

# Compiling and Linking

Praktikum „C-Programmierung“



---

Eugen Betke, Nathanael Hübbe, Michael Kuhn, Jakob Lüttgau, Jannek Squar

2019-12-16

Wissenschaftliches Rechnen  
Fachbereich Informatik  
Universität Hamburg

# GCC Compiler Toolchain: Eine Übersicht

binutils

GNU Compiler Collection

Linux Kernel Headers

Application Binary Interface (ABI)

C Library

Compilation Process mit GCC

Compiler Flags

Empfohlene Flags

Beispiel: `-fstack-protector`

Beispiel: `-D_FORTIFY_SOURCE`

- binutils
- GCC: GNU Compiler Collection
  - C Library
  - Runtime
- Linux Kernel Headers

## Alternative Toolchains:

- LLVM/Clang <sup>a</sup>
- Intel (C/C++, Fortran)
- Cray (C/C++, Fortran)
- IBM (C/C++)
- PGI
- NVIDIA (CUDA LLVM)
- AMD (AOCC: LLVM based)
- ARM (GCC or LLVM based)
- ...

---

<sup>a</sup> <https://clang.llvm.org/comparison.html>

## Sammlung von *binary tools*<sup>1</sup>

**as** Assembler Create object files

**ld** Linker Combine object files/libs

**addr2line** Convert addresses into filenames/line numbers.

**ar** Create, modify and extract from archives.

**c++filt** Filter to demangle encoded C++ symbols.

**dlltool** Creates files for building and using DLLs.

**gold** A new, faster, ELF only linker, still in beta test.

**gprof** Displays profiling information.

**nlmconv** Converts object code into an NLM.

**nm** Lists symbols from object files.

**objcopy** Copies and translates object files.

**objdump** Displays information from object files.

**ranlib** Generates index to contents of an archive.

**readelf** Show information for ELF format object file.

**size** Lists section sizes of an object/archive file.

**strings** Lists printable strings from files.

**strip** Discards symbols.

<sup>1</sup><https://www.gnu.org/software/binutils/>

Frontends:

- C, C++, Objective-C, Fortran, Ada, Go, and D, (sogar Java bis GCC7/2016)

Backends:

- (70+) Architekturen/Plattformen

Bestandteile von GCC<sup>2</sup>:

**cc1,cc1plus** Compiler, erzeugt Assembly Code

**gcc,g++** Compiler-Interface, Integration von binutils

**Header-Files** Deklarationen der Standard C-Library

**libgcc, listdc++, libfortran** Runtime-Libs

---

<sup>2</sup><https://www.gnu.org/software/gcc/>

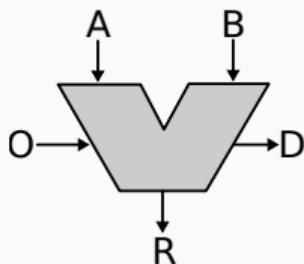
- Abstrahiert OS-Funktionalität über sog. System Calls
- Definition der Userspace-API in etwa 700 Header-Dateien (Linux 4.8)
- Nahtlose Integration in GCC Projekt
- Die Header für die Runtime (aber auch Dritt-Bibliotheken) befinden sich bei den meisten Distributionen unter /usr/include/

Das Application Binary Interface (ABI) definiert wie auf Datenstrukturen und Routinen in der Maschinenrepresentationen zugegriffen werden kann:

- Register Dateistruktur, Stack Organisation, Speicherlayout
- Größen, Aufbau und Alignments von Basistypen
- Aufruf-Konventionen: wie Argumente und Rückgabewerte übergeben werden
- Wie Systemaufrufe auszuführen sind
- Die Struktur der Object-Dateien

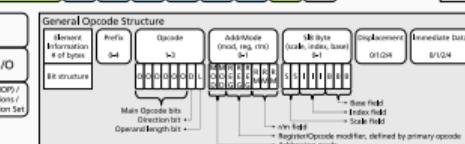
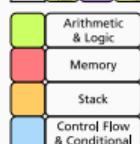
Die Linux ABI ist weitestgehend rückwärtskompatibel:

