

More than you ever wanted to  
know about GCC, GAS and ELF.

SHARE 98      Nashville, TN  
Session 8131      March 2002

David Bond  
Tachyon Software LLC  
dbond@tachyonsoft.com

© Tachyon Software LLC, 2002

Permission is granted to SHARE Incorporated to publish this material for SHARE activities and for others to copy, reproduce, or republish this material in a manor consistent with SHARE's By-laws and Canons of Conduct. Tachyon Software retains the right to publish this material elsewhere.

Tachyon Software is a registered trademark and Tachyon z/Assembler is a trademark of Tachyon Software LLC.

IBM, HLASM, OS/390, z/OS, z/Architecture, zSeries and System/390 are trademarks or registered trademarks of the International Business Machines Corporation.

LINUX is a registered trademark of Linus Torvalds.

Other trademarks belong to their respective owners.

# Terminology

---

## OS/390

C/C++ Compiler  
HLASM  
Binder  
TSO TEST, IPCS  
GOFF  
Program Object  
DLL  
Dump Data Set  
SYM records, ADATA  
Language Environment

## Linux

gcc  
as  
ld  
gdb  
ELF relocatable  
ELF executable  
ELF Shared Object  
ELF "core" file  
DWARF  
glibc

- ELF** Executable and Linkable Format  
Format of object, executable and dump files
- DWARF** Debug With Arbitrary Record Format  
Format of debugging records inside ELF files
- GCC** GNU Compiler Collection  
C, C++, JAVA, FORTRAN, ADA and other compilers
- GAS** GNU Assembler  
Assembler included with GNU tools

# Linux Development Packages

---


gcc	GNU Compiler Collection (C, C++, Objective C, FORTRAN, JAVA, ADA)
binutils	assembler, linker
gdb	debugger
glibc	GNU C/C++ library
make	make

Each of these packages are available at different levels. A given Linux distribution will provide a copy of these packages at some level. You can install other levels, as long as you watch for compatibility. For example, a given level of gcc may require a certain minimum binutils level.

For example, binutils 2.10 is required for gcc 3.0, at least for the S/390 version. gcc 3.0 generates some instruction operation codes that were not supported by the GNU assembler prior to binutils 2.10.

# C/C++ Compile Flow

---



gcc is the controller program for the compile process. Unless told to do otherwise, it will invoke the C/C++ preprocessor, C/C++ compiler, the assembler (as) and the linker (ld).

The C/C++ preprocessor reads the source and header files. The header files are provided by glibc, gcc, Linux and the application. The preprocessed C/C++ source code is passed to the C/C++ compiler.

The C/C++ compiler reads the preprocessed source code and generates an intermediate work file containing assembler source code (gas format). gcc can be told (via the -S option) to retain the assembler source file (filename.s) and to not invoke the assembler.

gcc normally invokes the GNU assembler and deletes the temporary assembler source file. The assembler produces an ELF relocatable object file, normally as a temporary work file that is passed to the linker. If gcc is invoked with the -c option, the object file (filename.o) will be retained and the linker will not be invoked.

If run without either -S or -c, gcc then invokes the linker and deletes the temporary object file. The linker combines the object file with object modules from various libraries (mostly glibc) and produces an ELF executable program.

You can run gcc with the -v option to see what programs are invoked.

# Why does GCC produce assembler source?

---

GCC is complex:

- Multiple languages (C, C++, FORTRAN, ADA...)
- Multiple operating systems (Linux, AIX, Solaris, Windows...)
- Multiple architectures (Intel/x86, PowerPC, Sparc, S/390...)
- Multiple object file formats (COFF, ELF...)
- Multiple debugging formats (DBX, DWARF...)

By producing assembler source, many of the details can be left to the assembler. The OS-supplied assembler can be used, or the GNU assembler (GAS) can be built for the target architecture and OS.

GCC can be built as a cross-compiler and GAS can be built as a cross-assembler, allowing compiles on one system which are targeted for another.

GCC must be aware of the target object file format and debugging information format in order to pass the correct information to the assembler. For instance, even though the assembler directives for producing DBX and DWARF debugging information are the same, the information content is different, so it is up to GCC to include the correct information. However, it is up to the assembler to actually produce the output files, so only the assembler has to know the physical file formats.

