

# GNAT Reference Manual

---

GNAT Reference Manual , Jan 09, 2026

AdaCore

Copyright © 2008-2026, Free Software Foundation

---

# Table of Contents

<b>1</b>	<b>About This Guide</b>	<b>2</b>
1.1	What This Reference Manual Contains	2
1.2	Conventions	3
1.3	Related Information	3
<b>2</b>	<b>Implementation Defined Pragas</b>	<b>5</b>
2.1	Pragma Abort_Defer	5
2.2	Pragma Abstract_State	5
2.3	Pragma Ada_83	6
2.4	Pragma Ada_95	7
2.5	Pragma Ada_05	7
2.6	Pragma Ada_2005	7
2.7	Pragma Ada_12	8
2.8	Pragma Ada_2012	8
2.9	Pragma Ada_2022	8
2.10	Pragma Aggregate_Individually_Assign	9
2.11	Pragma Allow_Integer_Address	9
2.12	Pragma Always_Terminates	10
2.13	Pragma Annotate	10
2.14	Pragma Assert	10
2.15	Pragma Assert_And_Cut	11
2.16	Pragma Assertion_Level	11
2.17	Pragma Assertion_Policy	12
2.18	Pragma Assume	13
2.19	Pragma Assume_No_Invalid_Values	14
2.20	Pragma Async_Readers	14
2.21	Pragma Async_Writers	14
2.22	Pragma Attribute_Definition	15
2.23	Pragma C_Pass_By_Copy	15
2.24	Pragma Check	15
2.25	Pragma Check_Float_Overflow	16
2.26	Pragma Check_Name	16
2.27	Pragma Check_Policy	17
2.28	Pragma Comment	18
2.29	Pragma Common_Object	18
2.30	Pragma Compile_Time_Error	19
2.31	Pragma Compile_Time_Warning	19
2.32	Pragma Complete_Representation	19
2.33	Pragma Complex_Representation	20
2.34	Pragma Component_Alignment	20
2.35	Pragma Constant_After_Elaboration	21
2.36	Pragma Contract_Cases	21
2.37	Pragma Convention_Identifier	22

2.38	Pragma CPP_Class	23
2.39	Pragma CPP_Constructor	23
2.40	Pragma CPP_Virtual	24
2.41	Pragma CPP_Vtable	24
2.42	Pragma CPU	24
2.43	Pragma Deadline_Floor	24
2.44	Pragma Debug	24
2.45	Pragma Debug_Policy	25
2.46	Pragma Default_Initial_Condition	25
2.47	Pragma Default_Scalar_Storage_Order	25
2.48	Pragma Default_Storage_Pool	26
2.49	Pragma Depends	26
2.50	Pragma Detect_Blocking	27
2.51	Pragma Disable_Atomic_Synchronization	27
2.52	Pragma Dispatching_Domain	28
2.53	Pragma Effective_Reads	28
2.54	Pragma Effective_Writes	28
2.55	Pragma Elaboration_Checks	28
2.56	Pragma Eliminate	28
2.57	Pragma Enable_Atomic_Synchronization	31
2.58	Pragma Exceptional_Cases	31
2.59	Pragma Exit_Cases	31
2.60	Pragma Export_Function	31
2.61	Pragma Export_Object	32
2.62	Pragma Export_Procedure	33
2.63	Pragma Export_Valued_Procedure	34
2.64	Pragma Extend_System	34
2.65	Pragma Extensions_Allowed	35
2.66	Pragma Extensions_Visible	35
2.67	Pragma External	36
2.68	Pragma External_Name_Casing	36
2.69	Pragma Fast_Math	37
2.70	Pragma Favor_Top_Level	37
2.71	Pragma Finalize_Storage_Only	37
2.72	Pragma Float_Representation	38
2.73	Pragma Ghost	38
2.74	Pragma Global	38
2.75	Pragma Ident	39
2.76	Pragma Ignore_Pragma	39
2.77	Pragma Implementation_Defined	39
2.78	Pragma Implemented	39
2.79	Pragma Implicit_Packing	40
2.80	Pragma Import_Function	41
2.81	Pragma Import_Object	42
2.82	Pragma Import_Procedure	42
2.83	Pragma Import_Valued_Procedure	43
2.84	Pragma Independent	44

2.85	Pragma Independent_Components	44
2.86	Pragma Initial_Condition	44
2.87	Pragma Initialize_Scalars	45
2.88	Pragma Initializes	46
2.89	Pragma Inline_Always	46
2.90	Pragma Inline_Generic	47
2.91	Pragma Interface	47
2.92	Pragma Interface_Name	47
2.93	Pragma Interrupt_Handler	47
2.94	Pragma Interrupt_State	48
2.95	Pragma Interrupts_System_By_Default	49
2.96	Pragma Invariant	49
2.97	Pragma Keep_Names	49
2.98	Pragma License	50
2.99	Pragma Link_With	51
2.100	Pragma Linker_Alias	51
2.101	Pragma Linker_Constructor	52
2.102	Pragma Linker_Destructor	52
2.103	Pragma Linker_Section	52
2.104	Pragma Lock_Free	53
2.105	Pragma Loop_Invariant	54
2.106	Pragma Loop_Optimize	55
2.107	Pragma Loop_Variant	55
2.108	Pragma Machine_Attribute	56
2.109	Pragma Main	56
2.110	Pragma Main_Storage	56
2.111	Pragma Max_Queue_Length	57
2.112	Pragma No_Body	57
2.113	Pragma No_Caching	57
2.114	Pragma No_Component_Reordering	57
2.115	Pragma No_Elaboration_Code_All	58
2.116	Pragma No_Heap_Finalization	58
2.117	Pragma No_Inline	58
2.118	Pragma No_Raise	59
2.119	Pragma No_Return	59
2.120	Pragma No_Strict_Aliasing	59
2.121	Pragma No_Tagged_Streams	59
2.122	Pragma Normalize_Scalars	60
2.123	Pragma Obsolescent	61
2.124	Pragma Optimize_Alignment	63
2.125	Pragma Ordered	64
2.126	Pragma Overflow_Mode	65
2.127	Pragma Overriding_Renamings	65
2.128	Pragma Part_Of	66
2.129	Pragma Partition_Elaboration_Policy	66
2.130	Pragma Passive	66
2.131	Pragma Persistent_BSS	67

2.132	Pragma Post	67
2.133	Pragma Postcondition	67
2.134	Pragma Post_Class	70
2.135	Pragma Pre	70
2.136	Pragma Precondition	70
2.137	Pragma Predicate	71
2.138	Pragma Predicate_Failure	72
2.139	Pragma Preelaborable_Initialization	72
2.140	Pragma Prefix_Exception_Messages	72
2.141	Pragma Pre_Class	72
2.142	Pragma Priority_Specific_Dispatching	73
2.143	Pragma Profile	73
2.144	Pragma Profile_Warnings	76
2.145	Pragma Program_Exit	76
2.146	Pragma Propagate_Exceptions	76
2.147	Pragma Provide_Shift_Operators	76
2.148	Pragma Psect_Object	77
2.149	Pragma Pure_Function	77
2.150	Pragma Rational	78
2.151	Pragma Ravenscar	78
2.152	Pragma Refined_Depends	78
2.153	Pragma Refined_Global	79
2.154	Pragma Refined_Post	79
2.155	Pragma Refined_State	79
2.156	Pragma Relative_Deadline	80
2.157	Pragma Remote_Access_Type	80
2.158	Pragma Rename_Pragma	80
2.159	Pragma Restricted_Run_Time	81
2.160	Pragma Restriction_Warnings	81
2.161	Pragma Reviewable	81
2.162	Pragma Secondary_Stack_Size	82
2.163	Pragma Share_Generic	83
2.164	Pragma Shared	83
2.165	Pragma Short_Circuit_And_Or	83
2.166	Pragma Short_Descriptors	84
2.167	Pragma Side_Effects	84
2.168	Pragma Simple_Storage_Pool_Type	84
2.169	Pragma Source_File_Name	85
2.170	Pragma Source_File_Name_Project	86
2.171	Pragma Source_Reference	87
2.172	Pragma SPARK_Mode	87
2.173	Pragma Static_Elaboration_Desired	88
2.174	Pragma Stream_Convert	88
2.175	Pragma Style_Checks	89
2.176	Pragma Subprogram_Variant	91
2.177	Pragma Subtitle	92
2.178	Pragma Suppress	92

2.179	Pragma Suppress_All .....	93
2.180	Pragma Suppress_Debug_Info .....	93
2.181	Pragma Suppress_Exception_Locations .....	93
2.182	Pragma Suppress_Initialization .....	94
2.183	Pragma Task_Name .....	94
2.184	Pragma Task_Storage .....	95
2.185	Pragma Test_Case .....	95
2.186	Pragma Thread_Local_Storage .....	96
2.187	Pragma Time_Slice .....	96
2.188	Pragma Title .....	97
2.189	Pragma Type_Invariant .....	97
2.190	Pragma Type_Invariant_Class .....	97
2.191	Pragma Unchecked_Union .....	98
2.192	Pragma Unevaluated_Use_Of_Old .....	98
2.193	Pragma User_Aspect_Definition .....	98
2.194	Pragma Unimplemented_Unit .....	99
2.195	Pragma Universal_Aliasing .....	99
2.196	Pragma Unmodified .....	99
2.197	Pragma Unreferenced .....	100
2.198	Pragma Unreferenced_Objects .....	101
2.199	Pragma Unreserve_All_Interrupts .....	101
2.200	Pragma Unsuppress .....	101
2.201	Pragma Unused .....	102
2.202	Pragma Use_VADS_Size .....	102
2.203	Pragma Validity_Checks .....	102
2.204	Pragma Volatile .....	103
2.205	Pragma Volatile_Full_Access .....	103
2.206	Pragma Volatile_Function .....	104
2.207	Pragma Warning_As_Error .....	104
2.208	Pragma Warnings .....	105
2.209	Pragma Weak_External .....	107
2.210	Pragma Wide_Character-Encoding .....	108
<b>3</b>	<b>Implementation Defined Aspects .....</b>	<b>109</b>
3.1	Aspect Abstract_State .....	109
3.2	Aspect Always_Terminates .....	109
3.3	Aspect Annotate .....	109
3.4	Aspect Async_Readers .....	110
3.5	Aspect Async_Writers .....	110
3.6	Aspect Constant_After_Elaboration .....	110
3.7	Aspect Contract_Cases .....	110
3.8	Aspect Depends .....	110
3.9	Aspect Default_Initial_Condition .....	110
3.10	Aspect Dimension .....	110
3.11	Aspect Dimension_System .....	111
3.12	Aspect Disable_Controlled .....	112

3.13	Aspect Effective_Reads .....	112
3.14	Aspect Effective_Writes .....	112
3.15	Aspect Exceptional_Cases .....	112
3.16	Aspect Exit_Cases .....	112
3.17	Aspect Extended_Access .....	112
3.18	Aspect Extensions_Visible .....	113
3.19	Aspect Favor_Top_Level .....	114
3.20	Aspect Ghost .....	114
3.21	Aspect Ghost_Predicate .....	114
3.22	Aspect Global .....	114
3.23	Aspect Initial_Condition .....	114
3.24	Aspect Initializes .....	114
3.25	Aspect Inline_Always .....	114
3.26	Aspect Invariant .....	114
3.27	Aspect Invariant'Class .....	114
3.28	Aspect Iterable .....	114
3.29	Aspect Linker_Section .....	115
3.30	Aspect Local_Restrictions .....	115
3.31	Aspect Lock_Free .....	116
3.32	Aspect Max_Queue_Length .....	116
3.33	Aspect No_Caching .....	116
3.34	Aspect No_Elaboration_Code_All .....	116
3.35	Aspect No_Inline .....	116
3.36	Aspect No_Raise .....	116
3.37	Aspect No_Tagged_Streams .....	117
3.38	Aspect No_Task_Parts .....	117
3.39	Aspect Object_Size .....	117
3.40	Aspect Obsolescent .....	117
3.41	Aspect Part_Of .....	117
3.42	Aspect Persistent_BSS .....	117
3.43	Aspect Potentially_Invalid .....	117
3.44	Aspect Predicate .....	117
3.45	Aspect Program_Exit .....	117
3.46	Aspect Pure_Function .....	117
3.47	Aspect Refined_Depends .....	118
3.48	Aspect Refined_Global .....	118
3.49	Aspect Refined_Post .....	118
3.50	Aspect Refined_State .....	118
3.51	Aspect Relaxed_Initialization .....	118
3.52	Aspect Remote_Access_Type .....	118
3.53	Aspect Scalar_Storage_Order .....	118
3.54	Aspect Secondary_Stack_Size .....	118
3.55	Aspect Shared .....	118
3.56	Aspect Side_Effects .....	118
3.57	Aspect Simple_Storage_Pool .....	118
3.58	Aspect Simple_Storage_Pool_Type .....	118
3.59	Aspect SPARK_Mode .....	119

3.60	Aspect Subprogram_Variant .....	119
3.61	Aspect Suppress_Debug_Info .....	119
3.62	Aspect Suppress_Initialization .....	119
3.63	Aspect Test_Case .....	119
3.64	Aspect Thread_Local_Storage .....	119
3.65	Aspect Universal_Aliasing .....	119
3.66	Aspect Unmodified .....	119
3.67	Aspect Unreferenced .....	119
3.68	Aspect Unreferenced_Objects .....	119
3.69	Aspect User_Aspect .....	119
3.70	Aspect Value_Size .....	120
3.71	Aspect Volatile_Full_Access .....	120
3.72	Aspect Volatile_Function .....	120
3.73	Aspect Warnings .....	120
<b>4</b>	<b>Implementation Defined Attributes .....</b>	<b>121</b>
4.1	Attribute Abort_Signal .....	121
4.2	Attribute Address_Size .....	121
4.3	Attribute Asm_Input .....	121
4.4	Attribute Asm_Output .....	121
4.5	Attribute Atomic_Always_Lock_Free .....	122
4.6	Attribute Bit .....	122
4.7	Attribute Bit_Position .....	122
4.8	Attribute Code_Address .....	122
4.9	Attribute Compiler_Version .....	123
4.10	Attribute Constrained .....	123
4.11	Attribute Default_Bit_Order .....	123
4.12	Attribute Default_Scalar_Storage_Order .....	123
4.13	Attribute Deref .....	123
4.14	Attribute Descriptor_Size .....	123
4.15	Attribute Elaborated .....	124
4.16	Attribute Elab_Body .....	124
4.17	Attribute Elab_Spec .....	124
4.18	Attribute Elab_Subp_Body .....	124
4.19	Attribute Emax .....	124
4.20	Attribute Enabled .....	125
4.21	Attribute Enum_Rep .....	125
4.22	Attribute Enum_Val .....	125
4.23	Attribute Epsilon .....	126
4.24	Attribute Fast_Math .....	126
4.25	Attribute Finalization_Size .....	126
4.26	Attribute Fixed_Value .....	126
4.27	Attribute From_Address .....	126
4.28	Attribute From_Any .....	127
4.29	Attribute Has_Access_Values .....	127
4.30	Attribute Has_Discriminants .....	127

4.31	Attribute Has_Tagged_Values .....	127
4.32	Attribute Img .....	127
4.33	Attribute Initialized .....	128
4.34	Attribute Integer_Value .....	128
4.35	Attribute Invalid_Value .....	128
4.36	Attribute Large .....	128
4.37	Attribute Library_Level .....	128
4.38	Attribute Loop_Entry .....	129
4.39	Attribute Machine_Size .....	129
4.40	Attribute Mantissa .....	129
4.41	Attribute Maximum_Alignment .....	129
4.42	Attribute Max_Integer_Size .....	129
4.43	Attribute Mechanism_Code .....	130
4.44	Attribute Null_Parameter .....	130
4.45	Attribute Object_Size .....	130
4.46	Attribute Old .....	131
4.47	Attribute Passed_By_Reference .....	131
4.48	Attribute Pool_Address .....	131
4.49	Attribute Range_Length .....	132
4.50	Attribute Restriction_Set .....	132
4.51	Attribute Result .....	133
4.52	Attribute Round .....	133
4.53	Attribute Safe_Emax .....	133
4.54	Attribute Safe_Large .....	133
4.55	Attribute Safe_Small .....	133
4.56	Attribute Scalar_Storage_Order .....	133
4.57	Attribute Simple_Storage_Pool .....	136
4.58	Attribute Small .....	136
4.59	Attribute Small_Denominator .....	137
4.60	Attribute Small_Numerator .....	137
4.61	Attribute Storage_Unit .....	137
4.62	Attribute Stub_Type .....	137
4.63	Attribute System_Allocator_Alignment .....	137
4.64	Attribute Target_Name .....	137
4.65	Attribute To_Address .....	138
4.66	Attribute To_Any .....	138
4.67	Attribute Type_Class .....	138
4.68	Attribute Type_Key .....	138
4.69	Attribute TypeCode .....	138
4.70	Attribute Unconstrained_Array .....	139
4.71	Attribute Universal_Literal_String .....	139
4.72	Attribute Unrestricted_Access .....	139
4.73	Attribute Update .....	142
4.74	Attribute Valid_Value .....	143
4.75	Attribute ValidScalars .....	143
4.76	Attribute VADS_Size .....	144
4.77	Attribute Value_Size .....	144

4.78	Attribute Wchar_T_Size .....	144
4.79	Attribute Word_Size .....	144

## 5 Standard and Implementation

### Defined Restrictions ..... 145

5.1	Partition-Wide Restrictions .....	145
5.1.1	Immediate_Reclamation .....	145
5.1.2	Max_Asynchronous_Select_Nesting .....	145
5.1.3	Max_Entry_Queue_Length .....	145
5.1.4	Max_Protected_Entries .....	145
5.1.5	Max_Select_Alternatives .....	145
5.1.6	Max_Storage_At_Blocking .....	146
5.1.7	Max_Task_Entries .....	146
5.1.8	Max_Tasks .....	146
5.1.9	No_Abort_Statements .....	146
5.1.10	No_Access_Parameter_Allocators .....	146
5.1.11	No_Access_Subprograms .....	146
5.1.12	No_Allocators .....	146
5.1.13	No_Anonymous_Allocators .....	146
5.1.14	No_Asynchronous_Control .....	146
5.1.15	No_Calendar .....	146
5.1.16	No_Coextensions .....	146
5.1.17	No_Default_Initialization .....	147
5.1.18	No_Delay .....	147
5.1.19	No_Dependence .....	147
5.1.20	No_Direct_Boolean_Operators .....	147
5.1.21	No_Dispatch .....	147
5.1.22	No_Dispatching_Calls .....	147
5.1.23	No_Dynamic_Attachment .....	149
5.1.24	No_Dynamic_Priorities .....	149
5.1.25	No_Entry_Calls_In_Elaboration_Code .....	149
5.1.26	No_Enumeration_Maps .....	149
5.1.27	No_Exception_Handlers .....	149
5.1.28	No_Exception_Propagation .....	149
5.1.29	No_Exception_Registration .....	150
5.1.30	No_Exceptions .....	150
5.1.31	No_Finalization .....	150
5.1.32	No_Fixed_Point .....	150
5.1.33	No_Floating_Point .....	150
5.1.34	No_Implicit_Conditionals .....	150
5.1.35	No_Implicit_Dynamic_Code .....	151
5.1.36	No_Implicit_Heap_Allocations .....	151
5.1.37	No_Implicit_Protected_Object_Allocations .....	151
5.1.38	No_Implicit_Task_Allocations .....	151
5.1.39	No_InitializeScalars .....	151
5.1.40	No_IO .....	151

5.1.41	No_Local_Allocators	151
5.1.42	No_Local_Protected_Objects	151
5.1.43	No_Local_Tagged_Types	151
5.1.44	No_Local_Timing_Events	152
5.1.45	No_Long_Long_Integers	152
5.1.46	No_Multiple_Elaboration	152
5.1.47	No_Nested_Finalization	152
5.1.48	No_Protected_Type_Allocators	152
5.1.49	No_Protected_Types	152
5.1.50	No_Recursion	152
5.1.51	No_Reentrancy	152
5.1.52	No_Relative_Delay	152
5.1.53	No_Requeue_Statements	152
5.1.54	No_Secondary_Stack	153
5.1.55	No_Select_Statements	153
5.1.56	No_Specific_Termination_Handlers	153
5.1.57	No_Specification_of_Aspect	153
5.1.58	No_Standard_Allocators_After_Elaboration	153
5.1.59	No_Standard_Storage_Pools	153
5.1.60	No_Stream_Optimizations	153
5.1.61	No_Streams	153
5.1.62	No_Tagged_Type_Registration	154
5.1.63	No_Task_Allocators	154
5.1.64	No_Task_At_Interrupt_Priority	154
5.1.65	No_Task_Attributes_Package	154
5.1.66	No_Task_Hierarchy	154
5.1.67	No_Task_Termination	154
5.1.68	No_Tasking	154
5.1.69	No_Terminate_Alternatives	154
5.1.70	No_Unchecked_Access	155
5.1.71	No_Unchecked_Conversion	155
5.1.72	No_Unchecked_Deallocation	155
5.1.73	No_Use_Of_Attribute	155
5.1.74	No_Use_Of_Entity	155
5.1.75	No_Use_Of_Pragma	155
5.1.76	Pure_Barriers	155
5.1.77	Simple_Barriers	156
5.1.78	Static_Priorities	156
5.1.79	Static_Storage_Size	156
5.2	Program Unit Level Restrictions	156
5.2.1	No_Elaboration_Code	156
5.2.2	No_Dynamic_Accessibility_Checks	157
5.2.3	No_Dynamic_Sized_Objects	157
5.2.4	No_Entry_Queue	158
5.2.5	No_Implementation_Aspect_Specifications	158
5.2.6	No_Implementation_Attributes	158
5.2.7	No_Implementation_Identifiers	158

5.2.8	No_Implementation_Pragmas .....	158
5.2.9	No_Implementation_Restrictions .....	158
5.2.10	No_Implementation_Units .....	158
5.2.11	No_Implicit_Aliasing .....	158
5.2.12	No_Implicit_Loops .....	159
5.2.13	No_Obsolescent_Features .....	159
5.2.14	No_Wide_Characters .....	159
5.2.15	Static_Dispatch_Tables .....	159
5.2.16	SPARK_05 .....	159
<b>6</b>	<b>Implementation Advice .....</b>	<b>160</b>
6.1	RM 1.1.3(20): Error Detection .....	160
6.2	RM 1.1.3(31): Child Units .....	160
6.3	RM 1.1.5(12): Bounded Errors .....	160
6.4	RM 2.8(16): Pragmas .....	160
6.5	RM 2.8(17-19): Pragmas .....	161
6.6	RM 3.5.2(5): Alternative Character Sets .....	161
6.7	RM 3.5.4(28): Integer Types .....	162
6.8	RM 3.5.4(29): Integer Types .....	162
6.9	RM 3.5.5(8): Enumeration Values .....	162
6.10	RM 3.5.7(17): Float Types .....	162
6.11	RM 3.6.2(11): Multidimensional Arrays .....	163
6.12	RM 9.6(30-31): Duration'Small .....	163
6.13	RM 10.2.1(12): Consistent Representation .....	163
6.14	RM 11.4.1(19): Exception Information .....	163
6.15	RM 11.5(28): Suppression of Checks .....	164
6.16	RM 13.1 (21-24): Representation Clauses .....	164
6.17	RM 13.2(6-8): Packed Types .....	164
6.18	RM 13.3(14-19): Address Clauses .....	165
6.19	RM 13.3(29-35): Alignment Clauses .....	165
6.20	RM 13.3(42-43): Size Clauses .....	166
6.21	RM 13.3(50-56): Size Clauses .....	166
6.22	RM 13.3(71-73): Component Size Clauses .....	167
6.23	RM 13.4(9-10): Enumeration Representation Clauses .....	167
6.24	RM 13.5.1(17-22): Record Representation Clauses .....	167
6.25	RM 13.5.2(5): Storage Place Attributes .....	168
6.26	RM 13.5.3(7-8): Bit Ordering .....	168
6.27	RM 13.7(37): Address as Private .....	168
6.28	RM 13.7.1(16): Address Operations .....	168
6.29	RM 13.9(14-17): Unchecked Conversion .....	168
6.30	RM 13.11(23-25): Implicit Heap Usage .....	169
6.31	RM 13.11.2(17): Unchecked Deallocation .....	169
6.32	RM 13.13.2(1.6): Stream Oriented Attributes .....	169
6.33	RM A.1(52): Names of Predefined Numeric Types .....	170
6.34	RM A.3.2(49): <b>Ada.Characters.Handling</b> .....	170
6.35	RM A.4.4(106): Bounded-Length String Handling .....	170

6.36	RM A.5.2(46-47): Random Number Generation .....	170
6.37	RM A.10.7(23): <code>Get_Immediate</code> .....	171
6.38	RM A.18: <code>Containers</code> .....	171
6.39	RM B.1(39-41): <code>Pragma Export</code> .....	171
6.40	RM B.2(12-13): <code>Package Interfaces</code> .....	172
6.41	RM B.3(63-71): Interfacing with C .....	172
6.42	RM B.4(95-98): Interfacing with COBOL .....	173
6.43	RM B.5(22-26): Interfacing with Fortran .....	173
6.44	RM C.1(3-5): Access to Machine Operations .....	174
6.45	RM C.1(10-16): Access to Machine Operations .....	174
6.46	RM C.3(28): Interrupt Support .....	175
6.47	RM C.3.1(20-21): Protected Procedure Handlers .....	175
6.48	RM C.3.2(25): <code>Package Interrupts</code> .....	175
6.49	RM C.4(14): Pre-elaboration Requirements .....	175
6.50	RM C.5(8): <code>Pragma Discard_Names</code> .....	175
6.51	RM C.7.2(30): The Package <code>Task_Attributes</code> .....	175
6.52	RM D.3(17): Locking Policies .....	176
6.53	RM D.4(16): Entry Queuing Policies .....	176
6.54	RM D.6(9-10): Preemptive Abort .....	176
6.55	RM D.7(21): Tasking Restrictions .....	176
6.56	RM D.8(47-49): Monotonic Time .....	176
6.57	RM E.5(28-29): Partition Communication Subsystem .....	177
6.58	RM F(7): COBOL Support .....	177
6.59	RM F.1(2): Decimal Radix Support .....	177
6.60	RM G: Numerics .....	177
6.61	RM G.1.1(56-58): Complex Types .....	178
6.62	RM G.1.2(49): Complex Elementary Functions .....	178
6.63	RM G.2.4(19): Accuracy Requirements .....	179
6.64	RM G.2.6(15): Complex Arithmetic Accuracy .....	179
6.65	RM H.6(15/2): <code>Pragma Partition_Elaboration_Policy</code> .....	179
<b>7</b>	<b>Implementation Defined Characteristics ....</b>	<b>180</b>
<b>8</b>	<b>Intrinsic Subprograms .....</b>	<b>199</b>
8.1	Intrinsic Operators .....	199
8.2	<code>Compilation_ISO_Date</code> .....	199
8.3	<code>Compilation_Date</code> .....	199
8.4	<code>Compilation_Time</code> .....	199
8.5	<code>Enclosing_Entity</code> .....	200
8.6	<code>Exception_Information</code> .....	200
8.7	<code>Exception_Message</code> .....	200
8.8	<code>Exception_Name</code> .....	200
8.9	<code>File</code> .....	200
8.10	<code>Line</code> .....	200
8.11	<code>Shifts and Rotates</code> .....	201
8.12	<code>Source_Location</code> .....	201

<b>9</b>	<b>Representation Clauses and Pragmas .....</b>	<b>202</b>
9.1	Alignment Clauses .....	202
9.2	Size Clauses .....	203
9.3	Storage_Size Clauses .....	204
9.4	Size of Variant Record Objects .....	205
9.5	Biased Representation .....	207
9.6	Value_Size and Object_Size Clauses .....	207
9.7	Component_Size Clauses .....	210
9.8	Bit_Order Clauses .....	211
9.9	Effect of Bit_Order on Byte Ordering .....	212
9.10	Pragma Pack for Arrays .....	216
9.11	Pragma Pack for Records .....	218
9.12	Record Representation Clauses .....	219
9.13	Handling of Records with Holes .....	220
9.14	Enumeration Clauses .....	221
9.15	Address Clauses .....	222
9.16	Use of Address Clauses for Memory-Mapped I/O .....	227
9.17	Effect of Convention on Representation .....	227
9.18	Conventions and Anonymous Access Types .....	228
9.19	Determining the Representations chosen by GNAT .....	230
<b>10</b>	<b>Standard Library Routines .....</b>	<b>233</b>
<b>11</b>	<b>The Implementation of Standard I/O .....</b>	<b>244</b>
11.1	Standard I/O Packages .....	244
11.2	FORM Strings .....	245
11.3	Direct_IO .....	245
11.4	Sequential_IO .....	245
11.5	Text_IO .....	246
11.5.1	Stream Pointer Positioning .....	247
11.5.2	Reading and Writing Non-Regular Files .....	247
11.5.3	Get_Immediate .....	248
11.5.4	Treating Text_IO Files as Streams .....	248
11.5.5	Text_IO Extensions .....	248
11.5.6	Text_IO Facilities for Unbounded Strings .....	248
11.6	Wide_Text_IO .....	249
11.6.1	Stream Pointer Positioning .....	251
11.6.2	Reading and Writing Non-Regular Files .....	252
11.7	Wide_Wide_Text_IO .....	252
11.7.1	Stream Pointer Positioning .....	253
11.7.2	Reading and Writing Non-Regular Files .....	253
11.8	Stream_IO .....	254
11.9	Text Translation .....	254
11.10	Shared Files .....	254
11.11	Filenames encoding .....	255

11.12	File content encoding .....	255
11.13	Open Modes .....	256
11.14	Operations on C Streams.....	256
11.15	Interfacing to C Streams .....	259
<b>12</b>	<b>The GNAT Library.....</b>	<b>262</b>
12.1	Ada.Characters.Latin_9 (a-chlat9.ads) .....	262
12.2	Ada.Characters.Wide_Latin_1 (a-cwila1.ads) .....	262
12.3	Ada.Characters.Wide_Latin_9 (a-cwila9.ads) .....	262
12.4	Ada.Characters.Wide_Wide_Latin_1 (a-chzla1.ads) .....	262
12.5	Ada.Characters.Wide_Wide_Latin_9 (a-chzla9.ads) .....	263
12.6	Ada.Containers.Bounded_Holders (a-coboho.ads).....	263
12.7	Ada.Command_Line.Environment (a-colien.ads).....	263
12.8	Ada.Command_Line.Remove (a-colire.ads).....	263
12.9	Ada.Command_Line.Response_File (a-clrefi.ads).....	263
12.10	Ada.Direct_IO.C_Streams (a-diocst.ads).....	263
12.11	Ada.Exceptions.Is_Null_Occurrence (a-einuoc.ads)....	263
12.12	Ada.Exceptions.Last_Chance_Handler (a-elchha.ads)...	263
12.13	Ada.Exceptions.Traceback (a-extra.ads).....	264
12.14	Ada.Sequential_IO.C_Streams (a-siocst.ads).....	264
12.15	Ada.Streams.Stream_IO.C_Streams (a-ssicst.ads) .....	264
12.16	Ada.Strings.Unbounded.Text_IO (a-suteio.ads).....	264
12.17	Ada.Strings.Wide_Unbounded.Wide_Text_IO (a-swuwti.ads)..	264
12.18	Ada.Strings.Wide_Wide_Unbounded.Wide_ Wide_Text_IO (a-szuzti.ads) .....	264
12.19	Ada.Task_Initialization (a-tasini.ads).....	264
12.20	Ada.Text_IO.C_Streams (a-tiocst.ads) .....	264
12.21	Ada.Text_IO.Reset_Standard_Files (a-tirsfi.ads) .....	264
12.22	Ada.Wide_Characters.Unicode (a-wichun.ads).....	265
12.23	Ada.Wide_Text_IO.C_Streams (a-wtcstr.ads) .....	265
12.24	Ada.Wide_Text_IO.Reset_Standard_Files (a-wrstfi.ads)..	265
12.25	Ada.Wide_Wide_Characters.Unicode (a-zchuni.ads) .....	265
12.26	Ada.Wide_Wide_Text_IO.C_Streams (a-ztcstr.ads) .....	265
12.27	Ada.Wide_Wide_Text_IO.Reset_ Standard_Files (a-zrstfi.ads) .....	265
12.28	GNAT.Altivec (g-altive.ads).....	265
12.29	GNAT.Altivec.Conversions (g-altcon.ads).....	265
12.30	GNAT.Altivec.Vector_Operations (g-alveop.ads) .....	266
12.31	GNAT.Altivec.Vector_Types (g-alvety.ads) .....	266
12.32	GNAT.Altivec.Vector_Views (g-alvevi.ads) .....	266
12.33	GNAT.Array_Split (g-arrspl.ads) .....	266
12.34	GNAT.AWK (g-awk.ads) .....	266
12.35	GNAT.Binary_Search (g-binsea.ads).....	266
12.36	GNAT.Bind_Environment (g-binenv.ads) .....	266
12.37	GNAT.Branch_Prediction (g-brapre.ads).....	266

12.38	GNAT.Bounded_Buffers (g-boubuf.ads) .....	266
12.39	GNAT.Bounded_Mailboxes (g-boumai.ads) .....	266
12.40	GNAT.Bubble_Sort (g-bubsor.ads) .....	267
12.41	GNAT.Bubble_Sort_A (g-busora.ads) .....	267
12.42	GNAT.Bubble_Sort_G (g-busorg.ads) .....	267
12.43	GNAT.Byte_Order_Mark (g-byorma.ads) .....	267
12.44	GNAT.Byte_Swapping (g-bytswa.ads) .....	267
12.45	GNAT.C_Time (g-c_time.ads) .....	267
12.46	GNAT.Calendar (g-calend.ads) .....	267
12.47	GNAT.Calendar.Time_IO (g-catiio.ads) .....	267
12.48	GNAT.CRC32 (g-crc32.ads) .....	267
12.49	GNAT.Case_Util (g-casuti.ads) .....	267
12.50	GNAT.CGI (g-cgi.ads) .....	268
12.51	GNAT.CGI.Cookie (g-cgicoo.ads) .....	268
12.52	GNAT.CGI.Debug (g-cgideb.ads) .....	268
12.53	GNAT.Command_Line (g-comlin.ads) .....	268
12.54	GNAT.Compiler_Version (g-comver.ads) .....	268
12.55	GNAT.Ctrl_C (g-ctrl_c.ads) .....	268
12.56	GNAT.Current_Exception (g-curexc.ads) .....	268
12.57	GNAT.Debug_Pools (g-debpoo.ads) .....	268
12.58	GNAT.Debug_Uutilities (g-debuti.ads) .....	268
12.59	GNAT.Decode_String (g-decstr.ads) .....	269
12.60	GNAT.Decode_UTF8_String (g-deutst.ads) .....	269
12.61	GNAT.Directory_Operations (g-dirope.ads) .....	269
12.62	GNAT.Directory_Operations.Iteration (g-diopit.ads) ..	269
12.63	GNAT.Dynamic_HTables (g-dynhta.ads) .....	269
12.64	GNAT.Dynamic_Tables (g-dyntab.ads) .....	269
12.65	GNAT.Encode_String (g-encstr.ads) .....	269
12.66	GNAT.Encode_UTF8_String (g-enutst.ads) .....	269
12.67	GNAT.Exception_Actions (g-exact.ads) .....	270
12.68	GNAT.Exception_Traces (g-extra.ads) .....	270
12.69	GNAT.Exceptions (g-except.ads) .....	270
12.70	GNAT.Expect (g-expect.ads) .....	270
12.71	GNAT.Expect.TTY (g-exptty.ads) .....	270
12.72	GNAT.Float_Control (g-flocon.ads) .....	270
12.73	GNAT.Formatted_String (g-forstr.ads) .....	270
12.74	GNAT.Generic_Fast_Math_Functions (g-gfmafu.ads) .....	270
12.75	GNAT.Heap_Sort (g-heasor.ads) .....	271
12.76	GNAT.Heap_Sort_A (g-hesora.ads) .....	271
12.77	GNAT.Heap_Sort_G (g-hesorg.ads) .....	271
12.78	GNAT.HTable (g-htable.ads) .....	271
12.79	GNAT.IO (g-io.ads) .....	271
12.80	GNAT.IO_Aux (g-io_aux.ads) .....	271
12.81	GNAT.Lock_Files (g-locfil.ads) .....	271
12.82	GNAT.MBBS_Discrete_Random (g-mbdira.ads) .....	272
12.83	GNAT.MBBS_Float_Random (g-mbflra.ads) .....	272
12.84	GNAT.MD5 (g-md5.ads) .....	272

12.85	GNAT.Memory_Dump (g-memdum.ads) .....	272
12.86	GNAT.Most_Recent_Exception (g-moreex.ads) .....	272
12.87	GNAT.OS_Lib (g-os_lib.ads) .....	272
12.88	GNAT.Perfect_Hash_Generators (g-pehage.ads) .....	272
12.89	GNAT.Random_Numbers (g-rannum.ads) .....	272
12.90	GNAT.Regexp (g-regexp.ads) .....	273
12.91	GNAT.Registry (g-regist.ads) .....	273
12.92	GNAT.Regpat (g-regpat.ads) .....	273
12.93	GNAT.Rewrite_Data (g-rewdat.ads) .....	273
12.94	GNAT.Secondary_Stack_Info (g-sestin.ads) .....	273
12.95	GNAT.Semaphores (g-semaph.ads) .....	273
12.96	GNAT.Serial_Communications (g-sercom.ads) .....	273
12.97	GNAT.SHA1 (g-sha1.ads) .....	273
12.98	GNAT.SHA224 (g-sha224.ads) .....	273
12.99	GNAT.SHA256 (g-sha256.ads) .....	274
12.100	GNAT.SHA384 (g-sha384.ads) .....	274
12.101	GNAT.SHA512 (g-sha512.ads) .....	274
12.102	GNAT.Signals (g-signal.ads) .....	274
12.103	GNAT.Sockets (g-socket.ads) .....	274
12.104	GNAT.Source_Info (g-souinf.ads) .....	274
12.105	GNAT.Spelling_Checker (g-speche.ads) .....	274
12.106	GNAT.Spelling_Checker_Generic (g-spchge.ads) .....	274
12.107	GNAT.Spitbol.Patterns (g-spiPAT.ads) .....	274
12.108	GNAT.Spitbol (g-spitbo.ads) .....	275
12.109	GNAT.Spitbol.Table_Boolean (g-sptabo.ads) .....	275
12.110	GNAT.Spitbol.Table_Integer (g-sptain.ads) .....	275
12.111	GNAT.Spitbol.Table_VString (g-sptavs.ads) .....	275
12.112	GNAT.SSE (g-sse.ads) .....	275
12.113	GNAT.SSE.Vector_Types (g-ssvety.ads) .....	275
12.114	GNAT.String_Hash (g-strhas.ads) .....	275
12.115	GNAT.Strings (g-string.ads) .....	275
12.116	GNAT.String_Split (g-strspl.ads) .....	275
12.117	GNAT.Table (g-table.ads) .....	276
12.118	GNAT.Task_Lock (g-tasloc.ads) .....	276
12.119	GNAT.Time_Stamp (g-timsta.ads) .....	276
12.120	GNAT.Threads (g-thread.ads) .....	276
12.121	GNAT.Traceback (g-traceb.ads) .....	276
12.122	GNAT.Traceback.Symbolic (g-trasym.ads) .....	276
12.123	GNAT.UTF_32 (g-utf_32.ads) .....	276
12.124	GNAT.UTF_32_Spelling_Checker (g-u3spch.ads) .....	276
12.125	GNAT.Wide_Spelling_Checker (g-wispch.ads) .....	277
12.126	GNAT.Wide_String_Split (g-wistsp.ads) .....	277
12.127	GNAT.Wide_Wide_Spelling_Checker (g-zspche.ads) .....	277
12.128	GNAT.Wide_Wide_String_Split (g-zistsp.ads) .....	277
12.129	Interfaces.C.Extensions (i-cexten.ads) .....	277
12.130	Interfaces.C.Streams (i-cstrea.ads) .....	277
12.131	Interfaces.Packed_Decimal (i-pacdec.ads) .....	277

12.132	Interfaces.VxWorks (i-vxwork.ads).....	277
12.133	Interfaces.VxWorks.IO (i-vxwoio.ads).....	277
12.134	System.Address_Image (s-addima.ads).....	277
12.135	System.Assertions (s-assert.ads).....	278
12.136	System.Atomic_Counters (s-atocou.ads).....	278
12.137	System.Memory (s-memory.ads).....	278
12.138	System.Multiprocessors (s-multip.ads).....	278
12.139	System.Multiprocessors.Dispatching_Domains (s-mudido.ads) ..	278
12.140	System.Partition_Interface (s-parint.ads).....	278
12.141	System.Pool_Global (s-pooglo.ads).....	278
12.142	System.Pool_Local (s-pooloc.ads).....	279
12.143	System.Restrictions (s-restri.ads).....	279
12.144	System.Rident (s-rident.ads).....	279
12.145	System.Strings.Stream_Ops (s-ststop.ads).....	279
12.146	System.Unsigned_Types (s-unstyp.ads).....	279
12.147	System.Wch_Cnv (s-wchcnv.ads).....	279
12.148	System.Wch_Con (s-wchcon.ads).....	279
<b>13</b>	<b>Interfacing to Other Languages .....</b>	<b>280</b>
13.1	Interfacing to C.....	280
13.2	Interfacing to C++.....	281
13.3	Interfacing to COBOL.....	284
13.4	Interfacing to Fortran.....	284
13.5	Interfacing to non-GNAT Ada code.....	284
<b>14</b>	<b>Specialized Needs Annexes .....</b>	<b>286</b>
<b>15</b>	<b>Implementation of Specific Ada Features ..</b>	<b>287</b>
15.1	Machine Code Insertions .....	287
15.2	GNAT Implementation of Tasking.....	289
15.2.1	Mapping Ada Tasks onto the Underlying Kernel Threads ..	289
15.2.2	Ensuring Compliance with the Real-Time Annex.....	290
15.2.3	Support for Locking Policies .....	290
15.3	GNAT Implementation of Shared Passive Packages .....	291
15.4	Code Generation for Array Aggregates.....	292
15.4.1	Static constant aggregates with static bounds.....	292
15.4.2	Constant aggregates with unconstrained nominal types ..	293
15.4.3	Aggregates with static bounds .....	293
15.4.4	Aggregates with nonstatic bounds .....	293
15.4.5	Aggregates in assignment statements .....	293
15.5	The Size of Discriminated Records with Default Discriminants ..	294
15.6	Image Values For Nonscalar Types .....	295
15.7	Strict Conformance to the Ada Reference Manual.....	295

<b>16</b>	<b>Implementation of Ada 2022 Features . . . . .</b>	<b>297</b>
<b>17</b>	<b>GNAT language extensions . . . . .</b>	<b>330</b>
17.1	How to activate the extended GNAT Ada superset . . . . .	330
17.2	Curated Extensions . . . . .	330
17.2.1	Local Declarations Without Block . . . . .	330
17.2.2	Deep delta Aggregates . . . . .	332
17.2.2.1	Syntax . . . . .	332
17.2.2.2	Legality Rules . . . . .	333
17.2.2.3	Dynamic Semantics . . . . .	333
17.2.2.4	Examples . . . . .	334
17.2.3	Fixed lower bounds for array types and subtypes . . . . .	334
17.2.4	Prefixed-view notation for calls to primitive subprograms of untagged types . . . . .	335
17.2.5	Expression defaults for generic formal functions . . . . .	336
17.2.6	String interpolation . . . . .	336
17.2.7	Constrained attribute for generic objects . . . . .	337
17.2.8	Static aspect on intrinsic functions . . . . .	337
17.2.9	First Controlling Parameter . . . . .	338
17.2.10	Generalized Finalization . . . . .	339
17.2.10.1	Finalizable tagged types . . . . .	341
17.2.10.2	Composite types . . . . .	341
17.2.10.3	Interoperability with controlled types . . . . .	341
17.3	Experimental Language Extensions . . . . .	341
17.3.1	Conditional when constructs . . . . .	341
17.3.2	Implicit With . . . . .	342
17.3.3	Storage Model . . . . .	342
17.3.3.1	Aspect Storage_Model_Type . . . . .	343
17.3.3.2	Aspect Designated_Storage_Model . . . . .	345
17.3.3.3	Legacy Storage Pools . . . . .	346
17.3.4	Attribute Super . . . . .	347
17.3.5	Simpler Accessibility Model . . . . .	347
17.3.5.1	Stand-alone objects . . . . .	348
17.3.5.2	Subprogram parameters . . . . .	348
17.3.5.3	Function results . . . . .	350
17.3.6	Case pattern matching . . . . .	352
17.3.7	Mutably Tagged Types with Size'Class Aspect . . . . .	354
17.3.8	No_Raise aspect . . . . .	356
17.3.9	Inference of Dependent Types in Generic Instantiations . . . . .	356
17.3.10	External_Initialization Aspect . . . . .	357
17.3.11	Finally construct . . . . .	358
17.3.11.1	Syntax . . . . .	358
17.3.11.2	Legality Rules . . . . .	358
17.3.11.3	Dynamic Semantics . . . . .	358
17.3.12	Continue statement . . . . .	359
17.3.13	Destructors . . . . .	359

17.3.14	Structural Generic Instantiation .....	360
17.3.14.1	Syntax .....	360
17.3.14.2	Legality Rules .....	360
17.3.14.3	Static Semantics .....	360
<b>18</b>	<b>Security Hardening Features .....</b>	<b>363</b>
18.1	Register Scrubbing .....	363
18.2	Stack Scrubbing .....	363
18.3	Hardened Conditionals .....	365
18.4	Hardened Booleans .....	367
18.5	Control Flow Redundancy .....	368
<b>19</b>	<b>Obsolescent Features .....</b>	<b>371</b>
19.1	PolyORB .....	371
19.2	pragma No_Run_Time .....	371
19.3	pragma Ravenscar .....	371
19.4	pragma Restricted_Run_Time .....	371
19.5	pragma Task_Info .....	371
19.6	package System.Task_Info (s-tasinf.ads) .....	372
<b>20</b>	<b>Compatibility and Porting Guide .....</b>	<b>373</b>
20.1	Writing Portable Fixed-Point Declarations .....	373
20.2	Compatibility with Ada 83 .....	374
20.2.1	Legal Ada 83 programs that are illegal in Ada 95 .....	374
20.2.2	More deterministic semantics .....	376
20.2.3	Changed semantics .....	376
20.2.4	Other language compatibility issues .....	376
20.3	Compatibility between Ada 95 and Ada 2005 .....	377
20.4	Implementation-dependent characteristics .....	378
20.4.1	Implementation-defined pragmas .....	378
20.4.2	Implementation-defined attributes .....	378
20.4.3	Libraries .....	378
20.4.4	Elaboration order .....	378
20.4.5	Target-specific aspects .....	379
20.5	Compatibility with Other Ada Systems .....	379
20.6	Representation Clauses .....	380
20.7	Compatibility with HP Ada 83 .....	381
<b>21</b>	<b>GNU Free Documentation License .....</b>	<b>382</b>
	<b>Index .....</b>	<b>389</b>

‘GNAT, The GNU Ada Development Environment’

GCC version 16.0.1

AdaCore

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, with the Front-Cover Texts being “GNAT Reference Manual”, and with no Back-Cover Texts. A copy of the license is included in the section entitled [GNU Free Documentation License], page 381.

