

C declaration	Intel data type	Assembly-code suffix	Size (bytes)
char	Byte	b	1
short	Word	w	2
int	Double word	l	4
long	Quad word	q	8
char *	Quad word	q	8
float	Single precision	s	4
double	Double precision	l	8

Figure 3.1 Sizes of C data types in x86-64. With a 64-bit machine, pointers are 8 bytes long.



Figure 3.2 Integer registers. The low-order portions of all 16 registers can be accessed as byte, word (16-bit), double word (32-bit), and quad word (64-bit) quantities.

Type	Form	Operand value	Name
Immediate	$\$Imm$	Imm	Immediate
Register	r_a	$R[r_a]$	Register
Memory	Imm	$M[Imm]$	Absolute
Memory	(r_a)	$M[R[r_a]]$	Indirect
Memory	$Imm(r_b)$	$M[Imm + R[r_b]]$	Base + displacement
Memory	(r_b, r_i)	$M[R[r_b] + R[r_i]]$	Indexed
Memory	$Imm(r_b, r_i)$	$M[Imm + R[r_b] + R[r_i]]$	Indexed
Memory	$(, r_i, s)$	$M[R[r_i] \cdot s]$	Scaled indexed
Memory	$Imm(, r_i, s)$	$M[Imm + R[r_i] \cdot s]$	Scaled indexed
Memory	(r_b, r_i, s)	$M[R[r_b] + R[r_i] \cdot s]$	Scaled indexed
Memory	$Imm(r_b, r_i, s)$	$M[Imm + R[r_b] + R[r_i] \cdot s]$	Scaled indexed

Figure 3.3 Operand forms. Operands can denote immediate (constant) values, register values, or values from memory. The scaling factor s must be either 1, 2, 4, or 8.

Instruction	Effect	Description
MOV S, D	$D \leftarrow S$	Move
movb		Move byte
movw		Move word
movl		Move double word
movq		Move quad word
movabsq I, R	$R \leftarrow I$	Move absolute quad word

Figure 3.4 Simple data movement instructions.

Instruction	Effect	Description
MOVZ S, R	$R \leftarrow \text{ZeroExtend}(S)$	Move with zero extension
movzbw		Move zero-extended byte to word
movzbl		Move zero-extended byte to double word
movzwl		Move zero-extended word to double word
movzbq		Move zero-extended byte to quad word
movzwq		Move zero-extended word to quad word

Figure 3.5 Zero-extending data movement instructions. These instructions have a register or memory location as the source and a register as the destination.

Instruction	Effect	Description
MOVS S, R	$R \leftarrow \text{SignExtend}(S)$	Move with sign extension
movsbw		Move sign-extended byte to word
movsbl		Move sign-extended byte to double word
movswl		Move sign-extended word to double word
movsbq		Move sign-extended byte to quad word
movswq		Move sign-extended word to quad word
movslq		Move sign-extended double word to quad word
c1tq	$\%rax \leftarrow \text{SignExtend}(\%eax)$	Sign-extend %eax to %rax

Figure 3.6 Sign-extending data movement instructions. The movs instructions have a register or memory location as the source and a register as the destination. The c1tq instruction is specific to registers %eax and %rax.

Instruction	Effect	Description
pushq <i>S</i>	$R[\%rsp] \leftarrow R[\%rsp] - 8;$ $M[R[\%rsp]] \leftarrow S$	Push quad word
popq <i>D</i>	$D \leftarrow M[R[\%rsp]];$ $R[\%rsp] \leftarrow R[\%rsp] + 8$	Pop quad word

Figure 3.8 Push and pop instructions.

