

## Hardware Locality (hwloc)

v1.11-20191017.0500.git1fd9c25

Generated by Doxygen 1.7.6.1

Thu Oct 17 2019 05:01:01



# Contents

<b>1</b>	<b>Hardware Locality</b>	<b>1</b>
1.1	Introduction	1
1.2	Installation	3
1.2.1	Basic Installation	3
1.2.2	Installing from a Git clone	3
1.2.3	Optional Dependencies	3
1.3	Command-line Examples	4
1.4	Programming Interface	7
1.4.1	Portability	7
1.4.2	API Example	8
1.5	Questions and Bugs	12
1.6	History / Credits	12
1.7	Further Reading	12
<b>2</b>	<b>Terms and Definitions</b>	<b>15</b>
2.1	Objects	15
2.2	Indexes and Sets	15
2.3	Hierarchy, Tree and Levels	16
<b>3</b>	<b>Command-Line Tools</b>	<b>19</b>
3.1	lstopo and lstopo-no-graphics	19
3.2	hwloc-bind	19
3.3	hwloc-calc	20
3.4	hwloc-info	20
3.5	hwloc-distrib	20
3.6	hwloc-ps	21

---

3.7	hwloc-distances	21
3.8	hwloc-annotate	21
3.9	hwloc-diff, hwloc-patch and hwloc-compress-dir	21
3.10	hwloc-assembler	21
3.11	hwloc-assembler-remote	21
3.12	hwloc-dump-hwdata	22
3.13	hwloc-gather-topology	22
<b>4</b>	<b>Environment Variables</b>	<b>23</b>
<b>5</b>	<b>CPU and Memory Binding Overview</b>	<b>27</b>
<b>6</b>	<b>I/O Devices</b>	<b>29</b>
6.1	Enabling and requirements	29
6.2	I/O objects	29
6.3	OS devices	30
6.4	PCI devices and bridges	31
6.5	Consulting I/O devices and binding	31
6.6	Examples	32
<b>7</b>	<b>Miscellaneous objects</b>	<b>35</b>
7.1	Misc objects added by hwloc	35
7.2	Annotating topologies with Misc objects	36
<b>8</b>	<b>Multi-node Topologies</b>	<b>37</b>
8.1	Multi-node Objects Specificities	37
8.2	Assembling topologies with command-line tools	38
8.3	Assembling topologies with the programming interface	38
8.4	Example of assembly with the programming interface	39
<b>9</b>	<b>Object attributes</b>	<b>41</b>
9.1	Normal attributes	41
9.2	Custom string infos	41
<b>10</b>	<b>Importing and exporting topologies from/to XML files</b>	<b>45</b>
10.1	libxml2 and minimalistic XML backends	46
10.2	XML import error management	46

---

<b>11 Synthetic topologies</b>	<b>47</b>
11.1 Synthetic description string . . . . .	47
11.2 Loading a synthetic topology . . . . .	48
11.3 Exporting a topology as a synthetic string . . . . .	48
<b>12 Interoperability With Other Software</b>	<b>51</b>
<b>13 Thread Safety</b>	<b>53</b>
<b>14 Components and plugins</b>	<b>55</b>
14.1 Components enabled by default . . . . .	55
14.2 Selecting which components to use . . . . .	56
14.3 Loading components from plugins . . . . .	56
14.4 Adding new discovery components and plugins . . . . .	57
14.4.1 Basics of discovery components . . . . .	57
14.4.2 Registering a new discovery component . . . . .	57
14.5 Existing components and plugins . . . . .	58
<b>15 Embedding hwloc in Other Software</b>	<b>61</b>
15.1 Using hwloc's M4 Embedding Capabilities . . . . .	62
15.2 Example Embedding hwloc . . . . .	63
<b>16 Frequently Asked Questions</b>	<b>65</b>
16.1 Concepts . . . . .	65
16.1.1 I only need binding, why should I use hwloc ? . . . . .	65
16.1.2 Should I use logical or physical/OS indexes? and how? . . . . .	65
16.1.3 hwloc is only a structural model, it ignores performance models, memory bandwidth, etc.? . . . .	66
16.1.4 hwloc only has a one-dimensional view of the architecture, it ignores distances? . . . . .	67
16.1.5 What are these Group objects in my topology? . . . . .	67
16.1.6 What happens if my topology is asymmetric? . . . . .	68
16.1.7 What happens to my topology if I disable symmetric multithread- ing, hyper-threading, etc. in the system? . . . . .	68
16.1.8 How may I ignore symmetric multithreading, hyper-threading, etc. in hwloc? . . . . .	69
16.2 Advanced . . . . .	70