# 1 About This Guide

This manual contains useful information in writing programs using the GNAT compiler. It includes information on implementation dependent characteristics of GNAT, including all the information required by Annex M of the Ada language standard.

GNAT implements Ada 95, Ada 2005, Ada 2012 and Ada 2022, and it may also be invoked in Ada 83 compatibility mode. By default, GNAT assumes Ada 2012, but you can override with a compiler switch to explicitly specify the language version. (Please refer to the ‘GNAT User’s Guide’ for details on these switches.) Throughout this manual, references to ‘Ada’ without a year suffix apply to all the Ada versions of the language.

Ada is designed to be highly portable. In general, a program will have the same effect even when compiled by different compilers on different platforms. However, since Ada is designed to be used in a wide variety of applications, it also contains a number of system dependent features to be used in interfacing to the external world.

Note: Any program that makes use of implementation-dependent features may be non-portable. You should follow good programming practice and isolate and clearly document any sections of your program that make use of these features in a non-portable manner.

## 1.1 What This Reference Manual Contains

This reference manual contains the following chapters:

- \* [Implementation Defined Pragmas], page 4, lists GNAT implementation-dependent pragmas, which can be used to extend and enhance the functionality of the compiler.
- \* [Implementation Defined Attributes], page 120, lists GNAT implementation-dependent attributes, which can be used to extend and enhance the functionality of the compiler.
- \* [Standard and Implementation Defined Restrictions], page 144, lists GNAT implementation-dependent restrictions, which can be used to extend and enhance the functionality of the compiler.
- \* [Implementation Advice], page 159, provides information on generally desirable behavior which are not requirements that all compilers must follow since it cannot be provided on all systems, or which may be undesirable on some systems.
- \* [Implementation Defined Characteristics], page 179, provides a guide to minimizing implementation dependent features.
- \* [Intrinsic Subprograms], page 198, describes the intrinsic subprograms implemented by GNAT, and how they can be imported into user application programs.
- \* [Representation Clauses and Pragmas], page 201, describes in detail the way that GNAT represents data, and in particular the exact set of representation clauses and pragmas that is accepted.
- \* [Standard Library Routines], page 232, provides a listing of packages and a brief description of the functionality that is provided by Ada’s extensive set of standard library routines as implemented by GNAT.
- \* [The Implementation of Standard I/O], page 243, details how the GNAT implementation of the input-output facilities.
- \* [The GNAT Library], page 261, is a catalog of packages that complement the Ada predefined library.

- \* [Interfacing to Other Languages], page 279, describes how programs written in Ada using GNAT can be interfaced to other programming languages.
- \* [Specialized Needs Annexes], page 285, describes the GNAT implementation of all of the specialized needs annexes.
- \* [Implementation of Specific Ada Features], page 286, discusses issues related to GNAT's implementation of machine code insertions, tasking, and several other features.
- \* [Implementation of Ada 2022 Features], page 296, describes the status of the GNAT implementation of the Ada 2022 language standard.
- \* [Security Hardening Features], page 362, documents GNAT extensions aimed at security hardening.
- \* [Obsolescent Features], page 370, documents implementation dependent features, including pragmas and attributes, which are considered obsolescent, since there are other preferred ways of achieving the same results. These obsolescent forms are retained for backwards compatibility.
- \* [Compatibility and Porting Guide], page 372, presents some guidelines for developing portable Ada code, describes the compatibility issues that may arise between GNAT and other Ada compilation systems (including those for Ada 83), and shows how GNAT can expedite porting applications developed in other Ada environments.
- \* [GNU Free Documentation License], page 381, contains the license for this document.

This reference manual assumes a basic familiarity with the Ada 95 language, as described in the *International Standard ANSI/ISO/IEC-8652:1995*. It does not require knowledge of the new features introduced by Ada 2005 or Ada 2012. All three reference manuals are included in the GNAT documentation package.

## 1.2 Conventions

Following are examples of the typographical and graphic conventions used in this guide:

- \* **Functions, utility program names, standard names, and classes.**
- \* **Option flags**
- \* **File names**
- \* **Variables**
- \* ‘Emphasis’
- \* [optional information or parameters]
- \* Examples are described by text  
    **and then shown this way.**
- \* Commands that are entered by the user are shown as preceded by a prompt string comprising the \$ character followed by a space.

## 1.3 Related Information

See the following documents for further information on GNAT:

- \* *GNAT User's Guide for Native Platforms*, which provides information on how to use the GNAT development environment.

- \* *Ada 95 Reference Manual*, the Ada 95 programming language standard.
- \* *Ada 95 Annotated Reference Manual*, which is an annotated version of the Ada 95 standard. The annotations describe detailed aspects of the design decision, and in particular contain useful sections on Ada 83 compatibility.
- \* *Ada 2005 Reference Manual*, the Ada 2005 programming language standard.
- \* *Ada 2005 Annotated Reference Manual*, which is an annotated version of the Ada 2005 standard. The annotations describe detailed aspects of the design decision.
- \* *Ada 2012 Reference Manual*, the Ada 2012 programming language standard.
- \* *DEC Ada, Technical Overview and Comparison on DIGITAL Platforms*, which contains specific information on compatibility between GNAT and DEC Ada 83 systems.
- \* *DEC Ada, Language Reference Manual*, part number AA-PYZAB-TK, which describes in detail the pragmas and attributes provided by the DEC Ada 83 compiler system.

## 2 Implementation Defined Pragmas

Ada defines a set of pragmas that can be used to supply additional information to the compiler. These language defined pragmas are implemented in GNAT and work as described in the Ada Reference Manual.

In addition, Ada allows implementations to define additional pragmas whose meaning is defined by the implementation. GNAT provides a number of these implementation-defined pragmas, which can be used to extend and enhance the functionality of the compiler. This section of the GNAT Reference Manual describes these additional pragmas.

Note that any program using these pragmas might not be portable to other compilers (although GNAT implements this set of pragmas on all platforms). Therefore if portability to other compilers is an important consideration, the use of these pragmas should be minimized.

### 2.1 Pragma Abort\_Defer

Syntax:

```
pragma Abort_Defer;
```

This pragma must appear at the start of the statement sequence of a handled sequence of statements (right after the `begin`). It has the effect of deferring aborts for the sequence of statements (but not for the declarations or handlers, if any, associated with this statement sequence). This can also be useful for adding a polling point in Ada code, where asynchronous abort of tasks is checked when leaving the statement sequence, and is lighter than, for example, using `delay 0.0;`, since with zero-cost exception handling, propagating exceptions (implicitly used to implement task abort) cannot be done reliably in an asynchronous way.

An example of usage would be:

```
-- Add a polling point to check for task aborts

begin
    pragma Abort_Defer;
end;
```

### 2.2 Pragma Abstract\_State

Syntax:

```
pragma Abstract_State (ABSTRACT_STATE_LIST);

ABSTRACT_STATE_LIST ::=
    null
  | STATE_NAME_WITH_OPTIONS
  | (STATE_NAME_WITH_OPTIONS {, STATE_NAME_WITH_OPTIONS} )

STATE_NAME_WITH_OPTIONS ::=
    STATE_NAME
  | (STATE_NAME with OPTION_LIST)
```

```

OPTION_LIST ::= OPTION {, OPTION}

OPTION ::=
    SIMPLE_OPTION
  | NAME_VALUE_OPTION

SIMPLE_OPTION ::= Ghost | Synchronous

NAME_VALUE_OPTION ::=
    Part_Of => ABSTRACT_STATE
  | External [=> EXTERNAL_PROPERTY_LIST]

EXTERNAL_PROPERTY_LIST ::=
    EXTERNAL_PROPERTY
  | (EXTERNAL_PROPERTY {, EXTERNAL_PROPERTY} )

EXTERNAL_PROPERTY ::=
    Async_Readers      [=> static_boolean_EXPRESSION]
  | Async_Writers      [=> static_boolean_EXPRESSION]
  | Effective_Reads    [=> static_boolean_EXPRESSION]
  | Effective_Writes   [=> static_boolean_EXPRESSION]
  | others              => static_boolean_EXPRESSION

STATE_NAME ::= defining_identifier

ABSTRACT_STATE ::= name

```

For the semantics of this pragma, see the entry for aspect `Abstract_State` in the SPARK 2014 Reference Manual, section 7.1.4.

## 2.3 Pragma `Ada_83`

Syntax:

```
pragma Ada_83;
```

A configuration pragma that establishes Ada 83 mode for the unit to which it applies, regardless of the mode set by the command line switches. In Ada 83 mode, GNAT attempts to be as compatible with the syntax and semantics of Ada 83, as defined in the original Ada 83 Reference Manual as possible. In particular, the keywords added by Ada 95 and Ada 2005 are not recognized, optional package bodies are allowed, and generics may name types with unknown discriminants without using the (<>) notation. In addition, some but not all of the additional restrictions of Ada 83 are enforced.

Like all configuration pragmas, if the pragma is placed before a library level package specification it is not propagated to the corresponding package body (see RM 10.1.5(8)); it must be added explicitly to the package body.

Ada 83 mode is intended for two purposes. Firstly, it allows existing Ada 83 code to be compiled and adapted to GNAT with less effort. Secondly, it aids in keeping code backwards

compatible with Ada 83. However, there is no guarantee that code that is processed correctly by GNAT in Ada 83 mode will in fact compile and execute with an Ada 83 compiler, since GNAT does not enforce all the additional checks required by Ada 83.

## 2.4 Pragma Ada\_95

Syntax:

```
pragma Ada_95;
```

A configuration pragma that establishes Ada 95 mode for the unit to which it applies, regardless of the mode set by the command line switches. This mode is set automatically for the `Ada` and `System` packages and their children, so you need not specify it in these contexts. This pragma is useful when writing a reusable component that itself uses Ada 95 features, but which is intended to be usable from either Ada 83 or Ada 95 programs.

Like all configuration pragmas, if the pragma is placed before a library level package specification it is not propagated to the corresponding package body (see RM 10.1.5(8)); it must be added explicitly to the package body.

## 2.5 Pragma Ada\_05

Syntax:

```
pragma Ada_05;  
pragma Ada_05 (local_NAME);
```

A configuration pragma that establishes Ada 2005 mode for the unit to which it applies, regardless of the mode set by the command line switches. This pragma is useful when writing a reusable component that itself uses Ada 2005 features, but which is intended to be usable from either Ada 83 or Ada 95 programs.

Like all configuration pragmas, if the pragma is placed before a library level package specification it is not propagated to the corresponding package body (see RM 10.1.5(8)); it must be added explicitly to the package body.

The one argument form (which is not a configuration pragma) is used for managing the transition from Ada 95 to Ada 2005 in the run-time library. If an entity is marked as `Ada_2005` only, then referencing the entity in `Ada_83` or `Ada_95` mode will generate a warning. In addition, in `Ada_83` or `Ada_95` mode, a preference rule is established which does not choose such an entity unless it is unambiguously specified. This avoids extra subprograms marked this way from generating ambiguities in otherwise legal pre-Ada\_2005 programs. The one argument form is intended for exclusive use in the GNAT run-time library.

## 2.6 Pragma Ada\_2005

Syntax:

```
pragma Ada_2005;
```

This configuration pragma is a synonym for `pragma Ada_05` and has the same syntax and effect.

## 2.7 Pragma Ada\_12

Syntax:

```
pragma Ada_12;  
pragma Ada_12 (local_NAME);
```

A configuration pragma that establishes Ada 2012 mode for the unit to which it applies, regardless of the mode set by the command line switches. This mode is set automatically for the `Ada` and `System` packages and their children, so you need not specify it in these contexts. This pragma is useful when writing a reusable component that itself uses Ada 2012 features, but which is intended to be usable from Ada 83, Ada 95, or Ada 2005 programs.

Like all configuration pragmas, if the pragma is placed before a library level package specification it is not propagated to the corresponding package body (see RM 10.1.5(8)); it must be added explicitly to the package body.

The one argument form, which is not a configuration pragma, is used for managing the transition from Ada 2005 to Ada 2012 in the run-time library. If an entity is marked as `Ada_2012` only, then referencing the entity in any pre-`Ada_2012` mode will generate a warning. In addition, in any pre-`Ada_2012` mode, a preference rule is established which does not choose such an entity unless it is unambiguously specified. This avoids extra subprograms marked this way from generating ambiguities in otherwise legal pre-`Ada_2012` programs. The one argument form is intended for exclusive use in the GNAT run-time library.

## 2.8 Pragma Ada\_2012

Syntax:

```
pragma Ada_2012;
```

This configuration pragma is a synonym for `pragma Ada_12` and has the same syntax and effect.

## 2.9 Pragma Ada\_2022

Syntax:

```
pragma Ada_2022;  
pragma Ada_2022 (local_NAME);
```

A configuration pragma that establishes Ada 2022 mode for the unit to which it applies, regardless of the mode set by the command line switches. This mode is set automatically for the `Ada` and `System` packages and their children, so you need not specify it in these contexts. This pragma is useful when writing a reusable component that itself uses Ada 2022 features, but which is intended to be usable from Ada 83, Ada 95, Ada 2005 or Ada 2012 programs.

Like all configuration pragmas, if the pragma is placed before a library level package specification it is not propagated to the corresponding package body (see RM 10.1.5(8)); it must be added explicitly to the package body.

The one argument form, which is not a configuration pragma, is used for managing the transition from Ada 2012 to Ada 2022 in the run-time library. If an entity is marked as `Ada_2022` only, then referencing the entity in any pre-`Ada_2022` mode will generate a

warning. In addition, in any pre-Ada\_2012 mode, a preference rule is established which does not choose such an entity unless it is unambiguously specified. This avoids extra subprograms marked this way from generating ambiguities in otherwise legal pre-Ada\_2022 programs. The one argument form is intended for exclusive use in the GNAT run-time library.

## 2.10 Pragma Aggregate\_Individually\_Assign

Syntax:

```
pragma Aggregate_Individually_Assign;
```

Where possible, GNAT will store the binary representation of a record aggregate in memory for space and performance reasons. This configuration pragma changes this behavior so that record aggregates are instead always converted into individual assignment statements.

## 2.11 Pragma Allow\_Integer\_Address

Syntax:

```
pragma Allow_Integer_Address;
```

In almost all versions of GNAT, `System.Address` is a private type in accordance with the implementation advice in the RM. This means that integer values, in particular integer literals, are not allowed as address values. If the configuration pragma `Allow_Integer_Address` is given, then integer expressions may be used anywhere a value of type `System.Address` is required. The effect is to introduce an implicit unchecked conversion from the integer value to type `System.Address`. The reverse case of using an address where an integer type is required is handled analogously. The following example compiles without errors:

```
pragma Allow_Integer_Address;
with System; use System;
package AddrAsInt is
  X : Integer;
  Y : Integer;
  for X'Address use 16#1240#;
  for Y use at 16#3230#;
  m : Address := 16#4000#;
  n : constant Address := 4000;
  p : constant Address := Address (X + Y);
  v : Integer := y'Address;
  w : constant Integer := Integer (Y'Address);
  type R is new integer;
  RR : R := 1000;
  Z : Integer;
  for Z'Address use RR;
end AddrAsInt;
```

Note that pragma `Allow_Integer_Address` is ignored if `System.Address` is not a private type. In implementations of GNAT where `System.Address` is a visible integer type, this pragma serves no purpose but is ignored rather than rejected to allow common sets of sources to be used in the two situations.

## 2.12 Pragma Always\_Terminates

Syntax:

```
pragma Always_Terminates [ (boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Always_Terminates` in the SPARK 2014 Reference Manual, section 6.1.11.

## 2.13 Pragma Annotate

Syntax:

```
pragma Annotate (IDENTIFIER [, IDENTIFIER {, ARG}] [, entity => local_NAME]);
```

```
ARG ::= NAME | EXPRESSION
```

This pragma is used to annotate programs. `IDENTIFIER` identifies the type of annotation. GNAT verifies that it is an identifier, but does not otherwise analyze it. The second optional identifier is also left unanalyzed, and by convention is used to control the action of the tool to which the annotation is addressed. The remaining `ARG` arguments can be either string literals or more generally expressions. String literals (and concatenations of string literals) are assumed to be either of type `Standard.String` or else `Wide_String` or `Wide_Wide_String` depending on the character literals they contain. All other kinds of arguments are analyzed as expressions, and must be unambiguous. The last argument if present must have the identifier `Entity` and GNAT verifies that a local name is given.

The analyzed pragma is retained in the tree, but not otherwise processed by any part of the GNAT compiler, except to generate corresponding note lines in the generated ALI file. For the format of these note lines, see the compiler source file `lib-writ.ads`. This pragma is intended for use by external tools. The use of pragma `Annotate` does not affect the compilation process in any way. This pragma may be used as a configuration pragma.

## 2.14 Pragma Assert

Syntax:

```
pragma Assert (
  boolean_EXPRESSION
  [, string_EXPRESSION]);
```

The effect of this pragma depends on whether the corresponding command line switch is set to activate assertions. The pragma expands into code equivalent to the following:

```
if assertions-enabled then
  if not boolean_EXPRESSION then
    System.Assertions.Raise_Assert_Failure
      (string_EXPRESSION);
  end if;
end if;
```

The string argument, if given, is the message that will be associated with the exception occurrence if the exception is raised. If no second argument is given, the default message is `file:nnn`, where `file` is the name of the source file containing the assert, and `nnn` is the line number of the assert.

Note that, as with the `if` statement to which it is equivalent, the type of the expression is either `Standard.Boolean`, or any type derived from this standard type.

Assert checks can be either checked or ignored. By default they are ignored. They will be checked if either the command line switch ‘-gnata’ is used, or if an `Assertion_Policy` or `Check_Policy` pragma is used to enable `Assert_Checks`.

If assertions are ignored, then there is no run-time effect (and in particular, any side effects from the expression will not occur at run time). (The expression is still analyzed at compile time, and may cause types to be frozen if they are mentioned here for the first time).

If assertions are checked, then the given expression is tested, and if it is `False` then `System.Assertions.Raise_Assert_Failure` is called which results in the raising of `Assert_Failure` with the given message.

You should generally avoid side effects in the expression arguments of this pragma, because these side effects will turn on and off with the setting of the assertions mode, resulting in assertions that have an effect on the program. However, the expressions are analyzed for semantic correctness whether or not assertions are enabled, so turning assertions on and off cannot affect the legality of a program.

Note that the implementation defined policy `DISABLE`, given in a pragma `Assertion_Policy`, can be used to suppress this semantic analysis.

Note: this is a standard language-defined pragma in versions of Ada from 2005 on. In GNAT, it is implemented in all versions of Ada, and the `DISABLE` policy is an implementation-defined addition.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.15 Pragma `Assert_And_Cut`

Syntax:

```
pragma Assert_And_Cut (
    boolean_EXPRESSION
    [, string_EXPRESSION]);
```

The effect of this pragma is identical to that of pragma `Assert`, except that in an `Assertion_Policy` pragma, the identifier `Assert_And_Cut` is used to control whether it is ignored or checked (or disabled).

The intention is that this be used within a subprogram when the given test expression sums up all the work done so far in the subprogram, so that the rest of the subprogram can be verified (informally or formally) using only the entry preconditions, and the expression in this pragma. This allows dividing up a subprogram into sections for the purposes of testing or formal verification. The pragma also serves as useful documentation.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.16 Pragma `Assertion_Level`

Syntax:

```
pragma Assertion_Level (LEVEL_IDENTIFIER
```

```
[, depends => DEPENDENCY_DESCRIPTOR]);
```

```
DEPENDENCY_DESCRIPTOR ::= LEVEL_IDENTIFIER | LEVEL_IDENTIFIER_LIST
```

```
LEVEL_IDENTIFIER_LIST ::= '[' LEVEL_IDENTIFIER {, LEVEL_IDENTIFIER} ']'
```

For the semantics of this pragma, see the SPARK 2014 Reference Manual, section 11.4.3.

## 2.17 Pragma Assertion\_Policy

Syntax:

```
pragma Assertion_Policy (CHECK | DISABLE | IGNORE | SUPPRESSIBLE);
```

```
pragma Assertion_Policy (
  ASSERTION_KIND => POLICY_IDENTIFIER
  {, ASSERTION_KIND => POLICY_IDENTIFIER});
```

```
ASSERTION_KIND ::= RM_ASSERTION_KIND | ID_ASSERTION_KIND | ASSERTION_LEVEL
```

```
RM_ASSERTION_KIND ::= Assert
                    Static_Predicate
                    Dynamic_Predicate
                    Pre
                    Pre'Class
                    Post
                    Post'Class
                    Type_Invariant
                    Type_Invariant'Class
                    Default_Initial_Condition
```

```
ID_ASSERTION_KIND ::= Assertions
                    Assert_And_Cut
                    Assume
                    Contract_Cases
                    Debug
                    Ghost
                    Initial_Condition
                    Invariant
                    Invariant'Class
                    Loop_Invariant
                    Loop_Variant
                    Postcondition
                    Precondition
                    Predicate
                    Refined_Post
                    Statement_Assertions
                    Subprogram_Variant
```

**POLICY\_IDENTIFIER ::= Check | Disable | Ignore | Suppressible**

This is a standard Ada 2012 pragma that is available as an implementation-defined pragma in earlier versions of Ada. The assertion kinds **RM\_ASSERTION\_KIND** are those defined in the Ada standard. The assertion kinds **ID\_ASSERTION\_KIND** are implementation defined additions recognized by the GNAT compiler.

Additionally the pragma can apply to an assertion level defined by the **Assertion\_Level** pragma. For more details see the SPARK 2014 Reference Manual, section 11.4.2.

The pragma applies in both cases to pragmas and aspects with matching names, e.g. **Pre** applies to the **Pre** aspect, and **Precondition** applies to both the **Precondition** pragma and the aspect **Precondition**. Note that the identifiers for pragmas **Pre\_Class** and **Post\_Class** are **Pre'Class** and **Post'Class** (not **Pre\_Class** and **Post\_Class**), since these pragmas are intended to be identical to the corresponding aspects.

If the policy is **CHECK**, then assertions are enabled, i.e. the corresponding pragma or aspect is activated. If the policy is **IGNORE**, then assertions are ignored, i.e. the corresponding pragma or aspect is deactivated. This pragma overrides the effect of the ‘-gnata’ switch on the command line. If the policy is **SUPPRESSIBLE**, then assertions are enabled by default, however, if the ‘-gnatp’ switch is specified all assertions are ignored.

The implementation defined policy **DISABLE** is like **IGNORE** except that it completely disables semantic checking of the corresponding pragma or aspect. This is useful when the pragma or aspect argument references subprograms in a with'ed package which is replaced by a dummy package for the final build.

The implementation defined assertion kind **Assertions** applies to all assertion kinds. The form with no assertion kind given implies this choice, so it applies to all assertion kinds (RM defined, and implementation defined).

The implementation defined assertion kind **Statement\_Assertions** applies to **Assert**, **Assert\_And\_Cut**, **Assume**, **Loop\_Invariant**, and **Loop\_Variant**.

## 2.18 Pragma Assume

Syntax:

```
pragma Assume (
  boolean_EXPRESSION
  [, string_EXPRESSION]);
```

The effect of this pragma is identical to that of pragma **Assert**, except that in an **Assertion\_Policy** pragma, the identifier **Assume** is used to control whether it is ignored or checked (or disabled).

The intention is that this be used for assumptions about the external environment. So you cannot expect to verify formally or informally that the condition is met, this must be established by examining things outside the program itself. For example, we may have code that depends on the size of **Long\_Long\_Integer** being at least 64. So we could write:

```
pragma Assume (Long_Long_Integer'Size >= 64);
```

This assumption cannot be proved from the program itself, but it acts as a useful run-time check that the assumption is met, and documents the need to ensure that it is met by reference to information outside the program.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.19 Pragma Assume\_No\_Invalid\_Values

Syntax:

```
pragma Assume_No_Invalid_Values (On | Off);
```

This is a configuration pragma that controls the assumptions made by the compiler about the occurrence of invalid representations (invalid values) in the code.

The default behavior (corresponding to an Off argument for this pragma), is to assume that values may in general be invalid unless the compiler can prove they are valid. Consider the following example:

```
V1 : Integer range 1 .. 10;
V2 : Integer range 11 .. 20;
...
for J in V2 .. V1 loop
  ...
end loop;
```

if V1 and V2 have valid values, then the loop is known at compile time not to execute since the lower bound must be greater than the upper bound. However in default mode, no such assumption is made, and the loop may execute. If `Assume_No_Invalid_Values (On)` is given, the compiler will assume that any occurrence of a variable other than in an explicit `'Valid` test always has a valid value, and the loop above will be optimized away.

The use of `Assume_No_Invalid_Values (On)` is appropriate if you know your code is free of uninitialized variables and other possible sources of invalid representations, and may result in more efficient code. A program that accesses an invalid representation with this pragma in effect is erroneous, so no guarantees can be made about its behavior.

It is peculiar though permissible to use this pragma in conjunction with validity checking (`-gnatVa`). In such cases, accessing invalid values will generally give an exception, though formally the program is erroneous so there are no guarantees that this will always be the case, and it is recommended that these two options not be used together.

## 2.20 Pragma Async\_Readers

Syntax:

```
pragma Async_Readers [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Async_Readers` in the SPARK 2014 Reference Manual, section 7.1.2.

## 2.21 Pragma Async\_Writers

Syntax:

```
pragma Async_Writers [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Async_Writers` in the SPARK 2014 Reference Manual, section 7.1.2.

## 2.22 Pragma Attribute\_Definition

Syntax:

```
pragma Attribute_Definition
  ([Attribute =>] ATTRIBUTE_DESIGNATOR,
   [Entity    =>] LOCAL_NAME,
   [Expression =>] EXPRESSION | NAME);
```

If **Attribute** is a known attribute name, this pragma is equivalent to the attribute definition clause:

```
for Entity'Attribute use Expression;
```

If **Attribute** is not a recognized attribute name, the pragma is ignored, and a warning is emitted. This allows source code to be written that takes advantage of some new attribute, while remaining compilable with earlier compilers.

## 2.23 Pragma C\_Pass\_By\_Copy

Syntax:

```
pragma C_Pass_By_Copy
  ([Max_Size =>] static_integer_EXPRESSION);
```

Normally the default mechanism for passing C convention records to C convention subprograms is to pass them by reference, as suggested by RM B.3(69). Use the configuration pragma **C\_Pass\_By\_Copy** to change this default, by requiring that record formal parameters be passed by copy if all of the following conditions are met:

- \* The size of the record type does not exceed the value specified for **Max\_Size**.
- \* The record type has **Convention C**.
- \* The formal parameter has this record type, and the subprogram has a foreign (non-Ada) convention.

If these conditions are met the argument is passed by copy; i.e., in a manner consistent with what C expects if the corresponding formal in the C prototype is a struct (rather than a pointer to a struct).

You can also pass records by copy by specifying the convention **C\_Pass\_By\_Copy** for the record type, or by using the extended **Import** and **Export** pragmas, which allow specification of passing mechanisms on a parameter by parameter basis.

## 2.24 Pragma Check

Syntax:

```
pragma Check (
  [Name    =>] CHECK_KIND,
  [Check   =>] Boolean_EXPRESSION
  [, [Message =>] string_EXPRESSION] );
```

```
CHECK_KIND ::= IDENTIFIER          |
               Pre'Class            |
               Post'Class           |
```

```

Type_Invariant'Class |
Invariant'Class

```

This pragma is similar to the predefined pragma `Assert` except that an extra identifier argument is present. In conjunction with pragma `Check_Policy`, this can be used to define groups of assertions that can be independently controlled. The identifier `Assertion` is special, it refers to the normal set of pragma `Assert` statements.

Checks introduced by this pragma are normally deactivated by default. They can be activated either by the command line option ‘-gnata’, which turns on all checks, or individually controlled using pragma `Check_Policy`.

The identifiers `Assertions` and `Statement_Assertions` are not permitted as check kinds, since this would cause confusion with the use of these identifiers in `Assertion_Policy` and `Check_Policy` pragmas, where they are used to refer to sets of assertions.

## 2.25 Pragma `Check_Float_Overflow`

Syntax:

```
pragma Check_Float_Overflow;
```

In Ada, the predefined floating-point types (`Short_Float`, `Float`, `Long_Float`, `Long_Long_Float`) are defined to be ‘unconstrained’. This means that even though each has a well-defined base range, an operation that delivers a result outside this base range is not required to raise an exception. This implementation permission accommodates the notion of infinities in IEEE floating-point, and corresponds to the efficient execution mode on most machines. GNAT will not raise overflow exceptions on these machines; instead it will generate infinities and NaN’s as defined in the IEEE standard.

Generating infinities, although efficient, is not always desirable. Often the preferable approach is to check for overflow, even at the (perhaps considerable) expense of run-time performance. This can be accomplished by defining your own constrained floating-point subtypes – i.e., by supplying explicit range constraints – and indeed such a subtype can have the same base range as its base type. For example:

```
subtype My_Float is Float range Float'Range;
```

Here `My_Float` has the same range as `Float` but is constrained, so operations on `My_Float` values will be checked for overflow against this range.

This style will achieve the desired goal, but it is often more convenient to be able to simply use the standard predefined floating-point types as long as overflow checking could be guaranteed. The `Check_Float_Overflow` configuration pragma achieves this effect. If a unit is compiled subject to this configuration pragma, then all operations on predefined floating-point types including operations on base types of these floating-point types will be treated as though those types were constrained, and overflow checks will be generated. The `Constraint_Error` exception is raised if the result is out of range.

This mode can also be set by use of the compiler switch ‘-gnateF’.

## 2.26 Pragma `Check_Name`

Syntax:

```
pragma Check_Name (check_name_IDENTIFIER);
```

This is a configuration pragma that defines a new implementation defined check name (unless IDENTIFIER matches one of the predefined check names, in which case the pragma has no effect). Check names are global to a partition, so if two or more configuration pragmas are present in a partition mentioning the same name, only one new check name is introduced.

An implementation defined check name introduced with this pragma may be used in only three contexts: `pragma Suppress`, `pragma Unsuppress`, and as the prefix of a `Check_Name'Enabled` attribute reference. For any of these three cases, the check name must be visible. A check name is visible if it is in the configuration pragmas applying to the current unit, or if it appears at the start of any unit that is part of the dependency set of the current unit (e.g., units that are mentioned in `with` clauses).

Check names introduced by this pragma are subject to control by compiler switches (in particular `-gnatp`) in the usual manner.

## 2.27 Pragma Check\_Policy

Syntax:

```
pragma Check_Policy
  ([Name    =>] CHECK_KIND,
   [Policy =>] POLICY_IDENTIFIER);

pragma Check_Policy (
  CHECK_KIND => POLICY_IDENTIFIER
  {, CHECK_KIND => POLICY_IDENTIFIER});

ASSERTION_KIND ::= RM_ASSERTION_KIND | ID_ASSERTION_KIND

CHECK_KIND ::= IDENTIFIER
              Pre'Class
              Post'Class
              Type_Invariant'Class |
              Invariant'Class
```

The identifiers `Name` and `Policy` are not allowed as `CHECK_KIND` values. This avoids confusion between the two possible syntax forms for this pragma.

```
POLICY_IDENTIFIER ::= ON | OFF | CHECK | DISABLE | IGNORE
```

This pragma is used to set the checking policy for assertions (specified by aspects or pragmas), the `Debug` pragma, or additional checks to be checked using the `Check` pragma. It may appear either as a configuration pragma, or within a declarative part of package. In the latter case, it applies from the point where it appears to the end of the declarative region (like pragma `Suppress`).

The `Check_Policy` pragma is similar to the predefined `Assertion_Policy` pragma, and if the check kind corresponds to one of the assertion kinds that are allowed by `Assertion_Policy`, then the effect is identical.

If the first argument is `Debug`, then the policy applies to `Debug` pragmas, disabling their effect if the policy is `OFF`, `DISABLE`, or `IGNORE`, and allowing them to execute with normal semantics if the policy is `ON` or `CHECK`. In addition if the policy is `DISABLE`, then the procedure call in `Debug` pragmas will be totally ignored and not analyzed semantically.

Finally the first argument may be some other identifier than the above possibilities, in which case it controls a set of named assertions that can be checked using pragma `Check`. For example, if the pragma:

```
pragma Check_Policy (Critical_Error, OFF);
```

is given, then subsequent `Check` pragmas whose first argument is also `Critical_Error` will be disabled.

The check policy is `OFF` to turn off corresponding checks, and `ON` to turn on corresponding checks. The default for a set of checks for which no `Check_Policy` is given is `OFF` unless the compiler switch ‘-gnata’ is given, which turns on all checks by default.

The check policy settings `CHECK` and `IGNORE` are recognized as synonyms for `ON` and `OFF`. These synonyms are provided for compatibility with the standard `Assertion_Policy` pragma. The check policy setting `DISABLE` causes the second argument of a corresponding `Check` pragma to be completely ignored and not analyzed.

## 2.28 Pragma Comment

Syntax:

```
pragma Comment (static_string_EXPRESSION);
```

This is almost identical in effect to pragma `Ident`. It allows the placement of a comment into the object file and hence into the executable file if the operating system permits such usage. The difference is that `Comment`, unlike `Ident`, has no limitations on placement of the pragma (it can be placed anywhere in the main source unit), and if more than one pragma is used, all comments are retained.

## 2.29 Pragma Common\_Object

Syntax:

```
pragma Common_Object (
    [Internal =>] LOCAL_NAME
    [, [External =>] EXTERNAL_SYMBOL]
    [, [Size      =>] EXTERNAL_SYMBOL] );
```

```
EXTERNAL_SYMBOL ::=
    IDENTIFIER
    | static_string_EXPRESSION
```

This pragma enables the shared use of variables stored in overlaid linker areas corresponding to the use of `COMMON` in Fortran. The single object `LOCAL_NAME` is assigned to the area designated by the `External` argument. You may define a record to correspond to a series of fields. The `Size` argument is syntax checked in GNAT, but otherwise ignored.

`Common_Object` is not supported on all platforms. If no support is available, then the code generator will issue a message indicating that the necessary attribute for implementation of this pragma is not available.

## 2.30 Pragma Compile\_Time\_Error

Syntax:

```
pragma Compile_Time_Error
    (boolean_EXPRESSION, static_string_EXPRESSION);
```

This pragma can be used to generate additional compile time error messages. It is particularly useful in generics, where errors can be issued for specific problematic instantiations. The first parameter is a boolean expression. The pragma ensures that the value of an expression is known at compile time, and has the value False. The set of expressions whose values are known at compile time includes all static boolean expressions, and also other values which the compiler can determine at compile time (e.g., the size of a record type set by an explicit size representation clause, or the value of a variable which was initialized to a constant and is known not to have been modified). If these conditions are not met, an error message is generated using the value given as the second argument. This string value may contain embedded ASCII.LF characters to break the message into multiple lines.

## 2.31 Pragma Compile\_Time\_Warning

Syntax:

```
pragma Compile_Time_Warning
    (boolean_EXPRESSION, static_string_EXPRESSION);
```

Same as pragma Compile\_Time\_Error, except a warning is issued instead of an error message. If switch ‘-gnatw-C’ is used, a warning is only issued if the value of the expression is known to be True at compile time, not when the value of the expression is not known at compile time. Note that if this pragma is used in a package that is with’ed by a client, the client will get the warning even though it is issued by a with’ed package (normally warnings in with’ed units are suppressed, but this is a special exception to that rule).

One typical use is within a generic where compile time known characteristics of formal parameters are tested, and warnings given appropriately. Another use with a first parameter of True is to warn a client about use of a package, for example that it is not fully implemented.

In previous versions of the compiler, combining ‘-gnatwe’ with Compile\_Time\_Warning resulted in a fatal error. Now the compiler always emits a warning. You can use [Pragma Compile\_Time\_Error], page 18, to force the generation of an error.

## 2.32 Pragma Complete\_Representation

Syntax:

```
pragma Complete_Representation;
```

This pragma must appear immediately within a record representation clause. Typical placements are before the first component clause or after the last component clause. The effect is to give an error message if any component is missing a component clause. This pragma may be used to ensure that a record representation clause is complete, and that this invariant is maintained if fields are added to the record in the future.

## 2.33 Pragma Complex\_Representation

Syntax:

```
pragma Complex_Representation
    ([Entity =>] LOCAL_NAME);
```

The **Entity** argument must be the name of a record type which has two fields of the same floating-point type. The effect of this pragma is to force gcc to use the special internal complex representation form for this record, which may be more efficient. Note that this may result in the code for this type not conforming to standard ABI (application binary interface) requirements for the handling of record types. For example, in some environments, there is a requirement for passing records by pointer, and the use of this pragma may result in passing this type in floating-point registers.

## 2.34 Pragma Component\_Alignment

Syntax:

```
pragma Component_Alignment (
    [Form =>] ALIGNMENT_CHOICE
    [, [Name =>] type_LOCAL_NAME]);

ALIGNMENT_CHOICE ::=
    Component_Size
  | Component_Size_4
  | Storage_Unit
  | Default
```

Specifies the alignment of components in array or record types. The meaning of the **Form** argument is as follows:

‘Component\_Size’

Aligns scalar components and subcomponents of the array or record type on boundaries appropriate to their inherent size (naturally aligned). For example, 1-byte components are aligned on byte boundaries, 2-byte integer components are aligned on 2-byte boundaries, 4-byte integer components are aligned on 4-byte boundaries and so on. These alignment rules correspond to the normal rules for C compilers on all machines except the VAX.

‘Component\_Size\_4’

Naturally aligns components with a size of four or fewer bytes. Components that are larger than 4 bytes are placed on the next 4-byte boundary.

‘Storage\_Unit’

Specifies that array or record components are byte aligned, i.e., aligned on boundaries determined by the value of the constant **System.Storage\_Unit**.

‘Default’

Specifies that array or record components are aligned on default boundaries, appropriate to the underlying hardware or operating system or both. The **Default** choice is the same as **Component\_Size** (natural alignment).

If the `Name` parameter is present, `type_LOCAL_NAME` must refer to a local record or array type, and the specified alignment choice applies to the specified type. The use of `Component_Alignment` together with a `pragma Pack` causes the `Component_Alignment` pragma to be ignored. The use of `Component_Alignment` together with a record representation clause is only effective for fields not specified by the representation clause.

If the `Name` parameter is absent, the pragma can be used as either a configuration pragma, in which case it applies to one or more units in accordance with the normal rules for configuration pragmas, or it can be used within a declarative part, in which case it applies to types that are declared within this declarative part, or within any nested scope within this declarative part. In either case it specifies the alignment to be applied to any record or array type which has otherwise standard representation.

If the alignment for a record or array type is not specified (using `pragma Pack`, `pragma Component_Alignment`, or a record rep clause), the GNAT uses the default alignment as described previously.

## 2.35 Pragma Constant\_After\_Elaboration

Syntax:

```
pragma Constant_After_Elaboration [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Constant_After_Elaboration` in the SPARK 2014 Reference Manual, section 3.3.1.