Since GCC supports multiple programming languages, each language has a unique front-end processor and uses the common back-end code generator. The front-end processor is almost completely independent of the target architecture and the back-end is relatively independent of the source language. Only the back-end of GCC needs to be ported to new architectures, and once the port is done, all of the GCC languages are available.

# How good is GCC?

---

GCC is open-source, so it is as good as the effort put into it.

- Intel/x86 optimization is very good.
- Dave Pitts' I370 port generates HLASM and works with OS/390 USS and LE, however there is no optimizer.
- S/390 and S/390x (z/Architecture) ports are excellent. They were contributed by the same IBM group that writes millicode for the Z900. They have a deep understanding of the S/390 and z/Architecture, so the optimizer is outstanding (and getting better).

Major improvements as of gcc 3.02 – try it!

Dave Pitts' I370 port is available in source and executable form from:  
<http://www.cozx.com/~dpitts/gcc.html>

The GCC home page is:

<http://www.gnu.org/software/gcc/index.html>

The GCC site includes information about recent and upcoming releases and links to sites from which the source can be downloaded. It is relatively easy to download, compile and install a new release of GCC. Patch files can be downloaded to upgrade the source of one release to another. (e.g. From 3.0.2 to 3.0.3)

At this time, GCC 3.03 is current. 3.03 is mostly C++ bug fixes for 3.02.

The next major release is 3.1, scheduled for April 15, 2002, with maintenance releases scheduled for June 15 and August 15.

GCC releases are controlled by the GCC Steering Committee. Resources are largely provided by Cygnus/Red Hat.

The primary maintainer of the Linux for S/390 and zSeries version of GCC is Hartmut Penner of IBM Germany.

## S/390 and z/Architecture Tools

GCC supports both s390 and s390x from the same program. The default is s390 (31-bit). Use the -m64 compiler option to generate s390x (64-bit z/Architecture) code.

binutils and glibc must be built for either s390 or s390x and the correct version must be used.

s390 version requires Relative and Immediate instruction set support. IEEE floating point is assumed, but can be provided by emulation in Linux for S/390.

Linux for S/390 can only run on machines with the Relative and Immediate instruction set support since Linux for S/390 is built using GCC. However, Linux for S/390 can emulate the S/390 advanced floating point instructions in software so it does not require IEEE floating point support.

Since the z/Architecture includes the Relative and Immediate instructions and IEEE floating point support, Linux for zSeries has no special hardware requirements beyond a z800 or z900 machine.

# C Source Program

---

```
>cat hello.c
#include <stdio.h>
int main(
    int argc,
    char * argv[])
{
    puts("Hello world.");
    return 0;
}
>gcc -O3 -S hello.c
```

This is the classic “Hello world” program in C.

The UNIX/Linux “cat” command can be used to print a file to the terminal.

The hello.c program is compiled with the -O3 flag for maximum optimization and the -S flag to cause the compiler to generate the assembler source file and stop. If the compile is successful, the assembler source will be put in a file named hello.s



## GAS Source – page 1 of 2

```
>cat hello.s
    .file      "hello.c"           #data for .note section
    .section .rodata              #start .rodata section
    .align    2                   #halfword alignment
.LC18:                             #internal label
    .string   "Hello world."      #define string constant
    .text                               #start .text section
    .align    4                   #fullword alignment
    .globl   main                 #entry point definition
    .type    main,@function
main:
    stm     %r13,%r15,52(%r15)    #save caller's registers
    bras   %r13,.LTN0_0          #R13=literal base
.LTN0_0:
.LC19:
    .long   .LC18                #address of string constant
```

This is the 31-bit GNU assembler code produced by GCC 3.02 for Linux for S/390. The comments generated by GCC are omitted. The comments to the right of the GNU assembler code were added by hand.

## GAS Source – page 2 of 2

```
.LC20:
    .long    puts                #address of puts function
.LTN0_0:
    lr      %r1,%r15            #stack frame setup
    ahi     %r15,-96
    st      %r1,0(%r15)
    l       %r3,.LC20-.LTN0_0(%r13) #call puts function
    l       %r2,.LC19-.LTN0_0(%r13)
    basr   %r14,%r3
    lhi     %r2,0                #set return value
    l       %r4,152(%r15)        #restore return address
    lm      %r13,%r15,148(%r15)  #restore caller's registers
    br     %r4                  #return to caller
.Lfel:
    .size   main,.Lfel-main
    .ident  "GCC: (GNU) 3.0.2"  #data for .comment section
```

# GAS Syntax

---

GAS is free-form. Continuation is indicated by ending a line with a backslash (\), just like C. Multiple statements are allowed per line, each ended with a semicolon (;).

