



IA-64 Assembly Language Reference Guide

January 2000

Order Number 245363-001

THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE.

Information in this document is provided in connection with Intel products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

Intel processors associated with the Assembly Language may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's website at <http://developer.intel.com/design/litcentr>.

Copyright © Intel Corporation, 2000

*Third-party brands and names are the property of their respective owners.



Contents

Chapter 1 Overview

About this Manual	1-1
Related Documentation	1-2
Notation Conventions	1-2

Chapter 2 Program Elements Overview

Identifiers	2-1
Name Spaces	2-2
Symbols	2-2
Symbol Names.....	2-3
Symbol Types.....	2-5
Symbol Values	2-6
Register Names	2-6
Mnemonics	2-7
Machine Instruction Mnemonics	2-8
Pseudo-op Mnemonics	2-8
Directive Mnemonics	2-9
Data Allocation Mnemonics	2-9
Constants	2-9
Numeric Constants	2-9
C Numeric Constants	2-10
MASM Numeric Constants.....	2-11
Characters in Numeric Constants	2-13

String Constants	2-14
Expressions	2-14
Absolute Expressions	2-14
Relocatable Expressions	2-14
Operators	2-15
Statements.....	2-16
Label Statements	2-16
Instruction Statements	2-17
Directive Statements.....	2-19
Assignment Statements	2-19
Symbol Assignment Statements	2-20
Register Assignment Statements	2-20
Equate Statements	2-21
Symbol Equate Statements.....	2-21
Register Equate Statements	2-21
Data Allocation Statements.....	2-22
Cross-section Data Allocation Statements.....	2-23

Chapter 3 Program Structure

Sections	3-1
Section Flags and Section Type Operands	3-2
Windows NT (COFF32) Specific Section Flag Operands .	3-3
Associated Section Name Flag	3-3
Section Definition Directive	3-4
Section Stack Directives	3-5
Absolute Sections	3-6
Section Return Directive	3-6
Predefined Section Directives.....	3-6
Using Section Directives	3-8
Include File Directive	3-8
Bundles.....	3-9
Implicit Bundling.....	3-10

Explicit Bundling	3-10
Auto-template Selection.....	3-10
Explicit-template Selection.....	3-11
Instruction Groups.....	3-12
Dependency Violations and Assembly Modes	3-13
Procedures.....	3-13
Procedure Directives	3-14
Procedure Label (PLabel)	3-14
Stack Unwind Directives	3-15
Syntax for the .save.x Directives.....	3-21
Stack Unwind Directives Usage Guidelines.....	3-22
Windows NT (COFF32) Symbolic Debug Directives	3-24

Chapter 4 Declarations

Symbol Scope Declaration	4-1
Local Scope Declaration Directive.....	4-2
Global Scope Declaration Directive.....	4-2
Weak Scope Declaration Directive	4-2
Weak Scope Declaration for UNIX (ELF).....	4-3
Weak Scope Declaration for Windows NT (COFF32)..	4-3
Symbol Type Directive.....	4-4
Symbol Size Directive	4-5
File Name Override Directive	4-6
Common Symbol Declarations.....	4-6
Common Symbol Directive	4-6
Local Common Symbol Directive	4-7
Alias Declaration Directives.....	4-8

Chapter 5 Data Allocation

Data Allocation Statements.....	5-1
Uninitialized Space Allocation	5-3
Alignment	5-4
Cross-section Data Allocation Statements.....	5-5

Chapter 6 Miscellaneous Directives	
Register Stack Directive.....	6-1
Stacked Registers in Assignment and Equate Statements	6-3
Rotating Register Directives	6-3
Using Rotating Register directives	6-5
Rotating Registers in Assignment and Equate Statements	6-6
Byte Order Specification Directives	6-7
String Specification Directive	6-7
Radix Indicator Directive	6-8
Preprocessor Support.....	6-8
Chapter 7 Annotations	
Predicate Relationship Annotation.....	7-1
Predicate Vector Annotation	7-2
Memory Offset Annotation.....	7-2
Entry Annotation	7-3
Appendix A Register Names by Type	
Appendix B Pseudo-ops	
Appendix C Link-relocation Operators	
Appendix D List of IA-64 Assembly Language Directives	
Glossary	
Index	



Overview

This manual describes the programming conventions used to write an assembly program for the IA-64 architecture.

As prerequisites, you should be familiar with the IA-64 architecture, and have assembly-language programming experience.

About this Manual

This manual contains the following chapters and appendixes:

- This chapter lists related documentation and notation conventions.
- [Chapter 2, “Program Elements Overview”](#), describes the basic elements and language specifications of an assembly-language program for the IA-64 architecture.
- [Chapter 3, “Program Structure”](#), describes the directives used to structure the program.
- [Chapter 4, “Declarations”](#), describes the directives used to declare symbols in the program.
- [Chapter 5, “Data Allocation”](#), describes the statements used to allocate initialized and uninitialized space for data objects, and align data objects in the program.
- [Chapter 6, “Miscellaneous Directives”](#), describes directives not used to structure a program or to declare symbols.
- [Chapter 7, “Annotations”](#), describes the assembler annotations.

- [Appendix A. "Register Names by Type"](#), lists the IA-64 architecture registers.
- [Appendix B. "Pseudo-ops"](#), lists the IA-64 architecture pseudo operations and their equivalent machine instructions, and pseudo-ops with missing operands.
- [Appendix C. "Link-relocation Operators"](#), lists the link-relocation operators and describes their functionality.
- [Appendix D. "List of IA-64 Assembly Language Directives"](#), lists the assembly-language directives according to category.

Related Documentation

The following documents, available at <http://developer.intel.com>, provide additional information:

- *IA-64 Architecture Software Developer's Manual*
Volume 1: IA-64 Processor Application Architecture,
order number 245317-001
Volume 2: IA-64 Processor System Architecture,
order number 245318-001
Volume 3: IA-64 Instruction Set Description,
order number 245319-001
Volume 4: IA-32 Instruction Set Description,
order number 245320-001
- *IA-64 Software Conventions and Runtime Architecture Guide*,
order number 245256-002

Notation Conventions

This notation is used in syntax descriptions:

`This type style`

Indicates an element of syntax, a reserved word, keyword, a filename, computer output, or part of a program example. The text appears in lowercase, unless uppercase is significant.

This type style	Indicates the text you enter as input.
<i>This type style</i>	Indicates a placeholder for an identifier, an expression, a string, a symbol or a value. Substitute one of these items for the placeholder.
[items]	Indicates optional items.
[items item]	Indicates the possible choices. A vertical bar () separates the items. Choose one of the items enclosed in brackets.



Program Elements Overview

2

This chapter describes the basic elements and language specifications of an assembly-language program for the IA-64 architecture. The basic program elements are identifiers, symbols, name spaces, constants, expressions, and statements.

Identifiers

In IA-64 assembly language, objects such as machine instructions, registers, memory locations, sections in the object file, and constants, have symbolic names. In the source code these names are represented syntactically by identifiers.

An identifier may contain letters, digits, and a few special characters. Identifiers may not begin with a digit.

Table 2-1 summarizes the rules for character usage in identifiers.

Table 2-1 Character Usage in Identifiers

Character Types	First Characters	Remaining Characters
Letters	a-z or A-Z	a-z or A-Z
Special characters	@ _ \$? .	@ _ \$? .
Digits	not allowed	0-9

The assembler may place a limit on the length of an identifier, but this limit must be no less than 256 characters.

Name Spaces

There are three classes of names in the IA-64 assembly language:

- Symbols, which refer to memory locations, sections, and symbolic constants. These names are case sensitive.
- Registers, which refer to registers defined in the IA-64 architecture. These names are not case sensitive. Some register names consist of multiple syntactic elements rather than a single identifier.
- Mnemonics, which refer to machine instructions, pseudo-ops, directives, and completers. These names are not case sensitive.

The assembler places names in three separate name spaces, according to their class. A name may not be defined twice in the same namespace, but it may be defined once in each namespace. When a name is defined in both the register and symbol namespaces, the register name takes precedence over the symbol unless the identifier is “protected” by terminating it with the # operator; this forces the assembler to look up the identifier in the symbol namespace.

The # operator in conjunction with a symbol is legal only when the symbol is an operand.

The following examples illustrate the correct use of the # operator:

```
r5: //label named r5, where label is the symbol name
movl r4=r5#//moves the r5 label address to register r4
.global r5#//declares label r5 as global
```

The # operator is unnecessary and illegal when included in the symbol definition, as shown:

```
r5#: //illegal
```

Symbols

A symbol refers to a location in memory, an object file section, a numeric constant, or a register. A symbol has the following attributes:

- name

- type
- value

The special symbols dollar sign (\$) and period (.) when used in expressions, always refer to the current location counter. The current location counter points to the address of a bundle containing the current instruction, or to the address of the first data object defined by the current assembly statement. There is no difference between these symbols, either can be used.

In the following example, the `movl` instruction moves the address of the bundle containing the current instruction (\$) into register `r1`:

```
movl    r1=$
```

In the following data allocation statement, the period (.) is the address of the first data object defined by the assembly statement:

```
data4   2, 3, .
```

Symbol Names

Symbol names are case-sensitive identifiers. Symbols whose names begin with a period (.) are temporary. Temporary symbols are not placed in the object file symbol table. Symbols whose names begin with two periods (..) are temporary, and local. Local symbols are scope restricted symbols. Local symbols are recognized only within the scope in which they are defined. See the "Symbol Scope Declaration" section in Chapter 4 for more information about local symbol scopes.

Table 2-2 summarizes the rules for using temporary and scope-restricted indicators in different types of symbol names.

Table 2-2 Temporary and Scope-restricted Indicators in Symbol Names

Symbol Type	Temporary (.)	Temporary and Scope Restricted (..)
Labels	Allowed	Allowed
Instruction tags	Allowed	Allowed
Function names	Not allowed	Not allowed

Table 2-2 Temporary and Scope-restricted Indicators in Symbol Names

Symbol Type	Temporary (.)	Temporary and Scope Restricted (. .)
Symbolic constants	Not allowed	Allowed
Section names	Allowed	Not allowed

Symbols whose names begin with an "at" sign (@) are reserved as predefined constants. The assembler provides predefined symbolic constants for special operand values for several instructions, for example, `fclass` and `mux` instructions. Table 2-3 and Table 2-4 list the predefined symbolic constant names for the operands of these instructions. These symbolic constants can be used in expressions as any user-defined symbolic constant.

Table 2-3 fclass Condition Predefined Operand Names

Category	fclass Conditions	Predefined Name
NaN test	NaN	@nat
Sign test	Positive	@pos
	Negative	@neg
Class test	Normalized	@norm
	Unnormalized	@unorm
	Signaling NaN	@snan
	Quiet NaN	@qnan
	Zero	@zero
	Infinity	@inf

Table 2-4 mux Bytes Operation Predefined Type Operand Names

mux Bytes Operation Type (<i>mbtype</i>)	Predefined Name
Reverse	@rev
Mix	@mix
Shuffle	@shuf
Alternate	@alt
Broadcast	@brcst

Symbol Types

A symbol's type indicates the class of object to which it refers. A symbol type can be any of the following:

label Refers to a location of code or data in memory. A label cannot refer to a procedure entry point. A code label refers to the address of a bundle. An instruction that follows a code label always starts a new bundle. The "Bundles" section in Chapter 3 provides more information about instruction bundling.

instruction tag A symbol that refers to an instruction. An instruction tag is used in branch prediction instructions, and in unwind information directives. Unlike a label, an instruction tag does not cause the instruction to start a new bundle.

function name A symbol that refers to a procedure entry point.

section name Represents an existing section that is active in the output object file.

symbolic constant A constant assigned or equated to a number, symbol, or expression.

Symbol Values

A symbol is defined when it is assigned a value. A symbol value can also be a number or expression assigned to a symbolic constant. The value of a symbol identifies the object to which it refers. If the symbol refers to a location in memory, the assigned value is the address of that memory location. In most cases, this address is resolved only in link time.

Register Names

All registers have predefined names, which are listed in Appendix A. Predefined register names are not case-sensitive. You can assign new register names to some of the predefined registers with a register assignment statement, or a rotating register directive. See the "Assignment Statements", and "Equate Statements" section in this chapter, and the "Rotating Register Directives" section in Chapter 6, for more details. Registers that use the value of a specified general-purpose register as an index into the register file consist of the register file name followed by the name of a general register enclosed in brackets, such as `pmc[r]`.

The assembler determines the register type according to the form of its name, as shown in Table 2-5. Some registers appear in name and number form. For example, `ar.bsp` is the name form of an application register, which also has a number form, `ar17`.

Table 2-5 Register Number and Name Forms

Register Form	Register Name	Register Type
Number form	r0 - r127	General-purpose 64-bit registers
	in0 - in95	
	loc0 - loc95	
	out0 - out95	
	f0 - f127	Floating-point registers
	p0 - p63	Predicate registers (1-bit)
	b0 - b7	Branch registers
	ar0 - ar127	Application registers
	cr0 - cr127	Control registers
Name form	e.g. ar.bsp	Named application registers
	e.g. cr.dcr	Named control registers
	pr	All-predicate register (64-bits)
	pr.rot	All rotating registers
	ip	Instruction pointer
	psr.l	Processor-status registers
	psr.um	
Indirect file registers	e.g. pmc[r2]	Register file with general-purpose register as index.
User-defined registers	<i>user-name</i>	Registers assigned new names with assignment statements or rotating register directives.

Mnemonics

Mnemonics are predefined assembly-language names for machine instructions, pseudo-ops, directives, and data allocation statements. Mnemonics are not case-sensitive.

Machine Instruction Mnemonics

Machine instruction mnemonics specify the operation to be performed. For example, `brp` is the mnemonic for the branch predict instruction.

Some instruction mnemonics include suffixes and optional completers that indicate variations on the basic operation. The suffixes and completers are separated from the basic mnemonic by a period (.). For example, the instructions `brp.pp` (predicate predict), and `brp.ret` (return) include suffixes, and are variations of the basic branch predict (`brp`) instruction.

In this manual, completers are italicized to distinguish them from the instruction mnemonic suffixes. For example, in the instruction `brp.ret.sptk.imp b0,L`, the optional completers appear in italics to set them apart from the `.ret` suffix.

For a full description of the instructions, see the *IA-64 Architecture Software Developer's Manual*.

Pseudo-op Mnemonics

Pseudo-op mnemonics represent assembler instructions that alias machine instructions. They are equivalent to instruction mnemonics and are provided for the convenience of the programmer. See Appendix B for a list of the assembler pseudo-ops.

The following is an example of a pseudo-op:

```
mov r5=2
```

The assembler translates this pseudo-op into the equivalent machine instruction:

```
addl r5=2,r0
```

For more details about the pseudo-ops, see the *IA-64 Architecture Software Developer's Manual*.

Directive Mnemonics

Directives are assembler instructions to the assembler during the assembly process; they do not produce executable code. To distinguish them from other instructions, directive mnemonics begin with a period (.).