## 2.36 Pragma Contract\_Cases

Syntax:

```
pragma Contract_Cases (CONTRACT_CASE {, CONTRACT_CASE});
```

```
CONTRACT_CASE ::= CASE_GUARD => CONSEQUENCE
```

```
CASE_GUARD ::= boolean_EXPRESSION | others
```

```
CONSEQUENCE ::= boolean_EXPRESSION
```

The `Contract_Cases` pragma allows defining fine-grain specifications that can complement or replace the contract given by a precondition and a postcondition. Additionally, the `Contract_Cases` pragma can be used by testing and formal verification tools. The compiler checks its validity and, depending on the assertion policy at the point of declaration of the pragma, it may insert a check in the executable. For code generation, the contract cases

```
pragma Contract_Cases (
  Cond1 => Pred1,
  Cond2 => Pred2);
```

are equivalent to

```
C1 : constant Boolean := Cond1; -- evaluated at subprogram entry
C2 : constant Boolean := Cond2; -- evaluated at subprogram entry
pragma Precondition ((C1 and not C2) or (C2 and not C1));
pragma Postcondition (if C1 then Pred1);
pragma Postcondition (if C2 then Pred2);
```

The precondition ensures that one and only one of the case guards is satisfied on entry to the subprogram. The postcondition ensures that for the case guard that was True on entry, the corresponding consequence is True on exit. Other consequence expressions are not evaluated.

A precondition *P* and postcondition *Q* can also be expressed as contract cases:

```
pragma Contract_Cases (P => Q);
```

The placement and visibility rules for `Contract_Cases` pragmas are identical to those described for preconditions and postconditions.

The compiler checks that boolean expressions given in case guards and consequences are valid, where the rules for case guards are the same as the rule for an expression in `Precondition` and the rules for consequences are the same as the rule for an expression in `Postcondition`. In particular, attributes `'Old` and `'Result` can only be used within consequence expressions. The case guard for the last contract case may be `others`, to denote any case not captured by the previous cases. The following is an example of use within a package spec:

```
package Math_Functions is
...
  function Sqrt (Arg : Float) return Float;
  pragma Contract_Cases (((Arg in 0.0 .. 99.0) => Sqrt'Result < 10.0,
                                Arg >= 100.0           => Sqrt'Result >= 10.0,
                                others                  => Sqrt'Result = 0.0));
...
end Math_Functions;
```

The meaning of contract cases is that only one case should apply at each call, as determined by the corresponding case guard evaluating to True, and that the consequence for this case should hold when the subprogram returns.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.37 Pragma Convention\_Identifier

Syntax:

```
pragma Convention_Identifier (
  [Name =>] IDENTIFIER,
  [Convention =>] convention_IDENTIFIER);
```

This pragma provides a mechanism for supplying synonyms for existing convention identifiers. The `Name` identifier can subsequently be used as a synonym for the given convention in other pragmas (including for example `pragma Import` or another `Convention_Identifier` pragma). As an example of the use of this, suppose you had legacy code which used `Fortran77` as the identifier for Fortran. Then the pragma:

```
pragma Convention_Identifier (Fortran77, Fortran);
```

would allow the use of the convention identifier `Fortran77` in subsequent code, avoiding the need to modify the sources. As another example, you could use this to parameterize convention requirements according to systems. Suppose you needed to use `Stdcall` on windows systems, and `C` on some other system, then you could define a convention identifier

**Library** and use a single **Convention\_Identifier** pragma to specify which convention would be used system-wide.

## 2.38 Pragma **CPP\_Class**

Syntax:

```
pragma CPP_Class ([Entity =>] LOCAL_NAME);
```

The argument denotes an entity in the current declarative region that is declared as a record type. It indicates that the type corresponds to an externally declared C++ class type, and is to be laid out the same way that C++ would lay out the type. If the C++ class has virtual primitives then the record must be declared as a tagged record type.

Types for which **CPP\_Class** is specified do not have assignment or equality operators defined (such operations can be imported or declared as subprograms as required). Initialization is allowed only by constructor functions (see pragma **CPP\_Constructor**). Such types are implicitly limited if not explicitly declared as limited or derived from a limited type, and an error is issued in that case.

See [Interfacing to C++], page 281, for related information.

Note: Pragma **CPP\_Class** is currently obsolete. It is supported for backward compatibility but its functionality is available using pragma **Import** with **Convention = CPP**.

## 2.39 Pragma **CPP\_Constructor**

Syntax:

```
pragma CPP_Constructor ([Entity =>] LOCAL_NAME
    [, [External_Name =>] static_string_EXPRESSION ]
    [, [Link_Name      =>] static_string_EXPRESSION ]);
```

This pragma identifies an imported function (imported in the usual way with pragma **Import**) as corresponding to a C++ constructor. If **External\_Name** and **Link\_Name** are not specified then the **Entity** argument is a name that must have been previously mentioned in a pragma **Import** with **Convention = CPP**. Such name must be of one of the following forms:

- \* 'function' **Fname** 'return' T'
- \* 'function' **Fname** 'return' T'Class
- \* 'function' **Fname** (...) 'return' T'
- \* 'function' **Fname** (...) 'return' T'Class

where T is a limited record type imported from C++ with pragma **Import** and **Convention = CPP**.

The first two forms import the default constructor, used when an object of type T is created on the Ada side with no explicit constructor. The latter two forms cover all the non-default constructors of the type. See the GNAT User's Guide for details.

If no constructors are imported, it is impossible to create any objects on the Ada side and the type is implicitly declared abstract.

Pragma **CPP\_Constructor** is intended primarily for automatic generation using an automatic binding generator tool (such as the **-fdump-ada-spec** GCC switch). See [Interfacing to C++], page 281, for more related information.

Note: The use of functions returning class-wide types for constructors is currently obsolete. They are supported for backward compatibility. The use of functions returning the type T leave the Ada sources more clear because the imported C++ constructors always return an object of type T; that is, they never return an object whose type is a descendant of type T.

## 2.40 Pragma CPP\_Virtual

This pragma is now obsolete and, other than generating a warning if warnings on obsolescent features are enabled, is completely ignored. It is retained for compatibility purposes. It used to be required to ensure compatibility with C++, but is no longer required for that purpose because GNAT generates the same object layout as the G++ compiler by default. See [Interfacing to C++], page 281, for related information.

## 2.41 Pragma CPP\_Vtable

This pragma is now obsolete and, other than generating a warning if warnings on obsolescent features are enabled, is completely ignored. It used to be required to ensure compatibility with C++, but is no longer required for that purpose because GNAT generates the same object layout as the G++ compiler by default.

See [Interfacing to C++], page 281, for related information.

## 2.42 Pragma CPU

Syntax:

```
pragma CPU (EXPRESSION);
```

This pragma is standard in Ada 2012, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

## 2.43 Pragma Deadline\_Floor

Syntax:

```
pragma Deadline_Floor (time_span_EXPRESSION);
```

This pragma applies only to protected types and specifies the floor deadline inherited by a task when the task enters a protected object. It is effective only when the EDF scheduling policy is used.

## 2.44 Pragma Debug

Syntax:

```
pragma Debug ([CONDITION, ]PROCEDURE_CALL_WITHOUT_SEMICOLON);

PROCEDURE_CALL_WITHOUT_SEMICOLON ::=
  PROCEDURE_NAME
  | PROCEDURE_PREFIX ACTUAL_PARAMETER_PART
```

The procedure call argument has the syntactic form of an expression, meeting the syntactic requirements for pragmas.

If debug pragmas are not enabled or if the condition is present and evaluates to False, this pragma has no effect. If debug pragmas are enabled, the semantics of the pragma is exactly equivalent to the procedure call statement corresponding to the argument with a terminating semicolon. Pragmas are permitted in sequences of declarations, so you can use pragma `Debug` to intersperse calls to debug procedures in the middle of declarations. Debug pragmas can be enabled either by use of the command line switch ‘-gnata’ or by use of the pragma `Check_Policy` with a first argument of `Debug`.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.45 Pragma `Debug_Policy`

Syntax:

```
pragma Debug_Policy (CHECK | DISABLE | IGNORE | ON | OFF);
```

This pragma is equivalent to a corresponding `Check_Policy` pragma with a first argument of `Debug`. It is retained for historical compatibility reasons.

## 2.46 Pragma `Default_Initial_Condition`

Syntax:

```
pragma Default_Initial_Condition [ (null | boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Default_Initial_Condition` in the SPARK 2014 Reference Manual, section 7.3.3.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.47 Pragma `Default_Scalar_Storage_Order`

Syntax:

```
pragma Default_Scalar_Storage_Order (High_Order_First | Low_Order_First);
```

Normally if no explicit `Scalar_Storage_Order` is given for a record type or array type, then the scalar storage order defaults to the ordinary default for the target. But this default may be overridden using this pragma. The pragma may appear as a configuration pragma, or locally within a package spec or declarative part. In the latter case, it applies to all subsequent types declared within that package spec or declarative part.

The following example shows the use of this pragma:

```
pragma Default_Scalar_Storage_Order (High_Order_First);
with System; use System;
package DSS01 is
  type H1 is record
    a : Integer;
  end record;

  type L2 is record
    a : Integer;
  end record;
```

```

    for L2'Scalar_Storage_Order use Low_Order_First;

    type L2a is new L2;

    package Inner is
        type H3 is record
            a : Integer;
        end record;

        pragma Default_Scalar_Storage_Order (Low_Order_First);

        type L4 is record
            a : Integer;
        end record;
    end Inner;

    type H4a is new Inner.L4;

    type H5 is record
        a : Integer;
    end record;
end DSS01;

```

In this example record types with names starting with ‘L’ have *Low\_Order\_First* scalar storage order, and record types with names starting with ‘H’ have *High\_Order\_First*. Note that in the case of H4a, the order is not inherited from the parent type. Only an explicitly set *Scalar\_Storage\_Order* gets inherited on type derivation.

If this pragma is used as a configuration pragma which appears within a configuration pragma file (as opposed to appearing explicitly at the start of a single unit), then the binder will require that all units in a partition be compiled in a similar manner, other than run-time units, which are not affected by this pragma. Note that the use of this form is discouraged because it may significantly degrade the run-time performance of the software, instead the default scalar storage order ought to be changed only on a local basis.

## 2.48 Pragma Default\_Storage\_Pool

Syntax:

```
pragma Default_Storage_Pool (storage_pool_NAME | null);
```

This pragma is standard in Ada 2012, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

## 2.49 Pragma Depends

Syntax:

```
pragma Depends (DEPENDENCY_RELATION);
```

```
DEPENDENCY_RELATION ::=
```

```

    null
  | (DEPENDENCY_CLAUSE {, DEPENDENCY_CLAUSE})

DEPENDENCY_CLAUSE ::=
    OUTPUT_LIST =>[+] INPUT_LIST
  | NULL_DEPENDENCY_CLAUSE

NULL_DEPENDENCY_CLAUSE ::= null => INPUT_LIST

OUTPUT_LIST ::= OUTPUT | (OUTPUT {, OUTPUT})

INPUT_LIST ::= null | INPUT | (INPUT {, INPUT})

OUTPUT ::= NAME | FUNCTION_RESULT
INPUT  ::= NAME

```

where `FUNCTION_RESULT` is a function `Result` `attribute_reference`

For the semantics of this pragma, see the entry for aspect `Depends` in the SPARK 2014 Reference Manual, section 6.1.5.

## 2.50 Pragma `Detect_Blocking`

Syntax:

```
pragma Detect_Blocking;
```

This is a standard pragma in Ada 2005, that is available in all earlier versions of Ada as an implementation-defined pragma.

This is a configuration pragma that forces the detection of potentially blocking operations within a protected operation, and to raise `Program_Error` if that happens.

## 2.51 Pragma `Disable_Atomic_Synchronization`

Syntax:

```
pragma Disable_Atomic_Synchronization [(Entity)];
```

```
pragma Enable_Atomic_Synchronization [(Entity)];
```

Ada requires that accesses (reads or writes) of an atomic variable be regarded as synchronization points in the case of multiple tasks. Particularly in the case of multi-processors this may require special handling, e.g. the generation of memory barriers. This synchronization is performed by default, but can be turned off using pragma `Disable_Atomic_Synchronization`. The `Enable_Atomic_Synchronization` pragma turns it back on.

The placement and scope rules for these pragmas are the same as those for `pragma Suppress`. In particular they can be used as configuration pragmas, or in a declaration sequence where they apply until the end of the scope. If an `Entity` argument is present, the action applies only to that entity.

## 2.52 Pragma Dispatching\_Domain

Syntax:

```
pragma Dispatching_Domain (EXPRESSION);
```

This pragma is standard in Ada 2012, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

## 2.53 Pragma Effective\_Reads

Syntax:

```
pragma Effective_Reads [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Effective_Reads` in the SPARK 2014 Reference Manual, section 7.1.2.

## 2.54 Pragma Effective\_Writes

Syntax:

```
pragma Effective_Writes [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Effective_Writes` in the SPARK 2014 Reference Manual, section 7.1.2.

## 2.55 Pragma Elaboration\_Checks

Syntax:

```
pragma Elaboration_Checks (Dynamic | Static);
```

This is a configuration pragma which specifies the elaboration model to be used during compilation. For more information on the elaboration models of GNAT, consult the chapter on elaboration order handling in the ‘GNAT User’s Guide’.

The pragma may appear in the following contexts:

- \* Configuration pragmas file
- \* Prior to the context clauses of a compilation unit’s initial declaration

Any other placement of the pragma will result in a warning and the effects of the offending pragma will be ignored.

If the pragma argument is `Dynamic`, then the dynamic elaboration model is in effect. If the pragma argument is `Static`, then the static elaboration model is in effect.

## 2.56 Pragma Eliminate

Syntax:

```
pragma Eliminate (
    [ Unit_Name      => ] IDENTIFIER | SELECTED_COMPONENT ,
    [ Entity         => ] IDENTIFIER |
                                SELECTED_COMPONENT |
                                STRING_LITERAL
    [, Source_Location => SOURCE_TRACE ] );
```

```
SOURCE_TRACE      ::= STRING_LITERAL
```

This pragma indicates that the given entity is not used in the program to be compiled and built, thus allowing the compiler to eliminate the code or data associated with the named entity. Any reference to an eliminated entity causes a compile-time or link-time error.

The pragma has the following semantics, where **U** is the unit specified by the **Unit\_Name** argument and **E** is the entity specified by the **Entity** argument:

- \* **E** must be a subprogram that is explicitly declared either:
  - \* Within **U**, or
  - \* Within a generic package that is instantiated in **U**, or
  - \* As an instance of generic subprogram instantiated in **U**.

Otherwise the pragma is ignored.

- \* If **E** is overloaded within **U** then, in the absence of a **Source\_Location** argument, all overloadings are eliminated.
- \* If **E** is overloaded within **U** and only some overloadings are to be eliminated, then each overloading to be eliminated must be specified in a corresponding pragma **Eliminate** with a **Source\_Location** argument identifying the line where the declaration appears, as described below.
- \* If **E** is declared as the result of a generic instantiation, then a **Source\_Location** argument is needed, as described below.

Pragma **Eliminate** allows a program to be compiled in a system-independent manner, so that unused entities are eliminated but without needing to modify the source text. Normally the required set of **Eliminate** pragmas is constructed automatically using the **gnatelim** tool.

Any source file change that removes, splits, or adds lines may make the set of **Eliminate** pragmas invalid because their **Source\_Location** argument values may get out of date.

Pragma **Eliminate** may be used where the referenced entity is a dispatching operation. In this case all the subprograms to which the given operation can dispatch are considered to be unused (are never called as a result of a direct or a dispatching call).

The string literal given for the source location specifies the line number of the declaration of the entity, using the following syntax for **SOURCE\_TRACE**:

```
SOURCE_TRACE      ::= SOURCE_REFERENCE [ LBRACKET SOURCE_TRACE RBRACKET ]
```

```
LBRACKET          ::= '['
```

```
RBRACKET          ::= ']'
```

```
SOURCE_REFERENCE ::= FILE_NAME : LINE_NUMBER
```

```
LINE_NUMBER       ::= DIGIT {DIGIT}
```

Spaces around the colon in a **SOURCE\_REFERENCE** are optional.

The source trace that is given as the **Source\_Location** must obey the following rules (or else the pragma is ignored), where **U** is the unit **U** specified by the **Unit\_Name** argument and **E** is the subprogram specified by the **Entity** argument:

- \* `FILE_NAME` is the short name (with no directory information) of the Ada source file for `U`, using the required syntax for the underlying file system (e.g. case is significant if the underlying operating system is case sensitive). If `U` is a package and `E` is a subprogram declared in the package specification and its full declaration appears in the package body, then the relevant source file is the one for the package specification; analogously if `U` is a generic package.
- \* If `E` is not declared in a generic instantiation (this includes generic subprogram instances), the source trace includes only one source line reference. `LINE_NUMBER` gives the line number of the occurrence of the declaration of `E` within the source file (as a decimal literal without an exponent or point).
- \* If `E` is declared by a generic instantiation, its source trace (from left to right) starts with the source location of the declaration of `E` in the generic unit and ends with the source location of the instantiation, given in square brackets. This approach is applied recursively with nested instantiations: the rightmost (nested most deeply in square brackets) element of the source trace is the location of the outermost instantiation, and the leftmost element (that is, outside of any square brackets) is the location of the declaration of `E` in the generic unit.

Examples:

```
pragma Eliminate (Pkg0, Proc);
-- Eliminate (all overloadings of) Proc in Pkg0

pragma Eliminate (Pkg1, Proc,
                  Source_Location => "pkg1.ads:8");
-- Eliminate overloading of Proc at line 8 in pkg1.ads

-- Assume the following file contents:
--   gen_pkg.ads
--   1: generic
--   2:   type T is private;
--   3: package Gen_Pkg is
--   4:   procedure Proc(N : T);
--   ...   ...
--   ... end Gen_Pkg;
--
--   q.adb
--   1: with Gen_Pkg;
--   2: procedure Q is
--   3:   package Inst_Pkg is new Gen_Pkg(Integer);
--   ...   -- No calls on Inst_Pkg.Proc
--   ... end Q;

-- The following pragma eliminates Inst_Pkg.Proc from Q
pragma Eliminate (Q, Proc,
                  Source_Location => "gen_pkg.ads:4[q.adb:3]");
```

## 2.57 Pragma Enable\_Atomic\_Synchronization

Syntax:

```
pragma Enable_Atomic_Synchronization [(Entity)];
```

Reenables atomic synchronization; see `pragma Disable_Atomic_Synchronization` for details.

## 2.58 Pragma Exceptional\_Cases

Syntax:

```
pragma Exceptional_Cases (EXCEPTIONAL_CASE_LIST);
```

```
EXCEPTIONAL_CASE_LIST ::= EXCEPTIONAL_CASE {, EXCEPTIONAL_CASE}
EXCEPTIONAL_CASE      ::= exception_choice {'|' exception_choice} => CONSEQUENCE
CONSEQUENCE           ::= Boolean_expression
```

For the semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.1.9.

## 2.59 Pragma Exit\_Cases

Syntax:

```
pragma Exit_Cases (EXIT_CASE_LIST);
```

```
EXIT_CASE_LIST ::= EXIT_CASE {, EXIT_CASE}
EXIT_CASE      ::= GUARD => EXIT_KIND
EXIT_KIND      ::= Normal_Return
                  | Exception_Raised
                  | (Exception_Raised => exception_name)
                  | Program_Exit
GUARD          ::= Boolean_expression
```

For the semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.1.10.

## 2.60 Pragma Export\_Function

Syntax:

```
pragma Export_Function (
    [Internal          =>] LOCAL_NAME
    [, [External       =>] EXTERNAL_SYMBOL]
    [, [Parameter_Types =>] PARAMETER_TYPES]
    [, [Result_Type    =>] result_SUBTYPE_MARK]
    [, [Mechanism      =>] MECHANISM]
    [, [Result_Mechanism =>] MECHANISM_NAME]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
  | static_string_EXPRESSION
  | ""
```

```

PARAMETER_TYPES ::=
    null
  | TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
    subtype_NAME
  | subtype_Name ' Access

MECHANISM ::=
    MECHANISM_NAME
  | (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
    [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::= Value | Reference

```

Use this pragma to make a function externally callable and optionally provide information on mechanisms to be used for passing parameter and result values. We recommend, for the purposes of improving portability, this pragma always be used in conjunction with a separate pragma `Export`, which must precede the pragma `Export_Function`. GNAT does not require a separate pragma `Export`, but if none is present, `Convention Ada` is assumed, which is usually not what is wanted, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention. Pragma `Export_Function` (and `Export`, if present) must appear in the same declarative region as the function to which they apply.

The `internal_name` must uniquely designate the function to which the pragma applies. If more than one function name exists of this name in the declarative part you must use the `Parameter_Types` and `Result_Type` parameters to achieve the required unique designation. The `subtype_marks` in these parameters must exactly match the subtypes in the corresponding function specification, using positional notation to match parameters with subtype marks. The form with an `'Access` attribute can be used to match an anonymous access parameter.

Special treatment is given if the `EXTERNAL` is an explicit null string or a static string expressions that evaluates to the null string. In this case, no external name is generated. This form still allows the specification of parameter mechanisms.

## 2.61 Pragma `Export_Object`

Syntax:

```

pragma Export_Object (
    [Internal =>] LOCAL_NAME
    [, [External =>] EXTERNAL_SYMBOL]
    [, [Size      =>] EXTERNAL_SYMBOL]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER

```

```
| static_string_EXPRESSION
```

This pragma designates an object as exported, and apart from the extended rules for external symbols, is identical in effect to the use of the normal `Export` pragma applied to an object. You may use a separate `Export` pragma (and you probably should from the point of view of portability), but it is not required. `Size` is syntax checked, but otherwise ignored by GNAT.

## 2.62 Pragma `Export_Procedure`

Syntax:

```
pragma Export_Procedure (
    [Internal      =>] LOCAL_NAME
    [, [External   =>] EXTERNAL_SYMBOL]
    [, [Parameter_Types =>] PARAMETER_TYPES]
    [, [Mechanism   =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
| static_string_EXPRESSION
| ""

PARAMETER_TYPES ::=
    null
| TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
    subtype_NAME
| subtype_Name ' Access

MECHANISM ::=
    MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
    [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::= Value | Reference
```

This pragma is identical to `Export_Function` except that it applies to a procedure rather than a function and the parameters `Result_Type` and `Result_Mechanism` are not permitted. GNAT does not require a separate pragma `Export`, but if none is present, `Convention Ada` is assumed, which is usually not what is wanted, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention.

Special treatment is given if the `EXTERNAL` is an explicit null string or a static string expressions that evaluates to the null string. In this case, no external name is generated. This form still allows the specification of parameter mechanisms.

## 2.63 Pragma Export\_Valued\_Procedure

Syntax:

```
pragma Export_Valued_Procedure (
    [Internal          =>] LOCAL_NAME
    [, [External       =>] EXTERNAL_SYMBOL]
    [, [Parameter_Types =>] PARAMETER_TYPES]
    [, [Mechanism      =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
  | static_string_EXPRESSION
  | ""

PARAMETER_TYPES ::=
    null
  | TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
    subtype_NAME
  | subtype_Name ' Access

MECHANISM ::=
    MECHANISM_NAME
  | (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
    [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::= Value | Reference
```

This pragma is identical to `Export_Procedure` except that the first parameter of `LOCAL_NAME`, which must be present, must be of mode `out`, and externally the subprogram is treated as a function with this parameter as the result of the function. GNAT provides for this capability to allow the use of `out` and `in out` parameters in interfacing to external functions (which are not permitted in Ada functions). GNAT does not require a separate pragma `Export`, but if none is present, `Convention Ada` is assumed, which is almost certainly not what is wanted since the whole point of this pragma is to interface with foreign language functions, so it is usually appropriate to use this pragma in conjunction with a `Export` or `Convention` pragma that specifies the desired foreign convention.

Special treatment is given if the `EXTERNAL` is an explicit null string or a static string expressions that evaluates to the null string. In this case, no external name is generated. This form still allows the specification of parameter mechanisms.

## 2.64 Pragma Extend\_System

Syntax:

```
pragma Extend_System ([Name =>] IDENTIFIER);
```

This pragma is used to provide backwards compatibility with other implementations that extend the facilities of package **System**. In GNAT, **System** contains only the definitions that are present in the Ada RM. However, other implementations, notably the DEC Ada 83 implementation, provide many extensions to package **System**.

For each such implementation accommodated by this pragma, GNAT provides a package **Aux\_xxx**, e.g., **Aux\_DEC** for the DEC Ada 83 implementation, which provides the required additional definitions. You can use this package in two ways. You can **with** it in the normal way and access entities either by selection or using a **use** clause. In this case no special processing is required.

However, if existing code contains references such as **System.xxx** where ‘xxx’ is an entity in the extended definitions provided in package **System**, you may use this pragma to extend visibility in **System** in a non-standard way that provides greater compatibility with the existing code. Pragma **Extend\_System** is a configuration pragma whose single argument is the name of the package containing the extended definition (e.g., **Aux\_DEC** for the DEC Ada case). A unit compiled under control of this pragma will be processed using special visibility processing that looks in package **System.Aux\_xxx** where **Aux\_xxx** is the pragma argument for any entity referenced in package **System**, but not found in package **System**.

You can use this pragma either to access a predefined **System** extension supplied with the compiler, for example **Aux\_DEC** or you can construct your own extension unit following the above definition. Note that such a package is a child of **System** and thus is considered part of the implementation. To compile it you will have to use the ‘-gnatg’ switch for compiling **System** units, as explained in the GNAT User’s Guide.

## 2.65 Pragma Extensions\_Allowed

Syntax:

```
pragma Extensions_Allowed (On | Off | All_Extensions);
```

This configuration pragma enables (via the “On” or “All\_Extensions” argument) or disables (via the “Off” argument) the implementation extension mode; the pragma takes precedence over the **-gnatX** and **-gnatX0** command switches.

If an argument of “On” is specified, the latest version of the Ada language is implemented (currently Ada 2022) and, in addition, a curated set of GNAT specific extensions are recognized. (See the list here [here], page 330)

An argument of “All\_Extensions” has the same effect except that some extra experimental extensions are enabled (See the list here [here], page 341)

## 2.66 Pragma Extensions\_Visible

Syntax:

```
pragma Extensions_Visible [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect **Extensions\_Visible** in the SPARK 2014 Reference Manual, section 6.1.7.

## 2.67 Pragma External

Syntax:

```
pragma External (
  [ Convention      =>] convention_IDENTIFIER,
  [ Entity          =>] LOCAL_NAME
  [, [External_Name =>] static_string_EXPRESSION ]
  [, [Link_Name     =>] static_string_EXPRESSION ] );
```

This pragma is identical in syntax and semantics to pragma **Export** as defined in the Ada Reference Manual. It is provided for compatibility with some Ada 83 compilers that used this pragma for exactly the same purposes as pragma **Export** before the latter was standardized.

## 2.68 Pragma External\_Name\_Casing

Syntax:

```
pragma External_Name_Casing (
  Uppercase | Lowercase
  [, Uppercase | Lowercase | As_Is]);
```

This pragma provides control over the casing of external names associated with Import and Export pragmas. There are two cases to consider:

\* Implicit external names

Implicit external names are derived from identifiers. The most common case arises when a standard Ada Import or Export pragma is used with only two arguments, as in:

```
pragma Import (C, C_Routine);
```

Since Ada is a case-insensitive language, the spelling of the identifier in the Ada source program does not provide any information on the desired casing of the external name, and so a convention is needed. In GNAT the default treatment is that such names are converted to all lower case letters. This corresponds to the normal C style in many environments. The first argument of pragma **External\_Name\_Casing** can be used to control this treatment. If **Uppercase** is specified, then the name will be forced to all uppercase letters. If **Lowercase** is specified, then the normal default of all lower case letters will be used.

This same implicit treatment is also used in the case of extended DEC Ada 83 compatible Import and Export pragmas where an external name is explicitly specified using an identifier rather than a string.

\* Explicit external names

Explicit external names are given as string literals. The most common case arises when a standard Ada Import or Export pragma is used with three arguments, as in:

```
pragma Import (C, C_Routine, "C_routine");
```

In this case, the string literal normally provides the exact casing required for the external name. The second argument of pragma **External\_Name\_Casing** may be used to modify this behavior. If **Uppercase** is specified, then the name will be forced to all uppercase letters. If **Lowercase** is specified, then the name will be forced to all

lowercase letters. A specification of `As_Is` provides the normal default behavior in which the casing is taken from the string provided.

This pragma may appear anywhere that a pragma is valid. In particular, it can be used as a configuration pragma in the `gnat.adc` file, in which case it applies to all subsequent compilations, or it can be used as a program unit pragma, in which case it only applies to the current unit, or it can be used more locally to control individual Import/Export pragmas.

It was primarily intended for use with OpenVMS systems, where many compilers convert all symbols to upper case by default. For interfacing to such compilers (e.g., the DEC C compiler), it may be convenient to use the pragma:

```
pragma External_Name_Casing (Uppercase, Uppercase);
```

to enforce the upper casing of all external symbols.

## 2.69 Pragma Fast\_Math

Syntax:

```
pragma Fast_Math;
```

This is a configuration pragma which activates a mode in which speed is considered more important for floating-point operations than absolutely accurate adherence to the requirements of the standard. Currently the following operations are affected:

‘Complex Multiplication’

The normal simple formula for complex multiplication can result in intermediate overflows for numbers near the end of the range. The Ada standard requires that this situation be detected and corrected by scaling, but in `Fast_Math` mode such cases will simply result in overflow. Note that to take advantage of this you must instantiate your own version of `Ada.Numerics.Generic_Complex_Types` under control of the pragma, rather than use the preinstantiated versions.

## 2.70 Pragma Favor\_Top\_Level

Syntax:

```
pragma Favor_Top_Level (type_LOCAL_NAME);
```

The argument of pragma `Favor_Top_Level` must be a named access-to-subprogram type. This pragma is an efficiency hint to the compiler, regarding the use of `'Access` or `'Unrestricted_Access` on nested (non-library-level) subprograms. The pragma means that nested subprograms are not used with this type, or are rare, so that the generated code should be efficient in the top-level case. When this pragma is used, dynamically generated trampolines may be used on some targets for nested subprograms. See restriction `No_Implicit_Dynamic_Code`.

## 2.71 Pragma Finalize\_Storage\_Only

Syntax:

```
pragma Finalize_Storage_Only (first_subtype_LOCAL_NAME);
```

The argument of pragma `Finalize_Storage_Only` must denote a local type which is derived from `Ada.Finalization.Controlled` or `Limited_Controlled`. The pragma suppresses the call to `Finalize` for declared library-level objects of the argument type. This is mostly useful for types where finalization is only used to deal with storage reclamation since in most environments it is not necessary to reclaim memory just before terminating execution, hence the name. Note that this pragma does not suppress `Finalize` calls for library-level heap-allocated objects (see pragma `No_Heap_Finalization`).

## 2.72 Pragma `Float_Representation`

Syntax:

```
pragma Float_Representation (FLOAT_REP[, float_type_LOCAL_NAME]);

FLOAT_REP ::= VAX_Float | IEEE_Float
```

In the one argument form, this pragma is a configuration pragma which allows control over the internal representation chosen for the predefined floating point types declared in the packages `Standard` and `System`. This pragma is only provided for compatibility and has no effect.

The two argument form specifies the representation to be used for the specified floating-point type. The argument must be `IEEE_Float` to specify the use of IEEE format, as follows:

- \* For a digits value of 6, 32-bit IEEE short format will be used.
- \* For a digits value of 15, 64-bit IEEE long format will be used.
- \* No other value of digits is permitted.

## 2.73 Pragma `Ghost`

Syntax:

```
pragma Ghost [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Ghost` in the SPARK 2014 Reference Manual, section 6.9.

## 2.74 Pragma `Global`

Syntax:

```
pragma Global (GLOBAL_SPECIFICATION);

GLOBAL_SPECIFICATION ::=
    null
  | (GLOBAL_LIST)
  | (MODED_GLOBAL_LIST {, MODED_GLOBAL_LIST})

MODED_GLOBAL_LIST ::= MODE_SELECTOR => GLOBAL_LIST

MODE_SELECTOR ::= In_Out | Input | Output | Proof_In
GLOBAL_LIST   ::= GLOBAL_ITEM | (GLOBAL_ITEM {, GLOBAL_ITEM})
```

`GLOBAL_ITEM ::= NAME`

For the semantics of this pragma, see the entry for aspect `Global` in the SPARK 2014 Reference Manual, section 6.1.4.

## 2.75 Pragma Ident

Syntax:

`pragma Ident (static_string_EXPRESSION);`

This pragma is identical in effect to pragma `Comment`. It is provided for compatibility with other Ada compilers providing this pragma.

## 2.76 Pragma Ignore\_Pragma

Syntax:

`pragma Ignore_Pragma (pragma_IDENTIFIER);`

This is a configuration pragma that takes a single argument that is a simple identifier. Any subsequent use of a pragma whose pragma identifier matches this argument will be silently ignored. Any preceding use of a pragma whose pragma identifier matches this argument will be parsed and then ignored. This may be useful when legacy code or code intended for compilation with some other compiler contains pragmas that match the name, but not the exact implementation, of a GNAT pragma. The use of this pragma allows such pragmas to be ignored, which may be useful in CodePeer mode, or during porting of legacy code.

## 2.77 Pragma Implementation\_Defined

Syntax:

`pragma Implementation_Defined (local_NAME);`

This pragma marks a previously declared entity as implementation-defined. For an overloaded entity, applies to the most recent homonym.

`pragma Implementation_Defined;`

The form with no arguments appears anywhere within a scope, most typically a package spec, and indicates that all entities that are defined within the package spec are `Implementation_Defined`.

This pragma is used within the GNAT runtime library to identify implementation-defined entities introduced in language-defined units, for the purpose of implementing the `No_Implementation_Identifiers` restriction.

## 2.78 Pragma Implemented

Syntax:

`pragma Implemented (procedure_LOCAL_NAME, implementation_kind);`

`implementation_kind ::= By_Entry | By_Protected_Procedure | By_Any`

This is an Ada 2012 representation pragma which applies to protected, task and synchronized interface primitives. The use of pragma `Implemented` provides a way to impose a

static requirement on the overriding operation by adhering to one of the three implementation kinds: entry, protected procedure or any of the above. This pragma is available in all earlier versions of Ada as an implementation-defined pragma.