- Daher, ältere Linux-Header i.d.R. weiterhin benutzbar:  
z.B. ein 3.4 Header funktioniert mit einem 4.5 Kernel
- Neue Header (insbesondere mit neuen Features) mit einem alten Kernel zu benutzen führt meistens zu Problemen



## x86 Opcode Structure and Instruction Overview

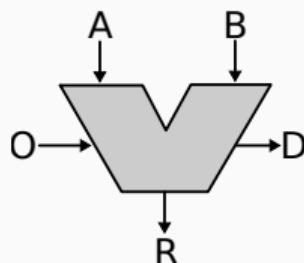
	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	ADD		PUSH	ES	ES		OR				TWO BYTE					
1	ADC		SS	SS			SBB				DS	PUSH	POP	DS		
2	AND		ES	DAA			SUB				CS	DAS				
3	XOR		SS	AAA			CMP				DS	DAS	DS	AAS		
4	INC						DEC									
5	PUSH						POP									
6	PUSHAD	POPAD	BOUND	ARPL	FS	GS	DECODED	ALTERNATE	PUSH	IMUL	PUSH	IMUL	INS	OUTS		
7	JO	JNO	JB	JNB	JE	JNE	JBE	JA	JS	JNS	JPE	JPO	JL	JGE	JLE	JG
8	ADD/ADC/AND/XOR OR/SBB/SUB/CMP				TEST	XCHG		MOV REG		MOV SREG	LEA	MOV SREG	POP			
9	NOP		XCHG EAX				CWD	CDO	CALL/FWAIT	PUSHFD	POPD	SAHF	LAHF			
A	MOV EAX		MOVS	CMPS		TEST	STOS	LODS	SCAS							
B			MOV													
C	SHIFT IMM	RETN	LES	LDS	MOV IMM	ENTER	LEAVE	RETF	INT3	INT IMM	INTO	IRETD				
D	SHIFT 1	SHIFT CL	AAM	AAD	SALC	XLAT										
E	LOOPNE	LOOP	LOOP	JECKZ	IN IMM	OUT IMM	CALL	JMP	JMPF	JMP SHORT	IN DX	OUT DX				
F	LOCK	ICE	REPE	REPZ	REPE	TEST	NOTZERO	REPNE	REPZ	CLC	STC	CLI	STI	CLD	STD	INC DEC



v1.0 – 30.08.2011  
Contact: Daniel Pöhmann – +49 228 73 54 228 – daniel.pöhmann@fkj.fraunhofer.de

	0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	3-SLD	LSLQD	LSLROT	LSLSHW	LAR	LSL		CLTS	INVD	WBINV	UD2	NOP				
1															HINT NOP	
2															SSE{1,2}	
3	WMMX	RDSC	RDPMC	RDPMC	SYSENTER	SYSENTR										
4															CMOV	
5															SSE{1,2}	
6															MMX, SSE2	
7															MMX, SSE{1,2}, VMX	MMX, SSE{2,3}
8	JO	JNO	JB	JNB	JE	JNE	JBE	JA	JS	JNS	JPE	JPO	JL	JGE	JLE	JG
9	SETO	SETNO	SETB	SETNB	SETE	SETNE	SETBE	SETA	SETS	SETNS	SETPE	SETPO	SETL	SETGE	SETLE	SETG
A	PUSH	POP	FS	FS	CPUID	BT	SHLD			PUSH	POP	GS	RSM	BTS	SHRD	FENCE
B	CMPXCHG	LSS	BTR	LFS	LGS	MOVZX	POPNOT	UD	BT	BT	BSF	BSR	MOVSX			
C	XADD														BSWAP	
D															MMX, SSE{1,2,3}	
E															MMX, SSE{1,2}	
F															MMX, SSE{1,2,3}	

Source: Intel x86 Instruction Set Reference  
Opcodes table presentation inspired by work of Ange Albertini