Comments start with a pound sign/hash mark (#) and continue for the rest of the line. Continuation is allowed.

Labels end with a colon (:). Assembler directives start with a period (.). Everything else is a machine instruction.

Symbols consist of the letters A-Z, a-z, digits 0-9, underscore (\_), dollar sign (\$) and period(.). Symbols are case sensitive.

GAS has no HLASM-like macro facility.

UNIX and LINUX are C-based, so it is natural that the GNU assembler syntax details are very C-like.

GCC allows assembler instructions to be included in C and C++ code using the asm language extension. GCC writes these assembler instructions directly to the GAS file.

The primary use of GAS is as a compiler back-end. Very little of LINUX code is written in assembler, even in the kernel. Because few humans code in GNU assembler, GAS does not need the level of sophistication found in HLASM (powerful macros, DSECTs, USINGs, good diagnostics).

# Registers, Labels, Numbers

## Registers:

`%r0 - %r15`    General registers  
`%f0 - %f15`    Floating point registers  
`.`                Current location counter (like HLASM `*`)

## Labels:

`.Lxxxx:`        Internal label – not visible in debugger  
`0: - 9:`        temporary labels, referenced by (for example)  
                  **0B** (backward reference) or **1F** (forward reference).

If not defined, symbols are assumed to be external references.

## Numbers:

C-like syntax, if starts with `0x`, remainder of number is hexadecimal, if starts with `0b`, remainder of number is binary, otherwise if starts with `0`, number is octal, otherwise decimal.

Temporary labels are very nice features of the GNU assembler language. They allow local labels to be defined without the need to resort to HLASM's `&SYNDX` variable symbol. For instance, a reference to **2F** is resolved to the next (forward) definition of the temporary label **2:** and a reference to **9B** is resolved to the nearest previous (backward) definition of the temporary label **9:**.

As an example, the following code is generated by GCC for the `strlen()` function:

```
sr    0,0
lr    %r12,%r11
0: srst 0,%r12
jo    0b
lr    %r12,0
sr    %r12,%r11
```

This same code sequence could be generated later without any conflict between the definitions of **0:**.

# Machine Instructions

---

GAS machine instructions are similar to HLASM format, except:

In RX instructions, the base and index registers are reversed.  
If only one register is specified, it is assumed to be a base register.

Additional branch mnemonics (GAS: HLASM)

JHE: BRC 10	JLE: BRC 12	JLH: BRC 6
JNHE: BRC 5	JNLE: BRC 3	JNLH: BRC 9

Same for BC and BCR mnemonics.

Different BRCL mnemonics (GAS: HLASM)

JG: JLU	JGE: JLE	JGH: JLH
JGHE: BRCL 10	JGL: JLL	JGLE: BRCL 12
JGLH: BRCL 6	JGM: JLM	JGNE: JLNE
JGNH: JLNH	JGNHE: BRCL 5	JGNL: JLNL

... etc.

GCC/GAS introduced extended mnemonics for the 6 condition code masks missing from the HLASM extended mnemonics. Unfortunately, the GCC/GAS JLE and JLH extended mnemonics for BRC conflict with the HLASM extended mnemonics for BRCL. GCC/GAS replaced the JLxxx (Jump Long) extended mnemonics for BRCL with the JGxxx (Jump Grande) mnemonics.

See <http://www.tachyonsoft.com/txac.htm> for instruction tables that include the GCC/GAS extended mnemonics. These tables include instructions from all of the Principles of Operation books as well instructions published in other IBM manuals or discovered from other sources.

# Storage Definition Directives

<u>GAS</u>	<u>HLASM</u>
.byte	DC X'xx'
.short	DC HL2'nn'
.long	DC FL4'nn' or DC AL4(symbol)
.quad	DC FDL8'nn' or DC ADL8(symbol)
.single	DC EBL4'nn'
.double	DC DBL8'nn'
.comm	COM, DS
.ascii	DC C'cccc'
.string	DC C'cccc',X'0'

GAS does not align constants based on type.  
GCC does not produce .single and .double – IEEE value is generated by .long 0XXXXXXXXX constants.  
.ascii and .string use C-like syntax for characters, including escape values. (e.g. "\n\0")

GAS does not align anything. Even machine instructions are not halfword aligned! However the .align directive can align to any power of 2.