```

type Synch_Iface is synchronized interface;
procedure Prim_Op (Obj : in out Iface) is abstract;
pragma Implemented (Prim_Op, By_Protected_Procedure);

protected type Prot_1 is new Synch_Iface with
  procedure Prim_Op; -- Legal
end Prot_1;

protected type Prot_2 is new Synch_Iface with
  entry Prim_Op;      -- Illegal
end Prot_2;

task type Task_Typ is new Synch_Iface with
  entry Prim_Op;      -- Illegal
end Task_Typ;

```

When applied to the `procedure_or_entry_NAME` of a `requeue` statement, pragma `Implemented` determines the runtime behavior of the `requeue`. Implementation kind `By_Entry` guarantees that the action of `requeueing` will proceed from an entry to another entry. Implementation kind `By_Protected_Procedure` transforms the `requeue` into a dispatching call, thus eliminating the chance of blocking. Kind `By_Any` shares the behavior of `By_Entry` and `By_Protected_Procedure` depending on the target's overriding subprogram kind.

## 2.79 Pragma `Implicit_Packing`

Syntax:

```
pragma Implicit_Packing;
```

This is a configuration pragma that requests implicit packing for packed arrays for which a size clause is given but no explicit pragma `Pack` or specification of `Component_Size` is present. It also applies to records where no record representation clause is present. Consider this example:

```

type R is array (0 .. 7) of Boolean;
for R'Size use 8;

```

In accordance with the recommendation in the RM (RM 13.3(53)), a `Size` clause does not change the layout of a composite object. So the `Size` clause in the above example is normally rejected, since the default layout of the array uses 8-bit components, and thus the array requires a minimum of 64 bits.

If this declaration is compiled in a region of code covered by an occurrence of the configuration pragma `Implicit_Packing`, then the `Size` clause in this and similar examples will cause implicit packing and thus be accepted. For this implicit packing to occur, the type in question must be an array of small components whose size is known at compile time, and the `Size` clause must specify the exact size that corresponds to the number of elements in the array multiplied by the size in bits of the component type (both single and multi-dimensioned arrays can be controlled with this pragma).

Similarly, the following example shows the use in the record case

```
type r is record
  a, b, c, d, e, f, g, h : boolean;
  chr                    : character;
end record;
for r'size use 16;
```

Without a pragma Pack, each Boolean field requires 8 bits, so the minimum size is 72 bits, but with a pragma Pack, 16 bits would be sufficient. The use of pragma Implicit\_Packing allows this record declaration to compile without an explicit pragma Pack.

## 2.80 Pragma Import\_Function

Syntax:

```
pragma Import_Function (
  [Internal          =>] LOCAL_NAME,
  [, [External       =>] EXTERNAL_SYMBOL]
  [, [Parameter_Types =>] PARAMETER_TYPES]
  [, [Result_Type    =>] SUBTYPE_MARK]
  [, [Mechanism      =>] MECHANISM]
  [, [Result_Mechanism =>] MECHANISM_NAME]);

EXTERNAL_SYMBOL ::=
  IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
  null
| TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
  subtype_NAME
| subtype_Name ' Access

MECHANISM ::=
  MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
  [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::=
  Value
| Reference
```

This pragma is used in conjunction with a pragma Import to specify additional information for an imported function. The pragma Import (or equivalent pragma Interface) must

precede the `Import_Function` pragma and both must appear in the same declarative part as the function specification.

The `Internal` argument must uniquely designate the function to which the pragma applies. If more than one function name exists of this name in the declarative part you must use the `Parameter_Types` and `Result_Type` parameters to achieve the required unique designation. Subtype marks in these parameters must exactly match the subtypes in the corresponding function specification, using positional notation to match parameters with subtype marks. The form with an `'Access` attribute can be used to match an anonymous access parameter.

You may optionally use the `Mechanism` and `Result_Mechanism` parameters to specify passing mechanisms for the parameters and result. If you specify a single mechanism name, it applies to all parameters. Otherwise you may specify a mechanism on a parameter by parameter basis using either positional or named notation. If the mechanism is not specified, the default mechanism is used.

## 2.81 Pragma Import\_Object

Syntax:

```
pragma Import_Object (
    [Internal =>] LOCAL_NAME
    [, [External =>] EXTERNAL_SYMBOL]
    [, [Size      =>] EXTERNAL_SYMBOL]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
    | static_string_EXPRESSION
```

This pragma designates an object as imported, and apart from the extended rules for external symbols, is identical in effect to the use of the normal `Import` pragma applied to an object. Unlike the subprogram case, you need not use a separate `Import` pragma, although you may do so (and probably should do so from a portability point of view). `size` is syntax checked, but otherwise ignored by GNAT.

## 2.82 Pragma Import\_Procedure

Syntax:

```
pragma Import_Procedure (
    [Internal      =>] LOCAL_NAME
    [, [External    =>] EXTERNAL_SYMBOL]
    [, [Parameter_Types =>] PARAMETER_TYPES]
    [, [Mechanism    =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
    | static_string_EXPRESSION

PARAMETER_TYPES ::=
    null
```

```

| TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
    subtype_NAME
| subtype_Name ' Access

MECHANISM ::=
    MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
    [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::= Value | Reference

```

This pragma is identical to `Import_Function` except that it applies to a procedure rather than a function and the parameters `Result_Type` and `Result_Mechanism` are not permitted.

## 2.83 Pragma `Import_Valued_Procedure`

Syntax:

```

pragma Import_Valued_Procedure (
    [Internal          =>] LOCAL_NAME
    [, [External       =>] EXTERNAL_SYMBOL]
    [, [Parameter_Types =>] PARAMETER_TYPES]
    [, [Mechanism      =>] MECHANISM]);

EXTERNAL_SYMBOL ::=
    IDENTIFIER
| static_string_EXPRESSION

PARAMETER_TYPES ::=
    null
| TYPE_DESIGNATOR {, TYPE_DESIGNATOR}

TYPE_DESIGNATOR ::=
    subtype_NAME
| subtype_Name ' Access

MECHANISM ::=
    MECHANISM_NAME
| (MECHANISM_ASSOCIATION {, MECHANISM_ASSOCIATION})

MECHANISM_ASSOCIATION ::=
    [formal_parameter_NAME =>] MECHANISM_NAME

MECHANISM_NAME ::= Value | Reference

```

This pragma is identical to `Import_Procedure` except that the first parameter of `LOCAL_NAME`, which must be present, must be of mode `out`, and externally the subprogram is treated as a function with this parameter as the result of the function. The purpose of this capability is to allow the use of `out` and `in out` parameters in interfacing to external functions (which are not permitted in Ada functions). You may optionally use the `Mechanism` parameters to specify passing mechanisms for the parameters. If you specify a single mechanism name, it applies to all parameters. Otherwise you may specify a mechanism on a parameter by parameter basis using either positional or named notation. If the mechanism is not specified, the default mechanism is used.

Note that it is important to use this pragma in conjunction with a separate pragma `Import` that specifies the desired convention, since otherwise the default convention is `Ada`, which is almost certainly not what is required.

## 2.84 Pragma Independent

Syntax:

```
pragma Independent (component_LOCAL_NAME);
```

This pragma is standard in Ada 2012 mode (which also provides an aspect of the same name). It is also available as an implementation-defined pragma in all earlier versions. It specifies that the designated object or all objects of the designated type must be independently addressable. This means that separate tasks can safely manipulate such objects. For example, if two components of a record are independent, then two separate tasks may access these two components. This may place constraints on the representation of the object (for instance prohibiting tight packing).

## 2.85 Pragma Independent\_Components

Syntax:

```
pragma Independent_Components (Local_NAME);
```

This pragma is standard in Ada 2012 mode (which also provides an aspect of the same name). It is also available as an implementation-defined pragma in all earlier versions. It specifies that the components of the designated object, or the components of each object of the designated type, must be independently addressable. This means that separate tasks can safely manipulate separate components in the composite object. This may place constraints on the representation of the object (for instance prohibiting tight packing).

## 2.86 Pragma Initial\_Condition

Syntax:

```
pragma Initial_Condition (boolean_EXPRESSION);
```

For the semantics of this pragma, see the entry for aspect `Initial_Condition` in the SPARK 2014 Reference Manual, section 7.1.6.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.87 Pragma InitializeScalars

Syntax:

```
pragma InitializeScalars
  [ ( TYPE_VALUE_PAIR {, TYPE_VALUE_PAIR} ) ];

TYPE_VALUE_PAIR ::=
  SCALAR_TYPE => static_EXPRESSION

SCALAR_TYPE :=
  Short_Float
| Float
| Long_Float
| Long_Long_Float
| Signed_8
| Signed_16
| Signed_32
| Signed_64
| Unsigned_8
| Unsigned_16
| Unsigned_32
| Unsigned_64
```

This pragma is similar to `NormalizeScalars` conceptually but has two important differences.

First, there is no requirement for the pragma to be used uniformly in all units of a partition. In particular, it is fine to use this just for some or all of the application units of a partition, without needing to recompile the run-time library. In the case where some units are compiled with the pragma, and some without, then a declaration of a variable where the type is defined in package `Standard` or is locally declared will always be subject to initialization, as will any declaration of a scalar variable. For composite variables, whether the variable is initialized may also depend on whether the package in which the type of the variable is declared is compiled with the pragma.

The other important difference is that the programmer can control the value used for initializing scalar objects. This effect can be achieved in several different ways:

- \* At compile time, the programmer can specify the invalid value for a particular family of scalar types using the optional arguments of the pragma.

The compile-time approach is intended to optimize the generated code for the pragma, by possibly using fast operations such as `memset`. Note that such optimizations require using values where the bytes all have the same binary representation.

- \* At bind time, the programmer has several options:
  - \* Initialization with invalid values (similar to `NormalizeScalars`, though for `InitializeScalars` it is not always possible to determine the invalid values in complex cases like signed component fields with nonstandard sizes).
  - \* Initialization with high values.
  - \* Initialization with low values.

- \* Initialization with a specific bit pattern.

See the GNAT User's Guide for binder options for specifying these cases.

The bind-time approach is intended to provide fast turnaround for testing with different values, without having to recompile the program.

- \* At execution time, the programmer can specify the invalid values using an environment variable. See the GNAT User's Guide for details.

The execution-time approach is intended to provide fast turnaround for testing with different values, without having to recompile and rebind the program.

Note that pragma `Initialize_Scalars` is particularly useful in conjunction with the enhanced validity checking that is now provided in GNAT, which checks for invalid values under more conditions. Using this feature (see description of the '-gnatV' flag in the GNAT User's Guide) in conjunction with pragma `Initialize_Scalars` provides a powerful new tool to assist in the detection of problems caused by uninitialized variables.

Note: the use of `Initialize_Scalars` has a fairly extensive effect on the generated code. This may cause your code to be substantially larger. It may also cause an increase in the amount of stack required, so it is probably a good idea to turn on stack checking (see description of stack checking in the GNAT User's Guide) when using this pragma.

## 2.88 Pragma `Initializes`

Syntax:

```
pragma Initializes (INITIALIZATION_LIST);

INITIALIZATION_LIST ::=
    null
  | (INITIALIZATION_ITEM {, INITIALIZATION_ITEM})

INITIALIZATION_ITEM ::= name [=> INPUT_LIST]

INPUT_LIST ::=
    null
  | INPUT
  | (INPUT {, INPUT})

INPUT ::= name
```

For the semantics of this pragma, see the entry for aspect `Initializes` in the SPARK 2014 Reference Manual, section 7.1.5.

## 2.89 Pragma `Inline_Always`

Syntax:

```
pragma Inline_Always (NAME [, NAME]);
```

Similar to pragma `Inline` except that inlining is unconditional. `Inline_Always` instructs the compiler to inline every direct call to the subprogram or else to emit a compilation error, independently of any option, in particular '-gnatn' or '-gnatN' or the optimization level. It

is an error to take the address or access of `NAME`. It is also an error to apply this pragma to a primitive operation of a tagged type. Thanks to such restrictions, the compiler is allowed to remove the out-of-line body of `NAME`.

## 2.90 Pragma Inline\_Generic

Syntax:

```
pragma Inline_Generic (GNAME {, GNAME});

GNAME ::= generic_unit_NAME | generic_instance_NAME
```

This pragma is provided for compatibility with Dec Ada 83. It has no effect in GNAT (which always inlines generics), other than to check that the given names are all names of generic units or generic instances.

## 2.91 Pragma Interface

Syntax:

```
pragma Interface (
    [Convention      =>] convention_identifier,
    [Entity          =>] local_NAME
    [, [External_Name =>] static_string_expression]
    [, [Link_Name     =>] static_string_expression]);
```

This pragma is identical in syntax and semantics to the standard Ada pragma `Import`. It is provided for compatibility with Ada 83. The definition is upwards compatible both with pragma `Interface` as defined in the Ada 83 Reference Manual, and also with some extended implementations of this pragma in certain Ada 83 implementations. The only difference between pragma `Interface` and pragma `Import` is that there is special circuitry to allow both pragmas to appear for the same subprogram entity (normally it is illegal to have multiple `Import` pragmas). This is useful in maintaining Ada 83/Ada 95 compatibility and is compatible with other Ada 83 compilers.

## 2.92 Pragma Interface\_Name

Syntax:

```
pragma Interface_Name (
    [Entity          =>] LOCAL_NAME
    [, [External_Name =>] static_string_EXPRESSION]
    [, [Link_Name     =>] static_string_EXPRESSION]);
```

This pragma provides an alternative way of specifying the interface name for an interfaced subprogram, and is provided for compatibility with Ada 83 compilers that use the pragma for this purpose. You must provide at least one of `External_Name` or `Link_Name`.

## 2.93 Pragma Interrupt\_Handler

Syntax:

```
pragma Interrupt_Handler (procedure_LOCAL_NAME);
```

This program unit pragma is supported for parameterless protected procedures as described in Annex C of the Ada Reference Manual.

## 2.94 Pragma `Interrupt_State`

Syntax:

```
pragma Interrupt_State
  ([Name =>] value,
   [State =>] SYSTEM | RUNTIME | USER);
```

Normally certain interrupts are reserved to the implementation. Any attempt to attach an interrupt causes `Program_Error` to be raised, as described in RM C.3.2(22). A typical example is the `SIGINT` interrupt used in many systems for an `Ctrl-C` interrupt. Normally this interrupt is reserved to the implementation, so that `Ctrl-C` can be used to interrupt execution. Additionally, signals such as `SIGSEGV`, `SIGABRT`, `SIGFPE` and `SIGILL` are often mapped to specific Ada exceptions, or used to implement run-time functions such as the `abort` statement and stack overflow checking.

Pragma `Interrupt_State` provides a general mechanism for overriding such uses of interrupts. It subsumes the functionality of pragma `Unreserve_All_Interrupts`. Pragma `Interrupt_State` is not available on Windows. On all other platforms than VxWorks, it applies to signals; on VxWorks, it applies to vectored hardware interrupts and may be used to mark interrupts required by the board support package as reserved.

Interrupts can be in one of three states:

- \* System

The interrupt is reserved (no Ada handler can be installed), and the Ada run-time may not install a handler. As a result you are guaranteed standard system default action if this interrupt is raised. This also allows installing a low level handler via C APIs such as `sigaction()`, outside of Ada control.

- \* Runtime

The interrupt is reserved (no Ada handler can be installed). The run time is allowed to install a handler for internal control purposes, but is not required to do so.

- \* User

The interrupt is unreserved. The user may install an Ada handler via `Ada.Interrupts` and pragma `Interrupt_Handler` or `Attach_Handler` to provide some other action.

These states are the allowed values of the `State` parameter of the pragma. The `Name` parameter is a value of the type `Ada.Interrupts.Interrupt_ID`. Typically, it is a name declared in `Ada.Interrupts.Names`.

This is a configuration pragma, and the binder will check that there are no inconsistencies between different units in a partition in how a given interrupt is specified. It may appear anywhere a pragma is legal.

The effect is to move the interrupt to the specified state.

By declaring interrupts to be `SYSTEM`, you guarantee the standard system action, such as a core dump.

By declaring interrupts to be `USER`, you guarantee that you can install a handler.

Note that certain signals on many operating systems cannot be caught and handled by applications. In such cases, the pragma is ignored. See the operating system documentation, or the value of the array `Reserved` declared in the spec of package `System.OS_Interface`. Overriding the default state of signals used by the Ada runtime may interfere with an application's runtime behavior in the cases of the synchronous signals, and in the case of the signal used to implement the `abort` statement.

## 2.95 Pragma `Interrupts_System_By_Default`

Syntax:

```
pragma Interrupts_System_By_Default;
```

Default all interrupts to the System state as defined above in pragma `Interrupt_State`. This is a configuration pragma.

## 2.96 Pragma `Invariant`

Syntax:

```
pragma Invariant
  ([Entity =>]      private_type_LOCAL_NAME,
   [Check  =>]     EXPRESSION
   [, [Message =>] String_Expression]);
```

This pragma provides exactly the same capabilities as the `Type_Invariant` aspect defined in AI05-0146-1, and in the Ada 2012 Reference Manual. The `Type_Invariant` aspect is fully implemented in Ada 2012 mode, but since it requires the use of the aspect syntax, which is not available except in 2012 mode, it is not possible to use the `Type_Invariant` aspect in earlier versions of Ada. However the `Invariant` pragma may be used in any version of Ada. Also note that the aspect `Invariant` is a synonym in GNAT for the aspect `Type_Invariant`, but there is no pragma `Type_Invariant`.

The pragma must appear within the visible part of the package specification, after the type to which its `Entity` argument appears. As with the `Invariant` aspect, the `Check` expression is not analyzed until the end of the visible part of the package, so it may contain forward references. The `Message` argument, if present, provides the exception message used if the invariant is violated. If no `Message` parameter is provided, a default message that identifies the line on which the pragma appears is used.

It is permissible to have multiple `Invariants` for the same type entity, in which case they are and'ed together. It is permissible to use this pragma in Ada 2012 mode, but you cannot have both an invariant aspect and an invariant pragma for the same entity.

For further details on the use of this pragma, see the Ada 2012 documentation of the `Type_Invariant` aspect.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.97 Pragma `Keep_Names`

Syntax:

```
pragma Keep_Names ([On =>] enumeration_first_subtype_LOCAL_NAME);
```

The `LOCAL_NAME` argument must refer to an enumeration first subtype in the current declarative part. The effect is to retain the enumeration literal names for use by `Image` and `Value` even if a global `Discard_Names` pragma applies. This is useful when you want to generally suppress enumeration literal names and for example you therefore use a `Discard_Names` pragma in the `gnat.adc` file, but you want to retain the names for specific enumeration types.

## 2.98 Pragma License

Syntax:

```
pragma License (Unrestricted | GPL | Modified_GPL | Restricted);
```

This pragma is provided to allow automated checking for appropriate license conditions with respect to the standard and modified GPL. A pragma `License`, which is a configuration pragma that typically appears at the start of a source file or in a separate `gnat.adc` file, specifies the licensing conditions of a unit as follows:

- \* `Unrestricted` This is used for a unit that can be freely used with no license restrictions. Examples of such units are public domain units, and units from the Ada Reference Manual.
- \* `GPL` This is used for a unit that is licensed under the unmodified GPL, and which therefore cannot be `with`d by a restricted unit.
- \* `Modified_GPL` This is used for a unit licensed under the GNAT modified GPL that includes a special exception paragraph that specifically permits the inclusion of the unit in programs without requiring the entire program to be released under the GPL.
- \* `Restricted` This is used for a unit that is restricted in that it is not permitted to depend on units that are licensed under the GPL. Typical examples are proprietary code that is to be released under more restrictive license conditions. Note that restricted units are permitted to `with` units which are licensed under the modified GPL (this is the whole point of the modified GPL).

Normally a unit with no `License` pragma is considered to have an unknown license, and no checking is done. However, standard GNAT headers are recognized, and license information is derived from them as follows.

A GNAT license header starts with a line containing 78 hyphens. The following comment text is searched for the appearance of any of the following strings.

If the string ‘GNU General Public License’ is found, then the unit is assumed to have GPL license, unless the string ‘As a special exception’ follows, in which case the license is assumed to be modified GPL.

If one of the strings ‘This specification is adapted from the Ada Semantic Interface’ or ‘This specification is derived from the Ada Reference Manual’ is found then the unit is assumed to be unrestricted.

These default actions means that a program with a restricted license pragma will automatically get warnings if a GPL unit is inappropriately `with`d. For example, the program:

```
with Sem_Ch3;
with GNAT.Sockets;
procedure Secret_Stuff is
```

```
...
end Secret_Stuff
```

if compiled with pragma License (Restricted) in a `gnat.adc` file will generate the warning:

```
1.  with Sem_Ch3;
    |
    >>> license of withed unit "Sem_Ch3" is incompatible

2.  with GNAT.Sockets;
3.  procedure Secret_Stuff is
```

Here we get a warning on `Sem_Ch3` since it is part of the GNAT compiler and is licensed under the GPL, but no warning for `GNAT.Sockets` which is part of the GNAT run time, and is therefore licensed under the modified GPL.

## 2.99 Pragma Link\_With

Syntax:

```
pragma Link_With (static_string_EXPRESSION {,static_string_EXPRESSION});
```

This pragma is provided for compatibility with certain Ada 83 compilers. It has exactly the same effect as pragma `Linker_Options` except that spaces occurring within one of the string expressions are treated as separators. For example, in the following case:

```
pragma Link_With ("-labc -ldef");
```

results in passing the strings `-labc` and `-ldef` as two separate arguments to the linker. In addition pragma `Link_With` allows multiple arguments, with the same effect as successive pragmas.

## 2.100 Pragma Linker\_Alias

Syntax:

```
pragma Linker_Alias (
  [Entity =>] LOCAL_NAME,
  [Target =>] static_string_EXPRESSION);
```

`LOCAL_NAME` must refer to an object that is declared at the library level. This pragma establishes the given entity as a linker alias for the given target. It is equivalent to `--attribute__((alias))` in GNU C and causes `LOCAL_NAME` to be emitted as an alias for the symbol `static_string_EXPRESSION` in the object file, that is to say no space is reserved for `LOCAL_NAME` by the assembler and it will be resolved to the same address as `static_string_EXPRESSION` by the linker.

The actual linker name for the target must be used (e.g., the fully encoded name with qualification in Ada, or the mangled name in C++), or it must be declared using the C convention with pragma `Import` or pragma `Export`.

Not all target machines support this pragma. On some of them it is accepted only if pragma `Weak_External` has been applied to `LOCAL_NAME`.

```
-- Example of the use of pragma Linker_Alias
```

```
package p is
```

```

    i : Integer := 1;
    pragma Export (C, i);

    new_name_for_i : Integer;
    pragma Linker_Alias (new_name_for_i, "i");
end p;

```

## 2.101 Pragma Linker\_Constructor

Syntax:

```
pragma Linker_Constructor (procedure_LOCAL_NAME);
```

`procedure_LOCAL_NAME` must refer to a parameterless procedure that is declared at the library level. A procedure to which this pragma is applied will be treated as an initialization routine by the linker. It is equivalent to `__attribute__((constructor))` in GNU C and causes `procedure_LOCAL_NAME` to be invoked before the entry point of the executable is called (or immediately after the shared library is loaded if the procedure is linked in a shared library), in particular before the Ada run-time environment is set up.

Because of these specific contexts, the set of operations such a procedure can perform is very limited and the type of objects it can manipulate is essentially restricted to the elementary types. In particular, it must only contain code to which pragma Restrictions (No\_Elaboration\_Code) applies.

This pragma is used by GNAT to implement auto-initialization of shared Stand Alone Libraries, which provides a related capability without the restrictions listed above. Where possible, the use of Stand Alone Libraries is preferable to the use of this pragma.

## 2.102 Pragma Linker\_Destructor

Syntax:

```
pragma Linker_Destructor (procedure_LOCAL_NAME);
```

`procedure_LOCAL_NAME` must refer to a parameterless procedure that is declared at the library level. A procedure to which this pragma is applied will be treated as a finalization routine by the linker. It is equivalent to `__attribute__((destructor))` in GNU C and causes `procedure_LOCAL_NAME` to be invoked after the entry point of the executable has exited (or immediately before the shared library is unloaded if the procedure is linked in a shared library), in particular after the Ada run-time environment is shut down.

See `pragma Linker_Constructor` for the set of restrictions that apply because of these specific contexts.

## 2.103 Pragma Linker\_Section

Syntax:

```

pragma Linker_Section (
  [Entity =>] LOCAL_NAME,
  [Section =>] static_string_EXPRESSION);

```

`LOCAL_NAME` must refer to an object, type, or subprogram that is declared at the library level. This pragma specifies the name of the linker section for the given entity. It is

equivalent to `__attribute__((section))` in GNU C and causes `LOCAL_NAME` to be placed in the `static_string_EXPRESSION` section of the executable (assuming the linker doesn't rename the section). GNAT also provides an implementation defined aspect of the same name.

In the case of specifying this aspect for a type, the effect is to specify the corresponding section for all library-level objects of the type that do not have an explicit linker section set. Note that this only applies to whole objects, not to components of composite objects.

In the case of a subprogram, the linker section applies to all previously declared matching overloaded subprograms in the current declarative part which do not already have a linker section assigned. The linker section aspect is useful in this case for specifying different linker sections for different elements of such an overloaded set.

Note that an empty string specifies that no linker section is specified. This is not quite the same as omitting the pragma or aspect, since it can be used to specify that one element of an overloaded set of subprograms has the default linker section, or that one object of a type for which a linker section is specified should have the default linker section.

The compiler normally places library-level entities in standard sections depending on the class: procedures and functions generally go in the `.text` section, initialized variables in the `.data` section and uninitialized variables in the `.bss` section.

Other, special sections may exist on given target machines to map special hardware, for example I/O ports or flash memory. This pragma is a means to defer the final layout of the executable to the linker, thus fully working at the symbolic level with the compiler.

Some file formats do not support arbitrary sections so not all target machines support this pragma. The use of this pragma may cause a program execution to be erroneous if it is used to place an entity into an inappropriate section (e.g., a modified variable into the `.text` section). See also `pragma Persistent_BSS`.

-- Example of the use of `pragma Linker_Section`

```
package IO_Card is
  Port_A : Integer;
  pragma Volatile (Port_A);
  pragma Linker_Section (Port_A, ".bss.port_a");

  Port_B : Integer;
  pragma Volatile (Port_B);
  pragma Linker_Section (Port_B, ".bss.port_b");

  type Port_Type is new Integer with Linker_Section => ".bss";
  PA : Port_Type with Linker_Section => ".bss.PA";
  PB : Port_Type; -- ends up in linker section ".bss"

  procedure Q with Linker_Section => "Qsection";
end IO_Card;
```

## 2.104 Pragma Lock\_Free

Syntax:

```
pragma Lock_Free [ (static_boolean_EXPRESSION) ];
```

This pragma may be specified for protected types or objects. It specifies that the implementation of protected operations must be implemented without locks. Compilation fails if the compiler cannot generate lock-free code for the operations.

The current conditions required to support this pragma are:

- \* Protected type declarations may not contain entries
- \* Protected subprogram declarations may not have nonelementary parameters

In addition, each protected subprogram body must satisfy:

- \* May reference only one protected component
- \* May not reference nonconstant entities outside the protected subprogram scope
- \* May not contain address representation items, allocators, or quantified expressions
- \* May not contain delay, goto, loop, or procedure-call statements
- \* May not contain exported and imported entities
- \* May not dereferenced access values
- \* Function calls and attribute references must be static

If the `Lock_Free` aspect is specified to be `True` for a protected unit and the `Ceiling_Locking` locking policy is in effect, then the run-time actions associated with the `Ceiling_Locking` locking policy (described in Ada RM D.3) are not performed when a protected operation of the protected unit is executed.

## 2.105 Pragma `Loop_Invariant`

Syntax:

```
pragma Loop_Invariant ( boolean_EXPRESSION );
```

The effect of this pragma is similar to that of `pragma Assert`, except that in an `Assertion_Policy` pragma, the identifier `Loop_Invariant` is used to control whether it is ignored or checked (or disabled).

`Loop_Invariant` can only appear as one of the items in the sequence of statements of a loop body, or nested inside block statements that appear in the sequence of statements of a loop body. The intention is that it be used to represent a “loop invariant” assertion, i.e. something that is true each time through the loop, and which can be used to show that the loop is achieving its purpose.

Multiple `Loop_Invariant` and `Loop_Variant` pragmas that apply to the same loop should be grouped in the same sequence of statements.

To aid in writing such invariants, the special attribute `Loop_Entry` may be used to refer to the value of an expression on entry to the loop. This attribute can only be used within the expression of a `Loop_Invariant` pragma. For full details, see documentation of attribute `Loop_Entry`.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.106 Pragma Loop\_Optimize

Syntax:

```
pragma Loop_Optimize (OPTIMIZATION_HINT {, OPTIMIZATION_HINT});
```

```
OPTIMIZATION_HINT ::= Ivdep | No_Unroll | Unroll | No_Vector | Vector
```

This pragma must appear immediately within a loop statement. It allows the programmer to specify optimization hints for the enclosing loop. The hints are not mutually exclusive and can be freely mixed, but not all combinations will yield a sensible outcome.

There are five supported optimization hints for a loop:

- \* Ivdep

The programmer asserts that there are no loop-carried dependencies which would prevent consecutive iterations of the loop from being executed simultaneously.

- \* No\_Unroll

The loop must not be unrolled. This is a strong hint: the compiler will not unroll a loop marked with this hint.

- \* Unroll

The loop should be unrolled. This is a weak hint: the compiler will try to apply unrolling to this loop preferably to other optimizations, notably vectorization, but there is no guarantee that the loop will be unrolled.

- \* No\_Vector

The loop must not be vectorized. This is a strong hint: the compiler will not vectorize a loop marked with this hint.

- \* Vector

The loop should be vectorized. This is a weak hint: the compiler will try to apply vectorization to this loop preferably to other optimizations, notably unrolling, but there is no guarantee that the loop will be vectorized.

These hints do not remove the need to pass the appropriate switches to the compiler in order to enable the relevant optimizations, that is to say ‘-funroll-loops’ for unrolling and ‘-ftree-vectorize’ for vectorization.

## 2.107 Pragma Loop\_Variant

Syntax:

```
pragma Loop_Variant ( LOOP_VARIANT_ITEM {, LOOP_VARIANT_ITEM } );
LOOP_VARIANT_ITEM ::= CHANGE_DIRECTION => discrete_EXPRESSION
CHANGE_DIRECTION ::= Increases | Decreases
```

**Loop\_Variant** can only appear as one of the items in the sequence of statements of a loop body, or nested inside block statements that appear in the sequence of statements of a loop body. It allows the specification of quantities which must always decrease or increase in successive iterations of the loop. In its simplest form, just one expression is specified, whose value must increase or decrease on each iteration of the loop.

In a more complex form, multiple arguments can be given which are interpreted in a nesting lexicographic manner. For example:

```
pragma Loop_Variant (Increases => X, Decreases => Y);
```

specifies that each time through the loop either X increases, or X stays the same and Y decreases. A `Loop_Variant` pragma ensures that the loop is making progress. It can be useful in helping to show informally or prove formally that the loop always terminates.

`Loop_Variant` is an assertion whose effect can be controlled using an `Assertion_Policy` with a check name of `Loop_Variant`. The policy can be `Check` to enable the loop variant check, `Ignore` to ignore the check (in which case the pragma has no effect on the program), or `Disable` in which case the pragma is not even checked for correct syntax.

Multiple `Loop_Invariant` and `Loop_Variant` pragmas that apply to the same loop should be grouped in the same sequence of statements.

The `Loop_Entry` attribute may be used within the expressions of the `Loop_Variant` pragma to refer to values on entry to the loop.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.108 Pragma Machine\_Attribute

Syntax:

```
pragma Machine_Attribute (
    [Entity          =>] LOCAL_NAME,
    [Attribute_Name =>] static_string_EXPRESSION
    [, [Info         =>] static_EXPRESSION {, static_EXPRESSION}] );
```

Machine-dependent attributes can be specified for types and/or declarations. This pragma is semantically equivalent to `__attribute__((attribute_name))` (if `info` is not specified) or `__attribute__((attribute_name(info)))` or `__attribute__((attribute_name(info,...)))` in GNU C, where ‘attribute\_name’ is recognized by the compiler middle-end or the `TARGET_ATTRIBUTE_TABLE` machine specific macro. Note that a string literal for the optional parameter `info` or the following ones is transformed by default into an identifier, which may make this pragma unusable for some attributes. For further information see *GNU Compiler Collection (GCC) Internals*.

## 2.109 Pragma Main

Syntax:

```
pragma Main
    (MAIN_OPTION [, MAIN_OPTION]);

MAIN_OPTION ::=
    [Stack_Size          =>] static_integer_EXPRESSION
    | [Task_Stack_Size_Default =>] static_integer_EXPRESSION
    | [Time_Slicing_Enabled  =>] static_boolean_EXPRESSION
```

This pragma is provided for compatibility with OpenVMS VAX Systems. It has no effect in GNAT, other than being syntax checked.

## 2.110 Pragma Main\_Storage

Syntax:

```
pragma Main_Storage
```

```
(MAIN_STORAGE_OPTION [, MAIN_STORAGE_OPTION]);

MAIN_STORAGE_OPTION ::=
    [WORKING_STORAGE =>] static_SIMPLE_EXPRESSION
  | [TOP_GUARD      =>] static_SIMPLE_EXPRESSION
```

This pragma is provided for compatibility with OpenVMS VAX Systems. It has no effect in GNAT, other than being syntax checked.

## 2.111 Pragma Max\_Queue\_Length

Syntax:

```
pragma Max_Queue_Length (static_integer_EXPRESSION);
```

This pragma is used to specify the maximum callers per entry queue for individual protected entries and entry families. It accepts a single integer (-1 or more) as a parameter and must appear after the declaration of an entry.

A value of -1 represents no additional restriction on queue length.

## 2.112 Pragma No\_Body

Syntax:

```
pragma No_Body;
```

There are a number of cases in which a package spec does not require a body, and in fact a body is not permitted. GNAT will not permit the spec to be compiled if there is a body around. The pragma No\_Body allows you to provide a body file, even in a case where no body is allowed. The body file must contain only comments and a single No\_Body pragma. This is recognized by the compiler as indicating that no body is logically present.

This is particularly useful during maintenance when a package is modified in such a way that a body needed before is no longer needed. The provision of a dummy body with a No\_Body pragma ensures that there is no interference from earlier versions of the package body.

## 2.113 Pragma No\_Caching

Syntax:

```
pragma No_Caching [ (static_boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect No\_Caching in the SPARK 2014 Reference Manual, section 7.1.2.

## 2.114 Pragma No\_Component\_Reordering

Syntax:

```
pragma No_Component_Reordering [[Entity =>] type_LOCAL_NAME];
```

type\_LOCAL\_NAME must refer to a record type declaration in the current declarative part. The effect is to preclude any reordering of components for the layout of the record, i.e. the record is laid out by the compiler in the order in which the components are declared textually. The form with no argument is a configuration pragma which applies to all record

types declared in units to which the pragma applies and there is a requirement that this pragma be used consistently within a partition.

### 2.115 Pragma `No_Elaboration_Code_All`

Syntax:

```
pragma No_Elaboration_Code_All [(program_unit_NAME)];
```

This is a program unit pragma (there is also an equivalent aspect of the same name) that establishes the restriction `No_Elaboration_Code` for the current unit and any extended main source units (body and subunits). It also has the effect of enforcing a transitive application of this aspect, so that if any unit is implicitly or explicitly with'ed by the current unit, it must also have the *No\_Elaboration\_Code\_All* aspect set. It may be applied to package or subprogram specs or their generic versions.

### 2.116 Pragma `No_Heap_Finalization`

Syntax:

```
pragma No_Heap_Finalization [ (first_subtype_LOCAL_NAME) ];
```

Pragma `No_Heap_Finalization` may be used as a configuration pragma or as a type-specific pragma.

In its configuration form, the pragma must appear within a configuration file such as `gnat.adc`, without an argument. The pragma suppresses the call to `Finalize` for heap-allocated objects created through library-level named access-to-object types in cases where the designated type requires finalization actions.

In its type-specific form, the argument of the pragma must denote a library-level named access-to-object type. The pragma suppresses the call to `Finalize` for heap-allocated objects created through the specific access type in cases where the designated type requires finalization actions.

It is still possible to finalize such heap-allocated objects by explicitly deallocating them.

A library-level named access-to-object type declared within a generic unit will lose its `No_Heap_Finalization` pragma when the corresponding instance does not appear at the library level.

### 2.117 Pragma `No_Inline`

Syntax:

```
pragma No_Inline (NAME {, NAME});
```

This pragma suppresses inlining for the callable entity or the instances of the generic subprogram designated by `NAME`, including inlining that results from the use of pragma `Inline`. This pragma is always active, in particular it is not subject to the use of option `'-gnatn'` or `'-gnatN'`. It is illegal to specify both pragma `No_Inline` and pragma `Inline_Always` for the same `NAME`.

## 2.118 Pragma No\_Raise

Syntax:

```
pragma No_Raise (subprogram_LOCAL_NAME {, subprogram_LOCAL_NAME});
```

Each `subprogram_LOCAL_NAME` argument must refer to one or more subprogram declarations in the current declarative part. A subprogram to which this pragma is applied may not raise an exception that is not caught within it. An implementation-defined check named *Raise\_Check* is associated with the pragma, and *Program\_Error* is raised upon its failure (see RM 11.5(19/5)).

## 2.119 Pragma No\_Return

Syntax:

```
pragma No_Return (procedure_LOCAL_NAME {, procedure_LOCAL_NAME});
```

Each `procedure_LOCAL_NAME` argument must refer to one or more procedure declarations in the current declarative part. A procedure to which this pragma is applied may not contain any explicit `return` statements. In addition, if the procedure contains any implicit returns from falling off the end of a statement sequence, then execution of that implicit return will cause *Program\_Error* to be raised.

One use of this pragma is to identify procedures whose only purpose is to raise an exception. Another use of this pragma is to suppress incorrect warnings about missing returns in functions, where the last statement of a function statement sequence is a call to such a procedure.

Note that in Ada 2005 mode, this pragma is part of the language. It is available in all earlier versions of Ada as an implementation-defined pragma.

## 2.120 Pragma No\_Strict\_Aliasing

Syntax:

```
pragma No_Strict_Aliasing [(Entity =>] type_LOCAL_NAME);
```

`type_LOCAL_NAME` must refer to an access type declaration in the current declarative part. The effect is to inhibit strict aliasing optimization for the given type. The form with no arguments is a configuration pragma which applies to all access types declared in units to which the pragma applies. For a detailed description of the strict aliasing optimization, and the situations in which it must be suppressed, see the section on Optimization and Strict Aliasing in the *GNAT User's Guide*.

This pragma currently has no effects on access to unconstrained array types.

## 2.121 Pragma No\_Tagged\_Streams

Syntax:

```
pragma No_Tagged_Streams [(Entity =>] tagged_type_LOCAL_NAME);
```

Normally when a tagged type is introduced using a full type declaration, part of the processing includes generating stream access routines to be used by stream attributes referencing the type (or one of its subtypes or derived types). This can involve the generation of significant amounts of code which is wasted space if stream routines are not needed for the type in question.

The `No_Tagged_Streams` pragma causes the generation of these stream routines to be skipped, and any attempt to use stream operations on types subject to this pragma will be statically rejected as illegal.

There are two forms of the pragma. The form with no arguments must appear in a declarative sequence or in the declarations of a package spec. This pragma affects all subsequent root tagged types declared in the declaration sequence, and specifies that no stream routines be generated. The form with an argument (for which there is also a corresponding aspect) specifies a single root tagged type for which stream routines are not to be generated.

Once the pragma has been given for a particular root tagged type, all subtypes and derived types of this type inherit the pragma automatically, so the effect applies to a complete hierarchy (this is necessary to deal with the class-wide dispatching versions of the stream routines).

When pragmas `Discard_Names` and `No_Tagged_Streams` are simultaneously applied to a tagged type its `Expanded_Name` and `External_Tag` are initialized with empty strings. This is useful to avoid exposing entity names at binary level but has a negative impact on the debuggability of tagged types.

## 2.122 Pragma `Normalize_Scalars`

Syntax:

```
pragma Normalize_Scalars;
```

This is a language defined pragma which is fully implemented in GNAT. The effect is to cause all scalar objects that are not otherwise initialized to be initialized. The initial values are implementation dependent and are as follows:

‘Standard.Character’

Objects whose root type is `Standard.Character` are initialized to `Character'Last` unless the subtype range excludes NUL (in which case NUL is used). This choice will always generate an invalid value if one exists.

‘Standard.Wide\_Character’

Objects whose root type is `Standard.Wide_Character` are initialized to `Wide_Character'Last` unless the subtype range excludes NUL (in which case NUL is used). This choice will always generate an invalid value if one exists.

‘Standard.Wide\_Wide\_Character’

Objects whose root type is `Standard.Wide_Wide_Character` are initialized to the invalid value `16#FFFF_FFFF#` unless the subtype range excludes NUL (in which case NUL is used). This choice will always generate an invalid value if one exists.

‘Integer types’

Objects of an integer type are treated differently depending on whether negative values are present in the subtype. If no negative values are present, then all one bits is used as the initial value except in the special case where zero is excluded from the subtype, in which case all zero bits are used. This choice will always generate an invalid value if one exists.

For subtypes with negative values present, the largest negative number is used, except in the unusual case where this largest negative number is in the subtype,

and the largest positive number is not, in which case the largest positive value is used. This choice will always generate an invalid value if one exists.

‘Floating-Point Types’

Objects of all floating-point types are initialized to all 1-bits. For standard IEEE format, this corresponds to a NaN (not a number) which is indeed an invalid value.

‘Fixed-Point Types’

Objects of all fixed-point types are treated as described above for integers, with the rules applying to the underlying integer value used to represent the fixed-point value.

‘Modular types’

Objects of a modular type are initialized to all one bits, except in the special case where zero is excluded from the subtype, in which case all zero bits are used. This choice will always generate an invalid value if one exists.

‘Enumeration types’

Objects of an enumeration type are initialized to all one-bits, i.e., to the value `2 ** typ'Size - 1` unless the subtype excludes the literal whose Pos value is zero, in which case a code of zero is used. This choice will always generate an invalid value if one exists.

## 2.123 Pragma Obsolescent

Syntax:

```
pragma Obsolescent;

pragma Obsolescent (
  [Message =>] static_string_EXPRESSION
  [, [Version =>] Ada_05]);

pragma Obsolescent (
  [Entity =>] NAME
  [, [Message =>] static_string_EXPRESSION
  [, [Version =>] Ada_05]]);
```

This pragma can occur immediately following a declaration of an entity, including the case of a record component. If no Entity argument is present, then this declaration is the one to which the pragma applies. If an Entity parameter is present, it must either match the name of the entity in this declaration, or alternatively, the pragma can immediately follow an enumeration type declaration, where the Entity argument names one of the enumeration literals.

This pragma is used to indicate that the named entity is considered obsolescent and should not be used. Typically this is used when an API must be modified by eventually removing or modifying existing subprograms or other entities. The pragma can be used at an intermediate stage when the entity is still present, but will be removed later.

The effect of this pragma is to output a warning message on a reference to an entity thus marked that the subprogram is obsolescent if the appropriate warning option in the compiler

is activated. If the **Message** parameter is present, then a second warning message is given containing this text. In addition, a reference to the entity is considered to be a violation of pragma **Restrictions** (**No\_Obsolescent\_Features**).

This pragma can also be used as a program unit pragma for a package, in which case the entity name is the name of the package, and the pragma indicates that the entire package is considered obsolescent. In this case a client **with**ing such a package violates the restriction, and the **with** clause is flagged with warnings if the warning option is set.

If the **Version** parameter is present (which must be exactly the identifier **Ada\_05**, no other argument is allowed), then the indication of obsolescence applies only when compiling in Ada 2005 mode. This is primarily intended for dealing with the situations in the predefined library where subprograms or packages have become defined as obsolescent in Ada 2005 (e.g., in **Ada.Characters.Handling**), but may be used anywhere.

The following examples show typical uses of this pragma:

```
package p is
  pragma Obsolescent (p, Message => "use pp instead of p");
end p;

package q is
  procedure q2;
  pragma Obsolescent ("use q2new instead");

  type R is new integer;
  pragma Obsolescent
    (Entity => R,
     Message => "use RR in Ada 2005",
     Version => Ada_05);

  type M is record
    F1 : Integer;
    F2 : Integer;
    pragma Obsolescent;
    F3 : Integer;
  end record;

  type E is (a, bc, 'd', quack);
  pragma Obsolescent (Entity => bc)
  pragma Obsolescent (Entity => 'd')

  function "+"
    (a, b : character) return character;
  pragma Obsolescent (Entity => "+");
end;
```

Note that, as for all pragmas, if you use a pragma argument identifier, then all subsequent parameters must also use a pragma argument identifier. So if you specify **Entity =>** for the **Entity** argument, and a **Message** argument is present, it must be preceded by **Message =>**.

## 2.124 Pragma Optimize\_Alignment

Syntax:

```
pragma Optimize_Alignment (TIME | SPACE | OFF);
```

This is a configuration pragma which affects the choice of default alignments for types and objects where no alignment is explicitly specified. There is a time/space trade-off in the selection of these values. Large alignments result in more efficient code, at the expense of larger data space, since sizes have to be increased to match these alignments. Smaller alignments save space, but the access code is slower. The normal choice of default alignments for types and individual alignment promotions for objects (which is what you get if you do not use this pragma, or if you use an argument of OFF), tries to balance these two requirements.

