format results (nonblocking).

Syntax

```
int64 GciNbExecuteStrFetchBytes(
    const char *      source,
    int64             sourceSize,
    OopType           sourceClass,
    OopType           contextObject,
    OopType           symbolList,
    ByteType *       result,
    int64             maxResultSize );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>sourceSize</i>	The number of bytes in the <i>source</i> , or -1. If <i>sourceSize</i> is -1, <code>strlen(source)</code> is used.
<i>sourceClass</i>	The OOP of the class that <i>source</i> should be converted to. Examples are <code>OOP_CLASS_STRING</code> , <code>OOP_CLASS_Utf8</code> , <code>OOP_CLASS_Unicode7</code> .
<i>contextObject</i>	The OOP of any GemStone object. The code to be executed is compiled as if it were an instance method in the class of <i>contextObject</i> . A value of <code>OOP_NO_CONTEXT</code> means no context.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of <code>SymbolListDictionary</code>). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of <code>OOP_NIL</code> means to use the default symbol list for the current GemStone session (that is, <code>System.myUserProfile.symbolList</code>).
<i>maxResultSize</i>	Maximum size of the resulting string.

Result Arguments

<i>result</i>	Array in which to put the results of executing the <i>source</i> .
---------------	--

Return Value

The `GciNbExecuteStrFetchBytes` function, unlike `GciExecuteStrFetchBytes`, does not have a return value. However, when the executed operation is complete, you can access a value identical in meaning to the return value of `GciExecuteStrFetchBytes` by using the argument to `GciNbEnd`.

Description

This function sends an expression (or sequence of expressions) to GemStone for execution. The execution result, which should be a byte format object is returned in the **result* buffer.

See Also

- [GciErr](#), page 148
- [GciExecuteStrFetchBytes](#), page 155
- [GciNbContinue](#), page 240
- [GciNbEnd](#), page 242
- [GciNbExecute](#), page 245
- [GciNbExecuteStr](#), page 246
- [GciNbExecuteStrFromContext](#), page 249

GciNbExecuteStrFromContext

GciNbExecuteStrFromContext_

Execute a Smalltalk expression contained in a C string as if it were a message sent to an object (nonblocking).

Syntax

```
void GciNbExecuteStrFromContext(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList );

void GciNbExecuteStrFromContext_(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    ushort             environmentId );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>contextObject</i>	The OOP of any GemStone object.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolListDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbExecuteStrFromContext** function, unlike **GciExecuteStrFromContext**, does not have a return value. However, when the executed operation is complete, you can access a value identical in meaning to the return value of **GciExecuteStrFromContext** by using the argument to **GciNbEnd**.

Description

The **GciNbExecuteStrFromContext** function is equivalent in effect to **GciExecuteStrFromContext**. However, **GciNbExecuteStrFromContext** permits the application to proceed with non-GemStone tasks while the Smalltalk expression is executed, and **GciExecuteStrFromContext** does not.

See Also

GciContinue, page 120
GciErr, page 148

GciExecuteStrFromContext, page 157

GciNbContinue, page 240

GciNbEnd, page 242

GciNbExecute, page 245

GciNbExecuteStr, page 246

GciNbExecuteStrTrav

GciNbExecuteStrTrav_

First execute a Smalltalk expression contained in a C string as if it were a message sent to an object, then traverse the result of the execution (nonblocking).

Syntax

```
void GciNbExecuteStrTrav(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    GciClampedTravArgsSType *travArgs );
```

```
void GciNbExecuteStrTrav_(
    const char          source[ ],
    OopType             contextObject,
    OopType             symbolList,
    GciClampedTravArgsSType *travArgs,
    ushort             environmentId );
```

Input Arguments

<i>source</i>	A null-terminated string containing a sequence of one or more statements to be executed.
<i>contextObject</i>	The OOP of any GemStone object. A value of OOP_ILLEGAL means no context.
<i>symbolList</i>	The OOP of a GemStone symbol list (that is, an Array of instances of SymbolDictionary). The compiler uses the <i>symbolList</i> to resolve symbolic references in the code in <i>source</i> . A value of OOP_NIL means to use the default symbol list for the current GemStone session (that is, <code>System myUserProfile symbolList</code>).
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType . See GciExecuteStrTrav on page 159 for field definitions.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbExecuteStrTrav** function, unlike **GciExecuteStrTrav**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciExecuteStrTrav** by using the argument to **GciNbEnd**.

Description

The **GciNbExecuteStrTrav** function is equivalent in effect to **GciExecuteStrTrav**. However, **GciNbExecuteStrTrav** permits the application to proceed with non-GemStone tasks while the traversal is completed, and **GciExecuteStrTrav** does not.

See Also

`GciExecuteStrTrav`, page 159
`GciExecuteStr`, page 153
`GciMoreTraversal`, page 234
`GciPerformTraverse`, page 306

GciNbMoreTraversal

Continue object traversal, reusing a given buffer (nonblocking).

Syntax

```
void GciNbMoreTraversal(  
    GciTravBufType *    travBuff);
```

Result Arguments

travBuff A buffer in which the results of the traversal will be placed.

Return Value

The **GciNbMoreTraversal** function, unlike **GciMoreTraversal**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciMoreTraversal** by using the argument to **GciNbEnd**.

Description

The **GciNbMoreTraversal** function is equivalent in effect to **GciMoreTraversal**. However, **GciNbMoreTraversal** permits the application to proceed with non-GemStone tasks while the traversal is completed, and **GciMoreTraversal** does not.

GciNbMoreTraversal provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see “Byte-Swizzling of Binary Floating-Point Values” on page 25.

See Also

GCL_ALIGN, page 94
GciFindObjRep, page 200
GciMoreTraversal, page 234
GciNbTraverseObjs, page 265
GciObjRepSize_, page 282
GciTraverseObjs, page 408

GciNbPerform

GciNbPerform_

Send a message to a GemStone object (nonblocking).

Syntax

```
void GciNbPerform(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs );

void GciNbPerform_(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    ushort          environmentId );
```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A string that defines the message selector. For keyword selectors, all keywords are concatenated in the string. (For example, at:put:).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbPerform** function, unlike **GciPerform**, does not have a return value. However, when the performed operation is complete, you can access a value identical in meaning to the return value of **GciPerform** by using the argument to **GciNbEnd**.

Description

The **GciNbPerform** function is equivalent in effect to **GciPerform**. However, **GciNbPerform** permits the application to proceed with non-GemStone tasks while the message is executed, and **GciPerform** does not.

See Also

- GciContinue**, page 120
- GciErr**, page 148
- GciExecute**, page 149
- GciNbContinue**, page 240
- GciNbExecute**, page 245
- GciNbPerformNoDebug**, page 256
- GciPerform**, page 298

GciNbPerformNoDebug

GciNbPerformNoDebug_

Send a message to a GemStone object, and temporarily disable debugging (nonblocking).

Syntax

```
void GciNbPerformNoDebug(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    int              flags );

void GciNbPerformNoDebug_(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    int              flags,
    ushort          environmentId );
```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A string that defines the message selector. For keyword selectors, all keywords are concatenated in the string. (For example, at:put:).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>flags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

The **GciNbPerformNoDebug** function, unlike **GciPerformNoDebug**, does not have a return value. However, when the performed operation is complete, you can access a value identical in meaning to the return value of **GciPerformNoDebug** by using the argument to **GciNbEnd**.

Description

The **GciNbPerformNoDebug** function is equivalent in effect to **GciPerformNoDebug**. However, **GciNbPerformNoDebug** permits the application to proceed with non-GemStone tasks while the message is executed, and **GciPerformNoDebug** does not.

See Also

- GciContinue**, page 120
- GciErr**, page 148
- GciExecute**, page 149
- GciNbContinue**, page 240
- GciNbExecute**, page 245
- GciNbPerform**, page 254
- GciPerformNoDebug**, page 300

GciNbPerformTrav

GciNbPerformTrav_

First send a message to a GemStone object, then traverse the result of the message (nonblocking).

Syntax

```
BoolType GciNbPerformTrav(
    OopType          receiver,
    const char *     selector,
    const OopType *  args,
    int              numArgs,
    GciClampedTravArgsSType *travArgs );
```

```
BoolType GciNbPerformTrav_(
    OopType          receiver,
    const char *     selector,
    const OopType *  args,
    int              numArgs,
    GciClampedTravArgsSType *travArgs,
    ushort          environmentId );
```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A pointer to a character collection containing the message selector. For keyword selectors, all keywords are concatenated in the String. (For example, at :put :).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType . See GciClampedTrav on page 108 for documentation of the fields in <i>travArgs</i> .
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80

Result Arguments

The result of the **GciNbPerformTrav** is the first object in the resulting *travBufs* field in *travArgs*.

Return Value

The **GciNbPerformTrav** function, unlike **GciPerformTrav**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciPerformTrav** by using the argument to **GciNbEnd**.

Description

The **GciNbPerformTrav** function is equivalent in effect to **GciPerformTrav**. However, **GciNbStoreTrav** permits the application to proceed with non-GemStone tasks while the traversal is done, and **GciPerformTrav** does not.

See Also

GciPerformTrav, page 304

GciPerform, page 298

GciClampedTrav, page 108

GciNbStoreTrav

Stores multiple traversal buffer values in objects (nonblocking).

Syntax

```
void GciNbStoreTrav(  
    GciTravBufType *   travBuff,  
    int                 behaviorFlag );
```

Input Arguments

<i>travBuff</i>	A buffer that contains the object reports to be stored. The first element in the buffer is an integer that indicates how many bytes are stored in the buffer. The remainder of the traversal buffer consists of a series of object reports, each of which is of type GciObjRepSType .
<i>behaviorFlag</i>	A flag specifying whether the values returned by GciStoreTrav should be added to the values in the traversal buffer or should replace the values in the traversal buffer. Flag values, predefined in the <code>gci.ht</code> header file, are <code>GCI_STORE_TRAV_NSC_ADD</code> (add to the traversal buffer) and <code>GCI_STORE_TRAV_NSC_REPLACE</code> (replace traversal buffer contents).

Description

The **GciNbStoreTrav** function is equivalent in effect to **GciStoreTrav**. However, **GciNbStoreTrav** permits the application to proceed with non-GemStone tasks while the traversals are stored, and **GciStoreTrav** does not.

GciNbStoreTrav provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see “Byte-Swizzling of Binary Floating-Point Values” on page 25.

See Also

GciMoreTraversal, page 234
GciNbMoreTraversal, page 253
GciNbTraverseObjs, page 265
GciStoreTrav, page 384
GciTraverseObjs, page 408

GciNbStoreTravDo_

Store multiple traversal buffer values in objects, execute the specified code, and return the resulting object (non-blocking).

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciNbStoreTravDo** (without the underscore).

Syntax

```
void GciNbStoreTravDo_(  
    GciStoreTravDoArgsSType *stdArgs);
```

Input Arguments

stdArgs An instance of **GciStoreTravDoArgsSType**. For details, refer to the discussion of **GciStoreTravDo_** on page 387.

Return Value

Unlike **GciStoreTravDo_**, the **GciNbStoreTravDo_** function does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciStoreTravDo_** by using the argument to **GciNbEnd**.

Description

The **GciNbStoreTravDo_** function is equivalent in effect to **GciStoreTravDo_**. However, **GciNbStoreTravDo_** permits the application to proceed with non-GemStone tasks while the traversal is done, and **GciStoreTravDo_** does not.

See Also

GciNbClampedTrav, page 238
GciNbEnd, page 242
GciNbStoreTrav, page 260
GciNbStoreTravDoTrav_, page 262
GciStoreTravDo_, page 387

GciNbStoreTravDoTrav_

Combine in a single function the calls to **GciNbStoreTravDo_** and **GciNbClampedTrav**, to store multiple traversal buffer values in objects, execute the specified code, and traverse the result object (non-blocking).

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciNbStoreTravDoTrav** (without the underscore).

Syntax

```
void GciNbStoreTravDoTrav_(  
    GciStoreTravDoArgsSType *stdArgs,  
    GciClampedTravArgsSType *ctArgs );
```

Input Arguments

<i>stdArgs</i>	An instance of GciStoreTravDoArgsSType . For details, refer to the discussion of GciStoreTravDo_ on page 387.
<i>ctArgs</i>	An instance of GciClampedTravArgsSType . For details, see the discussion of GciClampedTrav on page 108.

Return Value

The **GciNbStoreTravDoTrav_** function, unlike **GciStoreTravDoTrav_**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciStoreTravDoTrav_** by using the argument to **GciNbEnd**.

Description

This function allows the client to execute behavior on the Gem and return the traversal of the result object in a single network round-trip.

The **GciNbStoreTravDoTrav_** function is equivalent in effect to **GciStoreTravDoTrav_**. However, **GciNbStoreTravDoTrav_** permits the application to proceed with non-GemStone tasks while the traversals are stored, and **GciStoreTravDoTrav_** does not.

See Also

GciNbClampedTrav, page 238
GciNbEnd, page 242
GciNbStoreTrav, page 260
GciStoreTravDoTrav_, page 390

GciNbStoreTravDoTravRefs_

Combine in a single function modifications to session sets, traversal of objects to the server, optional Smalltalk execution, and traversal to the client of changed objects and (optionally) the result object (non blocking).

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciNbStoreTravDoTravRefs** (without the underscore).

Syntax

```
void GciNbStoreTravDoTravRefs_(
    const OopType *      oopsNoLongerReplicated,
    int                  numNotReplicated,
    const OopType *      oopsGcedOnClient,
    int                  numGced,
    GciStoreTravDoArgsSType *stdArgs,
    GciClampedTravArgsSType *ctArgs );
```

Input Arguments

<i>oopsNoLongerReplicated</i>	An Array of objects to be removed from the PureExportSet and added to the ReferencedSet.
<i>numNotReplicated</i>	The number of elements in <i>oopsNoLongerReplicated</i> .
<i>oopsGcedOnClient</i>	An Array of objects to be removed from both the PureExportSet and ReferencedSet.
<i>numGced</i>	The number of elements in <i>oopsGcedOnClient</i> .
<i>stdArgs</i>	An instance of GciStoreTravDoArgsSType . For details, refer to the discussion of GciStoreTravDo_ on page 387.
<i>ctArgs</i>	An instance of GciClampedTravArgsSType . For details, see the discussion of GciClampedTrav on page 108.

Return Value

The **GciNbStoreTravDoTravRefs_** function, unlike **GciStoreTravDoTravRefs_**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciStoreTravDoTravRefs_** by using the argument to **GciNbEnd**

Description

This function allows the client to modify the PureExportSet and ReferencedSet, modify or create any number of objects on the server, execute behavior on the Gem, and return the traversal of the result object, all in a single network round-trip.

The **GciNbStoreTravDoTravRefs_** function is equivalent in effect to **GciStoreTravDoTravRefs_**. However, **GciNbStoreTravDoTravRefs_** permits the application to proceed with non-GemStone tasks while the traversals are stored, and **GciStoreTravDoTravRefs_** does not.

See Also

GciNbClampedTrav, page 238

GciNbEnd, page 242

GciStoreTravDoTrav_, page 390

GciStoreTravDoTravRefs_, page 391

GciNbTraverseObjs

Traverse an array of GemStone objects (nonblocking).

Syntax

```
void GciNbTraverseObjs(  
    const OopType      theOops [ ],  
    int                numOops,  
    GciTravBufType *  travBuff,  
    int                level );
```

Input Arguments

<i>theOops</i>	An array of OOPs representing the objects to traverse.
<i>numOops</i>	The number of elements in theOops.
<i>travBuffSize</i>	The number of bytes allocated to the traversal buffer.
<i>level</i>	Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in theOops. When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.

Result Arguments

<i>travBuff</i>	A buffer in which the results of the traversal will be placed.
-----------------	--

Return Value

The **GciNbTraverseObjs** function, unlike **GciTraverseObjs**, does not have a return value. However, when the traversal operation is complete, you can access a value identical in meaning to the return value of **GciTraverseObjs** by using the argument to **GciNbEnd**.

Description

The **GciNbTraverseObjs** function is equivalent in effect to **GciTraverseObjs**. However, **GciNbTraverseObjs** permits the application to proceed with non-GemStone tasks while the traversal is completed, and **GciTraverseObjs** does not.

GciNbTraverseObjs provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see page 25.

See Also

GciMoreTraversal, page 234
GciNbMoreTraversal, page 253
GciNbStoreTrav, page 260
GciStoreTrav, page 384
GciTraverseObjs, page 408

GciNewByteObj

Create and initialize a new byte object.

Syntax

```
OopType GciNewByteObj(  
    OopType          aClass,  
    const ByteType * value,  
    int64            valueSize );
```

Input Arguments

<i>aClass</i>	The OOP of the class of which an instance is to be created.
<i>value</i>	Pointer to an array of byte values to be stored in the newly-created object.
<i>valueSize</i>	The number of byte values in value.

Return Value

The OOP of the newly created object.

Description

Returns a new instance of *aClass*, of size *valueSize*, and containing a copy of the bytes located at *value*. Equivalent to **GciNewOop** followed by **GciStoreBytes**. *aClass* must be a class whose format is Bytes.

GciNewCharObj

Create and initialize a new character object.

Syntax

```
OopType GciNewCharObj(  
    OopType          aClass,  
    const char *     cString );
```

Input Arguments

<i>aClass</i>	The OOP of the class of which an instance is to be created. <i>aClass</i> must be a class whose format is BYTE.
<i>cString</i>	Pointer to an array of characters to be stored in the newly-created object. The terminating '\0' character is not stored.

Return Value

The OOP of the newly-created object.

Description

Returns a new instance of *aClass* which has been initialized to contain the bytes of *cString*, excluding the null terminator.

GciNewDateTime

Create and initialize a new date-time object.

Syntax

```
OopType GciNewDateTime(  
    OopType          theClass,  
    const GciDateTimeSType *timeVal );
```

Input Arguments

<i>theClass</i>	The class of the object to be created. <i>theClass</i> must be OOP_CLASS_DATE_TIME or a subclass thereof.
<i>timeVal</i>	The time value to be assigned to the newly-created object.

Return Value

Returns the OOP of the newly-created object. If an error occurs, returns OOP_ILLEGAL.

Description

Creates a new instance of *theClass* having the value that *timeVal* points to.

GciNewOop

Create a new GemStone object.

Syntax

```
OopType GciNewOop(  
    OopType          oclass );
```

Input Arguments

<i>oclass</i>	The OOP of the class of which the new object is an instance. This may be the OOP of a class that you have created, or it may be one of the Smalltalk kernel classes, such as OOP_CLASS_STRING for an object of class String. It may not be Symbol or DoubleByteSymbol. Appendix A, "Reserved OOPs" lists the C constants that are defined for each of the Smalltalk kernel classes.
---------------	---

Return Value

Returns the OOP of the new object. In case of error, this function returns OOP_NIL.

Description

This function creates a new object of the specified class and returns the object's OOP. It cannot be used to create instances of Symbol or DoubleByteSymbol.

Example

```
OopType newOop_example(void)  
{  
    // create a new instance of String  
    OopType result = GciNewOop(OOP_CLASS_STRING);  
    return result;  
}
```

See Also

- [GciNewOops](#), page 270
- [GciNewOopUsingObjRep](#), page 272
- [GciReleaseAllOops](#), page 320
- [GciReleaseGlobalOops](#), page 322

GciNewOops

Create multiple new GemStone objects.

Syntax

```
void GciNewOops(  
    int                numOops,  
    const OopType      oclass[ ],  
    const int64        idxSize[ ],  
    OopType            result[ ] );
```

Input Arguments

<i>numOops</i>	The number of new objects to be created.
<i>oclass</i>	For each new object, the OOP of its class. This should not be the OOP of Symbol or DoubleByteSymbol.
<i>idxSize</i>	For each new object, the number of its indexed variables. If the specified oclass of an object is not indexable, its <i>idxSize</i> is ignored.

Result Arguments

<i>result</i>	An array of the OOPs of the new objects created with this function.
---------------	---

Return Value

If an error is encountered, this function will stop at the first error and the contents of the *result* array will be undefined.

Description

This function creates multiple objects of the specified classes and sizes, and returns the OOPs of the new objects.

Each OOP in *oclass* may be the OOP of a class that you have created, or it may be one of the Smalltalk kernel classes, such as OOP_CLASS_STRING for an object of class String. This function cannot be used to create instances of Symbol or DoubleByteSymbol. If *oclass* contains the OOP of a class with special implementation (such as Boolean), then the corresponding element in *result* is OOP_NIL. Appendix A, "Reserved OOPs" lists the C constants that are defined for each of the Smalltalk kernel classes.

GciNewOops generates an error when either of the following conditions is TRUE for any object:

- *idxSize* < 0
- (*idxSize* + number_of_named_instance_variables) > maxSmallInt

Example

```
void newOops_example(void)  
{
```

```
enum { num_objs = 3 };
OopType classes[num_objs];
classes[0] = OOP_CLASS_STRING;
classes[1] = OOP_CLASS_IDENTITY_SET;
classes[2] = OOP_CLASS_ARRAY;

int64 sizes[num_objs];
sizes[0] = 50;
sizes[1] = 0; /* ignored for NSCs anyway */
sizes[2] = 3;

OopType newObjs[num_objs];

GciNewOops(num_objs, classes, sizes, newObjs);
GciErrSType errInfo;
if (GciErr(&errInfo)) {
    printf("error category "FMT_OID" number %d, %s\n",
           errInfo.category, errInfo.number, errInfo.message);
} else {
    printf("objIds of new objects are "FMT_OID" "FMT_OID"
           "FMT_OID"\n",
           newObjs[0], newObjs[1], newObjs[2]);
}
}
```

See Also

[GciNewOop](#), page 269
[GciNewOopUsingObjRep](#), page 272
[GciReleaseAllOops](#), page 320
[GciReleaseGlobalOops](#), page 322
[GciStoreTrav](#), page 384

GciNewOopUsingObjRep

Create a new GemStone object from an existing object report.

Syntax

```
void GciNewOopUsingObjRep(  
    GciObjRepSType *    anObjectReport );
```

Input Arguments

anObjectReport A pointer to an object report.

Result Arguments

anObjectReport A modified object report that contains the OOP of the new object (*hdr.objId*), the ID of the object's security policy (*hdr.objectSecurityPolicyId*), the number of named instance variables in the object (*hdr.namedSize*), the updated number of the object's indexed variables (*hdr.idxSize*), and the object's complete size (the sum of its named and unnamed variables, *hdr.objSize*).

Description

This function allows you to submit an object report that creates a GemStone object and specifies the values of its instance variables. You can use this function to define a String, pointer, or NSC object with known OOPs.

The object report consists of two parts: a header (a **GciObjRepHdrSType** structure) followed by a value buffer (an array of values of the object's instance variables). For more information on object reports, see "The Object Report Structure" on page 74.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

GciNewOopUsingObjRep provides automatic byte swizzling for Float and SmallFloat objects. (For more about byte swizzling, see "Byte-Swizzling of Binary Floating-Point Values" on page 25.)

Error Conditions

In addition to general GemBuilder error conditions, this function generates an error if any of the following conditions exist:

- If `(idxSize < 0)`
- If `(idxSize + namedSize) > maxSmallInt`
- If `firstOffset > (objSize + 1)`
- For pointer objects and NSCs, if `valueBuffSize` is not divisible by 4
- If the specified `oclass` is not the OOP of a Smalltalk class object
- If the specified `oclass` and implementation (`objImpl`) do not agree
- If `objId` is a Float or SmallFloat, then `startIndex` must be one and `valueBuffSize` must be the actual size for the class of `objId`.

Note that you cannot use this function to create new special objects (instances of SmallInteger, Character, Boolean, SmallDouble, or UndefinedObject).

Example

```
void newOopUsingObjRep_example(void)
{
    int arrSize = 100;
    size_t bodySize = sizeof(OopType) * arrSize ;
    size_t rptSize = GCI_ALIGN(sizeof(GciObjRepSType) + bodySize );
    GciObjRepSType *rpt = (GciObjRepSType*) malloc(rptSize);
    if (rpt == NULL) {
        printf("malloc failure\n");
        return;
    }
    rpt->hdr.objId = OOP_NIL; // ignored by GciNewOopUsingObjRep
    rpt->hdr.oclass = OOP_CLASS_ARRAY;
    rpt->hdr.setObjImpl(GC_FORMAT_OOP);
    rpt->hdr.segmentId = WORLD_RW_SEGMENT_ID ;
    rpt->hdr.firstOffset = 1;
    rpt->hdr.namedSize = 0; // ignored by GciNewOopUsingObjRep
    rpt->hdr.setIdxSize( arrSize );
    rpt->hdr.valueBuffSize = bodySize ;

    OopType *body = rpt->valueBufferOops();
    for (int i = 0; i < arrSize; i += 1) {
        body[i] = GciI32ToOop(i);
    }
    GciNewOopUsingObjRep(rpt);

    GciErrSType errInfo;
    if (GciErr(&errInfo)) {
        printf("error category "FMT_OID" number %d, %s\n",
            errInfo.category, errInfo.number, errInfo.message);
    }
}
```

See Also

[GciNewOop](#), page 269

[GciReleaseAllOops](#), page 320

[GciReleaseGlobalOops](#), page 322

[GciTraverseObjs](#), page 408

GciNewString

Create a new String object from a C character string.

Syntax

```
oopType GciNewString(  
    const char *      cString );
```

Input Arguments

cString Pointer to a character string.

Return Value

The OOP of the newly created object.

Description

Returns a new instance of OOP_CLASS_STRING with the value that *cString* points to.

GciNewSymbol

Create a new Symbol object from a C character string.

Syntax

```
OopType GciNewSymbol(  
    const char *      cString );
```

Input Arguments

cString Pointer to a character string.

Return Value

The OOP of the newly-created object.

Description

Returns a new instance of OOP_CLASS_SYMBOL with the value that *cString* points to.

GciNewUtf8String

Create a new Unicode string object from a UTF-8 encoded C character string.