.ascii and .string operands are strings enclosed in double quotes. The .string directive includes a terminating X'00' byte in the string. .string "hello" is the same as .ascii "hello\0".

.ascii and .string operands can include “escape” values:

\a = X'07' (BEL)    \b = X'08' (BS)    \f = X'0C' (FF)  
\n = X'0A' (LF)    \r = X'0D' (CR)    \t = X'09' (TAB)  
\v = X'0B' (VT)    \" = double quote    \\ = back slash  
\x<hex digits> = X'<hex digits>'    \<octal digits>

.byte, .short, .long and .quad operands can be expressions. Numbers are decimal unless they start with a zero. Zero is the “escape” character, allowing octal, hexadecimal (0x) or binary values (0b). Character values can also be included (e.g. 'a') and support the same set of escape values as can be specified in .ascii and .string operands.

## Miscellaneous Directives

<u>GAS</u>	<u>HLASM</u>
.align	CNOP
.file	none
.globl	ENTRY
.ident	none
.local	none
.org	ORG
.set	EQU
.size	none
.stab	ADATA
.type	XATTR
.version	none
.weak	WXTRN

Notes: .org is forward only. .align fills with X'07' bytes by default.

Unlike HLASM's EQU, .set can be used more than once for the same symbol, allowing a symbol to have a different value for different parts of the assembly. .set can also be used to change the value of the location counter in a forward direction: `.set ., .+4`

.align fills with X'07' bytes generating “NOPR 7” instructions in any halfwords.

The inability of the GNU assembler to support ORG to a previous location vastly simplifies the assembler logic!

# Section Definitions

---

<u>GAS</u>	<u>HLASM</u>
.text	RSECT
.data	CSECT
.bss	COM
.section	RSECT/CSECT/LOCTR

GAS ELF sections are similar to HLASM GOFF classes.

There is no GAS equivalent to DXD or DSECT.

.text is read-only executable code and literals  
.data is modifiable initialized storage  
.bss is modifiable uninitialized storage

Other sections may be defined, e.g. .rodata is for read-only data.

.text, .data, .bss and .rodata are the only sections generated by GCC for normal C code.

For C++ code, GCC generates many additional sections that are recognized by the GNU linker for special processing. These special sections are to handle static constructors and destructors and other C++ features that must be processed at link time.



## Linux/390 Register Usage

R0-R1	Not saved
R2-R3	Not saved, parameters and return values
R4-R5	Not saved, parameters
R6	Saved, parameter
R7-R12	Saved
R13	Saved, often literal pool base register
R14	Not saved, return address
R15	Saved, stack pointer
F0	Not saved, parameter and return value
F2	Not saved, parameter
F4,F6	Saved, z/Architecture parameters
F1,F3,F5,F7-F15	Not saved
Access Registers	Not saved

### **Function return values:**

<u>Type</u>	<u>Linux for S/390</u>	<u>Linux for zSeries</u>
char, short, int, long, * or 1, 2 and 4-byte structures	R2	R2
long long and 8-byte structures	R2 and R3	R2
float, double	F0	F0

### **Function parameter values:**

First 5 integer (char, short, int, long, long long) and pointer parameters are passed in registers R2, R3, R4, R5 and R6. For Linux for S/390, long long parameters are passed in register pairs. Structures of 1, 2, 4 or 8 bytes are passed as integers.

In Linux for S/390, first 2 floating point parameters are passed in F0 and F2. In Linux for zSeries, first 4 floating point parameters are passed in F0, F2, F4 and F6.

All other parameters are passed on the stack. If the return value is not an integer, pointer, float, double or 1, 2, 4 or 8-byte structure, a “hidden” parameter in R2 will contain the address of the return area.

## Linux/390 Stack Frame

<u>S/390</u>	<u>zSeries</u>	<u>Purpose</u>
0-3	0-8	back chain
4-15	8-31	reserved
16-23	32-47	scratch area
24-63	48-127	saved r6-r15 of caller function
64-79	128-143	saved f4 and f6 of caller function
80-95	144-159	undefined
96	160	outgoing args passed from caller to callee
96+x	160+x	possible stack alignment to 8-bytes
96+x+y	160+x+y	alloca space of caller ( if used )
96+x+y+z	160+x+y+z	automatics of caller ( if used )

Only the registers that are modified by a function need to be saved.  
The stack grows toward lower addresses.