16.2.1	I do not want hwloc to rediscover my enormous machine topology every time I rerun a process . . . . .	70
16.2.2	How many topologies may I use in my program? . . . . .	70
16.2.3	How to avoid memory waste when manipulating multiple similar topologies? . . . . .	71
16.2.4	How do I annotate the topology with private notes? . . . . .	71
16.3	Caveats . . . . .	71
16.3.1	Why is hwloc slow? . . . . .	71
16.3.2	Does hwloc require privileged access? . . . . .	72
16.3.3	What should I do when hwloc reports "operating system" warnings? . . . . .	73
16.3.4	Why does Valgrind complain about hwloc memory leaks? . . . . .	73
16.3.5	How do I handle ABI breaks and API upgrades? . . . . .	74
16.4	Platform-specific . . . . .	75
16.4.1	How do I find the local MCDRAM NUMA node on Intel Xeon Phi processor? . . . . .	75
16.4.2	Why do I need hwloc-dump-hwdata for memory on Intel Xeon Phi processor? . . . . .	76
16.4.3	How do I build for Intel Xeon Phi coprocessor? . . . . .	76
16.4.4	How do I build hwloc for BlueGene/Q? . . . . .	77
16.4.5	How do I build hwloc for Windows? . . . . .	77
16.4.6	How to get useful topology information on NetBSD? . . . . .	77
16.4.7	Why does binding fail on AIX? . . . . .	77
<b>17</b>	<b>Module Index</b>	<b>79</b>
17.1	Modules . . . . .	79
<b>18</b>	<b>Data Structure Index</b>	<b>81</b>
18.1	Data Structures . . . . .	81
<b>19</b>	<b>Module Documentation</b>	<b>83</b>
19.1	API version . . . . .	83
19.1.1	Define Documentation . . . . .	83
19.1.1.1	HWLOC_API_VERSION . . . . .	83
19.1.1.2	HWLOC_COMPONENT_ABI . . . . .	83
19.1.2	Function Documentation . . . . .	83
19.1.2.1	hwloc_get_api_version . . . . .	84

---

19.2 Object Sets ( <code>hwloc_cpuset_t</code> and <code>hwloc_nodeset_t</code> ) . . . . .	85
19.2.1 Detailed Description . . . . .	85
19.2.2 Typedef Documentation . . . . .	85
19.2.2.1 <code>hwloc_const_cpuset_t</code> . . . . .	85
19.2.2.2 <code>hwloc_const_nodeset_t</code> . . . . .	85
19.2.2.3 <code>hwloc_cpuset_t</code> . . . . .	85
19.2.2.4 <code>hwloc_nodeset_t</code> . . . . .	85
19.3 Object Types . . . . .	87
19.3.1 Typedef Documentation . . . . .	87
19.3.1.1 <code>hwloc_obj_bridge_type_t</code> . . . . .	87
19.3.1.2 <code>hwloc_obj_cache_type_t</code> . . . . .	87
19.3.1.3 <code>hwloc_obj_osdev_type_t</code> . . . . .	87
19.3.2 Enumeration Type Documentation . . . . .	88
19.3.2.1 <code>hwloc_compare_types_e</code> . . . . .	88
19.3.2.2 <code>hwloc_obj_bridge_type_e</code> . . . . .	88
19.3.2.3 <code>hwloc_obj_cache_type_e</code> . . . . .	88
19.3.2.4 <code>hwloc_obj_osdev_type_e</code> . . . . .	88
19.3.2.5 <code>hwloc_obj_type_t</code> . . . . .	89
19.3.3 Function Documentation . . . . .	90
19.3.3.1 <code>hwloc_compare_types</code> . . . . .	90
19.4 Object Structure and Attributes . . . . .	91
19.4.1 Typedef Documentation . . . . .	91
19.4.1.1 <code>hwloc_obj_t</code> . . . . .	91
19.5 Topology Creation and Destruction . . . . .	92
19.5.1 Typedef Documentation . . . . .	92
19.5.1.1 <code>hwloc_topology_t</code> . . . . .	92
19.5.2 Function Documentation . . . . .	92
19.5.2.1 <code>hwloc_topology_check</code> . . . . .	92
19.5.2.2 <code>hwloc_topology_destroy</code> . . . . .	92
19.5.2.3 <code>hwloc_topology_dup</code> . . . . .	93
19.5.2.4 <code>hwloc_topology_init</code> . . . . .	93
19.5.2.5 <code>hwloc_topology_load</code> . . . . .	93
19.6 Topology Detection Configuration and Query . . . . .	95
19.6.1 Detailed Description . . . . .	96