Syntax

```
OopType GciNewUtf8String(  
  const char *      unicodeCString,  
  BoolType         utf8OrUnicode );
```

Input Arguments

<i>unicodeCString</i>	Pointer to a null-terminated UTF-8 encoded character string.
<i>utf8OrUnicode</i>	Boolean indicating whether to create an instance of Utf8 or of a Unicode class. If <i>utf8OrUnicode</i> = 0, return an instance of Utf8. If <i>utf8OrUnicode</i> =1, return an instance of Unicode7, Unicode16, or Unicode32, the minimal character size required to represent <i>unicodeCString</i>

Return Value

The OOP of the newly created object.

Description

Returns a new instance of Utf8, Unicode7, Unicode16, or Unicode32, with the value that the UTF-8 encoded *unicodeCString* points to.

GciNscIncludesOop

Determines whether the given OOP is present in the specified unordered collection.

Syntax

```
BoolType GciNscIncludesOop(  
    OopType      theNsc,  
    OopType      theOop);
```

Input Arguments

<i>theNsc</i>	The unordered collection in which to search.
<i>theOop</i>	The OOP to search for.

Return Value

True if the OOP was found; false otherwise.

Description

GciNscIncludesOop searches the specified unordered collection to determine if it includes the specified object. It is equivalent to the GemStone Smalltalk method `UnorderedCollection >> includesIdentical::`.

Example

```
BoolType nscIncludesOop_example(OopType nscOop, OopType anOop)  
{  
    if (!GciIsKindOfClass(nscOop, OOP_CLASS_IDENTITY_BAG) ) {  
        printf("first argument is not an Nsc\n");  
        return FALSE; /* error: nscOop is not an NSC */  
    }  
  
    return GciNscIncludesOop(nscOop, anOop);  
}
```

See Also

- [GciAddOopToNsc](#), page 92
- [GciAddOopsToNsc](#), page 93
- [GciRemoveOopFromNsc](#), page 326
- [GciRemoveOopsFromNsc](#), page 327

GciObjExists

Determine whether or not a GemStone object exists.

Syntax

```
BoolType GciObjExists(  
    OopType      theObject );
```

Input Arguments

theObject The OOP of an object.

Return Value

Returns TRUE if *theObject* exists, FALSE otherwise.

Description

This function tests an OOP to see if the object to which it points is a valid object.

GciObjInCollection

Determine whether or not a GemStone object is in a Collection.

Syntax

```
BoolType GciObjInCollection(  
    OopType          anObj,  
    OopType          aCollection );
```

Input Arguments

<i>anObj</i>	The OOP of an object for which to check.
<i>aCollection</i>	The OOP of a collection.

Return Value

Returns TRUE if *anObj* exists in *aCollection*, FALSE otherwise.

Description

Searches the specified collection for the specified object. If *aCollection* is an NSC (such as a Bag or Set), this is a tree lookup. If *aCollection* is a kind of Array or String, this is a sequential scan. This function is equivalent to the GemStone Smalltalk method `Object >> in:`.

GciObjIsCommitted

Determine whether or not an object is committed.

Syntax

```
BoolType GciObjIsCommitted(  
    OopType          oop );
```

Input Arguments

oop The OOP of an object.

Return Value

GciObjIsCommitted returns TRUE if the object is committed, FALSE otherwise.

Description

The **GciObjIsCommitted** function determines if the given object is committed or not.

See Also

GciObjExists, page 279

GciObjRepSize_

Find the number of bytes in an object report.

Syntax

```
size_t GciObjRepSize_(anObjectReport)
const GciObjRepHdrSType *anObjectReport;
```

Input Arguments

anObjectReport A pointer to an object report returned by **GciFindObjectRep**.

Return Value

Returns the size of the specified object report.

Description

This function calculates the number of bytes in an object report. Before your application allocates memory for a copy of the object report, it can call this function to obtain the size of the report.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

```
void objRepSize_example(void)
{
    // Assumes topaz code for GciFetchVaryingOops example has run

    OopType rootObj = GciResolveSymbol("AllComponents", OOP_NIL);
    GciTravBufType *buf = GciTravBufType::malloc(8000);

    GciTraverseObjs(&rootObj, 1, buf, 10);
    GciObjRepHdrSType *rpt = buf->firstReportHdr();
    GciObjRepHdrSType *limit = buf->readLimitHdr();
    if (rpt < limit) {
        size_t reportSize = GciObjRepSize_(rpt);
        printf("size of first report is %ld bytes\n", reportSize);
    } else {
        printf("error, GciTraverseObjs returned empty buffer\n");
    }
}
```

See Also

GciFindObjRep, page 200
GciMoreTraversal, page 234
GciTraverseObjs, page 408

GciOldOopToNewOop

Return a GemStone/S 64 Bit v2.0 OopType corresponding to a GemStone/S 64 Bit v1.1 OOP.

Syntax

```
OopType GciOldOopToNewOop(  
    unsigned int      oldOop);
```

Input Arguments

oldOop The GemStone/S 64 Bit v1.1.1 OOP.

Return Value

Returns an OopType that corresponds to the GemStone/S 64 Bit v1.1 OOP. Returns OOP_ILLEGAL if the argument is not a valid GemStone/S 64 Bit v1.1 OopType.

Description

This function converts a v1.1 OOP into the equivalent v2.0 OopType. If the result is not a special OOP, this function does not check for the existence of the object.

This function returns OOP_ILLEGAL if the argument is not a legal special OOP or if the current session is not valid.

This function does not convert LargeIntegers to new SmallIntegers.

GCI_OOP_IS_BOOL

(MACRO) Determine whether or not a GemStone object represents a Boolean value.

Syntax

`GCI_OOP_IS_BOOL(theOop)`

Input Arguments

theOop The OOP of the object to test.

Return Value

A C Boolean value. Returns TRUE if the object represents a Boolean, FALSE otherwise.

Description

This macro tests to see if *theOop* represents a Boolean value.

See Also

GCI_OOP_IS_SMALL_INT, page 286

GCI_OOP_IS_SPECIAL, page 287

GCI_OOP_IS_SMALL_INT

(MACRO) Determine whether or not a GemStone object represents a SmallInteger.

Syntax

`GCI_OOP_IS_SMALL_INT(theOop)`

Input Arguments

theOop The OOP of the object to test.

Return Value

A C Boolean value. Returns TRUE if the object represents a SmallInteger, FALSE otherwise.

Description

This macro tests to see if *theOop* represents a SmallInteger.

See Also

GCI_OOP_IS_BOOL, page 285
GCI_OOP_IS_SPECIAL, page 287

GCI_OOP_IS_SPECIAL

(MACRO) Determine whether or not a GemStone object has a special representation.

Syntax

`GCI_OOP_IS_SPECIAL(theOop)`

Input Arguments

theOop The OOP of the object to test.

Return Value

A C Boolean value. Returns TRUE if the object has a special representation, FALSE otherwise.

Description

This macro tests to see if *theOop* has a special representation.

See Also

`GCI_OOP_IS_BOOL`, page 285
`GCI_OOP_IS_SMALL_INT`, page 286

GciOopToBool

Convert a Boolean object to a C Boolean value.

Syntax

```
BoolType GciOopToBool(  
    OopType      theObject );
```

Input Arguments

theObject The OOP of the Boolean object to be translated into a C Boolean value.

Return Value

Returns the C Boolean value that corresponds to the GemStone object. In case of error, this function returns FALSE.

Description

This function translates a GemStone Boolean object into the equivalent C Boolean value.

Example

```
BoolType oopToBoolExample(OopType anObj)  
{  
    BoolType aBool = GciOopToBool(anObj);  
  
    GciErrSType errInfo;  
    if (GciErr(&errInfo)) {  
        // argument was not a Boolean  
        printf("error category \"FMT_OID\" number %d, %s\n",  
            errInfo.category, errInfo.number, errInfo.message);  
        return 0;  
    }  
    return aBool;  
}
```

See Also

GCL_BOOL_TO_OOP, page 103

GCI_OOP_TO_BOOL

(MACRO) Convert a Boolean object to a C Boolean value.

Syntax

`GCI_OOP_TO_BOOL(theObject)`

Input Arguments

theObject The OOP of the Boolean object to be translated into a C Boolean value.

Return Value

A C Boolean value. Returns the C Boolean value that corresponds to the GemStone object. In case of error, this macro returns FALSE.

Description

This macro translates a GemStone Boolean object into the equivalent C Boolean value.

Provided for compatibility only. New code should use **GciOopToBool** (page 288). For the definition of `GCI_OOP_TO_BOOL`, see `$GEMSTONE/include/gcicmn.ht`

See Also

`GCI_BOOL_TO_OOP`, page 103

GciOopToChar16

Convert a Character object to a 16-bit C character value.

Syntax

```
unsigned int GciOopToChar16(  
    OopType theObject );
```

Input Arguments

theObject The OOP of theCharacter or JisCharacter object to be translated into a 16-bit C character value.

Return Value

Returns the 16-bit C character value that corresponds to the GemStone object. In case of error, this function returns zero.

Description

This function translates a GemStone Character object into the equivalent 16-bit C character value.

See Also

GciOopToChar32, page 291
GciOopToChr, page 292

GciOopToChar32

Convert a Character object to a 32-bit C character value.

Syntax

```
unsigned int GciOopToChar32(  
    OopType theObject );
```

Input Arguments

theObject The OOP of the Character or JisCharacter object to be translated into a 32-bit C character value.

Return Value

Returns the 32-bit C character value that corresponds to the GemStone object. In case of error, this function returns zero.

Description

This function translates a GemStone Character object into the equivalent 32-bit C character value.

See Also

GciOopToChar16, page 290
GciOopToChr, page 292

GciOopToChr

Convert a Character object to a C character value.

Syntax

```
char GciOopToChr(  
    OopType          theObject );
```

Input Arguments

theObject The OOP of the Character object to be translated into a C character value.

Return Value

Returns the C character value that corresponds to the GemStone object. In case of error, this function returns zero.

Attempting to convert a GemStone Character that is outside the range of C characters will result in an error.

Description

This function translates a GemStone Character object into the equivalent C character value.

Example

```
char oopToChar_example(OopType arg)  
{  
    char aChar = GciOopToChr(arg);  
  
    GciErrSType errInfo;  
    if (GciErr(&errInfo)) {  
        // argument was not a Character  
        printf("error category \"FMT_OID\" number %d, %s\n",  
            errInfo.category, errInfo.number, errInfo.message);  
        return 0;  
    }  
    return aChar;  
}
```

See Also

GCL_CHR_TO_OOP, page 107
GciOopToChar16, page 290

GCI_OOP_TO_CHR

(MACRO) Convert a Character object to a C character value.

Syntax

`GCI_OOP_TO_CHR(theObject)`

Input Arguments

theObject The OOP of the Character object to be translated into a C character value.

Return Value

The `GCI_OOP_TO_CHR` macro returns the C character value that corresponds to the GemStone object. In case of error, it returns zero.

Description

Provided for compatibility only. New code should use `GciOopToChr` or `GciOopToChar16`.

See Also

`GciOopToChar16`, page 290
`GciOopToChr`, page 292

GciOopToFlt

Convert a SmallDouble, Float, or SmallFloat object to a C double.

Syntax

```
double GciOopToFlt(  
    OopType          theObject );
```

Input Arguments

theObject The OOP of the SmallDouble, Float, or SmallFloat object to be translated into a C floating point value.

Return Value

Returns the C double precision value that corresponds to the GemStone object. In case of any error other than HOST_ERR_INEXACT_PRECISION, this function returns a PlusQuietNaN.

Description

This function translates a GemStone Float object into the equivalent C double precision value.

If your C compiler's floating point package doesn't have a representation that corresponds to one of the values listed below, **GciOopToFlt** may generate the following errors when converting GemStone Float objects into C values:

HOST_ERR_INEXACT_PRECISION

when called to convert a number whose precision exceeds that of the C double type

HOST_ERR_MAGNITUDE_OUT_OF_RANGE

when called to convert a number whose exponent is too large (or small) to be held in a C double precision value

HOST_ERR_NO_PLUS_INFINITY

when called to convert a value of positive infinity

HOST_ERR_NO_MINUS_INFINITY

when called to convert a value of negative infinity

HOST_ERR_NO_PLUS_QUIET_NAN

when called to convert a positive quiet NaN

HOST_ERR_NO_MINUS_QUIET_NAN

when called to convert a negative quiet NaN

HOST_ERR_NO_PLUS_SIGNALING_NAN

when called to convert a positive signaling NaN

HOST_ERR_NO_MINUS_SIGNALING_NAN

when called to convert a negative signaling NaN

Example

```
double oopToFlt_example(OopType arg)
{
    double d = GciOopToFlt(arg);

    GciErrSType errInfo;
    if (GciErr(&errInfo)) {
        // argument was not a Float, SmallFloat or SmallDouble
        printf("error category "FMT_OID" number %d, %s\n",
            errInfo.category, errInfo.number, errInfo.message);
        return 0.0 ;
    }
    return d;
}
```

See Also

GciFltToOop, page 203
GciGetFreeOopsEncoded, page 208

GciOopToI32

GciOopToI32_

Convert a GemStone object to a C 32-bit integer value.

Syntax

```
int GciOopToI32(  
    OopType          theObject );
```

```
int GciOopToI32_(  
    OopType          theObject,  
    BoolType *       error );
```

Input Arguments

theObject The OOP of the object to be translated into a C 32-bit integer value.

Result Arguments

error TRUE if *theObject* does not fit in the result type or is not an Integer. Otherwise unchanged.

Return Value

The **GciOopToI32** and **GciOopToI32_** functions return the C 32-bit integer value that is equivalent to the value of *theObject*.

Description

The **GciOopToI32** and **GciOopToI32_** functions translate a GemStone object into the equivalent C 32-bit integer value. The GemStone object must be a SmallInteger within the range of C integers.

Otherwise, **GciOopToI32** generates an error; **GciOopToI32_** does not generate an error, but places a boolean in the *error* argument.

See Also

GciOopToI64, page 297

GciOopToI64

GciOopToI64_

Convert a GemStone object to a C 64-bit integer value.

Syntax

```
int64 GciOopToI64(  
    OopType          theObject );
```

```
int64 GciOopToI64_(  
    OopType          theObject,  
    BoolType *      error );
```

Input Arguments

theObject The OOP of the object to be translated into a C 64-bit integer value.

Result Arguments

error TRUE if *theObject* does not fit in the result type or is not an Integer. Otherwise unchanged.

Return Value

The **GciOopToI64** and **GciOopToI64_** functions return the C int64_t value that is equivalent to the value of *theObject*.

Description

The **GciOopToI64** and **GciOopToI64_** functions translate a GemStone object into the equivalent C 64-bit integer value.

The object identified by *theObject* must be a SmallInteger or a LargeInteger. If the object is not one of these kinds, **GciOopToI64** generates an error; **GciOopToI64_** places a boolean in the argument *error*.

See Also

GciOopToI32, page 296

GciPerform

GciPerform_

Send a message to a GemStone object.

Syntax

```

OopType GciPerform(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs );

```

```

OopType GciPerform_(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    ushort           environmentId );

```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A string that defines the message selector. For keyword selectors, all keywords are concatenated in the string. (For example, at :put :).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the result of Smalltalk execution. In case of error, this function returns OOP_NIL.

Description

This function sends a message (that is, the selector along with any keyword arguments and their corresponding values) to the specified receiver (an object in the GemStone database). Because **GciPerform** calls the virtual machine, you can issue a soft break while this function is executing. For more information, see “Interrupting GemStone Execution” on page 27.

See Also

GciContinue, page 120

GciErr, page 148

GciExecute, page 149

GciNbPerform, page 254

GciPerformNoDebug, page 300

GciPerformSymDbg, page 302

GciPerformNoDebug

GciPerformNoDebug_

Send a message to a GemStone object, and temporarily disable debugging.

Syntax

```

OopType GciPerformNoDebug(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    int              flags );

OopType GciPerformNoDebug_(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    int              flags,
    ushort          environmentId );

```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A string that defines the message selector. For keyword selectors, all keywords are concatenated in the string. (For example, at :put :).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>flags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the result of Smalltalk execution. In case of error, this function returns OOP_NIL.

Description

This function is a variant of **GciPerform** that is identical to it except for just one difference. **GciPerformNoDebug** disables any breakpoints and single step points that currently exist in GemStone while the message is executing. This feature is typically used while implementing a Smalltalk debugger.

The value of *flags* may be 0 for default behavior, or can be given by using one or more of these GemBuilder mnemonics:

- `GCL_PERFORM_FLAG_ENABLE_DEBUG` makes **GciPerformNoDebug** behave like **GciPerform** with respect to debugging.
- `GCL_PERFORM_FLAG_DISABLE_ASYNC_EVENTS` disables asynchronous events.
- `GCL_PERFORM_FLAG_SINGLE_STEP` places a single-step breakpoint at the start of the method to be performed, and then executes to hit that breakpoint.

These flags can either be used alone or logically “or”ed together.

See Also

GciContinue, page 120

GciErr, page 148

GciExecute, page 149

GciNbContinue, page 240

GciNbExecute, page 245

GciNbPerform, page 254

GciNbPerformNoDebug, page 256

GciPerform, page 298

GciPerformSymDbg, page 302

GciPerformSymDbg

GciPerformSymDbg_

Send a message to a GemStone object, using a String object as a selector.

Syntax

```

OopType GciPerformSymDbg(
    OopType      receiver,
    OopType      selector,
    const OopType args[ ],
    int          numArgs,
    int          flags );

OopType GciPerformSymDbg_(
    OopType      receiver,
    OopType      selector,
    const OopType args[ ],
    int          numArgs,
    int          flags,
    ushort      environmentId );

```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	The OOP of a String object that defines the message selector. For keyword selectors, all keywords are concatenated in the String. (For example, <code>at:put:</code>).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>flags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Return Value

Returns the OOP of the result of Smalltalk execution. In case of error, this function returns OOP_NIL.

Description

If the `isNoDebug` flag is FALSE, this function is a variant of **GciPerform**; if the flag is TRUE, this function is a variant of **GciPerformNoDebug**. In either case, its operation is identical to the other function. The difference is that **GciPerformSymDbg** takes an OOP as its selector instead of a C string.

See Also

GciContinue, page 120

GciErr, page 148

GciExecute, page 149

GciPerform, page 298

GciPerformNoDebug, page 300

GciPerformTrav

GciPerformTrav_

First send a message to a GemStone object, then traverse the result of the message.

Syntax

```
BoolType GciPerformTrav(
    OopType          receiver,
    const char *     selector,
    const OopType *  args,
    int              numArgs,
    GciClampedTravArgsSType *travArgs );
```

```
BoolType GciPerformTrav_(
    OopType          receiver,
    const char *     selector,
    const OopType *  args,
    int              numArgs,
    GciClampedTravArgsSType *travArgs,
    ushort          environmentId );
```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A pointer to a character collection that defines the message selector. For keyword selectors, all keywords are concatenated in the String. (For example, at :put :).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>travArgs</i>	Pointer to an instance of GciClampedTravArgsSType . See GciClampedTrav on page 108 (page 108) for documentation of the fields in <i>travArgs</i> .
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see “environmentId” on page 80.

Result Arguments

The result of the **GciPerform** is the first object in the resulting *travBuff* field in *travArgs*.

Return Value

Returns TRUE if the result is complete and no errors occurred. Returns FALSE if the traversal is not yet completed. You can then call **GciMoreTraversal** to proceed, if there is no GciError.

Description

This function is does the equivalent of a **GciPerform** using the first four arguments, and then performs a **GciClampedTrav**, starting from the result of the perform, and doing a traversal as specified by *travArgs*. In all GemBuilder traversals, objects are traversed post depth first.

See Also

GciPerform, page 298

GciClampedTrav, page 108

GciPerformTraverse

GciPerformTraverse_

First send a message to a GemStone object, then traverse the result of the message.

Syntax

```

BoolType GciPerformTraverse(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    GciTravBufType * travBuff,
    int              level );

BoolType GciPerformTraverse_(
    OopType          receiver,
    const char       selector[ ],
    const OopType    args[ ],
    int              numArgs,
    GciTravBufType * travBuff,
    int              level,
    ushort           environmentId );

```

Input Arguments

<i>receiver</i>	The OOP of the receiver of the message.
<i>selector</i>	A pointer to a character collection that defines the message selector. For keyword selectors, all keywords are concatenated in the String. (For example, at : put :).
<i>args</i>	An array of OOPs. Each element in the array corresponds to an argument for the message. If there are no message arguments, use a dummy OOP here.
<i>numArgs</i>	The number of arguments to the message. For unary selectors (messages with no arguments), <i>numArgs</i> is zero.
<i>level</i>	Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in the OOps. When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.
<i>environmentId</i>	The compilation environment for method lookup. Used with Ruby applications, but not with Smalltalk applications. For details, see "environmentId" on page 80.

Result Arguments

<i>travBuff</i>	A buffer in which the results of the traversal are placed.
-----------------	--

Return Value

Returns FALSE if the traversal is not yet completed, but further traversal would cause the *travBuffSize* to be exceeded. If the *travBuffSize* is reached before the traversal is complete, you can then call **GciMoreTraversal** to proceed from the point where *travBuffSize* was exceeded.

Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal**.

Description

Consider the following function call:

```
BoolType performTrav_1(void)
{
    OopType receiver = GciResolveSymbol("AllComponents", OOP_NIL);
    OopType arg = GciI32ToOop(1);
    GciTravBufType *buf = GciTravBufType::malloc(8000);

    BoolType atEnd = GciPerformTraverse(receiver, "at:", &arg, 1, buf,
10);
    return atEnd;
}
```

It is equivalent to the following code:

```
BoolType performTrav_2(void)
{
    OopType receiver = GciResolveSymbol("AllComponents", OOP_NIL);
    OopType arg = GciI32ToOop(1);
    OopType obj = GciPerform(receiver, "at:", &arg, 1);

    GciTravBufType *buf = GciTravBufType::malloc(8000);
    BoolType atEnd = GciTraverseObjs(&obj, 1, buf, 10);
    return atEnd;
}
```

GciPerformTraverse provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see page 25.

See Also

- GciContinue**, page 120
- GciErr**, page 148
- GciExecute**, page 149
- GciMoreTraversal**, page 234
- GciPerform**, page 298
- GciPerformNoDebug**, page 300
- GciPerformSymDbg**, page 302
- GciStoreTrav**, page 384
- GciTraverseObjs**, page 408

GciPointerToByteArray

Given a C pointer, return a SmallInteger or ByteArray containing the value of the pointer.

Syntax

```
oopType GciPointerToByteArray(  
    void * pointer );
```

Input Arguments

pointer A C pointer.

Return Value

Returns a GemStone SmallInteger or ByteArray containing the value of the pointer.

If the argument is a 64-bit pointer aligned on an 8-byte boundary, or is a 32-bit pointer, the result is a SmallInteger. Otherwise, the result is a ByteArray.

Description

The result has a machine-dependent byte order and is not intended to be committed.

See Also

[GciByteArrayToPointer](#), page 104

GciPollForSignal

Poll GemStone for signal errors without executing any Smalltalk methods.

Syntax

```
BoolType GciPollForSignal()
```

Return Value

This function returns TRUE if a signal error or an asynchronous error exists, and FALSE otherwise.

Description

GemStone permits selective response to signal errors: RT_ERR_SIGNAL_ABORT, RT_ERR_SIGNAL_COMMIT, and RT_ERR_SIGNAL_GEMSTONE_SESSION. The default condition is to leave them all invisible. GemStone responds to each single kind of signal error only after an associated method of class System has been executed: enableSignaledAbortError, enableSignaledObjectsError, and enableSignaledGemStoneSessionError respectively.

After **GciInit** executes successfully, the GemBuilder default condition also leaves all signal errors invisible. The **GciPollForSignal** function permits GemBuilder to check signal errors manually. However, GemStone must respond to each kind of error in order for GemBuilder to respond to it. Thus, if an application calls **GciPollForSignal**, then GemBuilder can check exactly the same kinds of signal errors as GemStone responds to. If GemStone has not executed any of the appropriate System methods, then this call has no effect until it does.

GemBuilder treats any signal errors that it finds just like any other errors, through **GciErr** or the **GciLongJump** mechanism, as appropriate. Instead of checking manually, these errors can be checked automatically by calling the **GciEnableSignaledErrors** function.

GciPollForSignal also detects any asynchronous errors whenever they occur, including but not limited to the following errors: ABORT_ERR_LOST_OT_ROOT, GS_ERR_SHRPC_CONNECTION_FAILURE, GS_ERR_STN_NET_LOST, GS_ERR_STN_SHUTDOWN, and GS_ERR_SESSION_SHUTDOWN.

See Also

GciEnableSignaledErrors, page 145

GciErr, page 148

GciPollSocketForRead

Wait for the specified socket to be read-ready.

Syntax

```
int GciPollSocketForRead(  
    int          socketFd,  
    int          timeoutMs);
```

Return Value

This function returns 0 if timed out, 1 if socket is ready for read, and an `int < 0` if an error occurred. The result in this cases is the negated `errno` value.

Description

Wait *timeoutMs* milliseconds for the specified socket to be read-ready or to have an error.

This function retries the poll on `EINTR`, even if a `SIGTERM` was received.

Thread safe function, has no relationship to the current GCI session .

See Also

[GciEnableSignaledErrors](#), page 145

[GciErr](#), page 148

GciPopErrJump

Discard a previously saved error jump buffer.

Syntax

```
void GciPopErrJump(  
    GciJmpBufSType *    jumpBuffer );
```

Input Arguments

jumpBuffer A pointer to a jump buffer specified in an earlier call to **GciPushErrJump**.

Description

This function discards one or more jump buffers that were saved with earlier calls to **GciPushErrJump**. Your program must call this function when a saved execution environment is no longer useful for error handling.

GemBuilder maintains a stack of error jump buffers. After your program calls **GciPopErrJump**, the jump buffer at the top of the stack will be used for subsequent GemBuilder error handling. If no jump buffers remain, your program will need to call **GciErr** and test for errors locally.

To pop multiple jump buffers in a single call to **GciPopErrJump**, specify the *jumpBuffer* argument from an earlier call to **GciPushErrJump**. See the following example.