Stack frames are always double-word aligned.

The stack of the process's initial thread starts at the high end of virtual storage and grows down. For Linux for S/390, the high end is X'7FFFFFFF' (2G-1).

In Linux for zSeries 2.4, 42 bits of the 64-bit possible virtual storage range are used, so the highest address is X'000003FF FFFFFFFF' (4T-1).

Linux threads are like OS/390 tasks. Each thread has its own stack.

## GCC Output – Notes

---

<pre>.section .rodata .align 2 .LC18: .string "Hello world."</pre>	A half-word aligned, null-terminated ASCII string constant is defined in the .rodata section.
<pre>.text .align 4 .globl main .type main,@function main:     stm %r13,%r15,52(%r15)</pre>	An externally visible function named “main” is defined in the .text section. It is fullword aligned. The function saves only r13, r14 and r15 since those are the only non-volatile registers modified.

GAS uses the strongest `.align` directive within a section to set the alignment for the section. When combining sections, the linker uses the strongest alignment contributed by any object module for the section alignment within the program.

GCC generates the minimum STM instruction required to save registers that must be preserved when control is returned to the caller. When generating code, GCC first tries to use the volatile registers (0-5), which are the registers that a caller does not expect to be preserved across a call. If all of the volatile registers are used, GCC uses the highest numbered register available. This allows GCC to generate minimal STM/LM instructions in the function prolog and epilog.

## GCC Output – Notes

---

<pre>    bras %r13, .LTN0_0 .LT0_0: .LC19:     .long    .LC18 .LC20:     .long    puts .LTN0_0:</pre>	<p>Register 13 is set up as the base register for the literal pool. Two address constants are in the literal pool: the address of the string constant in <code>.rodata</code> and the address of the external function named “puts”.</p>
<pre>    lr    %r1, %r15     ahi   %r15, -96     st    %r1, 0(%r15)</pre>	<p>A new stack frame of 96 bytes is allocated and the address of the caller's stack frame is saved (back chain).</p>

If a function does not call another function, it is called a “leaf” function. Leaf functions that do not modify any of the caller's non-volatile registers do not need to save and restore any registers, so no stack frame is needed.

The GCC “literal pool” is actually a set of constants generated near the start of each function and usually addressed via R13 which is set up via the BRAS instruction. The label `LT0_0` for this function is used as the base address of the literal pool when literals are referenced.

If local variables need to be allocated in the function's stack frame, more than 96 bytes would be needed and the AHI instruction would reflect this. The stack frame size is always a multiple of 8 to ensure that stack frames are doubleword aligned.

## GCC Output – Notes

---

<pre>l    %r3, .LC20-.LT0_0(%r13) l    %r2, .LC19-.LT0_0(%r13) basr %r14,%r3</pre>	<p>Register 3 is loaded with the address of the “puts” function; register 2 is loaded with the first parameter; “puts” function called.</p>
<pre>lhi  %r2,0 l    %r4,152(%r15) lm   %r13,%r15,148(%r15) br   %r4</pre>	<p>The return value is loaded into register 2; the return address is loaded into register 4; registers 13 and 15 are restored; control is returned to the caller.</p>

GCC generates explicit displacement values for literal references by subtracting the base address of the literal pool (LT0\_0 in this case) from the address of the literal.

Unlike OS/360 standard linkage conventions, Linux/390 does not require that the called routine's address be in any specific register. Called routines cannot expect that the starting address is in a register.

In the call to the “puts” function, the load for the address of “puts” is performed as far in advance of the BASR as possible to reduce pipeline stalls. The load instructions for R2 and R3 can run in parallel.