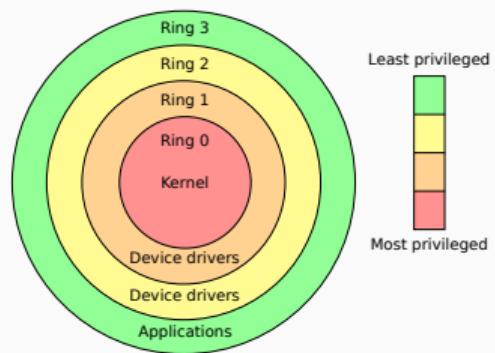
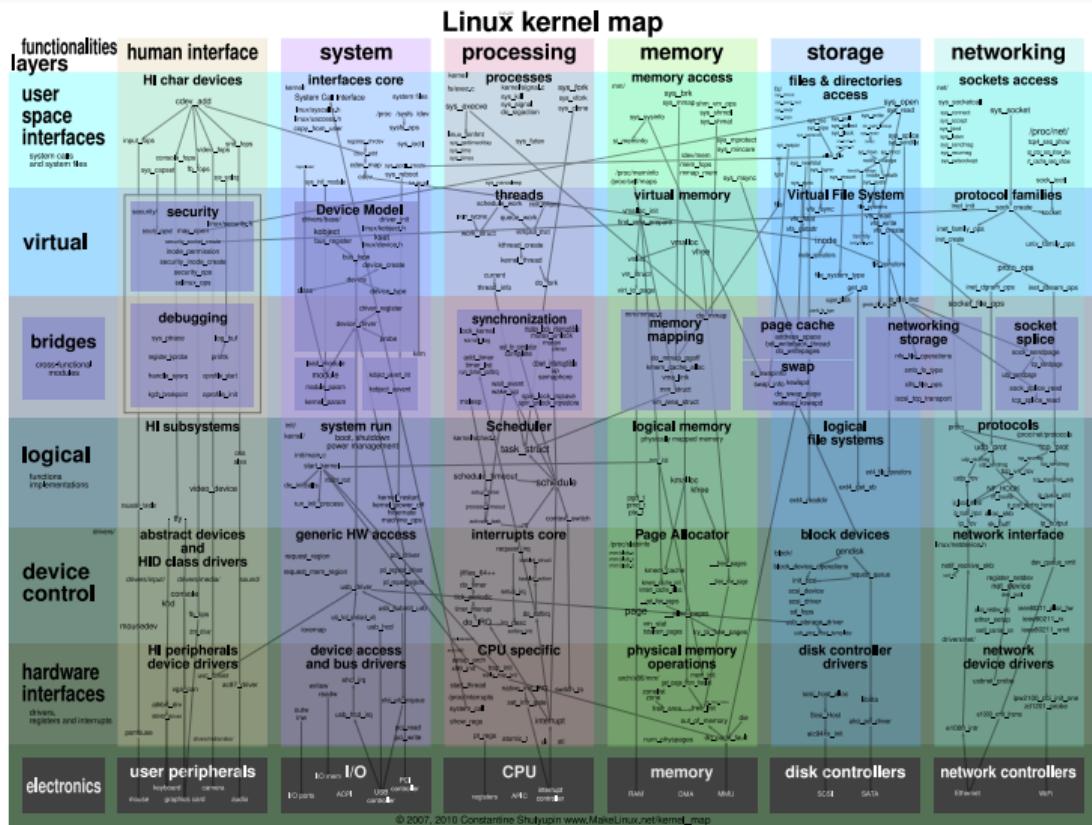