Chapter 3 through Chapter 7 describe the assembler directives and explain how to use them.

Data Allocation Mnemonics

Data allocation mnemonics specify the types of data objects assembled in data allocation statements. See the "Data Allocation" section in Chapter 5 for a list of these mnemonics. Data allocation statements are used to allocate initialized memory areas.

Constants

Constants can be numeric or string.

- Numeric constants contain integers and floating-point numbers.
- String constants contain one or more characters.

Numeric Constants

A numeric constant contains integer and floating-point numbers. The assembler supports C and Microsoft Macro Assembly language (MASM) numeric constants. C numeric constants are the default.

C Numeric Constants

C numeric constants can be any of the following:

- **Decimal integer constants** (base 10) consist of one or more digits, 0 through 9, where 0 cannot be used as the first digit.
- **Binary constants** (base 2) begin with a 0b or 0B prefix, followed by one or more binary digits (0, 1).
- **Octal constants** (base 8) consist of one or more digits 0 through 7, where the first digit is 0.
- **Hexadecimal constants** (base 16) begin with a 0x or 0X prefix, followed by a hexadecimal number represented by a combination of digits 0 through 9, and characters A through F.
- **Floating-point constants** consist of:
 - an optional sign - or +
 - an integer part a combination of digits 0 through 9
 - a period .
 - a fractional part a sequence of digits 0 through 9
 - an optional exponent e or E, followed by an optionally signed sequence of one or more digits

For example, the following floating-point constant contains both the optional and required parts: +1.15e-12.

The following floating-point constant contains only the required parts: 1.0

The following formal grammar summarizes the rules for the C numeric constants:

C-constant:

C-integer-constant
floating-point-constant
character-constant

C-integer-constant:

```
[1-9][0-9]*  
0[bB][01]*  
0[0-7]*  
0[xX][0-9a-fA-F]*
```

floating-point-constant:

```
integer-part.[ fractional-part ] [ exponent-part ]
```

integer-part:

```
[0-9]*
```

fractional-part:

```
[0-9]*
```

exponent-part:

```
[eE][+-][0-9]*
```

```
[eE][0-9]*
```

MASM Numeric Constants

MASM numeric constants can be any of following:

- **Radix constants** are numeric constants that also specify the radix of the value. They consist of one or more digits, 0 through 9, followed by a radix indicator. The radix indicators of MASM numeric constants define them as decimal (D), hexadecimal (H), octal (O), or binary (B). If the current radix is hexadecimal, the letters B and D are interpreted as digits. In this case, T specifies a decimal radix, and Y specifies a binary radix. See Table 2-6.

Radix indicators are not case-sensitive.

See the "Radix Indicator Directive" section in Chapter 6, for more information about how to specify a radix.

- **Integer constants** in the current radix consist of one or more digits, 0 through 9, A through F. If the current radix is not hexadecimal, the characters A through F are not applicable.
- **Floating-point constants** have the same syntax as in C. See the "C Numeric Constants" section on page 2-10.

Table 2-6 MASM Radix Indicators

Radix	Radix Indicator Suffix
Decimal	D (d), or T (t) when the current radix is hex
Hexadecimal	H (h)
Octal	O (o) or Q (q)
Binary	B (b), or Y (y) when the current radix is hex

The following formal grammar summarizes the rules for the MASM numeric constants:

MASM-constant :

MASM-integer-constant

MASM-radix-constant

floating-point-constant

character-constant

MASM-integer-constant :

[0-9][0-9a-fA-F]*

MASM-radix-constant

[0-9][0-9a-fA-F]*[tTdDhHOoqQbByY]

floating-point-constant : (as in C, see page 2-11.)

Characters in Numeric Constants

An underscore (`_`) can be inserted in a numeric constant to improve readability, as follows `1_000_000`. An underscore can be inserted anywhere except before the first character. The assembler ignores underscores.

Characters can represent numeric constants. For instance, a single ASCII character can represent a numeric constant by enclosing it in single quotes (`' '`). The numeric constant is the ASCII value of the specified character. To use other special characters to represent numeric constants, use the character escapes defined in the ANSI C language, and enclose them in single quotes. Table 2-7 lists the common character escapes. To use the single quote (`' '`) to represent a numeric constant, insert a backslash (`\`) before it, and enclose both in single quotes (`' '`), as such, `'\''`.

Table 2-7 Common Character Escapes

Escape Character	Definition	ASCII Value
<code>\'</code>	Single quote	39
<code>\"</code>	Double quote	34
<code>\b</code>	Backspace	8
<code>\t</code>	Tab	9
<code>\n</code>	New line	10
<code>\f</code>	Form feed	12
<code>\r</code>	Carriage return	13
<code>\\</code>	Backslash	92
<code>\num</code>	Character with octal value <i>num</i> (maximum three digits)	-
<code>\Xhh</code>	Character with the hexadecimal value <i>hh</i> (maximum two digits)	-

String Constants

String constants consist of a sequence of characters enclosed in double quotes (" ").

To specify double-quotes (") in a string constant, insert a backslash (\) before it, as such, "\\ ".

To include other special characters in a string constant, use the character escapes defined in the ANSI C language. See Table 2-7 for a list of common character escapes.

Expressions

An expression is a combination of symbols, numeric constants, and operators that uses standard arithmetic notation to yield a result. Expressions can be absolute or relocatable.

Absolute Expressions

An expression is absolute when it is not subject to link-time relocation. An absolute expression may contain relocatable symbols, but they must reduce to pairs of the form $(R_1 - R_2)$, where R_1 and R_2 are relocatable symbols defined in the same section in the current source file.

Relocatable Expressions

An expression is relocatable when it is subject to link-time relocation. A relocatable expression contains a relocatable symbol, and may contain an absolute expression. If a relocatable expression contains an absolute expression, it must be reducible to the form $(R + K)$, where R is either a relocatable symbol defined in the current source file, or an undefined symbol, and K is an absolute expression. The address of the relocatable symbol is defined in link time.

Operators

The assembly operators indicate arithmetic or bitwise-logic calculations. Parentheses (()) determine the order in which calculations occur. The assembler evaluates all operators of the same precedence from left to right.

The assembler evaluates all operators according to their level of precedence. Table 2-8 lists the operator precedence rules from lowest to highest.

Table 2-8 Precedence of Arithmetic and Bitwise Logic Operations

Precedence	Operator Symbol	Operation
0 (Low)	+	Addition
	-	Subtraction
		Bitwise inclusive OR
	^	Bitwise exclusive OR
1 (Medium)	*	Multiplication
	/	Division
	%	Remainder
	<<	Shift Left
	>>	Arithmetic shift right
	&	Bitwise AND
2 (High)	-	Unary negation
	~	Unary one's complement

Link-relocation Operators

Link-relocation operators generate link-relocation entries in expressions. See Appendix C for a list of the link-relocation operators.

Statements

An assembly-language program consists of a series of statements separated by a semicolon (;). Multiple statements may be on the same line.

To separate lines, use the standard line termination convention on the local host system, typically CR (carriage return) and LF (line feed). To separate elements within a statement, use the CR, LF, FF (form feed), VT (vertical tab), Space, or Tab that represent white space.

To separate a comment from the code at the end of a statement, insert the comment before the semi colon (;) and precede it with a double-backslash (//). The assembler ignores comments.

The assembler may place a limit on the length of an input line, but this limit must be no less than 256 characters.

The types of assembly-language statements are as follows:

- label statements
- instruction statements
- directive statements
- assignment statements
- equate statements
- data allocation statements
- cross-section data allocation statements

The following sections detail each of the statement types, their components and syntax, and provide an example of each.

Label Statements

A label statement has the following syntax:

```
[label]:      // comments
```

Where:

label Defines a symbol whose value is the address of the current location counter. If the assembler inserts padding to align the location counter to an implied alignment boundary, the value of the label is not affected.

The assembler interprets a label followed by a double colon (: :) as a global symbol. See the "Symbol Scope Declaration" section in Chapter 4 for more information about global symbols.

The following is an example of a global label statement:

```
foo::
```

Instruction Statements

An instruction statement has the following syntax:

```
[label:] [[tag:]] [(qp)] mnemonic[.completers]  
dests=sources //comments
```

Where:

label Defines a symbol whose value is the address of a bundle. When a label is present, the assembler always starts a new bundle.

If the assembler inserts padding to align the location counter to a bundle boundary, the label is assigned the address of the newly-aligned bundle.

The assembler interprets a label followed by a double colon (: :) as a global symbol. See the "Symbol Scope Declaration" section in Chapter 4 for more information about global symbols.

[*tag*] Defines a symbol whose value is the bundle address and slot number of the current instruction.

(*qp*) Represents a predicate register symbol, which must be enclosed in parentheses. If this field is not defined, predicate register 0 (*p0*) is the default.

mnemonic.completers

Represents the instruction mnemonic or pseudo-op. Instructions may optionally include one or more completers. Completers must appear in the specified order in the instruction syntax.

Mnemonics and completer mnemonics are not case-sensitive.

Refer to the *IA-64 Architecture Software Developer's Manual* for a description of the machine instructions, pseudo-ops, and completers.

dests =sources

Represents the destination and source operands. The operands are register names, expressions, or keywords, depending on the instruction. Some instructions can have two destination operands, and one or more source operands. When there are multiple operands they are separated by a comma (,). In cases where all operands are destination operands or all operands are source operands, the equal (=) sign is omitted.

The following is an example of an instruction statement with a *label* and (*qp*):

```
L5: (p7) addl r14 = @gprel(L0), r1
```

The following is an example of an instruction statement with a *tag*:

```
[t1:] fclass.m.unc p4, p5 = f6, @pos
```

@pos is a predefined constant representing the *fclass* operation. p4 is true if f6 is positive.

Directive Statements

A directive statement has the following syntax:

```
.directive    [operands]    // comments
```

Where:

.directive Represents the directive mnemonic. Directives always begin with a period (.). Directive mnemonics are not case-sensitive.

operands The operands are optional and determined by the directive. Where multiple operands are present in directives, separate them with commas.

The following is an example of a directive statement:

```
.proc foo
```

Assignment Statements

Assignment statements enable the programmer to define or redefine a symbol by assigning it a value. This value may be a reference to another symbol, register name, or expression. The new value takes effect immediately and remains in effect until the symbol is redefined. Symbols defined in assignment statements do not have forward references.

In addition, symbols defined in assignment statements cannot:

- appear in the symbol table of an output object file.
- be declared global.
- be defined in an equate statement.

There are two types of assignment statements:

- Symbol assignment statements, which define or redefine a symbol in the symbol name space.
- Register assignment statements, which define or redefine a register name in the symbol name space.

Symbol Assignment Statements

A symbol assignment statement has the following syntax:

```
identifier=expression      // comments
```

Where:

identifier Represents a symbol in the symbol name space.
expression Specifies the type and value of the identifier. The expression cannot contain forward references.

The following is an example of an assignment statement that defines a symbol:

```
C = L0+2
```

Register Assignment Statements

A register assignment statement has the following syntax:

```
identifier=register name      // comments
```

Where:

identifier Represents a register name in the symbol name space.
register name Specifies an alternate register name. If the register name is a stack or rotating register name, the new register name continues to reference the previously-defined register name, even if the name is no longer in effect. See the "Register Stack Directive" section and "Rotating Register Directives" section in Chapter 6.

The following is an example of an assignment statement that defines a register name:

```
A = r1
```

Equate Statements

Equate statements enable the programmer to define a symbol by assigning it a value. This value may be a reference to another symbol, register name, or expression. In equate statements, a symbol can be defined only once throughout the source file. These symbols may have forward references, except when referencing a register name. A symbol name defined in an equate statement cannot be defined in an assignment statement.

Equate statements have the same syntax as assignment statements, except for the operator.

There are two types of equate statements:

- symbol equate statements
- register equate statements

Symbol Equate Statements

A symbol equate statement has the following syntax:

```
identifier==expression // comments
```

Where:

<i>identifier</i>	Represents a symbol in the symbol name space.
<i>expression</i>	Specifies the type and value of the identifier. The expression can contain forward references.

The following is an example of an equate statement that defines a symbol:

```
A == 5
```

Register Equate Statements

A register equate statement has the following syntax:

```
identifier==register name // comments
```

Where:

<i>identifier</i>	Represents a register name in the symbol name space.
-------------------	--

register name Specifies an alternate register name. The register name cannot contain forward references. If the register name is a stack or rotating register name, the new register name continues to refer to the previously-defined register, even if the name is no longer in effect. See the "Register Stack Directive" section and "Rotating Register Directives" section in Chapter 6.

The following is an example of an equate statement that defines a register name:

```
A == r1
```

Data Allocation Statements

A data allocation statement has the following syntax:

```
[label:] dataop operands // comments
```

Where:

label Defines a symbol whose value is the address of the first data object defined by the statement. If the assembler inserts padding to align the location counter to an implied alignment boundary, the label is assigned the value of the newly-aligned address.

The assembler interprets a label followed by a double-colon (: :) as a global symbol. See the "Symbol Scope Declaration" section in Chapter 4 for more information about global symbols.

dataop Defines the type and size of data objects that are assembled. Data object mnemonics are not case-sensitive. The "Data Allocation Statements" section in Chapter 5 lists the data object mnemonics.

operands Contain multiple expressions separated by commas. Each expression defines a separate data object of the same type and size. The assembler puts the data objects into consecutive locations in memory, and automatically aligns each to its natural boundary.

The following is an example of a data-allocation statement with a *label*:

```
L2: data4.ua L1, L1+7, .t1+0x34, $-15
```

Cross-section Data Allocation Statements

A cross-section data allocation statement has the following syntax:

```
xdataop section-name, operands //comments
```

Where:

xdataop Defines the type and size of data objects that are assembled. Cross-section data object mnemonics are not case-sensitive.

section-name Refers to a predefined name of an existing section in the object file.

operands Contain multiple expressions that are separated by commas. Each expression defines a separate data object of the same type and size. The assembler puts the data objects into consecutive locations in memory, and automatically aligns each to its natural boundary.

The following is an example of a cross-section data allocation statement:

```
.xdata8 .data, 0x123, L1
```


Program Structure

This chapter describes the overall structure of the IA-64 assembly-language program and explains how to:

- Define sections
- Include contents from other files in the current source file
- Bundle and group instructions
- Define procedures
- Generate symbolic debug information (Windows NT* specific)

Sections

The output object file of an assembly program is made up of named sections that contain code and data. The assembler allows any number of sections to be created in parallel within the output object file, one of which can be accessed at a time. The section currently accessed is referred to as the current section.

The assembler maintains a separate location counter for each existing section. The assembler always adds new code or data to the end of the current section, moving the location counter in that section ahead to incorporate the new code or data. The "Cross-section Data Allocation Statements" section in Chapter 5 explains how to add data to a section that is not the current section.

Section directives and predefined section directives are used to define and switch between sections. Some section directives have flag and type operands that specify the flag and type attributes of a section.