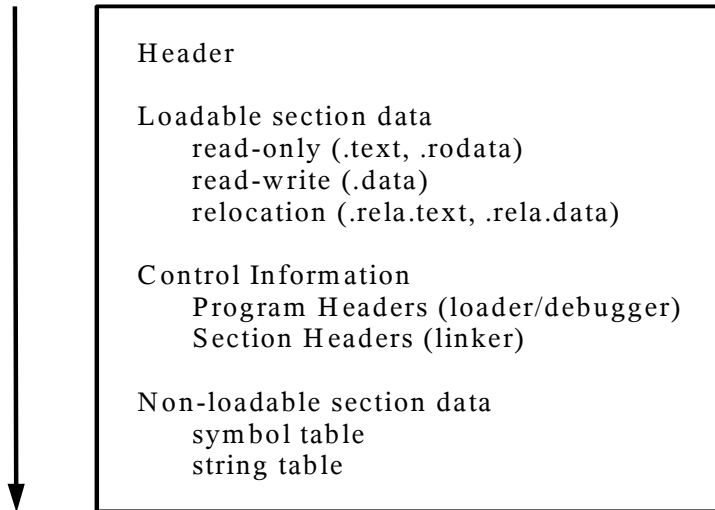
In the function epilog, the three instruction sequence (L,LM,BR) is used instead of the minimal two instruction sequence (LM 13,15,148(15);BR 14) to allow the maximum amount of parallelism. In the three instruction sequence, the LM instruction can execute in parallel with the BR and subsequent instructions.

*Session 8158 (8:00AM Thursday) discusses these performance issues.*

Because GCC understands how to exploit instruction parallelism, it generates code that is unusual (and faster) than most human-coded assembler. Human-generated assembler should be maintainable, so it is not good practice to separate related instructions in non-obvious ways.

# ELF File Contents

---



ELF files, like UNIX files in general, are byte oriented, not record oriented.

ELF is based on structures. A program reading an ELF file can navigate through the structures via offset values within the structures. All offsets are in bytes from the start of the file – useful for `fseek/lseek` C library functions.

The ELF header contains offsets, sizes and counts of Program Headers and Section Headers. The Program Headers and Section Headers contain offsets, sizes and counts of everything else.

Program Headers are useful for the loader and debugger. They contain the information needed to describe a program in memory. Executable programs, Shared Objects and “core” files (but not relocatable object files) contain Program Headers.

Section Headers are used by the linker to find the various parts and symbols to be linked together. Relocatable object files, executable programs and Shared Objects (but not “core” files) contain Section Headers. Like OS/390 Load Modules and Program Objects, ELF executable programs and Shared Objects can be relinked.

GAS and the linker both write section data in roughly the same order.

See <http://www.tachyonsoft.com/elf.com> for links to information about ELF.

# ELF Header Contents

---

**Fixed Part:**

- File Type (Relocatable, Executable, Shared Object, Core)
- Word Size (32 or 64 bit)
- Endian (MSB: S/390,PPC,Sparc; LSB: Intel)

**Variable Part:** (depending on word size)

- Architecture (Intel, S/390, PPC ...)
- Offset to, sizes and number of Program/Section Headers
- Offsets, virtual storage addresses and most length fields are the width of and aligned to the word size.

With the information in the header, a program running on one architecture can read an ELF file for a different architecture. All of the information needed to decode an ELF file is in the header.

The only difference between s390 (31-bit) and s390x (64-bit) is the word size indicator in the header and the consequent field size/location changes. Linux for zSeries will load and execute both 31-bit and 64-bit programs.

Before the S/390 machine type was formalized, code X'A390' was used. The official S/390 machine type code is now X'0016'. Many (all?) of the current binutils programs recognize both. Older versions may not recognize X'0016'.

# Section Header Contents

---

One entry per section, containing the name, type, attributes, file offset, size.

## Typical Sections:

<code>.text</code>	executable code
<code>.rodata</code>	constant data
<code>.data</code>	initialized non-constant data (writable static)
<code>.bss</code>	uninitialized data (requires no room in ELF file)
<code>.rela.text</code>	relocation information for the <code>.text</code> section
<code>.rela.data</code>	relocation information for the <code>.data</code> section
<code>.stab</code>	debugging information (DWARF)
<code>.symtab</code>	symbols for linking (and debugging)
<code>.strtab</code>	strings: symbol names
<code>.note</code>	miscellaneous loadable information
<code>.comment</code>	miscellaneous non-loadable information

The linker simply concatenates most sections with like names from different modules.

`.rela` sections can exist for any initialized section. The name of the section containing the addresses to be relocated is suffixed to `.rela` (e.g. `.rela.text`) and the index number of the containing section is in the `.rela` section header entry.

The `.symtab` section is created from the merged symbol table, combining the external references with the global symbols.