---

19.6.2	Enumeration Type Documentation	96
19.6.2.1	hwloc_topology_flags_e	96
19.6.3	Function Documentation	97
19.6.3.1	hwloc_topology_get_flags	97
19.6.3.2	hwloc_topology_get_support	98
19.6.3.3	hwloc_topology_get_userdata	98
19.6.3.4	hwloc_topology_ignore_all_keep_structure	98
19.6.3.5	hwloc_topology_ignore_type	98
19.6.3.6	hwloc_topology_ignore_type_keep_structure	98
19.6.3.7	hwloc_topology_is_thissystem	99
19.6.3.8	hwloc_topology_set_custom	99
19.6.3.9	hwloc_topology_set_distance_matrix	99
19.6.3.10	hwloc_topology_set_flags	100
19.6.3.11	hwloc_topology_set_fsroot	100
19.6.3.12	hwloc_topology_set_pid	100
19.6.3.13	hwloc_topology_set_synthetic	101
19.6.3.14	hwloc_topology_set_userdata	101
19.6.3.15	hwloc_topology_set_xml	102
19.6.3.16	hwloc_topology_set_xmlbuffer	102
19.7	Object levels, depths and types	104
19.7.1	Detailed Description	104
19.7.2	Enumeration Type Documentation	104
19.7.2.1	hwloc_get_type_depth_e	104
19.7.3	Function Documentation	105
19.7.3.1	hwloc_get_depth_type	105
19.7.3.2	hwloc_get_nbobjs_by_depth	105
19.7.3.3	hwloc_get_nbobjs_by_type	105
19.7.3.4	hwloc_get_next_obj_by_depth	105
19.7.3.5	hwloc_get_next_obj_by_type	105
19.7.3.6	hwloc_get_obj_by_depth	106
19.7.3.7	hwloc_get_obj_by_type	106
19.7.3.8	hwloc_get_root_obj	106
19.7.3.9	hwloc_get_type_depth	106
19.7.3.10	hwloc_get_type_or_above_depth	107



---

19.7.3.11	hwloc_get_type_or_below_depth	107
19.7.3.12	hwloc_topology_get_depth	107
19.8	Converting between Object Types, Sets and Attributes, and Strings	108
19.8.1	Function Documentation	108
19.8.1.1	hwloc_obj_attr_sprintf	108
19.8.1.2	hwloc_obj_cpuset_sprintf	108
19.8.1.3	hwloc_obj_type_sprintf	109
19.8.1.4	hwloc_obj_type_sscanf	109
19.8.1.5	hwloc_obj_type_string	110
19.9	Consulting and Adding Key-Value Info Attributes	111
19.9.1	Function Documentation	111
19.9.1.1	hwloc_obj_add_info	111
19.9.1.2	hwloc_obj_get_info_by_name	111
19.10	CPU binding	112
19.10.1	Detailed Description	112
19.10.2	Enumeration Type Documentation	113
19.10.2.1	hwloc_cpupbind_flags_t	113
19.10.3	Function Documentation	114
19.10.3.1	hwloc_get_cpupbind	114
19.10.3.2	hwloc_get_last_cpu_location	114
19.10.3.3	hwloc_get_proc_cpupbind	114
19.10.3.4	hwloc_get_proc_last_cpu_location	115
19.10.3.5	hwloc_get_thread_cpupbind	115
19.10.3.6	hwloc_set_cpupbind	115
19.10.3.7	hwloc_set_proc_cpupbind	116
19.10.3.8	hwloc_set_thread_cpupbind	116
19.11	Memory binding	117
19.11.1	Detailed Description	118
19.11.2	Enumeration Type Documentation	119
19.11.2.1	hwloc_membind_flags_t	119
19.11.2.2	hwloc_membind_policy_t	120
19.11.3	Function Documentation	121
19.11.3.1	hwloc_alloc	121
19.11.3.2	hwloc_alloc_membind	121

19.11.3.3	<a href="#">hwloc_alloc_membind_nodeset</a>	121
19.11.3.4	<a href="#">hwloc_alloc_membind_policy</a>	122
19.11.3.5	<a href="#">hwloc_alloc_membind_policy_nodeset</a>	122
19.11.3.6	<a href="#">hwloc_free</a>	122
19.11.3.7	<a href="#">hwloc_get_area_membind</a>	122
19.11.3.8	<a href="#">hwloc_get_area_membind_nodeset</a>	123
19.11.3.9	<a href="#">hwloc_get_area_memlocation</a>	123
19.11.3.10	<a href="#">hwloc_get_membind</a>	124
19.11.3.11	<a href="#">hwloc_get_membind_nodeset</a>	125
19.11.3.12	<a href="#">hwloc_get_proc_membind</a>	125
19.11.3.13	<a href="#">hwloc_get_proc_membind_nodeset</a>	126
19.11.3.14	<a href="#">hwloc_set_area_membind</a>	127
19.11.3.15	<a href="#">hwloc_set_area_membind_nodeset</a>	127
19.11.3.16	<a href="#">hwloc_set_membind</a>	127
19.11.3.17	<a href="#">hwloc_set_membind_nodeset</a>	128
19.11.3.18	<a href="#">hwloc_set_proc_membind</a>	128
19.11.3.19	<a href="#">hwloc_set_proc_membind_nodeset</a>	128
19.12	<a href="#">Modifying a loaded Topology</a>	130
19.12.1	<a href="#">Enumeration Type Documentation</a>	130
19.12.1.1	<a href="#">hwloc_restrict_flags_e</a>	130
19.12.2	<a href="#">Function Documentation</a>	130
19.12.2.1	<a href="#">hwloc_topology_insert_misc_object_by_cpuset</a>	130
19.12.2.2	<a href="#">hwloc_topology_insert_misc_object_by_parent</a>	131
19.12.2.3	<a href="#">hwloc_topology_restrict</a>	131
19.13	<a href="#">Building Custom Topologies</a>	133
19.13.1	<a href="#">Detailed Description</a>	133
19.13.2	<a href="#">Function Documentation</a>	133
19.13.2.1	<a href="#">hwloc_custom_insert_group_object_by_parent</a>	133
19.13.2.2	<a href="#">hwloc_custom_insert_topology</a>	133
19.14	<a href="#">Exporting Topologies to XML</a>	135
19.14.1	<a href="#">Function Documentation</a>	135
19.14.1.1	<a href="#">hwloc_export_obj_userdata</a>	135
19.14.1.2	<a href="#">hwloc_export_obj_userdata_base64</a>	135
19.14.1.3	<a href="#">hwloc_free_xmlbuffer</a>	136