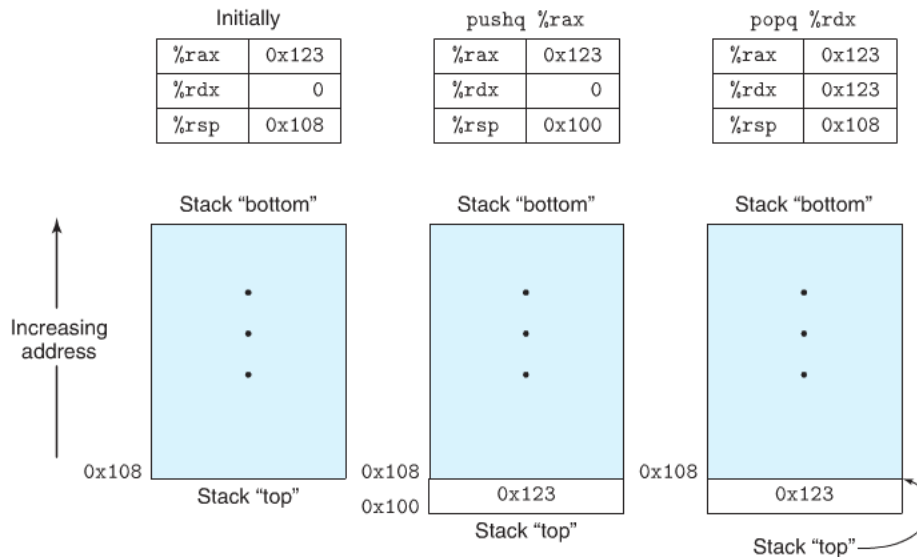


Figure 3.9 Illustration of stack operation. By convention, we draw stacks upside down, so that the “top” of the stack is shown at the bottom. With x86-64, stacks grow toward lower addresses, so pushing involves decrementing the stack pointer (register `%rsp`) and storing to memory, while popping involves reading from memory and incrementing the stack pointer.

Instruction	Effect	Description
leaq <i>S, D</i>	$D \leftarrow \&S$	Load effective address
INC <i>D</i>	$D \leftarrow D + 1$	Increment
DEC <i>D</i>	$D \leftarrow D - 1$	Decrement
NEG <i>D</i>	$D \leftarrow -D$	Negate
NOT <i>D</i>	$D \leftarrow \sim D$	Complement
ADD <i>S, D</i>	$D \leftarrow D + S$	Add
SUB <i>S, D</i>	$D \leftarrow D - S$	Subtract
IMUL <i>S, D</i>	$D \leftarrow D * S$	Multiply
XOR <i>S, D</i>	$D \leftarrow D \wedge S$	Exclusive-or
OR <i>S, D</i>	$D \leftarrow D \vee S$	Or
AND <i>S, D</i>	$D \leftarrow D \& S$	And
SAL <i>k, D</i>	$D \leftarrow D \ll k$	Left shift
SHL <i>k, D</i>	$D \leftarrow D \ll k$	Left shift (same as SAL)
SAR <i>k, D</i>	$D \leftarrow D \gg_A k$	Arithmetic right shift
SHR <i>k, D</i>	$D \leftarrow D \gg_L k$	Logical right shift

Figure 3.10 Integer arithmetic operations. The load effective address (`leaq`) instruction is commonly used to perform simple arithmetic. The remaining ones are more standard unary or binary operations. We use the notation \gg_A and \gg_L to denote arithmetic and logical right shift, respectively. Note the nonintuitive ordering of the operands with ATT-format assembly code.

Instruction		Effect	Description
imulq	<i>S</i>	$R[\%rdx]:R[\%rax] \leftarrow S \times R[\%rax]$	Signed full multiply
mulq	<i>S</i>	$R[\%rdx]:R[\%rax] \leftarrow S \times R[\%rax]$	Unsigned full multiply
cqto		$R[\%rdx]:R[\%rax] \leftarrow \text{SignExtend}(R[\%rax])$	Convert to oct word
idivq	<i>S</i>	$R[\%rdx] \leftarrow R[\%rdx]:R[\%rax] \bmod S;$ $R[\%rax] \leftarrow R[\%rdx]:R[\%rax] \div S$	Signed divide
divq	<i>S</i>	$R[\%rdx] \leftarrow R[\%rdx]:R[\%rax] \bmod S;$ $R[\%rax] \leftarrow R[\%rdx]:R[\%rax] \div S$	Unsigned divide

Figure 3.12 Special arithmetic operations. These operations provide full 128-bit multiplication and division, for both signed and unsigned numbers. The pair of registers %rdx and %rax are viewed as forming a single 128-bit oct word.

Condition Flag Registers

CF: Carry flag. The most recent operation generated a carry out of the most significant bit. Used to detect overflow for unsigned operations.

ZF: Zero flag. The most recent operation yielded zero.

SF: Sign flag. The most recent operation yielded a negative value.

OF: Overflow flag. The most recent operation caused a two's-complement overflow—either negative or positive.

Instruction		Based on	Description
CMP	S_1, S_2	$S_2 - S_1$	Compare
cmpb			Compare byte
cmpw			Compare word
cmpl			Compare double word
cmpq			Compare quad word
TEST	S_1, S_2	$S_1 \& S_2$	Test
testb			Test byte
testw			Test word
testl			Test double word
testq			Test quad word

Figure 3.13 Comparison and test instructions. These instructions set the condition codes without updating any other registers.