Specifying SPACE causes smaller default alignments to be chosen in two cases. First any packed record is given an alignment of 1. Second, if a size is given for the type, then the alignment is chosen to avoid increasing this size. For example, consider:

```
type R is record
  X : Integer;
  Y : Character;
end record;

for R'Size use 5*8;
```

In the default mode, this type gets an alignment of 4, so that access to the Integer field X are efficient. But this means that objects of the type end up with a size of 8 bytes. This is a valid choice, since sizes of objects are allowed to be bigger than the size of the type, but it can waste space if for example fields of type R appear in an enclosing record. If the above type is compiled in `Optimize_Alignment (Space)` mode, the alignment is set to 1.

However, there is one case in which SPACE is ignored. If a variable length record (that is a discriminated record with a component which is an array whose length depends on a discriminant), has a pragma Pack, then it is not in general possible to set the alignment of such a record to one, so the pragma is ignored in this case (with a warning).

Specifying SPACE also disables alignment promotions for standalone objects, which occur when the compiler increases the alignment of a specific object without changing the alignment of its type.

Specifying SPACE also disables component reordering in unpacked record types, which can result in larger sizes in order to meet alignment requirements.

Specifying TIME causes larger default alignments to be chosen in the case of small types with sizes that are not a power of 2. For example, consider:

```
type R is record
  A : Character;
  B : Character;
  C : Boolean;
end record;

pragma Pack (R);
for R'Size use 17;
```

The default alignment for this record is normally 1, but if this type is compiled in `Optimize_Alignment (Time)` mode, then the alignment is set to 4, which wastes space for objects of the type, since they are now 4 bytes long, but results in more efficient access when the whole record is referenced.

As noted above, this is a configuration pragma, and there is a requirement that all units in a partition be compiled with a consistent setting of the optimization setting. This would normally be achieved by use of a configuration pragma file containing the appropriate setting. The exception to this rule is that units with an explicit configuration pragma in the same file as the source unit are excluded from the consistency check, as are all predefined units. The latter are compiled by default in pragma `Optimize_Alignment (Off)` mode if no pragma appears at the start of the file.

## 2.125 Pragma Ordered

Syntax:

```
pragma Ordered (enumeration_first_subtype_LOCAL_NAME);
```

Most enumeration types are from a conceptual point of view unordered. For example, consider:

```
type Color is (Red, Blue, Green, Yellow);
```

By Ada semantics `Blue > Red` and `Green > Blue`, but really these relations make no sense; the enumeration type merely specifies a set of possible colors, and the order is unimportant.

For unordered enumeration types, it is generally a good idea if clients avoid comparisons (other than equality or inequality) and explicit ranges. (A ‘client’ is a unit where the type is referenced, other than the unit where the type is declared, its body, and its subunits.) For example, if code buried in some client says:

```
if Current_Color < Yellow then ...
if Current_Color in Blue .. Green then ...
```

then the client code is relying on the order, which is undesirable. It makes the code hard to read and creates maintenance difficulties if entries have to be added to the enumeration type. Instead, the code in the client should list the possibilities, or an appropriate subtype should be declared in the unit that declares the original enumeration type. E.g., the following subtype could be declared along with the type `Color`:

```
subtype RGB is Color range Red .. Green;
```

and then the client could write:

```
if Current_Color in RGB then ...
if Current_Color = Blue or Current_Color = Green then ...
```

However, some enumeration types are legitimately ordered from a conceptual point of view. For example, if you declare:

```
type Day is (Mon, Tue, Wed, Thu, Fri, Sat, Sun);
```

then the ordering imposed by the language is reasonable, and clients can depend on it, writing for example:

```
if D in Mon .. Fri then ...
if D < Wed then ...
```

The pragma ‘Ordered’ is provided to mark enumeration types that are conceptually ordered, alerting the reader that clients may depend on the ordering. GNAT provides a pragma to mark enumerations as ordered rather than one to mark them as unordered, since in our experience, the great majority of enumeration types are conceptually unordered.

The types `Boolean`, `Character`, `Wide_Character`, and `Wide_Wide_Character` are considered to be ordered types, so each is declared with a pragma `Ordered` in package `Standard`.

Normally pragma `Ordered` serves only as documentation and a guide for coding standards, but GNAT provides a warning switch ‘-gnatw.u’ that requests warnings for inappropriate uses (comparisons and explicit subranges) for unordered types. If this switch is used, then any enumeration type not marked with pragma `Ordered` will be considered as unordered, and will generate warnings for inappropriate uses.

Note that generic types are not considered ordered or unordered (since the template can be instantiated for both cases), so we never generate warnings for the case of generic enumerated types.

For additional information please refer to the description of the ‘-gnatw.u’ switch in the GNAT User’s Guide.

## 2.126 Pragma `Overflow_Mode`

Syntax:

```
pragma Overflow_Mode
( [General      =>] MODE
  [, [Assertions =>] MODE] );
```

MODE ::= STRICT | MINIMIZED | ELIMINATED

This pragma sets the current overflow mode to the given setting. For details of the meaning of these modes, please refer to the ‘Overflow Check Handling in GNAT’ appendix in the GNAT User’s Guide. If only the `General` parameter is present, the given mode applies to all expressions. If both parameters are present, the `General` mode applies to expressions outside assertions, and the `Eliminated` mode applies to expressions within assertions.

The case of the `MODE` parameter is ignored, so `MINIMIZED`, `Minimized` and `minimized` all have the same effect.

The `Overflow_Mode` pragma has the same scoping and placement rules as pragma `Suppress`, so it can occur either as a configuration pragma, specifying a default for the whole program, or in a declarative scope, where it applies to the remaining declarations and statements in that scope.

The pragma `Suppress (Overflow_Check)` suppresses overflow checking, but does not affect the overflow mode.

The pragma `Unsuppress (Overflow_Check)` unsuppresses (enables) overflow checking, but does not affect the overflow mode.

## 2.127 Pragma `Overriding_Renamings`

Syntax:

```
pragma Overriding_Renamings;
```

This is a GNAT configuration pragma to simplify porting legacy code accepted by the Rational Ada compiler. In the presence of this pragma, a renaming declaration that renames an inherited operation declared in the same scope is legal if selected notation is used as in:

```
pragma Overriding_Renamings;
...
package R is
  function F (...);
  ...
  function F (...) renames R.F;
end R;
```

even though RM 8.3 (15) stipulates that an overridden operation is not visible within the declaration of the overriding operation.

## 2.128 Pragma Part\_Of

Syntax:

```
pragma Part_Of (ABSTRACT_STATE);

ABSTRACT_STATE ::= NAME
```

For the semantics of this pragma, see the entry for aspect `Part_Of` in the SPARK 2014 Reference Manual, section 7.2.6.

## 2.129 Pragma Partition\_Elaboration\_Policy

Syntax:

```
pragma Partition_Elaboration_Policy (POLICY_IDENTIFIER);

POLICY_IDENTIFIER ::= Concurrent | Sequential
```

This pragma is standard in Ada 2005, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

## 2.130 Pragma Passive

Syntax:

```
pragma Passive [(Semaphore | No)];
```

Syntax checked, but otherwise ignored by GNAT. This is recognized for compatibility with DEC Ada 83 implementations, where it is used within a task definition to request that a task be made passive. If the argument `Semaphore` is present, or the argument is omitted, then DEC Ada 83 treats the pragma as an assertion that the containing task is passive and that optimization of context switch with this task is permitted and desired. If the argument `No` is present, the task must not be optimized. GNAT does not attempt to optimize any tasks in this manner (since protected objects are available in place of passive tasks).

For more information on the subject of passive tasks, see the section ‘Passive Task Optimization’ in the GNAT Users Guide.

### 2.131 Pragma Persistent\_BSS

Syntax:

```
pragma Persistent_BSS [(object_LOCAL_NAME)]
```

This pragma allows selected objects to be placed in the `.persistent_bss` section. On some targets the linker and loader provide for special treatment of this section, allowing a program to be reloaded without affecting the contents of this data (hence the name persistent).

There are two forms of usage. If an argument is given, it must be the local name of a library-level object, with no explicit initialization and whose type is potentially persistent. If no argument is given, then the pragma is a configuration pragma, and applies to all library-level objects with no explicit initialization of potentially persistent types.

A potentially persistent type is a scalar type, or an untagged, non-discriminated record, all of whose components have no explicit initialization and are themselves of a potentially persistent type, or an array, all of whose constraints are static, and whose component type is potentially persistent.

If this pragma is used on a target where this feature is not supported, then the pragma will be ignored. See also `pragma Linker_Section`.

### 2.132 Pragma Post

Syntax:

```
pragma Post (Boolean_Expression);
```

The `Post` pragma is intended to be an exact replacement for the language-defined `Post` aspect, and shares its restrictions and semantics. It must appear either immediately following the corresponding subprogram declaration (only other pragmas may intervene), or if there is no separate subprogram declaration, then it can appear at the start of the declarations in a subprogram body (preceded only by other pragmas).

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

### 2.133 Pragma Postcondition

Syntax:

```
pragma Postcondition (
  [Check    =>] Boolean_Expression
  [, [Message =>] String_Expression]);
```

The `Postcondition` pragma allows specification of automatic postcondition checks for subprograms. These checks are similar to assertions, but are automatically inserted just prior to the return statements of the subprogram with which they are associated (including implicit returns at the end of procedure bodies and associated exception handlers).

In addition, the boolean expression which is the condition which must be true may contain references to `function'Result` in the case of a function to refer to the returned value.

`Postcondition` pragmas may appear either immediately following the (separate) declaration of a subprogram, or at the start of the declarations of a subprogram body. Only other pragmas may intervene (that is appear between the subprogram declaration and its postconditions, or appear before the postcondition in the declaration sequence in a subprogram

body). In the case of a postcondition appearing after a subprogram declaration, the formal arguments of the subprogram are visible, and can be referenced in the postcondition expressions.

The postconditions are collected and automatically tested just before any return (implicit or explicit) in the subprogram body. A postcondition is only recognized if postconditions are active at the time the pragma is encountered. The compiler switch ‘gnata’ turns on all postconditions by default, and pragma `Check_Policy` with an identifier of `Postcondition` can also be used to control whether postconditions are active.

The general approach is that postconditions are placed in the spec if they represent functional aspects which make sense to the client. For example we might have:

```
function Direction return Integer;
pragma Postcondition
  (Direction'Result = +1
   or else
   Direction'Result = -1);
```

which serves to document that the result must be +1 or -1, and will test that this is the case at run time if postcondition checking is active.

Postconditions within the subprogram body can be used to check that some internal aspect of the implementation, not visible to the client, is operating as expected. For instance if a square root routine keeps an internal counter of the number of times it is called, then we might have the following postcondition:

```
Sqrt_Calls : Natural := 0;

function Sqrt (Arg : Float) return Float is
  pragma Postcondition
    (Sqrt_Calls = Sqrt_Calls'Old + 1);
  ...
end Sqrt
```

As this example, shows, the use of the `Old` attribute is often useful in postconditions to refer to the state on entry to the subprogram.

Note that postconditions are only checked on normal returns from the subprogram. If an abnormal return results from raising an exception, then the postconditions are not checked.

If a postcondition fails, then the exception `System.Assertions.Assert_Failure` is raised. If a message argument was supplied, then the given string will be used as the exception message. If no message argument was supplied, then the default message has the form “Postcondition failed at file\_name:line”. The exception is raised in the context of the subprogram body, so it is possible to catch postcondition failures within the subprogram body itself.

Within a package spec, normal visibility rules in Ada would prevent forward references within a postcondition pragma to functions defined later in the same package. This would introduce undesirable ordering constraints. To avoid this problem, all postcondition pragmas are analyzed at the end of the package spec, allowing forward references.

The following example shows that this even allows mutually recursive postconditions as in:

```
package Parity_Functions is
```

```

function Odd (X : Natural) return Boolean;
pragma Postcondition
  (Odd'Result =
    (x = 1
     or else
     (x /= 0 and then Even (X - 1))));

function Even (X : Natural) return Boolean;
pragma Postcondition
  (Even'Result =
    (x = 0
     or else
     (x /= 1 and then Odd (X - 1))));

end Parity_Functions;

```

There are no restrictions on the complexity or form of conditions used within `Postcondition` pragmas. The following example shows that it is even possible to verify performance behavior.

```

package Sort is

  Performance : constant Float;
  -- Performance constant set by implementation
  -- to match target architecture behavior.

  procedure Treesort (Arg : String);
  -- Sorts characters of argument using N*logN sort
  pragma Postcondition
    (Float (Clock - Clock'Old) <=
      Float (Arg'Length) *
      log (Float (Arg'Length)) *
      Performance);

end Sort;

```

Note: `postcondition` pragmas associated with subprograms that are marked as `Inline_Always`, or those marked as `Inline` with front-end inlining (`-gnatN` option set) are accepted and legality-checked by the compiler, but are ignored at run-time even if `postcondition` checking is enabled.

Note that `pragma Postcondition` differs from the language-defined `Post` aspect (and corresponding `Post` pragma) in allowing multiple occurrences, allowing occurrences in the body even if there is a separate spec, and allowing a second string parameter, and the use of the pragma identifier `Check`. Historically, `pragma Postcondition` was implemented prior to the development of Ada 2012, and has been retained in its original form for compatibility purposes.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

### 2.134 Pragma Post\_Class

Syntax:

```
pragma Post_Class (Boolean_Expression);
```

The `Post_Class` pragma is intended to be an exact replacement for the language-defined `Post'Class` aspect, and shares its restrictions and semantics. It must appear either immediately following the corresponding subprogram declaration (only other pragmas may intervene), or if there is no separate subprogram declaration, then it can appear at the start of the declarations in a subprogram body (preceded only by other pragmas).

Note: This pragma is called `Post_Class` rather than `Post'Class` because the latter would not be strictly conforming to the allowed syntax for pragmas. The motivation for providing pragmas equivalent to the aspects is to allow a program to be written using the pragmas, and then compiled if necessary using an Ada compiler that does not recognize the pragmas or aspects, but is prepared to ignore the pragmas. The assertion policy that controls this pragma is `Post'Class`, not `Post_Class`.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

### 2.135 Pragma Pre

Syntax:

```
pragma Pre (Boolean_Expression);
```

The `Pre` pragma is intended to be an exact replacement for the language-defined `Pre` aspect, and shares its restrictions and semantics. It must appear either immediately following the corresponding subprogram declaration (only other pragmas may intervene), or if there is no separate subprogram declaration, then it can appear at the start of the declarations in a subprogram body (preceded only by other pragmas).

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

### 2.136 Pragma Precondition

Syntax:

```
pragma Precondition (  
    [Check    =>] Boolean_Expression  
    [, [Message =>] String_Expression]);
```

The `Precondition` pragma is similar to `Postcondition` except that the corresponding checks take place immediately upon entry to the subprogram, and if a precondition fails, the exception is raised in the context of the caller, and the attribute `'Result` cannot be used within the precondition expression.

Otherwise, the placement and visibility rules are identical to those described for postconditions. The following is an example of use within a package spec:

```
package Math_Functions is  
    ...  
    function Sqrt (Arg : Float) return Float;  
    pragma Precondition (Arg >= 0.0)
```

```
...
end Math_Functions;
```

**Precondition** pragmas may appear either immediately following the (separate) declaration of a subprogram, or at the start of the declarations of a subprogram body. Only other pragmas may intervene (that is appear between the subprogram declaration and its postconditions, or appear before the postcondition in the declaration sequence in a subprogram body).

Note: precondition pragmas associated with subprograms that are marked as `Inline_Always`, or those marked as `Inline` with front-end inlining (`-gnatN` option set) are accepted and legality-checked by the compiler, but are ignored at run-time even if precondition checking is enabled.

Note that pragma **Precondition** differs from the language-defined **Pre** aspect (and corresponding **Pre** pragma) in allowing multiple occurrences, allowing occurrences in the body even if there is a separate spec, and allowing a second string parameter, and the use of the pragma identifier **Check**. Historically, pragma **Precondition** was implemented prior to the development of Ada 2012, and has been retained in its original form for compatibility purposes.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.137 Pragma Predicate

Syntax:

```
pragma Predicate
  ([Entity =>] type_LOCAL_NAME,
   [Check  =>] EXPRESSION);
```

This pragma (available in all versions of Ada in GNAT) encompasses both the **Static\_Predicate** and **Dynamic\_Predicate** aspects in Ada 2012. A predicate is regarded as static if it has an allowed form for **Static\_Predicate** and is otherwise treated as a **Dynamic\_Predicate**. Otherwise, predicates specified by this pragma behave exactly as described in the Ada 2012 reference manual. For example, if we have

```
type R is range 1 .. 10;
subtype S is R;
pragma Predicate (Entity => S, Check => S not in 4 .. 6);
subtype Q is R
pragma Predicate (Entity => Q, Check => F(Q) or G(Q));
```

the effect is identical to the following Ada 2012 code:

```
type R is range 1 .. 10;
subtype S is R with
  Static_Predicate => S not in 4 .. 6;
subtype Q is R with
  Dynamic_Predicate => F(Q) or G(Q);
```

Note that there are no pragmas **Dynamic\_Predicate** or **Static\_Predicate**. That is because these pragmas would affect legality and semantics of the program and thus do not have a neutral effect if ignored. The motivation behind providing pragmas equivalent to

corresponding aspects is to allow a program to be written using the pragmas, and then compiled with a compiler that will ignore the pragmas. That doesn't work in the case of static and dynamic predicates, since if the corresponding pragmas are ignored, then the behavior of the program is fundamentally changed (for example a membership test `A in B` would not take into account a predicate defined for subtype `B`). When following this approach, the use of predicates should be avoided.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

### 2.138 Pragma Predicate\_Failure

Syntax:

```
pragma Predicate_Failure
  ([Entity =>] type_LOCAL_NAME,
   [Message =>] String_Expression);
```

The `Predicate_Failure` pragma is intended to be an exact replacement for the language-defined `Predicate_Failure` aspect, and shares its restrictions and semantics.

### 2.139 Pragma Preelaborable\_Initialization

Syntax:

```
pragma Preelaborable_Initialization (DIRECT_NAME);
```

This pragma is standard in Ada 2005, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

### 2.140 Pragma Prefix\_Exception\_Messages

Syntax:

```
pragma Prefix_Exception_Messages;
```

This is an implementation-defined configuration pragma that affects the behavior of raise statements with a message given as a static string constant (typically a string literal). In such cases, the string will be automatically prefixed by the name of the enclosing entity (giving the package and subprogram containing the raise statement). This helps to identify where messages are coming from, and this mode is automatic for the run-time library.

The pragma has no effect if the message is computed with an expression other than a static string constant, since the assumption in this case is that the program computes exactly the string it wants. If you still want the prefixing in this case, you can always call `GNAT.Source_Info.Enclosing_Entity` and prepend the string manually.

### 2.141 Pragma Pre\_Class

Syntax:

```
pragma Pre_Class (Boolean_Expression);
```

The `Pre_Class` pragma is intended to be an exact replacement for the language-defined `Pre'Class` aspect, and shares its restrictions and semantics. It must appear either immediately following the corresponding subprogram declaration (only other pragmas may

intervene), or if there is no separate subprogram declaration, then it can appear at the start of the declarations in a subprogram body (preceded only by other pragmas).

Note: This pragma is called `Pre_Class` rather than `Pre'Class` because the latter would not be strictly conforming to the allowed syntax for pragmas. The motivation for providing pragmas equivalent to the aspects is to allow a program to be written using the pragmas, and then compiled if necessary using an Ada compiler that does not recognize the pragmas or aspects, but is prepared to ignore the pragmas. The assertion policy that controls this pragma is `Pre'Class`, not `Pre_Class`.

This is an assertion kind pragma that can associate a set of its arguments with an assertion level. See SPARK 2014 Reference Manual, section 11.4.2.

## 2.142 Pragma `Priority_Specific_Dispatching`

Syntax:

```
pragma Priority_Specific_Dispatching (
    POLICY_IDENTIFIER,
    first_priority_EXPRESSION,
    last_priority_EXPRESSION)
```

```
POLICY_IDENTIFIER ::=
    EDF_Across_Priorities           |
    FIFO_Within_Priorities         |
    Non_Preemptive_Within_Priorities |
    Round_Robin_Within_Priorities
```

This pragma is standard in Ada 2005, but is available in all earlier versions of Ada as an implementation-defined pragma. See Ada 2012 Reference Manual for details.

## 2.143 Pragma `Profile`

Syntax:

```
pragma Profile (Ravenscar | Restricted | Rational | Jorvik |
    GNAT_Extended_Ravenscar | GNAT_Ravenscar_EDF );
```

This pragma is standard in Ada 2005, but is available in all earlier versions of Ada as an implementation-defined pragma. This is a configuration pragma that establishes a set of configuration pragmas that depend on the argument. `Ravenscar` is standard in Ada 2005. `Jorvik` is standard in Ada 202x. The other possibilities (`Restricted`, `Rational`, `GNAT_Extended_Ravenscar`, `GNAT_Ravenscar_EDF`) are implementation-defined. `GNAT_Extended_Ravenscar` is an alias for `Jorvik`.

The set of configuration pragmas is defined in the following sections.

### \* Pragma `Profile` (`Ravenscar`)

The `Ravenscar` profile is standard in Ada 2005, but is available in all earlier versions of Ada as an implementation-defined pragma. This profile establishes the following set of configuration pragmas:

### \* `Task_Dispatching_Policy` (`FIFO_Within_Priorities`)

[RM D.2.2] Tasks are dispatched following a preemptive priority-ordered scheduling policy.

- \* `Locking_Policy (Ceiling_Locking)`  
[RM D.3] While tasks and interrupts execute a protected action, they inherit the ceiling priority of the corresponding protected object.
- \* `Detect_Blocking`  
This pragma forces the detection of potentially blocking operations within a protected operation, and to raise `Program_Error` if that happens.

plus the following set of restrictions:

- \* `Max_Entry_Queue_Length => 1`  
No task can be queued on a protected entry.
- \* `Max_Protected_Entries => 1`
- \* `Max_Task_Entries => 0`  
No rendezvous statements are allowed.
- \* `No_Abort_Statements`
- \* `No_Dynamic_Attachment`
- \* `No_Dynamic_Priorities`
- \* `No_Implicit_Heap_Allocations`
- \* `No_Local_Protected_Objects`
- \* `No_Local_Timing_Events`
- \* `No_Protected_Type_Allocators`
- \* `No_Relative_Delay`
- \* `No_Requeue_Statements`
- \* `No_Select_Statements`
- \* `No_Specific_Termination_Handlers`
- \* `No_Task_Allocators`
- \* `No_Task_Hierarchy`
- \* `No_Task_Termination`
- \* `Simple_Barriers`

The Ravenscar profile also includes the following restrictions that specify that there are no semantic dependencies on the corresponding predefined packages:

- \* `No_Dependence => Ada.Asynchronous_Task_Control`
- \* `No_Dependence => Ada.Calendar`
- \* `No_Dependence => Ada.Execution_Time.Group_Budget`
- \* `No_Dependence => Ada.Execution_Time.Timers`
- \* `No_Dependence => Ada.Task_Attributes`
- \* `No_Dependence => System.Multiprocessors.Dispatching_Domains`

This set of configuration pragmas and restrictions correspond to the definition of the ‘Ravenscar Profile’ for limited tasking, devised and published by the *International Real-Time Ada Workshop, 1997*. A description is also available at ‘<http://www-users.cs.york.ac.uk/~burns/ravenscar.ps>’.

The original definition of the profile was revised at subsequent IRTAW meetings. It has been included in the ISO *Guide for the Use of the Ada Programming Language in High Integrity Systems*, and was made part of the Ada 2005 standard. The formal definition given by the Ada Rapporteur Group (ARG) can be found in two Ada Issues (AI-249 and AI-305) available at ‘<http://www.ada-auth.org/cgi-bin/cvsweb.cgi/ais/ai-00249.txt>’ and ‘<http://www.ada-auth.org/cgi-bin/cvsweb.cgi/ais/ai-00305.txt>’.

The above set is a superset of the restrictions provided by pragma `Profile (Restricted)`, it includes six additional restrictions (`Simple_Barriers`, `No_Select_Statements`, `No_Calendar`, `No_Implicit_Heap_Allocations`, `No_Relative_Delay` and `No_Task_Termination`). This means that pragma `Profile (Ravenscar)`, like the pragma `Profile (Restricted)`, automatically causes the use of a simplified, more efficient version of the tasking run-time library.

\* Pragma `Profile (Jorvik)`

Jorvik is the new profile added to the Ada 202x draft standard, previously implemented under the name `GNAT_Extended_Ravenscar`.

The `No_Implicit_Heap_Allocations` restriction has been replaced by `No_Implicit_Task_Allocations` and `No_Implicit_Protected_Object_Allocations`.

The `Simple_Barriers` restriction has been replaced by `Pure_Barriers`.

The `Max_Protected_Entries`, `Max_Entry_Queue_Length`, and `No_Relative_Delay` restrictions have been removed.

Details on the rationale for Jorvik and implications for use may be found in *A New Ravenscar-Based Profile* by P. Rogers, J. Ruiz, T. Gingold and P. Bernardi, in *Reliable Software Technologies – Ada Europe 2017*, Springer-Verlag Lecture Notes in Computer Science, Number 10300.

\* Pragma `Profile (GNAT_Ravenscar_EDF)`

This profile corresponds to the Ravenscar profile but using `EDF_Across_Priority` as the `Task_Scheduling_Policy`.

\* Pragma `Profile (Restricted)`

This profile corresponds to the GNAT restricted run time. It establishes the following set of restrictions:

- \* `No_Abort_Statements`
- \* `No_Entry_Queue`
- \* `No_Task_Hierarchy`
- \* `No_Task_Allocators`
- \* `No_Dynamic_Priorities`
- \* `No_Terminate_Alternatives`
- \* `No_Dynamic_Attachment`
- \* `No_Protected_Type_Allocators`
- \* `No_Local_Protected_Objects`
- \* `No_Requeue_Statements`
- \* `No_Task_Attributes_Package`

```

* Max_Asynchronous_Select_Nesting = 0
* Max_Task_Entries = 0
* Max_Protected_Entries = 1
* Max_Select_Alternatives = 0

```

This set of restrictions causes the automatic selection of a simplified version of the run time that provides improved performance for the limited set of tasking functionality permitted by this set of restrictions.

\* **Pragma Profile (Rational)**

The Rational profile is intended to facilitate porting legacy code that compiles with the Rational APEX compiler, even when the code includes non- conforming Ada constructs. The profile enables the following three pragmas:

```

* pragma Implicit_Packing
* pragma Overriding_Renamings
* pragma Use_VADS_Size

```

## 2.144 Pragma Profile\_Warnings

Syntax:

```
pragma Profile_Warnings (Ravenscar | Restricted | Rational);
```

This is an implementation-defined pragma that is similar in effect to `pragma Profile` except that instead of generating `Restrictions` pragmas, it generates `Restriction_Warnings` pragmas. The result is that violations of the profile generate warning messages instead of error messages.

## 2.145 Pragma Program\_Exit

Syntax:

```
pragma Program_Exit [ (boolean_EXPRESSION) ];
```

For the semantics of this pragma, see the entry for aspect `Program_Exit` in the SPARK 2014 Reference Manual, section 6.1.10.

## 2.146 Pragma Propagate\_Exceptions

Syntax:

```
pragma Propagate_Exceptions;
```

This pragma is now obsolete and, other than generating a warning if warnings on obsolescent features are enabled, is ignored. It is retained for compatibility purposes. It used to be used in connection with optimization of a now-obsolete mechanism for implementation of exceptions.

## 2.147 Pragma Provide\_Shift\_Operators

Syntax:

```
pragma Provide_Shift_Operators (integer_first_subtype_LOCAL_NAME);
```











To obtain the required output specified in RM H.3.1, the compiler must be run with various special switches as follows:

- \* ‘Where compiler-generated run-time checks remain’  
The switch ‘-gnatGL’ may be used to list the expanded code in pseudo-Ada form. Runtime checks show up in the listing either as explicit checks or operators marked with {} to indicate a check is present.
- \* ‘An identification of known exceptions at compile time’  
If the program is compiled with ‘-gnatwa’, the compiler warning messages will indicate all cases where the compiler detects that an exception is certain to occur at run time.
- \* ‘Possible reads of uninitialized variables’  
The compiler warns of many such cases, but its output is incomplete.

A supplemental static analysis tool may be used to obtain a comprehensive list of all possible points at which uninitialized data may be read.

- \* ‘Where run-time support routines are implicitly invoked’  
In the output from ‘-gnatGL’, run-time calls are explicitly listed as calls to the relevant run-time routine.
- \* ‘Object code listing’  
This may be obtained either by using the ‘-S’ switch, or the objdump utility.
- \* ‘Constructs known to be erroneous at compile time’  
These are identified by warnings issued by the compiler (use ‘-gnatwa’).
- \* ‘Stack usage information’  
Static stack usage data (maximum per-subprogram) can be obtained via the ‘-fstack-usage’ switch to the compiler. Dynamic stack usage data (per task) can be obtained via the ‘-u’ switch to gnatbind
- \* ‘Object code listing of entire partition’  
This can be obtained by compiling the partition with ‘-S’, or by applying objdump to all the object files that are part of the partition.
- \* ‘A description of the run-time model’  
The full sources of the run-time are available, and the documentation of these routines describes how these run-time routines interface to the underlying operating system facilities.
- \* ‘Control and data-flow information’

A supplemental static analysis tool may be used to obtain complete control and data-flow information, as well as comprehensive messages identifying possible problems based on this information.

## 2.162 Pragma Secondary\_Stack\_Size

Syntax:

```
pragma Secondary_Stack_Size (integer_EXPRESSION);
```

This pragma appears within the task definition of a single task declaration or a task type declaration (like pragma `Storage_Size`) and applies to all task objects of that type. The















```

-- switch on layout checks
pragma Style_Checks ("1");
-- set the number of maximum allowed nesting levels to 15
pragma Style_Checks ("L15");

*

gcc -c -gnatyl -gnatyL15 ...

```

The string literal values can be cumulatively switched on and off by prefixing the value with + or -, where:

- \* + is equivalent to no prefix. It applies the check referenced by the literal value;
- \* - switches the referenced check off.

```

-- allow misaligned block by disabling layout check
pragma Style_Checks ("-1");
declare
  msg : constant String := "Hello";
begin
  Put_Line (msg);
end;

-- enable the layout check again
pragma Style_Checks ("1");
declare
  msg : constant String := "Hello";
begin
  Put_Line (msg);
end;

```

The code above contains two layout errors, however, only the last line is picked up by the compiler.

Similarly, the switches containing a numeric value can be applied in sequence. In the example below, the permitted nesting level is reduced in in the middle block and the compiler raises a warning on the highlighted line.

```

-- Permit 3 levels of nesting
pragma Style_Checks ("L3");

procedure Main is
begin
  if True then
    if True then
      null;
    end if;
  end if;
  -- Reduce permitted nesting levels to 2.
  -- Note that "+L2" and "L2" are equivalent.
  pragma Style_Checks ("L2");
  if True then
    if True then

```

```

        null;
    end if;
end if;
-- Disable checking permitted nesting levels.
-- Note that the number after "-L" is insignificant,
-- "-L", "-L3" and "-Lx" are all equivalent.
pragma Style_Checks ("-L3");
if True then
    if True then
        null;
    end if;
end if;
end Main;

```

The form `ALL_CHECKS` activates all standard checks (its use is equivalent to the use of the `gnaty` switch with no options. See the *GNAT User's Guide* for details.)

Note: the behavior is slightly different in GNAT mode (`-gnatg` used). In this case, `ALL_CHECKS` implies the standard set of GNAT mode style check options (i.e. equivalent to `-gnatyg`).

The forms with `Off` and `On` can be used to temporarily disable style checks as shown in the following example:

```

pragma Style_Checks ("k"); -- requires keywords in lower case
pragma Style_Checks (Off); -- turn off style checks
NULL;                      -- this will not generate an error message
pragma Style_Checks (On);  -- turn style checks back on
NULL;                      -- this will generate an error message

```

Finally the two argument form is allowed only if the first argument is `On` or `Off`. The effect is to turn of semantic style checks for the specified entity, as shown in the following example:

```

pragma Style_Checks ("r"); -- require consistency of identifier casing
Arg : Integer;
Rf1 : Integer := ARG;      -- incorrect, wrong case
pragma Style_Checks (Off, Arg);
Rf2 : Integer := ARG;      -- OK, no error

```

## 2.176 Pragma Subprogram\_Variant

Syntax:

```

pragma Subprogram_Variant (SUBPROGRAM_VARIANT_LIST);

SUBPROGRAM_VARIANT_LIST ::=
    STRUCTURAL_SUBPROGRAM_VARIANT_ITEM
| NUMERIC_SUBPROGRAM_VARIANT_ITEMS

NUMERIC_SUBPROGRAM_VARIANT_ITEMS ::=
    NUMERIC_SUBPROGRAM_VARIANT_ITEM {, NUMERIC_SUBPROGRAM_VARIANT_ITEM}

NUMERIC_SUBPROGRAM_VARIANT_ITEM ::=

```



- \* Additional check names previously introduced by use of the **Check\_Name** pragma are also allowed.

Note that pragma **Suppress** gives the compiler permission to omit checks, but does not require the compiler to omit checks. The compiler will generate checks if they are essentially free, even when they are suppressed. In particular, if the compiler can prove that a certain check will necessarily fail, it will generate code to do an unconditional ‘raise’, even if checks are suppressed. The compiler warns in this case.

Of course, run-time checks are omitted whenever the compiler can prove that they will not fail, whether or not checks are suppressed.

## 2.179 Pragma **Suppress\_All**

Syntax:

```
pragma Suppress_All;
```

This pragma can appear anywhere within a unit. The effect is to apply **Suppress (All\_Checks)** to the unit in which it appears. This pragma is implemented for compatibility with DEC Ada 83 usage where it appears at the end of a unit, and for compatibility with Rational Ada, where it appears as a program unit pragma. The use of the standard Ada pragma **Suppress (All\_Checks)** as a normal configuration pragma is the preferred usage in GNAT.

## 2.180 Pragma **Suppress\_Debug\_Info**

Syntax:

```
pragma Suppress_Debug_Info ([Entity =>] LOCAL_NAME);
```

This pragma can be used to suppress generation of debug information for the specified entity. It is intended primarily for use in debugging the debugger, and navigating around debugger problems.

## 2.181 Pragma **Suppress\_Exception\_Locations**

Syntax:

```
pragma Suppress_Exception_Locations;
```

In normal mode, a raise statement for an exception by default generates an exception message giving the file name and line number for the location of the raise. This is useful for debugging and logging purposes, but this entails extra space for the strings for the messages. The configuration pragma **Suppress\_Exception\_Locations** can be used to suppress the generation of these strings, with the result that space is saved, but the exception message for such raises is null. This configuration pragma may appear in a global configuration pragma file, or in a specific unit as usual. It is not required that this pragma be used consistently within a partition, so it is fine to have some units within a partition compiled with this pragma and others compiled in normal mode without it.



```

        pragma Task_Name (Name.all);
    end Task_Typ;

    task body Task_Typ is
        Nam : constant String := Image (Current_Task);
    begin
        Put_Line ("-->" & Nam (1 .. 14) & "<--");
    end Task_Typ;

    type Ptr_Task is access Task_Typ;
    Task_Var : Ptr_Task;

begin
    Task_Var :=
        new Task_Typ (new String'("This is task 1"));
    Task_Var :=
        new Task_Typ (new String'("This is task 2"));
end;
```

## 2.184 Pragma Task\_Storage

Syntax:

```

pragma Task_Storage (
    [Task_Type =>] LOCAL_NAME,
    [Top_Guard =>] static_integer_EXPRESSION);
```

This pragma specifies the length of the guard area for tasks. The guard area is an additional storage area allocated to a task. A value of zero means that either no guard area is created or a minimal guard area is created, depending on the target. This pragma can appear anywhere a `Storage_Size` attribute definition clause is allowed for a task type.

## 2.185 Pragma Test\_Case

Syntax:

```

pragma Test_Case (
    [Name      =>] static_string_Expression
    , [Mode     =>] (Nominal | Robustness)
    [, Requires => Boolean_Expression]
    [, Ensures  => Boolean_Expression]);
```

The `Test_Case` pragma allows defining fine-grain specifications for use by testing tools. The compiler checks the validity of the `Test_Case` pragma, but its presence does not lead to any modification of the code generated by the compiler.

`Test_Case` pragmas may only appear immediately following the (separate) declaration of a subprogram in a package declaration, inside a package spec unit. Only other pragmas may intervene (that is appear between the subprogram declaration and a test case).

The compiler checks that boolean expressions given in `Requires` and `Ensures` are valid, where the rules for `Requires` are the same as the rule for an expression in `Precondition`













external names of tagged types and wants to ensure that the duplicated tag check occurs even if all run-time checks are suppressed by a compiler switch, the following configuration pragma will ensure this test is not suppressed:

```
pragma Unsuppress (Duplicated_Tag_Check);
```

This pragma is standard in Ada 2005. It is available in all earlier versions of Ada as an implementation-defined pragma.

Note that in addition to the checks defined in the Ada RM, GNAT recognizes a number of implementation-defined check names. See the description of pragma `Suppress` for full details.

## 2.201 Pragma Unused

Syntax:

```
pragma Unused (LOCAL_NAME {, LOCAL_NAME});
```

This pragma signals that the assignable entities (variables, `out` parameters, and `in out` parameters) whose names are listed deliberately do not get assigned or referenced in the current source unit after the occurrence of the pragma in the current source unit. This suppresses warnings about the entities that are unreferenced and/or not assigned, and, in addition, a warning will be generated if one of these entities gets assigned or subsequently referenced in the same unit as the pragma (in the corresponding body or one of its subunits).

This is particularly useful for clearly signaling that a particular parameter is not modified or referenced, even though the spec suggests that it might be.

For the variable case, warnings are never given for unreferenced variables whose name contains one of the substrings `DISCARD`, `DUMMY`, `IGNORE`, `JUNK`, `UNUSE`, `TMP`, `TEMP` in any casing. Such names are typically to be used in cases where such warnings are expected. Thus it is never necessary to use `pragma Unused` for such variables, though it is harmless to do so.

## 2.202 Pragma Use\_VADS\_Size

Syntax:

```
pragma Use_VADS_Size;
```

This is a configuration pragma. In a unit to which it applies, any use of the ‘Size attribute is automatically interpreted as a use of the ‘VADS\_Size attribute. Note that this may result in incorrect semantic processing of valid Ada 95 or Ada 2005 programs. This is intended to aid in the handling of existing code which depends on the interpretation of Size as implemented in the VADS compiler. See description of the VADS\_Size attribute for further details.

## 2.203 Pragma Validity\_Checks

Syntax:

```
pragma Validity_Checks (string_LITERAL | ALL_CHECKS | On | Off);
```

This pragma is used in conjunction with compiler switches to control the built-in validity checking provided by GNAT. The compiler switches, if set provide an initial setting for the switches, and this pragma may be used to modify these settings, or the settings may be

provided entirely by the use of the pragma. This pragma can be used anywhere that a pragma is legal, including use as a configuration pragma (including use in the `gnat.adc` file).

The form with a string literal specifies which validity options are to be activated. The validity checks are first set to include only the default reference manual settings, and then a string of letters in the string specifies the exact set of options required. The form of this string is exactly as described for the ‘-gnatVx’ compiler switch (see the GNAT User’s Guide for details). For example the following two methods can be used to enable validity checking for mode `in` and `in out` subprogram parameters:

```
*
    pragma Validity_Checks ("im");
*

$ gcc -c -gnatVim ...
```

The form `ALL_CHECKS` activates all standard checks (its use is equivalent to the use of the `gnatVa` switch).

The forms with `Off` and `On` can be used to temporarily disable validity checks as shown in the following example:

```
pragma Validity_Checks ("c"); -- validity checks for copies
pragma Validity_Checks (Off); -- turn off validity checks
A := B;                       -- B will not be validity checked
pragma Validity_Checks (On);  -- turn validity checks back on
A := C;                       -- C will be validity checked
```

## 2.204 Pragma Volatile

Syntax:

```
pragma Volatile (LOCAL_NAME);
```

This pragma is defined by the Ada Reference Manual, and the GNAT implementation is fully conformant with this definition. The reason it is mentioned in this section is that a pragma of the same name was supplied in some Ada 83 compilers, including DEC Ada 83. The Ada 95 / Ada 2005 implementation of pragma `Volatile` is upwards compatible with the implementation in DEC Ada 83.

## 2.205 Pragma Volatile\_Full\_Access

Syntax:

```
pragma Volatile_Full_Access (LOCAL_NAME);
```

This is similar in effect to pragma `Volatile`, except that any reference to the object is guaranteed to be done only with instructions that read or write all the bits of the object. Furthermore, if the object is of a composite type, then any reference to a subcomponent of the object is guaranteed to read and/or write all the bits of the object.

The intention is that this be suitable for use with memory-mapped I/O devices on some machines. Note that there are two important respects in which this is different from pragma `Atomic`. First a reference to a `Volatile_Full_Access` object is not a sequential action in



```

2. function Warnerr return String is
3.   X : Integer;
   |
   >>> error: variable "X" is never read and
       never assigned [-gnatwv] [warning-as-error]

4.   Y : Integer;
   |
   >>> warning: variable "Y" is assigned but
       never read [-gnatwu]

5. begin
6.   Y := 0;
7.   return %ABC%;
   |
   >>> error: use of "%" is an obsolescent
       feature (RM J.2(4)), use "" instead
       [-gnatwj] [warning-as-error]

8. end;
```

8 lines: No errors, 3 warnings (2 treated as errors)

Note that this pragma does not affect the set of warnings issued in any way, it merely changes the effect of a matching warning if one is produced as a result of other warnings options. As shown in this example, if the pragma results in a warning being treated as an error, the tag is changed from “warning:” to “error:” and the string “[warning-as-error]” is appended to the end of the message.