---

19.14.1.4	<a href="#">hwloc_topology_export_xml</a>	136
19.14.1.5	<a href="#">hwloc_topology_export_xmlbuffer</a>	136
19.14.1.6	<a href="#">hwloc_topology_set_userdata_export_callback</a>	137
19.14.1.7	<a href="#">hwloc_topology_set_userdata_import_callback</a>	137
19.15	Exporting Topologies to Synthetic	139
19.15.1	Enumeration Type Documentation	139
19.15.1.1	<a href="#">hwloc_topology_export_synthetic_flags_e</a>	139
19.15.2	Function Documentation	139
19.15.2.1	<a href="#">hwloc_topology_export_synthetic</a>	139
19.16	Finding Objects inside a CPU set	141
19.16.1	Function Documentation	141
19.16.1.1	<a href="#">hwloc_get_first_largest_obj_inside_cpuset</a>	141
19.16.1.2	<a href="#">hwloc_get_largest_objs_inside_cpuset</a>	141
19.16.1.3	<a href="#">hwloc_get_nbobjs_inside_cpuset_by_depth</a>	142
19.16.1.4	<a href="#">hwloc_get_nbobjs_inside_cpuset_by_type</a>	142
19.16.1.5	<a href="#">hwloc_get_next_obj_inside_cpuset_by_depth</a>	142
19.16.1.6	<a href="#">hwloc_get_next_obj_inside_cpuset_by_type</a>	143
19.16.1.7	<a href="#">hwloc_get_obj_index_inside_cpuset</a>	143
19.16.1.8	<a href="#">hwloc_get_obj_inside_cpuset_by_depth</a>	143
19.16.1.9	<a href="#">hwloc_get_obj_inside_cpuset_by_type</a>	144
19.17	Finding Objects covering at least CPU set	145
19.17.1	Function Documentation	145
19.17.1.1	<a href="#">hwloc_get_child_covering_cpuset</a>	145
19.17.1.2	<a href="#">hwloc_get_next_obj_covering_cpuset_by_depth</a>	145
19.17.1.3	<a href="#">hwloc_get_next_obj_covering_cpuset_by_type</a>	146
19.17.1.4	<a href="#">hwloc_get_obj_covering_cpuset</a>	146
19.18	Looking at Ancestor and Child Objects	147
19.18.1	Detailed Description	147
19.18.2	Function Documentation	147
19.18.2.1	<a href="#">hwloc_get_ancestor_obj_by_depth</a>	147
19.18.2.2	<a href="#">hwloc_get_ancestor_obj_by_type</a>	147
19.18.2.3	<a href="#">hwloc_get_common_ancestor_obj</a>	147
19.18.2.4	<a href="#">hwloc_get_next_child</a>	147
19.18.2.5	<a href="#">hwloc_obj_is_in_subtree</a>	148