ZMM0	YMM0	XMM0	ZMM1	YMM1	XMM1	ST(0) MM0	ST(1) MM1	RAX	R8	R12	CR0	CR4							
ZMM2	YMM2	XMM2	ZMM3	YMM3	XMM3	ST(2) MM2	ST(3) MM3	RBX	R9	R13	CR1	CR5							
ZMM4	YMM4	XMM4	ZMM5	YMM5	XMM5	ST(4) MM4	ST(5) MM5	RCX	R10	R14	CR2	CR6							
ZMM6	YMM6	XMM6	ZMM7	YMM7	XMM7	ST(6) MM6	ST(7) MM7	RDX	R11	R15	CR3	CR7							
ZMM8	YMM8	XMM8	ZMM9	YMM9	XMM9	BPL	BPH	RBP	DI	EDI	RDI	IP	EIP	RIP					
ZMM10	YMM10	XMM10	ZMM11	YMM11	XMM11	CW	FP_IP	FP_DP	FP_CS	SIL	SI	ESI	RSI	SP	ESP	RSP	MSW	CR9	
ZMM12	YMM12	XMM12	ZMM13	YMM13	XMM13	SW	■ 8-bit register ■ 16-bit register											CR10	
ZMM14	YMM14	XMM14	ZMM15	YMM15	XMM15	TW	■ 32-bit register ■ 64-bit register											CR11	
ZMM16	ZMM17	ZMM18	ZMM19	ZMM20	ZMM21	ZMM22	ZMM23	FP_DS	CS	SS	DS	GDTR	IDTR	DR0	DR6	CR12			
ZMM24	ZMM25	ZMM26	ZMM27	ZMM28	ZMM29	ZMM30	ZMM31	FP_OPC	FP_DP	FP_IP	ES	FS	GS	TR	LDTR	DR1	DR7	CR13	
											FLAGS	EFFLAGS	RFLAGS	DR2	DR8	CR14			
														DR3	DR9	DR4	DR10	DR12	DR14
														DR5	DR11	DR13	DR15	MXCSR	

Details: <https://en.wikipedia.org/wiki/X86#32-bit>

# Betriebssystem Perspektive

## Application Binary Interface (ABI)



Details/Grafik:

[https://en.wikipedia.org/wiki/Protected\\_mode](https://en.wikipedia.org/wiki/Protected_mode)

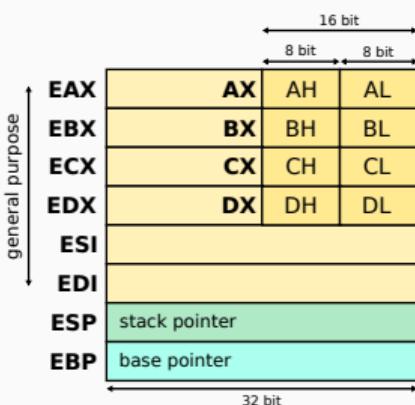
Linux Kernel Map:

<https://makelinux.github.io/kernel/map/>

```

1 int callee(int, int);
2 int caller(int a, int b, ..)
{
3     return callee(1, 2) + 5;
4 }
5

```



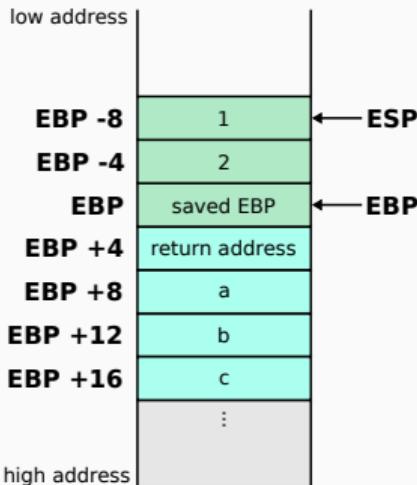
```

1 caller:
2 ; make new call frame (some compilers may produce an 'enter' instruction)
3 push    ebp      ; save old call frame
4 mov     ebp, esp ; initialize new call frame
5 ; push call arguments in reverse, or alternatively:
6 ; sub esp, 8      : 'enter' instruction might do this for us
7 ; mov [ebp-4], 2  : or mov [esp+8], 2
8 ; mov [ebp-8], 1  : or mov [esp+4], 1
9 push    2
10 push   1
11 call   callee   ; call subroutine 'callee'
12 add    eax, 5   ; modify subroutine result
13 ; (eax is the return value of our callee, thus no extra mov)
14 ; restore old call frame (some compilers may produce a 'leave' instruction)
15 ; add    esp, 8   ; remove arguments from frame, ebp - esp = 8.
16 ; compilers will usually produce the following instead,
17 ; which unlike the add instruction, also works for variable
18 ; length arguments/arrays
19 mov    esp, ebp ; most calling conventions dictate ebp be callee-saved
20 ; -> make sure the callee doesn't modify (or restores) ebp:
21 pop    ebp      ; restore old call frame
22 ret
23 ; return