Section Flags and Section Type Operands

The *flags* operand specifies one or more flag attributes of a section. The *flags* operand is a string constant composed of one or more characters. [Table 3-1](#) lists the valid flag characters. The *flags* operand is case-sensitive. The assembler does not detect invalid specifications made by the programmer, such as stores to a section that is a non-writable section. A non-writable section is not flagged by the *w* flag character.

Table 3-1 Section Flag Characters

Flag Characters	Description
w	Write access allowed.
a	Section is allocated in memory.
x	Section contains executable instructions.
s	Section contains "short" data.

The *type* operand specifies a section's type attribute. The *type* operand is a string constant containing one of the valid section types listed in [Table 3-2](#). The section types listed in the table correspond directly to ELF (UNIX*) section types, except for the "comdat" section type, which corresponds to COFF32 (Windows NT). The *type* operand is case-sensitive.

Table 3-2 Section Types

Section Type	Description
"progbits"	Sections with initialized data or code.
"nobits"	Sections with uninitialized data (bss).
"comdat" (Windows NT specific)	COMDAT sections. See the "Windows NT (COFF32) Specific Section Flag Operands" section.
"note"	Note sections.

Windows NT (COFF32) Specific Section Flag Operands

In addition to the section flags described in the ["Section Flags and Section Type Operands" section](#), the assembler recognizes the flags listed in [Table 3-3](#) when the section type is "comdat" and the object file format is COFF32 (Windows NT).

These flags represent link-time selection criteria, and are case-sensitive.

Table 3-3 COMDAT Section Flag Characters

Flag	Description
D	Allow only one instance of this section.
Y	Select any one instance of this section.
S	Select any one instance of this section; all instances must be the same size.
E	Select any one instance of this section; all instances must have identical contents.
N	Select the newest instance of this section.
L	Select the largest instance of this section.
T	Section contains thread local storage (tls) data.
A	Select an instance of this section only if the associated section name is selected. See the "Associated Section Name Flag" section .

Associated Section Name Flag

When the A flag is present, the assembler identifies an associated section name. Use the A flag in conjunction with an associated section operand. The associated section operand is a section name. A section name can only be loaded in link time if the associated section is already loaded.

To select the A flag, use the `.section` or `.pushsection` directive with an additional *assoc-section* operand in one of the following formats:

```
.section section-name [,"flags","type" [,assoc-section]]
.section section-name = "flags","type" [,assoc-section]
.pushsection section-name [,"flags","type"
    [,assoc-section]]
.pushsection section-name = "flags","type" [,assoc-section]
```

Where:

<i>section-name</i>	Represents a user-defined name using any valid identifier. Section names are case-sensitive.
<i>flags</i>	Represents a string constant composed of one or more characters that specify the attributes of a section. See the Table 3-1 for a list of the valid flag characters.
<i>type</i>	Represents a string constant specifying a type attribute of a section. See Table 3-2 for a list of the section types
<i>assoc-section</i>	Represents a user-defined section name.

Section Definition Directive

The `.section` directive defines new sections, switches from one section to another, and sets the current section. The `.section` directive has the following formats, with a different functionality for each format:

```
.section      section-name
.section      section-name, "flags", "type"
.section      section-name = "flags", "type"
```

Where:

<i>section-name</i>	Represents a user-defined name using any valid identifier. Section names are case-sensitive.
<i>flags</i>	Represents a string constant composed of one or more characters that specify the attributes of a section. See Table 3-1 for a list of the valid flag characters.
<i>type</i>	Represents a string constant specifying a type attribute of a section. See Table 3-2 for a list of the section types.

In the first format, the `.section` directive sets the *section-name* as the current section.

In the second format, the `.section` directive defines a new section, assigns *flags* and *type* attributes, and makes the newly-defined section the current section. If the newly-defined section has the same name, *flag* attributes, and *type* attribute as a previously-defined existing section, the assembler switches to the previously-defined section without defining a

new one. For example, the following `.section` directive defines a new section (`my_section`), assigns `flags` ("aw") and `type` ("progbits") attributes, and makes it the current section.

```
.section    my_section, "aw", "progbits"
```

In the third format, the `.section` directive creates a new section with a previously-defined section name, and assigns it new `flags` and `type` attributes. The newly-created section becomes the current section; any reference to this section name refers to the newly-created section.

The ["Using Section Directives" section](#) illustrates how to use the `.section` directive.

Section Stack Directives

The assembler maintains a section stack, which is defined by the `.pushsection` and `.popsection` directives. These directives push and pop previously-defined sections to and from the section stack. The assembler may limit the depth of a section stack, but it must allow at least ten levels.

The `.pushsection` directive pushes the current section onto the stack and switches to the section specified in the directive. The `.pushsection` directive, like the `.section` directive, has one of the following formats:

```
.pushsection    section-name
.pushsection    section-name, "flags", "type"
.pushsection    section-name = "flags", "type"
```

Where:

<i>section-name</i>	Represents a user-defined name using any valid identifier. Section names are case-sensitive.
<i>flags</i>	Represents a string constant composed of one or more characters that specify the attributes of a section. See Table 3-1 for a list of the valid flag characters.
<i>type</i>	Represents a string constant specifying a type attribute of a section. See Table 3-2 for a list of the section types

The `.popsection` directive pops the previously-pushed section from the top of the stack, and makes it the current one.

The ["Using Section Directives" section](#) illustrates how to use the `.pushsection` and `.popsection` directives.

Absolute Sections

Absolute sections are only supported by ELF object file formats. To define an absolute section with a fixed starting address, use the `.section` and `.pushsection` directives with an optional *origin* operand. The *origin* operand must be an absolute expression. See the "Absolute Expressions" section in Chapter 2 for more information about absolute expressions. Absolute section addresses cannot overlap. The linker does not merge absolute sections with other section types, or with other absolute sections.

The following example defines a new section name and assigns it new *flags* and *type* attributes, with a starting address specified by the *origin* parameter.

```
.section    new_name, "aw", "progbits", 0x1000
```

Section Return Directive

The `.previous` directive returns to the previously-defined section of the current section and makes it the current section. This directive does not affect the section stack. The ["Using Section Directives" section](#) illustrates how to use this directive.

Predefined Section Directives

The predefined section directives define and switch between commonly-used sections. A predefined section directive creates a new section with the default *flags* and *type* attributes, and makes that section the current section.

The predefined section directive mnemonics are the same as the section names. The assembler generates section names in lower case, even though directive mnemonics are not case-sensitive.

On some platforms the assembler automatically creates a local symbol with a "section" type attribute for each defined section in the object file. See the "Symbol Type Directive" section in Chapter 4 for more information about symbol types.

The linker combines sections with the same name, *flags* and *type* attributes. The linker creates two separate output sections for sections with the same name, but different *flags* and *type* attributes.

To define a section without the default flags and type attributes, use the `.section` directive.

The predefined section directives cannot define a new section using the same name as a previously-defined section.

[Table 3-4](#) lists the predefined section directives, and their default *flags* and *type* attributes. A predefined section directive can have the same name as a section name.

Table 3-4 Predefined Section Directives

Directive/ Section Name	Flags	Type	Usage
<code>.text</code>	"ax"	"progbits"	Read-only object code.
<code>.data</code>	"wa"	"progbits"	Read-write initialized long data.
<code>.sdata</code>	"was"	"progbits"	Read-write initialized short data.
<code>.bss</code>	"wa"	"nobits"	Read-write uninitialized long data.
<code>.sbss</code>	"was"	"nobits"	Read-write uninitialized short data.
<code>.rodata</code>	"a"	"progbits"	Read-only long data (literals). ELF format only.
<code>.comment</code>	" "	"progbits"	Comments in the object file. ELF format, and COFF format only when used with the <code>-Q y</code> command-line option.

Using Section Directives

The following code illustrates the use and behavior of the section directives `.text`, `.section`, `.pushsection`, `.popsection`, and `.previous`:

Example 3-1 Code Sequence using Section Directives

```
.text          //Default
.section A     //Makes A the current section..text is A's previous
              //section
.pushsection B //Pushes A onto the stack and makes B the current
              //section. A is B's previous section
.pushsection C //Pushes B onto stack and makes C the current section,
              // B is C's previous section.
.popsection    //Pops B from stack and makes it current
.popsection    //Pops A from stack and makes it current..text is
              // A's previous section
.previous      //Makes A's previously current section .text
              // the current section. A becomes .text's
              // previous section
.previous      //Makes A the current section, .text becomes A's
              // previous section
```

Include File Directive

To include the contents of another file in the current source file, use the `.include` directive in the following format:

```
.include      "filename"
```

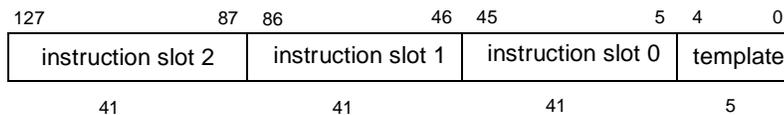
Where:

"*filename*" Specifies a string constant. If the specified filename is an absolute pathname, the file is included. If the specified filename is a relative pathname, the assembler performs a platform-dependent search to locate the include file.

Bundles

IA-64 architecture instructions are grouped together in 128-bit aligned containers called bundles. Each bundle contains three 41-bit instruction slots, and a 5-bit template field. The template field specifies which type of execution unit processes each instruction in the bundle. Bit 0 is set to 1 if there is a stop at the end of a bundle. There is no fixed relation between the boundaries of an instruction group and the boundaries of a bundle. Figure 3-1 illustrates the format of a bundle.

Figure 3-1 Bundle Format



Multiway branch bundles contain more than one branch instruction. When the first branch instruction of a multiway bundle is taken, the subsequent branch instruction does not execute.

Bundles are always aligned at 16-byte boundaries. The assembler automatically aligns sections containing bundles to at least 16-bytes.

Bundling can be:

- implicit (automatically performed by the assembler)
- explicit (specified by the programmer)
 - with automatic selection of the template
 - with explicit selection of the template

Refer to the *IA-64 Architecture Software Developer's Manual* for more details about bundles.

Implicit Bundling

The assembler bundles instructions automatically by default.

In the implicit-bundling mode, section directives do not terminate a partially-filled bundle of a previously-defined section. This means that the assembler can return to the previous section and continue to fill the bundle.

In implicit-bundling mode, a label forces the assembler to start a new bundle.

Explicit Bundling

The programmer can explicitly assemble bundles by grouping together up to three instructions, and enclosing them in braces ({}). The assembler places these instructions in one bundle, separate from all preceding and subsequent instructions. Stops at the end of an explicit bundle can be placed before or after the closing brace.

Section directives and data allocation statements cannot be used within an explicit bundle. Cross-section data allocation statements can be used within an explicit bundle. See the "Cross-section Data Allocation Statements" section in Chapter 5 for more information.

In explicit-bundling mode, labels can be inserted only as the first statement of an explicit bundle. Instruction tags can be applied to any instruction.

When using explicit-bundling, the appropriate template can be selected in one of the following ways:

- automatically by the assembler.
- explicitly by the programmer, using the explicit-template directives.

Auto-template Selection

By default, the assembler searches and selects a matching template for a bundle. The template fields specify intra-bundle instruction stops. When two templates consist of the same sequence of instruction types, they are

distinguished by stops. The assembler selects the appropriate template field based on the stops within the bundle. If no template is found, the assembler produces a diagnostic message. Instruction group stops may occur in a bundle.

Explicit-template Selection

To explicitly select a specific template, use one of the directives listed in [Table 3-5](#), as the first statement of your code within the braces. For example, the `.mii` directive selects the memory-integer-integer (`mii`) template.

Table 3-5 Explicit Template Selection Directives

Directive	Template Selection		
	Slot 0	Slot 1	Slot 2
<code>.mmi</code>	memory	integer	integer
<code>.mfi</code>	memory	floating-point	integer
<code>.bbb</code>	branch	branch	branch
<code>.mlx</code>	memory	long immediate	
<code>.mib</code>	memory	integer	branch
<code>.mmb</code>	memory	memory	branch
<code>.mmi</code>	memory	memory	integer
<code>.mhb</code>	memory	branch	branch
<code>.mfb</code>	memory	floating-point	branch
<code>.mmf</code>	memory	memory	floating-point

Refer to the *IA-64 Architecture Software Developer's Manual* for more information about template field encoding and instruction slot mapping.



NOTE. Select the `.mlx` directive for the move long immediate instruction. This instruction operates on 64-bit data types and is too large to fit into one of the 41-bit bundle slots. This directive selects the `mlx` template and inserts the instruction in slot 1 and slot 2 of the bundle.

Example 3-2 is code that shows an explicit bundle using explicit template selection, and a stop.

Example 3-2 Bundle with Explicit Template Selection and a Stop

```
{.mmi          //use the mmi template for this bundle
  m inst      //memory instruction
  ;;         //stop
  m inst      //memory instruction
  i inst      //integer instruction
}
```

Instruction Groups

IA-64 architecture instructions are organized in instruction groups. Each instruction group contains one or more statically contiguous instruction(s) that can execute in parallel. An instruction group must contain at least one instruction; there is no upper limit on the number of instructions in an instruction group.

An instruction group is terminated statically by a stop, and dynamically by taken branches. Stops are represented by a double semi-colon (; ;). The programmer can explicitly define stops. Stops immediately follow an instruction, or appear on a separate line. They can be inserted between two instructions on the same line.

Refer to the *IA-64 Architecture Software Developer's Manual* for more detailed information about instruction groups.

Dependency Violations and Assembly Modes

Dependency violations occur when instructions within an instruction group access the same resource register, including registers that appear as implicit operands. Dependency violations result in architecturally undefined behavior. The assembler can detect and eliminate dependency violations that occur within instruction groups, depending on its mode.

The assembler reads and processes assembly code in one of two modes: explicit and automatic.

Use explicit mode if you are an expert user with profound knowledge of IA-64 architecture or performance is important. In explicit mode, you are responsible for bundling and stops (`; ;`), and the assembler generates errors where it finds dependency violations.

Use automatic mode if you are a novice user or performance is not the highest consideration. In automatic mode, the assembler bundles the code and adds stops to avoid dependency violations. It ignores existing stops and annotations.

You can mix code from both modes in the one file. Set the mode using the command-line option or the directives `.auto` and `.explicit`. The directive `.default` causes the assembler to revert to the mode of operation defined in the command line.

For a complete description of the rules of data dependencies, see the *IA-64 Architecture Software Developer's Manual*.

This feature may not be currently supported by all assemblers.

Procedures

Software conventions require that instructions belong to a declared procedure, and that procedure prologues be separated from the main body within the procedure. These conventions ensure that the proper stack unwind information is placed in the object file. Refer to the *IA-64 Software Conventions and Runtime Architecture Guide* for details about the software conventions.

Procedure Directives

The `.proc` and `.endp` directives combine code belonging to the same procedure. The `.proc` directive marks the beginning of a procedure, and the `.endp` directive marks the end of a procedure. A single procedure may consist of several disjointed blocks of code. Each block should be individually bracketed with these directives. Name operands within a procedure can be used only for that specific procedure.