Instruction	Synonym	Effect	Set condition
sete <i>D</i>	setz	$D \leftarrow \text{ZF}$	Equal / zero
setne <i>D</i>	setnz	$D \leftarrow \sim \text{ZF}$	Not equal / not zero
sets <i>D</i>		$D \leftarrow \text{SF}$	Negative
setns <i>D</i>		$D \leftarrow \sim \text{SF}$	Nonnegative
setg <i>D</i>	setnle	$D \leftarrow \sim (\text{SF} \wedge \text{OF}) \& \sim \text{ZF}$	Greater (signed >)
setge <i>D</i>	setnl	$D \leftarrow \sim (\text{SF} \wedge \text{OF})$	Greater or equal (signed >=)
setl <i>D</i>	setnge	$D \leftarrow \text{SF} \wedge \text{OF}$	Less (signed <)
setle <i>D</i>	setng	$D \leftarrow (\text{SF} \wedge \text{OF}) \mid \text{ZF}$	Less or equal (signed <=)
seta <i>D</i>	setnbe	$D \leftarrow \sim \text{CF} \& \sim \text{ZF}$	Above (unsigned >)
setae <i>D</i>	setnb	$D \leftarrow \sim \text{CF}$	Above or equal (unsigned >=)
setb <i>D</i>	setnae	$D \leftarrow \text{CF}$	Below (unsigned <)
setbe <i>D</i>	setna	$D \leftarrow \text{CF} \mid \text{ZF}$	Below or equal (unsigned <=)

Figure 3.14 The SET instructions. Each instruction sets a single byte to 0 or 1 based on some combination of the condition codes. Some instructions have “synonyms,” that is, alternate names for the same machine instruction.

Instruction	Synonym	Jump condition	Description
jmp <i>Label</i>		1	Direct jump
jmp <i>*Operand</i>		1	Indirect jump
je <i>Label</i>	jz	ZF	Equal / zero
jne <i>Label</i>	jnz	~ZF	Not equal / not zero
js <i>Label</i>		SF	Negative
jns <i>Label</i>		~SF	Nonnegative
jg <i>Label</i>	jnle	~(SF ^ OF) & ~ZF	Greater (signed >)
jge <i>Label</i>	jnl	~(SF ^ OF)	Greater or equal (signed >=)
jl <i>Label</i>	jnge	SF ^ OF	Less (signed <)
jle <i>Label</i>	jng	(SF ^ OF) ZF	Less or equal (signed <=)
ja <i>Label</i>	jnbe	~CF & ~ZF	Above (unsigned >)
jae <i>Label</i>	jnb	~CF	Above or equal (unsigned >=)
jb <i>Label</i>	jnae	CF	Below (unsigned <)
jbe <i>Label</i>	jna	CF ZF	Below or equal (unsigned <=)

Figure 3.15 The jump instructions. These instructions jump to a labeled destination when the jump condition holds. Some instructions have “synonyms,” alternate names for the same machine instruction.

Instruction	Synonym	Move condition	Description
cmovz <i>S, R</i>	cmovz	ZF	Equal / zero
cmovne <i>S, R</i>	cmovnz	~ZF	Not equal / not zero
cmovs <i>S, R</i>		SF	Negative
cmovns <i>S, R</i>		~SF	Nonnegative
cmovg <i>S, R</i>	cmovnle	~(SF ^ OF) & ~ZF	Greater (signed >)
cmovge <i>S, R</i>	cmovnl	~(SF ^ OF)	Greater or equal (signed >=)
cmovl <i>S, R</i>	cmovnge	SF ^ OF	Less (signed <)
cmovle <i>S, R</i>	cmovng	(SF ^ OF) ZF	Less or equal (signed <=)
cmova <i>S, R</i>	cmovnbe	~CF & ~ZF	Above (unsigned >)
cmovae <i>S, R</i>	cmovnb	~CF	Above or equal (Unsigned >=)
cmovb <i>S, R</i>	cmovnae	CF	Below (unsigned <)
cmovbe <i>S, R</i>	cmovna	CF ZF	Below or equal (unsigned <=)

Figure 3.18 The conditional move instructions. These instructions copy the source value *S* to its destination *R* when the move condition holds. Some instructions have “synonyms,” alternate names for the same machine instruction.