## 2.208 Pragma Warnings

Syntax:

```
pragma Warnings ([TOOL_NAME,] DETAILS [, REASON]);
```

```
DETAILS ::= On | Off
```

```
DETAILS ::= On | Off, local_NAME
```

```
DETAILS ::= static_string_EXPRESSION
```

```
DETAILS ::= On | Off, static_string_EXPRESSION
```

```
TOOL_NAME ::= GNAT | GNATprove
```

```
REASON ::= Reason => STRING_LITERAL {& STRING_LITERAL}
```

Note: in Ada 83 mode, a string literal may be used in place of a static string expression (which does not exist in Ada 83).

Note if the second argument of `DETAILS` is a `local_NAME` then the second form is always understood. If the intention is to use the fourth form, then you can write `NAME & ""` to force the interpretation as a ‘static\_string\_EXPRESSION’.





**LOCAL\_NAME** must refer to an object that is declared at the library level. This pragma specifies that the given entity should be marked as a weak symbol for the linker. It is equivalent to `__attribute__((weak))` in GNU C and causes **LOCAL\_NAME** to be emitted as a weak symbol instead of a regular symbol, that is to say a symbol that does not have to be resolved by the linker if used in conjunction with a pragma **Import**.

When a weak symbol is not resolved by the linker, its address is set to zero. This is useful in writing interfaces to external modules that may or may not be linked in the final executable, for example depending on configuration settings.

If a program references at run time an entity to which this pragma has been applied, and the corresponding symbol was not resolved at link time, then the execution of the program is erroneous. It is not erroneous to take the Address of such an entity, for example to guard potential references, as shown in the example below.

Some file formats do not support weak symbols so not all target machines support this pragma.

```
-- Example of the use of pragma Weak_External
```

```
package External_Module is
  key : Integer;
  pragma Import (C, key);
  pragma Weak_External (key);
  function Present return boolean;
end External_Module;

with System; use System;
package body External_Module is
  function Present return boolean is
  begin
    return key'Address /= System.Null_Address;
  end Present;
end External_Module;
```

## 2.210 Pragma **Wide\_Character-Encoding**

Syntax:

```
pragma Wide_Character-Encoding (IDENTIFIER | CHARACTER_LITERAL);
```

This pragma specifies the wide character encoding to be used in program source text appearing subsequently. It is a configuration pragma, but may also be used at any point that a pragma is allowed, and it is permissible to have more than one such pragma in a file, allowing multiple encodings to appear within the same file.

However, note that the pragma cannot immediately precede the relevant wide character, because then the previous encoding will still be in effect, causing “illegal character” errors.

The argument can be an identifier or a character literal. In the identifier case, it is one of **HEX**, **UPPER**, **SHIFT\_JIS**, **EUC**, **UTF8**, or **BRACKETS**. In the character literal case it is correspondingly one of the characters **h**, **u**, **s**, **e**, **8**, or **b**.

Note that when the pragma is used within a file, it affects only the encoding within that file, and does not affect withed units, specs, or subunits.

### 3 Implementation Defined Aspects

Ada defines (throughout the Ada 2012 reference manual, summarized in Annex K) a set of aspects that can be specified for certain entities. These language defined aspects are implemented in GNAT in Ada 2012 mode and work as described in the Ada 2012 Reference Manual.

In addition, Ada 2012 allows implementations to define additional aspects whose meaning is defined by the implementation. GNAT provides a number of these implementation-defined aspects which can be used to extend and enhance the functionality of the compiler. This section of the GNAT reference manual describes these additional aspects.

Note that any program using these aspects may not be portable to other compilers (although GNAT implements this set of aspects on all platforms). Therefore if portability to other compilers is an important consideration, you should minimize the use of these aspects.

Note that for many of these aspects, the effect is essentially similar to the use of a pragma or attribute specification with the same name applied to the entity. For example, if we write:

```
type R is range 1 .. 100
  with Value_Size => 10;
```

then the effect is the same as:

```
type R is range 1 .. 100;
for R'Value_Size use 10;
```

and if we write:

```
type R is new Integer
  with Shared => True;
```

then the effect is the same as:

```
type R is new Integer;
pragma Shared (R);
```

In the documentation below, such cases are simply marked as being boolean aspects equivalent to the corresponding pragma or attribute definition clause.

#### 3.1 Aspect `Abstract_State`

This aspect is equivalent to `[pragma Abstract_State]`, page 5.

#### 3.2 Aspect `Always_Terminates`

This boolean aspect is equivalent to `[pragma Always_Terminates]`, page 9.

#### 3.3 Aspect `Annotate`

There are three forms of this aspect (where ID is an identifier, and ARG is a general expression), corresponding to `[pragma Annotate]`, page 10.

`'Annotate => ID'`

Equivalent to `pragma Annotate (ID, Entity => Name);`

```

‘Annotate => (ID)’
    Equivalent to pragma Annotate (ID, Entity => Name);
‘Annotate => (ID ,ID {, ARG})’
    Equivalent to pragma Annotate (ID, ID {, ARG}, Entity => Name);

```

### 3.4 Aspect Async\_Readers

This boolean aspect is equivalent to `[pragma Async_Readers]`, page 14.

### 3.5 Aspect Async\_Writers

This boolean aspect is equivalent to `[pragma Async_Writers]`, page 14.

### 3.6 Aspect Constant\_After\_Elaboration

This aspect is equivalent to `[pragma Constant_After_Elaboration]`, page 21.

### 3.7 Aspect Contract\_Cases

This aspect is equivalent to `[pragma Contract_Cases]`, page 21, the sequence of clauses being enclosed in parentheses so that syntactically it is an aggregate.

### 3.8 Aspect Depends

This aspect is equivalent to `[pragma Depends]`, page 26.

### 3.9 Aspect Default\_Initial\_Condition

This aspect is equivalent to `[pragma Default_Initial_Condition]`, page 25.

### 3.10 Aspect Dimension

The `Dimension` aspect is used to specify the dimensions of a given subtype of a dimensioned numeric type. The aspect also specifies a symbol used when doing formatted output of dimensioned quantities. The syntax is:

```

with Dimension =>
    ([Symbol =>] SYMBOL, DIMENSION_VALUE {, DIMENSION_Value})

SYMBOL ::= STRING_LITERAL | CHARACTER_LITERAL

DIMENSION_VALUE ::=
    RATIONAL
  | others                => RATIONAL
  | DISCRETE_CHOICE_LIST => RATIONAL

RATIONAL ::= [-] NUMERIC_LITERAL [/ NUMERIC_LITERAL]

```

This aspect can only be applied to a subtype whose parent type has a `Dimension_System` aspect. The aspect must specify values for all dimensions of the system. The rational

values are the powers of the corresponding dimensions that are used by the compiler to verify that physical (numeric) computations are dimensionally consistent. For example, the computation of a force must result in dimensions ( $L \Rightarrow 1$ ,  $M \Rightarrow 1$ ,  $T \Rightarrow -2$ ). For further examples of the usage of this aspect, see package `System.Dim.Mks`. Note that when the dimensioned type is an integer type, then any dimension value must be an integer literal.

### 3.11 Aspect `Dimension_System`

The `Dimension_System` aspect is used to define a system of dimensions that will be used in subsequent subtype declarations with `Dimension` aspects that reference this system. The syntax is:

```
with Dimension_System => (DIMENSION {, DIMENSION});

DIMENSION ::= ([Unit_Name    =>] IDENTIFIER,
               [Unit_Symbol =>] SYMBOL,
               [Dim_Symbol  =>] SYMBOL)

SYMBOL ::= CHARACTER_LITERAL | STRING_LITERAL
```

This aspect is applied to a type, which must be a numeric derived type (typically a floating-point type), that will represent values within the dimension system. Each `DIMENSION` corresponds to one particular dimension. A maximum of 7 dimensions may be specified. `Unit_Name` is the name of the dimension (for example `Meter`). `Unit_Symbol` is the shorthand used for quantities of this dimension (for example `m` for `Meter`). `Dim_Symbol` gives the identification within the dimension system (typically this is a single letter, e.g. `L` standing for length for unit name `Meter`). The `Unit_Symbol` is used in formatted output of dimensioned quantities. The `Dim_Symbol` is used in error messages when numeric operations have inconsistent dimensions.

GNAT provides the standard definition of the International MKS system in the run-time package `System.Dim.Mks`. You can easily define similar packages for cgs units or British units, and define conversion factors between values in different systems. The MKS system is characterized by the following aspect:

```
type Mks_Type is new Long_Long_Float with
  Dimension_System => (
    (Unit_Name => Meter,    Unit_Symbol => 'm',   Dim_Symbol => 'L'),
    (Unit_Name => Kilogram, Unit_Symbol => "kg",   Dim_Symbol => 'M'),
    (Unit_Name => Second,   Unit_Symbol => 's',    Dim_Symbol => 'T'),
    (Unit_Name => Ampere,    Unit_Symbol => 'A',    Dim_Symbol => 'I'),
    (Unit_Name => Kelvin,   Unit_Symbol => 'K',    Dim_Symbol => '@'),
    (Unit_Name => Mole,      Unit_Symbol => "mol",  Dim_Symbol => 'N'),
    (Unit_Name => Candela,   Unit_Symbol => "cd",   Dim_Symbol => 'J'));
```

Note that in the above type definition, we use the `at` symbol (`@`) to represent a theta character (avoiding the use of extended Latin-1 characters in this context).

See section ‘Performing Dimensionality Analysis in GNAT’ in the GNAT Users Guide for detailed examples of use of the dimension system.

### 3.12 Aspect `Disable_Controlled`

The aspect `Disable_Controlled` is defined for controlled record types. If active, this aspect causes suppression of all related calls to `Initialize`, `Adjust`, and `Finalize`. The intended use is for conditional compilation, where for example you might want a record to be controlled or not depending on whether some run-time check is enabled or suppressed.

### 3.13 Aspect `Effective_Reads`

This aspect is equivalent to `[pragma Effective_Reads]`, page 28.

### 3.14 Aspect `Effective_Writes`

This aspect is equivalent to `[pragma Effective_Writes]`, page 28.

### 3.15 Aspect `Exceptional_Cases`

This aspect may be specified for procedures and functions with side effects; it can be used to list exceptions that might be propagated by the subprogram with side effects in the context of its precondition, and associate them with a specific postcondition.

For the syntax and semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.1.9.

### 3.16 Aspect `Exit_Cases`

This aspect may be specified for procedures and functions with side effects; it can be used to partition the input state into a list of cases and specify, for each case, how the subprogram is allowed to terminate (i.e. return normally or propagate an exception).

For the syntax and semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.1.10.

### 3.17 Aspect `Extended_Access`

This nonoverridable boolean-valued type-related representation aspect can be specified as part of a `full_type_declaration` for a general access type designating an unconstrained array subtype.

The absence of an `Extended_Access` aspect specification for such a `full_type_declaration` is equivalent to an explicit “`Extended_Access => False`” specification. This implies that the aspect is never unspecified for an eligible access type. An access type for which this aspect is `True` is said to be an extended access type; this includes the case of a type derived from an extended access type. Similarly, a value of such a type is said to be an extended access value.

The representation of an extended access value is different than that of other access values. This representation makes it possible to designate objects that cannot be designated using the usual “thin” or “fat” access representations for an access type designating an unconstrained array subtype (notably slices and array objects imported from other languages).

In particular, two rules are modified in determining the legality of an `Access` or `Unchecked_Access` attribute reference if the expected access type is an extended access type:

- \* A slice of an aliased array object of a non-bitpacked type (more precisely, of an array type having independently addressable components) is considered to be aliased (and the accessibility level of a slice of an array object is defined to be that of the array object); this also applies to renamings of such slices, slices of such renamings, etc.
- \* The requirement that the nominal subtype of the prefix shall statically match the designated subtype of the access type need not be met.

The Size aspect (and other aspects including Stream\_Size, Object\_Size, and Alignment) of an extended access type may depend on the properties of the designated type. Further details of this dependence are not documented.

An extended access value is not convertible to a non-extended access type, although conversions in the opposite direction are allowed. We don't want to allow

```
type Big_Ref is access all String with Extended_Access;
type Small_Ref is access all String;
Obj : aliased String := "abcde";
Big_Ptr : Big_Ref := Obj (2 .. 4)'Access; -- OK
Small_Ptr : Small_Ref := Small_Ref (Big_Ptr); -- ERROR: illegal conversion
```

because there is no way to represent the result of such a conversion.

A dereference of an extended access value (or a reference to a renaming thereof) shall not occur in any of the following contexts:

- \* as an operative constituent of the prefix of an Access or Unchecked\_Access attribute reference whose expected type is not extended; or
- \* as an operative constituent of an actual parameter in a call where the corresponding formal parameter is explicitly aliased.

For the same reasons that explicit conversions from an extended access type to a non-extended access type are forbidden, we also need to disallow getting the same effect via a Extended\_Ptr.all'Access reference; this includes the case of passing Extended\_Ptr.all as an actual parameter in a call where the corresponding formal parameter is explicitly aliased (because the callee could evaluate Formal\_Parameter'Access). This goal is accomplished by adjusting the definition of the term “aliased”. A dereference of an extended value occurring in one of these contexts is defined to denote a nonaliased view. This has the desired effect because these contexts require an aliased view. Continuing the preceding example, this rule disallows

```
Sneaky_1 : Small_Ptr := Big_Ptr.all'Access; -- ERROR: illegal 'Access prefix

function Make (Str : aliased in out String) return Small_Ptr
is (Str'Access); -- OK

Sneaky_2 : Small_Ptr := Make (Str => Big_Ptr.all); -- ERROR: bad parameter
```

for the same reason given above in the case of an explicit type conversion.

### 3.18 Aspect Extensions\_Visible

This aspect is equivalent to [pragma Extensions\_Visible], page 35.

### 3.19 Aspect Favor\_Top\_Level

This boolean aspect is equivalent to [pragma Favor\_Top\_Level], page 37.

### 3.20 Aspect Ghost

This aspect is equivalent to [pragma Ghost], page 38.

### 3.21 Aspect Ghost\_Predicate

This aspect introduces a subtype predicate that can reference ghost entities. The subtype cannot appear as a subtype\_mark in a membership test.

For the detailed semantics of this aspect, see the entry for subtype predicates in the SPARK Reference Manual, section 3.2.4.

### 3.22 Aspect Global

This aspect is equivalent to [pragma Global], page 38.

### 3.23 Aspect Initial\_Condition

This aspect is equivalent to [pragma Initial\_Condition], page 44.

### 3.24 Aspect Initializes

This aspect is equivalent to [pragma Initializes], page 46.

### 3.25 Aspect Inline\_Always

This boolean aspect is equivalent to [pragma Inline\_Always], page 46.

### 3.26 Aspect Invariant

This aspect is equivalent to [pragma Invariant], page 49. It is a synonym for the language defined aspect `Type_Invariant` except that it is separately controllable using pragma `Assertion_Policy`.

### 3.27 Aspect Invariant'Class

This aspect is equivalent to [pragma Type\_Invariant\_Class], page 97. It is a synonym for the language defined aspect `Type_Invariant'Class` except that it is separately controllable using pragma `Assertion_Policy`.

### 3.28 Aspect Iterable

This aspect provides a light-weight mechanism for loops and quantified expressions over container types, without the overhead imposed by the tampering checks of standard Ada 2012 iterators. The value of the aspect is an aggregate with six named components, of which the last three are optional: `First`, `Next`, `Has_Element`, `Element`, `Last`, and `Previous`. When only the first three components are specified, only the `for .. in` form of iteration



include (only) `No_Heap_Allocations` and `No_Secondary_Stack`. `No_Secondary_Stack` corresponds to the GNAT-defined (global) restriction of the same name. `No_Heap_Allocations` corresponds to the conjunction of the Ada-defined restrictions `No_Allocators` and `No_Implicit_Heap_Allocations`.

Additional requirements are imposed in order to ensure that restriction violations cannot be achieved via overriding dispatching operations, calling through an access-to-subprogram value, calling a generic formal subprogram, or calling through a subprogram renaming. For a dispatching operation, an overrider must be subject to (at least) the same restrictions as the overridden inherited subprogram; similarly, the actual subprogram corresponding to a generic formal subprogram in an instantiation must be subject to (at least) the same restrictions as the formal subprogram. A call through an access-to-subprogram value is conservatively assumed to violate all local restrictions; tasking-related constructs (notably entry calls) are treated similarly. A renaming-as-body is treated like a subprogram body containing a call to the renamed subprogram.

The `Local_Restrictions` aspect can be specified for a package specification, in which case the aspect specification also applies to all eligible entities declared with the package. This includes types. Default initialization of an object of a given type is treated like a call to an implicitly-declared initialization subprogram. Such a “call” is subject to the same local restriction checks as any other call. If a type is subject to a local restriction, then any violations of that restriction within the default initialization expressions (if any) of the type are rejected. This may include “calls” to the default initialization subprograms of other types.

`Local_Restrictions` aspect specifications are additive (for example, in the case of a declaration that occurs within nested packages that each have a `Local_Restrictions` specification).

### 3.31 Aspect `Lock_Free`

This boolean aspect is equivalent to `[pragma Lock_Free]`, page 53.

### 3.32 Aspect `Max_Queue_Length`

This aspect is equivalent to `[pragma Max_Queue_Length]`, page 57.

### 3.33 Aspect `No_Caching`

This boolean aspect is equivalent to `[pragma No_Caching]`, page 57.

### 3.34 Aspect `No_Elaboration_Code_All`

This aspect is equivalent to `[pragma No_Elaboration_Code_All]`, page 58, for a program unit.

### 3.35 Aspect `No_Inline`

This boolean aspect is equivalent to `[pragma No_Inline]`, page 58.

### 3.36 Aspect `No_Raise`

This boolean aspect is equivalent to `[pragma No_Raise]`, page 58.

### 3.37 Aspect No\_Tagged\_Streams

This aspect is equivalent to [pragma No\_Tagged\_Streams], page 59, with an argument specifying a root tagged type (thus this aspect can only be applied to such a type).

### 3.38 Aspect No\_Task\_Parts

Applies to a type. If True, requires that the type and any descendants do not have any task parts. The rules for this aspect are the same as for the language-defined No\_Controlled\_Parts aspect (see RM-H.4.1), replacing “controlled” with “task”.

If No\_Task\_Parts is True for a type T, then the compiler can optimize away certain tasking-related code that would otherwise be needed for T'Class, because descendants of T might contain tasks.

### 3.39 Aspect Object\_Size

This aspect is equivalent to [attribute Object\_Size], page 130.

### 3.40 Aspect Obsolescent

This aspect is equivalent to [pragma Obsolescent], page 61. Note that the evaluation of this aspect happens at the point of occurrence, it is not delayed until the freeze point.

### 3.41 Aspect Part\_Of

This aspect is equivalent to [pragma Part\_Of], page 66.

### 3.42 Aspect Persistent\_BSS

This boolean aspect is equivalent to [pragma Persistent\_BSS], page 66.

### 3.43 Aspect Potentially\_Invalid

For the syntax and semantics of this aspect, see the SPARK 2014 Reference Manual, section 13.9.1.

### 3.44 Aspect Predicate

This aspect is equivalent to [pragma Predicate], page 71. It is thus similar to the language defined aspects `Dynamic_Predicate` and `Static_Predicate` except that whether the resulting predicate is static or dynamic is controlled by the form of the expression. It is also separately controllable using pragma `Assertion_Policy`.

### 3.45 Aspect Program\_Exit

This boolean aspect is equivalent to [pragma Program\_Exit], page 76.

### 3.46 Aspect Pure\_Function

This boolean aspect is equivalent to [pragma Pure\_Function], page 77.

### 3.47 Aspect Refined\_Depends

This aspect is equivalent to [pragma Refined\_Depends], page 78.

### 3.48 Aspect Refined\_Global

This aspect is equivalent to [pragma Refined\_Global], page 79.

### 3.49 Aspect Refined\_Post

This aspect is equivalent to [pragma Refined\_Post], page 79.

### 3.50 Aspect Refined\_State

This aspect is equivalent to [pragma Refined\_State], page 79.

### 3.51 Aspect Relaxed\_Initialization

For the syntax and semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.10.

### 3.52 Aspect Remote\_Access\_Type

This aspect is equivalent to [pragma Remote\_Access\_Type], page 80.

### 3.53 Aspect Scalar\_Storage\_Order

This aspect is equivalent to a [attribute Scalar\_Storage\_Order], page 133.

### 3.54 Aspect Secondary\_Stack\_Size

This aspect is equivalent to [pragma Secondary\_Stack\_Size], page 82.

### 3.55 Aspect Shared

This boolean aspect is equivalent to [pragma Shared], page 83, and is thus a synonym for aspect `Atomic`.

### 3.56 Aspect Side\_Effects

This aspect is equivalent to [pragma Side\_Effects], page 84.

### 3.57 Aspect Simple\_Storage\_Pool

This aspect is equivalent to [attribute Simple\_Storage\_Pool], page 136.

### 3.58 Aspect Simple\_Storage\_Pool\_Type

This boolean aspect is equivalent to [pragma Simple\_Storage\_Pool\_Type], page 84.

### 3.59 Aspect SPARK\_Mode

This aspect is equivalent to `[pragma SPARK_Mode]`, page 87, and may be specified for either or both of the specification and body of a subprogram or package.

### 3.60 Aspect Subprogram\_Variant

For the syntax and semantics of this aspect, see the SPARK 2014 Reference Manual, section 6.1.8.

### 3.61 Aspect Suppress\_Debug\_Info

This boolean aspect is equivalent to `[pragma Suppress_Debug_Info]`, page 93.

### 3.62 Aspect Suppress\_Initialization

This boolean aspect is equivalent to `[pragma Suppress_Initialization]`, page 93.

### 3.63 Aspect Test\_Case

This aspect is equivalent to `[pragma Test_Case]`, page 95.

### 3.64 Aspect Thread\_Local\_Storage

This boolean aspect is equivalent to `[pragma Thread_Local_Storage]`, page 96.

### 3.65 Aspect Universal\_Aliasing

This boolean aspect is equivalent to `[pragma Universal_Aliasing]`, page 99.

### 3.66 Aspect Unmodified

This boolean aspect is equivalent to `[pragma Unmodified]`, page 99.

### 3.67 Aspect Unreferenced

This boolean aspect is equivalent to `[pragma Unreferenced]`, page 100.

When using the `-gnat2022` switch, this aspect is also supported on formal parameters, which is in particular the only form possible for expression functions.

### 3.68 Aspect Unreferenced\_Objects

This boolean aspect is equivalent to `[pragma Unreferenced_Objects]`, page 100.

### 3.69 Aspect User\_Aspect

This aspect takes an argument that is the name of an aspect defined by a `User_Aspect_Definition` configuration pragma. A `User_Aspect` aspect specification is semantically equivalent to replicating the set of aspect specifications associated with the named pragma-defined aspect.

### **3.70 Aspect Value\_Size**

This aspect is equivalent to [attribute Value\_Size], page 144.

### **3.71 Aspect Volatile\_Full\_Access**

This boolean aspect is equivalent to [pragma Volatile\_Full\_Access], page 103.

### **3.72 Aspect Volatile\_Function**

This boolean aspect is equivalent to [pragma Volatile\_Function], page 104.

### **3.73 Aspect Warnings**

This aspect is equivalent to the two argument form of [pragma Warnings], page 105, where the first argument is **ON** or **OFF** and the second argument is the entity.

## 4 Implementation Defined Attributes

Ada defines (throughout the Ada reference manual, summarized in Annex K), a set of attributes that provide useful additional functionality in all areas of the language. These language defined attributes are implemented in GNAT and work as described in the Ada Reference Manual.

In addition, Ada allows implementations to define additional attributes whose meaning is defined by the implementation. GNAT provides a number of these implementation-dependent attributes which can be used to extend and enhance the functionality of the compiler. This section of the GNAT reference manual describes these additional attributes. It also describes additional implementation-dependent features of standard language-defined attributes.

Note that any program using these attributes may not be portable to other compilers (although GNAT implements this set of attributes on all platforms). Therefore if portability to other compilers is an important consideration, you should minimize the use of these attributes.

### 4.1 Attribute `Abort_Signal`

`Standard'Abort_Signal` (`Standard` is the only allowed prefix) provides the entity for the special exception used to signal task abort or asynchronous transfer of control. Normally this attribute should only be used in the tasking runtime (it is highly peculiar, and completely outside the normal semantics of Ada, for a user program to intercept the abort exception).

### 4.2 Attribute `Address_Size`

`Standard'Address_Size` (`Standard` is the only allowed prefix) is a static constant giving the number of bits in an `Address`. It is the same value as `System.Address'Size`, but has the advantage of being static, while a direct reference to `System.Address'Size` is nonstatic because `Address` is a private type.

### 4.3 Attribute `Asm_Input`

The `Asm_Input` attribute denotes a function that takes two parameters. The first is a string, the second is an expression of the type designated by the prefix. The first (string) argument is required to be a static expression, and is the constraint for the parameter, (e.g., what kind of register is required). The second argument is the value to be used as the input argument. The possible values for the constant are the same as those used in the RTL, and are dependent on the configuration file used to build the GCC back end. [Machine Code Insertions], page 287,

### 4.4 Attribute `Asm_Output`

The `Asm_Output` attribute denotes a function that takes two parameters. The first is a string, the second is the name of a variable of the type designated by the attribute prefix. The first (string) argument is required to be a static expression and designates the constraint for the parameter (e.g., what kind of register is required). The second argument is the variable to be updated with the result. The possible values for constraint are the same as

those used in the RTL, and are dependent on the configuration file used to build the GCC back end. If there are no output operands, then this argument may either be omitted, or explicitly given as `No_Output_Operands`. [Machine Code Insertions], page 287,

## 4.5 Attribute `Atomic_Always_Lock_Free`

The prefix of the `Atomic_Always_Lock_Free` attribute is a type. The result indicates whether atomic operations are supported by the target for the given type.

## 4.6 Attribute `Bit`

`obj'Bit`, where `obj` is any object, yields the bit offset within the storage unit (byte) that contains the first bit of storage allocated for the object. The value of this attribute is of the type ‘universal\_integer’ and is always a nonnegative number smaller than `System.Storage_Unit`.

For an object that is a variable or a constant allocated in a register, the value is zero. (The use of this attribute does not force the allocation of a variable to memory).

For an object that is a formal parameter, this attribute applies to either the matching actual parameter or to a copy of the matching actual parameter.

For an access object the value is zero. Note that `obj.all'Bit` is subject to an `Access_Check` for the designated object. Similarly for a record component `X.C'Bit` is subject to a discriminant check and `X(I).Bit` and `X(I1..I2)'Bit` are subject to index checks.

This attribute is designed to be compatible with the DEC Ada 83 definition and implementation of the `Bit` attribute.

## 4.7 Attribute `Bit_Position`

`R.C'Bit_Position`, where `R` is a record object and `C` is one of the fields of the record type, yields the bit offset within the record contains the first bit of storage allocated for the object. The value of this attribute is of the type ‘universal\_integer’. The value depends only on the field `C` and is independent of the alignment of the containing record `R`.

## 4.8 Attribute `Code_Address`

The `'Address` attribute may be applied to subprograms in Ada 95 and Ada 2005, but the intended effect seems to be to provide an address value which can be used to call the subprogram by means of an address clause as in the following example:

```
procedure K is ...

procedure L;
for L'Address use K'Address;
pragma Import (Ada, L);
```

A call to `L` is then expected to result in a call to `K`. In Ada 83, where there were no access-to-subprogram values, this was a common work-around for getting the effect of an indirect call. GNAT implements the above use of `Address` and the technique illustrated by the example code works correctly.

However, for some purposes, it is useful to have the address of the start of the generated code for the subprogram. On some architectures, this is not necessarily the same as the `Address` value described above. For example, the `Address` value may reference a subprogram descriptor rather than the subprogram itself.

The `'Code_Address` attribute, which can only be applied to subprogram entities, always returns the address of the start of the generated code of the specified subprogram, which may or may not be the same value as is returned by the corresponding `'Address` attribute.

## 4.9 Attribute `Compiler_Version`

`Standard'Compiler_Version` (`Standard` is the only allowed prefix) yields a static string identifying the version of the compiler being used to compile the unit containing the attribute reference.

## 4.10 Attribute `Constrained`

In addition to the usage of this attribute in the Ada RM, GNAT also permits the use of the `'Constrained` attribute in a generic template for any type, including types without discriminants. The value of this attribute in the generic instance when applied to a scalar type or a record type without discriminants is always `True`. This usage is compatible with older Ada compilers, including notably DEC Ada.

## 4.11 Attribute `Default_Bit_Order`

`Standard'Default_Bit_Order` (`Standard` is the only allowed prefix), provides the value `System.Default_Bit_Order` as a `Pos` value (0 for `High_Order_First`, 1 for `Low_Order_First`). This is used to construct the definition of `Default_Bit_Order` in package `System`.

## 4.12 Attribute `Default_Scalar_Storage_Order`

`Standard'Default_Scalar_Storage_Order` (`Standard` is the only allowed prefix), provides the current value of the default scalar storage order (as specified using pragma `Default_Scalar_Storage_Order`, or equal to `Default_Bit_Order` if unspecified) as a `System.Bit_Order` value. This is a static attribute.

## 4.13 Attribute `Deref`

The attribute `typ'Deref(expr)` where `expr` is of type `System.Address` yields the variable of type `typ` that is located at the given address. It is similar to `(totyp(expr)).all`, where `totyp` is an unchecked conversion from address to a named access-to-`typ` type, except that it yields a variable, so it can be used on the left side of an assignment.

## 4.14 Attribute `Descriptor_Size`

Nonstatic attribute `Descriptor_Size` returns the size in bits of the descriptor allocated for a type. The result is non-zero only for unconstrained array types and the returned value is of type universal integer. In GNAT, an array descriptor contains bounds information and is located immediately before the first element of the array.

```
type Unconstr_Array is array (Short_Short_Integer range <>) of Positive;
```

```
Put_Line ("Descriptor size = " & Unconstr_Array'Descriptor_Size'Img);
```

The attribute takes into account any padding due to the alignment of the component type. In the example above, the descriptor contains two values of type `Short_Short_Integer` representing the low and high bound. But, since `Positive` has an alignment of 4, the size of the descriptor is `2 * Short_Short_Integer'Size` rounded up to the next multiple of 32, which yields a size of 32 bits, i.e. including 16 bits of padding.

## 4.15 Attribute Elaborated

The prefix of the `'Elaborated` attribute must be a unit name. The value is a Boolean which indicates whether or not the given unit has been elaborated. This attribute is primarily intended for internal use by the generated code for dynamic elaboration checking, but it can also be used in user programs. The value will always be `True` once elaboration of all units has been completed. An exception is for units which need no elaboration, the value is always `False` for such units.

## 4.16 Attribute Elab\_Body

This attribute can only be applied to a program unit name. It returns the entity for the corresponding elaboration procedure for elaborating the body of the referenced unit. This is used in the main generated elaboration procedure by the binder and is not normally used in any other context. However, there may be specialized situations in which it is useful to be able to call this elaboration procedure from Ada code, e.g., if it is necessary to do selective re-elaboration to fix some error.

## 4.17 Attribute Elab\_Spec

This attribute can only be applied to a program unit name. It returns the entity for the corresponding elaboration procedure for elaborating the spec of the referenced unit. This is used in the main generated elaboration procedure by the binder and is not normally used in any other context. However, there may be specialized situations in which it is useful to be able to call this elaboration procedure from Ada code, e.g., if it is necessary to do selective re-elaboration to fix some error.

## 4.18 Attribute Elab\_Subp\_Body

This attribute can only be applied to a library level subprogram name and is only allowed in CodePeer mode. It returns the entity for the corresponding elaboration procedure for elaborating the body of the referenced subprogram unit. This is used in the main generated elaboration procedure by the binder in CodePeer mode only and is unrecognized otherwise.

## 4.19 Attribute Emax

The `Emax` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

## 4.20 Attribute Enabled

The **Enabled** attribute allows an application program to check at compile time to see if the designated check is currently enabled. The prefix is a simple identifier, referencing any predefined check name (other than **All\_Checks**) or a check name introduced by **pragma Check\_Name**. If no argument is given for the attribute, the check is for the general state of the check, if an argument is given, then it is an entity name, and the check indicates whether an **Suppress** or **Unsuppress** has been given naming the entity (if not, then the argument is ignored).

Note that instantiations inherit the check status at the point of the instantiation, so a useful idiom is to have a library package that introduces a check name with **pragma Check\_Name**, and then contains generic packages or subprograms which use the **Enabled** attribute to see if the check is enabled. A user of this package can then issue a **pragma Suppress** or **pragma Unsuppress** before instantiating the package or subprogram, controlling whether the check will be present.

## 4.21 Attribute Enum\_Rep

Note that this attribute is now standard in Ada 202x and is available as an implementation defined attribute for earlier Ada versions.

For every enumeration subtype **S**, **S'Enum\_Rep** denotes a function with the following spec:

```
function S'Enum_Rep (Arg : S'Base) return <Universal_Integer>;
```

It is also allowable to apply **Enum\_Rep** directly to an object of an enumeration type or to a non-overloaded enumeration literal. In this case **S'Enum\_Rep** is equivalent to **typ'Enum\_Rep(S)** where **typ** is the type of the enumeration literal or object.

The function returns the representation value for the given enumeration value. This will be equal to value of the **Pos** attribute in the absence of an enumeration representation clause. This is a static attribute (i.e., the result is static if the argument is static).

**S'Enum\_Rep** can also be used with integer types and objects, in which case it simply returns the integer value. The reason for this is to allow it to be used for (<>) discrete formal arguments in a generic unit that can be instantiated with either enumeration types or integer types. Note that if **Enum\_Rep** is used on a modular type whose upper bound exceeds the upper bound of the largest signed integer type, and the argument is a variable, so that the universal integer calculation is done at run time, then the call to **Enum\_Rep** may raise **Constraint\_Error**.

## 4.22 Attribute Enum\_Val

Note that this attribute is now standard in Ada 202x and is available as an implementation defined attribute for earlier Ada versions.

For every enumeration subtype **S**, **S'Enum\_Val** denotes a function with the following spec:

```
function S'Enum_Val (Arg : <Universal_Integer>) return S'Base;
```

The function returns the enumeration value whose representation matches the argument, or raises **Constraint\_Error** if no enumeration literal of the type has the matching value. This will be equal to value of the **Val** attribute in the absence of an enumeration representation clause. This is a static attribute (i.e., the result is static if the argument is static).

### 4.23 Attribute Epsilon

The `Epsilon` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.24 Attribute Fast\_Math

`Standard'Fast_Math` (`Standard` is the only allowed prefix) yields a static Boolean value that is `True` if pragma `Fast_Math` is active, and `False` otherwise.

### 4.25 Attribute Finalization\_Size

The prefix of attribute `Finalization_Size` must be an object or a non-class-wide type. This attribute returns the size of any hidden data reserved by the compiler to handle finalization-related actions. The type of the attribute is ‘universal\_integer’.

`Finalization_Size` yields a value of zero for a type with no controlled parts, an object whose type has no controlled parts, or an object of a class-wide type whose tag denotes a type with no controlled parts.

Note that only heap-allocated objects contain finalization data.

### 4.26 Attribute Fixed\_Value

For every fixed-point type `S`, `S'Fixed_Value` denotes a function with the following specification:

```
function S'Fixed_Value (Arg : <Universal_Integer>) return S;
```

The value returned is the fixed-point value `V` such that:

```
V = Arg * S'Small
```

The effect is thus similar to first converting the argument to the integer type used to represent `S`, and then doing an unchecked conversion to the fixed-point type. The difference is that there are full range checks, to ensure that the result is in range. This attribute is primarily intended for use in implementation of the input-output functions for fixed-point values.

### 4.27 Attribute From\_Address

The prefix of this attribute must be a general access-to-array type (or subtype); the attribute takes a `System.Address` argument and possibly some additional arguments (described below) and yields a value of the given access type that designates an array object located at the given address. In the case of a non-null array this means that the given address is the address of the first element of the array object (not the address of any sort of bounds descriptor). This allows associating bounds with an address that is, for example, passed in from C code.

If the designated array subtype is unconstrained (which is the usual case), then for each dimension (in order) the attribute takes either one or two additional arguments of the corresponding index type - one if the index subtype is a fixed lower bound subtype, two (low bound first) otherwise. In this case, the access type shall be an extended access type (see the description of the `Extended_Access` aspect). These additional arguments specify the bounds of the designated array object.

If the designated array subtype is constrained then no additional arguments are provided and the bounds of the designated object are those of the designated subtype.

Roughly speaking, `My_Access_Type'From_Address (Addr, Lo, Hi)` is equivalent to a declare expression:

```
(declare
  Obj : aliased Designated_Array_Type (Lo .. Hi)
  with Import, Address => Addr;
begin
  My_Access_Type'(Obj'Unchecked_Access)
end)
```

## 4.28 Attribute `From_Any`

This internal attribute is used for the generation of remote subprogram stubs in the context of the Distributed Systems Annex.

## 4.29 Attribute `Has_Access_Values`

The prefix of the `Has_Access_Values` attribute is a type. The result is a Boolean value which is True if the is an access type, or is a composite type with a component (at any nesting depth) that is an access type, and is False otherwise. The intended use of this attribute is in conjunction with generic definitions. If the attribute is applied to a generic private type, it indicates whether or not the corresponding actual type has access values.

## 4.30 Attribute `Has_Discriminants`

The prefix of the `Has_Discriminants` attribute is a type. The result is a Boolean value which is True if the type has discriminants, and False otherwise. The intended use of this attribute is in conjunction with generic definitions. If the attribute is applied to a generic private type, it indicates whether or not the corresponding actual type has discriminants.

## 4.31 Attribute `Has_Tagged_Values`

The prefix of the `Has_Tagged_Values` attribute is a type. The result is a Boolean value which is True if the type is a composite type (array or record) that is either a tagged type or has a subcomponent that is tagged, and is False otherwise. The intended use of this attribute is in conjunction with generic definitions. If the attribute is applied to a generic private type, it indicates whether or not the corresponding actual type has access values.

## 4.32 Attribute `Img`

The `Img` attribute differs from `Image` in that, while both can be applied directly to an object, `Img` cannot be applied to types.

Example usage of the attribute:

```
Put_Line ("X = " & X'Img);
```

which has the same meaning as the more verbose:

```
Put_Line ("X = " & T'Image (X));
```

where  $T$  is the (sub)type of the object  $X$ .

Note that technically, in analogy to `Image`, `X'Img` returns a parameterless function that returns the appropriate string when called. This means that `X'Img` can be renamed as a function-returning-string, or used in an instantiation as a function parameter.

### 4.33 Attribute Initialized

For the syntax and semantics of this attribute, see the SPARK 2014 Reference Manual, section 6.10.

### 4.34 Attribute Integer\_Value

For every integer type  $S$ , `S'Integer_Value` denotes a function with the following spec:

```
function S'Integer_Value (Arg : <Universal_Fixed>) return S;
```

The value returned is the integer value  $V$ , such that:

```
Arg = V * T'Small
```

where  $T$  is the type of `Arg`. The effect is thus similar to first doing an unchecked conversion from the fixed-point type to its corresponding implementation type, and then converting the result to the target integer type. The difference is that there are full range checks, to ensure that the result is in range. This attribute is primarily intended for use in implementation of the standard input-output functions for fixed-point values.

### 4.35 Attribute Invalid\_Value

For every scalar type  $S$ , `S'Invalid_Value` returns an undefined value of the type. If possible this value is an invalid representation for the type. The value returned is identical to the value used to initialize an otherwise uninitialized value of the type if `pragma Initialize Scalars` is used, including the ability to modify the value with the binder `-Sxx` flag and relevant environment variables at run time.

### 4.36 Attribute Large

The `Large` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.37 Attribute Library\_Level

`P'Library_Level`, where  $P$  is an entity name, returns a Boolean value which is `True` if the entity is declared at the library level, and `False` otherwise. Note that within a generic instantiation, the name of the generic unit denotes the instance, which means that this attribute can be used to test if a generic is instantiated at the library level, as shown in this example:

```
generic
  ...
package Gen is
  pragma Compile_Time_Error
    (not Gen'Library_Level,
```

```
        "Gen can only be instantiated at library level");  
    ...  
end Gen;
```

### 4.38 Attribute Loop\_Entry

Syntax:

```
X'Loop_Entry [(loop_name)]
```

The `Loop_Entry` attribute is used to refer to the value that an expression had upon entry to a given loop in much the same way that the `Old` attribute in a subprogram postcondition can be used to refer to the value an expression had upon entry to the subprogram. The relevant loop is either identified by the given loop name, or it is the innermost enclosing loop when no loop name is given.

A `Loop_Entry` attribute can only occur within an `Assert`, `Assert_And_Cut`, `Assume`, `Loop_Variant` or `Loop_Invariant` pragma. In addition, such a pragma must be one of the items in the sequence of statements of a loop body, or nested inside block statements that appear in the sequence of statements of a loop body. A common use of `Loop_Entry` is to compare the current value of objects with their initial value at loop entry, in a `Loop_Invariant` pragma.

The effect of using `X'Loop_Entry` is the same as declaring a constant initialized with the initial value of `X` at loop entry. This copy is not performed if the loop is not entered, or if the corresponding pragmas are ignored or disabled.

### 4.39 Attribute Machine\_Size

This attribute is identical to the `Object_Size` attribute. It is provided for compatibility with the DEC Ada 83 attribute of this name.