The `.proc` directive declares a symbol as a function. The `.proc` directive does not define the symbol by assigning it a value. Symbols must be defined as a label within the procedure. When `name` is defined, it is automatically assigned a "function" type.

The following code sequence shows the basic format of a procedure:

```
.proc  name, ...  
    name:      //label  
...           //instructions in procedure  
.endp  name, ...
```

Where:

`name` Represents one or more entry points of the procedure. Each entry point has a different name. The assembler ignores the `name` operands of the `.endp` directive.

Procedure Label (PLabel)

When the object file format is COFF32 (Windows NT), the assembler creates two symbols for a defined procedure. One symbol represents the procedure entry point and appears in the object file symbol table with the original symbol name preceded by a dot. For example, the label named `foo` becomes `.foo` in the object file symbol table. The other symbol represents the procedure label, also referred to as the function descriptor or PLabel, and is implicitly generated by the assembler using the original symbol name. Refer to the *IA-64 Software Conventions and Runtime Architecture Guide* for more information about the procedure label.

Stack Unwind Directives

Stack unwind directives are used to generate unwind information for a procedure.

The *IA-64 Software Conventions and Runtime Architecture Guide* describes stack unwind elements and their semantics. Refer to this document for information about the semantics of the stack unwind directives described in this section.

Procedures are bound by the `.proc` and `.endp` directives. See the ["Procedure Directives" section](#) for more information about these directives. Procedures are section-sensitive. The assembler interprets stack unwind directives according to the procedure in which they appear.

Procedures contain prologue and body regions that are divided by headers. These headers are specified using the `.prologue` and `.body` directives.

The `.prologue` directive introduces a prologue region within a procedure. Each prologue region must be introduced by the `.prologue` directive.

The `.body` directive separates the procedure prologue from the main body of the procedure. You can use the `.body` directive more than once within procedures with multiple body regions.

For language specific data, use the `.handlerdata` directive followed by handler data allocations with the `.endp` directive after the handler data allocations. The assembler places the handler data in the `.xdata` section. See the "Stack Unwind Directives Usage Guidelines" section on page 3-22 for more information about using this directive.

These directives may not be currently supported by all assemblers.

Example 3-3 illustrates the format of a procedure with two prologues, two body regions, and language specific data.

Example 3-3 Procedure Format in a Code Sequence

```
.proc      name,...      //start of procedure
```

Example 3-3 Procedure Format in a Code Sequence (continued)

```
.prologue                //instructions in first prologue
.body                   //instructions in first body region
.prologue               //instructions in second prologue
.body                   //instructions in second body region
.handlerdata           //data allocations go to .xdata section
.endp                  name,... //end of procedure
```

Stack unwind directives, except for the `.endp` directive, do not break bundles. When a `tag` operand is present in a stack unwind directive, the tag refers to a location of an instruction slot. If the `tag` is omitted, the location default is the location counter of the next instruction. More than one directive can refer to the same location of an instruction slot.

Generally, functions have unwind table entries. A stack unwind directive must be present between the `.proc` and `.endp` directives to write function entries and unwind information to the unwind table.

To create a function entry for unwind information when there is no stack unwind information, use the `.unwentry` directive.

[Table 3-6](#) lists the stack unwind directives and their operands. The right-most column of the table summarizes the records and fields that are affected by these directives. For more information about the affected records and fields, refer to the *IA-64 Software Conventions and Runtime Architecture Guide*.

Table 3-6 Stack Unwind Directives

Directive Name	First Operand	Second Operand	Third Operand	Affected Records and Fields
<code>.proc</code>	<code>symbol</code>			entry-start
<code>.endp</code>				entry-end
<code>.handlerdata</code>				handler data allocation
<code>.unwentry</code>				entry generation

Table 3-6 Stack Unwind Directives (continued)

Directive Name	First Operand	Second Operand	Third Operand	Affected Records and Fields
.prologue				prologue header previous header
.prologue	<i>imm-mask</i>	<i>grsave</i>		prologue header previous header
.body				body header previous header
.personality	<i>symbol</i>	[<i>phases</i>]		personality
.fframe	<i>size</i>	[<i>tag</i>]		mem_stack_f
.vframe	<i>gr-location</i>	[<i>tag</i>]		mem_stack_v psp_gr
.vframesp	<i>spoff</i>	[<i>tag</i>]		mem_stack_v psp_sprel
.vframepsp	<i>pspoff</i>	[<i>tag</i>]		mem_stak_v pso_psprel
.restore	<i>sp</i>	[<i>ecount</i>]	[<i>tag</i>]	epilogue
.copy_state	<i>state_no</i>			copy_state
.label_state	<i>state_no</i>			label_state
.save	<i>rp</i>	<i>gr-location</i>	[<i>tag</i>]	rp_when rp_gr
.altrp	<i>br-location</i>			rp_br
.savesp	<i>rp</i>	<i>imm-location</i>	[<i>tag</i>]	rp_when rp_sprel
.savepsp	<i>rp</i>	<i>imm-location</i>	[<i>tag</i>]	rp_when rp_psprel
.save	<i>ar.fpsr</i>	<i>gr_location</i>	[<i>tag</i>]	fpsr_when fpsr_gr
.savesp	<i>ar.fpsr</i>	<i>gr_location</i>	[<i>tag</i>]	fpsr_when fpsr_sprel

Table 3-6 Stack Unwind Directives (continued)

Directive Name	First Operand	Second Operand	Third Operand	Affected Records and Fields
<code>.savepsp</code>	<code>ar.fpsr</code>	<code>imm_location</code>	<code>[tag]</code>	<code>fpsr_when</code> <code>fpsr_ksprel</code>
<code>.save</code>	<code>ar.bsp</code>	<code>gr_location</code>	<code>[tag]</code>	<code>bsp_when</code> <code>bsp_gr</code>
<code>.savesp</code>	<code>ar.bsp</code>	<code>imm_location</code>	<code>[tag]</code>	<code>bsp_when</code> <code>bsp_sprel</code>
<code>.savepsp</code>	<code>ar.bsp</code>	<code>imm_location</code>	<code>[tag]</code>	<code>bsp_when</code> <code>bsp_ksprel</code>
<code>.save</code>	<code>ar.bspstore</code>	<code>gr_location</code>	<code>[tag]</code>	<code>bspstore_when</code> <code>bspstore_gr</code>
<code>.savesp</code>	<code>ar.bspstore</code>	<code>imm_location</code>	<code>[tag]</code>	<code>bspstore_when</code> <code>bspstore_sprel</code>
<code>.savepsp</code>	<code>ar.bspstore</code>	<code>imm_location</code>	<code>[tag]</code>	<code>bspstore_when</code> <code>bspstore_ksprel</code>
<code>.save</code>	<code>ar.rnat</code>	<code>gr_location</code>	<code>[tag]</code>	<code>rnat_when</code> <code>rnat_gr</code>
<code>.savesp</code>	<code>ar.rnat</code>	<code>imm_location</code>	<code>[tag]</code>	<code>rnat_when</code> <code>rnat_sprel</code>
<code>.savepsp</code>	<code>ar.rnat</code>	<code>imm_location</code>	<code>[tag]</code>	<code>rnat_when</code> <code>rnat_ksprel</code>
<code>.save</code>	<code>ar.pfs</code>	<code>gr_location</code>	<code>[tag]</code>	<code>pfs_when</code> <code>pfs_gr</code>
<code>.savesp</code>	<code>ar.pfs</code>	<code>imm_location</code>	<code>[tag]</code>	<code>pfs_when</code> <code>pfs_sprel</code>
<code>.savepsp</code>	<code>ar.pfs</code>	<code>imm_location</code>	<code>[tag]</code>	<code>pfs_when</code> <code>pfs_ksprel</code>
<code>.save</code>	<code>ar.unat</code>	<code>gr_location</code>	<code>[tag]</code>	<code>natcr_when</code> <code>natcr_gr</code>

Table 3-6 Stack Unwind Directives (continued)

Directive Name	First Operand	Second Operand	Third Operand	Affected Records and Fields
.savesp	ar.unat	imm-location	[tag]	natcr_when natcr_sprel
.savepsp	ar.unat	imm-location	[tag]	natcr_when natcr_psprel
.save	ar.lc	gr-location	[tag]	lc_when lc_gr
.savesp	ar.lc	imm-location	[tag]	lc_when lc_sprel
.savepsp	ar.lc	imm-location	[tag]	lc_when lc_psprel
.save	pr	gr-location	[tag]	preds_when preds_gr
.savesp	pr	imm-location	[tag]	preds_when preds_sprel
.savepsp	pr	imm-location	[tag]	preds_when preds_psprel
.save	@priunat	gr_location	[tag]	priunat_when priunat_gr
.savesp	@priunat	imm_location	[tag]	priunat_when priunat_sprel
.savepsp	@priunat	imm_location	[tag]	priunat_when priunat_psprel
.save.g	imm-grmask			gr_mem spill_imask
.save.g	imm_grmask	gr_location	[tag]	gr_gr imask
.save.f	imm-frmask			fr_mem spill_imask

Table 3-6 Stack Unwind Directives (continued)

Directive Name	First Operand	Second Operand	Third Operand	Affected Records and Fields
<code>.save.b</code>	<code>imm-brmask</code>			<code>br_mem</code> <code>spill_imask</code>
<code>.save.gf</code>	<code>imm-grmask</code>	<code>imm-frmask</code>		<code>frgr_mem</code> <code>spill_imask</code>
<code>.save.b</code>	<code>imm-brmask</code>	<code>gr-location</code>		<code>br_gr</code> <code>spill_imask</code>
<code>.spill</code>	<code>imm-location</code>			<code>spill_base</code>
<code>.unwabi</code>	<code>@svr4</code>	<code>imm_context</code>		<code>abi</code>
<code>.unwabi</code>	<code>@hpux</code>	<code>imm_context</code>		<code>abi</code>
<code>.unwabi</code>	<code>@nt</code>	<code>imm_context</code>		<code>abi</code>

The following alphabetical list defines the stack unwind directive operands listed in [Table 3-6](#):

- `br-location` is the alternative branch register used to get the return link. By default, `b0` is the return link.
- `ecount` is the number of prologues -1 specified by the assembler if this field is not specified by the user.
- `gr-location` is a general-purpose register that specifies the destination of the save operation. For example, registers `r1` and `loc1`.
- `grsave` saves the `rp`, `ar.pfs`, `psp`, and `pr` register contents to the first general-purpose register.
- `imm-location` (immediate location) is the offset between the `sp` or `psp`, and the `save_address`, specified in bytes. This offset is always positive and specified as follows:
- `imm-mask` (immediate mask) is an integer constant specifying a bit pattern for the preserved registers, as follows:
 - The immediate mask (`imm-mask`) of the `.prologue` directive is specified as follows: `rp` (return link) (bit 3), `ar.pfs` register (bit 2), `psp` (previous stack pointer) (bit 1), `pr` register (bit 0)

- The immediate mask (*imm-fmask*) of the `.save.f` and `.save.gf` directives refer to the preserved floating-point registers.
- The immediate mask (*imm-grmask*) of the `.save.g` and `.save.gf` directives refer to the preserved general registers.
- The immediate mask (*imm-brmask*) of the `.save.b` directive refers to the preserved branch registers.

sp_offset: imm-location = save_address - sp_address

psp_offset: imm-location = psp_address - save_address

- *phases* is the number of phases (0 to 3).
- `@priunat` is a predefined symbol and indicates a primary unit.
- *size* is the fixed frame size in bytes.
- `sp`, `rp`, `ar.pfs`, `ar.unat`, `ar.lc`, and `pr` are explicit register names.
- *state_no* is the state copied or restored.
- `@svr4`, `@hpux`, and `@nt` specify the operating system type.
- *symbol* is an assembly label.
- *tag* is an optional operand, which specifies a "when" attribute of the operation described by the directive.

Syntax for the `.save.x` Directives

The directives, `.save.f`, `.save.g`, `.save.gf`, and `.save.b`, define 2-bit fields for each save operation in the `imask` descriptor. The assembler interprets the instruction that immediately follows a save directive as a save instruction.

Example 3-4 illustrates the use of the `.save.g` directive. Each `.save.g` directive describes the subsequent store instruction. The operand is a mask where only one bit is set. This bit specifies the preserved saved register. The assembler produces a `gr_mem` descriptor with a `0x5` mask. In addition, the assembler marks the 2-bit fields of the `imask` descriptor, corresponding to the slots of the two store instructions.

Example 3-4 Code Sequence Using the .save.g Directive

```
.save.g 0x1  
st8... = r4  
...  
.save.g 0x4  
st8... = r6
```

Example 3-5 illustrates the use of the `.save.gf` directive. The `.save.gf` directive describes the subsequent store instruction. The operands is a mask where only one bit is set. This bit specifies the preserved saved register. The assembler produces a `frgr_mem` descriptor with a `0x42` mask for the floating-point registers and a `0x2` mask for the general-purpose registers. In addition, the assembler marks the 2-bit fields of the `imask` descriptor, corresponding to the slots of the three store instructions.

Example 3-5 Code Sequence Using the .save.gf Directive

```
.save.gf 0, 0x2  
fst... = f3  
...  
.save.gf 0, 0x40  
fst... = f18  
...  
.save.gf 0x2, 0  
st8... = r5
```

Stack Unwind Directives Usage Guidelines

Follow these guidelines when using the stack unwind directives:

- Place stack unwind directives between the unwind entry point of the function declared in `.proc` and `.endp`.
- The first directive in each region in a procedure must be one of the following region header directives, `.prologue` or `.body`.
- The first directive in the procedure must point to the same address as the first unwind entry point of the function.

- Region header directives should alternate. No two consecutive regions of the same type are allowed.
- When none of the stack unwind directives listed in [Table 3-6](#) are specified, optionally use the `.unwentry` directive to create an unwind entry for the function. Do not use this directive if the unwind records are filled by the compiler.
- Use tags only within the current region. A tag operand cannot be specified out of the scope region. If a tag is omitted, the directive refers to the next instruction, which resides in the same region.
- Use only one `.personality` directive at any point within each procedure.
- Always precede the `.handlerdata` directive with the `.personality` directive.
- Follow these guidelines for prologue regions:
 - Use one of the following frame directives: `.fframe`, `.vframe`, or `.vframesp`.
 - Use each of the `.save` directives only once. For example: `.save rp, ar.pfs, ar.unat, ar.lc, and pr`.
 - Multiple usage of the directives, `.save.g`, `.save.f`, `.save.b`, and `.save.gf` is allowed. The number of bits set in the bit-mask operand specifies the number of the consecutive save instructions that immediately follow the directive.
 - A single unwind record is built for one or more occurrences of the following directives: `.save.g`, `.save.f`, `.save.b`, and `.save.gf`. The bit-mask field of the record is a bitwise OR of all the masks that appear in the directives.
 - Use only one `.save.b` with the `gr-location` operand.
 - Use only one `.spill` directive.
- Use only one `.restore` directive for body regions.