Other special sections can be created, especially for C++. These sections handle virtual method tables and static constructors and destructors. The linker performs special processing for many of these special sections.



# Program Header Contents

One entry per loadable area, containing the attributes, file offset, virtual storage location, virtual storage size and initialized size.

Each loadable area contains all of the sections with the same attributes (read-only, read/write, ...)

For the modifiable data area, the virtual storage size is probably more than the initialized size – the difference is the total uninitialized data area size (.bss section).

Virtual storage areas are page aligned. Executable programs do not need to be relocated since they are linked at the correct load address. Executable programs can be paged-in as needed.

Shared Objects must be relocated since all are linked to the same address.

In an executable program or Shared Object, the Program Headers provide a different view of the same data provided by the Section Headers. The information in the program headers is organized to be easy to load, whereas the Section Headers provide all of the details required to link a program.

In Linux, a process (address space) consists of exactly one program and any number of Shared Objects. In Linux for S/390, programs are always loaded at X'00400000'. In Linux for zSeries, programs are always loaded at X'00000000 80000000'.

In executable programs referencing Shared Objects and in Shared Objects, one of the Program Header entries points to the information needed to resolve cross-module references.

Another Program Header entry specifies the name of the “interpreter” for the program. This is the program actually given control after the program file is loaded. For executable programs, the “interpreter” is the name of the dynamic linker, which will load the referenced Shared Objects, resolve the references and give control to the loaded program. Potentially, ELF files could contain the name of a different “interpreter” program.

# Core File Contents

---

Core files are produced when a process is terminated by certain signals – e.g. SIGSEGV (S0C4).

A core file contains one Program Header entry for each area of virtual storage dumped. Each entry contains the file offset, size, virtual storage address and attributes of the storage area.

One special Program Header is for a “notes” area, which contains process status information, PSW, registers, etc. The “notes” area format is specific to both the Linux release and hardware architecture.

core files can be read by the gdb program.

If you do not see core files being created, issue the **ulimit -a** command. It will probably show that the core file limit is 0. This can be changed by issuing the **ulimit -c unlimited** command.

The elfdump program can format the contents of S/390 and Intel/x86 ELF files, including core files, object files, executable programs and shared objects. The elfdump program is available via <http://www.tachyonsoft.com/elf.html>

# GCC for z/OS?

Why can't GCC be used for z/OS?

- GCC assumes ASCII.
- GCC needs GLIBC. GLIBC is for UNIX/Linux and ASCII.
- GCC generates GNU assembler code. GAS generates ELF.
- GCC debugging information is DWARF, which is not (yet) supported by any z/OS debuggers.
- z/OS Program Objects must be created by the binder. The binder does not read ELF.
- C++ ELF object modules must be processed by the GNU linker. The GNU linker reads and writes ELF. The z/OS loader cannot load ELF load modules.

One solution: Dave Pitts altered GCC to accept EBCDIC, produce HLASM and interface with LE.

Another solution would be to provide an ELF loader for z/OS. This would allow all of the GNU tools to be used: GCC, GAS, GNU linker and most of GLIBC. The ASCII problem would still need to be solved.

Our solution: assemble the GNU assembler code produced by GCC directly to GOFF and provide a replacement for GLIBC, called LIB390.

# GCC for z/OS!

## **GAS/ELF Problem**

- GAS syntax assembler can be converted to HLASM syntax for GOFF. The Tachyon z/Assembler can automatically assemble GAS source and create GOFF.
- GOFF can be linked by the z/OS binder to create normal z/OS load modules.

## **ASCII Problem**

- For now, use patches from Dave Pitts' I370 port or use the ASCII<->EBCDIC translation support in LIB390.
- Integrated GCC support for EBCDIC from Cygnus/Red Hat is in beta test and will be released soon.

The Tachyon z/Assembler can be used for free when used to assemble the output of the GCC compiler. The free version can generate either GOFF (31-bit) or ELF (31-bit or 64-bit) object files from GNU assembler source. The free version of the Tachyon z/Assembler can be downloaded from the Tachyon Software web site at <http://www.tachyonsoft.com>

By assembling GAS to GOFF, no changes to GCC are required. This takes advantage of all of the work done by the GCC maintainers, including the excellent zSeries instruction optimizer.