Example

```
void popErr_example(void)  
{  
    GciJmpBufSType jumpBuff1, jumpBuff2, jumpBuff3, jumpBuff4;  
  
    GciPushErrJump(&jumpBuff1);  
  
    GciPushErrJump(&jumpBuff2);  
  
    GciPushErrJump(&jumpBuff3);  
  
    GciPushErrJump(&jumpBuff4);  
  
    GciPopErrJump(&jumpBuff1); /* pops buffers 1-4 */  
}
```

See Also

GciErr, page 148
GciPushErrJump, page 314
GciSetErrJump, page 343
GciLongJmp, page 233

GciProcessDeferredUpdates_

Process deferred updates to objects that do not allow direct structural update.

Syntax

```
int64 GciProcessDeferredUpdates_()
```

Return Value

Returns the number of objects that had deferred updates.

Description

This function processes updates to instances of classes that have the `noStructuralUpdate` bit set, including `AbstractDictionary`, `Bag`, `Set`, and their subclasses. After operations that modify an instance of one of these classes, either **GciProcessDeferredUpdates_** must be called, or the final **GciStoreTrav** must have `GCI_STORE_TRAV_FINISH_UPDATES` set.

The following GemBuilder calls operate on instances whose classes have `noStructuralUpdate` set: **GciCreateOopObj**, **GciStoreTrav**, **GciStore...Oops**, **GciAdd...Oops**, **GciReplace...Oops**. Behavior of other GemBuilder update calls on such instances is undefined.

An attempt to commit automatically executes a deferred update.

Executing a deferred update before all forward references are resolved can produce errors that require the application to recover by doing a **GciAbort** or **GciLogout**.

An OOP buffer used to update the varying portion of an object with `noStructuralUpdate` must contain the OOPs to be added to the varying portion of the object, with two exceptions:

- If the object is a kind of `KeyValueDictionary` that does not store `Associations`, the buffer must contain (key, value) pairs.
- If the object is a kind of `AbstractDictionary` that stores `Associations` or (key, `Association`) pairs, the value buffer must contain `Associations`.

See Also

GciStoreTrav, page 384

GciProduct

Return an 8-bit unsigned integer that indicates the GemStone/S product.

Syntax

```
unsigned char GciProduct();
```

Return Value

Returns an 8-bit unsigned integer indicating the GemStone/S product to which the client library belongs. Currently-defined integers are:

- 1 – GemStone/S
- 2 – GemStone/S 2G
- 3 – GemStone/S 64 Bit

Description

GciProduct allows a GemBuilder client to determine which GemStone/S product it is talking to. Combined with GciVersion, it allows the client to adapt to differences between GemBuilder features across different products and versions.

Although GciProduct can be used by any GemBuilder client, it is specifically provided for the use of GemBuilder for Smalltalk.

Future products in the GemStone/S line will be assigned integers beginning with 4.

The integer zero is reserved, and will never be assigned to any product.

See Also

[GciVersion](#), page 415

GciPushErrJump

Associate GemBuilder error handling with a jump buffer by pushing a jump buffer onto the stack.

Syntax

```
void GciPushErrJump(  
    GciJumpBufSType *    jumpBuffer );
```

Result Arguments

jumpBuffer A pointer to a jump buffer, as described below. The *jumpBuffer* must have been initialized by passing it as the argument to the macro `Gci_SETJMP`.

Description

Associate GemBuilder error handling with a jump buffer by pushing a jump buffer onto the stack.

This function allows your application program to take advantage of the `setjmp/longjmp` style of error-handling mechanism from within any GemBuilder function call. However, you cannot use this mechanism to handle errors within **GciPushErrJump** itself, or within the related functions **GciPopErrJump** and **GciSetErrJump**.

Rather than using `setjmp` and `longjmp` directly, this style of error handling in GemBuilder requires you to use **Gci_SETJMP** and **GciLongJump**.

When your program calls **Gci_SETJMP**, the context of the C environment is saved in a jump buffer that you designate. To associate subsequent GemBuilder error handling with that jump buffer, you would then call **GciPushErrJump**.

GemBuilder maintains a stack of up to 20 error jump buffers. A buffer is pushed onto the stack when **GciPushErrJump** is called, and popped when **GciPopErrJump** is called. When an error occurs during a GemBuilder call, the GemBuilder implementation calls **GciLongJump** using the buffer currently at the top of GemBuilder's error jump stack, and pops that buffer from the stack.

For functions with local error recovery, your program can call **GciSetErrJump** to temporarily disable the **GciLongJump** mechanism (and to re-enable it afterwards).

Whenever the jump stack is empty, the application must use **GciErr** to poll for GBC errors.

Example

For an example of how **GciPushErrJump** is used, see **GciPopErrJump** on page 311.

See Also

GciErr, page 148

GciPopErrJump, page 311

GciSetErrJump, page 343

GciRaiseException

Signal an error, synchronously, within a user action.

Syntax

```
void GciRaiseException(  
    const GciErrSType *    err);
```

Input Arguments

err A pointer to the error type to raise.

Description

When executed from within a user action, this function raises an exception and passes the given error to the error signaling mechanism, causing control to return to Smalltalk. In order to signal an error on the Smalltalk client, this function must be invoked.

This function has no effect when executed outside of a user action.

Example

```
OopType res = GciNewOop(anOopClass);  
GciErrSType theErr;  
// the term res==OOP_NIL is a performance optimization  
if (res == OOP_NIL && GciErr(&theErr)) {  
    GciRaiseException(&theErr);  
}
```

See Also

[GciErr](#), page 148

GciReadSharedCounter

Lock and fetch the value of a shared counter.

Syntax

```
BoolType GciReadSharedCounter(  
    int counterIdx,  
    int64_t * value);
```

Input Arguments

counterIdx The offset into the shared counters array of the value to fetch.

Result Arguments

value Pointer to a value that indicates the value at this shared counter.

Return Value

Returns a C Boolean value indicating whether the value was successfully read. Returns TRUE if successful, FALSE if an error occurred.

Description

Lock the shared counter indicated by *counterIdx*, and fetch its value. The contents of the *value* pointer will be set to the value of the shared counter.

Not supported for remote GCI interfaces.

See Also

[GciFetchNumSharedCounters](#), page 177
[GciDecSharedCounter](#), page 134
[GciIncSharedCounter](#), page 214
[GciSetSharedCounter](#), page 353
[GciReadSharedCounterNoLock](#), page 317
[GciFetchSharedCounterValuesNoLock](#), page 190

GciReadSharedCounterNoLock

Fetch the value of a shared counter without locking it.

Syntax

```
BoolType GciReadSharedCounterNoLock(  
    int counterIdx,  
    int64_t * value);
```

Input Arguments

counterIdx The offset into the shared counters array of the value to fetch.

Result Arguments

value Pointer to a value at this shared counter.

Return Value

Returns a C Boolean value indicating whether the value was successfully read. Returns TRUE if successful, FALSE if an error occurred.

Description

Fetch the value of the shared counter indicated by *counterIdx*. The contents of the *value* pointer will be set to the value of the shared counter. This function is faster than **GciReadSharedCounter**, but may be less accurate.

Not supported for remote GCI interfaces.

See Also

GciFetchNumSharedCounters, page 177
GciDecSharedCounter, page 134
GciIncSharedCounter, page 214
GciSetSharedCounter, page 353
GciReadSharedCounter, page 316
GciFetchSharedCounterValuesNoLock, page 190

GciRealloc

Reallocates memory.

Syntax

```
void* GciRealloc(  
    void *           p,  
    size_t          length,  
    int             lineNumber,  
    const char *    fileName,  
    );
```

Description

Return NULL if the underlying realloc() fails.

GciReleaseAllGlobalOops

Remove all OOPS from the PureExportSet, making these objects eligible for garbage collection.

Syntax

```
void GciReleaseAllGlobalOops( )
```

Description

The **GciReleaseAllGlobalOops** function removes all OOPs from the PureExportSet, thus permitting GemStone to consider removing them as a result of garbage collection. Objects that are referenced from persistent objects are not removed during garbage collection, even if they are not in PureExportSet. If invoked from a user action, this function does not affect the user action's export set.

GciReleaseAllGlobalOops is similar to **GciReleaseAllOops**, with the exception that OOPs are removed from the PureExportSet regardless of whether it is called from within a user action or not.

The **GciSaveGlobalObjs** or **GciSaveGlobalObjs** functions may be used to make objects ineligible for garbage collection. Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, and **GciExecute...** functions are automatically added to the PureExportSet. You must release those objects explicitly if they are to be eligible for garbage collection.

CAUTION

Before releasing all objects, be sure that you do not need to retain any of them for any reason.

See Also

“Garbage Collection” on page 40
GciReleaseAllOops, page 320
GciReleaseGlobalOops, page 322
GciSaveGlobalObjs, page 337
GciSaveObjs, page 338

GciReleaseAllOops

Remove all OOPS from the PureExportSet, or if in a user action, from the user action's export set, making these objects eligible for garbage collection.

Syntax

```
void GciReleaseAllOops()
```

Description

The **GciReleaseAllOops** function removes all OOPs from the applicable export set, thus permitting GemStone to consider removing them as a result of garbage collection. If called from within a user action, **GciReleaseAllOops** releases only those objects that have been saved since the beginning of the user action and are therefore in the user action's export set. If not called from within a user action, **GciReleaseAllOops** removes all OOPs from the PureExportSet. To remove all objects from the PureExportSet, regardless of user action context, use **GciReleaseAllGlobalOops**.

Objects that are referenced by persistent objects are not removed during garbage collection, even if they are not in an export set. It is typical usage to call **GciReleaseAllOops** after successfully committing a transaction.

The **GciSaveObjs** or **GciSaveGlobalObjs** functions may be used to make objects ineligible for garbage collection. Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, and **GciExecute...** functions are automatically ineligible. You must release those objects explicitly if they are to be eligible.

CAUTION

Before releasing all objects, be sure that you do not need to retain any of them for any reason.

See Also

“Garbage Collection” on page 40
GciReleaseAllGlobalOops, page 319
GciReleaseGlobalOops, page 322
GciReleaseOops, page 323
GciSaveGlobalObjs, page 337
GciSaveObjs, page 338

GciReleaseAllTrackedOops

Clear the GciTrackedObjs set, making all tracked OOPs eligible for garbage collection.

Syntax

```
void GciReleaseAllTrackedOops( )
```

Description

The **GciReleaseAllTrackedOops** function removes all OOPs from the user session's GciTrackedObjs set, thus making them eligible to be garbage collected. This function does not affect the export sets; objects that are also in an export set will remain protected from garbage collection.

CAUTION

Before releasing any of your objects, be sure that you do not need to retain them for any reason.

See Also

GciHiddenSetIncludesOop, page 211

GciReleaseAllGlobalOops, page 319

GciReleaseAllOops, page 320

GciReleaseTrackedOops, page 325

GciSaveAndTrackObjs, page 336

GciReleaseGlobalOops

Remove an array of GemStone OOPs from the PureExportSet, making them eligible for garbage collection.

Syntax

```
void GciReleaseGlobalOops(  
    const OopType    theOops[],  
    int              numOops);
```

Input Arguments

<i>theOops</i>	An array of OOPs. Each element of the array corresponds to an object to be released.
<i>numOops</i>	The number of elements in <i>theOops</i> .

Description

The **GciReleaseGlobalOops** function removes the specified OOPs from the PureExportSet, thus making them eligible to be garbage collected.

This function differs from **GciReleaseOops** in that it operates the same if invoked from within a user action or not.

The **GciSaveObjs** or **GciSaveGlobalObjs** functions may be used to make objects ineligible for garbage collection. Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, and **GciExecute...** functions are automatically ineligible. You must release those objects explicitly if they are to be eligible.

CAUTION

Before releasing any of your objects, be sure that you do not need to retain them for any reason.

See Also

“Garbage Collection” on page 40
GciReleaseAllGlobalOops, page 319
GciReleaseOops, page 323
GciSaveGlobalObjs, page 337

GciReleaseOops

Remove an array of GemStone OOPs from the PureExportSet, or if in a user action, remove them from the user action's export set, making them eligible for garbage collection.

Syntax

```
void GciReleaseOops(
    const OopType    theOops[ ],
    int              numOops );
```

Input Arguments

<i>theOops</i>	An array of OOPs. Each element of the array corresponds to an object to be released.
<i>numOops</i>	The number of elements in <i>theOops</i> .

Description

The **GciReleaseOops** function removes the specified OOPs from the applicable export set, thus making them eligible to be garbage collected. If invoked from within a user action, the specified OOPs are removed from the user action's export set, otherwise the OOPs are removed from the PureExportSet.

To remove OOPs from the PureExportSet, regardless of user action context, use **GciReleaseGlobalOops**.

The **GciSaveObjs** or **GciSaveGlobalObjs** functions may be used to make objects ineligible for garbage collection. Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, and **GciExecute...** functions are automatically ineligible. You must release those objects explicitly if they are to be eligible.

CAUTION

Before releasing any of your objects, be sure that you do not need to retain them for any reason.

Example

```
void releaseOops_example(void)
{
    // assumes topaz code for GciFetchVaryingOops example has run.

    OopType oClass = GciResolveSymbol("Component", OOP_NIL);

    OopType namedIvs[3];
    namedIvs[0] = GciI32ToOop(5699); // a SmallInteger , don't need to
    release
    namedIvs[1] = GciNewString("cfm56-99");
    namedIvs[2] = GciFltToOop(9.0e6); // a Float or SmallDouble

    OopType newComp = GciNewOop(oClass);
```

```
GciStoreOops(newComp, 1, namedIvs, 3);

OopType oColl = GciResolveSymbol("AllComponents", OOP_NIL);
GciAddOopToNsc(oColl, newComp); // new objects now reachable from
AllComponents

// release newly created objects so that if aComp is removed from
// AllComponents by other application code, these new objects can
// be garbage collected.
OopType releaseBuf[3];
releaseBuf[0] = namedIvs[1]; // a String
releaseBuf[1] = namedIvs[1]; // might be a Float
releaseBuf[2] = newComp; // a Component
GciReleaseOops(releaseBuf, 3);
}
```

See Also

"Garbage Collection" on page 40
GciReleaseAllGlobalOops, page 319
GciReleaseAllOops, page 320
GciReleaseGlobalOops, page 322
GciSaveGlobalObjs, page 337
GciSaveObjs, page 338

GciReleaseTrackedOops

Remove an array of OOPs from the GciTrackedObjs set, making them eligible for garbage collection.

Syntax

```
void GciReleaseTrackedOops(  
    const OopType    theOops [ ],  
    int              numOops );
```

Input Arguments

<i>theOops</i>	An array of OOPs. Each element of the array corresponds to an object to be released.
<i>numOops</i>	The number of elements in <i>theOops</i> .

Description

The **GciReleaseTrackedOops** function removes the specified OOPs from the user session's GciTrackedObjs set, thus making them eligible to be garbage collected. This function does not affect the export sets; objects that also appear in an export set will remain protected from garbage collection.

CAUTION

Before releasing any of your objects, be sure that you do not need to retain them for any reason.

See Also

GciHiddenSetIncludesOop, page 211
GciReleaseAllTrackedOops, page 321
GciSaveAndTrackObjs, page 336
GciTrackedObjsInit, page 407

GciRemoveOopFromNsc

Remove an OOP from an NSC.

Syntax

```
BoolType GciRemoveOopFromNsc(  
    OopType      theNsc,  
    OopType      theOop);
```

Input Arguments

<i>theNsc</i>	The OOP of the NSC from which to remove an OOP.
<i>theOop</i>	The OOP of the object to be removed.

Result Arguments

<i>theNsc</i>	The OOP of the modified NSC.
---------------	------------------------------

Return Value

Returns FALSE if *theOop* was not present in the NSC. Returns TRUE if *theOop* was present.

Description

This function removes an OOP from the unordered variables of an NSC, using structural access.

Example

```
BoolType removeOop_example(void)  
{  
    // assumes topaz code for GciFetchVaryingOop has run  
    OopType aComponent = GciExecuteStr(  
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);  
  
    OopType aColl = GciResolveSymbol("AllComponents", OOP_NIL);  
  
    BoolType wasPresent = GciRemoveOopFromNsc(aColl, aComponent);  
  
    GciReleaseOops(&aComponent, 1); // release because it was a result  
        // from an execute  
    return wasPresent;  
}
```

See Also

[GciAddOopToNsc](#), page 92
[GciAddOopsToNsc](#), page 93

[GciNscIncludesOop](#), page 278
[GciRemoveOopsFromNsc](#), page 327

GciRemoveOopsFromNsc

Remove one or more OOPs from an NSC.

Syntax

```
BoolType GciRemoveOopsFromNsc(
    OopType      theNsc,
    const OopType theOops[ ],
    int          numOops );
```

Input Arguments

<i>theNsc</i>	The OOP of the NSC from which to remove the OOPs.
<i>theOops</i>	The array of OOPs to be removed from the NSC.
<i>numOops</i>	The number of OOPs to remove.

Result Arguments

<i>theNsc</i>	The OOP of the modified NSC.
---------------	------------------------------

Return Value

Returns FALSE if any element of *theOops* was not present in the NSC. Returns TRUE if all elements of *theOops* were present in the NSC.

Description

This function removes multiple OOPs from the unordered variables of an NSC, using structural access. If any individual OOP is not present in the NSC, this function returns FALSE, but it still removes all OOPs that it finds in the NSC.

Example

```
BoolType removeOops_example(void)
{
    // assumes topaz code for GciFetchVaryingOop has run
    OopType subColl = GciExecuteStr(
        "AllComponents select:[i|i partNumber > 1000 ]", OOP_NIL);

    OopType buf[10];
    int numRet = GciFetchVaryingOops(subColl, 1, buf, 10);
    // buf contains at most 10 components with partNumber > 1000 .

    OopType aColl = GciResolveSymbol("AllComponents", OOP_NIL);
    BoolType allPresent = GciRemoveOopsFromNsc(aColl, buf, numRet);

    GciReleaseOops(&subColl, 1); // release because it was result of an
    execute
```

```
    return allPresent;  
}
```

See Also

[GciAddOopToNsc](#), page 92
[GciAddOopsToNsc](#), page 93
[GciNscIncludesOop](#), page 278
[GciRemoveOopFromNsc](#), page 326

GciReplaceOops

Replace all instance variables in a GemStone object.

Syntax

```
void GciReplaceOops(  
    OopType          theObj,  
    const OopType    theOops[],  
    int              numOops );
```

Input Arguments

<i>theOops</i>	The array of OOPs used as the replacements.
<i>numOops</i>	The number of OOPs in <i>theOops</i> .

Result Arguments

<i>theObj</i>	The object whose instance variables are replaced.
---------------	---

Description

GciReplaceOops uses structural access to replace *all* the instance variables in the object. However, it does so in a context that is external to the object. Hence, it completely ignores private named instance variables in its operation.

If *theObj* is of fixed size, then it is an error for *numOops* to be of a different size. If *theObj* is of a variable size, then it is an error for *numOops* to be of a size smaller than the number of named instance variables (*namedSize*) of the object. For variable-sized objects, **GciReplaceOops** resets the number of unnamed variables to *numOops* - *namedSize*.

GciReplaceOops is not recommended for use with variable-sized objects unless they are indexable or are NSCs. Other variable-sized objects, such as KeyValue dictionaries, do not store values at fixed offsets.

See Also

GciReplaceVaryingOops, page 330

GciStoreIdxOops, page 369

GciStoreNamedOops, page 373

GciStoreOops, page 377

GciReplaceVaryingOops

Replace all unnamed instance variables in an NSC object.

Syntax

```
void GciReplaceVaryingOops(  
    OopsType      theNsc,  
    const OopsType theOops[ ],  
    int           numOops );
```

Input Arguments

theOops The array of objects used as the replacements.
numOops The number of objects in *theOops*.

Result Arguments

theNsc The NSC object whose unnamed instance variables are replaced.

Description

GciReplaceVaryingOops uses structural access to replace all unnamed instance variables in the NSC object.

See Also

GciReplaceOops, page 329
GciStoreIdxOops, page 369
GciStoreNamedOops, page 373
GciStoreOops, page 377

GciResolveSymbol

Find the OOP of the object to which a symbol name refers, in the context of the current session's user profile.

Syntax

```
OopType GciResolveSymbol(  
    const char *      cString,  
    OopType           symbolList );
```

Input Arguments

<i>cString</i>	The name of a symbol as a character string.
<i>symbolList</i>	The OOP of an instance of OOP_CLASS_SYMBOL_LIST or OOP_NIL.

Return Value

The OOP of the object that corresponds to the specified symbol.

Description

Attempts to resolve the symbol name *cString* using symbol list *symbolList*. If *symbolList* is OOP_NIL, this function searches the symbol list in the user's UserProfile. If the symbol is not found or an error is generated, the result is OOP_ILLEGAL. If result is OOP_ILLEGAL and **GciErr** reports no error, then the symbol could not be resolved using the given *symbolList*. If an error such as an authorization error occurs, the result is OOP_ILLEGAL and the error is accessible by **GciErr**.

This function is similar to **GciResolveSymbolObj**, except that the symbol argument is a C string instead of an object identifier.

See Also

GciResolveSymbolObj, page 332

GciResolveSymbolObj

Find the OOP of the Symbol to which a String object refers, in the context of the current session's user profile.

Syntax

```
OopType GciResolveSymbolObj(  
    OopType          aStringObj,  
    OopType          symbolList );
```

Input Arguments

<i>aStringObj</i>	The OOP of a kind of String or MultiByteString. That is, this object's class must be OOP_CLASS_STRING or a subclass thereof.
<i>symbolList</i>	The OOP of an instance of OOP_CLASS_SYMBOL_LIST, OOP_CLASS_MultiByteString, or OOP_NIL.

Return Value

The OOP of the Symbol that corresponds to the specified String.

Description

Attempts to resolve *aStringObj* using symbol list *symbolList*. If *symbolList* is OOP_NIL, this function searches the symbol list in the user's UserProfile. If the symbol is not found or an error is generated, the result is OOP_ILLEGAL. If the result is OOP_ILLEGAL and **GciErr** reports no error, then the symbol could not be resolved using the given *symbolList*. If an error such as an authorization error occurs, the result is OOP_ILLEGAL and the error is accessible by **GciErr**.

This function is similar to **GciResolveSymbol**, except that the symbol argument is an object identifier for a String instead of a C string.

See Also

GciResolveSymbol, page 331

GciRtlIsLoaded

Report whether a GemBuilder library is loaded.

Syntax

```
BoolType GciRtlIsLoaded()
```

Return Value

Returns TRUE if a GemBuilder library is loaded and FALSE if not.

Description

The **GciRtlIsLoaded** function reports whether an executable has loaded one of the versions of GemBuilder. The GemBuilder library files are dynamically loaded at run time. See “The GemBuilder for C Shared Libraries” on page 44 for more information.

See Also

GciRtlLoad, page 334

GciRtlUnload, page 335

GciRtlLoad

Load a GemBuilder library.

Syntax

```
BoolType GciRtlLoad(  
    BoolType          useRpc,  
    const char *      path,  
    char              errBuf[ ],  
    size_t            errBufSize );
```

Input Arguments

<i>useRpc</i>	A flag to specify the RPC or linked version of GemBuilder.
<i>path</i>	A list of directories (separated by ;) to search for the GemBuilder library.
<i>errBuf</i>	A buffer to store any error message.
<i>errBufSize</i>	The size of <i>errBuf</i> .

Return Value

Returns TRUE if a GemBuilder library loads successfully. If the load fails, the return value is FALSE, and a null-terminated error message is stored in *errBuf*, unless *errBuf* is NULL.

Description

The **GciRtlLoad** function attempts to load one of the GemBuilder libraries. If *useRpc* is TRUE, the RPC version of GemBuilder is loaded. If *useRpc* is FALSE, the linked version of GemBuilder is loaded. See “The GemBuilder for C Shared Libraries” on page 44 for more information.

If *path* is not NULL, it must point to a list of directories to search for the library to load. If *path* is NULL, then a default path is searched.

If a GemBuilder library is already loaded, the call fails.

See Also

GciRtlIsLoaded, page 333

GciRtlUnload, page 335

GciRtlUnload

Unload a GemBuilder library.

Syntax

```
void GciRtlUnload()
```

Description

The **GciRtlUnload** function causes the library loaded by **GciRtlLoad** to be unloaded. Once the current library is unloaded, **GciRtlLoad** can be called again to load a different GemBuilder library. See “The GemBuilder for C Shared Libraries” on page 44 for more information.

See Also

GciRtlLoad, page 334

GciRtlIsLoaded, page 333

GciSaveAndTrackObjs

Add objects to GemStone's internal GciTrackedObjs set to prevent them from being garbage collected.

Syntax

```
void GciSaveAndTrackObjs(  
    const OopType    theOops [ ],  
    int              numOops );
```

Input Arguments

<i>theOops</i>	An array of OOPs.
<i>numOops</i>	The number of elements in theOops.

Description

The **GciSaveAndTrackOops** function adds the specified OOPS to GemStone's GciTrackedObjs set. This prevents the GemStone garbage collector from causing the objects to disappear during a session if they become unreferenced, and enables changes to these objects to show up in the TrackedDirtyObjs set.

This function does *not* cause the objects to be referenced from a permanent object; there is no guarantee that they will be saved to disk at commit.

The results of **GciNew...**, **GciCreate...**, **GciSend...**, **GciPerform...**, and **GciExecute...** calls are automatically added to the export set, which also prevents them from being garbage collected.

This function may only be called after **GciTrackedObjsInit** has been executed.

You can use **GciReleaseTrackedOops** or **GciReleaseAllTrackedOops** calls to cancel the effect of a **GciSaveAndTrackOops** call, thereby making objects eligible for garbage collection. Objects that have been added to the GciTrackedObjs set and have been modified can be retrieved using **GciTrackedDirtyObjs**, **GciDirtySaveObjs**, or **GciTrackedObjsFetchAllDirty**.

See Also

- GciHiddenSetIncludesOop**, page 211
- GciDirtySaveObjs**, page 138
- GciDirtyTrackedObjs**, page 140
- GciReleaseAllTrackedOops**, page 321
- GciReleaseTrackedOops**, page 325
- GciTrackedObjsInit**, page 407
- GciTrackedObjsFetchAllDirty**, page 405

GciSaveGlobalObjs

Add an array of OOPs to the PureExportSet, making them ineligible for garbage collection.

Syntax

```
void GciSaveGlobalObjs(  
    const OopType    theOops [ ],  
    int              numOops );
```

Input Arguments

<i>theOops</i>	An array of OOPs.
<i>numOops</i>	The number of elements in <i>theOops</i> .