### 4.40 Attribute Mantissa

The `Mantissa` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.41 Attribute Maximum\_Alignment

`Standard'Maximum_Alignment` (`Standard` is the only allowed prefix) provides the maximum default alignment value for the target, that is to say the maximum alignment that the compiler may choose by default for a type or an object. Larger alignments are supported up to some maximum value dependent on the target, but may require specific mechanisms that are not needed up to `Standard'Maximum_Alignment`.

### 4.42 Attribute Max\_Integer\_Size

`Standard'Max_Integer_Size` (`Standard` is the only allowed prefix) provides the size of the largest supported integer type for the target. The result is a static constant.

### 4.43 Attribute Mechanism\_Code

`func 'Mechanism_Code` yields an integer code for the mechanism used for the result of function `func`, and `subprog 'Mechanism_Code (n)` yields the mechanism used for formal parameter number 'n' (a static integer value, with 1 meaning the first parameter) of subprogram `subprog`. The code returned is:

```
'1'
    by copy (value)
'2'
    by reference
```

### 4.44 Attribute Null\_Parameter

A reference `T'Null_Parameter` denotes an imaginary object of type or subtype `T` allocated at machine address zero. The attribute is allowed only as the default expression of a formal parameter, or as an actual expression of a subprogram call. In either case, the subprogram must be imported.

The identity of the object is represented by the address zero in the argument list, independent of the passing mechanism (explicit or default).

This capability is needed to specify that a zero address should be passed for a record or other composite object passed by reference. There is no way of indicating this without the `Null_Parameter` attribute.

### 4.45 Attribute Object\_Size

The size of an object is not necessarily the same as the size of the type of an object. This is because by default object sizes are increased to be a multiple of the alignment of the object. For example, `Natural'Size` is 31, but by default objects of type `Natural` will have a size of 32 bits. Similarly, a record containing an integer and a character:

```
type Rec is record
  I : Integer;
  C : Character;
end record;
```

will have a size of 40 (that is `Rec'Size` will be 40). The alignment will be 4, because of the integer field, and so the default size of record objects for this type will be 64 (8 bytes).

If the alignment of the above record is specified to be 1, then the object size will be 40 (5 bytes). This is true by default, and also an object size of 40 can be explicitly specified in this case.

A consequence of this capability is that different object sizes can be given to subtypes that would otherwise be considered in Ada to be statically matching. But it makes no sense to consider such subtypes as statically matching. Consequently, GNAT adds a rule to the static matching rules that requires object sizes to match. Consider this example:

```
1. procedure BadAVConvert is
2.   type R is new Integer;
3.   subtype R1 is R range 1 .. 10;
```

```

4.    subtype R2 is R range 1 .. 10;
5.    for R1'Object_Size use 8;
6.    for R2'Object_Size use 16;
7.    type R1P is access all R1;
8.    type R2P is access all R2;
9.    R1PV : R1P := new R1'(4);
10.   R2PV : R2P;
11. begin
12.   R2PV := R2P (R1PV);
      |
      >>> target designated subtype not compatible with
          type "R1" defined at line 3

13. end;
```

In the absence of lines 5 and 6, types R1 and R2 statically match and hence the conversion on line 12 is legal. But since lines 5 and 6 cause the object sizes to differ, GNAT considers that types R1 and R2 are not statically matching, and line 12 generates the diagnostic shown above.

Similar additional checks are performed in other contexts requiring statically matching subtypes.

## 4.46 Attribute Old

In addition to the usage of `Old` defined in the Ada 2012 RM (usage within `Post` aspect), GNAT also permits the use of this attribute in implementation defined pragmas `Postcondition`, `Contract_Cases` and `Test_Case`. Also usages of `Old` which would be illegal according to the Ada 2012 RM definition are allowed under control of implementation defined pragma `Unevaluated_Use_Of_Old`.

## 4.47 Attribute Passed\_By\_Reference

`typ'Passed_By_Reference` for any subtype `typ` returns a value of type `Boolean` value that is `True` if the type is normally passed by reference and `False` if the type is normally passed by copy in calls. For scalar types, the result is always `False` and is static. For non-scalar types, the result is nonstatic.

## 4.48 Attribute Pool\_Address

`X'Pool_Address` for any object `X` returns the address of `X` within its storage pool. This is the same as `X'Address`, except that for an unconstrained array whose bounds are allocated just before the first component, `X'Pool_Address` returns the address of those bounds, whereas `X'Address` returns the address of the first component.

Here, we are interpreting ‘storage pool’ broadly to mean **wherever the object is allocated**, which could be a user-defined storage pool, the global heap, on the stack, or in a static memory area. For an object created by `new`, `Ptr.all'Pool_Address` is what is passed to `Allocate` and returned from `Deallocate`.

## 4.49 Attribute Range\_Length

`typ'Range_Length` for any discrete type *typ* yields the number of values represented by the subtype (zero for a null range). The result is static for static subtypes. `Range_Length` applied to the index subtype of a one dimensional array always gives the same result as `Length` applied to the array itself.

## 4.50 Attribute Restriction\_Set

This attribute allows compile time testing of restrictions that are currently in effect. It is primarily intended for specializing code in the run-time based on restrictions that are active (e.g. don't need to save fpt registers if restriction `No_Floating_Point` is known to be in effect), but can be used anywhere.

There are two forms:

```
System'Restriction_Set (partition_boolean_restriction_NAME)
System'Restriction_Set (No_Dependence => library_unit_NAME);
```

In the case of the first form, the only restriction names allowed are parameterless restrictions that are checked for consistency at bind time. For a complete list see the subtype `System.Rident.Partition_Boolean_Restrictions`.

The result returned is `True` if the restriction is known to be in effect, and `False` if the restriction is known not to be in effect. An important guarantee is that the value of a `Restriction_Set` attribute is known to be consistent throughout all the code of a partition.

This is trivially achieved if the entire partition is compiled with a consistent set of restriction pragmas. However, the compilation model does not require this. It is possible to compile one set of units with one set of pragmas, and another set of units with another set of pragmas. It is even possible to compile a spec with one set of pragmas, and then `WITH` the same spec with a different set of pragmas. Inconsistencies in the actual use of the restriction are checked at bind time.

In order to achieve the guarantee of consistency for the `Restriction_Set` pragma, we consider that a use of the pragma that yields `False` is equivalent to a violation of the restriction.

So for example if you write

```
if System'Restriction_Set (No_Floating_Point) then
  ...
else
  ...
end if;
```

And the result is `False`, so that the `else` branch is executed, you can assume that this restriction is not set for any unit in the partition. This is checked by considering this use of the restriction pragma to be a violation of the restriction `No_Floating_Point`. This means that no other unit can attempt to set this restriction (if some unit does attempt to set it, the binder will refuse to bind the partition).

Technical note: The restriction name and the unit name are interpreted entirely syntactically, as in the corresponding `Restrictions` pragma, they are not analyzed semantically, so they do not have a type.

### 4.51 Attribute Result

`function'Result` can only be used with in a Postcondition pragma for a function. The prefix must be the name of the corresponding function. This is used to refer to the result of the function in the postcondition expression. For a further discussion of the use of this attribute and examples of its use, see the description of pragma Postcondition.

### 4.52 Attribute Round

In addition to the usage of this attribute in the Ada RM, GNAT also permits the use of the `'Round` attribute for ordinary fixed point types.

### 4.53 Attribute Safe\_Emax

The `Safe_Emax` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.54 Attribute Safe\_Large

The `Safe_Large` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.55 Attribute Safe\_Small

The `Safe_Small` attribute is provided for compatibility with Ada 83. See the Ada 83 reference manual for an exact description of the semantics of this attribute.

### 4.56 Attribute Scalar\_Storage\_Order

For every array or record type `S`, the representation attribute `Scalar_Storage_Order` denotes the order in which storage elements that make up scalar components are ordered within `S`. The value given must be a static expression of type `System.Bit_Order`. The following is an example of the use of this feature:

```
-- Component type definitions

subtype Yr_Type is Natural range 0 .. 127;
subtype Mo_Type is Natural range 1 .. 12;
subtype Da_Type is Natural range 1 .. 31;

-- Record declaration

type Date is record
  Years_Since_1980 : Yr_Type;
  Month            : Mo_Type;
  Day_Of_Month     : Da_Type;
end record;

-- Record representation clause
```

```

for Date use record
  Years_Since_1980 at 0 range 0 .. 6;
  Month           at 0 range 7 .. 10;
  Day_Of_Month    at 0 range 11 .. 15;
end record;

-- Attribute definition clauses

for Date'Bit_Order use System.High_Order_First;
for Date'Scalar_Storage_Order use System.High_Order_First;
-- If Scalar_Storage_Order is specified, it must be consistent with
-- Bit_Order, so it's best to always define the latter explicitly if
-- the former is used.

```

Other properties are as for the standard representation attribute `Bit_Order` defined by Ada RM 13.5.3(4). The default is `System.Default_Bit_Order`.

For a record type `T`, if `T'Scalar_Storage_Order` is specified explicitly, it shall be equal to `T'Bit_Order`. Note: this means that if a `Scalar_Storage_Order` attribute definition clause is not confirming, then the type's `Bit_Order` shall be specified explicitly and set to the same value.

Derived types inherit an explicitly set scalar storage order from their parent types. This may be overridden for the derived type by giving an explicit scalar storage order for it. However, for a record extension, the derived type must have the same scalar storage order as the parent type.

A component of a record type that is itself a record or an array and that does not start and end on a byte boundary must have the same scalar storage order as the record type. A component of a bit-packed array type that is itself a record or an array must have the same scalar storage order as the array type.

No component of a type that has an explicit `Scalar_Storage_Order` attribute definition may be aliased.

A confirming `Scalar_Storage_Order` attribute definition clause (i.e. with a value equal to `System.Default_Bit_Order`) has no effect.

If the opposite storage order is specified, then whenever the value of a scalar component of an object of type `S` is read, the storage elements of the enclosing machine scalar are first reversed (before retrieving the component value, possibly applying some shift and mask operations on the enclosing machine scalar), and the opposite operation is done for writes.

In that case, the restrictions set forth in 13.5.1(10.3/2) for scalar components are relaxed. Instead, the following rules apply:

- \* the underlying storage elements are those at positions `(position + first_bit / storage_element_size) .. (position + (last_bit + storage_element_size - 1) / storage_element_size)`
- \* the sequence of underlying storage elements shall have a size no greater than the largest machine scalar
- \* the enclosing machine scalar is defined as the smallest machine scalar starting at a position no greater than `position + first_bit / storage_element_size` and covering

storage elements at least up to `position + (last_bit + storage_element_size - 1) / storage_element_size`

\* the position of the component is interpreted relative to that machine scalar.

If no scalar storage order is specified for a type (either directly, or by inheritance in the case of a derived type), then the default is normally the native ordering of the target, but this default can be overridden using pragma `Default_Scalar_Storage_Order`.

If a component of `T` is itself of a record or array type, the specified `Scalar_Storage_Order` does ‘not’ apply to that nested type: an explicit attribute definition clause must be provided for the component type as well if desired.

Representation changes that explicitly or implicitly toggle the scalar storage order are not supported and may result in erroneous execution of the program, except when performed by means of an instance of `Ada.Unchecked_Conversion`.

In particular, overlays are not supported and a warning is given for them:

```

type Rec_LE is record
  I : Integer;
end record;

for Rec_LE use record
  I at 0 range 0 .. 31;
end record;

for Rec_LE'Bit_Order use System.Low_Order_First;
for Rec_LE'Scalar_Storage_Order use System.Low_Order_First;

type Rec_BE is record
  I : Integer;
end record;

for Rec_BE use record
  I at 0 range 0 .. 31;
end record;

for Rec_BE'Bit_Order use System.High_Order_First;
for Rec_BE'Scalar_Storage_Order use System.High_Order_First;

R_LE : Rec_LE;

R_BE : Rec_BE;
for R_BE'Address use R_LE'Address;

```

warning: overlay changes scalar storage order [enabled by default]

In most cases, such representation changes ought to be replaced by an instantiation of a function or procedure provided by `GNAT.Byte_Swapping`.

Note that the scalar storage order only affects the in-memory data representation. It has no effect on the representation used by stream attributes.

Note that debuggers may be unable to display the correct value of scalar components of a type for which the opposite storage order is specified.

## 4.57 Attribute `Simple_Storage_Pool`

For every nonformal, nonderived access-to-object type `Acc`, the representation attribute `Simple_Storage_Pool` may be specified via an `attribute_definition_clause` (or by specifying the equivalent aspect):

```
My_Pool : My_Simple_Storage_Pool_Type;

type Acc is access My_Data_Type;

for Acc'Simple_Storage_Pool use My_Pool;
```

The name given in an `attribute_definition_clause` for the `Simple_Storage_Pool` attribute shall denote a variable of a ‘simple storage pool type’ (see pragma `Simple_Storage_Pool_Type`).

The use of this attribute is only allowed for a prefix denoting a type for which it has been specified. The type of the attribute is the type of the variable specified as the simple storage pool of the access type, and the attribute denotes that variable.

It is illegal to specify both `Storage_Pool` and `Simple_Storage_Pool` for the same access type.

If the `Simple_Storage_Pool` attribute has been specified for an access type, then applying the `Storage_Pool` attribute to the type is flagged with a warning and its evaluation raises the exception `Program_Error`.

If the `Simple_Storage_Pool` attribute has been specified for an access type `S`, then the evaluation of the attribute `S'Storage_Size` returns the result of calling `Storage_Size` (`S'Simple_Storage_Pool`), which is intended to indicate the number of storage elements reserved for the simple storage pool. If the `Storage_Size` function has not been defined for the simple storage pool type, then this attribute returns zero.

If an access type `S` has a specified simple storage pool of type `SSP`, then the evaluation of an allocator for that access type calls the primitive `Allocate` procedure for type `SSP`, passing `S'Simple_Storage_Pool` as the pool parameter. The detailed semantics of such allocators is the same as those defined for allocators in section 13.11 of the *Ada Reference Manual*, with the term ‘simple storage pool’ substituted for ‘storage pool’.

If an access type `S` has a specified simple storage pool of type `SSP`, then a call to an instance of the `Ada.Unchecked_Deallocation` for that access type invokes the primitive `Deallocate` procedure for type `SSP`, passing `S'Simple_Storage_Pool` as the pool parameter. The detailed semantics of such unchecked deallocations is the same as defined in section 13.11.2 of the *Ada Reference Manual*, except that the term ‘simple storage pool’ is substituted for ‘storage pool’.

## 4.58 Attribute `Small`

The `Small` attribute is defined in Ada 95 (and Ada 2005) only for fixed-point types. GNAT also allows this attribute to be applied to floating-point types for compatibility with Ada 83.

See the Ada 83 reference manual for an exact description of the semantics of this attribute when applied to floating-point types.

### 4.59 Attribute `Small_Denominator`

`typ'Small_Denominator` for any fixed-point subtype *typ* yields the denominator in the representation of `typ'Small` as a rational number with coprime factors (i.e. as an irreducible fraction).

### 4.60 Attribute `Small_Numerator`

`typ'Small_Numerator` for any fixed-point subtype *typ* yields the numerator in the representation of `typ'Small` as a rational number with coprime factors (i.e. as an irreducible fraction).

### 4.61 Attribute `Storage_Unit`

`Standard'Storage_Unit` (`Standard` is the only allowed prefix) provides the same value as `System.Storage_Unit`.

### 4.62 Attribute `Stub_Type`

The GNAT implementation of remote access-to-classwide types is organized as described in AARM section E.4 (20.t): a value of an RACW type (designating a remote object) is represented as a normal access value, pointing to a “stub” object which in turn contains the necessary information to contact the designated remote object. A call on any dispatching operation of such a stub object does the remote call, if necessary, using the information in the stub object to locate the target partition, etc.

For a prefix *T* that denotes a remote access-to-classwide type, `T'Stub_Type` denotes the type of the corresponding stub objects.

By construction, the layout of `T'Stub_Type` is identical to that of type `RACW_Stub_Type` declared in the internal implementation-defined unit `System.Partition_Interface`. Use of this attribute will create an implicit dependency on this unit.

### 4.63 Attribute `System_Allocator_Alignment`

`Standard'System_Allocator_Alignment` (`Standard` is the only allowed prefix) provides the observable guaranteed to be honored by the system allocator (malloc). This is a static value that can be used in user storage pools based on malloc either to reject allocation with alignment too large or to enable a realignment circuitry if the alignment request is larger than this value.

### 4.64 Attribute `Target_Name`

`Standard'Target_Name` (`Standard` is the only allowed prefix) provides a static string value that identifies the target for the current compilation. For GCC implementations, this is the standard gcc target name without the terminating slash (for example, GNAT 5.0 on windows yields “i586-pc-mingw32msv”).

### 4.65 Attribute To\_Address

The `System'To_Address` (`System` is the only allowed prefix) denotes a function identical to `System.Storage_Elements.To_Address` except that it is a static attribute. This means that if its argument is a static expression, then the result of the attribute is a static expression. This means that such an expression can be used in contexts (e.g., preelaborable packages) which require a static expression and where the function call could not be used (since the function call is always nonstatic, even if its argument is static). The argument must be in the range  $-(2^{m-1}) \dots 2^{m-1}$ , where  $m$  is the memory size (typically 32 or 64). Negative values are interpreted in a modular manner (e.g., -1 means the same as `16#FFFF_FFFF#` on a 32 bits machine).

### 4.66 Attribute To\_Any

This internal attribute is used for the generation of remote subprogram stubs in the context of the Distributed Systems Annex.

### 4.67 Attribute Type\_Class

`typ'Type_Class` for any type or subtype *typ* yields the value of the type class for the full type of *typ*. If *typ* is a generic formal type, the value is the value for the corresponding actual subtype. The value of this attribute is of type `System.Aux_DEC.Type_Class`, which has the following definition:

```
type Type_Class is
  (Type_Class_Enumeration,
   Type_Class_Integer,
   Type_Class_Fixed_Point,
   Type_Class_Floating_Point,
   Type_Class_Array,
   Type_Class_Record,
   Type_Class_Access,
   Type_Class_Task,
   Type_Class_Address);
```

Protected types yield the value `Type_Class_Task`, which thus applies to all concurrent types. This attribute is designed to be compatible with the DEC Ada 83 attribute of the same name.

### 4.68 Attribute Type\_Key

The `Type_Key` attribute is applicable to a type or subtype and yields a value of type `Standard.String` containing encoded information about the type or subtype. This provides improved compatibility with other implementations that support this attribute.

### 4.69 Attribute TypeCode

This internal attribute is used for the generation of remote subprogram stubs in the context of the Distributed Systems Annex.

## 4.70 Attribute Unconstrained\_Array

The `Unconstrained_Array` attribute can be used with a prefix that denotes any type or subtype. It is a static attribute that yields `True` if the prefix designates an unconstrained array, and `False` otherwise. In a generic instance, the result is still static, and yields the result of applying this test to the generic actual.

## 4.71 Attribute Universal\_Literal\_String

The prefix of `Universal_Literal_String` must be a named number. The static result is the string consisting of the characters of the number as defined in the original source. This allows the user program to access the actual text of named numbers without intermediate conversions and without the need to enclose the strings in quotes (which would preclude their use as numbers).

For example, the following program prints the first 50 digits of pi:

```
with Text_IO; use Text_IO;
with Ada.Numerics;
procedure Pi is
begin
  Put (Ada.Numerics.Pi'Universal_Literal_String);
end;
```

## 4.72 Attribute Unrestricted\_Access

The `Unrestricted_Access` attribute is similar to `Access` except that all accessibility and aliased view checks are omitted. This is a user-beware attribute.

For objects, it is similar to `Address`, for which it is a desirable replacement where the value desired is an access type. In other words, its effect is similar to first applying the `Address` attribute and then doing an unchecked conversion to a desired access type.

For subprograms, `P'Unrestricted_Access` may be used where `P'Access` would be illegal, to construct a value of a less-nested named access type that designates a more-nested subprogram. This value may be used in indirect calls, so long as the more-nested subprogram still exists; once the subprogram containing it has returned, such calls are erroneous. For example:

```
package body P is

  type Less_Nested is access procedure;
  Global : Less_Nested;

  procedure P1 is
  begin
    Global.all;
  end P1;

  procedure P2 is
    Local_Var : Integer;
```

```

        procedure More_Nested is
        begin
            ... Local_Var ...
        end More_Nested;
    begin
        Global := More_Nested'Unrestricted_Access;
        P1;
    end P2;

end P;

```

When P1 is called from P2, the call via Global is OK, but if P1 were called after P2 returns, it would be an erroneous use of a dangling pointer.

For objects, it is possible to use `Unrestricted_Access` for any type. However, if the result is of an access-to-unconstrained array subtype, then the resulting pointer has the same scope as the context of the attribute, and must not be returned to some enclosing scope. For instance, if a function uses `Unrestricted_Access` to create an access-to-unconstrained-array and returns that value to the caller, the result will involve dangling pointers. In addition, it is only valid to create pointers to unconstrained arrays using this attribute if the pointer has the normal default ‘fat’ representation where a pointer has two components, one points to the array and one points to the bounds. If a size clause is used to force ‘thin’ representation for a pointer to unconstrained where there is only space for a single pointer, then the resulting pointer is not usable.

In the simple case where a direct use of `Unrestricted_Access` attempts to make a thin pointer for a non-aliased object, the compiler will reject the use as illegal, as shown in the following example:

```

with System; use System;
procedure SliceUA2 is
    type A is access all String;
    for A'Size use Standard'Address_Size;

    procedure P (Arg : A) is
    begin
        null;
    end P;

    X : String := "hello world!";
    X2 : aliased String := "hello world!";

    AV : A := X'Unrestricted_Access;    -- ERROR
    |
>>> illegal use of Unrestricted_Access attribute
>>> attempt to generate thin pointer to unaliased object

begin
    P (X'Unrestricted_Access);          -- ERROR
    |

```

```

>>> illegal use of Unrestricted_Access attribute
>>> attempt to generate thin pointer to unaliased object

      P (X(7 .. 12)'Unrestricted_Access); -- ERROR
      |
>>> illegal use of Unrestricted_Access attribute
>>> attempt to generate thin pointer to unaliased object

      P (X2'Unrestricted_Access);          -- OK
end;

```

but other cases cannot be detected by the compiler, and are considered to be erroneous. Consider the following example:

```

with System; use System;
with System; use System;
procedure SliceUA is
  type AF is access all String;

  type A is access all String;
  for A'Size use Standard'Address_Size;

  procedure P (Arg : A) is
  begin
    if Arg'Length /= 6 then
      raise Program_Error;
    end if;
  end P;

  X : String := "hello world!";
  Y : AF := X (7 .. 12)'Unrestricted_Access;

begin
  P (A (Y));
end;

```

A normal unconstrained array value or a constrained array object marked as aliased has the bounds in memory just before the array, so a thin pointer can retrieve both the data and the bounds. But in this case, the non-aliased object **X** does not have the bounds before the string. If the size clause for type **A** were not present, then the pointer would be a fat pointer, where one component is a pointer to the bounds, and all would be well. But with the size clause present, the conversion from fat pointer to thin pointer in the call loses the bounds, and so this is erroneous, and the program likely raises a **Program\_Error** exception.

In general, it is advisable to completely avoid mixing the use of thin pointers and the use of **Unrestricted\_Access** where the designated type is an unconstrained array. The use of thin pointers should be restricted to cases of porting legacy code that implicitly assumes the size of pointers, and such code should not in any case be using this attribute.

Another erroneous situation arises if the attribute is applied to a constant. The resulting pointer can be used to access the constant, but the effect of trying to modify a constant in this manner is not well-defined. Consider this example:

```
P : constant Integer := 4;
type R is access all Integer;
RV : R := P'Unrestricted_Access;
..
RV.all := 3;
```

Here we attempt to modify the constant P from 4 to 3, but the compiler may or may not notice this attempt, and subsequent references to P may yield either the value 3 or the value 4 or the assignment may blow up if the compiler decides to put P in read-only memory. One particular case where `Unrestricted_Access` can be used in this way is to modify the value of an `in` parameter:

```
procedure K (S : in String) is
  type R is access all Character;
  RV : R := S (3)'Unrestricted_Access;
begin
  RV.all := 'a';
end;
```

In general this is a risky approach. It may appear to “work” but such uses of `Unrestricted_Access` are potentially non-portable, even from one version of GNAT to another, so are best avoided if possible.

## 4.73 Attribute Update

The `Update` attribute creates a copy of an array or record value with one or more modified components. The syntax is:

```
PREFIX'Update ( RECORD_COMPONENT_ASSOCIATION_LIST )
PREFIX'Update ( ARRAY_COMPONENT_ASSOCIATION {, ARRAY_COMPONENT_ASSOCIATION } )
PREFIX'Update ( MULTIDIMENSIONAL_ARRAY_COMPONENT_ASSOCIATION
                {, MULTIDIMENSIONAL_ARRAY_COMPONENT_ASSOCIATION } )

MULTIDIMENSIONAL_ARRAY_COMPONENT_ASSOCIATION ::= INDEX_EXPRESSION_LIST_LIST => EXPRESSION
INDEX_EXPRESSION_LIST_LIST                    ::= INDEX_EXPRESSION_LIST { | INDEX_EXPRESSION_LIST
INDEX_EXPRESSION_LIST                         ::= ( EXPRESSION {, EXPRESSION } )
```

where `PREFIX` is the name of an array or record object, the association list in parentheses does not contain an `others` choice and the box symbol `<>` may not appear in any expression. The effect is to yield a copy of the array or record value which is unchanged apart from the components mentioned in the association list, which are changed to the indicated value. The original value of the array or record value is not affected. For example:

```
type Arr is Array (1 .. 5) of Integer;
...
Avar1 : Arr := (1,2,3,4,5);
Avar2 : Arr := Avar1'Update (2 => 10, 3 .. 4 => 20);
```

yields a value for `Avar2` of 1,10,20,20,5 with `Avar1` begin unmodified. Similarly:

```
type Rec is A, B, C : Integer;
```

```

...
Rvar1 : Rec := (A => 1, B => 2, C => 3);
Rvar2 : Rec := Rvar1'Update (B => 20);

```

yields a value for `Rvar2` of `(A => 1, B => 20, C => 3)`, with `Rvar1` being unmodified. Note that the value of the attribute reference is computed completely before it is used. This means that if you write:

```
Avar1 := Avar1'Update (1 => 10, 2 => Function_Call);
```

then the value of `Avar1` is not modified if `Function_Call` raises an exception, unlike the effect of a series of direct assignments to elements of `Avar1`. In general this requires that two extra complete copies of the object are required, which should be kept in mind when considering efficiency.

The `Update` attribute cannot be applied to prefixes of a limited type, and cannot reference discriminants in the case of a record type. The accessibility level of an `Update` attribute result object is defined as for an aggregate.

In the record case, no component can be mentioned more than once. In the array case, two overlapping ranges can appear in the association list, in which case the modifications are processed left to right.

Multi-dimensional arrays can be modified, as shown by this example:

```

A : array (1 .. 10, 1 .. 10) of Integer;
..
A := A'Update ((1, 2) => 20, (3, 4) => 30);

```

which changes element (1,2) to 20 and (3,4) to 30.

## 4.74 Attribute Valid\_Value

The `'Valid_Value` attribute is defined for enumeration types other than those in package `Standard` or types derived from those types. This attribute is a function that takes a `String`, and returns `Boolean`. `T'Valid_Value (S)` returns `True` if and only if `T'Value (S)` would not raise `Constraint_Error`.

## 4.75 Attribute Valid\_Scalars

The `'Valid_Scalars` attribute is intended to make it easier to check the validity of scalar subcomponents of composite objects. The attribute is defined for any prefix `P` which denotes an object. Prefix `P` can be any type except for tagged private or `Unchecked_Union` types. The value of the attribute is of type `Boolean`.

`P'Valid_Scalars` yields `True` if and only if the evaluation of `C'Valid` yields `True` for every scalar subcomponent `C` of `P`, or if `P` has no scalar subcomponents. Attribute `'Valid_Scalars` is equivalent to attribute `'Valid` for scalar types.

It is not specified in what order the subcomponents are checked, nor whether any more are checked after any one of them is determined to be invalid. If the prefix `P` is of a class-wide type `T'Class` (where `T` is the associated specific type), or if the prefix `P` is of a specific tagged type `T`, then only the subcomponents of `T` are checked; in other words, components of extensions of `T` are not checked even if `T'Class (P) 'Tag /= T'Tag`.

The compiler will issue a warning if it can be determined at compile time that the prefix of the attribute has no scalar subcomponents.

Note: `ValidScalars` can generate a lot of code, especially in the case of a large variant record. If the attribute is called in many places in the same program applied to objects of the same type, it can reduce program size to write a function with a single use of the attribute, and then call that function from multiple places.

## 4.76 Attribute `VADS_Size`

The `'VADS_Size` attribute is intended to make it easier to port legacy code which relies on the semantics of `'Size` as implemented by the VADS Ada 83 compiler. GNAT makes a best effort at duplicating the same semantic interpretation. In particular, `'VADS_Size` applied to a predefined or other primitive type with no `Size` clause yields the `Object_Size` (for example, `Natural'Size` is 32 rather than 31 on typical machines). In addition `'VADS_Size` applied to an object gives the result that would be obtained by applying the attribute to the corresponding type.

## 4.77 Attribute `Value_Size`

`type'Value_Size` is the number of bits required to represent a value of the given subtype. It is the same as `type'Size`, but, unlike `Size`, may be set for non-first subtypes.

## 4.78 Attribute `Wchar_T_Size`

`Standard'Wchar_T_Size` (`Standard` is the only allowed prefix) provides the size in bits of the C `wchar_t` type primarily for constructing the definition of this type in package `Interfaces.C`. The result is a static constant.

## 4.79 Attribute `Word_Size`

`Standard'Word_Size` (`Standard` is the only allowed prefix) provides the value `System.Word_Size`. The result is a static constant.

## 5 Standard and Implementation Defined Restrictions

All Ada Reference Manual-defined Restriction identifiers are implemented:

- \* language-defined restrictions (see 13.12.1)
- \* tasking restrictions (see D.7)
- \* high integrity restrictions (see H.4)

GNAT implements additional restriction identifiers. All restrictions, whether language defined or GNAT-specific, are listed in the following.

### 5.1 Partition-Wide Restrictions

There are two separate lists of restriction identifiers. The first set requires consistency throughout a partition (in other words, if the restriction identifier is used for any compilation unit in the partition, then all compilation units in the partition must obey the restriction).

#### 5.1.1 Immediate\_Reclamation

[RM H.4] This restriction ensures that, except for storage occupied by objects created by allocators and not deallocated via unchecked deallocation, any storage reserved at run time for an object is immediately reclaimed when the object no longer exists.

#### 5.1.2 Max\_Asynchronous\_Select\_Nesting

[RM D.7] Specifies the maximum dynamic nesting level of asynchronous selects. Violations of this restriction with a value of zero are detected at compile time. Violations of this restriction with values other than zero cause `Storage_Error` to be raised.

#### 5.1.3 Max\_Entry\_Queue\_Length

[RM D.7] This restriction is a declaration that any protected entry compiled in the scope of the restriction has at most the specified number of tasks waiting on the entry at any one time, and so no queue is required. Note that this restriction is checked at run time. Violation of this restriction results in the raising of `Program_Error` exception at the point of the call.

The restriction `Max_Entry_Queue_Depth` is recognized as a synonym for `Max_Entry_Queue_Length`. This is retained for historical compatibility purposes (and a warning will be generated for its use if warnings on obsolescent features are activated).

#### 5.1.4 Max\_Protected\_Entries

[RM D.7] Specifies the maximum number of entries per protected type. The bounds of every entry family of a protected unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static.

#### 5.1.5 Max\_Select\_Alternatives

[RM D.7] Specifies the maximum number of alternatives in a selective accept.

### 5.1.6 Max\_Storage\_At\_Blocking

[RM D.7] Specifies the maximum portion (in storage elements) of a task's `Storage_Size` that can be retained by a blocked task. A violation of this restriction causes `Storage_Error` to be raised.

### 5.1.7 Max\_Task\_Entries

[RM D.7] Specifies the maximum number of entries per task. The bounds of every entry family of a task unit shall be static, or shall be defined by a discriminant of a subtype whose corresponding bound is static.

### 5.1.8 Max\_Tasks

[RM D.7] Specifies the maximum number of task that may be created, not counting the creation of the environment task. Violations of this restriction with a value of zero are detected at compile time. Violations of this restriction with values other than zero cause `Storage_Error` to be raised.

### 5.1.9 No\_Abort\_Statements

[RM D.7] There are no `abort_statements`, and there are no calls to `Task_Identification.Abort_Task`.

### 5.1.10 No\_Access\_Parameter\_Allocators

[RM H.4] This restriction ensures at compile time that there are no occurrences of an allocator as the actual parameter to an access parameter.

### 5.1.11 No\_Access\_Subprograms

[RM H.4] This restriction ensures at compile time that there are no declarations of access-to-subprogram types.

### 5.1.12 No\_Allocators

[RM H.4] This restriction ensures at compile time that there are no occurrences of an allocator.

### 5.1.13 No\_Anonymous\_Allocators

[RM H.4] This restriction ensures at compile time that there are no occurrences of an allocator of anonymous access type.

### 5.1.14 No\_Asynchronous\_Control

[RM J.13] This restriction ensures at compile time that there are no semantic dependences on the predefined package `Asynchronous_Task_Control`.

### 5.1.15 No\_Calendar

[GNAT] This restriction ensures at compile time that there are no semantic dependences on package `Calendar`.

### 5.1.16 No\_Coextensions

[RM H.4] This restriction ensures at compile time that there are no coextensions. See 3.10.2.

### 5.1.17 No\_Default\_Initialization

[GNAT] This restriction prohibits any instance of default initialization of variables or components. The binder implements a consistency check that prevents any unit without the restriction from with'ing a unit with the restriction (this allows the generation of initialization procedures to be skipped, since you can be sure that no call is ever generated to an initialization procedure in a unit with the restriction active). If used in conjunction with `Initialize_Scalars` or `Normalize_Scalars`, the effect is to prohibit all cases of variables declared without a specific initializer (including the case of OUT scalar parameters).

### 5.1.18 No\_Delay

[RM H.4] This restriction ensures at compile time that there are no delay statements and no semantic dependences on package `Calendar`.

### 5.1.19 No\_Dependence

[RM 13.12.1] This restriction ensures at compile time that there are no dependences on a library unit. For GNAT, this includes implicit implementation dependences on units of the runtime library that are created by the compiler to support specific constructs of the language. Here are some examples:

- \* `System.Arith_64`: 64-bit arithmetics for 32-bit platforms,
- \* `System.Arith_128`: 128-bit arithmetics for 64-bit platforms,
- \* `System.Memory`: heap memory allocation routines,
- \* `System.Memory_Compare`: memory comparison routine (aka `memcmp` for C),
- \* `System.Memory_Copy`: memory copy routine (aka `memcpy` for C),
- \* `System.Memory_Move`: memory move routine (aka `memmove` for C),
- \* `System.Memory_Set`: memory set routine (aka `memset` for C),
- \* `System.Stack_Checking[.Operations]`: stack checking without MMU,
- \* `System.GCC`: support routines from the GCC library.

### 5.1.20 No\_Direct\_Boolean\_Operators

[GNAT] This restriction ensures that no logical operators (and/or/xor) are used on operands of type `Boolean` (or any type derived from `Boolean`). This is intended for use in safety critical programs where the certification protocol requires the use of short-circuit (and then, or else) forms for all composite boolean operations.

### 5.1.21 No\_Dispatch

[RM H.4] This restriction ensures at compile time that there are no occurrences of `T'Class`, for any (tagged) subtype `T`.

### 5.1.22 No\_Dispatching\_Calls

[GNAT] This restriction ensures at compile time that the code generated by the compiler involves no dispatching calls. The use of this restriction allows the safe use of record extensions, classwide membership tests and other classwide features not involving implicit dispatching. This restriction ensures that the code contains no indirect calls through a dispatching mechanism. Note that this includes internally-generated calls created by the

compiler, for example in the implementation of class-wide objects assignments. The membership test is allowed in the presence of this restriction, because its implementation requires no dispatching. This restriction is comparable to the official Ada restriction `No_Dispatch` except that it is a bit less restrictive in that it allows all classwide constructs that do not imply dispatching. The following example indicates constructs that violate this restriction.

```

package Pkg is
  type T is tagged record
    Data : Natural;
  end record;
  procedure P (X : T);

  type DT is new T with record
    More_Data : Natural;
  end record;
  procedure Q (X : DT);
end Pkg;

with Pkg; use Pkg;
procedure Example is
  procedure Test (O : T'Class) is
    N : Natural := O'Size; -- Error: Dispatching call
    C : T'Class := O;      -- Error: implicit Dispatching Call
  begin
    if O in DT'Class then -- OK    : Membership test
      Q (DT (O));         -- OK    : Type conversion plus direct call
    else
      P (O);              -- Error: Dispatching call
    end if;
  end Test;

  Obj : DT;
begin
  P (Obj);                -- OK    : Direct call
  P (T (Obj));            -- OK    : Type conversion plus direct call
  P (T'Class (Obj));      -- Error: Dispatching call

  Test (Obj);             -- OK    : Type conversion

  if Obj in T'Class then  -- OK    : Membership test
    null;
  end if;
end Example;

```

### 5.1.23 No\_Dynamic\_Attachment

[RM D.7] This restriction ensures that there is no call to any of the operations defined in package `Ada.Interrupts` (`Is_Reserved`, `Is_Attached`, `Current_Handler`, `Attach_Handler`, `Exchange_Handler`, `Detach_Handler`, and `Reference`).

The restriction `No_Dynamic_Interrupts` is recognized as a synonym for `No_Dynamic_Attachment`. This is retained for historical compatibility purposes (and a warning will be generated for its use if warnings on obsolescent features are activated).

### 5.1.24 No\_Dynamic\_Priorities

[RM D.7] There are no semantic dependencies on the package `Dynamic_Priorities`.

### 5.1.25 No\_Entry\_Calls\_In\_Elaboration\_Code

[GNAT] This restriction ensures at compile time that no task or protected entry calls are made during elaboration code. As a result of the use of this restriction, the compiler can assume that no code past an `accept` statement in a task can be executed at elaboration time.

### 5.1.26 No\_Enumeration\_Maps

[GNAT] This restriction ensures at compile time that no operations requiring enumeration maps are used (that is `Image` and `Value` attributes applied to enumeration types).

### 5.1.27 No\_Exception\_Handlers

[GNAT] This restriction ensures at compile time that there are no explicit exception handlers. It also indicates that no exception propagation will be provided. In this mode, exceptions may be raised but will result in an immediate call to the last chance handler, a routine that the user must define with the following profile:

```
procedure Last_Chance_Handler
  (Source_Location : System.Address; Line : Integer);
pragma Export (C, Last_Chance_Handler,
               "__gnat_last_chance_handler");
```

The `Source_Location` parameter is a C null-terminated string representing a message to be associated with the exception (typically the source location of the `raise` statement generated by the compiler). The `Line` parameter when nonzero represents the line number in the source program where the `raise` occurs.

### 5.1.28 No\_Exception\_Propagation

[GNAT] This restriction guarantees that exceptions are never propagated to an outer subprogram scope. The only case in which an exception may be raised is when the handler is statically in the same subprogram, so that the effect of a `raise` is essentially like a `goto` statement. Any other `raise` statement (implicit or explicit) will be considered unhandled. Exception handlers are allowed, but may not contain an exception occurrence identifier (exception choice). In addition, use of the package `GNAT.Current_Exception` is not permitted, and `raise` statements (`raise` with no operand) are not permitted.

### 5.1.29 No\_Exception\_Registration

[GNAT] This restriction ensures at compile time that no stream operations for types `Exception_Id` or `Exception_Occurrence` are used. This also makes it impossible to pass exceptions to or from a partition with this restriction in a distributed environment. If this restriction is active, the generated code is simplified by omitting the otherwise-required global registration of exceptions when they are declared.

### 5.1.30 No\_Exceptions

[RM H.4] This restriction ensures at compile time that there are no raise statements and no exception handlers and also suppresses the generation of language-defined run-time checks.

### 5.1.31 No\_Finalization

[GNAT] This restriction disables the language features described in chapter 7.6 of the Ada 2005 RM as well as all form of code generation performed by the compiler to support these features. The following types are no longer considered controlled when this restriction is in effect:

- \* `Ada.Finalization.Controlled`
- \* `Ada.Finalization.Limited_Controlled`
- \* Derivations from `Controlled` or `Limited_Controlled`
- \* Class-wide types
- \* Protected types
- \* Task types
- \* Array and record types with controlled components

The compiler no longer generates code to initialize, finalize or adjust an object or a nested component, either declared on the stack or on the heap. The deallocation of a controlled object no longer finalizes its contents.

### 5.1.32 No\_Fixed\_Point

[RM H.4] This restriction ensures at compile time that there are no occurrences of fixed point types and operations.

### 5.1.33 No\_Floating\_Point

[RM H.4] This restriction ensures at compile time that there are no occurrences of floating point types and operations.

### 5.1.34 No\_Implicit\_Conditionals

[GNAT] This restriction ensures that the generated code does not contain any implicit conditionals, either by modifying the generated code where possible, or by rejecting any construct that would otherwise generate an implicit conditional. Note that this check does not include run time constraint checks, which on some targets may generate implicit conditionals as well. To control the latter, constraint checks can be suppressed in the normal manner. Constructs generating implicit conditionals include comparisons of composite objects and the `Max/Min` attributes.