---

19.19	Looking at Cache Objects . . . . .	149
19.19.1	Function Documentation . . . . .	149
19.19.1.1	hwloc_get_cache_covering_cpuset . . . . .	149
19.19.1.2	hwloc_get_cache_type_depth . . . . .	149
19.19.1.3	hwloc_get_shared_cache_covering_obj . . . . .	150
19.20	Finding objects, miscellaneous helpers . . . . .	151
19.20.1	Detailed Description . . . . .	151
19.20.2	Function Documentation . . . . .	151
19.20.2.1	hwloc_get_closest_objs . . . . .	151
19.20.2.2	hwloc_get_numanode_obj_by_os_index . . . . .	151
19.20.2.3	hwloc_get_obj_below_array_by_type . . . . .	152
19.20.2.4	hwloc_get_obj_below_by_type . . . . .	152
19.20.2.5	hwloc_get_pu_obj_by_os_index . . . . .	152
19.21	Distributing items over a topology . . . . .	153
19.21.1	Enumeration Type Documentation . . . . .	153
19.21.1.1	hwloc_distrib_flags_e . . . . .	153
19.21.2	Function Documentation . . . . .	153
19.21.2.1	hwloc_distrib . . . . .	153
19.22	CPU and node sets of entire topologies . . . . .	155
19.22.1	Function Documentation . . . . .	155
19.22.1.1	hwloc_topology_get_allowed_cpuset . . . . .	155
19.22.1.2	hwloc_topology_get_allowed_nodeset . . . . .	155
19.22.1.3	hwloc_topology_get_complete_cpuset . . . . .	156
19.22.1.4	hwloc_topology_get_complete_nodeset . . . . .	156
19.22.1.5	hwloc_topology_get_online_cpuset . . . . .	156
19.22.1.6	hwloc_topology_get_topology_cpuset . . . . .	157
19.22.1.7	hwloc_topology_get_topology_nodeset . . . . .	157
19.23	Converting between CPU sets and node sets . . . . .	158
19.23.1	Detailed Description . . . . .	158
19.23.2	Function Documentation . . . . .	158
19.23.2.1	hwloc_cpuset_from_nodeset . . . . .	158
19.23.2.2	hwloc_cpuset_from_nodeset_strict . . . . .	158
19.23.2.3	hwloc_cpuset_to_nodeset . . . . .	159
19.23.2.4	hwloc_cpuset_to_nodeset_strict . . . . .	159

---

19.24	Manipulating Distances	160
19.24.1	Function Documentation	160
19.24.1.1	<a href="#">hwloc_get_distance_matrix_covering_obj_by_depth</a>	160
19.24.1.2	<a href="#">hwloc_get_latency</a>	160
19.24.1.3	<a href="#">hwloc_get_whole_distance_matrix_by_depth</a>	161
19.24.1.4	<a href="#">hwloc_get_whole_distance_matrix_by_type</a>	161
19.25	Finding I/O objects	163
19.25.1	Function Documentation	163
19.25.1.1	<a href="#">hwloc_bridge_covers_pcibus</a>	163
19.25.1.2	<a href="#">hwloc_get_hostbridge_by_pcibus</a>	163
19.25.1.3	<a href="#">hwloc_get_next_bridge</a>	163
19.25.1.4	<a href="#">hwloc_get_next_osdev</a>	163
19.25.1.5	<a href="#">hwloc_get_next_pcidev</a>	164
19.25.1.6	<a href="#">hwloc_get_non_io_ancestor_obj</a>	164
19.25.1.7	<a href="#">hwloc_get_pcidev_by_busid</a>	164
19.25.1.8	<a href="#">hwloc_get_pcidev_by_busidstring</a>	164
19.26	The bitmap API	165
19.26.1	Detailed Description	166
19.26.2	Define Documentation	167
19.26.2.1	<a href="#">hwloc_bitmap_foreach_begin</a>	167
19.26.2.2	<a href="#">hwloc_bitmap_foreach_end</a>	167
19.26.3	Typedef Documentation	167
19.26.3.1	<a href="#">hwloc_bitmap_t</a>	167
19.26.3.2	<a href="#">hwloc_const_bitmap_t</a>	167
19.26.4	Function Documentation	168
19.26.4.1	<a href="#">hwloc_bitmap_allbut</a>	168
19.26.4.2	<a href="#">hwloc_bitmap_alloc</a>	168
19.26.4.3	<a href="#">hwloc_bitmap_alloc_full</a>	168
19.26.4.4	<a href="#">hwloc_bitmap_and</a>	168
19.26.4.5	<a href="#">hwloc_bitmap_andnot</a>	168
19.26.4.6	<a href="#">hwloc_bitmap_asprintf</a>	168
19.26.4.7	<a href="#">hwloc_bitmap_clr</a>	168
19.26.4.8	<a href="#">hwloc_bitmap_clr_range</a>	169
19.26.4.9	<a href="#">hwloc_bitmap_compare</a>	169