Windows NT (COFF32) Symbolic Debug Directives

When the object file format is COFF32 (Windows NT), the symbolic debug directive `.ln` stores the line number table entry of a function in the symbolic debug information. The symbolic debug directive `.ln` must be enclosed within a function defined by the `.bf` and `.ef` directives. The `.bf` and `.ef` directives define the beginning and the end of a function.

The `.ln` directive has the following format:

```
.ln    line-number [,function]
```

Where:

line-number Specifies the source line number associated with the next assembled instruction.

function Is the name of the current function.

The `.bf` and `.ef` directives have the following format:

```
.bf    function,line
```

```
.ef    function,line,code-size
```

Where:

function Represents the function name.

line Is an integer number corresponding to the first source line of the function.

code-size Is an integer number representing line group code size, which is written as debug information.

Declarations

This chapter describes the IA-64 assembly-language directives associated with symbol declarations. These directives can be used to perform the following functions:

- Declare symbol scopes
- Specify symbol types
- Specify symbol sizes
- Override default file names
- Declare common symbols
- Declare aliases for labels, function names, symbolic constants, or sections

Symbol Scope Declaration

Symbols are declared as global, weak, or local scopes. Symbol scopes are used to resolve symbol references within one object file or between multiple object files. The symbol scope attribute is placed in the object file symbol table and any reference to a symbol is resolved in link time. By default, symbols have a local scope, where they are available only to the current assembly-language source file in which they are defined.

Local Scope Declaration Directive

References to symbols with a local scope are resolved from within the object file in which the symbols are declared. Local symbols with the same name in different object files do not refer to the same entity. Symbols have a local scope by default, so it is not necessary to declare symbols with local scopes. However, the `.local` directive is available for completeness. The `.local` directive has the following format:

```
.local      name,name, ...
```

Where:

`name` Represents a symbol name.

Global Scope Declaration Directive

References to symbols with a global scope are resolved within the object file in which the symbols are declared, and within other object files. Global symbols with the same name in different object files refer to the same entity.

To declare one or more symbols with a global scope, use the `.global` directive. These symbols are flagged as global symbols for the linkage editor. The `.global` directive has the following format:

```
.global      name,name, ...
```

Where:

`name` Represents a symbol name.

Weak Scope Declaration Directive

References to symbols with a weak scope are resolved within the object file in which the symbols are declared, and within other object files. Weak symbols with the same name in different object files may not refer to the same entity. When a symbol name is declared with a weak scope as well as a global or local scope, the global or local scope will take precedence over the weak scope in link time.

To declare one or more symbols with a weak scope, use the `.weak` directive. These symbols are flagged as weak symbols for the linkage editor. The weak scope declaration format for UNIX* (ELF) and Windows NT (COFF32) differ and are described in the sections that follow.

Weak Scope Declaration for UNIX (ELF)

For UNIX (ELF), use the `.weak` directive in the following format:

```
.weak    name1, name2, ...
```

Where:

name Represents a symbol name.

The following example illustrates how to declare an undefined symbol with a weak scope. The defined symbol `x`: has a local scope. `y` has the attributes of `x` and has a local scope. The symbol `y` can then be declared with a weak scope using the `.weak` directive while keeping the other attributes of `x`.

```
x:  
    y == x  
    .weak y
```

Weak Scope Declaration for Windows NT (COFF32)

For Windows NT (COFF32), use the `.weak` directive in the following format to declare a symbol with a weak scope and search for defined symbols within other object files and libraries:

```
.weak    identifier1 = identifier2
```

identifier1 Represents a symbol name that is assigned a weak symbol scope, which is resolved in link time.

identifier2 Represents a symbol name that holds the symbol definition.

Use the following syntax to declare a symbol with a weak scope and search for defined symbols within other object files and **not** within libraries:

```
.weak    identifier1 == identifier2
```

Where:

- identifier1* Represents a symbol name that is assigned a weak symbol scope, which is resolved in link time.
- identifier2* Represents a symbol name that holds the symbol definition.

The following example illustrates a weak scope declaration where *x*: is a local defined symbol. *x* is the associated symbol for *y*. The `.weak` directive assigns *y* a weak scope.

```
x:
.weak y = x
```

Symbol Type Directive

The default type of a symbol in an object file is based on the assembly-time type of the symbol. See Table 4-1 for a list of the symbol types and their predefined names. To explicitly specify a symbol's type, use the `.type` directive in the following format:

```
.type      name, type
```

Where:

- name* Represents a symbol name.
- type* Specifies the symbol type using one of the predefined symbols listed in Table 4-1.

Table 4-1 Symbol Types

Symbol Types	Predefined Symbol Name of Type
Symbolic constants and undefined symbols	@notype
Labels and common symbols	@object
Function names	@function
Section names	Created by the assembler.



NOTE. *The assembler automatically creates a symbol of type name for section names.*

When the object file format is COFF32 (Windows NT) the assembler creates a function symbol name for @function. For more information see the "Procedure Label (PLabel)" section in Chapter 3.

Symbol Size Directive

To explicitly specify the size attribute of a symbol, use the `.size` directive.

The `.size` directive has the following format:

```
.size    name, size
```

Where:

<code>name</code>	Represents a symbol name.
<code>size</code>	Represents an absolute integer expression with no forward references.

To implicitly specify the default size attribute of a symbol, use a data allocation statement. The default symbol size is written to the symbol table. See the "Data Allocation Statements" section in Chapter 5 for more information.



NOTE. *When the object file format is COFF32 (Windows NT), the `.size` directive is only effective for common symbols.*

File Name Override Directive

By default, the file name is the name of the source file. To override the default file name use the `.file` directive. If you use the `.file` directive more than once in a source file, the assembler places multiple file names in the output object file. The `.file` directive has the following format:

```
.file      "name"
```

Where:

"name" Represents a string constant specifying a source file name.

Common Symbol Declarations

Common and local common symbol declarations enable you to define a symbol with the same name in different object files. The difference between a common symbol and local common symbol is as follows:

- The linker merges two or more common symbol declarations for the same symbol.
- The assembler merges two or more local common symbol declarations for the same symbol.

If a symbol is declared as both common and local common, the common declaration overrides the local common declaration. Any definition of a symbol supersedes either type of common declaration.

Common Symbol Directive

To declare a symbol as a common symbol, use the `.common` directive. Common symbols have a global scope, and do not necessarily have the same size and alignment attributes. The `.common` directive has the following format:

```
.common      name,size,alignment
```

Where:

<i>name</i>	Represents a symbol name.
<i>size</i>	Represents an absolute integer expression.
<i>alignment</i>	Represents an absolute integer expression to the power of two. Not supported in COFF32 format.



NOTE. *When the object file format is COFF32 (Windows NT), the *alignment* operand is not supported.*

Local Common Symbol Directive

To declare a symbol as a local common symbol use the `.lcomm` directive. The `.lcomm` directive has the following format:

```
.lcomm    name, size, alignment
```

Where:

<i>name</i>	Represents a symbol name.
<i>size</i>	Represents an absolute integer expression.
<i>alignment</i>	Represents an absolute integer expression to the power of two.

The assembler allocates storage in the `.bss` or `.sbss` sections for undefined symbols declared as local common. The `.bss` or `.sbss` sections are chosen according to the size of the local common symbol. The assembler defines the symbol with the relocatable address of the allocated storage. The symbol is declared with a local scope, and assigned the largest size and alignment attributes of the local common declarations for that symbol.

Alias Declaration Directives

The `.alias` directive declares an alias for a label, a function name, or a symbolic constant. This directive can be used to reference an external symbol whose name is not legal in the assembly language. The `.alias` directive has the following format:

```
.alias    symbol, "alias"
```

Where:

`symbol` Represents a symbol name that the assembler can recognize. This name must be a valid name for the type of symbol.

`"alias"` Represents a string constant, which is the name the assembler exports to the object file symbol table.

The `.secalias` directive declares an alias for a section name. This directive can be used to reference an external section whose name is not legal in the assembly language. The `.secalias` directive has the following format:

```
.secalias    section-name, "alias"
```

Where:

`section-name` Represents a section name that the assembler can recognize. This name must be a valid name for the type of section.

`"alias"` Represents a string constant, which is the name the assembler exports to the object file symbol table.

Data Allocation

This chapter describes the IA-64 assembly language statements used to allocate initialized and uninitialized space for data objects in current sections and in cross sections, and to align data objects in sections of the code.

Data Allocation Statements

Data allocation statements allocate space for data objects in the current section, and initialize the space by assigning it a value. Data objects can be integer numbers, floating-point numbers, or strings. Integer numbers and floating point numbers are aligned according to their size. A data allocation statement with a *label*, defines a symbol of @object type, and sets the size attribute for that symbol.

Data allocation statements have any of the following formats:

```
[label:] data1 expression, ...
[label:] data2 expression, ...
[label:] data4 expression, ...
[label:] data8 expression, ...
[label:] real4 expression, ...
[label:] real8 expression, ...
[label:] real10 expression, ...
[label:] real16 expression, ...
[label:] string "string", ...
```

```
[label: ] stringz "string", ...
```

Where:

<i>label</i>	Specifies the data allocation address of the first data object.
<i>expression</i>	Represents any of the valid expression types listed in Table 5-1. Data allocation statements can have more than one expression operand.
<i>string</i>	Represents any of the valid string expression type values listed in Table 5-1.

Table 5-1 summarizes the data allocation mnemonics, and their expression type, memory format, data-object size, and alignment boundary for each.

Table 5-1 Data Allocation Statements

Mnemonic	Expression Type	Memory Format	Size (in bytes)	Alignment
data1	Integer	Integer	1	1
data2	Integer	Integer	2	2
data4	Integer	Integer	4	4
data8	Integer	Integer	8	8
real4	Floating-point or integer	IEEE single precision floating-point	4	4
real8	Floating-point or integer	IEEE double precision floating-point	8	8
real10	Floating-point or integer	IEEE extended precision floating-point (80-bit)	10	16
real16	Floating-point or integer	IEEE extended precision floating-point (80-bit)	16	16
string	String constant	Array of ASCII characters	Length of string	1
stringz	String constant	Array of ASCII characters, with null terminator	Length of string + 1	1

To disable the automatic alignment of data objects in data allocation statements, add the `.ua` completer after the mnemonic, for example, `data4.ua`. These statements allocate unaligned data objects at the current location within the current section.

The default byte order for data allocation statements is platform dependant. To specify the byte order for data allocation statements, use the `.msb`, or `.lsb` directives described in the "Byte Order Specification Directives" section in Chapter 6.

Uninitialized Space Allocation

The `.skip` and `.org` statements reserve uninitialized space in a section without assigning it a value.

The `.skip` and `.org` statements enable the assembler to reserve space in any section type, including a "nobits" section. During program execution, the contents of a "nobits" section are initialized as zero by the operating system program loader. When using the `.skip` and `.org` statements in any other section type, the assembler initializes the reserved space with zeros.

The `.skip` statement reserves a block of space in the current section. The size of the block is specified in bytes, and is determined by an *expression* operand. The *expression* operand specifies the size of space reserved in the current section. The `.skip` statement with a *label*, defines a symbol of @object type, and sets the size attribute for that symbol.

The `.skip` statement has the following format:

```
[label:] .skip expression
```

Where:

label Specifies the data allocation address of the beginning of the reserved block.

expression Represents an absolute integer expression with no forward references. The location counter advances to a location relative to the current location within the section. This operand cannot have a negative value since the location counter cannot be reversed.

The `.org` statement reserves a block of space in the current section. The `.org` statement advances the location counter to the location specified by the *expression* operand. The `.org` statement with a *label* defines a symbol of `@object` type, and sets the size attribute for that symbol.

The `.org` statement has the following format:

```
[label:] .org expression
```

Where:

label Specifies the data allocation address of the beginning of the reserved block.

expression Represents an integer, or a relocatable expression, with no forward references. If the expression is relocatable, it must be reducible to the form $R+K$, where R is a symbol previously defined in the current section, and K is an absolute constant. The location counter is set to the indicated offset relative to the beginning of the section. Since the location counter cannot be reversed, this operand must be greater than, or equal to, the current location counter.

Alignment

Instructions and data objects are aligned on natural alignment boundaries within a section. To disable automatic alignment of data objects in data allocation statements, add the `.ua` completer after the data allocation mnemonic, for example, `data4.ua`. Bundles are aligned at 16-byte boundaries, and data objects are aligned according to their size. The assembler does not align string data, since they are byte arrays.

Each section has an alignment attribute, which is determined by the largest aligned object within the section.

Section location counters are not aligned automatically. To align the location counter in the current section to a specified alignment boundary use the `.align` statement. The `.align` statement has the following format:

```
.align    expression
```

Where:

expression Is an integer number that specifies the alignment boundary of the location counter in the current section. The integer must be a power of two.

The `.align` statement enables the assembler to reserve space in any section type, including a "nobits" section. During program execution time the contents of a "nobits" section are initialized as zero by the operating system program loader. When using the `.align` statement in any other section type, the assembler initializes the reserved space with zeros for non-executable sections, and with a NOP pattern for executable sections.



NOTE. *When the object file format is COFF32 (Windows NT) the section alignment boundary is limited to 8KB. The assembler does not guarantee alignment for requests above 8KB.*

Cross-section Data Allocation Statements

Cross-section data allocation statements add data to a section that is not the current section. These statements save the overhead of switching between sections using the `.section` directive. See the "Sections" section in Chapter 3 for more information about switching between sections. Cross-section data allocation statements may be used within an explicit

bundle. All data objects are aligned to their natural boundaries in the cross section. Cross-section data allocation statements have any of the following formats:

```
.xdata1    section,expression, ...
.xdata2    section,expression, ...
.xdata4    section,expression, ...
.xdata8    section,expression, ...
.xstring   section,"string", ...
.xstringz  section,"string", ...
```

Where:

<i>section</i>	Represents the name of a previously-defined section that is not the current section.
<i>expression</i>	Represents an absolute or relocatable integer expression. When these expressions reference a location counter, they refer to the location counter within the cross section, not within the current section.
<i>string</i>	Represents any of the valid string expression type values listed in Table 5-1.

To disable automatic alignment of data objects in a cross-section data allocation statement, add the `.ua` completer to the statement, for example, `.xdata4.ua`. These statements allocate unaligned data objects at the current location counter of the cross section, not the current section.

The default byte order for cross-section data allocation statements is platform dependent. The byte order is determined by the cross section, not by the current section.

Miscellaneous Directives

6

This chapter describes the following IA-64 assembly-language directives:

- Register stack directive
- Rotating register directives
- Byte-order specification directive
- Ident string specification directive
- Radix indicator directive
- Preprocessor support

Register Stack Directive

The IA-64 architecture provides a mechanism for register renaming. Register renaming is implemented by allocating a register stack frame consisting of input, local, and output registers. These registers can be renamed. These renamable registers map to the general registers `r32` through `r127`. The assembler provides predefined alternate register names for the input, local, and output register areas of the register stack frame. The mapping of these registers to the general registers is determined by the nearest preceding `alloc` instruction.