```

```

1 int callee(int, int);
2 int caller(int a, int b, ..)
3 {
4     return callee(1, 2) + 5;
5 }
```



```

1 caller:
2 ; make new call frame (some compilers may produce an 'enter' instruction)
3 push    ebp      ; save old call frame
4 mov     ebp, esp ; initialize new call frame
5 ; push call arguments in reverse, or alternatively:
6 ; sub esp, 8      : 'enter' instruction might do this for us
7 ; mov [ebp-4], 2  : or mov [esp+8], 2
8 ; mov [ebp-8], 1  : or mov [esp+4], 1
9 push    2
10 push   1
11 call   callee   ; call subroutine 'callee'
12 add    eax, 5   ; modify subroutine result
13 ; (eax is the return value of our callee, thus no extra mov)
14 ; restore old call frame (some compilers may produce a 'leave' instruction)
15 ; add esp, 8      ; remove arguments from frame, ebp - esp = 8.
16 ; compilers will usually produce the following instead,
17 ; which unlike the add instruction, also works for variable
18     ; length arguments/arrays
19     ; mov esp, ebp   ; most calling conventions dictate ebp be callee-saved
20     ; -> make sure the callee doesn't modify (or restores) ebp:
21 pop    ebp      ; restore old call frame
22 ret
```

vgl. Grafik: <https://eli.thegreenplace.net/2011/02/04/where-the-top-of-the-stack-is-on-x86/>

Stellt POSIX Standard Funktionen bereit, und versteckt an vielen Stellen z.B. die Interaktion mit low-level Linux Systemcalls.

Es gibt diverse Implementationen der C Library jeweils mit unterschiedlichen Schwerpunkten:

- glibc
- uClibc-ng
- musl
- bionic (Android)
- newlib, dietlib, klibc (for very minimal systems)

Beispiel: glibc Source Code: [https://sourceware.org/git/gitweb.cgi?p=glibc.git;a=tree;f=sysdeps/x86\\_64/nptl;hb=HEAD](https://sourceware.org/git/gitweb.cgi?p=glibc.git;a=tree;f=sysdeps/x86_64/nptl;hb=HEAD)

# GCC Compiler Toolchain: Eine Übersicht

binutils

GNU Compiler Collection

Linux Kernel Headers

Application Binary Interface (ABI)