---

19.26.4.10hwloc_bitmap_compare_first . . . . .	169
19.26.4.11hwloc_bitmap_copy . . . . .	170
19.26.4.12hwloc_bitmap_dup . . . . .	170
19.26.4.13hwloc_bitmap_fill . . . . .	170
19.26.4.14hwloc_bitmap_first . . . . .	170
19.26.4.15hwloc_bitmap_free . . . . .	170
19.26.4.16hwloc_bitmap_from_ith_ulong . . . . .	170
19.26.4.17hwloc_bitmap_from_ulong . . . . .	170
19.26.4.18hwloc_bitmap_intersects . . . . .	170
19.26.4.19hwloc_bitmap_isequal . . . . .	171
19.26.4.20hwloc_bitmap_isfull . . . . .	171
19.26.4.21hwloc_bitmap_isincluded . . . . .	171
19.26.4.22hwloc_bitmap_isset . . . . .	171
19.26.4.23hwloc_bitmap_iszero . . . . .	172
19.26.4.24hwloc_bitmap_last . . . . .	172
19.26.4.25hwloc_bitmap_list_asprintf . . . . .	172
19.26.4.26hwloc_bitmap_list_snprintf . . . . .	172
19.26.4.27hwloc_bitmap_list_sscanf . . . . .	172
19.26.4.28hwloc_bitmap_next . . . . .	173
19.26.4.29hwloc_bitmap_not . . . . .	173
19.26.4.30hwloc_bitmap_only . . . . .	173
19.26.4.31hwloc_bitmap_or . . . . .	173
19.26.4.32hwloc_bitmap_set . . . . .	173
19.26.4.33hwloc_bitmap_set_ith_ulong . . . . .	173
19.26.4.34hwloc_bitmap_set_range . . . . .	173
19.26.4.35hwloc_bitmap_singlify . . . . .	174
19.26.4.36hwloc_bitmap_snprintf . . . . .	174
19.26.4.37hwloc_bitmap_sscanf . . . . .	174
19.26.4.38hwloc_bitmap_taskset_asprintf . . . . .	174
19.26.4.39hwloc_bitmap_taskset_snprintf . . . . .	175
19.26.4.40hwloc_bitmap_taskset_sscanf . . . . .	175
19.26.4.41hwloc_bitmap_to_ith_ulong . . . . .	175
19.26.4.42hwloc_bitmap_to_ulong . . . . .	175
19.26.4.43hwloc_bitmap_weight . . . . .	175

---

19.26.4.4	<a href="#">hwloc_bitmap_xor</a>	175
19.26.4.5	<a href="#">hwloc_bitmap_zero</a>	176
19.27	<a href="#">Linux-specific helpers</a>	177
19.27.1	<a href="#">Detailed Description</a>	177
19.27.2	<a href="#">Function Documentation</a>	177
19.27.2.1	<a href="#">hwloc_linux_get_tid_cpupbind</a>	177
19.27.2.2	<a href="#">hwloc_linux_get_tid_last_cpu_location</a>	177
19.27.2.3	<a href="#">hwloc_linux_parse_cpumap_file</a>	177
19.27.2.4	<a href="#">hwloc_linux_set_tid_cpupbind</a>	178
19.28	<a href="#">Interoperability with Linux libnuma unsigned long masks</a>	179
19.28.1	<a href="#">Detailed Description</a>	179
19.28.2	<a href="#">Function Documentation</a>	179
19.28.2.1	<a href="#">hwloc_cpuset_from_linux_libnuma_ulongs</a>	179
19.28.2.2	<a href="#">hwloc_cpuset_to_linux_libnuma_ulongs</a>	180
19.28.2.3	<a href="#">hwloc_nodeset_from_linux_libnuma_ulongs</a>	180
19.28.2.4	<a href="#">hwloc_nodeset_to_linux_libnuma_ulongs</a>	180
19.29	<a href="#">Interoperability with Linux libnuma bitmask</a>	181
19.29.1	<a href="#">Detailed Description</a>	181
19.29.2	<a href="#">Function Documentation</a>	181
19.29.2.1	<a href="#">hwloc_cpuset_from_linux_libnuma_bitmask</a>	181
19.29.2.2	<a href="#">hwloc_cpuset_to_linux_libnuma_bitmask</a>	181
19.29.2.3	<a href="#">hwloc_nodeset_from_linux_libnuma_bitmask</a>	182
19.29.2.4	<a href="#">hwloc_nodeset_to_linux_libnuma_bitmask</a>	182
19.30	<a href="#">Interoperability with glibc sched affinity</a>	183
19.30.1	<a href="#">Detailed Description</a>	183
19.30.2	<a href="#">Function Documentation</a>	183
19.30.2.1	<a href="#">hwloc_cpuset_from_glibc_sched_affinity</a>	183
19.30.2.2	<a href="#">hwloc_cpuset_to_glibc_sched_affinity</a>	183
19.31	<a href="#">Interoperability with OpenCL</a>	184
19.31.1	<a href="#">Detailed Description</a>	184
19.31.2	<a href="#">Function Documentation</a>	184
19.31.2.1	<a href="#">hwloc_opencl_get_device_cpuset</a>	184
19.31.2.2	<a href="#">hwloc_opencl_get_device_osdev</a>	184
19.31.2.3	<a href="#">hwloc_opencl_get_device_osdev_by_index</a>	185