In 64-bit code generated by GCC, the operands of the BRASL, BRCL and LARL instructions are often external references which can be resolved by the GNU linker. Unlike ELF, GOFF does not have defined relocation types for the operands of these instructions, so there is no way to assemble the 64-bit code generated by GCC into a GOFF object file. Until IBM adds support in GOFF and the binder, this problem can be fixed using a prelinker.

# GCC for z/OS!

## **GLIBC Problem**

- Tachyon Software has started an open-source version of GLIBC for 31-bit OS/390 and z/OS, called LIB390.
- LIB390 is LGPL code, so it cannot be statically linked with non-open-source products. A goal of the project is to make LIB390 into a “DLL” so that GCC and LIB390 can be used in commercial products.

## **C++ Problem**

- C++ support will require a prelinker to perform the “magic” currently performed by the GNU linker for C++.

The source and object files of LIB390 can be downloaded from the Tachyon Software web site at <http://www.tachyonsoft.com/lib390.html>

LIB390 is based on GLIBC with changes and replacement routines as required for OS/390 and z/OS. Since GLIBC is distributed under the GNU Library General Public License (LGPL), LIB390 is also distributed under the same license. One provision of the LGPL is that any user of a program linked with GLIBC must be allowed to replace the GLIBC routines, usually by relinking. This usually requires the program to be distributed in object module form so the user can relink it, or else the program should be dynamically linked to GLIBC. It is intended that a future version of LIB390 can be dynamically linked with programs, allowing commercial products to be built with LIB390 while complying with the LGPL.

# GCC for z/OS!

## **Debugger Problem**

A DWARF-based debugger is needed. Anyone want to contribute?

## **GCC and GAS do not run on z/OS**

- For now, GCC can be run on Linux/390 or as a cross-compiler on Windows or Linux/x86. Linux/390 can be run on Windows or Linux/x86 under Hercules. The Tachyon z/Assembler runs on Linux/390, Linux/x86, Windows, AIX and Solaris.
- A project goal is to allow GCC to build itself to run on z/OS.

Using cross-platform development tools like GCC and the Tachyon z/Assembler, you can build programs on one platform to be executed on another. For instance, you can build the object files for a program on Linux and then upload them to z/OS where they can be bound into Program Objects or Load Modules and executed.

With the limited TSO access on z/OS.e, cross-platform development is probably the preferred method. The only other choice would seem to be a telnet session into z/OS Unix System Services.

Since IBM FORTRAN and COBOL programs cannot be run under z/OS.e, GCC FORTRAN (g77) and GNU COBOL could be used to build programs for z/OS.e since LE would not be needed. The GNU COBOL home page is: <http://www.gnu.org/software/cobol/cobol.html>

Hercules is an open-source System/370, System/390 and z/Architecture emulator. Using Hercules, you can have the fun of installing and running Linux for S/390 and zSeries on your PC, Macintosh or whatever!

You can download Hercules from: <http://www.conmicro.cx/hercules>

*Hercules will be discussed in sessions 2881 (6:00PM Monday), 2880 (6:00PM Tuesday) and 2861 (4:30PM Wednesday).*

# How Can You Help?

---

Cygnus/Red Hat is providing EBCDIC support in GCC.

Tachyon Software is contributing a free version of the Tachyon z/Assembler for GCC and the start of the runtime library support.

**Your help is needed:**

- Runtime library
- C++, COBOL and FORTRAN support
- Debugger and Profiler
- VSE and CMS support

The goal is to have a common set of open-source compilers and other development tools across all IBM mainframe operating systems: z/OS, OS/390, z/VM, VSE and Linux for S/390 and zSeries.

There is already enough of the infrastructure in place to build trivial USS C programs. The near-term milestone is to be able to build real-world 31-bit z/OS C programs by next SHARE.

This is not a toy. C/C++ for z/OS and GCC are two great mainframe compilers from IBM.

# How to Get Started?

---

ELF and DWARF for S/390 Links:

<http://www.tachyonsoft.com/elf.html>

GCC:

<http://www.gnu.org/software/gcc/index.html>

COBOL for GCC:

<http://www.gnu.org/software/cobol/cobol.html>

LIB390:

<http://www.tachyonsoft.com/lib390.html>

Tachyon z/Assembler for GCC:

<http://www.tachyonsoft.com>

Or write to: [dbond@tachyonsoft.com](mailto:dbond@tachyonsoft.com)