### 5.1.35 No\_Implicit\_Dynamic\_Code

[GNAT] This restriction prevents the compiler from building ‘trampolines’. This is a structure that is built on the stack and contains dynamic code to be executed at run time. On some targets, a trampoline is built for the following features: **Access**, **Unrestricted\_Access**, or **Address** of a nested subprogram; nested task bodies; primitive operations of nested tagged types. Trampolines do not work on machines that prevent execution of stack data. For example, on windows systems, enabling DEP (data execution protection) will cause trampolines to raise an exception. Trampolines are also quite slow at run time.

On many targets, trampolines have been largely eliminated. Look at the version of `system.ads` for your target — if it has `Always-Compatible_Rep` equal to `False`, then trampolines are largely eliminated. In particular, a trampoline is built for the following features: **Address** of a nested subprogram; **Access** or **Unrestricted\_Access** of a nested subprogram, but only if `pragma Favor_Top_Level` applies, or the access type has a foreign-language convention; primitive operations of nested tagged types.

### 5.1.36 No\_Implicit\_Heap\_Allocations

[RM D.7] No constructs are allowed to cause implicit heap allocation.

### 5.1.37 No\_Implicit\_Protected\_Object\_Allocations

[GNAT] No constructs are allowed to cause implicit heap allocation of a protected object.

### 5.1.38 No\_Implicit\_Task\_Allocations

[GNAT] No constructs are allowed to cause implicit heap allocation of a task.

### 5.1.39 No\_InitializeScalars

[GNAT] This restriction ensures that no unit in the partition is compiled with `pragma InitializeScalars`. This allows the generation of more efficient code, and in particular eliminates dummy null initialization routines that are otherwise generated for some record and array types.

### 5.1.40 No\_IO

[RM H.4] This restriction ensures at compile time that there are no dependences on any of the library units `Sequential_IO`, `Direct_IO`, `Text_IO`, `Wide_Text_IO`, `Wide_Wide_Text_IO`, or `Stream_IO`.

### 5.1.41 No\_Local\_Allocators

[RM H.4] This restriction ensures at compile time that there are no occurrences of an allocator in subprograms, generic subprograms, tasks, and entry bodies.

### 5.1.42 No\_Local\_Protected\_Objects

[RM D.7] This restriction ensures at compile time that protected objects are only declared at the library level.

### 5.1.43 No\_Local\_Tagged\_Types

[GNAT] This restriction ensures at compile time that tagged types are only declared at the library level.

#### 5.1.44 No\_Local\_Timing\_Events

[RM D.7] All objects of type `Ada.Real_Time.Timing_Events.Timing_Event` are declared at the library level.

#### 5.1.45 No\_Long\_Long\_Integers

[GNAT] This partition-wide restriction forbids any explicit reference to type `Standard.Long_Long_Integer`, and also forbids declaring range types whose implicit base type is `Long_Long_Integer`, and modular types whose size exceeds `Long_Integer'Size`.

#### 5.1.46 No\_Multiple\_Elaboration

[GNAT] When this restriction is active and the static elaboration model is used, and `-fpreserve-control-flow` is not used, the compiler is allowed to suppress the elaboration counter normally associated with the unit, even if the unit has elaboration code. This counter is typically used to check for access before elaboration and to control multiple elaboration attempts. If the restriction is used, then the situations in which multiple elaboration is possible, including non-Ada main programs and Stand Alone libraries, are not permitted and will be diagnosed by the binder.

#### 5.1.47 No\_Nested\_Finalization

[RM D.7] All objects requiring finalization are declared at the library level.

#### 5.1.48 No\_Protected\_Type\_Allocators

[RM D.7] This restriction ensures at compile time that there are no allocator expressions that attempt to allocate protected objects.

#### 5.1.49 No\_Protected\_Types

[RM H.4] This restriction ensures at compile time that there are no declarations of protected types or protected objects.

#### 5.1.50 No\_Recursion

[RM H.4] A program execution is erroneous if a subprogram is invoked as part of its execution.

#### 5.1.51 No\_Reentrancy

[RM H.4] A program execution is erroneous if a subprogram is executed by two tasks at the same time.

#### 5.1.52 No\_Relative\_Delay

[RM D.7] This restriction ensures at compile time that there are no delay relative statements and prevents expressions such as `delay 1.23;` from appearing in source code.

#### 5.1.53 No\_Requeue\_Statements

[RM D.7] This restriction ensures at compile time that no requeue statements are permitted and prevents keyword `requeue` from being used in source code.

The restriction `No_Requeue` is recognized as a synonym for `No_Requeue_Statements`. This is retained for historical compatibility purposes (and a warning will be generated for its use if warnings on oNobsolescent features are activated).

#### 5.1.54 `No_Secondary_Stack`

[GNAT] This restriction ensures at compile time that the generated code does not contain any reference to the secondary stack. The secondary stack is used to implement functions returning unconstrained objects (arrays or records) on some targets. Suppresses the allocation of secondary stacks for tasks (excluding the environment task) at run time.

#### 5.1.55 `No_Select_Statements`

[RM D.7] This restriction ensures at compile time no select statements of any kind are permitted, that is the keyword `select` may not appear.

#### 5.1.56 `No_Specific_Termination_Handlers`

[RM D.7] There are no calls to `Ada.Task_Termination.Set_Specific_Handler` or to `Ada.Task_Termination.Specific_Handler`.

#### 5.1.57 `No_Specification_of_Aspect`

[RM 13.12.1] This restriction checks at compile time that no aspect specification, attribute definition clause, or pragma is given for a given aspect.

#### 5.1.58 `No_Standard_Allocators_After_Elaboration`

[RM D.7] Specifies that an allocator using a standard storage pool should never be evaluated at run time after the elaboration of the library items of the partition has completed. Otherwise, `Storage_Error` is raised.

#### 5.1.59 `No_Standard_Storage_Pools`

[GNAT] This restriction ensures at compile time that no access types use the standard default storage pool. Any access type declared must have an explicit `Storage_Pool` attribute defined specifying a user-defined storage pool.

#### 5.1.60 `No_Stream_Optimizations`

[GNAT] This restriction affects the performance of stream operations on types `String`, `Wide_String` and `Wide_Wide_String`. By default, the compiler uses block reads and writes when manipulating `String` objects due to their superior performance. When this restriction is in effect, the compiler performs all IO operations on a per-character basis.

#### 5.1.61 `No_Streams`

[GNAT] This restriction ensures at compile/bind time that there are no stream objects created and no use of stream attributes. This restriction does not forbid dependences on the package `Ada.Streams`. So it is permissible to with `Ada.Streams` (or another package that does so itself) as long as no actual stream objects are created and no stream attributes are used.

Note that the use of restriction allows optimization of tagged types, since they do not need to worry about dispatching stream operations. To take maximum advantage of this

space-saving optimization, any unit declaring a tagged type should be compiled with the restriction, though this is not required.

When pragmas `Discard_Names` and `Restrictions (No_Streams)` simultaneously apply to a tagged type, its `Expanded_Name` and `External_Tag` are also initialized with empty strings. In particular, both these pragmas can be applied as configuration pragmas to avoid exposing entity names at binary level for the entire partition.

#### 5.1.62 `No_Tagged_Type_Registration`

[GNAT] If this restriction is active, then class-wide streaming attributes are not supported. In addition, the subprograms in `Ada.Tags` are not supported. If this restriction is active, the generated code is simplified by omitting the otherwise-required global registration of tagged types when they are declared. This restriction may be necessary in order to also apply the `No_Elaboration_Code` restriction.

#### 5.1.63 `No_Task_Allocators`

[RM D.7] There are no allocators for task types or types containing task subcomponents.

#### 5.1.64 `No_Task_At_Interrupt_Priority`

[GNAT] This restriction ensures at compile time that there is no `Interrupt_Priority` aspect or pragma for a task or a task type. As a consequence, the tasks are always created with a priority below that an interrupt priority.

#### 5.1.65 `No_Task_Attributes_Package`

[GNAT] This restriction ensures at compile time that there are no implicit or explicit dependencies on the package `Ada.Task_Attributes`.

The restriction `No_Task_Attributes` is recognized as a synonym for `No_Task_Attributes_Package`. This is retained for historical compatibility purposes (and a warning will be generated for its use if warnings on obsolescent features are activated).

#### 5.1.66 `No_Task_Hierarchy`

[RM D.7] All (non-environment) tasks depend directly on the environment task of the partition.

#### 5.1.67 `No_Task_Termination`

[RM D.7] Tasks that terminate are erroneous.

#### 5.1.68 `No_Tasking`

[GNAT] This restriction prevents the declaration of tasks or task types throughout the partition. It is similar in effect to the use of `Max_Tasks => 0` except that violations are caught at compile time and cause an error message to be output either by the compiler or binder.

#### 5.1.69 `No_Terminate_Alternatives`

[RM D.7] There are no selective accepts with terminate alternatives.

**5.1.70 No\_Unchecked\_Access**

[RM H.4] This restriction ensures at compile time that there are no occurrences of the `Unchecked_Access` attribute.

**5.1.71 No\_Unchecked\_Conversion**

[RM J.13] This restriction ensures at compile time that there are no semantic dependences on the predefined generic function `Unchecked_Conversion`.

**5.1.72 No\_Unchecked\_Deallocation**

[RM J.13] This restriction ensures at compile time that there are no semantic dependences on the predefined generic procedure `Unchecked_Deallocation`.

**5.1.73 No\_Use\_Of\_Attribute**

[RM 13.12.1] This is a standard Ada 2012 restriction that is GNAT defined in earlier versions of Ada.

**5.1.74 No\_Use\_Of\_Entity**

[GNAT] This restriction ensures at compile time that there are no references to the entity given in the form

```
No_Use_Of_Entity => Name
```

where `Name` is the fully qualified entity, for example

```
No_Use_Of_Entity => Ada.Text_IO.Put_Line
```

**5.1.75 No\_Use\_Of\_Pragma**

[RM 13.12.1] This is a standard Ada 2012 restriction that is GNAT defined in earlier versions of Ada.

**5.1.76 Pure\_Barriers**

[GNAT] This restriction ensures at compile time that protected entry barriers are restricted to:

- \* components of the protected object (excluding selection from dereferences),
- \* constant declarations,
- \* named numbers,
- \* enumeration literals,
- \* integer literals,
- \* real literals,
- \* character literals,
- \* implicitly defined comparison operators,
- \* uses of the Standard.”not” operator,
- \* short-circuit operator,
- \* the `Count` attribute

This restriction is a relaxation of the `Simple_Barriers` restriction, but still ensures absence of side effects, exceptions, and recursion during the evaluation of the barriers.

### 5.1.77 Simple\_Barriers

[RM D.7] This restriction ensures at compile time that barriers in entry declarations for protected types are restricted to either static boolean expressions or references to simple boolean variables defined in the private part of the protected type. No other form of entry barriers is permitted.

The restriction `Boolean_Entry_Barriers` is recognized as a synonym for `Simple_Barriers`. This is retained for historical compatibility purposes (and a warning will be generated for its use if warnings on obsolescent features are activated).

### 5.1.78 Static\_Priorities

[GNAT] This restriction ensures at compile time that all priority expressions are static, and that there are no dependences on the package `Ada.Dynamic_Priorities`.

### 5.1.79 Static\_Storage\_Size

[GNAT] This restriction ensures at compile time that any expression appearing in a `Storage_Size` pragma or attribute definition clause is static.

## 5.2 Program Unit Level Restrictions

The second set of restriction identifiers does not require partition-wide consistency. The restriction may be enforced for a single compilation unit without any effect on any of the other compilation units in the partition.

### 5.2.1 No\_Elaboration\_Code

[GNAT] This restriction ensures at compile time that no elaboration code is generated. Note that this is not the same condition as is enforced by pragma `Preelaborate`. There are cases in which pragma `Preelaborate` still permits code to be generated (e.g., code to initialize a large array to all zeroes), and there are cases of units which do not meet the requirements for pragma `Preelaborate`, but for which no elaboration code is generated. Generally, it is the case that preelaborable units will meet the restrictions, with the exception of large aggregates initialized with an `others_clause`, and exception declarations (which generate calls to a run-time registry procedure). This restriction is enforced on a unit by unit basis, it need not be obeyed consistently throughout a partition.

In the case of aggregates with `others`, if the aggregate has a dynamic size, there is no way to eliminate the elaboration code (such dynamic bounds would be incompatible with `Preelaborate` in any case). If the bounds are static, then use of this restriction actually modifies the code choice of the compiler to avoid generating a loop, and instead generate the aggregate statically if possible, no matter how many times the data for the `others` clause must be repeatedly generated.

It is not possible to precisely document the constructs which are compatible with this restriction, since, unlike most other restrictions, this is not a restriction on the source code, but a restriction on the generated object code. For example, if the source contains a declaration:

```
Val : constant Integer := X;
```

where X is not a static constant, it may be possible, depending on complex optimization circuitry, for the compiler to figure out the value of X at compile time, in which case this

initialization can be done by the loader, and requires no initialization code. It is not possible to document the precise conditions under which the optimizer can figure this out.

Note that this the implementation of this restriction requires full code generation. If it is used in conjunction with “semantics only” checking, then some cases of violations may be missed.

When this restriction is active, we are not requesting control-flow preservation with `-fpreserve-control-flow`, and the static elaboration model is used, the compiler is allowed to suppress the elaboration counter normally associated with the unit. This counter is typically used to check for access before elaboration and to control multiple elaboration attempts.

### 5.2.2 No\_Dynamic\_Accessibility\_Checks

[GNAT] No dynamic accessibility checks are generated when this restriction is in effect. Instead, dangling references are prevented via more conservative compile-time checking. More specifically, existing compile-time checks are enforced but with more conservative assumptions about the accessibility levels of the relevant entities. These conservative assumptions eliminate the need for dynamic accessibility checks.

These new rules for computing (at compile-time) the accessibility level of an anonymous access type `T` are as follows:

- \* If `T` is a function result type then, from the caller’s perspective, its level is that of the innermost master enclosing the function call. From the callee’s perspective, the level of parameters and local variables of the callee is statically deeper than the level of `T`.  
For any other accessibility level `L` such that the level of parameters and local variables of the callee is statically deeper than `L`, the level of `T` (from the callee’s perspective) is also statically deeper than `L`.
- \* If `T` is the type of a formal parameter then, from the caller’s perspective, its level is at least as deep as that of the type of the corresponding actual parameter (whatever that actual parameter might be). From the callee’s perspective, the level of parameters and local variables of the callee is statically deeper than the level of `T`.
- \* If `T` is the type of a discriminant then its level is that of the discriminated type.
- \* If `T` is the type of a stand-alone object then its level is the level of the object.
- \* In all other cases, the level of `T` is as defined by the existing rules of Ada.

### 5.2.3 No\_Dynamic\_Sized\_Objects

[GNAT] This restriction disallows certain constructs that might lead to the creation of dynamic-sized composite objects (or array or discriminated type). An array subtype indication is illegal if the bounds are not static or references to discriminants of an enclosing type. A discriminated subtype indication is illegal if the type has discriminant-dependent array components or a variant part, and the discriminants are not static. In addition, array and record aggregates are illegal in corresponding cases. Note that this restriction does not forbid access discriminants. It is often a good idea to combine this restriction with `No_Secondary_Stack`.

### 5.2.4 No\_Entry\_Queue

[GNAT] This restriction is a declaration that any protected entry compiled in the scope of the restriction has at most one task waiting on the entry at any one time, and so no queue is required. This restriction is not checked at compile time. A program execution is erroneous if an attempt is made to queue a second task on such an entry.

### 5.2.5 No\_Implementation\_Aspect\_Specifications

[RM 13.12.1] This restriction checks at compile time that no GNAT-defined aspects are present. With this restriction, the only aspects that can be used are those defined in the Ada Reference Manual.

### 5.2.6 No\_Implementation\_Attributes

[RM 13.12.1] This restriction checks at compile time that no GNAT-defined attributes are present. With this restriction, the only attributes that can be used are those defined in the Ada Reference Manual.

### 5.2.7 No\_Implementation\_Identifiers

[RM 13.12.1] This restriction checks at compile time that no implementation-defined identifiers (marked with pragma `Implementation_Defined`) occur within language-defined packages.

### 5.2.8 No\_Implementation\_Pragmas

[RM 13.12.1] This restriction checks at compile time that no GNAT-defined pragmas are present. With this restriction, the only pragmas that can be used are those defined in the Ada Reference Manual.

### 5.2.9 No\_Implementation\_Restrictions

[GNAT] This restriction checks at compile time that no GNAT-defined restriction identifiers (other than `No_Implementation_Restrictions` itself) are present. With this restriction, the only other restriction identifiers that can be used are those defined in the Ada Reference Manual.

### 5.2.10 No\_Implementation\_Units

[RM 13.12.1] This restriction checks at compile time that there is no mention in the context clause of any implementation-defined descendants of packages `Ada`, `Interfaces`, or `System`.

### 5.2.11 No\_Implicit\_Aliasing

[GNAT] This restriction, which is not required to be partition-wide consistent, requires an explicit aliased keyword for an object to which `'Access`, `'Unchecked_Access`, or `'Address` is applied, and forbids entirely the use of the `'Unrestricted_Access` attribute for objects. Note: the reason that `Unrestricted_Access` is forbidden is that it would require the prefix to be aliased, and in such cases, it can always be replaced by the standard attribute `Unchecked_Access` which is preferable.

### 5.2.12 No\_Implicit\_Loops

[GNAT] This restriction ensures that the generated code of the unit marked with this restriction does not contain any implicit `for` loops, either by modifying the generated code where possible, or by rejecting any construct that would otherwise generate an implicit `for` loop. If this restriction is active, it is possible to build large array aggregates with all static components without generating an intermediate temporary, and without generating a loop to initialize individual components. Otherwise, a loop is created for arrays larger than about 5000 scalar components. Note that if this restriction is set in the spec of a package, it will not apply to its body.

### 5.2.13 No\_Obsolescent\_Features

[RM 13.12.1] This restriction checks at compile time that no obsolescent features are used, as defined in Annex J of the Ada Reference Manual.

### 5.2.14 No\_Wide\_Characters

[GNAT] This restriction ensures at compile time that no uses of the types `Wide_Character` or `Wide_String` or corresponding wide wide types appear, and that no wide or wide wide string or character literals appear in the program (that is literals representing characters not in type `Character`).

### 5.2.15 Static\_Dispatch\_Tables

[GNAT] This restriction checks at compile time that all the artifacts associated with dispatch tables can be placed in read-only memory.

### 5.2.16 SPARK\_05

[GNAT] This restriction no longer has any effect and is superseded by SPARK 2014, whose restrictions are checked by the tool GNATprove. To check that a codebase respects SPARK 2014 restrictions, mark the code with pragma or aspect `SPARK_Mode`, and run the tool GNATprove at Stone assurance level, as follows:

```
gnatprove -P project.gpr --mode=stone
```

or equivalently:

```
gnatprove -P project.gpr --mode=check_all
```

## 6 Implementation Advice

The main text of the Ada Reference Manual describes the required behavior of all Ada compilers, and the GNAT compiler conforms to these requirements.

In addition, there are sections throughout the Ada Reference Manual headed by the phrase ‘Implementation advice’. These sections are not normative, i.e., they do not specify requirements that all compilers must follow. Rather they provide advice on generally desirable behavior. They are not requirements, because they describe behavior that cannot be provided on all systems, or may be undesirable on some systems.

As far as practical, GNAT follows the implementation advice in the Ada Reference Manual. Each such RM section corresponds to a section in this chapter whose title specifies the RM section number and paragraph number and the subject of the advice. The contents of each section consists of the RM text within quotation marks, followed by the GNAT interpretation of the advice. Most often, this simply says ‘followed’, which means that GNAT follows the advice. However, in a number of cases, GNAT deliberately deviates from this advice, in which case the text describes what GNAT does and why.

### 6.1 RM 1.1.3(20): Error Detection

“If an implementation detects the use of an unsupported Specialized Needs Annex feature at run time, it should raise `Program_Error` if feasible.”

Not relevant. All specialized needs annex features are either supported, or diagnosed at compile time.

### 6.2 RM 1.1.3(31): Child Units

“If an implementation wishes to provide implementation-defined extensions to the functionality of a language-defined library unit, it should normally do so by adding children to the library unit.”

Followed.

### 6.3 RM 1.1.5(12): Bounded Errors

“If an implementation detects a bounded error or erroneous execution, it should raise `Program_Error`.”

Followed in all cases in which the implementation detects a bounded error or erroneous execution. Not all such situations are detected at runtime.

### 6.4 RM 2.8(16): Pragmas

“Normally, implementation-defined pragmas should have no semantic effect for error-free programs; that is, if the implementation-defined pragmas are removed from a working program, the program should still be legal, and should still have the same semantics.”

The following implementation defined pragmas are exceptions to this rule:

<b>Pragma</b>	<b>Explanation</b>
<code>'Abort_Defer'</code>	Affects semantics
<code>'Ada_83'</code>	Affects legality
<code>'Assert'</code>	Affects semantics
<code>'CPP_Class'</code>	Affects semantics
<code>'CPP_Constructor'</code>	Affects semantics
<code>'Debug'</code>	Affects semantics
<code>'Interface_Name'</code>	Affects semantics
<code>'Machine_Attribute'</code>	Affects semantics
<code>'Unimplemented_Unit'</code>	Affects legality
<code>'Unchecked_Union'</code>	Affects semantics

In each of the above cases, it is essential to the purpose of the pragma that this advice not be followed. For details see [Implementation Defined Pragmas], page 4.

## 6.5 RM 2.8(17-19): Pragmas

“Normally, an implementation should not define pragmas that can make an illegal program legal, except as follows:

- \* A pragma used to complete a declaration, such as a pragma `Import`;
- \* A pragma used to configure the environment by adding, removing, or replacing `library_items`.”

See [RM 2.8(16); Pragmas], page 160.

## 6.6 RM 3.5.2(5): Alternative Character Sets

“If an implementation supports a mode with alternative interpretations for `Character` and `Wide_Character`, the set of graphic characters of `Character` should nevertheless remain a proper subset of the set of graphic characters of `Wide_Character`. Any character set ‘localizations’ should be reflected in the results of the subprograms defined in the language-defined package `Characters.Handling` (see A.3) available in such a mode. In a mode with an alternative interpretation of `Character`, the implementation should also support a corresponding change in what is a legal `identifier_letter`.”

Not all wide character modes follow this advice, in particular the JIS and IEC modes reflect standard usage in Japan, and in these encoding, the upper half of the Latin-1 set is not part of the wide-character subset, since the most significant bit is used for wide character encoding. However, this only applies to the external forms. Internally there is no such restriction.

## 6.7 RM 3.5.4(28): Integer Types

“An implementation should support `Long_Integer` in addition to `Integer` if the target machine supports 32-bit (or longer) arithmetic. No other named integer subtypes are recommended for package `Standard`. Instead, appropriate named integer subtypes should be provided in the library package `Interfaces` (see B.2).”

`Long_Integer` is supported. Other standard integer types are supported so this advice is not fully followed. These types are supported for convenient interface to C, and so that all hardware types of the machine are easily available.

## 6.8 RM 3.5.4(29): Integer Types

“An implementation for a two’s complement machine should support modular types with a binary modulus up to `System.Max_Int*2+2`. An implementation should support a non-binary modules up to `Integer'Last`.”

Followed.

## 6.9 RM 3.5.5(8): Enumeration Values

“For the evaluation of a call on `S'Pos` for an enumeration subtype, if the value of the operand does not correspond to the internal code for any enumeration literal of its type (perhaps due to an un-initialized variable), then the implementation should raise `Program_Error`. This is particularly important for enumeration types with noncontiguous internal codes specified by an `enumeration_representation_clause`.”

Followed.

## 6.10 RM 3.5.7(17): Float Types

“An implementation should support `Long_Float` in addition to `Float` if the target machine supports 11 or more digits of precision. No other named floating point subtypes are recommended for package `Standard`. Instead, appropriate named floating point subtypes should be provided in the library package `Interfaces` (see B.2).”

`Short_Float` and `Long_Long_Float` are also provided. The former provides improved compatibility with other implementations supporting this type. The latter corresponds to the highest precision floating-point type supported by the hardware. On most machines, this will be the same as `Long_Float`, but on some machines, it will correspond to the IEEE extended form. The notable case is all x86 implementations, where `Long_Long_Float` corresponds to the 80-bit extended precision format supported in hardware on this processor.

Note that the 128-bit format on SPARC is not supported, since this is a software rather than a hardware format.

### 6.11 RM 3.6.2(11): Multidimensional Arrays

“An implementation should normally represent multidimensional arrays in row-major order, consistent with the notation used for multidimensional array aggregates (see 4.3.3). However, if a pragma `Convention (Fortran, ...)` applies to a multidimensional array type, then column-major order should be used instead (see B.5, ‘Interfacing with Fortran’).”

Followed.

### 6.12 RM 9.6(30-31): Duration'Small

“Whenever possible in an implementation, the value of `Duration'Small` should be no greater than 100 microseconds.”

Followed. (`Duration'Small = 10*(-9)`).

“The time base for `delay_relative_statements` should be monotonic; it need not be the same time base as used for `Calendar.Clock`.”

Followed.

### 6.13 RM 10.2.1(12): Consistent Representation

“In an implementation, a type declared in a pre-elaborated package should have the same representation in every elaboration of a given version of the package, whether the elaborations occur in distinct executions of the same program, or in executions of distinct programs or partitions that include the given version.”

Followed, except in the case of tagged types. Tagged types involve implicit pointers to a local copy of a dispatch table, and these pointers have representations which thus depend on a particular elaboration of the package. It is not easy to see how it would be possible to follow this advice without severely impacting efficiency of execution.

### 6.14 RM 11.4.1(19): Exception Information

“`Exception_Message` by default and `Exception_Information` should produce information useful for debugging. `Exception_Message` should be short, about one line. `Exception_Information` can be long. `Exception_Message` should not include the `Exception_Name`. `Exception_Information` should include both the `Exception_Name` and the `Exception_Message`.”

Followed. For each exception that doesn't have a specified `Exception_Message`, the compiler generates one containing the location of the raise statement. This location has the form ‘file\_name:line’, where file\_name is the short file name (without path information) and line is the line number in the file. Note that in the case of the Zero Cost Exception mechanism, these messages become redundant with the `Exception_Information` that contains a full backtrace of the calling sequence, so they are disabled. To disable explicitly the generation of the source location message, use the `Pragma Discard_Names`.

### 6.15 RM 11.5(28): Suppression of Checks

“The implementation should minimize the code executed for checks that have been suppressed.”

Followed.

### 6.16 RM 13.1 (21-24): Representation Clauses

“The recommended level of support for all representation items is qualified as follows:

An implementation need not support representation items containing nonstatic expressions, except that an implementation should support a representation item for a given entity if each nonstatic expression in the representation item is a name that statically denotes a constant declared before the entity.”

Followed. In fact, GNAT goes beyond the recommended level of support by allowing nonstatic expressions in some representation clauses even without the need to declare constants initialized with the values of such expressions. For example:

```
X : Integer;
Y : Float;
for Y'Address use X'Address;
```

is accepted directly by GNAT.

“An implementation need not support a specification for the **Size** for a given composite subtype, nor the size or storage place for an object (including a component) of a given composite subtype, unless the constraints on the subtype and its composite subcomponents (if any) are all static constraints.”

Followed. Size Clauses are not permitted on nonstatic components, as described above.

“An aliased component, or a component whose type is by-reference, should always be allocated at an addressable location.”

Followed.

### 6.17 RM 13.2(6-8): Packed Types

“If a type is packed, then the implementation should try to minimize storage allocated to objects of the type, possibly at the expense of speed of accessing components, subject to reasonable complexity in addressing calculations.

The recommended level of support pragma **Pack** is:

For a packed record type, the components should be packed as tightly as possible subject to the Sizes of the component subtypes, and subject to any ‘record-representation-clause’ that applies to the type; the implementation may, but need not, reorder components or cross aligned word boundaries to improve the packing. A component whose **Size** is greater than the word size may be allocated an integral number of words.”

Followed. Tight packing of arrays is supported for all component sizes up to 64-bits. If the array component size is 1 (that is to say, if the component is a boolean type or an enumeration type with two values) then values of the type are implicitly initialized to zero. This happens both for objects of the packed type, and for objects that have a subcomponent of the packed type.

## 6.18 RM 13.3(14-19): Address Clauses

“For an array **X**, **X'Address** should point at the first component of the array, and not at the array bounds.”

Followed.

“The recommended level of support for the **Address** attribute is:

**X'Address** should produce a useful result if **X** is an object that is aliased or of a by-reference type, or is an entity whose **Address** has been specified.”

Followed. A valid address will be produced even if none of those conditions have been met. If necessary, the object is forced into memory to ensure the address is valid.

“An implementation should support **Address** clauses for imported subprograms.”

Followed.

“Objects (including subcomponents) that are aliased or of a by-reference type should be allocated on storage element boundaries.”

Followed.

“If the **Address** of an object is specified, or it is imported or exported, then the implementation should not perform optimizations based on assumptions of no aliases.”

Followed.

## 6.19 RM 13.3(29-35): Alignment Clauses

“The recommended level of support for the **Alignment** attribute for subtypes is:

An implementation should support specified Alignments that are factors and multiples of the number of storage elements per word, subject to the following:”

Followed.

“An implementation need not support specified Alignments for combinations of Sizes and Alignments that cannot be easily loaded and stored by available machine instructions.”

Followed.

“An implementation need not support specified Alignments that are greater than the maximum **Alignment** the implementation ever returns by default.”

Followed.

“The recommended level of support for the **Alignment** attribute for objects is:

Same as above, for subtypes, but in addition:”

Followed.

“For stand-alone library-level objects of statically constrained subtypes, the implementation should support all alignments supported by the target linker. For example, page alignment is likely to be supported for such objects, but not for subtypes.”

Followed.

## 6.20 RM 13.3(42-43): Size Clauses

“The recommended level of support for the **Size** attribute of objects is:

A **Size** clause should be supported for an object if the specified **Size** is at least as large as its subtype’s **Size**, and corresponds to a size in storage elements that is a multiple of the object’s **Alignment** (if the **Alignment** is nonzero).”

Followed.

## 6.21 RM 13.3(50-56): Size Clauses

“If the **Size** of a subtype is specified, and allows for efficient independent addressability (see 9.10) on the target architecture, then the **Size** of the following objects of the subtype should equal the **Size** of the subtype:

Aliased objects (including components).”

Followed.

“*Size* clause on a composite subtype should not affect the internal layout of components.”

Followed. But note that this can be overridden by use of the implementation pragma `Implicit_Packing` in the case of packed arrays.

“The recommended level of support for the **Size** attribute of subtypes is:

The **Size** (if not specified) of a static discrete or fixed point subtype should be the number of bits needed to represent each value belonging to the subtype using an unbiased representation, leaving space for a sign bit only if the subtype contains negative values. If such a subtype is a first subtype, then an implementation should support a specified **Size** for it that reflects this representation.”

Followed.

“For a subtype implemented with levels of indirection, the **Size** should include the size of the pointers, but not the size of what they point at.”

Followed.

## 6.22 RM 13.3(71-73): Component Size Clauses

“The recommended level of support for the `Component_Size` attribute is:  
An implementation need not support specified `Component_Sizes` that are less than the `Size` of the component subtype.”

Followed.

“An implementation should support specified `Component_Sizes` that are factors and multiples of the word size. For such `Component_Sizes`, the array should contain no gaps between components. For other `Component_Sizes` (if supported), the array should contain no gaps between components when packing is also specified; the implementation should forbid this combination in cases where it cannot support a no-gaps representation.”

Followed.

## 6.23 RM 13.4(9-10): Enumeration Representation Clauses

“The recommended level of support for enumeration representation clauses is:

An implementation need not support enumeration representation clauses for boolean types, but should at minimum support the internal codes in the range `System.Min_Int .. System.Max_Int`.”

Followed.

## 6.24 RM 13.5.1(17-22): Record Representation Clauses

“The recommended level of support for ‘`record_representation_clause`’s is:

An implementation should support storage places that can be extracted with a load, mask, shift sequence of machine code, and set with a load, shift, mask, store sequence, given the available machine instructions and run-time model.”

Followed.

“A storage place should be supported if its size is equal to the `Size` of the component subtype, and it starts and ends on a boundary that obeys the `Alignment` of the component subtype.”

Followed.

“If the default bit ordering applies to the declaration of a given type, then for a component whose subtype’s `Size` is less than the word size, any storage place that does not cross an aligned word boundary should be supported.”

Followed.

“An implementation may reserve a storage place for the tag field of a tagged type, and disallow other components from overlapping that place.”

Followed. The storage place for the tag field is the beginning of the tagged record, and its size is `Address'Size`. GNAT will reject an explicit component clause for the tag field.

“An implementation need not support a ‘component\_clause’ for a component of an extension part if the storage place is not after the storage places of all components of the parent type, whether or not those storage places had been specified.”

Followed. The above advice on record representation clauses is followed, and all mentioned features are implemented.

## 6.25 RM 13.5.2(5): Storage Place Attributes

“If a component is represented using some form of pointer (such as an offset) to the actual data of the component, and this data is contiguous with the rest of the object, then the storage place attributes should reflect the place of the actual data, not the pointer. If a component is allocated discontinuously from the rest of the object, then a warning should be generated upon reference to one of its storage place attributes.”

Followed. There are no such components in GNAT.

## 6.26 RM 13.5.3(7-8): Bit Ordering

“The recommended level of support for the non-default bit ordering is: The implementation should support the nondefault bit ordering in addition to the default bit ordering.”

Followed.

## 6.27 RM 13.7(37): Address as Private

“*Address* should be of a private type.”

Followed.

## 6.28 RM 13.7.1(16): Address Operations

“Operations in `System` and its children should reflect the target environment semantics as closely as is reasonable. For example, on most machines, it makes sense for address arithmetic to ‘wrap around’. Operations that do not make sense should raise `Program_Error`.”

Followed. Address arithmetic is modular arithmetic that wraps around. No operation raises `Program_Error`, since all operations make sense.

## 6.29 RM 13.9(14-17): Unchecked Conversion

“The `Size` of an array object should not include its bounds; hence, the bounds should not be part of the converted data.”

Followed.

“The implementation should not generate unnecessary run-time checks to ensure that the representation of `S` is a representation of the target

type. It should take advantage of the permission to return by reference when possible. Restrictions on unchecked conversions should be avoided unless required by the target environment.”

Followed. There are no restrictions on unchecked conversion. A warning is generated if the source and target types do not have the same size since the semantics in this case may be target dependent.

“The recommended level of support for unchecked conversions is:

Unchecked conversions should be supported and should be reversible in the cases where this clause defines the result. To enable meaningful use of unchecked conversion, a contiguous representation should be used for elementary subtypes, for statically constrained array subtypes whose component subtype is one of the subtypes described in this paragraph, and for record subtypes without discriminants whose component subtypes are described in this paragraph.”

Followed.

### 6.30 RM 13.11(23-25): Implicit Heap Usage

“An implementation should document any cases in which it dynamically allocates heap storage for a purpose other than the evaluation of an allocator.”

Followed, the only other points at which heap storage is dynamically allocated are as follows:

- \* At initial elaboration time, to allocate dynamically sized global objects.
- \* To allocate space for a task when a task is created.
- \* To extend the secondary stack dynamically when needed. The secondary stack is used for returning variable length results.

“A default (implementation-provided) storage pool for an access-to-constant type should not have overhead to support deallocation of individual objects.”

Followed.

“A storage pool for an anonymous access type should be created at the point of an allocator for the type, and be reclaimed when the designated object becomes inaccessible.”

Followed.

### 6.31 RM 13.11.2(17): Unchecked Deallocation

“For a standard storage pool, `Free` should actually reclaim the storage.”

Followed.

### 6.32 RM 13.13.2(1.6): Stream Oriented Attributes

“If not specified, the value of `Stream_Size` for an elementary type should be the number of bits that corresponds to the minimum number of stream elements required by the first subtype of the type, rounded up to the

nearest factor or multiple of the word size that is also a multiple of the stream element size.”

Followed, except that the number of stream elements is 1, 2, 3, 4 or 8. The `Stream_Size` may be used to override the default choice.

The default implementation is based on direct binary representations and is therefore target- and endianness-dependent. To address this issue, GNAT also supplies an alternate implementation of the stream attributes `Read` and `Write`, which uses the target-independent XDR standard representation for scalar types. This XDR alternative can be enabled via the binder switch `-xdr`.

### 6.33 RM A.1(52): Names of Predefined Numeric Types

“If an implementation provides additional named predefined integer types, then the names should end with `Integer` as in `Long_Integer`. If an implementation provides additional named predefined floating point types, then the names should end with `Float` as in `Long_Float`.”

Followed.

### 6.34 RM A.3.2(49): `Ada.Characters.Handling`

“If an implementation provides a localized definition of `Character` or `Wide_Character`, then the effects of the subprograms in `Characters.Handling` should reflect the localizations. See also 3.5.2.”

Followed. GNAT provides no such localized definitions.

### 6.35 RM A.4.4(106): Bounded-Length String Handling

“Bounded string objects should not be implemented by implicit pointers and dynamic allocation.”

Followed. No implicit pointers or dynamic allocation are used.

### 6.36 RM A.5.2(46-47): Random Number Generation

“Any storage associated with an object of type `Generator` should be reclaimed on exit from the scope of the object.”

Followed.

“If the generator period is sufficiently long in relation to the number of distinct initiator values, then each possible value of `Initiator` passed to `Reset` should initiate a sequence of random numbers that does not, in a practical sense, overlap the sequence initiated by any other value. If this is not possible, then the mapping between initiator values and generator states should be a rapidly varying function of the initiator value.”

Followed. The generator period is sufficiently long for the first condition here to hold true.

### 6.37 RM A.10.7(23): `Get_Immediate`

“The `Get_Immediate` procedures should be implemented with unbuffered input. For a device such as a keyboard, input should be available if a key has already been typed, whereas for a disk file, input should always be available except at end of file. For a file associated with a keyboard-like device, any line-editing features of the underlying operating system should be disabled during the execution of `Get_Immediate`.”

Followed on all targets except VxWorks. For VxWorks, there is no way to provide this functionality that does not result in the input buffer being flushed before the `Get_Immediate` call. A special unit `Interfaces.Vxworks.IO` is provided that contains routines to enable this functionality.

### 6.38 RM A.18: Containers

All implementation advice pertaining to `Ada.Containers` and its child units (that is, all implementation advice occurring within section A.18 and its subsections) is followed except for A.18.24(17):

“Bounded ordered set objects should be implemented without implicit pointers or dynamic allocation. “

The implementations of the two `Reference_Preserving_Key` functions of the generic package `Ada.Containers.Bounded_Ordered_Sets` each currently make use of dynamic allocation; other operations on bounded ordered set objects follow the implementation advice.

### 6.39 RM B.1(39-41): `Pragma Export`

“If an implementation supports pragma `Export` to a given language, then it should also allow the main subprogram to be written in that language. It should support some mechanism for invoking the elaboration of the Ada library units included in the system, and for invoking the finalization of the environment task. On typical systems, the recommended mechanism is to provide two subprograms whose link names are `adainit` and `adafinal`. `adainit` should contain the elaboration code for library units. `adafinal` should contain the finalization code. These subprograms should have no effect the second and subsequent time they are called.”

Followed.

“Automatic elaboration of pre-elaborated packages should be provided when pragma `Export` is supported.”

Followed when the main program is in Ada. If the main program is in a foreign language, then `adainit` must be called to elaborate pre-elaborated packages.

“For each supported convention ‘L’ other than `Intrinsic`, an implementation should support `Import` and `Export` pragmas for objects of ‘L’-compatible types and for subprograms, and pragma `Convention` for ‘L’-eligible types and for subprograms, presuming the other language has corresponding features. Pragma `Convention` need not be supported for scalar types.”

Followed.

## 6.40 RM B.2(12-13): Package Interfaces

“For each implementation-defined convention identifier, there should be a child package of package `Interfaces` with the corresponding name. This package should contain any declarations that would be useful for interfacing to the language (implementation) represented by the convention. Any declarations useful for interfacing to any language on the given hardware architecture should be provided directly in `Interfaces`.”

Followed.

“An implementation supporting an interface to C, COBOL, or Fortran should provide the corresponding package or packages described in the following clauses.”

Followed. GNAT provides all the packages described in this section.

## 6.41 RM B.3(63-71): Interfacing with C

“An implementation should support the following interface correspondences between Ada and C.”

Followed.

“An Ada procedure corresponds to a void-returning C function.”

Followed.

“An Ada function corresponds to a non-void C function.”

Followed.

“An Ada `in` scalar parameter is passed as a scalar argument to a C function.”

Followed.

“An Ada `in` parameter of an access-to-object type with designated type `T` is passed as a `t*` argument to a C function, where `t` is the C type corresponding to the Ada type `T`.”

Followed.

“An Ada access `T` parameter, or an Ada `out` or `in out` parameter of an elementary type `T`, is passed as a `t*` argument to a C function, where `t` is the C type corresponding to the Ada type `T`. In the case of an elementary `out` or `in out` parameter, a pointer to a temporary copy is used to preserve by-copy semantics.”

Followed.

“An Ada parameter of a record type `T`, of any mode, is passed as a `t*` argument to a C function, where `t` is the C structure corresponding to the Ada type `T`.”

Followed. This convention may be overridden by the use of the `C_Pass_By_Copy` pragma, or `Convention`, or by explicitly specifying the mechanism for a given call using an extended import or export pragma.

“An Ada parameter of an array type with component type `T`, of any mode, is passed as a `t*` argument to a C function, where `t` is the C type corresponding to the Ada type `T`.”





































































































































































































































































































































































































































