---

19.32 Interoperability with the CUDA Driver API . . . . .	186
19.32.1 Detailed Description . . . . .	186
19.32.2 Function Documentation . . . . .	186
19.32.2.1 hwloc_cuda_get_device_cpuset . . . . .	186
19.32.2.2 hwloc_cuda_get_device_osdev . . . . .	186
19.32.2.3 hwloc_cuda_get_device_osdev_by_index . . . . .	187
19.32.2.4 hwloc_cuda_get_device_pci_ids . . . . .	187
19.32.2.5 hwloc_cuda_get_device_pcidev . . . . .	187
19.33 Interoperability with the CUDA Runtime API . . . . .	188
19.33.1 Detailed Description . . . . .	188
19.33.2 Function Documentation . . . . .	188
19.33.2.1 hwloc_cudart_get_device_cpuset . . . . .	188
19.33.2.2 hwloc_cudart_get_device_osdev_by_index . . . . .	188
19.33.2.3 hwloc_cudart_get_device_pci_ids . . . . .	189
19.33.2.4 hwloc_cudart_get_device_pcidev . . . . .	189
19.34 Interoperability with the NVIDIA Management Library . . . . .	190
19.34.1 Detailed Description . . . . .	190
19.34.2 Function Documentation . . . . .	190
19.34.2.1 hwloc_nvml_get_device_cpuset . . . . .	190
19.34.2.2 hwloc_nvml_get_device_osdev . . . . .	190
19.34.2.3 hwloc_nvml_get_device_osdev_by_index . . . . .	191
19.35 Interoperability with OpenGL displays . . . . .	192
19.35.1 Detailed Description . . . . .	192
19.35.2 Function Documentation . . . . .	192
19.35.2.1 hwloc_gl_get_display_by_osdev . . . . .	192
19.35.2.2 hwloc_gl_get_display_osdev_by_name . . . . .	192
19.35.2.3 hwloc_gl_get_display_osdev_by_port_device . . . . .	193
19.36 Interoperability with Intel Xeon Phi (MIC) . . . . .	194
19.36.1 Detailed Description . . . . .	194
19.36.2 Function Documentation . . . . .	194
19.36.2.1 hwloc_intel_mic_get_device_cpuset . . . . .	194
19.36.2.2 hwloc_intel_mic_get_device_osdev_by_index . . . . .	194
19.37 Interoperability with OpenFabrics . . . . .	195
19.37.1 Detailed Description . . . . .	195



---

19.37.2 Function Documentation	195
19.37.2.1 hwloc_ibv_get_device_cpuset	195
19.37.2.2 hwloc_ibv_get_device_osdev	195
19.37.2.3 hwloc_ibv_get_device_osdev_by_name	196
19.38 Interoperability with Myrinet Express	197
19.38.1 Detailed Description	197
19.38.2 Function Documentation	197
19.38.2.1 hwloc_mx_board_get_device_cpuset	197
19.38.2.2 hwloc_mx_endpoint_get_device_cpuset	197
19.39 Topology differences	198
19.39.1 Detailed Description	199
19.39.2 Typedef Documentation	199
19.39.2.1 hwloc_topology_diff_obj_attr_type_t	199
19.39.2.2 hwloc_topology_diff_t	199
19.39.2.3 hwloc_topology_diff_type_t	199
19.39.3 Enumeration Type Documentation	199
19.39.3.1 hwloc_topology_diff_apply_flags_e	199
19.39.3.2 hwloc_topology_diff_obj_attr_type_e	200
19.39.3.3 hwloc_topology_diff_type_e	200
19.39.4 Function Documentation	200
19.39.4.1 hwloc_topology_diff_apply	200
19.39.4.2 hwloc_topology_diff_build	201
19.39.4.3 hwloc_topology_diff_destroy	201
19.39.4.4 hwloc_topology_diff_export_xml	201
19.39.4.5 hwloc_topology_diff_export_xmlbuffer	202
19.39.4.6 hwloc_topology_diff_load_xml	202
19.39.4.7 hwloc_topology_diff_load_xmlbuffer	202
19.40 Components and Plugins: Discovery components	204
19.40.1 Typedef Documentation	204
19.40.1.1 hwloc_disc_component_type_t	204
19.40.2 Enumeration Type Documentation	204
19.40.2.1 hwloc_disc_component_type_e	204
19.41 Components and Plugins: Discovery backends	205
19.41.1 Enumeration Type Documentation	205