Refer to the *IA-64 Architecture Software Developer's Manual* for detailed information about register renaming and for a full description of the `alloc` instruction.

The `.regstk` directive replaces the default register mappings defined by a preceding `alloc` instruction with new mappings. The `.regstk` directive does not allocate a new register stack frame.

The `.regstk` directive has the following format:

```
.regstk    ins, locals, outs, rotators
```

Where:

<i>ins</i>	Represents the number of input registers in the general register stack frame. <i>in</i> ₀ through <i>in</i> _{<i>ins</i>-1} represent <i>r</i> ₃₂ through <i>r</i> _{31+<i>ins</i>} for <i>ins</i> > 0.
<i>locals</i>	Represents the number of local registers in the general register stack frame state. <i>loc</i> ₀ through <i>loc</i> _{<i>locals</i>-1} represent <i>r</i> _{32+<i>ins</i>} through <i>r</i> _{31+<i>ins</i>+<i>locals</i>} for <i>locals</i> > 0.
<i>outs</i>	Represents the number of output registers in the general register stack frame. <i>out</i> ₀ through <i>out</i> _{<i>outs</i>-1} represent <i>r</i> _{32+<i>ins</i>+<i>locs</i>} through <i>r</i> _{31+<i>ins</i>+<i>locals</i>+<i>outs</i>} for <i>outs</i> > 0.
<i>rotators</i>	Represents the number of rotating registers in the general register frame. <i>rotators</i> must be ≤ <i>ins</i> + <i>locals</i> + <i>outs</i> .

The *in*, *loc*, and *out* register names defined by a previous `.regstk` directive or `alloc` instruction are visible by all subsequent instructions until the next `.regstk` directive or `alloc` instruction is specified.

The alternate register names specified by the operands of the `.regstk` directive refer to registers in the current register stack frame. If you reference input, local, or output registers using the alternate register names that are not within the current stack frame, the assembler produces an error message.

To prevent referencing the alternate register names, use the `.regstk` directive without the operands. The operands of a subsequent `.regstk` directive or `alloc` instruction redefine the mappings of the alternate register names.

The `alloc` instruction and `.regstk` directive do not affect the names of the general registers, `r32` through `r127`.

Stacked Registers in Assignment and Equate Statements

To define an alternate register name for a stacked register, use an assignment statement. The alternate register name is not affected by any subsequent changes to the rotating register. See the "Assignment Statements" and "Equate Statements" sections in Chapter 2 for more details about assignment and equate statements.

Example 6-1 illustrates how to define an alternate register name using an assignment statement, so that the alternate register name is not affected by a subsequent `.regstk` directive. The local register name `loc0` maps to the general register `r36`. `loc0` is assigned to `tmp`. The subsequent `add` instruction refers to `loc0`, which is currently mapped to `r40`. The next `add` instruction refers to `tmp` which is mapped to `r36`, not `r40`.

Example 6-1 Defining a Stacked Register in an Assignment Statement

```
tmp    .regstk    4,4,2,0
      = loc0                                //loc0 is currently r36
      ...
      .regstk    8,1,3,0

      add       loc0 = r1,r7                //loc0 is currently r40
      add       r1 = r2,tmp                 // tmp = r36!
```

Rotating Register Directives

General registers, floating-point registers, and predicate registers contain a subset of rotating registers. This subset of rotating registers can be renamed.

The following directives enable the programmer to provide names for one or more registers within each rotating register region:

- `.rotr` for general registers
- `.rotf` for floating-point registers
- `.rotp` for predicate registers

The `.rotx` directives assign alternate names and generation numbers for the rotating registers. One generation corresponds to one iteration of a software-pipelined loop. Each copied register is numbered with an index, where the most recent copy of a register has a zero index, such as `b[0]`. For every loop iteration, the registers within the group are renamed, and become one generation older by incrementing the index by one.

The `.rotx` directives define the number of instances of each pipeline variable and allocate them in the appropriate rotating register region. You can use an arbitrary name with a subscript-like notation for referencing the current and previous generations of each variable.

The rotating register directives have the following format:

```
.rotr    name [expression], ...  
.rotf    name [expression], ...  
.rotp    name [expression], ...
```

Where:

<i>name</i>	Represents a register name specified by the user, and represents a pipelined variable.
<i>expression</i>	Specifies the number of generations needed for the variable. The <i>expression</i> must be an absolute integer expression with no forward references.

When the alias rotating register names are used as instruction operands, they have the following format:

```
name[expression]
```

Where:

<i>name</i>	Represents an alias rotating register name defined by one of the rotating register directives.
<i>expression</i>	Represents an absolute integer expression with no forward references. The index must be between 0 and $(n-1)$, where n is the number of generations defined for that name. If the index is negative, or greater than $(n-1)$, the assembler produces an error message.

The `.rotr`, `.rotf`, and `.rotp` directives cancel all previous alias names associated with the appropriate register file, before defining new register names. The register files include the general, floating-point, and predicate registers.

If the number of rotating general registers implied by a `.rotr` directive exceeds the number of rotating registers declared by the nearest preceding `alloc` instruction, or `.regstk` directive, the assembler issues a warning.

Using Rotating Register directives

Example 6-2 and Example 6-3 illustrate the behaviour of the `.rotp` and `.rotf` directives, respectively.

Example 6-2 illustrates how the `.rotp` directive declares alternate rotating predicate register names for two predicate registers, `p[2]`, and three predicate registers `q[3]`. Instructions subsequent to the `.rotp` directive refer to `p[0]` for the current generation of `p`, and `p[1]` for the previous generation of `p`. For the current generation of `q`, the subsequent instructions refer to `q[0]`, `q[1]` for the previous generation, and `q[2]` for the one before the previous generation.

Example 6-2 Using the `.rotp` Directive

```
.rotp p[2],q[3]
```

Example 6-2 Using the .rotp Directive (continued)

The alternate predicate register names map to the predicate registers as follows:

```
p[0] = p16; p[1] = p17
q[0] = p18; q[1] = p19; q[2] = p20
```

Example 6-3 illustrates how the `.rotf` directive declares alternate floating-point register names for three floating-point registers `x[3]`, two floating-point registers `y[2]`, and three floating-point registers `z[3]`.

Example 6-3 Using the .rotf Directive

```
.rotf x[3],y[2],z[3]
```

The alternate floating-point register names map to the floating-point registers as follows:

```
x[0]=f32;x[1]=f33;x[2]=f34
y[0]=f35;y[1]=f36
z[0]=f37;z[1]=f38;z[2]=f39
```

Rotating Registers in Assignment and Equate Statements

To define an alias name for a rotating register, use an assignment statement. The alias register name is not affected by any subsequent changes to the rotating register. See the "Assignment Statements" and "Equate Statements" sections in Chapter 2 for more details about assignment and equate statements.

Example 6-4 illustrates how to define an alias name using an assignment statement so that the alias name is not affected by a subsequent `.rotr` directive. The `.rotr` directive maps `b[1]` to general register `r36`. `b[1]` is assigned to `tmp`. The second `.rotr` directive defines the new mapping of `b[1]` to `r33`. The subsequent `add` instruction that refers to `b[1]` is currently mapped to `r33`. The second `add` instruction refers to `tmp`, which is mapped to `r36`, not `r33`.

Example 6-4 Defining an Alias Name in an Assignment Statement

```
.rotr    a[3],b[2],c[4]
tmp = b[1]                                //b[1] is currently r36
...
.rotr    b[4],c[3],d[2]
add      b[1] = r1,r7                      //b[1] is currently r33
add      r1 = r2,tmp                       // tmp = r36!
```

Byte Order Specification Directives

The `.msb` and `.lsb` directives determine the byte order of data assembled by the `data n` , `real n` , and `.xdata n` data allocation statements. The values of n for `data` and `.xdata` are 1, 2, 4, and 8. The values of n for `real` are 4, 8, 10, and 16. See Chapter 5 for more information about data allocation statements.

The `.msb` and `.lsb` directives change the byte order for current sections only. They do not affect the instructions that are assembled. They only affect the data created. The default byte order is little-endian.

The `.msb` directive switches to MSB, where the **most-significant byte** is stored at the lowest address (big-endian). The `.lsb` directive switches to LSB, where the **least-significant byte** is stored at the lowest address (little-endian).

The byte order is a property of each section. If the byte order is changed in one section, it remains in effect for that section until the byte order is redefined. This change does not affect the byte order of other sections in the assembly program.

String Specification Directive

The `.ident` directive places a null terminated string in the `.comment` section of an output object file. See the use of `.comment` in [Chapter 3, “Program Structure”](#). The `.ident` directive has the following format:

```
.ident    "string"
```

Where:

"string" Represents a string.

Radix Indicator Directive

The `.radix` directive selects the numeric constant style.

To select a MASM numeric constant and specify a radix indicator, use the `.radix` directive in the following format:

```
.radix [radix-indicator]
```

Where:

radix-indicator

Indicates a MASM (Microsoft* macro assembler) numeric constant and specifies the radix. See Table 2-6 in Chapter 2, for a list of the radix indicators.

The MASM numeric constant and radix remain in effect until redefined.

To select a C numeric constant, use the `.radix` directive in the following format:

```
.radix [C]
```

Where:

c Indicates a C numeric constant.

The `.radix` directive used with an operand, pushes the previous numeric constant style and radix onto a radix stack. The `.radix` directive without the *radix-indicator* operand, pops and restores the previous style and radix from the stack. The assembler may limit the depth of a radix stack, but this limit must be no less than 10 levels.

Preprocessor Support

The assembler recognizes a special filename and the line number directive (`#line`) inserted by the standard C preprocessor, and sets its record as the current filename and line number accordingly. The `#line` directive has the following format:

`#line` *line_number, filename*

Where:

line_number Specifies the source line number

filename Identifies the name of the current filename.

Additionally, the assembler supports the following built-in symbols:

@line Current line number

@filename Current filename

@filepath Current file path

Annotations

Annotations are a subset of the assembler directives. They explicitly provide additional information for the assembler during the assembly process. These annotations have the same format and syntax as all other directives. This chapter describes these annotations and their functionality.

The annotations covered in this chapter include

- `.pred.rel`
- `.pred.vector`
- `.mem.offset`
- `.entry`

Predicate Relationship Annotation

The predicate relationship annotation `.pred.rel` provides information for the assembler about a logical relationship between the values of predicate registers. It is relevant only for explicit code.

The annotation `.pred.rel` takes the following forms:

"mutex"	mutual exclusion
"imply"	implication
"clear"	clear existing relations

When conflicting instructions follow an entry point, IAS ignores all existing predicate relationships defined before the entry point.

Predicate Vector Annotation

The predicate vector annotation `.pred.vector` explicitly specifies the predicate register contents using a user-defined value. The user-defined value is represented by a 64-bit binary number and each bit corresponds to a predicate register, respectively. A second optional operand can be used as a mask to selectively set only some of the predicate registers. Currently this annotation is ignored by the IA-64 Assembler.

This annotation takes effect at the point of insertion and the assembler may use this information for further analysis. The `.pred.vector` annotation has the following syntax:

```
.pred.vector  val [,mask]
```

Where:

<i>val</i>	Specifies a number represented as a 64-bit binary number. Each bit represents a 1-bit value in each of the corresponding 64 predicate registers. If <i>val</i> is not within the 64-bit range, this annotation is ignored.
<i>mask</i>	Represents an optional mask value used to define a subset of the predicate register file.

Example 7-1 illustrates a predicate vector annotation that sets the predicate registers according the specified value `0x9`, and uses a mask of `0xffff` to define a subset of the predicate register file.

Example 7-1 Using a Predicate Vector Annotation with a Mask

```
.pred.vector 0x9, 0xffff    //only refers to lowest 16-bits that  
                           are set in the mask. Values of p0-p15  
                           are defined.
```

Memory Offset Annotation

The memory offset annotation `.mem.offset` provides hints about the address that memory operations address, when the exact address is unknown. The annotation is useful for avoiding false reports of dependency violations. The annotation affects the instruction that follows.

The `.mem.offset` annotation has the following syntax:

```
.mem.offset    off_val,base_ind
```

Where:

off_val The relative offset for the memory region where the data is stored or retrieved.

base_ind A number that identifies the memory region where the information is stored or retrieved. The number is an arbitrary method of distinguishing between different memory regions.

Example 7-2 illustrates a `.mem.offset` annotation.

Example 7-2 Using the Memory Offset Annotation

```
.proc foo
foo::
FOO_STACK_INDEX=0
...
.mem.offset 0,FOO_STACK_INDEX    //code...
st8.spill [r3]=r32,8             //Suppose r3 contains the stack pointer
.mem.offset 8,FOO_STACK_INDEX    //We want to save r32-r34
st8.spill [r3]=r33,8
.mem.offset 16,FOO_STACK_INDEX
st8.spill [r3]=r34,8

.endp
.proc bar
bar::
.BAR_STACK_INDEX=1
...
.mem.offset 0,BAR_STACK_INDEX    //code...
st8.spill [r3]=r40               //Suppose r3 contains the stack pointer
                                 //We want to save r40
```

Entry Annotation

The entry annotation `.entry` notifies the assembler that a label can be entered from another function. By default, only global labels, designated by `<label>::`, are considered entry points. The annotation and the label need not be consecutive.

The `.entry` annotation has the following syntax:

```
.entry      label [, labels...]
```

Where:

`label` Represents the associated label.

Example 7-3 Using the Entry Annotation

```
.entry A                    //entry annotation  
A: mov r1=r2
```

Register Names by Type

Table A-1 through Table A-8 list the IA-64 architecture registers and their names.