Description

The **GciSaveGlobalObjs** function places the specified OOPs in the PureExportSet, thus preventing GemStone from removing them as a result of garbage collection. **GciSaveGlobalObjs** can add any OOP to the PureExportSet. It differs from **GciSaveObjs** in that OOPs are placed in the PureExportSet regardless of user action context.

The **GciSaveGlobalObjs** function does *not* itself make objects persistent, and it does *not* create a reference to them from a persistent object so that the next commit operation will try to do so either. It only protects them from garbage collection.

Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, **GciExecute...**, and **GciResolve...** functions are automatically added to the export set. The **GciRelease...** functions may be used to make objects eligible for garbage collection.

See Also

- “Garbage Collection” on page 40
- GciReleaseAllGlobalOops, page 319
- GciReleaseAllOops, page 320
- GciReleaseGlobalOops, page 322
- GciReleaseOops, page 323
- GciSaveObjs, page 338

GciSaveObjs

Add an array of OOPs to the PureExportSet, or if in a user action to the user action's export set, making them ineligible for garbage collection.

Syntax

```
void GciSaveObjs(  
    const OopType      theOops [ ],  
    int                numOops );
```

Input Arguments

<i>theOops</i>	An array of OOPs.
<i>numOops</i>	The number of elements in <i>theOops</i> .

Description

The **GciSaveObjs** function places the specified OOPs in the applicable export set, thus preventing GemStone from removing them as a result of garbage collection. If invoked from within a user action, the OOPs are added to the user action's export set; otherwise the OOPs are added to the PureExportSet. To add OOPS to the PureExportSet, regardless of the user action context, use **GciSaveGlobalObjs**. **GciSaveObjs** can add any OOP to the export set.

The **GciSaveObjs** function does *not* itself make objects persistent, and it does *not* create a reference to them from a persistent object so that the next commit operation will try to do so either. It only protects them from garbage collection.

Note that results of the **GciNew...**, **GciCreate...**, **GciPerform...**, **GciExecute...**, and **GciResolve...** functions are automatically added to the export set. The **GciRelease...** functions may be used to make objects eligible for garbage collection.

See Also

- "Garbage Collection" on page 40
- GciReleaseGlobalOops, page 322
- GciReleaseOops, page 323
- GciSaveGlobalObjs, page 337

GciServerIsBigEndian

Determine whether or not the server process is big-endian.

Syntax

```
BoolType GciServerIsBigEndian();
```

Return Value

Returns TRUE if the session is RPC and the server process is big-endian, or if the session is linked and this process is big-endian. Returns FALSE otherwise.

Description

This function determines whether the server process is big-endian. If the current session is invalid, this generates an error.

GciSessionIsRemote

Determine whether or not the current session is using a Gem on another machine.

Syntax

```
BoolType GciSessionIsRemote()
```

Return Value

Returns TRUE if the current GemBuilder session is connected to a remote Gem. It returns FALSE if the current GemBuilder session is connected to a linked Gem.

GciSessionIsRemote raises an error if the current session is invalid.

GciSetCacheName_

Set the name that a linked application will be known by in the shared cache.

Syntax

```
BoolType GciSetCacheName_(  
    const char *      name );
```

Input Arguments

name The processName reported by System cacheStatistics.

Return Value

Returns FALSE if called before GciInit and GciIsRemote returns FALSE.

Description

This function sets the name that a linked application will be known by in the shared cache. This function has no effect if GciIsRemote returns TRUE.

GciSetDynLib

Swap the byte order of an array of uint.

Syntax

```
void GciSetDynLib(  
    void * handle );
```

Description

Used by the topaz.c main program to save the result of dlopen() which loaded the GCI shared library.

GciSetErrJump

Enable or disable the current error handler.

Syntax

```
BoolType GciSetErrJump(  
    BoolType          aBoolean );
```

Input Arguments

aBoolean TRUE enables error jumps to the execution environment saved by the most recent **GciPushErrJump**; FALSE disables error jumps.

Return Value

Returns TRUE if error handling was previously enabled for the jump buffer at the top of the error jump stack. Returns FALSE if error handling was previously disabled. If your program has no buffers saved in its error jump stack, this function returns FALSE. (This function cannot generate an error.)

For most GemBuilder functions, calling **GciErr** after a successful function call will return zero (that is, false). In such cases, the **GciErrSType** error report structure will contain some default values. (See **GciErr** on page 148 for details.) However, a successful call to **GciSetErrJump** does not alter any previously existing error report information. That is, calling **GciErr** after a successful call to **GciSetErrJump** will return the same error information that was present before this function was called.

Description

This function enables or disables the error handler at the top of GemBuilder's error jump stack.

Example

```
void setErrJump_example(void)  
{  
    GciJumpBufSType jumpBuf1;  
    GciPushErrJump(&jumpBuf1);  
  
    if (Gci_SETJMP(&jumpBuf1)) {  
        GciErrSType errInfo;  
        if (GciErr(&errInfo)) {  
            printf("LONGJMP, error category \"FMT_OID\" number %d, %s\n",  
                errInfo.category, errInfo.number, errInfo.message);  
        } else {  
            printf("GCI longjmp, but no error found\n"); // should not  
            happen  
        }  
        GciPopErrJump(&jumpBuf1);  
        return;  
    }  
}
```

```
    }
    BoolType prevVal = GciSetErrJump(FALSE); // disable error jumps
    printf("error jumps previously %s\n", prevVal ? "enabled" :
"disabled");

    OopType oRcvr = GciI32ToOop(3);
    GciPerform(oRcvr, "frob", NULL, 0); // expect does-not-understand
error
    GciErrSType errInfo;
    if (GciErr(&errInfo)) {
        printf("error category "FMT_OID" number %d, %s\n",
            errInfo.category, errInfo.number, errInfo.message);
    } else {
        printf("expected error but found none\n");
    }

    GciSetErrJump(TRUE);
    GciPerform(oRcvr, "frob", NULL, 0); // expect a longjmp

    printf("GCI longjmp did not happen\n"); // should not reach here
}
```

See Also

[GciErr](#), page 148

[GciPopErrJump](#), page 311

[GciPushErrJump](#), page 314

GciSetHaltOnError

Halt the current session when a specified error occurs.

Syntax

```
int GciSetHaltOnError(  
    int errNum );
```

Input Arguments

errNum When this error occurs, halt the current session.

Return Value

Returns the previous error number on which the session was to halt.

Description

The **GciSetHaltOnError** function causes the current session to halt for internal debugging when the specified GemBuilder error occurs. When *errNum* is zero, halt on error is disabled.

See Also

GciErr, page 148

Gci_SETJMP

(MACRO) Save a jump buffer in GemBuilder's error jump stack.

Syntax

```
void Gci_SETJMP(  
    GciJumpBufSType *    jumpBuffer );
```

Input Arguments

jumpBuffer A pointer to a jump buffer.

Description

When your program calls this macro, the context of the C environment is saved in a jump buffer that you designate. GemBuilder maintains a stack of up to 20 error jump buffers.

Except for the difference in argument type, the semantics of this function are the same as for `setjmp()` on Solaris and `_setjmp()` on HP-UX.

See Also

GciErr, page 148

GciLongJump, page 233

GciPopErrJump, page 311

GciPushErrJump, page 314

GciSetErrJump, page 343

GciSetNet

Set network parameters for connecting the user to the Gem and Stone processes.

Syntax

```
void GciSetNet(
    const char      StoneName[ ],
    const char      HostUserId[ ],
    const char      HostPassword[ ],
    const char      GemService[ ] );
```

Input Arguments

<i>StoneName</i>	Network resource string for the database monitor process.
<i>HostUserId</i>	UNIX host login name.
<i>HostPassword</i>	password of the UNIX user.
<i>GemService</i>	Network resource string for the GemStone service.

Description

Your application, your GemStone session (Gem), and the database monitor (Stone) can all run in separate processes, on separate machines in your network. **GciSetNet** specifies the network parameters that are used to connect the current user to GemStone on the host, whenever **GciLogin** is called. Network resource strings specify the information needed to establish communications between these processes. See the *System Administration Guide for GemStone/S 64 Bit* for complete information on NRS Syntax and the network environment.

StoneName identifies the name and network location of the database monitor process (Stone), which is the final arbiter of all sessions that access a specific database. Every session must communicate with a Stone, in both linked and remote applications. Hence, *StoneName* is a required argument.

A Stone process called "gs64stone" on node "lichen" could be described in a network resource string as:

```
!@lichen!gs64stone
```

A Stone of the same name that is running on the same machine as the application could be described in shortened form simply as:

```
gs64stone
```

GemService identifies the name and network location of the GemStone service that creates a session process (Gem), which then arbitrates data access between the database and the application. Every GemStone session requires a Gem. In linked applications, one Gem is present within the same process as the application; in remote applications the Gem is a separate process specific to that login session. Therefore, each time an application user logs in to GemStone (after the first time in linked applications), the GemStone service must create a new Gem. Hence, *GemService* is a required argument, except in the special case of a linked application that limits itself to one GemStone login per application process. In this special case, specify *GemService* as an empty string.

For most installations, the GemStone service name is *gemnetobject*. Specify, for example:

```
!@lichen!gemnetobject
```

HostUserId and *HostPassword* are your login name and password, respectively, on the machines that host the Gem and Stone processes. Do not confuse these values with your GemStone username and password - the GemStone username and password will be provided as arguments to **GciLogin**. *HostUserId* and *HostPassword* provide authentication for such tasks as creating a Gem and establishing communications with a Stone, and are optional in some configuration. When such authentication is required, an application user cannot login to GemStone until the host login is verified for the machine running the Stone or Gem, in addition to the GemStone login itself.

Authentication is always required if the NetLDI process that is related to the Stone is running in secure mode. In this case, it makes no difference whether the application is linked or remote. Authentication is also required to create a remote Gem, unless the NetLDI process is running in guest mode.

If the *HostUserId* argument is set to an empty C string or a NULL pointer, GemBuilder will try to find a username and password for authentication on a host machine in your network initialization file. To prevent GemBuilder from looking for authentication information in the network initialization file, supply a valid non-empty C string for the *HostUserId* argument, and a non-empty string for the *HostPassword* argument to provide a password. An empty string and a NULL pointer both mean that no password will be used for authentication.

Example

For an example of how **GciSetNet** is used, see the example for **GciLogin** on page 229.

See Also

GciLogin, page 229

GciLoginEx, page 231

GciSetNetEx, page 349

GciSetNetEx

GciSetNetEx_

Set network parameters for connecting the user to the Gem and Stone processes, allowing encryption.

Syntax

```
void GciSetNetEx(
    const char      StoneName[ ],
    const char      HostUserId[ ],
    const char      HostPassword[ ],
    const char      GemService[ ],
    BoolType        passwordIsEncrypted);

BoolType GciSetNetEx_(
    const char      StoneName[ ],
    const char      HostUserId[ ],
    const char      HostPassword[ ],
    const char      GemService[ ],
    BoolType        passwordIsEncrypted,
    char *          errString,
    size_t          maxErrSize);
```

Input Arguments

<i>StoneName</i>	Network resource string for the database monitor process.
<i>HostUserId</i>	UNIX host login name.
<i>HostPassword</i>	password of the UNIX user.
<i>GemService</i>	Network resource string for the GemStone service.
<i>passwordIsEncrypted</i>	Whether the password has been encrypted.

Result Arguments

<i>errString</i>	If there is a syntax error in the NRS for the <i>StoneName</i> or <i>GemService</i> , GciSetNetEx_ returns FALSE and sets <i>errString</i> to the details of the error. .
<i>maxErrSize</i>	Maximum size of <i>errString</i> .

Return Value

GciSetNetEx has no return value. **GciSetNetEx_** returns TRUE if the network parameters were set correctly, FALSE if there was a syntax error in *StoneName* or *GemService*.

Description

This function is similar to **GciSetNet**, but allows specifying additional behavior. For details on how to specify the *StoneName* and *GemService* using NRS, and the requirements for *HostUserId*, *HostPassword*, see **GciSetNet** on page 347.

GciSetNetEx allows you to specify that the host password you send is encrypted. The host password may be encrypted using **GciEncrypt**.

Example

For an example of how **GciSetNet** is used, see **GciLogin** on page 229.

See Also

GciEncrypt, page 147
GciLogin, page 229
GciLoginEx, page 231
GciSetNet, page 347

GciSetSessionId

Set an active session to be the current one.

Syntax

```
void GciSetSessionId(
    GciSessionIdType    sessionId);
```

Input Arguments

sessionId The session ID of an active (logged-in) GemStone session.

Description

This function can be used to switch between multiple GemStone sessions in an application program with multiple logins.

Example

```
void setSession_example(void)
{
    // assume topaz code for GciFetchVaryingOop has run
    // see GciLogin for login_example()
    if (! login_example())
        return;
    GciSessionIdType sess1 = GciGetSessionId();

    if (! login_example())
        return;
    GciSessionIdType sess2 = GciGetSessionId();

    { OopType aColl = GciResolveSymbol("AllComponents", OOP_NIL);
      OopType aComponent = GciExecuteStr(
          "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
      GciRemoveOopFromNsc(aColl, aComponent);
      GciReleaseOops(&aComponent, 1);
      printf("session %d , size after removal "FMT_I64"\n",
          sess2, GciFetchVaryingSize_(aColl));
    }
    // other session will still see the original size before removal
    // because it has an independent transactional view of the
    repository.
    GciSetSessionId(sess1);
    { OopType aColl = GciResolveSymbol("AllComponents", OOP_NIL);
      printf("session %d , current size "FMT_I64"\n",
          sess1, GciFetchVaryingSize_(aColl));
    }
}
```

See Also

[GciGetSessionId](#), page 209

[GciLogin](#), page 229

GciSetSharedCounter

Set the value of a shared counter.

Syntax

```
BoolType GciSetSharedCounter(  
    int counterIdx,  
    int64_t * value);
```

Input Arguments

<i>counterIdx</i>	The offset into the shared counters array of the value to modify.
<i>value</i>	Pointer to a value that containing the new value for this shared counter.

Return Value

Returns a C Boolean value indicating whether the value was successfully changed. Returns TRUE if the modification succeeded, FALSE if it failed.

Description

Set the value of the shared counter indicated by *counterIdx*. The contents of the *value* pointer indicate the new value of the shared counter.

Not supported for remote GCI interfaces.

See Also

[GciFetchNumSharedCounters](#), page 177
[GciDecSharedCounter](#), page 134
[GciIncSharedCounter](#), page 214
[GciReadSharedCounter](#), page 316
[GciReadSharedCounterNoLock](#), page 317
[GciFetchSharedCounterValuesNoLock](#), page 190

GciSetTraversalBufSwizzling

Control swizzling of the traversal buffers.

Syntax

```
BoolType GciSetTraversalBufSwizzling(  
    BoolType enabled );
```

Input Arguments

enabled If TRUE, enable normal byte-order swizzling of traversal buffers for the current RPC session. This is the default state for a session created by successful GciLogin().
If FALSE, the application program (for example, GemBuilder for Smalltalk) is responsible for subsequent swizzling of traversal buffers if needed.

Return Value

Returns the previous value of swizzling of traversal buffers. When called on a linkable session, returns FALSE and has no effect. If the current session is invalid, generates an error and returns FALSE.

Description

GciSetTraversalBufSwizzling controls swizzling of the traversal buffers used by these calls in an RPC session:

GciStoreTrav, GciNbStoreTrav
GciStoreTravDo_, GciNbStoreTravDo_
GciStoreTravDoTrav_, GciNbStoreTravDoTrav_
GciClampedTrav, GciNbClampedTrav
GciMoreTraversal, GciNbMoreTraversal
GciPerformTrav, GciNbPerformTrav
GciExecuteStrTrav, GciNbExecuteStrTrav

GciSetVaryingSize

Set the size of a collection.

Syntax

```
void GciSetVaryingSize(  
    OopType      collection,  
    int64        size );
```

Input Arguments

<i>collection</i>	The OOP of the collection whose size you are specifying.
<i>size</i>	The desired number of elements in the collection.

Description

GciSetVaryingSize changes the size of a collection, adding nils to grow it, or truncating it, as necessary. It is equivalent to the Smalltalk method `Object >> size:`. It does not change the number of any named instance variables.

Example

```
void setVaryingSize_example(void)  
{  
    OopType oArr = GciNewOop(OOP_CLASS_ARRAY); // create new Array of size 0  
  
    GciSetVaryingSize(oArr, 1000000);  
    // logical size now 1 million  
  
    GciStoreOop(oArr, 500000, GciI32ToOop(5678));  
}
```

See Also

[GciFetchVaryingSize_](#), page 199

GciShutdown

Logout from all sessions and deactivate GemBuilder.

Syntax

```
void GciShutdown()
```

Description

This function is intended to be called by image exit routines, such as the **on_exit** system call. In the linkable GemBuilder, **GciShutdown** calls **GciLogout**. In the RPC version, it logs out all sessions connected to the Gem process and shuts down the networking layer, thus releasing all memory allocated by GemBuilder.

It is especially important to call this function explicitly on any computer whose operating system does not automatically deallocate resources when a process quits. This effect is found on certain small, single-user systems.

GciSoftBreak

Interrupt the execution of Smalltalk code, but permit it to be restarted.

Syntax

```
void GciSoftBreak()
```

Description

This function sends a soft break to the current user session (set by the last **GciLogin** or **GciSetSessionId**).

GemBuilder allows users of your application to terminate Smalltalk execution. This is useful, for example, if the a Smalltalk method is invoked that enters an infinite loop.

GciSoftBreak interrupts only the Smalltalk virtual machine (if it is running), and does so in such a way that the it can be restarted. The only GemBuilder functions that can recognize a soft break include **GciSendMessage**, **GciPerform**, and **GciContinue**, and the **GciExecute...** functions.

GciHardBreak has no effect if called from within a User Action.

In order for GemBuilder functions in your program to recognize interrupts, your program will need a signal handler that can call the functions **GciSoftBreak** and **GciHardBreak**. Since GemBuilder does not relinquish control to an application until it has finished its processing, soft and hard breaks must be initiated from another thread.

If GemStone is executing when it receives the break, it replies with the error message `RT_ERR_SOFT_BREAK`. Otherwise, it ignores the break.

Example

```
#include "signal.h"

extern "C" {
    static void doSoftBreak(int sigNum, siginfo_t* info, void* ucArg)
    {
        GciSoftBreak();
    }
}

void softBreakExample(void)
{
    // save previous SIGINT handler and install ours
    struct sigaction oldHandler;
    struct sigaction newHandler;
    newHandler.sa_handler = SIG_DFL;
    newHandler.sa_sigaction = doSoftBreak;
    newHandler.sa_flags = SA_SIGINFO | SA_RESTART ;
    sigaction(SIGINT, &newHandler, &oldHandler);
}
```

```
// execute a loop that will take 120 seconds to execute and
// return the SmallInteger with value 11 .
OopType result = GciExecuteStr(
    "| a | a := 1 . 10 timesRepeat:[ System sleep:10. a := a + 1]. ^ a",
    OOP_NIL/*use default symbolList for execution*/);

BoolType done = FALSE;
int breakCount = 0;
do {
    // assume the user may type ctl-C or issue kill -INT from
    // another shell process during the 120 seconds .
    GciErrSType errInfo;
    if ( GciErr(&errInfo)) {
        if (errInfo.number == RT_ERR_SOFT_BREAK) {
            // GciExecuteStr was interrupted by a GciSoftBreak .
            breakCount++ ;
            // now continue the execution to finish the computation
            result = GciContinue(errInfo.context);
        } else {
            // FMT_OID format string is defined in gci.ht
            printf("unexpected error category \"FMT_OID\" number %d, %s\n",
                errInfo.category, errInfo.number, errInfo.message);
            // terminate the execution
            GciClearStack(errInfo.context);
            done = TRUE;
        }
    } else {
        // GciExecuteStr or GciContinue completed without error
        done = TRUE;
        BoolType conversionErr = FALSE;
        int val = GciOopToI32_(result, &conversionErr);
        if (conversionErr) {
            printf("Error converting result to C int\n");
        } else {
            printf("Got %d interrupts, result = %d\n", breakCount, val);
        }
    }
} while (! done);

// restore previous SIGINT handler
sigaction(SIGINT, &oldHandler, NULL);
}
```

See Also

[GciClearStack](#), page 113

[GciContinue](#), page 120

[GciExecute](#), page 149

[GciHardBreak](#), page 210

[GciPerform](#), page 298

GciStep

GciStep_

Continue code execution in GemStone with specified single-step semantics.

Syntax

```

OopType GciStep(
    OopType          process,
    int              level );

OopType GciStep_(
    OopType          process,
    int              level
    BoolType        through);

```

Input Arguments

<i>process</i>	The OOP of a GsProcess object (obtained as the value of the context field of an error report returned by GciErr).
<i>level</i>	One of the following values: 0 – step-into semantics starting from top of stack 1 – step-over semantics starting from top of stack > 1 – step-over semantics from specified level on stack
<i>through</i>	When <i>level</i> =1 and this argument is TRUE, provides step through semantics, stopping in blocks for which the top of the stack frame is the home method.

Return Value

Returns the OOP of the result of the Smalltalk execution. Returns OOP_ILLEGAL in case of error.

Description

This function continues code execution in GemStone using the specified single-step semantics. This function is intended for use by debuggers.

If you specify a *level* that is either less than zero or greater than the value represented by `GciPerform(process, "stackDepth", NULL, 0)`, this function generates an error.

See Also

`GciPerform`, page 298

GciStoreByte

Store one byte in a byte object.

Syntax

```
void GciStoreByte(  
    OopType          theObject,  
    int64            atIndex,  
    ByteType         theByte );
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone byte object.
<i>atIndex</i>	The index into theObject at which to store the byte.
<i>theByte</i>	The 8-bit value to be stored.

Result Arguments

<i>theObject</i>	The resulting GemStone byte object.
------------------	-------------------------------------

Description

This function stores a single element in a byte object at a specified index, using structural access.

GciStoreByte raises an error if *theObject* is a Float or SmallFloat. You must store all the bytes of a Float or SmallFloat if you store any.

Example

```
void storeByte_example(void)  
{  
    OopType oString = GciNewOop(OOP_CLASS_STRING);  
  
    for (int j = 0; j < 200; j++) {  
        ByteType val = j;  
        GciStoreByte(oString, j + 1 , val );  
    }  
}
```

See Also

[GciFetchByte](#), page 161
[GciFetchBytes_](#), page 162
[GciStoreBytes](#), page 361

GciStoreBytes

(MACRO) Store multiple bytes in a byte object.

Syntax

```
void GciStoreBytes(
    OopType          theObject,
    int64            startIndex,
    const ByteType   theBytes[],
    int64            numBytes);
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone byte object.
<i>startIndex</i>	The index into theObject at which to begin storing bytes.
<i>theBytes</i>	The array of bytes to be stored.
<i>numBytes</i>	The number of elements to store.

Result Arguments

<i>theObject</i>	The resulting GemStone byte object.
------------------	-------------------------------------

Description

This macro uses structural access to store multiple elements from a C array in a byte object, beginning at a specified index. A common application of **GciStoreBytes** would be to store a text string. For an object with multiple bytes per character or digit, *theBytes* is expected to be in client native byte order, and will be swizzled if needed by the server.

Error Conditions

GciStoreBytes raises an error if *theObject* is a Float or SmallFloat. Use **GciStoreBytesInstanceOf** instead for Float or SmallFloat objects.

Example

```
void storeBytes_example(void)
{
    OopType oString = GciNewOop(OOP_CLASS_STRING);

    enum { buf_size = 2000 };
    ByteType buf[buf_size];
    for (int j = 0; j < buf_size; j++) {
        buf[j] = (ByteType)j;
    }
    GciStoreBytes(oString, 1, buf, buf_size);
}
```

See Also

`GciFetchByte`, page 161
`GciFetchBytes_`, page 162
`GciStoreByte`, page 360
`GciStoreBytesInstanceOf`, page 363
`GciStoreChars`, page 365

GciStoreBytesInstanceOf

Store multiple bytes in a byte object.

Syntax

```
void GciStoreBytesInstanceOf(
    OopType          theClass,
    OopType          theObject,
    int64            startIndex,
    const ByteType   theBytes[],
    int64            numBytes);
```

Input Arguments

<i>theClass</i>	The OOP of the class of the GemStone byte object.
<i>theObject</i>	The OOP of the GemStone byte object.
<i>startIndex</i>	The index into theObject at which to begin storing bytes.
<i>theBytes</i>	The array of bytes to be stored.
<i>numBytes</i>	The number of elements to store.

Result Arguments

<i>theObject</i>	The resulting GemStone byte object.
------------------	-------------------------------------

Description

This function uses structural access to store multiple elements from a C array into a byte object, beginning at a specified index. A common application of **GciStoreBytesInstanceOf** would be to store a Float or LargeInteger object.

GciStoreBytesInstanceOf provides automatic byte swizzling for objects such as Float, LargeInteger, and DoubleByteString that use multiple bytes per digit or character. For more about byte swizzling, see “Byte-Swizzling of Binary Floating-Point Values” on page 25. For these objects, *theBytes* is assumed to be in client native byte order. For DoubleByteStrings, *startIndex* must be aligned on character boundaries and *numbytes* must be a multiple of the number of bytes per character; for numeric objects *startIndex* must be one and the *numBytes* the size of the numeric class.

The presence of the argument *theClass* enables the swizzling to be implemented more efficiently. If *theObject* is a Float or SmallFloat, then *theClass* must match the actual class of *theObject*, *startIndex* must be one, and *numBytes* must be the actual size for *theClass*. If any of these conditions are not met, then **GciStoreBytesInstanceOf** raises an error as a safety check.

If *theObject* is not a Float or SmallFloat, then *theClass* is ignored. Hence, you must supply the correct class for *theClass* if *theObject* is a Float or SmallFloat, but you can use OOP_NIL otherwise.