C Library

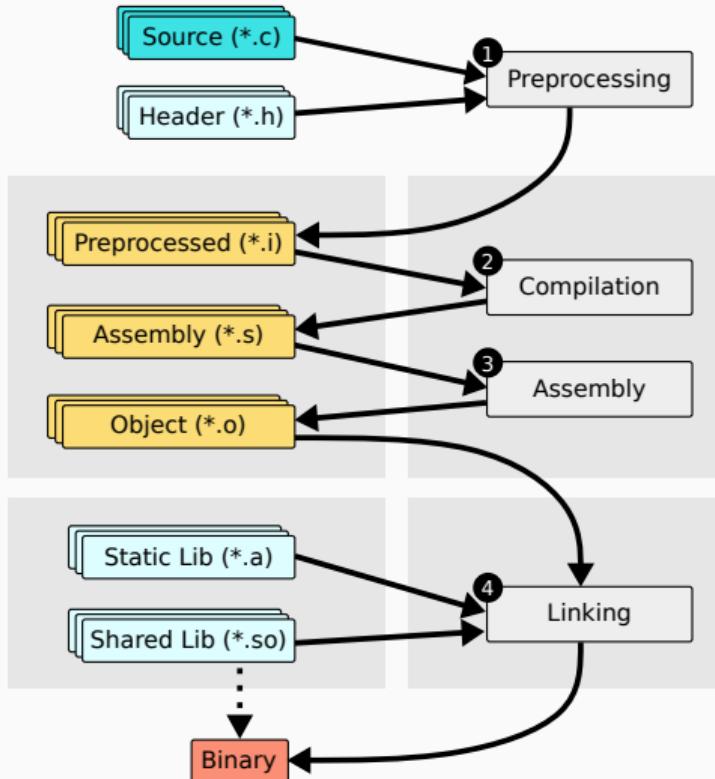
## Compilation Process mit GCC

Compiler Flags

Empfohlene Flags

Beispiel: `-fstack-protector`

Beispiel: `-D_FORTIFY_SOURCE`



## Ablauf <sup>a</sup>

1. `cpp hello.c > hello.i`
2. `gcc -S hello.i`
3. `as -o hello.o hello.s`
4. `ld -o hello hello.o -lc ...`

<sup>a</sup>Mehr Details: `gcc -v -o hello hello.c`

# GCC Compiler Toolchain: Eine Übersicht

binutils

GNU Compiler Collection

Linux Kernel Headers

Application Binary Interface (ABI)

C Library

Compilation Process mit GCC

## Compiler Flags

Empfohlene Flags

Beispiel: `-fstack-protector`

Beispiel: `-D_FORTIFY_SOURCE`

## Compiler Flags

---

GCC kommt mit vielen Optionen und Einstellungen die i.d.R. über sog. Compiler-Flags gesteuert werden.

- 1 -Wl werden an den linker (ld) weitergereicht (siehe "man ld")

```
1 # Sicherheit
2 -D_FORTIFY_SOURCE=2           Laufzeit Overflow Erkennung
3 -fpie -Wl,-pie               Address Space Layout Randomization (ASLR)
4 -fstack-clash-protection    Increased reliability of stack overflow detection
5 -fstack-protector           Overflow Erkennug via Canary (variants: all, strong) (RHEL6+)
6 -mcet -fcf-protection       Control flow integrity protection (future)
7 # Optimierung
8 -O2                         Recommended optimizations
9 -pipe                        Compile time optimization (avoid temporary files)
10 # Linker
11 -WL,-z,defs                Detect and reject unterlinking
12 -WL,-z,now                  Disable lazy binding (RHEL7+)
13 -WL,-z,relro                Read-only segments ofter relation (RHEL6+)
14 # Fehlerbehandlung
15 -fasynchronous-unwind-tables Increased reliability of backtraces
16 -fexceptions                Enable table-based thread cancellation
17 # Object Structure / Introspection
18 -fpic -shared                No text relocations for shared libraries
19 -fplugin=annobin             Inquire about hardening options, ABI compatability
20 # Debugging Informationen
21 -g                           Add debuggin information and labels
22 -grecord-gcc-switches        Compilerflags Metadata als debugging info
23 # Warnungen und Hinweise
24 -Wall                        Recommended compiler warnings
25 -Werror=format-security      Reject potentially unsafe format strings
26 -Werror=implicit-function-declaration Reject missing function prototypes
```