Table A-1 **General Registers**

Register	Register Name
Fixed general registers	r0 - r31
Stacked general registers	r32 - r127
Alternate names for input registers	in0 - in95
Alternate names for local registers	loc0 - loc95
Alternate names for output registers	out0 - out95
Global pointer (r1)	gp
Return value registers (r8-r11)	ret0 - ret3
Stack pointer (r12)	sp

Table A-2 **Floating-point Registers**

Register	Register Name
Floating-point registers	f0 - f127
Argument registers (f8-f15)	farg0 - farg7
Return value registers (f8-f15)	fret0 - fret7

Table A-3 Predicate Registers

Register	Register Name
Predicates	p0 - p63
All predicates	pr
Rotating predicates	pr.rot

Table A-4 Branch Registers

Register	Register Name
Branch registers	b0 - b7
Return pointer (b0)	rp

Table A-5 Application Registers

Register	Register Number	Register Name
Application registers by number	0 - 127	ar0 - ar127
Kernel registers	0 - 7	ar.k0 - ar.k7
RSE control register	16	ar.rsc
Backing store pointer	17	ar.bsp
Backing store "store" pointer	18	ar.bspstore
RSE NaT collection register	19	ar.rnat
Compare & Exchange comparison value	32	ar.ccv
User NaT collection register	36	ar.unat
Floating-point status register	40	ar.fpsr
Interval time counter	44	ar.itc
Previous frame state	64	ar.pfs
Loop counter	65	ar.lc
Epilog counter	66	ar.ec

Table A-6 Control Registers

Register	Register Number	Register Name
Control registers by number	0 - 127	cr0 - cr127
Default control register	0	cr.dcr
Interval time match	1	cr.itm
Interrupt vector address	2	cr.iva
Page table address	8	cr.pta
Interrupt processor status register	16	cr.ipsr
Interrupt status register	17	cr.isr
Interrupt data address	18	cr.ida
Interrupt instruction pointer	19	cr.iip
Interrupt data translation register	20	cr.idtr
Interrupt instruction translation register	21	cr.iitr
Interrupt instruction previous address	22	cr.iipa
Interrupt frame state	23	cr.ifs
Interrupt immediate	24	cr.iim
Interrupt hash address	25	cr.iha
External interrupt registers	66	cr.lid
	71	cr.ivr
	72	cr.tpr
	75	cr.eoi
	96,98,100,102	cr.irr0 - cr.irr3
	114	cr.itv
	116	cr.pmv
	117 - 118	cr.lrr0 - cr.lrr1
	119	cr.cmcv

Table A-7 Other Registers

Register	Register Name
Processor status register	psr
Processor status register, lower 32 bits	psr.l
User mask	psr.um
Instruction pointer	ip

Table A-8 Indirect-register Files

Register	Register Name
Performance monitor control registers	pmc[r]
Performance monitor data registers	pmd[r]
Protection key registers	pkc[r]
Region registers	rr[r]
Instruction breakpoint registers	ibr[r]
Data breakpoint registers	dbp[r]
Instruction translation registers	itr[r]
Data translation registers	dtr[r]
Processor identification register	CPUID[r]

Pseudo-ops

Table B-1 lists the assembly language pseudo-ops for the IA-64 architecture according to their opcodes. Table B-2 lists pseudo-ops with missing operands.

The opcodes are listed alphabetically, with their operands, and the equivalent machine instructions. The table lists mnemonics converted to other mnemonics.

Table B-1 Pseudo-ops Listed by Opcode

Opcode	Instruction Description	Operands	Equivalent Machine Instruction
add	Add immediate	$r1 =imm, r3$	adds $r1 =imm14, r3$ addl $r1 =imm22, r3$
break	Break	$imm21$	break.b $imm21$ (B) break.i $imm21$ (I) break.m $imm21$ (M) break.f $imm21$ (F)
chk.s	Speculation check	$r2, target25$	chk.s.i $r2, target25(I)$ chk.s.m $r2, target25(M)$
fabs	Floating-point absolute value	$f1 =f3$	fmerge.s $f1 =f0, f3$
fadd.pc	Floating-point add	$f1 =f3, f2$	fma.pc.sf $f1 =f3, f1, f2$
.sf			
fcvt.xu f	Convert integer to float unsigned	$f1 =f3$	fma.pc.sf $f1 =f3, f1, f0$

Table B-1 Pseudo-ops Listed by Opcode (continued)

Opcode	Instruction Description	Operands	Equivalent Machine Instruction
<i>fmpy.pc</i> <i>.sf</i>	Floating-point multiply	<i>f1 =f3, f4</i>	<i>fma.pc.sf f1 =f3, f4, f0</i>
<i>fneg</i>	Floating-point negate	<i>f1 =f3</i>	<i>fmerge.ns f1 =f3, f3</i>
<i>fnegabs</i>	Floating-point negate absolute value	<i>f1 =f3</i>	<i>fmerge.ns f1 =f0, f3</i>
<i>fnorm.pc</i> <i>.sf</i>	Floating-point normalize	<i>f1 =f3</i>	<i>fma.pc.sf f1 =f3, f1, f0</i>
<i>fsub.pc</i> <i>.sf</i>	Floating-point subtract	<i>f1 =f3, f2</i>	<i>fms.pc.sf f1 =f3, f1, f2</i>
<i>mov</i>	Move to application register immediate	<i>ar3 =imm8</i>	<i>mov.i ar3 =imm8 (I)</i> <i>mov.m ar3 =imm8 (M)</i>
<i>mov</i>	Move to application register	<i>ar3 =r2</i>	<i>mov.i ar3 =r2 (I)</i> <i>mov.m ar3 =r2 (M)</i>
<i>mov</i>	Move floating-point register	<i>f1 =f3</i>	<i>fmerge.s f1 =f3, f3</i>
<i>mov</i>	Move from application register	<i>r1 =ar3</i>	<i>mov.i r1 =ar3 (I)</i> <i>mov.m r1 =ar3 (M)</i>
<i>mov</i>	Move immediate	<i>r1 =imm22</i>	<i>addl r1 =imm22, r0</i>
<i>mov</i>	Move general register	<i>r1 =r2</i>	<i>adds r1 =0, r2</i>
<i>mov</i>	Move to branch register	<i>b1 =r2</i>	<i>mov b1 =r2</i>
<i>nop</i>	No operation	<i>imm21</i>	<i>nop.b imm21 (B)</i> <i>nop.i imm21 (I)</i> <i>nop.m imm21 (M)</i> <i>nop.f imm21 (F)</i>

Table B-1 Pseudo-ops Listed by Opcode (continued)

Opcode	Instruction Description	Operands	Equivalent Machine Instruction
shl	Shift left	<i>r1</i> <i>=r2, count6</i>	dep.z <i>r1</i> <i>=r2, count6, 64-count6</i>
shr	Shift right signed	<i>r1</i> <i>=r3, count6</i>	extr <i>r1</i> <i>=r3, count6, 64-count6</i>
shr.u	Shift right unsigned	<i>r1</i> <i>=r3, count6</i>	extr.u <i>r1</i> <i>=r3, count6, 64-count6</i>
xma.lu	Fixed-point multiply low unsigned	<i>f1</i> <i>=f2, f3, f4,</i>	xma.l <i>f1 =f2, f3, f4</i>

Table B-2 lists pseudo-ops that omit one or more operands of the machine instruction. The assembler substitutes the missing operand with a predefined value. The missing operand(s) appear as bold text.

In addition to omitting many operands, many completers may also be omitted.

Table B-2 Pseudo-ops with Missing Operands

Pseudo-op	Missing Operand(s)	Substitute Value
alloc	alloc <i>r1=ar.pfs, i, l, o, r</i>	ar.pfs
cmp	cmp.crel.ctype <i>p1, p2=imm8, r3</i>	p0
cmp	cmp.crel.ctype <i>p1, p2=r2, r3</i>	p0
cmp4	cmp4.crel.ctype <i>p1, p2=imm8, r3</i>	p0
cmp4	cmp4.crel.ctype <i>p1, p2=r2, r3</i>	p0
cmpxchg	cmpxchgsz.sem.ldhint <i>r1=[r3], r2, ar.ccv</i>	ar.ccv
fclass	fclass.m.fctype <i>p1, p2=f2, f3</i> fclass.nm.fctype <i>p1, p2=f2, f3</i>	p0
fcmp	fcmp.fcrel.fctype.sf <i>p1, p2=f2, f3</i>	p0
mov	mov <i>pr=r2, mask17</i>	all ones
tbit	tbit.trel.ctype <i>p1, p2=r3, pos6</i>	p0

Table B-2 Pseudo-ops with Missing Operands (continued)

Pseudo-op	Missing Operand(s)	Substitute Value
tnat	tbit.trel ctype p1,p2=r3	p0

Link-relocation Operators

Table C-1 lists and describes the link-relocation operators, and their usage.:

Table C-1 Link-relocation Operators

Operator	Generates a Relocation For:	Usage:
@gp _{rel} (<i>expr</i>)	The current instruction or data object that calculates the <i>gp</i> -relative offset to the address given by <i>expr</i> .	<i>data8</i> statements and add long immediate instructions.
@secrel(<i>expr</i>)	The current data object that calculates the offset, relative to the beginning of the section, to the address given by <i>expr</i> .	<i>data4</i> and <i>data8</i> statements, and the <i>addl</i> instruction.
@segrel(<i>expr</i>)	The current data object that calculates the offset, relative to the beginning of the segment, to the address given by <i>expr</i> .	<i>data4</i> and <i>data8</i> statements, in <i>ELF</i> format.
@imagerel(<i>expr</i>)	The current data object that calculates the offset, relative to the beginning of the image, to the address given by <i>expr</i> .	<i>data4</i> statements, in <i>ELF</i> format.
@ltoff(<i>expr</i>)	The current instruction that instructs the linker to create a linkage table entry for <i>expr</i> , and calculates the <i>gp</i> -relative offset to the new linkage table entry.	add long immediate instructions.

Table C-1 Link-relocation Operators (continued)

Operator	Generates a Relocation For:	Usage:
<code>@fptr(<i>sym</i>)</code>	The current instruction or data object that calculates the address of the official <i>plabel</i> descriptor for the symbol <i>sym</i> , which must be a procedure label (function descriptor) name.	<code>data4</code> and <code>data8</code> statements, and move long immediate instructions. Requires function symbol in COFF format. It can be used in add long immediate instructions when combined with the <code>@ltoff</code> operator in the <code>@ltoff(@fptr(<i>sym</i>))</code> form.
<code>@plttoff(<i>sym</i>)</code>	The current instruction or data object that calculates the <code>gp</code> -relative offset to the procedure linkage table entry for the symbol <i>sym</i> , which must be a function name.	<code>data8</code> statements and add long immediate instructions. The <code>PLT</code> entry referenced by this operator should be used only for a direct procedure call. It does not serve as a function descriptor name
<code>@ltv(<i>expr</i>)</code>	The current data object that calculates the address of the relocatable expression <i>expr</i> , with one exception; while it is expected that the addresses created will need further relocation at run-time, the linker should not create a corresponding relocation in the output executable or shared object file. The runnn-time consumer of the information provided is expected to relocate these values.	<code>data4</code> statements in ELF format.
<code>@section(<i>sec</i>)</code>	The current data object that provides the section header number of section <i>sec</i> . Used for debug information.	<code>data2</code> statements in COFF format.



List of IA-64 Assembly Language Directives



Table D-1 summarizes the IA-64 assembly language directives according to category.

Table D-1 IA-64 Assembly Language Directives

Category	Directives
Alias declaration directives	<code>.alias</code> <code>.secalias</code>
Assembler annotations	<code>.pred.rel</code> <code>.pred.vector</code> <code>.mem.offset</code> <code>.entry</code>
Assembler modes	<code>.auto</code> <code>.explicit</code> <code>.default</code>
Byte order specification directive	<code>.msb</code> <code>.lsb</code>
Common symbol declaration directives	<code>.common</code> <code>.lcomm</code>
Cross-section data allocation statements	<code>.xdata1</code> <code>.xdata2</code> <code>.xdata4</code> <code>.xdata8</code> <code>.xstring</code> <code>.xstringz</code>

Table D-1 IA-64 Assembly Language Directives (continued)

Category	Directives
Data-allocation statements	data1 data2 data4 data8 real4 real8 real10 real16 string stringz
Explicit template selection directives	.mi .mf .bbb .mlx .mib .mmb .mmi .mbb .mfb .mmf
File symbol declaration directive	.file
Ident string directive	.ident
Include file directive	.include
Language specific data directive (Windows NT* specific)	.handlerdata
Procedure declaration directives	.proc .endp
Radix indicator directive	.radix
Register stack directive	.regstk
Reserving uninitialized space statements	.skip .org
Rotating register directives	.rotr .rotp .rotf

Table D-1 IA-64 Assembly Language Directives (continued)

Category	Directives
Section directives	.section .pushsection .popsection .previous .text .data .sdata .bss .sbss .rodata .comment
Section and data alignment directive	.align
Stack unwind information directives	See Table 3-6 in Chapter 3
Symbol scope declaration directives	.global .weak .local
Symbol type and size directives	.type .size
Symbolic debug directive	.ln
Symbolic debug directive Windows NT specific	.bf .ef
Virtual register allocation directives	.vreg.allocatable .vreg.safe_across_calls .vreg.family .vreg.var .vreg.undef



Glossary

absolute address	A virtual (not physical) address within the process' address space that is computed as an absolute number.
absolute expression	An expression that is not subject to link-time relocation.
alias	Two identifiers referring to the same element.
assembler	A program that translates assembly language into machine language.
assembly language	A low level symbolic language closely resembling machine-code language.
binding	The process of resolving a symbolic reference in one module by finding the definition of the symbol in another module, and substituting the address of the definition in place of the symbolic reference. The linker binds relocatable object modules together, and the DLL loader binds executable load modules. When searching for a definition, the linker and DLL loader search each module in a certain order, so that a definition of a symbol in one module has precedence over a definition of the same symbol in a later module. This order is called the binding order.

bundle	128 bits that include three instructions and a template field.
COFF	Common Object File Format, an object-module format.
directive	An assembler instruction that does not produce executable code.
execution time	The time during which a program is actually executing, not including the time during which the program and its DLLs are being loaded.
expression	A sequence of symbols that represents a value.
function name	A label that refers to a procedure entry point.
global symbol	Symbol visible outside the source file in which it is defined.
IA-32	Intel Architecture-32: the name for Intel's current 32-bit Instruction Set Architecture (ISA).
identifier	Syntactic representation of symbol names using alphabetic or special characters, and digits.
instruction	An operation code that performs a specific machine operation.
instruction group	IA-64 architecture instructions are organized in instruction groups. Each instruction group contains one or more statically contiguous instructions that execute in parallel. An instruction group must contain at least one instruction; there is no upper limit on the number of instructions in an instruction group. An instruction group is terminated statically by a stop, and dynamically by taken branches.