Example

```
void storeBytesInstof_example(void)
{
    double pi = 3.1415926;
```

```
OopType oFloat = GciNewOop(OOP_CLASS_FLOAT);
GciStoreBytesInstanceOf(OOP_CLASS_FLOAT, oFloat, 1,
    (ByteType *)&pi, sizeof(pi));
}
```

See Also

[GciFetchByte](#), page 161
[GciFetchBytes_](#), page 162
[GciStoreByte](#), page 360
[GciStoreBytes](#), page 361
[GciStoreChars](#), page 365

GciStoreChars

Store multiple ASCII characters in a byte object.

Syntax

```
void GciStoreChars(  
    OopType          theObject,  
    int64            startIndex,  
    const char *     aString );
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone byte object.
<i>startIndex</i>	The index into theObject at which to begin storing the string.
<i>aString</i>	The string to be stored.

Result Arguments

<i>theObject</i>	The resulting GemStone byte object.
------------------	-------------------------------------

Description

This function uses structural access to store a C string in a byte object, beginning at a specified index.

GciStoreChars raises an error if *theObject* is a Float or SmallFloat. ASCII characters have no meaning as bytes in a Float or SmallFloat object.

Example

```
void storeChars_example(void)  
{  
    OopType oString = GciNewOop(OOP_CLASS_STRING);  
  
    GciStoreChars(oString, 1, "some string data");  
}
```

See Also

GciFetchByte, page 161
GciFetchBytes_, page 162
GciStoreByte, page 360
GciStoreBytes, page 361

GciStoreDynamicIv

Create or change the value of an object's dynamic instance variable.

Syntax

```
void GciStoreDynamicIv(  
    OopType          theObject,  
    OopType          aSymbol,  
    OopType          value);
```

Input Arguments

<i>theObject</i>	The OOP of the GemStone object.
<i>aSymbol</i>	Specifies the dynamic instance variable of the object.
<i>value</i>	The value to store in the dynamic instance variable.

Return Value

Creates or changes the value of the dynamic instance variable specified by *aSymbol* within *theObject*.

Description

This function stores a value into the dynamic instance variable specified by *aSymbol*.

Dynamic instance variables are not allowed in instances of ExecBlock, Behavior, GsNMethod, or special objects.

To delete a dynamic instance variable, pass OOP_REMOTE_NIL as the value.

See Also

[GciFetchDynamicIv](#), page 168

[GciFetchDynamicIvs](#), page 169

GciStoreIdxOop

Store one OOP in an indexable pointer object's unnamed instance variable.

Syntax

```
void GciStoreIdxOop(
    OopType          theObject,
    int64            atIndex,
    OopType          theOop);
```

Input Arguments

<i>theObject</i>	The pointer object.
<i>atIndex</i>	The index into <i>theObject</i> at which to store the object.
<i>theOop</i>	The OOP to be stored.

Result Arguments

<i>theObject</i>	The resulting pointer object.
------------------	-------------------------------

Description

This function stores a single OOP into an indexed variable of a pointer object at the specified index, using structural access. Note that this function cannot be used for NSCs. (To add an OOP to an NSC, use **GciAddOopToNsc** on page 92.)

Example

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void storeIdxOop_example(void)
{
    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);

    OopType otherComp = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1333]", OOP_NIL);

    // store new value into 3rd element of aComponent's parts list
    GciStoreIdxOop(aComponent, 3, otherComp);

    GciReleaseOops(&aComponent, 1); // release results of execution
    GciReleaseOops(&otherComp, 1);
}
```

See Also

[GciAddOopToNsc](#), page 92

[GciFetchVaryingOop](#), page 195

[GciFetchVaryingOops](#), page 197

[GciStoreIdxOops](#), page 369

GciStoreIdxOops

Store one or more OOPs in an indexable pointer object's unnamed instance variables.

Syntax

```
void GciStoreIdxOops(
    OopType          theObject,
    int64            startIndex,
    const OopType    theOops[],
    int              numOops);
```

Input Arguments

<i>theObject</i>	The pointer object.
<i>startIndex</i>	The index into <i>theObject</i> at which to begin storing OOPs.
<i>theOops</i>	The array of OOPs to be stored.
<i>numOops</i>	The number of OOPs to store.

Result Arguments

<i>theObject</i>	The resulting pointer object.
------------------	-------------------------------

Description

This function uses structural access to store multiple OOPs from a C array into the indexed variables of a pointer object, beginning at the specified index. Note that this call cannot be used with NSCs. (To add multiple OOPs to an NSC, use **GciAddOopsToNsc** on page 93.)

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void storeIdxOops_example(void)
{
    // retrieve a random instance of class Component
    OopType firstC = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);

    OopType secondC = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1333]", OOP_NIL);

    // make first component's parts list be identical to second
    component's list
    enum { buf_size = 100 };
    OopType buf[buf_size];
    int64 firstSize = GciFetchVaryingSize_(firstC);
    int64 idx = 1;
    while (idx <= firstSize) {
```

```
    int numRet = GciFetchVaryingOops(firstC, idx, buf, buf_size);
    GciStoreIdxOops(secondC, idx, buf, numRet);
    idx += numRet;
}
// truncate second component's parts list if it was larger than
first's
GciSetVaryingSize(secondC, firstSize);

GciReleaseOops(&firstC, 1); // release results of executions
GciReleaseOops(&secondC, 1);
}
```

See Also

GciAddOopsToNsc, page 93
GciFetchVaryingOops, page 195
GciFetchVaryingOops, page 197
GciReplaceOops, page 329
GciReplaceVaryingOops, page 330
GciStoreIdxOops, page 367
GciStoreIdxOops, page 369
GciStoreNamedOops, page 373
GciStoreOops, page 377

GciStoreNamedOop

Store one OOP into an object's named instance variable.

Syntax

```
void GciStoreNamedOop(
    OopType          theObject,
    int64            atIndex,
    OopType          theOop);
```

Input Arguments

<i>theObject</i>	The object in which to store the OOP.
<i>atIndex</i>	The index into <i>theObject</i> 's named instance variables at which to store the OOP.
<i>theOop</i>	The OOP to be stored.

Result Arguments

<i>theObject</i>	The resulting object with the new OOP.
------------------	--

Description

This function stores a single OOP into an object's named instance variable at the specified index, using structural access.

Example

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for `GciFetchVaryingOop` on page 195.

```
void storeNamedOop_example(void)
{
    // C constants to match Smalltalk class definition
    enum { COMPONENT_OFF_PARTNUMBER = 1,
          COMPONENT_OFF_NAME       = 2,
          COMPONENT_OFF_COST       = 3 };

    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // error during execution or detect found nothing
        return;
    }

    // assign a new value to the name instance variable of aComponent
    OopType newName = GciNewString("compressor blade");
    GciStoreNamedOop(aComponent, COMPONENT_OFF_NAME, newName);
}
```

```
// alternate approach: assign a new value to a named instance
//variable without knowing its offset at compile time
GciStoreNamedOop(aComponent,
GciIvNameToIdx(GciFetchClass(aComponent), "name"), newName);

GciReleaseOops(&newName, 1);
GciReleaseOops(&aComponent, 1);
}
```

See Also

GciFetchNamedOop, page 170
GciFetchNamedOops, page 172
GciStoreIdxOop, page 367
GciStoreNamedOops, page 373

GciStoreNamedOops

Store one or more OOPs into an object's named instance variables.

Syntax

```
void GciStoreNamedOops(
    OopType          theObject,
    int64            startIndex,
    const OopType    theOops[],
    int              numOops);
```

Input Arguments

<i>theObject</i>	The object in which to store the OOPs.
<i>startIndex</i>	The index into <i>theObject</i> 's named instance variables at which to begin storing OOPs.
<i>theOops</i>	The array of OOPs to be stored.
<i>numOops</i>	The number of OOPs to store. If (numOops+startIndex) exceeds the number of named instance variables in <i>theObject</i> , an error is generated.

Result Arguments

<i>theObject</i>	The resulting object with the new OOPs.
------------------	---

Description

This function uses structural access to store multiple OOPs from a C array into an object's named instance variables, beginning at the specified index.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void storeNamedOops_example(void)
{
    // retrieve a random instance of class Component
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        // execution error, or detect: found nothing
        return;
    }

    // fetch name instance variables without knowing offset at C
    // compile time
    int namedSize = GciFetchNamedSize(aComponent);
    if (namedSize == 0) {
        // error during fetch
    }
}
```

```
        return;
    }
    OopType *oBuffer = (OopType*) malloc( sizeof(OopType) * namedSize
);
    if (oBuffer == NULL) {
        printf("malloc failure\n");
        return;
    }
    int numRet = GciFetchNamedOops(aComponent, 1, oBuffer, namedSize);
    if (numRet != namedSize) {
        printf("error during fetch\n");
        return;
    }

    // alter one of the instVars and then store them all
    OopType newName = GciNewString("compressor blade");
    int ivOffset = GciIvNameToIdx(GciFetchClass(aComponent), "name");
    if (ivOffset <= 0) {
        printf("error during GciIvNameToIdx\n");
        return;
    }
    oBuffer[ivOffset - 1 ] = newName;
    GciStoreNamedOops(aComponent, 1, oBuffer, namedSize);

    GciReleaseOops(&newName, 1);
    GciReleaseOops(&aComponent, 1);
}
```

See Also

- [GciFetchNamedOop](#), page 170
- [GciFetchNamedOops](#), page 172
- [GciReplaceOops](#), page 329
- [GciReplaceVaryingOops](#), page 330
- [GciStoreIdxOop](#), page 367
- [GciStoreIdxOops](#), page 369
- [GciStoreNamedOop](#), page 371
- [GciStoreNamedOops](#), page 373
- [GciStoreOops](#), page 377

GciStoreOop

Store one OOP into an object's instance variable.

Syntax

```
void GciStoreOop(
    OopType          theObject,
    int64            atIndex,
    OopType          theOop);
```

Input Arguments

<i>theObject</i>	The object in which to store the OOP.
<i>atIndex</i>	The index into <i>theObject</i> at which to store the OOP. This function does not distinguish between named and unnamed instance variables. Indices are based at the beginning of an object's array of instance variables. In that array, the object's named instance variables are followed by its unnamed instance variables.
<i>theOop</i>	The OOP to be stored.

Result Arguments

<i>theObject</i>	The resulting object.
------------------	-----------------------

Description

This function stores a single OOP into an object at the specified index, using structural access. Note that this function cannot be used for NSCs. To add an object to an NSC, use **GciAddOopToNsc** on page 92.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void storeOop_example(void)
{
    /* C constants to match Smalltalk class definition */
    enum { COMPONENT_OFF_NAME = 2 };

    /* retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);

    OopType newName = GciNewString("vane");

    /* Two ways to assign new value to name instance variable of
    aComponent */
    GciStoreOop(aComponent, COMPONENT_OFF_NAME, newName);
```

```
GciStoreNamedOop(aComponent, COMPONENT_OFF_NAME, newName);

OopType subPart = GciExecuteStr(
"AllComponents detect:[i|i partNumber = 1333]", OOP_NIL);

/* Two ways to assign a new value to the 3rd element of
aComponent's parts list without knowing exactly how many named
instance variables exist */

GciStoreOop(aComponent, GciFetchNamedSize(aComponent) + 3,
subPart);
GciStoreIdxOop(aComponent, 3, subPart);
}
```

See Also

GciAddOopToNsc, page 92
GciFetchVaryingOop, page 195
GciFetchVaryingOops, page 197
GciFetchOops, page 183
GciStoreOops, page 377

GciStoreOops

Store one or more OOPs into an object's instance variables.

Syntax

```
void GciStoreOops(
    OopType          theObject,
    int64            startIndex,
    const OopType    theOops[],
    int              numOops);
```

Input Arguments

<i>theObject</i>	The object in which to store the OOPs.
<i>startIndex</i>	The index into <i>theObject</i> at which to begin storing OOPs. This function does not distinguish between named and unnamed instance variables. Indices are based at the beginning of an object's array of instance variables. In that array, the object's named instance variables are followed by its unnamed instance variables.
<i>theOops</i>	The array of OOPs to be stored.
<i>numOops</i>	The number of OOPs to store.

Result Arguments

<i>theObject</i>	The resulting object.
------------------	-----------------------

Description

This function uses structural access to store multiple OOPs from a C array into a pointer object, beginning at the specified index. Note that this call cannot be used with NSCs. To add multiple OOPs to an NSC, use **GciAddOopsToNsc** on page 93.

Example

In the following example, assume that you've defined the class Component and populated the set AllComponents, as shown in the example for **GciFetchVaryingOop** on page 195.

```
void storeOops_example(void)
{
    /* retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);

    int namedSize = GciFetchNamedSize(aComponent);
    int64 instSize = GciFetchSize_(aComponent);
    // allow space in buffer for storing into first varying instVar
    plus
    // appending a new varying instVar
    int64 bufVaryingSize = instSize - namedSize + 1;
```

```
if (bufVaryingSize < 2)
    bufVaryingSize = 2;

int64 bufSize = namedSize + bufVaryingSize;
OopsType *buf = (OopsType*) malloc(sizeof(OopsType) * bufSize);
if (buf == NULL) {
    printf("malloc failure");
    return;
}
GciFetchOops(aComponent, 1, buf, instSize);

OopsType newName = GciNewString("strut");
int nameOfs = GciIvNameToIdx(GciFetchClass(aComponent), "name");
buf[nameOfs - 1] = newName;

OopsType firstSubPart = GciExecuteStr(
"AllComponents detect:[i|i partNumber = 1333]", OOP_NIL);

OopsType lastSubPart = GciExecuteStr(
"AllComponents detect:[i|i partNumber = 1555]", OOP_NIL);

// assign first element of parts list
buf[namedSize] = firstSubPart;

// append lastSubPart to aComponent's parts list
int64 newSize = instSize + 1;
buf[newSize - 1] = lastSubPart;

// now store all the instVars back to the repository
GciStoreOops(aComponent, 1, buf, newSize);
}
```

See Also

- GciAddOopsToNsc**, page 93
- GciFetchNamedOops**, page 172
- GciFetchOops**, page 181
- GciFetchOops**, page 183
- GciFetchVaryingOops**, page 195
- GciReplaceOops**, page 329
- GciReplaceVaryingOops**, page 330
- GciStoreIdxOops**, page 369
- GciStoreNamedOops**, page 373
- GciStoreOops**, page 375
- GciStoreOops**, page 377

GciStorePaths

Store selected multiple OOPs into an object tree.

Syntax

```
BoolType GciStorePaths(
    const OopType    theOops [ ],
    int              numOops,
    const int        paths [ ],
    const int        pathSizes [ ],
    int              numPaths,
    const OopType    newValues [ ],
    int *            failCount );
```

Input Arguments

<i>theOops</i>	A collection of OOPs into which you want to store new values.
<i>numOops</i>	The size of <i>theOops</i> .
<i>paths</i>	An array of integers. This one-dimensional array contains the elements of all constituent paths, laid end to end.
<i>pathSizes</i>	An array of integers. Each element of this array is the length of the corresponding path in the <i>paths</i> array (that is, the number of elements in each constituent path).
<i>numPaths</i>	The number of paths in the <i>paths</i> array. This should be the same as the number of integers in the <i>pathSizes</i> array.
<i>newValues</i>	An array containing the new values to be stored into <i>theOops</i> .

Result Arguments

<i>failCount</i>	A pointer to an integer that indicates which element of the <i>newValues</i> array could not be successfully stored. If all values were successfully stored, <i>failCount</i> is 0. If the <i>i</i> th store failed, <i>failCount</i> is <i>i</i> . If any of the objects in <i>newValues</i> does not exist, or is not an OOP allocated to GemBuilder, <i>failCount</i> is 1.
------------------	--

Return Value

Returns TRUE if all values were successfully stored. Returns FALSE if the store on any path fails for any reason.

Description

This function allows you to store multiple objects at selected positions in an object tree with a single GemBuilder call, exporting only the desired information to the database.

NOTE

This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the

*"linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

Each path in the *paths* array is itself an array of longs. Those longs are offsets that specify a path along which to store objects. In each path, a positive integer *x* refers to an offset within an object's named instance variables, while a negative integer *-x* refers to an offset within an object's indexed instance variables.

The *newValues* array contains (*numOps* * *numPaths*) elements, stored in the following order:

```
[ 0, 0 ] .. [ 0, numPaths-1 ] .. [ 1, 0 ] .. [ 1, numPaths-1 ] ..
[ numOps-1, 0 ] .. [ numOps-1, numPaths-1 ]
```

The first element of this *newValues* array is stored along the first path into the first element of *theOps*. New values are then stored into the first element of *theOps* along each remaining element of the *paths* array. Similarly, new values are stored into each subsequent element of *theOps*, until all paths have been applied to all its elements.

The new value to be stored into object *i* along path *j* is thus represented as:

```
newValues[ ((i-1) * numPaths) + (j-1) ]
```

The expressions *i-1* and *j-1* are used because C has zero-based arrays.

If the store on any path fails for any reason, this function stops and generates a GemBuilder error. Any objects that were successfully stored before the error occurred will remain stored.

Example 1: Calling sequence for a single object and a single path

```
void storePath1(void)
{
    enum { path_size = 5 };
    int    aPath[path_size]; /* the path itself */
    int    aSize = path_size; /* the size of the path */

    OopType anOop; // the OOP to use as the root of the path
    anOop = GciExecuteStr("AllComponents detect[:i|i partNumber =
1234]", OOP_NIL);
    if (anOop == OOP_NIL) {
        return; // error during resolve
    }

    OopType newValue = GciNewString("a new value");
    int    failCount;

    GciStorePaths(&anOop, 1, aPath, &aSize, 1, &newValue, &failCount);
}
```

Example 2: Calling sequence for multiple objects with a single path

```
void storePath2(void)
{
```

```

OopType coll = GciResolveSymbol("AllComponents", OOP_NIL);
if (coll == OOP_NIL) {
    return ; // error during resolve
}
enum { num_roots = 3 ,
       path_size = 5 };
OopType oops[num_roots];
int numRet = GciFetchVaryingOops(coll, 1, oops, num_roots);
if (numRet != num_roots) {
    return; // error during fetch or collection too small
}

int  aPath[path_size];
int  aSize = path_size;
for (int j = 0; j < path_size; j++) {
    aPath[j] = 1; // arbitrary offsets
}

OopType newValues[num_roots];
for (int j = 0; j < num_roots; j++) {
    newValues[j] = GciI32ToOop(1345600 + j);
}
int    failCount;
GciStorePaths(oops, num_roots, aPath, &aSize, 1, newValues,
&failCount);
}

```

Example 3: Calling sequence for a single object with multiple paths

```

void storePath3(void)
{
    OopType anOop; // the OOP to use as the root of the path
    anOop = GciExecuteStr("AllComponents detect[:i|i partNumber =
1234]", OOP_NIL);
    if (anOop == OOP_NIL) {
        return; // error during execution
    }

    enum { num_paths = 10,
          path_size = 5 };

    int pathSizes[num_paths];
    int paths[path_size * num_paths ];
    int idx = 0;
    for (int j = 0; j < num_paths; j++) {
        for (int k = 0; k < path_size; k++) {
            paths[idx++] = k + 1; // arbitrary offset
        }
    }
    OopType newValues[num_paths];
}

```

```
    for (int j = 0; j < num_paths; j++) {
        newValues[j] = GciI32ToOop(1345600 + j);
    }
    int    failCount;

    GciStorePaths(&anOop, 1, paths, pathSizes, num_paths, newValues,
&failCount);
}
```

Example 4: Calling sequence for multiple objects with multiple paths

```
void storePaths4(void)
{
    OopType coll = GciResolveSymbol("AllComponents", OOP_NIL);
    if (coll == OOP_NIL) {
        return ; // error during resolve
    }

    enum { num_roots = 10,
          num_paths = 3,
          path_size = 5 ,
          num_new_values = num_roots * num_paths
        };
    OopType oops[num_roots];
    int numRet = GciFetchVaryingOops(coll, 1, oops, num_roots);
    if (numRet != num_roots) {
        return; // error during fetch or collection too small
    }

    int pathSizes[num_paths];
    int paths[path_size * num_paths ];
    int idx = 0;
    for (int j = 0; j < num_paths; j++) {
        for (int k = 0; k < path_size; k++) {
            paths[idx++] = k + 1; // arbitrary offset
        }
    }

    OopType newValues[num_new_values];
    for (int j = 0; j < num_new_values; j++) {
        newValues[j] = GciI32ToOop(1345600 + j);
    }
    int    failCount;
    GciStorePaths(oops, num_roots, paths, pathSizes, num_paths,
newValues,
&failCount);
}
```

Example 5: Integrated Code

In the following example, assume that you've defined the class `Component` and populated the set `AllComponents`, as shown in the example for `GciFetchVaryingOop` on page 195.

```
void storePaths5(void)
{
    // retrieve a random instance of class Component */
    OopType aComponent = GciExecuteStr(
        "AllComponents detect:[i|i partNumber = 1234]", OOP_NIL);
    if (aComponent == OOP_NIL) {
        return; // error in execute, or detect: found nothing
    }

    // assign a new value to the name instVar of 5th element of
    // aComponent's parts list
    enum { path_size = 2 };
    int path[path_size];
    path[0] = -5; // 5th varying instVar
    path[1] = GciIvNameToIdx(GciFetchClass(aComponent), "name");
    int pathSizes = path_size;

    OopType newValue = GciNewString("pump");
    int failCount;
    GciStorePaths(&aComponent, 1, path, &pathSizes, 1, &newValue,
        &failCount);
}
```

See Also

`GciFetchPaths`, page 185

GciStoreTrav

Store multiple traversal buffer values in objects.

Syntax

```
void GciStoreTrav(
    GciTravBufType *   travBuff,
    int                behaviorFlag );
```

Input Arguments

<i>travBuff</i>	A traversal buffer, which contains object data to be stored.
<i>behaviorFlag</i>	A flag that determines how the objects should be handled.

Description

This function stores data from the traversal buffer *travBuff* (a C-language structural description) into multiple GemStone objects. The first element in the traversal buffer is an integer that indicates how many bytes are stored in the buffer. The remainder of the traversal buffer consists of a series of object reports. Each object report is a C structure of type **GciObjRepSType**, which includes a variable-length data area. **GciStoreTrav** stores data object by object, using one object report at a time. **GciStoreTrav** raises an error if the traversal buffer contains a report for any object of special implementation format.

GciStoreTrav allows you to reduce the number of GemBuilder calls that are required for your application program to store complex objects in the database.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

The value of *behaviorFlag* should be given by using one or more of the following GemBuilder mnemonics: `GCI_STORE_TRAV_DEFAULT`, `GCI_STORE_TRAV_NSC_REP`, `GCI_STORE_TRAV_CREATE`, and `GCI_STORE_TRAV_FINISH_UPDATES`. The first of these must be used alone. The others can either be used alone or can be logically "or"ed together. The effect of the mnemonics depends somewhat upon the implementation format of the objects that are stored.

GciStoreTrav can create new objects and store data into them, or it can modify existing objects with the data in their object reports, or a combination of the two. By default (`GCI_STORE_TRAV_DEFAULT`), it can only modify existing objects, and it raises an error if an object does not already exist.

When `GCI_STORE_TRAV_CREATE` is used, it modifies any object that already exists and creates a new object when an object does not exist. Naturally, any new object is initialized with the data in its object report.

When `GCI_STORE_TRAV_FINISH_UPDATES` is used, **GciStoreTrav** automatically executes **GciProcessDeferredUpdates_** after processing the last object report in the traversal buffer.

When **GciStoreTrav** modifies an existing object of byte or pointer format, it replaces that object's data with the data in its object report, regardless of *behaviorFlag*. All instance variables, named (if any) or indexed (if any), receive new values. Named instance variables for which values are not given in the object report are initialized to nil or to zero. Indexable objects may change in size; the object report determines the new number of indexed variables.

Contrast byte and pointer object handling with the default when **GciStoreTrav** modifies an existing NSC. It replaces all named instance variables of the NSC (if any), but adds further data in its object report to the unordered variables, increasing its size. If *behaviorFlag* indicates `GCI_STORE_TRAV_NSC_REP`, then it removes all existing unordered variables and adds new unordered variables with values from the object report.

GciStoreTrav provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see "Byte-Swizzling of Binary Floating-Point Values" on page 25.

Use of Object Reports

GciStoreTrav stores values in GemStone objects according to the object reports contained in *travBuff*. Each object report is an instance of the C++ class **GciObjRepSType** (described in "The Object Report Structure" on page 74). **GciStoreTrav** uses the fields in each object report as follows:

rpt->hdr.valueBuffSize

The size (in bytes) of the value buffer, where object data is stored. If *objId* is a Float or SmallFloat and *valueBuffSize* differs from the actual size for objects of *objId*'s class, then **GciStoreTrav** raises an error.

rpt->hdr.namedSize

Ignored by this function.

rpt->hdr.setIdxSize()

Only needs to be called if the object is indexable. The number of indexed variables in the object stored by **GciStoreTrav** is never less than this quantity. It may be more if the value buffer contains enough data. **GciStoreTrav** stores all the indexed variables that it finds in the value buffer. If an existing object has more indexed variables, then it also retains the extras, up to a total of *idxSize*, and removes any beyond *idxSize*. If *idxSize* is larger than the number of indexed variables in both the current object and the value buffer, then **GciStoreTrav** creates slots for elements in the stored object up to index *idxSize* and initializes any added elements to nil.

rpt->hdr.firstOffset

Ignored for NSC objects. The absolute offset into the target object at which to begin storing values from the value buffer. The absolute offset of the object's first named instance variable (if any) is one; the offset of its first indexed variable (if any) is one more than the number of its named instance variables. Values are stored into the object in the order that they appear in the value buffer, ignoring the boundary between named and indexed variables. Variables whose offset is less than *firstOffset* (if any) are initialized to nil or zero. For nonindexable objects, **GciStoreTrav** raises an error if *valueBuffSize* and *firstOffset* imply a size that exceeds the actual size of the object. If *objId* is a Float or SmallFloat and *firstOffset* is not one, then **GciStoreTrav** raises an error.

rpt->hdr.objId

The OOP of the object to be stored.

rpt->hdr.oclass

Used only when creating a new object, to identify its intended class.

rpt->hdr.objectSecurityPolicyId

The ID of the object's security policy.

rpt->hdr.clearBits()

Must be called before any of the following:

rpt->hdr.setObjImpl()

You must call *rpt->hdr.setObjImpl* to set this field to be consistent with the object's implementation.

rpt->hdr.setInvariant()

Boolean value. Call *rpt->hdr.setInvariant(TRUE)* if you want this object to be made invariant after the store specified by *rpt** is completed.

rpt->hdr.setIndexable()

Ignored by this function.

rpt->valueBufferBytes()

The value buffer of an object of byte format.

rpt->valueBufferOops()

The value buffer of an object of pointer or NSC format.

Handling Error Conditions

If you get a runtime error while executing **GciStoreTrav**, the recommended course of action is to abort the current transaction.

See Also

GciMoreTraversal, page 234

GciNbMoreTraversal, page 253

GciNbStoreTrav, page 260

GciNbTraverseObjs, page 265

GciNewOopUsingObjRep, page 272

GciProcessDeferredUpdates_, page 312

GciStoreTravDo_, page 387

GciTraverseObjs, page 408

GciStoreTravDo_

Store multiple traversal buffer values in objects, execute the specified code, and return the resulting object.

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciStoreTravDo** (without the underscore).

Syntax