Alignment

The x86-64 hardware will work correctly regardless of the alignment of data. However, Intel recommends that data be aligned to improve memory system performance. Their alignment rule is based on the principle that any primitive object of *K* bytes must have an address that is a multiple of *K*. We can see that this rule leads to the following alignments:

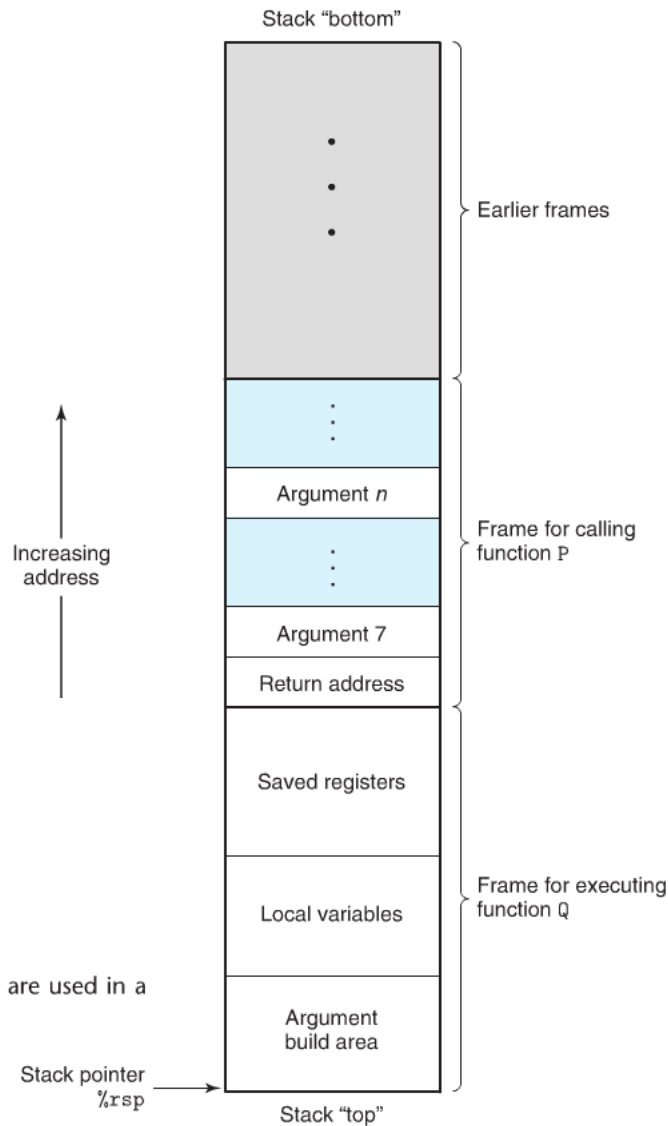
<i>K</i>	Types
1	char
2	short
4	int, float
8	long, double, char *

Procedures/Calling Conventions

Figure 3.25
General stack frame structure. The stack can be used for passing arguments, for storing return information, for saving registers, and for local storage. Portions may be omitted when not needed.

Operand size (bits)	Argument number					
	1	2	3	4	5	6
64	%rdi	%rsi	%rdx	%rcx	%r8	%r9
32	%edi	%esi	%edx	%ecx	%r8d	%r9d
16	%di	%si	%dx	%cx	%r8w	%r9w
8	%dil	%sil	%dl	%cl	%r8b	%r9b

Figure 3.28 Registers for passing function arguments. The registers are used in a specified order and named according to the argument sizes.



Instruction	Description
call <i>Label</i>	Procedure call
call <i>*Operand</i>	Procedure call
ret	Return from call

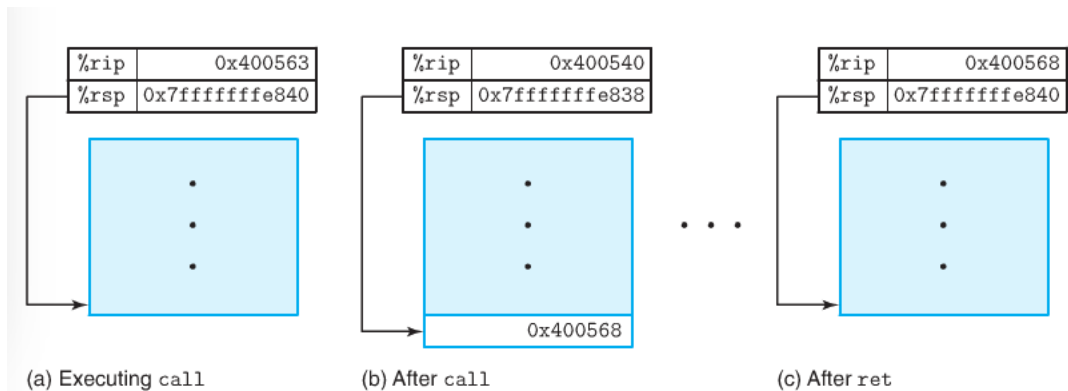


Figure 3.26 Illustration of call and ret functions. The call instruction transfers control to the start of a function, while the ret instruction returns back to the instruction following the call.