	<p>Stops are represented by a double semi-colon (; ;). You can explicitly define stops. Stops immediately follow an instruction, or appear on a separate line. They can be inserted between two instructions on the same line, as a semi-colon (;) is used to separate two instructions.</p>
Instruction Set Architecture	<p>The architecture that defines application level resources which include: user-level instructions, addressing modes, segmentation, and user visible register files.</p>
instruction tag	<p>A label that refers to an instruction.</p>
ISA	<p>See Instruction Set Architecture</p>
Itanium processor	<p>Name of Intel's first IA-64 processor.</p>
label	<p>A location in memory of code or data.</p>
link time	<p>The time when a program, dynamic-link library (DLL), or started object is processed by the linker. Any activity taking place at link time is static.</p>
linkage table	<p>A table containing text, unwind information, constants, literals, and pointers to imported data symbols and functions.</p>
local symbol	<p>Symbol visible only within the source file in which it is defined.</p>
location counter	<p>Keeps track of the current address when assembling a program. It starts at zero at the beginning of each segment and increments appropriately as each instruction is assembled. To adjust the location counter of a section, use the <code>.align</code> directive, or the <code>.org</code> directive.</p>

memory stack	A contiguous array of memory locations, commonly referred to as “the stack”, used in many processors to save the state of the calling procedure, pass parameters to the called procedure and store local variables for the currently executing procedure.
mnemonic	A predefined assembly-language name for machine instructions, pseudo-ops, directives, and data-allocation statements.
multiway branch bundle	A bundle that contains more than one branch instruction.
name space	A virtual (not physical) file. The assembler assigns names to a symbol, register, or mnemonic name space. Usually a name is defined only once in each separate name space. A name can be defined twice, in the symbol and register name space. In this case the register name takes precedence over the symbol name.
operator	The assembly-language operators indicate arithmetic or bitwise-logic calculations.
plabel	See procedure label.
predicate registers	64 1-bit predicate registers that control the execution of instructions. The first register, <code>p0</code> , is always treated as 1.
predication	The conditional execution of an instruction used to remove branches from code.
procedure label	A reference or pointer to a procedure. A procedure label (PLabel) is a special descriptor that uniquely identifies the procedure. The PLabel descriptor contains the address of the function’s actual entry point, and the linkage table pointer.

pseudo-op	An instruction aliasing a machine instruction, provided for the convenience of the programmer.
qualifying predicate	The execution of most instructions is gated by a qualifying predicate. If the predicate is true, the instruction executes normally; if the instruction is false the instruction does not modify architectural state or affect program behaviour.
register rotation	Software renaming of registers to provide every loop iteration with its own set of registers.
register stack configuration	A 64-bit register used to control the register stack engine (RSE).
relocatable expression	An expression that is subject to link-time relocation
rotating registers	Registers which are rotated by one register position at each loop execution so that the content of register X is in register $X+1$ after one rotation. The predicate, floating-point, and general registers can be rotated. The registers are rotated in a wrap-around fashion.
section	Portions of an object file, such as code or data, bound to one unit.
software pipelining	Pipelining of a loop by way of allowing the processor to execute, in any given time, several instructions in various instructions of the loop.
stacked registers	Stacked general registers, starting at $r32$, used to pass parameters to the called procedure and store local variables for the currently executing procedure.

statement	An assembly-language program consists of a series of statements. The following are primary types of assembly-language statements: <ul style="list-style-type: none">• label statements• instruction statements• directive statements• assignment statements• equate statements• data allocation statements• cross-data allocation statements
stop	Indicates the boundary of an instruction group. It is placed in the code by the assembly writer or compiler.
symbol declaration	The symbol address is resolved, not necessarily based on the current module. Declare symbols using a <code>.global</code> or <code>.weak</code> directive.
symbol definition	The symbol address is resolved based on the current module. A symbol is defined by assigning it a type and value. You can define a symbol either in an assignment statement, by using it as a label, or with a <code>.common</code> directive.
temporary symbol	A symbol name that is not placed in the object-file symbol table. To define a temporary symbol name, precede the name with a period (<code>.</code>).
weak symbol	Undefined symbol in object file, resolved during link time.



Index

Symbols

- " double quotes, 2-14
- "comdat", 3-2, 3-3
- "nobits", 3-2
- "note", 3-2
- "progbits", 3-2
- # line, C preprocessor support, 6-9
- # number sign, 2-2
- \$ dollar sign, 2-3
- () parentheses, 2-15
- . period, 2-3, Glossary-6
- . . two periods, 2-3
- .auto directive, 3-13
- . common directive, 6-7
- .default directive, 3-13
- .explicit directive, 3-13
- .mem.offset annotation, 7-2
- .ua completer, 5-3, 5-4
- @ "at" sign, 2-4
- @alt, 2-4
- @brcst, 2-4
- @filename, 6-9
- @filepath, 6-9
- @fptr, C-2
- @function, 4-4
- @gprel, C-1
- @imagerel, C-1
- @inf, 2-4
- @line, 6-9
- @ltoff, C-1
- @mix, 2-4
- @nat, 2-4
- @neg, 2-4
- @norm, 2-4
- @notype, 4-4
- @object, 4-4
- @pltoff, C-2
- @pos, 2-4
- @qnan, 2-4
- @rev, 2-4
- @secrel, C-1
- @segrel, C-1
- @shuf, 2-4
- @snan, 2-4
- @unorm, 2-4
- @zero, 2-4
- \ backslash, 2-14
- _ underscore, 2-13
- { } braces, 3-10
- ' single quote, 2-13

A

absolute expressions, 2-14
 absolute sections, 3-6
 alias declaration, 4-8
 .alias directive, 4-8
 .align statement, 5-5
 alignment, 5-4
 alloc instruction, 6-1
 annotations, 7-1
 .entry, 7-3
 .pred.rel, 7-1
 .pred.vector, 7-2
 application registers, A-2
 assembly modes, 3-13
 assignment statements, 2-19
 automatic mode, 3-13

B

.bbb directive, 3-11
 .bf directive, 3-24
 .body directive, 3-15
 body regions, 3-15
 branch registers, A-2
 .bss directive, 3-7
 bundles
 implicit and explicit bundling, 3-10
 template selection, 3-11
 with multiway branching, 3-9
 byte order specification, 5-3, 6-7

C

C numeric constants, 2-10
 character escapes, 2-13
 clear form, 7-1
 COMDAT section flags, 3-3

.comment directive, 3-7
 .common directive, 4-6
 common symbols, 4-6
 completers, 2-8, 2-18, 5-3, 5-4
 constants
 C numeric constants, 2-10
 MASM numeric constants, 2-11
 numeric constants, 2-9
 string constants, 2-14
 control registers, A-3
 .copy_state directive, 3-17
 cross-section data allocation statements, 2-23,
 5-5
 current location counter, 2-3
 current sections, 3-1

D

data allocation statements, 2-22, 5-1
 .align, 5-5
 .org, 5-3
 .skip, 5-3
 list of, 5-2
 specifying byte order, 5-3, 6-7
 .data directive, 3-7
 data object alignment, 5-4
 debug information, 3-24
 declarations, 4-1
 dependency violations, 3-13
 directive statements, 2-19
 directives
 .alias, 4-8
 .bbb, 3-11
 .body, 3-15
 .bss, 3-7
 .comment, 3-7
 .common, 4-6, 6-7
 .data, 3-7
 .endp, 3-14, 3-15, 3-16
 .frame, 3-17

- .file, 4-6
- .handlerdata, 3-15, 3-23
- .ident, 6-7
- .include, 3-8
- .lcomm, 4-7
- .ln, 3-24
- .local, 4-2
- .lsb, 5-3, 6-7
- .mbb, 3-11
- .mfb, 3-11
- .mib, 3-11
- .mlx, 3-11
- .mmb, 3-11
- .mmf, 3-11
- .mmi, 3-11
- .msb, 5-3, 6-7
- .personality, 3-17
- .popsection, 3-5
- .previous, 3-6
- .proc, 3-14, 3-15
- .prologue, 3-15
- .pushsection, 3-3, 3-5
- .radix, 6-8
- .regstk, 6-2
- .restore, 3-17
- .rodata, 3-7
- .rotf, 6-4
- .rotp, 6-4
- .rotr, 6-4
- .save, 3-17
- .save.b, 3-21
- .save.f, 3-21
- .save.g, 3-21
- .save.gf, 3-21
- .savepsp, 3-17
- .saveesp, 3-17
- .sbss, 3-7
- .sdata, 3-7
- .secalias, 4-8
- .section, 3-3, 3-4, 5-5
- .size, 4-5
- .spill, 3-20
- .text, 3-7
- .type, 4-4

- .unwabi, 3-20
- .unwentry, 3-16
- .vframe, 3-17
- .vframepsp, 3-17
- .vframesp, 3-17
- copy_state, 3-17
- label_state, 3-17
- listed by category, D-1 thru D-3
- Windows* NT* specific directives
 - .bf, 3-24
 - .ef, 3-24

E

- .ef directive, 3-24
- .endp directive, 3-14, 3-15, 3-16
- .entry annotation, 7-3
- equate statements, 2-21
- explicit mode, 3-13
- expressions, 2-14, Glossary-2
 - absolute expressions, 2-14
 - relocatable expressions, 2-14, Glossary-5
- external symbols, 4-8

F

- .fframe directive, 3-17
- .file directive, 4-6
- file name, overriding, 4-6
- floating-point registers, A-1
- function descriptor, 3-14
- function names, 2-5, Glossary-2

G

- general registers, A-1
- global symbols, 2-17, 4-2, Glossary-2

H

.`handlerdata` directive, 3-15, 3-23
headers, 3-15, 3-23

I

IA-32, Glossary-2
. `ident` directive, 6-7
ident strings, 6-7
identifiers, 2-1, Glossary-2
imply form, 7-1
. `include` directive, 3-8
include files, 3-8
indirect-register files, A-4
instruction bundles, 3-9
instruction completers, 2-8, 2-18
instruction group stops, 3-12
instruction groups, 3-12, Glossary-2
instruction statements, 2-17
instruction suffixes, 2-8
instruction tags, 2-5, Glossary-3
ISA, Glossary-3
Itanium™ processor, Glossary-3

L

label statements, 2-16
. `label_state` directive, 3-17
labels, 2-5, 2-16, Glossary-3
. `lcomm` directive, 4-7
link time, Glossary-3
linkage table, Glossary-3
link-relocation operators, 2-15
 list of, C-1 thru C-2
. `ln` directive, 3-24
local common symbols, 4-7
. `local` directive, 4-2

local symbols, 4-2, Glossary-3
location counter, 2-3, 3-1, Glossary-3
. `lsb` directive, 5-3, 6-7

M

machine instructions, 2-8, 3-9, 3-12, Glossary-2
 `alloc` instruction, 6-1
 tags, 2-5
MASM numeric constants, 2-11
memory stack, Glossary-4
. `mfb` directive, 3-11
. `mib` directive, 3-11
. `mlx` directive, 3-11
. `mmb` directive, 3-11
. `mmf` directive, 3-11
. `mbb` directive, 3-11
. `mmi` directive, 3-11
mnemonics, 2-7, Glossary-4
 cross-section data allocation mnemonics, 5-6
 data allocation mnemonics, 2-9, 5-2
 directive mnemonics, 2-9
 machine instruction mnemonics, 2-8
 pseudo-op mnemonics, 2-8
. `msb` directive, 5-3, 6-7
multiway branch bundles, 3-9, Glossary-4
mutex form, 7-1

N

name spaces, 2-2, 2-19, Glossary-4
nobits, 3-2, 3-7, 5-3
null terminated string, 6-7
numeric constants, 2-9, 6-8

O

opcodes, 2-18

operands, 2-19
operators, 2-15, Glossary-4
.org statement, 5-3

P

.personality directive, 3-17
PLabel, 3-14
plabel, Glossary-4
.popsection directive, 3-5
.pred.rel annotation, 7-1
.pred.vector annotation, 7-2
predefined section directives, 3-6, 3-7
predicate registers, A-2, Glossary-4
predication, Glossary-4
preprocessor built-in symbols, 6-9
.previous directive, 3-6
.proc directive, 3-14, 3-15
procedure label, Glossary-4
procedure label (PLabel), 3-14
procedures, 3-13
 declaring, 3-14
 prologues, 3-15
progbits, 3-2, 3-5
.prologue directive, 3-15
pseudo-ops, 2-8, Glossary-5
 list of, B-1 thru B-3
.pushsection directive, 3-3, 3-5

Q

qualifying predicate, 2-17, Glossary-5

R

radix constants and indicators, 6-8
.radix directive, 6-8
region headers, 3-15, 3-23

register assignment statements, 2-20
register equate statements, 2-21
register rotation, Glossary-5
register stack, 6-1
registers, 2-6
 application registers, A-2
 assigning new register names, 2-6, 2-20, 2-21
 branch registers, A-2
 control registers, A-3
 defining stacked registers, 6-3
 floating-point registers, A-1
 general registers, A-1
 indirect-register files, A-4
 predicate registers, A-2
 register forms, 2-7
 register renaming, 6-1
 rotating registers, 6-3
 stacked registers
 defining, 6-3
.regstk directive, 6-2
relocatable expressions, 2-14, Glossary-5
.restore directive, 3-17
.rodata directive, 3-7
rotating registers, Glossary-5
rotating registers, defining, 6-3, 6-6
.rotf directive, 6-4
.rotp directive, 6-4
.rotr directive, 6-4

S

.save directive, 3-17
.save.b directive, 3-21
.save.f directive, 3-21
.save.g directive, 3-21
.save.gf directive, 3-21
.savepsp directive, 3-17
.savesp directive, 3-17
.sbss directive, 3-7

- scope declarations, 4-1
 - declaring a global symbol, 4-2
 - declaring a local symbol, 4-2
 - declaring a weak symbol, 4-2
- `.sdata` directive, 3-7
- `.secalias` directive, 4-8
- `.section` directive, 3-3, 3-4, 5-5
- section flag characters, 3-2
- section flags and types, 3-6
- section names, 2-5
- section return directive, 3-6
- sections, 3-1, Glossary-5
 - defining, 3-4
 - defining a section stack, 3-5
 - predefined, 3-6
 - returning to, 3-6
 - section types, 3-2
 - specifying section flags and types, 3-2
- `.size` directive, 4-5
- `.skip` statement, 5-3
- software pipelining, 6-4, Glossary-5
- `.spill` directive, 3-20
- stack unwind directives, 3-15
 - list of, 3-16
 - using, 3-22
- stack unwind information, 3-13
- stacked registers, 6-1, Glossary-5
 - defining, 6-3
- statements, 2-16, Glossary-6
 - assignment statements, 2-19
 - defining register names, 2-20
 - defining symbol names, 2-20
 - cross-section data allocation statements, 2-23, 5-5
 - data allocation, 5-1
 - data allocation statements, 2-22, 5-2
 - directive statements, 2-19
 - equate statements, 2-21
 - defining register names, 2-21
 - defining symbol names, 2-21
 - instruction statements, 2-17
 - label statements, 2-16
- stop, Glossary-6
 - in bundles, 3-9
 - in explicit bundles, 3-10
 - in instruction groups, 3-12
- stop bit, 3-9
- string constants, 2-14
- symbol assignment statements, 2-20
- symbol definition, 2-6, Glossary-6
- symbol equate statements, 2-21
- symbolic constants, 2-5
- symbolic debug information
 - Windows* NT* specific, 3-24
- symbols, 2-2
 - assigning a value, 2-6
 - declaring, Glossary-6
 - declaring common symbols, 4-6
 - declaring local common symbols, 4-7
 - declaring symbol scopes, 4-1
 - defining, 2-6, Glossary-6
 - predefined symbol names, 2-4
 - referencing external symbols, 4-8
 - specifying symbol size, 4-5
 - specifying symbol types, 4-4
 - symbol names, 2-3
 - symbol types, 2-5

T

- tag operand, 2-5, 3-16
- temporary symbols, 2-3, Glossary-6
- `.text` directive, 3-7
- thread local storage, 3-3
- `.type` directive, 4-4

U

- unaligned data objects, 5-3, 5-4
- undefined symbols, Glossary-6
- uninitialized data allocation, 5-3

.unwabi directive, 3-20
.unwentry directive, 3-16

V

.vframe directive, 3-17
.vframepsp directive, 3-17
.vframesp directive, 3-17

W

weak symbols, 4-2, Glossary-6