Siehe auch: <https://developers.redhat.com/blog/2018/03/21/compiler-and-linker-flags-gcc/>

```
1 void fun() {  
2     char *buf = alloca(0x100);  
3     /* Don't allow gcc to optimise away the buf */  
4     asm volatile("") :: "m" (buf));  
5 }
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 08048404 <fun>:  
2 push %ebp ; prologue  
3 mov %esp,%ebp  
4  
5 sub $0x128,%esp ; reserve 0x128B on the stack  
6 lea 0xf(%esp),%eax ; eax = esp + 0xf  
7 and $0xffffffff0,%eax ; align eax  
8 mov %eax,-0xc(%ebp) ; save eax in the stack frame  
9  
10 leave ; epilogue  
11 ret
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 08048464 <fun >:  
2 push %ebp ; prologue  
3 mov %esp,%ebp  
4  
5 sub $0x128,%esp ; reserve 0x128B on the stack  
6  
7 mov %gs:0x14,%eax ; load stack canary using gs  
8 mov %eax,-0xc(%ebp) ; save it in the stack frame  
9 xor %eax,%eax ; clear the register  
10  
11 lea 0xf(%esp),%eax ; eax = esp + 0xf  
12 and $0xfffffffff0,%eax ; align eax  
13 mov %eax,-0x10(%ebp) ; save eax in the stack frame  
14  
15 mov -0xc(%ebp),%eax ; load canary  
16 xor %gs:0x14,%eax ; compare against one in gs  
17 je 8048493 <fun+0x2f>  
18 call 8048340 <__stack_chk_fail@plt>  
19  
20 leave ; epilogue  
21 ret
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 void fun(char *s) {  
2     char buf[0x100];  
3     strcpy(buf, s); // Though you should prefer strncpy anyways! ;)  
4     /* Don't allow gcc to optimise away the buf */  
5     asm volatile("") :: "m" (buf));  
6 }
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 08048450 <fun>:  
2 push %ebp ; prologue  
3 mov %esp,%ebp  
4  
5 sub $0x118,%esp ; reserve 0x118B on the stack  
6 mov 0x8(%ebp),%eax ; load parameter s to eax  
7 mov %eax,0x4(%esp) ; save parameter for strcpy  
8 lea -0x108(%ebp),%eax ; count buf in eax  
9 mov %eax,(%esp) ; save parameter for strcpy  
10 call 8048320 <strcpy@plt>  
11  
12 leave ; epilogue  
13 ret
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 08048470 <fun>:  
2 push %ebp ; prologue  
3 mov %esp,%ebp  
4  
5 sub $0x118,%esp ; reserve 0x118B on the stack  
6 movl $0x100,0x8(%esp) ; save value 0x100 as parameter  
7 mov 0x8(%ebp),%eax ; load parameter s to eax  
8 mov %eax,0x4(%esp) ; save parameter for strcpy  
9 lea -0x108(%ebp),%eax ; count buf in eax  
10 mov %eax,(%esp) ; save parameter for strcpy  
11 call 8048370 <_strcpy_chk@plt>  
12  
13 leave ; epilogue  
14 ret
```

Siehe auch: [https://idea.popcount.org/2013-08-15-fortify\\_source/](https://idea.popcount.org/2013-08-15-fortify_source/)

```
1 /* Copyright (C) 1991–2018 Free Software Foundation, Inc.  
2 This file is part of the GNU C Library.  
3 The GNU C Library is free software; you can redistribute it and/or  
4 modify it under the terms of the GNU Lesser General Public  
5 License as published by the Free Software Foundation; either  
6 version 2.1 of the License, or (at your option) any later version.  
7 The GNU C Library is distributed in the hope that it will be useful,  
8 but WITHOUT ANY WARRANTY; without even the implied warranty of  
9 MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU  
10 Lesser General Public License for more details.  
11 You should have received a copy of the GNU Lesser General Public  
12 License along with the GNU C Library; if not, see  
13 <http://www.gnu.org/licenses/>. */  
14  
15 #include <stddef.h>  
16 #include <string.h>  
17 #include <memcopy.h>  
18  
19 #undef strcpy  
20  
21 /* Copy SRC to DEST with checking of destination buffer overflow. */  
22 char * __strcpy_chk (char *dest, const char *src, size_t destlen) {  
23     size_t len = strlen (src);  
24     if (len >= destlen)  
25         __chk_fail ();  
26     return memcopy (dest, src, len + 1);  
27 }
```

Siehe auch: [http://sourceware.org/git/?p=glibc.git;a=blob\\_plain;f=debug/strcpy\\_chk.c;hb=HEAD](http://sourceware.org/git/?p=glibc.git;a=blob_plain;f=debug/strcpy_chk.c;hb=HEAD)