```

OopType GciStoreTravDo_(
    GciStoreTravDoArgsSType *args );

```

Input Arguments

<i>args</i>	An instance of <code>GciStoreTravDoArgsSType</code> (as described in <code>\$GEMSTONE/include/gcicmn.ht</code>) containing the following fields:
<code>GciTravBufType*</code> <i>storeTravBuff</i>	The traversal buffer. For details, see “GciStoreTrav” on page 384.
<code>int</code> <i>storeTravFlags</i>	A flag that determines how the objects should be handled. For details, see “GciStoreTrav” on page 384.
<code>int</code> <i>doPerform</i>	If this field is 0, this function executes a string using <i>args->u.executestr</i> , with the semantics of “GciExecuteStrFromContext” on page 157. If this field is 1, then executes a perform using <i>args->u.perform</i> , with the semantics of “GciPerformNoDebug” on page 300. Other values of this field are only for use with <code>GciStoreTravDoTravRefs_</code> or <code>GciNbStoreTravDoTravRefs_</code> .
<code>int</code> <i>doFlags</i>	Flags to disable or permit asynchronous events and debugging in Smalltalk, as described in “GciPerformNoDebug” on page 300. These flags apply whatever the value of <i>doPerform</i> .
<code>union</code> <i>u</i>	One of two structures containing appropriate input fields for the specified operation. The structure <i>u.perform</i> should be used when <i>doPerform</i> is set to 1, and <i>u.executestr</i> should be used when <i>doPerform</i> is set to 0. For more information on these structs and how to use them, see <code>gcicmn.ht</code> .
<code>OopType*</code> <i>alteredTheOops</i>	An array allocating memory for OOPs of objects that will be modified as a consequence of executing the specified code. For more information, see “GciAlteredObjs” on page 97.

int	<i>alteredNumOops</i> The number of OOPs in the previous array. On input, the caller must set this to the maximum number of OOPs that will fit in <i>alteredTheOops</i> . Upon completion, this field indicates the number of OOPs actually written to <i>alteredTheOops</i> .
BoolType	<i>alteredCompleted</i> Upon output, TRUE if the <i>alteredTheOops</i> contains the complete set of objects modified as a result of executing the specified code; false otherwise. If FALSE, call GciAlteredObjs for the rest of the modified objects.
const OopType*	<i>execBlock_args</i> This field is ignored.
int	<i>execBlock_numArgs</i> This field is ignored.

Return Value

Returns the OOP of the result of executing the specified code. In case of error, this function returns OOP_NIL.

Description

GciStoreTravDo_ works exactly as “GciStoreTrav” on page 384, and also executes the supplied code in the same network round-trip.

The description of “GciStoreTrav” on page 384 explains the first two arguments. If the value of the third argument is 1, see “GciPerformNoDebug” on page 300 for details of the next five arguments—flags to enable or disable asynchronous events, and the first nested structure.

If the value of the third argument is 2, see “GciExecuteStrFromContext” on page 157 for details on next set of arguments—flags to enable or disable asynchronous events, and the second nested structure of five arguments.

If the value of the third argument is 3, the arguments are similar to those for **GciExecuteStrFromContext**, but *source* must be a String that when compiled will return a Block. In this case, the last two arguments also are used, which provide the arguments, and the count of arguments, to be used to execute the compiled block.

The next five input arguments supply needed output after the function has completed. Read *alteredTheOops* to get the OOPs of the objects that were modified; read *alteredSymbolBuf* to get the pairs of symbols and symbol dictionaries for symbol canonicalization; finally, read *alteredCompleted* to determine if the array as originally allocated was large enough to hold all the modified objects. If the value is false, the array was too small and holds only some of the modified objects; in this case, call **GciAlteredObjs** for the rest.

Similarly to **GciStoreTrav**, **GciStoreTravDo_** provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see page 25.

Handling Error Conditions

If you get a run time error while executing `GciStoreTravDo_`, we recommend that you abort the current transaction.

See Also

- `GciAlteredObjs`, page 97
- `GciExecuteStrFromContext`, page 157
- `GciMoreTraversal`, page 234
- `GciNbMoreTraversal`, page 253
- `GciNbStoreTrav`, page 260
- `GciNbTraverseObjs`, page 265
- `GciNewOopUsingObjRep`, page 272
- `GciPerformNoDebug`, page 300
- `GciProcessDeferredUpdates_`, page 312
- `GciStoreTrav`, page 384
- `GciTraverseObjs`, page 408

GciStoreTravDoTrav_

Combine in a single function the calls to **GciStoreTravDo_** and **GciClampedTrav**, to store multiple traversal buffer values in objects, execute the specified code, and traverse the result object.

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciStoreTravDoTrav** (without the underscore).

Syntax

```
BoolType GciStoreTravDoTrav_(  
    GciStoreTravDoArgsSType *stdArgs,  
    GciClampedTravArgsSType *ctArgs );
```

Input Arguments

<i>stdArgs</i>	An instance of GciStoreTravDoArgsSType. For details, see page 387.
<i>ctArgs</i>	An instance of GciClampedTravArgsSType. For details, see page 108.

Return Value

Returns FALSE if the traversal is not yet completed. Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal** (that is, an object report was constructed for each object, minus the special objects).

Description

This function allows the client to execute behavior on the Gem and return the traversal of the result object in a single network round-trip. See the descriptions for **GciStoreTravDo_** on page 387 and **GciClampedTrav** on page 108 for details.

See Also

GciClampedTrav, page 108
GciStoreTrav, page 384
GciStoreTravDo_, page 387

GciStoreTravDoTravRefs_

Combine in a single function modifications to session sets, traversal of objects to the server, optional Smalltalk execution, and traversal to the client of changed objects and (optionally) the result object.

In GemStone/S 64 Bit releases earlier than 3.0, this function was named **GciStoreTravDoTravRefs** (without the underscore).

Syntax

```
int GciStoreTravDoTravRefs_(
    const OopType *      oopsNoLongerReplicated,
    int                  numNotReplicated,
    const OopType *      oopsGcedOnClient,
    int                  numGced,
    GciStoreTravDoArgsSType *stdArgs,
    GciClampedTravArgsSType *ctArgs );
```

Input Arguments

<i>oopsNoLongerReplicated</i>	An Array of objects to be removed from the PureExportSet and added to the ReferencedSet.
<i>numNotReplicated</i>	The number of elements in <i>oopsNoLongerReplicated</i> .
<i>oopsGcedOnClient</i>	An Array of objects to be removed from both the PureExportSet and ReferencedSet.
<i>numGced</i>	The number of elements in <i>oopsGcedOnClient</i> .
<i>stdArgs</i>	An instance of GciStoreTravDoArgsSType (as described in \$GEMSTONE/include/gcicmn.ht) containing the following fields:
GciTravBufType* <i>storeTravBuff</i>	The traversal buffer. For details, see “GciStoreTrav” on page 384.
int <i>storeTravFlags</i>	A flag that determines how the objects should be handled. For details, see “GciStoreTrav” on page 384.
int <i>doPerform</i>	<p>If this field is 0, this function executes a string using <i>args->u.executestr</i>, with the semantics of “GciExecuteStrFromContext” on page 157.</p> <p>If this field is 1, then executes a perform using <i>args->u.perform</i>, with the semantics of “GciPerformNoDebug” on page 300. If this field is 2, execute a string that is the source code for a Smalltalk block using <i>stdArgs->u.executestr</i>, passing the block arguments in <i>execBlock_args</i>. If this field is 3, perform no server Smalltalk execution, but traverse the object specified in <i>stdArgs->u.perform.receiver</i> as if it was the results of execution. If this field is 4, resume execution of a</p>

		suspended Smalltalk Process using <i>stdArgs->u.continueArgs</i> , with the semantics of “GciContinueWith” on page 121.
	int	<i>doFlags</i> Flags to disable or permit asynchronous events and debugging in Smalltalk, as described in “GciPerformNoDebug” on page 300. These flags apply whatever the value of <i>doPerform</i> .
	union	<i>u</i> One of three structures containing appropriate input fields for the specified operation. The structure <i>u.perform</i> should be used when <i>doPerform</i> is set to 1 or 3, <i>u.executestr</i> should be used when <i>doPerform</i> is set to 0 or 2, and <i>u.continueArgs</i> should be used when <i>doPerform</i> is set to 4. For more information on these structs and how to use them, see <i>gcicmn.ht</i> .
	OopType*	<i>alteredTheOops</i> This field is ignored.
	int	<i>alteredNumOops</i> This field is ignored.
	BoolType	<i>alteredCompleted</i> This field is not used.
	const OopType*	<i>execBlock_args</i> An array of the arguments to the block to be executed. Only applies if <i>doPerform</i> is 2, ignored otherwise.
	int	<i>execBlock_numArgs</i> The number of the arguments provided in <i>execBlock_args</i> . This must match the declared number of arguments in the block source string. Only applies if <i>doPerform</i> is 2, ignored otherwise.
<i>ctArgs</i>		An instance of GciClampedTravArgsSType . For details, see the discussion of GciClampedTrav on page 108, with one exception. The valid <i>retrievalFlags</i> are limited to: GCI_RETRIEVE_DEFAULT GCI_TRAV_REFS_EXCLUDE_RESULT_OBJ will suppress traversal of the result object. The altered objects will still be traversed to the specified level. No other <i>retrievalFlags</i> values should be used with this function.

Return Value

Returns an int with the following meaning:

- 0 – traversal of both altered objects and execution result completed.
- 1 – traversal buffer became full. You must call **GciMoreTraversal** to finish traversal of the altered and result objects.

Description

This function allows the client to modify the PureExportSet and ReferencedSet, modify or create any number of objects on the server, execute behavior on the Gem, and return the traversal of the changed objects and the result object, all in a single network round-trip.

The elements in *oopsGcedOnClient* are removed from both PureExportSet and ReferencedSet, and the elements in *oopsNoLongerReplicated* are removed from the PureExportSet and added to the ReferencedSet.

Objects in the ReferencedSet are protected from garbage collection, but may be faulted out of memory. Dirty tracking is not done on objects in the ReferencedSet.

Then per the *stdArgs*, a **GciStoreTrav** is done, which may modify or create any number of objects on the server. Newly created objects are added to the PureExportSet.

Then, if specified, Smalltalk execution is performed as in **GciPerformNoDebug**, **GciExecuteStrFromContext**, or executing the block code with the given arguments.

Finally, this function does a special **GciClampedTrav** starting with altered objects, followed by the execution result from the previous step. If no execution was specified, the specified object is traversed as if it was an execution result. Altered objects are those that would be returned from a **GciAlteredObjs** after the code execution step. This traversal both relies on the contents of the PureExportSet and ReferencedSet does not, and also modifies those sets in ways that **GciClampedTrav** does not. For details, see the comments in `gci.hf`.

GciStoreTravDoTravRefs_ is not intended for use within a user action.

See Also

GciClampedTrav, page 108

GciStoreTrav, page 384

GciStoreTravDo_, page 387

GciStoreTravDoTrav_, page 390

GciStringToInteger

Convert a C string to a GemStone SmallInteger or LargeInteger object.

Syntax

```
OopType GciStringToInteger(  
    const char*      string,  
    int64            stringSize );
```

Input Arguments

<i>string</i>	The C string to be translated into a GemStone SmallInteger or LargeInteger object.
<i>stringSize</i>	The length of <i>string</i> .

Return Value

Returns the OOP of the GemStone SmallInteger or LargeInteger object. If *string* has an invalid format, this function returns OOP_NIL without an error.

Description

The **GciStringToInteger** function translates a C string to a GemStone SmallInteger or LargeInteger object that has the same value.

Leading blanks are ignored. Trailing non-digits are ignored.

GciStrKeyValueDictAt

Find the value in a symbol KeyValue dictionary at the corresponding string key.

Syntax

```
void GciStrKeyValueDictAt(  
    OopType          theDict,  
    const char *     keyString,  
    OopType *        value );
```

Input Arguments

<i>theDict</i>	The OOP of a SymbolKeyValueDictionary.
<i>keyString</i>	The OOP of a key in the SymbolKeyValueDictionary.

Result Arguments

<i>value</i>	A pointer to the variable that is to receive the OOP of the returned value.
--------------	---

Description

Returns the value in symbol KeyValue dictionary *theDict* that corresponds to key *keyString*. If an error occurs or *keyString* is not found, *value* is OOP_ILLEGAL. KeyValue dictionaries do not have associations, so no association is returned. **GciStrKeyValueDictAt** is equivalent to **GciStrKeyValueDictAtObj** except that the key is a character string, not an object.

See Also

- GciStrKeyValueDictAtObj**, page 396
- GciStrKeyValueDictAtObjPut**, page 397
- GciStrKeyValueDictAtPut**, page 398

GciStrKeyValueDictAtObj

Find the value in a symbol KeyValue dictionary at the corresponding object key.

Syntax

```
void GciStrKeyValueDictAtObj(  
    OopType          theDict,  
    OopType          keyObj,  
    OopType*         value );
```

Input Arguments

<i>theDict</i>	The OOP of a SymbolKeyValueDictionary.
<i>keyObj</i>	The OOP of a key in the SymbolKeyValueDictionary.

Result Arguments

<i>value</i>	A pointer to the variable that is to receive the OOP of the returned value.
--------------	---

Description

Returns the value in symbol KeyValue dictionary *theDict* that corresponds to key *keyObj*. If an error occurs or *keyObj* is not found, *value* is OOP_ILLEGAL. KeyValue dictionaries do not have associations, so no association is returned. Equivalent to the GemStone Smalltalk expression:

```
^ { theDict at:keyObj }
```

See Also

- [GciStrKeyValueDictAt](#), page 395
- [GciStrKeyValueDictAtObjPut](#), page 397
- [GciStrKeyValueDictAtPut](#), page 398

GciStrKeyValueDictAtObjPut

Store a value into a symbol KeyValue dictionary at the corresponding object key.

Syntax

```
void GciStrKeyValueDictAtObjPut(  
    OopType          theDict,  
    OopType          keyObj,  
    OopType          theValue );
```

Input Arguments

<i>theDict</i>	The OOP of the SymbolKeyValueDictionary into which the object is to be stored.
<i>keyObj</i>	The OOP of the key under which the object is to be stored.
<i>theValue</i>	The OOP of the object to be stored in the SymbolKeyValueDictionary.

Description

Adds object *theValue* to symbol KeyValue dictionary *theDict* with key *keyObj*. Equivalent to the Smalltalk expression:

```
theDict at: keyObj put: theValue
```

See Also

[GciStrKeyValueDictAt](#), page 395
[GciStrKeyValueDictAtObj](#), page 396
[GciStrKeyValueDictAtPut](#), page 398

GciStrKeyValueDictAtPut

Store a value into a symbol KeyValue dictionary at the corresponding string key.

Syntax

```
void GciStrKeyValueDictAtPut(  
    OopType          theDict,  
    const char *     keyString,  
    OopType          theValue );
```

Input Arguments

<i>theDict</i>	The OOP of the SymbolKeyValueDictionary into which the object is to be stored.
<i>keyString</i>	The string key under which the object is to be stored.
<i>theValue</i>	The OOP of the object to be stored in the SymbolKeyValueDictionary.

Description

Adds object *theValue* to symbol KeyValue dictionary *theDict* with key *keyString*.

GciStrKeyValueDictAtPut is equivalent to **GciStrKeyValueDictAtObjPut**, except the key is a character string, not an object.

See Also

GciStrKeyValueDictAt, page 395
GciStrKeyValueDictAtObj, page 396
GciStrKeyValueDictAtObjPut, page 397

GciSwapBytesUint

Swap the byte order of an array of uint.

Syntax

```
void GciSwapBytesUint(  
    uint *          buf,  
    intptr_t       numChars );
```

Input Arguments

<i>buf</i>	An array of uint.
<i>numChars</i>	The size of the array.

Description

Swaps the byte order of the specified array of uint.

See Also

GciSwapBytesUshort, page 400

GciSwapBytesUshort

Swap the byte order of an array of ushort.

Syntax

```
void GciSwapBytesUshort(  
    ushort *      buf,  
    intptr_t      numChars );
```

Input Arguments

<i>buf</i>	An array of ushort.
<i>numChars</i>	The size of the array.

Description

Swaps the byte order of the specified array of ushort.

See Also

[GciSwapBytesUInt](#), page 399

GciSymDictAt

Find the value in a symbol dictionary at the corresponding string key.

Syntax

```
void GciSymDictAt(  
    OopType          theDict,  
    const char *     keyString,  
    OopType *        value,  
    OopType *        association );
```

Input Arguments

<i>theDict</i>	The OOP of a SymbolDictionary.
<i>keyString</i>	The OOP of a key in the SymbolDictionary.

Result Arguments

<i>value</i>	A pointer to the variable that is to receive the OOP of the returned value.
<i>association</i>	A pointer to the variable that is to receive the OOP of the association.

Description

Returns the value in symbol dictionary *theDict* that corresponds to key *keyString*. If an error occurs or *keyString* is not found, *value* is OOP_ILLEGAL. If *association* is not NULL and an error does not occur, stores the OOP of the association for *keyString* at **association*, or stores OOP_ILLEGAL if *keyString* was not found. Equivalent to **GciSymDictAtObj** except that the key is a character string, not an object.

To operate on kinds of Dictionary other than SymbolDictionary, such as KeyValueDictionary, use **GciPerform**, since the KeyValueDictionary class is implemented in Smalltalk. If your dictionary will be large (greater than 20 elements) a KeyValueDictionary is more efficient than a SymbolDictionary.

See Also

[GciSymDictAtObj](#), page 402
[GciSymDictAtObjPut](#), page 403
[GciSymDictAtPut](#), page 404

GciSymDictAtObj

Find the value in a symbol dictionary corresponding to the key object.

Syntax

```
void GciSymDictAtObj(  
    OopType          theDict,  
    OopType          keyObj,  
    OopType *        value,  
    OopType *        association );
```

Input Arguments

<i>theDict</i>	The OOP of a SymbolDictionary.
<i>keyObj</i>	The OOP of a key in the SymbolDictionary.

Result Arguments

<i>value</i>	A pointer to the variable that is to receive the OOP of the returned value.
<i>association</i>	A pointer to the variable that is to receive the OOP of the association.

Description

Fetches the value in symbol dictionary *theDict* that corresponds to key *keyObj*. If an error occurs or *keyObj* is not found, *value* is OOP_ILLEGAL. If *association* is not NULL and an error does not occur, stores the OOP of the association for *keyObj* at **association*, or stores OOP_ILLEGAL if *keyObj* was not found. Similar to the GemStone Smalltalk expression:

```
^ { theDict at: keyObj . theDict associationAt: keyObj }
```

See Also

[GciSymDictAt](#), page 401
[GciSymDictAtObjPut](#), page 403
[GciSymDictAtPut](#), page 404

GciSymDictAtObjPut

Store a value into a symbol dictionary at the corresponding object key.

Syntax

```
void GciSymDictAtObjPut(  
    OopType      theDict,  
    OopType      keyObj,  
    OopType      theValue );
```

Input Arguments

<i>theDict</i>	The OOP of the SymbolDictionary into which the value is to be stored.
<i>keyObj</i>	The OOP of the key under which the value is to be stored.
<i>theValue</i>	The OOP of the object to be stored in the SymbolDictionary.

Description

Adds object *theValue* to symbol dictionary *theDict* with key *keyObj*. Equivalent to the Smalltalk expression:

```
theDict at: keyObj put: theValue
```

See Also

[GciSymDictAt](#), page 401
[GciSymDictAtObj](#), page 402
[GciSymDictAtPut](#), page 404

GciSymDictAtPut

Store a value into a symbol dictionary at the corresponding string key.

Syntax

```
void GciSymDictAtPut(  
    OopType          theDict,  
    const char *     keyString,  
    OopType          theValue );
```

Input Arguments

<i>theDict</i>	The OOP of the SymbolDictionary into which the object is to be stored.
<i>keyString</i>	The string key under which the object is to be stored.
<i>theValue</i>	The OOP of the object to be stored in the SymbolDictionary.

Description

Adds object *theValue* to symbol dictionary *theDict* with key *keyString*. Equivalent to **GciSymDictAtObjPut**, except the key is a character string, not an object.

See Also

GciSymDictAt, page 401
GciSymDictAtObj, page 402
GciSymDictAtObjPut, page 403

GciTrackedObjsFetchAllDirty

Find all exported or tracked objects that have changed and are therefore in the ExportedDirtyObjs or TrackedDirtyObjs sets.

Syntax

```
void GciTrackedObjsFetchAllDirty(
    OopType          exportedDirty,
    int64 *          numExportedDirty,
    OopType          trackedDirty,
    int64 *          numTrackedDirty );
```

Input Arguments

<i>exportedDirty</i>	OOP of the collection (an instance of either IdentitySet or IdentityBag) that will contain the objects in the ExportedDirtyObjs set.
<i>trackedDirty</i>	OOP of the collection (an instance of either IdentitySet or IdentityBag) that will contain the objects in the TrackedDirtyObjs set.

Result Arguments

<i>numExportedDirty</i>	Pointer to an integer that returns the number of objects in the <i>exportedDirty</i> collection.
<i>numTrackedDirty</i>	Pointer to an integer that returns the number of objects in the <i>trackedDirty</i> collection.

Description

GciTrackedObjsFetchAllDirty fetches all dirty objects and sorts them into two categories:

- ▶ Objects in the ExportedDirtyObjs set - objects in the PureExportSet that have been changed since the ExportedDirtyObjs set was initialized or cleared.
- ▶ Objects in the TrackedDirtyObjs set - objects in the GciTrackedObjs set that have been changed since the TrackedDirtyObjs set was initialized or cleared.

The ExportedDirtyObjs set is initialized by **GciDirtyObjsInit**; it is cleared by calls to **GciDirtyAlteredObjs**, **GciDirtyExportedObjs**, **GciDirtySaveObjs**, or **GciTrackedObjsFetchAllDirty** (this function). The TrackedDirtyObjs set is initialized by **GciTrackedObjsInit** and cleared by calls to **GciDirtyAlteredObjs**, **GciDirtySaveObjs**, **GciDirtyTrackedObjs**, or **GciTrackedObjsFetchAllDirty** (this function).

An object is considered dirty (changed) under one or more of the following conditions:

- ▶ The object was changed by Smalltalk execution.
- ▶ The object was changed by a call to any GemBuilder function from within a user action.
- ▶ The object was changed by a call to one or more of the following functions: **GciStorePaths**, **GciSymDictAtObjPut**, **GciSymDictAtPut**, **GciStrKeyValueDictAtObjPut**, or **GciStrKeyValueDictAtPut**.

- ▶ A change to the object was committed by another transaction since it was read by this one.
- ▶ The object is persistent, but was modified in the current session before the session aborted the transaction. (When the transaction is aborted, the modifications are destroyed, thus changing the state of the object in memory).

You must call both **GciDirtyObjsInit** and **GciTrackedObjsInit** once after **GciLogin** before calling **GciTrackedObjsFetchAllDirty**.

Note that the `ExportedDirtyObjs` and `TrackedDirtyObjs` sets are cleared when this function is executed.

See Also

Garbage Collection, page 40
GciDirtyExportedObjs, page 135
GciDirtyObjsInit, page 137
GciDirtySaveObjs, page 138
GciDirtyTrackedObjs, page 140
GciTrackedObjsInit, page 407

GciTrackedObjsInit

Reinitialize the set of tracked objects maintained by GemStone.

Syntax

```
void GciTrackedObjsInit();
```

Description

The **GciTrackedObjsInit** function permits an application to request GemStone to maintain a set of tracked objects. **GciTrackedObjsInit** must be called once after **GciLogin** before other tracked objects functions in order for those functions to operate properly, because they depend upon GemStone's set of tracked objects.

See Also

GciDirtySaveObjs, page 138
GciDirtyTrackedObjs, page 140
GciHiddenSetIncludesOop, page 211
GciTrackedObjsFetchAllDirty, page 405

GciTraverseObjs

Traverse an array of GemStone objects.

Syntax

```
BoolType GciTraverseObjs(  
    const OopType      theOops [ ],  
    int                numOops,  
    GciTravBufType *  travBuff,  
    int                level );
```

Input Arguments

<i>theOops</i>	An array of OOPs representing the objects to traverse.
<i>numOops</i>	The number of elements in <i>theOops</i> .
<i>level</i>	Maximum traversal depth. When the level is 1, an object report is written to the traversal buffer for each element in <i>theOops</i> . When level is 2, an object report is also obtained for the instance variables of each level-1 object. When level is 0, the number of levels in the traversal is not restricted.

Result Arguments

<i>travBuff</i>	A buffer in which the results of the traversal will be placed.
-----------------	--

Return Value

Returns FALSE if the traversal is not yet completed. Returns TRUE if there are no more objects to be returned by subsequent calls to **GciMoreTraversal**.

Description

This function allows you to reduce the number of GemBuilder calls that are required for your application program to obtain information about complex objects in the database.

NOTE

*This function is most useful with applications that are linked with the "remote procedure call" version of GemBuilder. If your application will be linked with the "linkable" GemBuilder, you'll usually achieve best performance by using the simple **GciFetch...** and **GciStore...** functions instead. For more information, see "RPC and Linked applications" on page 43.*

There are no built-in limits on how much information can be obtained in the traversal. You can use the *level* argument to restrict the size of the traversal.

GciTraverseObjs provides automatic byte swizzling, unless **GciSetTraversalBufSwizzling** is used to disable swizzling. For more about byte swizzling, see "Byte-Swizzling of Binary Floating-Point Values" on page 25.

Organization of the Traversal Buffer

The first element placed in a traversal buffer is an integer that indicates how many bytes were actually stored in the buffer by this function. The remainder of the traversal buffer consists of a series of object reports, each of which is of type **GciObjRepSType**, as described on page 385.

In order for the traversal buffer to accommodate m objects, each of which is of size n bytes, your application should allocate at least enough memory so that the traversal buffer's size can be assigned according to the following formula:

```
GciTravBufType* travBufAllocation_example(void)
{
    int numObjs = 100;
    int bodyBytesPerObj = 1000;
    size_t allocationSize =
numObjs * GCI_ALIGN(sizeof(GciObjRepHdrSType) + bodyBytesPerObj);

    GciTravBufType *buf = GciTravBufType::malloc(allocationSize);
    return buf;
}
```

The macro `GCI_ALIGN` ensures that the value buffer portion of each object report begins at a word boundary.

This function ensures that each object report header and value buffer begins on a word boundary. To provide proper alignment, 0 to 7 bytes may be inserted between each header and value buffer.

The Value Buffer

The object report's *value buffer* begins at the first byte following the object report header. For byte objects, the value buffer `rpt->valueBufferBytes()` is an array of type **ByteType**; for pointer objects and NSCs, the buffer `rpt->valueBufferOops()` is an array of type **OopType**. The size of the report's value buffer (`rpt->hdr.valueBuffSize`) is the number of bytes of the object's value returned by this traversal. That number is no greater than the size of the object.

How This Function Works

This section explains how **GciTraverseObjs** stores object reports in the traversal buffer and values in the value buffer.

1. First, **GciTraverseObjs** verifies that the traversal buffer is large enough to accommodate at least one object report header (**GciObjRepHdrSType**). If the buffer is too small, GemBuilder returns an error.
2. For each object in the traversal, **GciTraverseObjs** discovers if there is enough space left in the traversal buffer to store both the object report header and the object's values. If there isn't enough space remaining, the function returns `FALSE`, and your program can call **GciMoreTraversal** to continue the traversal. Otherwise (if there is enough space), the object's values are stored in the traversal buffer.
3. When there are no more objects left to traverse, **GciTraverseObjs** returns `TRUE` to indicate that the traversal is complete.

Special Objects

For each occurrence of an object with a special implementation (that is, an instance of `SmallInteger`, `Character`, `Boolean`, or `UndefinedObject`) contained in *theOops*, this function will return an accurate object report. For any special object encountered at some deeper point in the traversal, no object report will be generated.

Authorization Violations

If the user is not authorized to read some object encountered during the traversal, the traversal will continue. No value will be placed in the object report's value buffer, but the report for the forbidden object will contain the following values:

```

hdr.valueBuffSize    0
hdr.namedSize        0
hdr.idxSize          0
hdr.firstOffset      1
hdr.objId            theOop
hdr.oclass           OOP_NIL
hdr.objectSecurityPolicyId  0
hdr.objImpl          GC_FORMAT_SPECIAL
hdr.isInvariant      0

```

Incomplete Object Reports

It is possible for an object report to not contain all the instance variables of an object, due to traversal specifications or buffer size limitations. The value buffer is incomplete when *hdr.isPartial()* returns non-zero.

Continuing the Traversal

When the amount of information obtained in a traversal exceeds the amount of available memory, your application can break the traversal into manageable amounts of information by issuing repeated calls to **GciMoreTraversal**. Generally speaking, an application can continue to call **GciMoreTraversal** until it has obtained all requested information.

During the entire sequence of **GciTraverseObjs** and **GciMoreTraversal** calls that constitute a traversal, any single object report will be returned exactly once. Regardless of the connectivity of objects in the GemStone database, only one report will be generated for any non-special object.

When Traversal Can't Be Continued

Naturally, GemStone will not continue an incomplete traversal if there is any chance that changes to the database in the intervening period might have invalidated the previous report or changed the connectivity of the objects in the path of the traversal. Specifically, GemStone will refuse to continue a traversal if, in the interval before attempting to continue, you:

- Modify the objects in the database directly, by calling any of the **GciStore...** or **GciAdd...** functions;
- Call one of the Smalltalk message-sending functions **GciPerform**, **GciContinue**, or any of the **GciExecute...** functions;
- Abort your transaction, thus invalidating any subsequent information from that traversal.

Any attempt to call **GciMoreTraversal** after one of these actions will generate an error.

Note that this holds true across multiple GemBuilder applications sharing the same GemStone session. Suppose, for example, that you were holding on to an incomplete traversal buffer and the user moved from the current application to another, did some work that required executing Smalltalk code, and then returned to the original application. You would be unable to continue the interrupted traversal.

Example

For an example of how **GciTraverseObjs** is used, see **GciMoreTraversal** on page 234.

See Also

- GciFindObjRep, page 200
- GciMoreTraversal, page 234
- GciNbMoreTraversal, page 253
- GciNbStoreTrav, page 260
- GciNbTraverseObjs, page 265
- GciNewOopUsingObjRep, page 272
- GciObjRepSize_, page 282
- GciStoreTrav, page 384

GciUncompress

Uncompress the supplied data, assumed to have been compressed with **GciCompress**.

Syntax

```
int GciUncompress(  
    char *          dest,  
    uint *         destLen,  
    const char *   source,  
    uint           sourceLen );
```

Input Arguments

<i>dest</i>	Pointer to the buffer intended to hold the resulting uncompressed data.
<i>destLen</i>	Length, in bytes, of the buffer intended to hold the uncompressed data.
<i>source</i>	Pointer to the source data to uncompress.
<i>sourceLen</i>	Length, in bytes, of the source data.

Result Arguments

<i>dest</i>	The resulting uncompressed data.
-------------	----------------------------------

Return Value

GciUncompress returns `Z_OK` (equal to 0) if the decompression succeeded, or various error values if it failed; see the documentation for the `uncompress` function in the GNU `zlib` library at <http://www.gzip.org>.

Description

GciUncompress passes the supplied inputs unchanged to the `uncompress` function in the GNU `zlib` library Version 1.2.3, and returns the result exactly as the GNU `uncompress` function returns it.

See Also

`GciCompress`, page 118

GciUserActionInit

Declare user actions for GemStone.

Syntax

```
void GciUserActionInit()
```

Description

GciUserActionInit is implemented by the application developer, but it is called by **GciInit**. It enables Smalltalk to find the entry points for the application's user actions, so that they can be executed from the database.

See Also

Chapter 3, "Writing C Functions To Be Called from GemStone", starting on page 47
"The User Action Information Structure" on page 78
[GciDeclareAction](#), page 132
[GciInstallUserAction](#), page 218
[GciUserActionShutdown](#), page 414

GciUserActionShutdown

Enable user-defined clean-up for user actions.

Syntax

```
void GciUserActionShutdown()
```

Description

GciUserActionShutdown is implemented by the application developer, and is called when a session user action library is unloaded. It enables user-defined clean-up for the application's user actions.

See Also

Chapter 3, "Writing C Functions To Be Called from GemStone", starting on page 47
"The User Action Information Structure" on page 78
GciDeclareAction, page 132
GciInstallUserAction, page 218
GciUserActionInit, page 413

GciVersion

Return a string that describes the GemBuilder version.

Syntax

```
const char* GciVersion()
```

Description

GciVersion returns a string terminated by 0, containing fields that describe the specific release of GemBuilder. Fields in the string are delimited by a period (.).

For more version information, use the methods in class System in the Version Management category.

See Also

GciProduct, page 313

Reserved OOPs

The GemStone/S 64 Bit distribution includes the file `gcioop.ht`, which defines C mnemonics for the OOPs of certain GemStone objects that are already defined in your GemStone software package. Your C application can compare all these mnemonics with any value of type `OopType`. However, the value of any mnemonic is subject to change without notice in future software releases. Your C application should refer to the OOPs of predefined GemStone objects *by mnemonic name only*.

The following mnemonic names for predefined GemStone objects are available to C programs. Refer to `$GEMSTONE/include/gcioop.ht` for more information.

- ▶ A value that, strictly speaking, is not an object at all, but that represents a value that is never used to represent any object in the database. You can use this mnemonic to test whether or not an OOP is valid, that is, whether or not it actually points to any GemStone object.
 - ▶ `OOP_ILLEGAL`
- ▶ Special objects
 - ▶ `OOP_NIL` (*nil*)
 - ▶ `OOP_FALSE` (*FALSE*)
 - ▶ `OOP_TRUE` (*true*)
- ▶ Instances of `SmallInteger`
 - ▶ `OOP_MinusOne`
 - ▶ `OOP_Zero`
 - ▶ `OOP_One`
 - ▶ `OOP_Two`
- ▶ Instances of `Character`
 - ▶ `OOP_ASCII_NUL` represents the first ASCII character OOP. Other Characters have specific OOPs, but no mnemonics.

- ▶ The GemStone Smalltalk kernel classes
 - ▶ `OOP_CLASS_className`. In reserved OOP mnemonics, most class names are in capital letters, with words separated by underscore characters; most recently added class names are in the same case as their Smalltalk class name. Refer to `gcioop.ht` for specific names.
 - ▶ `OOP_LAST_KERNEL_OOP` (which has the same value as the last class)
- ▶ The GemStone error dictionary
 - ▶ `OOP_GEMSTONE_ERROR_CAT`
- ▶ The global cluster bucket
 - ▶ `OOP_ALL_CLUSTER_BUCKETS`

GemStone C Statistics Interface

This appendix describes the GemStone C Statistics Interface (GCSI), a library of functions that allow your C application to collect GemStone statistics directly from the shared page cache without starting a database session.

B.1 Developing a GCSI Application

The command lines in this appendix assume that you have set the `GEMSTONE` environment variable to your GemStone installation directory.

Required Header Files

Your GCSI program must include the following header files:

- ▶ `$GEMSTONE/include/shrpcstats.ht` – Defines all cache statistics. (For a list of cache statistics, refer to the “Monitoring GemStone” chapter of the *System Administration Guide*.)
- ▶ `$GEMSTONE/include/gcsi.hf` – Prototypes for all GCSI functions.
- ▶ `$GEMSTONE/include/gcsierr.ht` – GCSI error numbers.

Your program must define a `main()` function somewhere.

The GCSI Shared Library

GemStone provides a shared library, `$GEMSTONE/lib/libgcsi-3.3.0-64.so`, that your program will load at runtime.

- ▶ Make sure that `$GEMSTONE/lib` is included in your `LD_LIBRARY_PATH` environment variable, so that the runtime loader can find the GCSI library. For example:

```
export LD_LIBRARY_PATH=$GEMSTONE/lib:$LD_LIBRARY_PATH
```

- ▶ `$GEMSTONE/lib/libgcsi-3.3.0-64.so` is a multi-threaded library, so your program must also be compiled and linked as a multi-threaded program.

Compiling and Linking

The `$GEMSTONE/examples` directory includes the sample GCSI program `gsstat.cc`, along with a set of sample makefiles that show how to compile the sample GCSI program, using the compilers that are used to build the GemStone product.

NOTE

It may still be possible to build your program with another compiler, as long as you specify the appropriate flags to enable multi-threading.

Whenever you upgrade to a new GemStone version, you must re-compile and re-link all your GCSI programs. This is because the internal structure of the shared cache may change from version to version. Assuming you've created a makefile, all you should need to do is change `$GEMSTONE` and rebuild.

Connecting to the Shared Page Cache

The GCSI library allows your program to connect to a single GemStone shared page cache. Once the connection is made, a thread is started to monitor the cache and disconnect from it if the cache monitor process dies. This thread is needed to prevent your program from "holding on" to the shared cache after all other processes have detached from it. In this way, your program can safely sleep for a long time without preventing the operating system from freeing and recycling shared memory should the Stone be unexpectedly shut down.

The Sample Program

The sample program `gsstat.cc` (in `$GEMSTONE/examples`) monitors a running GemStone repository by printing out a set of statistics at a regular interval that you specify. The program prints the following statistics:

- ▶ `Sess` – `TotalSessionsCount`; the total number of sessions currently logged in to the system.
- ▶ `CR` – `CommitRecordCount`; the number of outstanding commit records that are currently being maintained by the system.
- ▶ `PNR` – `PagesNeedReclaimSize`; the amount of reclamation work that is pending, that is, the backlog waiting for the `GcGem` reclaim task.
- ▶ `PD` – `PossibleDeadSize`; the number of objects previously marked as dereferenced in the repository, but for which sessions currently in a transaction might have created a reference in their object space.
- ▶ `DNR` – `DeadNotReclaimedSize`; the number of objects that have been determined to be dead (current sessions have indicated they do not have a reference to these objects) but have not yet been reclaimed.
- ▶ `FP` – The number of free pages in the Stone.
- ▶ `OCS` – `OldestCrSession`; the session ID of the session referencing the oldest commit record. Prints 0 if the oldest commit record is not referenced by any session, or if there is only one commit record.

▶ FF – FreeFrameCount; the number of unused page frames in the shared page cache.

To invoke `gsstat`, supply the name of a running Stone (or shared page cache, if running on a Gem server) and a time interval in seconds. For example:

```
% gsstat myStone 2
```

To stop the `gsstat` program and detach from the cache, issue a CTRL-C.

B.2 GCSI Data Types

The following C types are used by GCSI functions. The file `shrpcstats.ht` defines each of the GCSI types (shown in capital letters below). That file is in the `$GEMSTONE/include` directory.

ShrPcMonStatSType

Shared page cache monitor statistics.

ShrPcStnStatSType

Stone statistics.

ShrPcPgsvrStatSType

Page server statistics.

ShrPcGemStatSType

Gem session statistics.

ShrPcStatUnion

The union of all four statistics structured types: shared page cache monitor, page server, Stone, and Gem.

ShrPcCommonStatSType

Common statistics collected for all processes attached to the shared cache.

The Structure for Representing the GCSI Function Result

The structured type `GcsiResultSType` provides a C representation of the result of executing a GCSI function. This structure contains the following fields:

```
typedef struct {
    signed int          processId;
    signed int          sessionId;
    ShrPcCommonStatSType cmn;
    union ShrPcStatUnion u;
} ShrPcStatsSType;

class GcsiResultSType {
public:
    char                vsdName[SHRPC_PROC_NAME_SIZE + 1];
    unsigned int        statType;
    ShrPcStatsSType    stats;
};
```

In addition, a set of C mnemonics support representation of the count of each process-specific structured type.

```
#define COMMON_STAT_COUNT (sizeof(ShrPcCommonStatSType)/  
sizeof(int))  
  
#define SHRPC_STAT_COUNT (sizeof(ShrPcMonStatSType)/sizeof(int)  
+ COMMON_STAT_COUNT)  
  
#define GEM_STAT_COUNT (sizeof(ShrPcGemStatSType)/sizeof(int) +  
COMMON_STAT_COUNT)  
  
#define PGSVR_STAT_COUNT (sizeof(ShrPcPgsvrStatSType)/  
sizeof(int) + COMMON_STAT_COUNT)  
  
#define STN_STAT_COUNT (sizeof(ShrPcStnStatSType)/sizeof(int) +  
COMMON_STAT_COUNT)
```

GcsiAllStatsForMask

Get all cache statistics for a specified set of processes.

Syntax

```
int GcsiAllStatsForMask(mask, result, resultSize);
    unsigned int      mask;
    GcsiResultSType * result;
    int *             resultSize;
```

Input Arguments

<i>mask</i>	Indicates what types of processes to collect statistics for.
<i>result</i>	Address of an array of kind GcsiResultSType where statistics will be stored.
<i>resultSize</i>	Pointer to an integer that indicates the size of the result in elements (not bytes). On return, indicates the number of that were stored into <i>result</i> . Indicates the maximum number of processes for which statistics can be returned.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

Example

Mask bits should be set by a bitwise OR of the desired process types. For example, to get statistics for the stone and Shared Page Cache Monitor:

```
unsigned int mask = SHRPC_MONITOR | SHRPC_STONE;
```

GcsiAttachSharedCache

Attach to the specified shared page cache.

Syntax

```
int GcsiAttachSharedCache(fullCacheName, errBuf, errBufSize);
    const char *      fullCacheName;
    char *           errBuf;
    size_t           errBufSize;
```

Input Arguments

<i>fullCacheName</i>	Full name of the shared page cache, in the format <i>stoneName@stoneHostIpAddress</i> . To determine the full name of the shared cache, use the <code>gslis -x</code> utility.
<i>errBuf</i>	A buffer that will contain a string describing an error.
<i>errBufSize</i>	Size (in bytes) of <i>errBuf</i> .

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in `gcsierr.ht`.

See Also

[GcsiAttachSharedCacheForStone](#), page 425
[GcsiDetachSharedCache](#), page 426

GcsiAttachSharedCacheForStone

Attaches this process to the specified shared page cache.

Syntax

```
int GcsiAttachSharedCacheForStone( stoneName, errBuf, errBufSize);  
    const char *      stoneName;  
    char *           errBuf;  
    size_t          errBufSize;
```

Input Arguments

<i>stoneName</i>	Name of the Stone process.
<i>errBuf</i>	A buffer that will contain a string describing an error.
<i>errBufSize</i>	Size (in bytes) of <i>errBuf</i> .

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in `gcsierr.ht`.

Description

This function assumes that the cache name is `<stoneName>@<thisIpAddress>` where `thisIpAddress` is the IP address of the local machine. This function may fail if the host is multi-homed (has more than one network interface). In that case, use `GcsiAttachSharedCache` (page 424) to specify the full name of the shared cache.

See Also

`GcsiAttachSharedCache`, page 424
`GcsiDetachSharedCache`, page 426

GcsiDetachSharedCache

Detach from the shared page cache.

Syntax

```
int GcsiDetachSharedCache (void);
```

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

GcsiAttachSharedCache, page 424

GcsiAttachSharedCacheForStone, page 425

GcsiFetchMaxProcessesInCache

Return the maximum number of processes that can be attached to this shared cache at any instant. The result may be used to allocate memory for a calls to the GcsiFetchStatsForAll* family of functions.

Syntax

```
int GcsiFetchMaxProcessesInCache(maxProcesses);  
int * maxProcesses;
```

Input Arguments

<i>maxProcesses</i>	The maximum number of processes that can be attached to this shared cache at any instant.
---------------------	---

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

Gcsilnit

Initialize the library. This function must be called before all other GCSI functions.

Syntax

```
GcsiInit(void);
```

GcsiStatsForGemSessionId

Get the cache statistics for the given Gem session id.

Syntax

```
int GcsiStatsForGemSessionId(sessionId, result);  
    int                sessionId;  
    GcsiResultSType *  result;
```

Input Arguments

<i>sessionId</i>	Session ID of the Gem for which statistics are requested.
<i>result</i>	Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

- GcsiStatsForGemSessionWithName, page 430
- GcsiStatsForPgsvrSessionId, page 431
- GcsiStatsForProcessId, page 432
- GcsiStatsForShrPcMon, page 433
- GcsiStatsForStone, page 434

GcsiStatsForGemSessionWithName

Get the cache statistics for the first Gem in the cache with the given cache name.

Syntax

```
int GcsiStatsForGemSessionWithName(gemName, result);
char *                               gemName;
GcsiResultSType *                    result;
```

Input Arguments

<i>gemName</i>	The case-sensitive name of the Gem for which statistics are requested.
<i>result</i>	Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

, page 428
GcsiStatsForPgsvrSessionId, page 431
GcsiStatsForProcessId, page 432
GcsiStatsForShrPcMon, page 433
GcsiStatsForStone, page 434

GcsiStatsForPgsvrSessionId

Get the cache statistics for the given page server with the given session id. Remote Gems have page servers on the Stone's cache that assume the same session ID as the remote Gem.

Syntax

```
int GcsiStatsForPgsvrSessionId(sessionId, result);
    int                sessionId;
    GcsiResultSType *  result;
```

Input Arguments

<i>sessionId</i>	Session ID of the page server for which statistics are requested.
<i>result</i>	Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

- GcsiStatsForGemSessionId, page 429
- GcsiStatsForGemSessionWithName, page 430
- GcsiStatsForProcessId, page 432
- GcsiStatsForShrPcMon, page 433
- GcsiStatsForStone, page 434

GcsiStatsForProcessId

Get the cache statistics for the given process ID.

Syntax

```
int GcsiStatsForProcessId(pid, result);  
    int          pid;  
    GcsiResultSType * result;
```

Input Arguments

<i>pid</i>	Process ID for which statistics are requested.
<i>result</i>	Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

- GcsiStatsForGemSessionId, page 429
- GcsiStatsForGemSessionWithName, page 430
- GcsiStatsForPgsvrSessionId, page 431
- GcsiStatsForShrPcMon, page 433
- GcsiStatsForStone, page 434

GcsiStatsForShrPcMon

Get the cache statistics for the shared page cache monitor process for this shared page cache.

Syntax

```
int GcsiStatsForShrPcMon(result);  
GcsiResultSType * result;
```

Input Arguments

result Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

GcsiInit, page 428
GcsiStatsForGemSessionId, page 429
GcsiStatsForGemSessionWithName, page 430
GcsiStatsForPgsvrSessionId, page 431
GcsiStatsForProcessId, page 432
GcsiStatsForStone, page 434

GcsiStatsForStone

Get the cache statistics for the Stone if there is a Stone attached to this shared page cache.

Syntax

```
int GcsiStatsForStone(result);  
    GcsiResultSType * result;
```

Input Arguments

result Pointer to a GcsiResultSType structure.

Return Value

Returns 0 if successful; otherwise returns an error code, as defined in gcsierr.ht.

See Also

, page 428
GcsiStatsForGemSessionId, page 429
GcsiStatsForGemSessionWithName, page 430
GcsiStatsForPgsvrSessionId, page 431
GcsiStatsForProcessId, page 432
GcsiStatsForStone, page 434

GCSI Errors

The following errors are defined for the GemStone C Statistics Interface.

Table 1 GCSI Errors

Error Name	Definition
GCSI_SUCCESS	The requested operation was successful.
GCSI_ERR_NO_INIT	GcsiInit() must be called before any other Gcsi functions.
GCSI_ERR_CACHE_ALREADY_ATTACHED	The requested shared cache is already attached.
GCSI_ERR_NOT_FOUND	The requested session or process was not found.
GCSI_ERR_BAD_ARG	An invalid argument was passed to a Gcsi function.
GCSI_ERR_CACHE_CONNECTION_SEVERED	The connection to the shared cache was lost.
GCSI_ERR_NO_STONE	Stone statistics were requested on a cache with no stone process.
GCSI_ERR_CACHE_NOT_ATTACHED	No shared page cache is currently attached.
GCSI_ERR_NO_MORE_HANDLES	The maximum number of shared caches are attached.
GCSI_ERR_CACHE_ATTACH_FAILED	The attempt to attach the shared cache has failed.
GCSI_ERR_WATCHER_THREAD_FAILED	The cache watcher thread could not be started.
GCSI_ERR_CACHE_WRONG_VERSION	The shared cache version does not match that of the libgcsixx.so library.

Index

A

aborting transactions 21
adding
 OOBs to an indexable object (Collection) 375, 377
 OOBs to an NSC 92, 93
alignment of traversal buffer 94
allocating
 multiple OOBs 206, 208
 OOBs 204
altered objects
 finding 97
appending
 to a byte object 99, 100
 to a collection 101
application
 binding 44, 66
 compiling 63
 improving performance 34, 36, 185, 234, 253, 260, 261, 265, 319, 320, 321, 322, 325, 379, 384, 387, 390, 408
 linking 18
authorization
 traversal 410
 violation, what to do 21

B

beginning
 a transaction 102
binding to GemBuilder 44, 66
boolean

 converting to an object 103
 represented as a special object 23
byte array
 converting to a C pointer 104
byte object
 creating 266
 fetching bytes from 161, 162
 fetching characters from 164
 fetching the size 174, 191
 implementation type 29, 30, 180
 initializing 266
 storing bytes in 360, 361, 363
bytes
 appending 99

C

C mnemonic
 sizes and offsets into objects 70
C types defined for GemBuilder functions 70
call
 determining if in progress 105
call stack
 clearing 40, 113
calling
 the virtual machine 120, 149, 153, 157, 240, 245, 246, 249, 254, 298, 300, 302, 304, 306
 user actions from GemStone 53
changed object, and re-reading 106
changing class definitions 28
character
 converting to an object 107

- instance defined in GemStone 417
 - represented as a special object 23
 - character object
 - creating 267
 - initializing 267
 - characters
 - converting to objects 107
 - checking for GemBuilder errors 39, 148
 - clamped object traversal 108, 110, 239, 262, 390, 391
 - structured type 108, 160
 - class
 - compiling methods 19, 110, 116
 - fetching an object's 29, 165
 - modifying 28
 - object report 200
 - clearing the call stack 40
 - cluster bucket
 - mnemonic for category 418
 - committing transactions 21, 115, 239
 - compiling
 - applications 63
 - C code 62
 - class methods 19, 110, 116
 - instance methods 19, 219
 - methods 19, 219
 - user actions 63
 - compressing objects 118, 412
 - concurrency conflict 115, 239
 - what to do 21
 - configuration files 216
 - constraint violation, what to do 21
 - context
 - call stack 113
 - error handling 40
 - of GemStone system 120, 121, 240, 241
 - continuable error 120, 121, 240, 241
 - continuing
 - after an error 39, 120, 121, 240, 241
 - traversal 234, 410
 - controlling
 - sessions 115, 229, 231, 232, 239, 347
 - transactions 91, 115, 236, 239
 - converting between
 - byte arrays and pointers 104, 308
 - objects and booleans 103, 288, 289
 - objects and characters 107, 290, 291, 292, 293
 - objects and floating point numbers 142, 203, 294
 - objects and integers 213, 296, 297
 - special objects and C values 23
 - strings and integers 394
 - v1.1.1 oops and OopTypes 284
 - creating
 - byte object 266
 - character object 267
 - class methods 110, 116
 - database objects 18
 - DateTime object 268
 - GemStone sessions 21, 229, 231, 347
 - instances of a GemStone class 28
 - objects 33, 269, 270, 272
 - OOPs 269, 270
 - current session, defined 21
- ## D
- date, time
 - structured type 71
 - DateTime object
 - creating 268
 - fetching contents of 167
 - initializing 268
 - debugging 345
 - function, enabling 128, 130, 131
 - information, finding 314, 339, 341, 343
 - use GciRpc 43, 66
 - user action 54
 - decoding
 - an OOP array 133
 - decrementing shared counter 134
 - default
 - directory, host file access 39
 - login parameter value 215, 347
 - defining
 - new methods 19
 - deprecated functions
 - GciPathToStr 298
 - GciStrToPath 399
 - developing a user action 48
 - dirty objects 97, 135, 137, 138, 140, 405, 407
 - disabling
 - error handlers 339, 341, 343
 - dynamic instance variable
 - fetching contents of 168, 169, 366
- ## E
- enabling
 - debugging functions 128, 130, 131
 - error handlers 339, 341, 343
 - full compression 144
 - run-length encoding 143
 - signaled errors 145

- encoding
 - an OOP array 146
 - free OOPs 143
- enumerating named instance variables 112, 227
- environmentId 80
- error
 - checking 39, 148, 345
 - continuing execution after 120, 121, 240, 241
 - dictionary 23, 418
 - handling 311, 314, 339, 341, 343
 - jump buffer 40, 311, 314, 339, 341, 343
 - mnemonics 69, 70
 - polling 39, 148, 309, 310
- error report
 - structured type 72
- errors 233, 345, 346
 - signaled 145
- executing code in
 - GemBuilder, advantages over GemStone 19
 - GemStone 19, 26, 149, 151, 153, 155, 157, 159, 245, 246, 249, 251, 258, 304
 - advantages over GemBuilder 19
 - host file access method 39
- executing user action 54
- execution environment
 - and Ruby 80
- export set 97, 135, 137, 138, 140, 336, 337, 338, 405, 407
- exporting objects to GemStone 18, 28

F

- false, GemStone special object 23, 24, 103, 417
- fetching
 - bytes from a byte object 161, 162
 - characters from a byte object 164
 - class 29, 165
 - DateTime 167
 - dynamic instance variables 168, 169
 - number of shared counters 177
 - object implementation format 29, 180
 - object information 178
 - object size 174, 191, 199
 - objects by using paths 36, 185
 - OOPs from a pointer object 170, 172, 195, 197
 - OOPs from an indexable object (Collection) 31, 181, 183
 - OOPs from an NSC 181, 183
 - shared counter value 190, 316, 317
- finding
 - debugging information 314, 339, 341, 343
 - object reports in a traversal buffer 200
 - objects in a traversal buffer 36

- float kind 202
- floating point number
 - as a byte object 30
 - converting to an object 203
- format of an object
 - fetching 180
- format of an object, fetching 29
- free OOPs
 - run-length encoding 143
- full compression
 - enabling 144

G

- garbage collection 319, 320, 321, 322, 325, 336, 337, 338
 - saving and releasing objects 40
- GciAbort 91
- GciAddOopsToNsc 93
- GciAddOopToNsc 92
- GCI_ALIGN 94
- GciAll7Bit 95
- GciAllocTravBuf 96
- GciAlteredObjs 97
- GciAppendBytes 25, 26, 27, 33, 99
- GciAppendChars 100
- GciAppendOops 101
- GciBegin 102
- GCI_BOOL_TO_OOP 103
- GciByteArrayToPointer 104
- GciCallInProgress 38
- GciCheckAuth 106
- GCI_CHR_TO_OOP 107
- GciClampedTrav 108
- GciClampedTravArgsSType 108, 160
- GciClassMethodForClass 110
- GciClassNamedSize 112
- GciClearStack 113
- GciCommit 115
- GciCompileMethod 116
- GciCompress 118
- GciContinue 120
- GciContinueWith 121
- GciCreateByteObj 122
- GciCreateOopObj 124
- GciCTimeToDateTime 126
- GciDateTimeSType 71
- GciDateTimeToCTime 127
- GciDbgEstablish 128
- GciDbgEstablishToFile 130
- GciDbgLogString 131
- GciDeclareAction 50, 132

GciDecodeOopArray 133
 GciDecSharedCounter 134
 GciDirtyExportedObjs 135
 GciDirtyObjsInit 137
 GciDirtySaveObjs 138
 GciDirtyTrackedObjs 140
 Gci_doubleToSmallDouble 142
 GciEnableFreeOopEncoding 143
 GciEnableFullCompression 144
 GciEnableSignaledErrors 145
 GciEncodeOopArray 146
 GciErr 148
 GciErrSType 72
 GciExecute 149
 GciExecute_ 149
 GciExecuteFromContext 151
 GciExecuteFromContext_ 151
 GciExecuteStr 153
 GciExecuteStr_ 153
 GciExecuteStrFetchBytes 155
 GciExecuteStrFromContext 157
 GciExecuteStrFromContext_ 157
 GciExecuteStrTrav 159
 GciFetchByte 161
 GciFetchBytes_ 162
 GciFetchChars_ 164
 GciFetchClass 165
 GciFetchDateTime 167
 GciFetchDynamicIv 168
 GciFetchDynamicIvs 169
 GciFetchNamedOop 170
 GciFetchNamedOops 172
 GciFetchNamedSize 174
 GciFetchNameOfClass 175
 GciFetchNumEncodedOops 176
 GciFetchNumSharedCounters 177
 GciFetchObjectInfo 178
 GciFetchObjImpl 180
 GciFetchObjInfoArgsSType 178
 GciFetchOop 31, 181
 GciFetchOops 183
 GciFetchPaths 185
 GciFetchSharedCounterValuesNoLock 190
 GciFetchSize_ 191
 GciFetchVaryingOop 195
 GciFetchVaryingOops 197
 GciFetchVaryingSize_ 199
 GciFindObjRep 200
 GciFloatKind 202
 GciFloatKindEType 202
 GciFltToOop 203
 GciGetFreeOop 204
 GciGetFreeOops 206
 GciGetFreeOopsEncoded 208
 GciGetSessionId 209
 GciHardBreak 38, 210
 GciHiddenSetIncludesOop 211
 GCI_I64_IS_SMALL_INT 212
 GciI64ToOop 213
 GciIncSharedCounter 214
 GciInit 38, 215
 GciInstallUserAction 218
 GciInstallUserAction_ 218
 GciInstMethodForClass 219
 GciInUserAction 221
 GciIsKindOf 222
 GciIsKindOfClass 223
 GciIsRemote 224
 GciIsSubclassOf 225
 GciIsSubclassOfClass 226
 GciIvNameToIdx 227
 GciLnk
 configuration file 216
 GciIsRemote 224
 Loading 333, 334, 335
 object traversal function 44
 path access function 44
 use only with debugged applications 66
 user action 56
 GciLoadUserActionLibrary 228
 GciLogin 229, 231
 GciLogout 232
 GciLongJmp 40, 314
 GciMoreTraversal 234
 GciNbAbort 236
 GciNbBegin 237
 GciNbClampedTrav 238
 GciNbCommit 239
 GciNbContinue 240
 GciNbContinueWith 241
 GciNbEnd 242
 GciNbEnd_ 242
 GciNbExecute 245
 GciNbExecute_ 245
 GciNbExecuteStr 246
 GciNbExecuteStr_ 246
 GciNbExecuteStrFromContext 249
 GciNbExecuteStrFromContext_ 249
 GciNbExecuteStrTrav 251
 GciNbMoreTraversal 253
 GciNbPerform 254
 GciNbPerformNoDebug 256
 GciNbPerformNoDebug_ 256
 GciNbPerformTrav 258

GciNbStoreTrav 260
GciNbStoreTravDoTrav 262
GciNbStoreTravDoTravRefs 263
GciNbTraverseObjs 265
GciNewByteObj 266
GciNewCharObj 267
GciNewDateTime 268
GciNewOop 269
GciNewOops 270
GciNewOopUsingObjRep 272
GciNewString 275
GciNewSymbol 276
GciNscIncludesOop 278
GciObjExists 279
GciObjInCollection 280
GciObjInfoSType 72
GciObjIsCommitted 281
GciObjRepHdrSType 75
GciObjRepSize_ 282
GciObjRepSType 74, 160
GciOldOopToNewOop 284
GCI_OOP_IS_BOOL 285
GCI_OOP_IS_SMALL_INT 286
GCI_OOP_IS_SPECIAL 287
GCI_OOP_TO_BOOL 289
GciOopToBool 288
GciOopToChar16 290
GciOopToChar32 291
GCI_OOP_TO_CHR 293
GciOopToChr 292
GciOopToFlt 294
GciOopToI32 296
GciOopToI32_ 296
GciOopToI64 297
GciOopToI64_ 297
GciPerform 298
GciPerform_ 298
GciPerformNoDebug 300
GciPerformNoDebug_ 300
GciPerformSymDbg 302
GciPerformSymDbg_ 302
GciPerformTrav 304
GciPerformTraverse 304, 306
GciPointerToByteArray 308
GciPollForSignal 309
GciPollSocketForRead 310
GciPopErrJump 311
GciProcessDeferredUpdates_ 312
GciProduct 313
GciPushErrJump 314
GciRaiseException 315
GciReadSharedCounter 316
GciReadSharedCounterNoLock 317
GciRealloc 318
GciReleaseAllGlobalOops 319
GciReleaseAllOops 320
GciReleaseAllTrackedOops 321
GciReleaseGlobalOops 322
GciReleaseOops 323
GciReleaseTrackedOops 325
GciRemoveOopFromNsc 326
GciRemoveOopsFromNsc 327
GciReplaceOops 329
GciReplaceVaryingOops 330
GciResolveSymbol 331
GciResolveSymbolObj 332
GciRpc 144
 GciIsRemote 224
 loading 333, 334, 335
 multiple GemStone sessions 44
 object traversal function 43
 path access function 43
GciRtlIsLoaded 333
GciRtlLoad 334
GciRtlUnLoad 335
GciSaveAndTrackObjs 336
GciSaveGlobalObjs 337
GciSaveObjs 338
 in user actions 49
GciServerIsBigEndian 339
GciSessionIsRemote 340
GciSetCacheName_ 341
GciSetDynLib 342
GciSetErrJump 343
GciSetHaltOnError 345
Gci_SETJMP 40, 314
GciSetNet 347
GciSetSessionId 351
GciSetSharedCounter 353
GciSetTraversalBufSwizzling 354
GciSetVaryingSize 355
GciShutdown 356
GciSoftBreak 38, 357
GciStep 359
GciStep_ 359
GciStoreByte 360
GciStoreBytes 25, 26, 27, 361
GciStoreBytesInstanceOf 363
GciStoreChars 25, 365
GciStoreDynamicIv 366
GciStoreIdxOop 367
GciStoreIdxOops 369
GciStoreNamedOop 371
GciStoreNamedOops 373

- GciStoreOop 375
 - GciStoreOops 377
 - GciStorePaths 379
 - GciStoreTrav 384
 - GciStoreTravDo_ 387
 - GciStoreTravDoTrav_ 390
 - GciStoreTravDoTravRefs 391
 - GciStringToInteger 394
 - GciStrKeyValueDictAt 395
 - GciStrKeyValueDictAtObj 396
 - GciStrKeyValueDictAtObjPut 397
 - GciStrKeyValueDictAtPut 398
 - GciSwapBytesUint 399
 - GciSwapBytesUshort 400
 - GciSymDictAt 401
 - GciSymDictAtObj 402
 - GciSymDictAtObjPut 403
 - GciSymDictAtPut 404
 - GciTrackedObjs 211
 - and garbage collection 321, 325
 - GciTrackedObjsFetchAllDirty 405
 - GciTrackedObjsInit 407
 - GciTravBufType 79
 - allocating 96
 - GciTraverseObjs 408
 - GciUncompress 412
 - GciUserActionInit 50
 - GCIUSER_ACTION_INIT_DEF 50
 - GciUserActionShutdown 50, 51
 - GCIUSER_ACTION_SHUTDOWN_DEF 51
 - GciUserActionSType 78
 - GciVersion 415
 - GCSI
 - compiling and linking 420
 - connecting to shared page cache 420
 - data types 421
 - errors 435
 - function library 419
 - sample program
 - explained 420
 - introduced 420
 - shared library 419
 - GcsiAllStatsForMask 423
 - GcsiAttachSharedCache 424
 - GcsiDetachSharedCache 426
 - GcsiFetchMaxProcessesInCache 427
 - GcsiInit 428
 - GcsiResultSType 421
 - GcsiResultSType (structured type) 421
 - GcsiStatsForGemSessionId 429
 - GcsiStatsForGemSessionWithName 430
 - GcsiStatsForPgsvrSessionId 431
 - GcsiStatsForProcessId 432
 - GcsiStatsForShrPcMon 433
 - GcsiStatsForStone 434
 - GemBuilder
 - initializing 215
 - libraries 44
 - library file
 - gcirpc50.* 44
 - libgcklnk.* 44
 - loading 333, 334, 335
 - run-time binding 44, 66
 - starting 215
 - stopping 356
 - GemBuilder errors 233, 345, 346
 - GemRpc, user action 57
 - GemStone C Statistics Interface, see GCSI 419
 - GemStone service name 347
 - GemStone-defined object, making available to applications 23
 - gssstat.cc, sample GCSI program 420
- ## H
- handling errors 311, 314, 339, 341, 343
 - hard break 91, 210, 236
 - defined 27
 - hidden set 211
 - host
 - file access, default directory 39
 - password 347
 - username 347
 - host-specific C definition 70
- ## I
- implementation of an object
 - fetching 29, 180
 - object report 110, 200, 239, 262, 306
 - implementing a user action 48
 - importing objects
 - from GemStone 18, 28
 - improving application performance 34, 36, 185, 234, 253, 260, 261, 265, 319, 320, 321, 322, 325, 379, 384, 387, 390, 408
 - include file (GCI)
 - gci.ht* 70
 - gcioop.ht* 417
 - include file (GemBuilder)
 - flag.ht* 70
 - gcicmn.ht* 70
 - gciterr.ht* 70
 - gcifloat.hf* 69
 - gci.hf* 69

- gci.ht* 22, 70
- gcioc.ht* 70
- gcioop.ht* 23, 24, 70
- gcirtl.hf* 69, 70
- gcisend.hf* 70
- gciua.hf* 69
- gciuser.hf* 70
- version.ht* 70
- incomplete
 - object report 410
- incrementing shared counter 214
- indexable instance variable, fetching the value of an object's 197
- indexable object (Collection)
 - adding OOPs to 375, 377
 - fetching OOPs from 31, 181, 183
- initializing
 - byte object 266
 - character object 267
 - Date/Time object 268
 - objects 269, 270
 - OOPs 269, 270
- initializing GemBuilder 20, 215
- initiating a GemStone session 229, 231
- installing a user action 132, 218
- instance
 - GemStone-defined 417
 - method, compiling 19, 219
 - variable 30
 - enumerating for a class 112, 227
- interrupt
 - GemStone (hard break) 27
 - handling 27, 38
 - issuing 27, 91, 210, 236, 357
 - virtual machine (soft break) 27, 120, 240, 357
- interrupts 105
- J**
- jump buffer, error handling in GemBuilder 40, 311, 314, 339, 341, 343
- K**
- kernel class 269, 270
 - mnemonics 23, 418
- kind
 - of a class 222, 223
- L**
- level traversal 35, 234, 253, 265, 408
- library
 - GemBuilder 44
 - run-time loading 52
 - search 45
 - user action 49
- linkable GemBuilder (GciLnk)
 - GciIsRemote 224
- linking
 - applications 18
 - applications and user actions 56
- loading
 - user action 52
- loading GemBuilder 333, 334, 335
- logging in to GemStone 21, 229, 231, 347
- logging out from GemStone 21, 232
- logical access to objects 19, 26
- login parameter 229, 231
- longjmp, setjmp
 - equivalent functionality 40, 233, 314, 339, 341, 343, 346
- M**
- macros defined 70
- memory
 - reallocating 318
- message
 - GemBuilder function 26
 - sending 26, 254, 298, 300, 302, 304, 306
- method
 - calling C functions from 132, 218
 - compiling 19, 110, 116, 219
- mnemonic
 - GemStone error 39, 69, 70
- modifying
 - objects directly in C 18, 28
 - caution 29
- multiple GemStone sessions 57
 - GciRpc 44
 - switching among 209, 351, 354
- multiple objects
 - defining 272
 - exporting 200, 234, 253, 260, 261, 262, 265, 379, 384, 387, 390, 391, 408
 - importing 185, 200, 234, 253, 260, 261, 262, 265, 384, 387, 390, 391, 408
- N**
- named instance variable
 - fetching 170, 172
 - number of 112, 174, 227
 - pointer object 30

- network 347
 - minimizing traffic 34, 36, 185, 234, 253, 260, 261, 265, 379, 384, 387, 390, 408
 - node 347
 - parameter 229, 231, 347
 - traffic, minimizing 36
 - nil, GemStone special object 23, 24, 417
 - node name, network 347
 - nonblocking functions 36
 - non-sequenceable collection
 - searching 278
 - non-sequenceable collection (NSC) 32
 - adding OOPs to 92, 93
 - fetching OOPs from 181, 183
 - fetching the size 174, 191
 - implementation type 29, 32, 180
 - removing OOPs from 326, 327
 - number of an object's instance variables
 - object report 110, 200, 239, 262, 306
 - number of named instance variables in a class 112
 - number, converting to an object 203
 - numeric representation of a path 185, 379
- O**
- object
 - byte implementation type 30
 - control function 319, 320, 321, 322, 325, 336, 337, 338
 - converting to
 - boolean 288, 289
 - character 290, 291, 292, 293
 - floating-point number 294
 - integer 213, 296, 297
 - creating 33, 272
 - identity 22
 - importing or exporting multiple 34
 - mnemonic 23
 - NSC implementation type 32
 - pointer implementation type 30
 - releasing 40, 232, 319, 320, 321, 322, 325
 - report 110, 200, 234, 239, 253, 260, 261, 262, 265, 272, 306, 384, 387, 408
 - finding in a traversal buffer 200
 - incomplete 410
 - size 282
 - special objects 410
 - structure summary 35
 - traversal buffer 35
 - word alignment 94
 - representation in C 18, 28
 - saving 40
 - sending messages 254, 298, 300, 302, 304, 306
 - object information
 - fetching 178
 - structured type 72, 178
 - object report
 - structured type 74, 160
 - object report header
 - structured type 75
 - object traversal 304
 - objects
 - creating 269, 270
 - initializing 269, 270
 - saving from garbage collection 336, 337, 338
 - objectSecurityPolicyId, in GciObjRepHdrSType 75
 - OOP (object-oriented pointer)
 - adding to an indexable object (Collection) 375, 377
 - adding to an NSC 92, 93
 - defined 22
 - fetching from an indexable object (Collection) 31, 181, 183
 - fetching from an NSC 181, 183, 195, 197
 - removing from an NSC 326, 327
 - searching an NSC for 278
 - searching for 211
 - OOP array
 - decoding 133
 - encoding 146
 - obtaining the size of 176
 - OOPs
 - allocating 204, 206, 208
 - appending 101
 - creating 269, 270
 - initializing 269, 270
 - operating system considerations 38
- P**
- password
 - GemStone 229, 231, 347
 - host 347
 - path access
 - function, GciLnk 44
 - function, GciRpc 43
 - to objects 36, 185, 379
 - pause message 120, 121, 240, 241
 - performance, improving application 34, 36, 185, 234, 253, 260, 261, 265, 319, 320, 321, 322, 325, 379, 384, 387, 390, 408
 - pointer
 - converting to a byte array 308
 - pointer object

- fetching OOPs from 170, 172, 195, 197
- fetching the size 174, 191
- implementation type 29, 30, 180
- storing OOPs in 367, 369, 371, 373, 375, 377
- polling for GemBuilder errors 39, 148
- polling for signal errors 309, 310
- primitive, user-defined 132, 218
- private method, compilation restrictions 110, 116, 219
- ProfMonitor
 - and signals 80
- PureExportSet 211

R

- reallocating
 - memory 318
- reclaiming storage 319, 320, 321, 322, 325
- ReferencedSet 393
- releasing objects 40, 232, 319, 320, 321, 322, 325
- remote procedure call GemBuilder (GciRpc) 224
- removing OOPs from an NSC 326, 327
- report, of an object 110, 200, 234, 239, 253, 260, 261, 262, 265, 272, 282, 306, 384, 387, 408
- re-reading objects from the database 21
- reserved OOP 23
- resolving symbols 27, 110, 116, 149, 153, 157, 159, 219, 245, 246, 249, 251
- Ruby
 - and GCI execution environment 80
- run-length encoding 133, 146, 176
 - enabling 143
- run-time binding
 - GemBuilder 44, 66
- run-time loading 52

S

- saving objects 40
 - export set 41
- schema 18
- security policy
 - GciCheckAuth function 106
- security policy ID in Object Information Structure 73
- security policy, in data structure 75
- sending messages to GemStone objects 19, 26, 254, 298, 300, 302, 304, 306
- service name, GemStone 347
- session
 - control 20, 229, 231, 232, 347
 - creating (logging in) 21, 347
 - current 21

- defined 20
- finding the current ID number 209
- halting on error 345
- setting the current ID number 351, 354
- switching among multiple 209, 351, 354
- terminating (logging out) 21, 232
- setjmp, longjmp
 - equivalent functionality 40, 233, 314, 339, 341, 343, 346
- shared counter
 - decrementing 134
 - fetching value 190, 316, 317
 - finding how many 177
 - incrementing 214
 - setting value 353
- shared libraries
 - GemBuilder 44
 - user action 47
- SIGIO 38
- signal
 - handling 38
- signal** (system function) 38
- signal errors 309, 310
- signal handling
 - UNIX 80
- signaled errors 145
- single-step execution 359
- size of an object
 - fetching 174, 191
 - object report 200
- size of an object report, calculating 282
- size of an OOP array 176
- SmallInteger
 - represented as a special object 23
- soft break 120, 240, 357
 - defined 27
- special object
 - implementation type 29, 180
 - object report 410
 - traversal of 410
- stack
 - clearing the call 113
- starting GemBuilder 20, 215
- statistics, collecting directly from shared page cache 419
- stopping GemBuilder 20, 356
- storing
 - bytes in a byte object 360, 361, 363
 - dynamic instance variables 366
 - objects by using paths 36, 379
 - OOPs in a pointer object 367, 369, 371, 373, 375, 377

- string
 - appending 100
 - as a byte object 30
 - converting to an integer 394
 - fetching 161, 162, 164
 - storing 360, 361, 365
 - structural access 28, 92, 93, 112, 161, 162, 164, 165, 170, 172, 174, 180, 181, 183, 191, 195, 197, 222, 223, 227, 269, 270, 272, 326, 327, 360, 361, 363, 365, 367, 369, 371, 373, 375, 377
 - function 80
 - caution when using 80, 87
 - structured types
 - GciClampedTravArgsSType 108, 160
 - GciDateTimeSType 71
 - GciErrSType 72
 - GciFetchObjInfoArgsSType 178
 - GciObjInfoSType 72
 - GciObjRepHdrSType 75
 - GciObjRepSType 74, 160
 - GciTravBufType 79
 - GciUserActionSType 78
 - subclass
 - determining 225, 226
 - switching among multiple GemStone sessions 209, 351, 354
 - symbol
 - as a byte object 30
 - creating 276
 - resolution 27, 110, 116, 149, 151, 153, 157, 159, 219, 245, 246, 249, 251
- T**
- terminating GemStone sessions 232
 - testing an application
 - use GciRpc 66
 - tracing a GemBuilder call while debugging 128, 130, 131
 - tracked objects 405, 407
 - transaction
 - aborting 21, 91, 236
 - beginning a 102
 - committing 21, 115, 239
 - control 91, 115, 236, 239
 - management 106
 - workspace, creating 229, 231, 347
 - workspace, terminating 232
 - traversal 34, 108, 110, 200, 234, 239, 251, 253, 258, 260, 261, 262, 265, 304, 306, 384, 387, 390, 391, 408
 - buffer 35, 234, 253, 260, 261, 262, 265, 384, 387, 390, 391, 408
 - finding object reports 36, 200
 - word alignment 94
 - function
 - GciLnk 44
 - GciRpc 43
 - inability to continue 234, 410
 - level 35, 234, 253, 265, 408
 - special object 410
 - structured type 108, 160
 - threshold 110, 234, 239, 253, 260, 261, 262, 265, 304, 384, 387, 390, 391, 408
 - word alignment 94
 - traversal buffer
 - allocating 96
 - structured type 79
 - true, GemStone special object 23, 24, 103, 417
- U**
- uncommitted object, releasing 40, 232, 319, 320, 321, 322, 325
 - uncompressing objects 412
 - underscore character, private method 110, 116, 219
 - UNIX signal handling 80
 - unnamed instance variable, fetching 195, 197
 - updating the C representation of database objects 96, 97, 106, 107
 - user action 47–59, 221
 - calling from GemStone 53
 - compiling 63
 - debugging 54
 - defined 47
 - executing 54
 - GciDeclareAction 132
 - GciInstallUserAction 218
 - GciUserActionInit 413
 - GciUserActionShutdown 414
 - implementing 48
 - include file 70
 - installation macro defined 70
 - installing 132, 218
 - kinds of 54
 - library 49
 - loading 228
 - linked application 56
 - loading 52
 - making results persistent 49
 - RPC application 56
 - run-time loading 52
 - structured type 78
 - user action libraries 47

user name

 GemStone 229, 231, 347
 host 347

user profile, searching the symbol list in 110, 116,
 149, 153, 157, 219, 245, 246, 249

user session

 creating 229, 231, 347
 terminating 232

user, searching the symbol list for 27

V

value buffer

 object report 200, 234, 253, 260, 261, 262, 265,
 272, 384, 387, 408

word alignment 94

value of an instance variable, object report 110,
 200, 239, 262, 306

version

 GemBuilder 415

virtual machine

 call stack 40

 clearing 40, 113

 control function 120, 121, 149, 153, 157, 240,
 241, 245, 246, 249, 254, 298, 300, 302,
 304, 306

W

word alignment 94