19.41.1.1	<code>hwloc_backend_flag_e</code> . . . . .	205
19.41.2	Function Documentation . . . . .	205
19.41.2.1	<code>hwloc_backend_alloc</code> . . . . .	205
19.41.2.2	<code>hwloc_backend_enable</code> . . . . .	205
19.41.2.3	<code>hwloc_backends_get_obj_cpuset</code> . . . . .	206
19.41.2.4	<code>hwloc_backends_notify_new_object</code> . . . . .	206
19.42	Components and Plugins: Generic components . . . . .	207
19.42.1	Typedef Documentation . . . . .	207
19.42.1.1	<code>hwloc_component_type_t</code> . . . . .	207
19.42.2	Enumeration Type Documentation . . . . .	207
19.42.2.1	<code>hwloc_component_type_e</code> . . . . .	207
19.43	Components and Plugins: Core functions to be used by components . . . . .	208
19.43.1	Typedef Documentation . . . . .	208
19.43.1.1	<code>hwloc_report_error_t</code> . . . . .	208
19.43.2	Function Documentation . . . . .	208
19.43.2.1	<code>hwloc__insert_object_by_cpuset</code> . . . . .	208
19.43.2.2	<code>hwloc_alloc_setup_object</code> . . . . .	208
19.43.2.3	<code>hwloc_fill_object_sets</code> . . . . .	209
19.43.2.4	<code>hwloc_hide_errors</code> . . . . .	209
19.43.2.5	<code>hwloc_insert_object_by_cpuset</code> . . . . .	209
19.43.2.6	<code>hwloc_insert_object_by_parent</code> . . . . .	209
19.43.2.7	<code>hwloc_plugin_check_namespace</code> . . . . .	209
19.43.2.8	<code>hwloc_report_os_error</code> . . . . .	210
19.44	Components and Plugins: PCI functions to be used by components . . . . .	211
19.44.1	Function Documentation . . . . .	211
19.44.1.1	<code>hwloc_insert_pci_device_list</code> . . . . .	211
19.44.1.2	<code>hwloc_pci_find_cap</code> . . . . .	211
19.44.1.3	<code>hwloc_pci_find_linkspeed</code> . . . . .	211
19.44.1.4	<code>hwloc_pci_prepare_bridge</code> . . . . .	211
<b>20</b>	<b>Data Structure Documentation</b> . . . . .	<b>213</b>
20.1	<code>hwloc_backend</code> Struct Reference . . . . .	213
20.1.1	Detailed Description . . . . .	213
20.1.2	Field Documentation . . . . .	213

20.1.2.1	disable	214
20.1.2.2	discover	214
20.1.2.3	flags	214
20.1.2.4	get_obj_cpuset	214
20.1.2.5	is_custom	214
20.1.2.6	is_thissystem	214
20.1.2.7	notify_new_object	214
20.1.2.8	private_data	214
20.2	hwloc_obj_attr_u::hwloc_bridge_attr_s Struct Reference	215
20.2.1	Detailed Description	215
20.2.2	Field Documentation	215
20.2.2.1	depth	215
20.2.2.2	domain	215
20.2.2.3	downstream	215
20.2.2.4	downstream_type	215
20.2.2.5	pci	215
20.2.2.6	pci	215
20.2.2.7	secondary_bus	215
20.2.2.8	subordinate_bus	215
20.2.2.9	upstream	216
20.2.2.10	upstream_type	216
20.3	hwloc_obj_attr_u::hwloc_cache_attr_s Struct Reference	216
20.3.1	Detailed Description	216
20.3.2	Field Documentation	216
20.3.2.1	associativity	216
20.3.2.2	depth	216
20.3.2.3	linesize	216
20.3.2.4	size	216
20.3.2.5	type	217
20.4	hwloc_component Struct Reference	217
20.4.1	Detailed Description	217
20.4.2	Field Documentation	217
20.4.2.1	abi	217
20.4.2.2	data	217

---

20.4.2.3	finalize	217
20.4.2.4	flags	218
20.4.2.5	init	218
20.4.2.6	type	218
20.5	hwloc_disc_component Struct Reference	218
20.5.1	Detailed Description	219
20.5.2	Field Documentation	219
20.5.2.1	excludes	219
20.5.2.2	instantiate	219
20.5.2.3	name	219
20.5.2.4	priority	219
20.5.2.5	type	220
20.6	hwloc_distances_s Struct Reference	220
20.6.1	Detailed Description	220
20.6.2	Field Documentation	220
20.6.2.1	latency	220
20.6.2.2	latency_base	221
20.6.2.3	latency_max	221
20.6.2.4	nbobjs	221
20.6.2.5	relative_depth	221
20.7	hwloc_obj_attr_u::hwloc_group_attr_s Struct Reference	221
20.7.1	Detailed Description	221
20.7.2	Field Documentation	222
20.7.2.1	depth	222
20.7.2.2	dont_merge	222
20.8	hwloc_obj Struct Reference	222
20.8.1	Detailed Description	223
20.8.2	Field Documentation	223
20.8.2.1	allowed_cpuset	223
20.8.2.2	allowed_nodeset	223
20.8.2.3	arity	223
20.8.2.4	attr	224
20.8.2.5	children	224
20.8.2.6	complete_cpuset	224

---

20.8.2.7	<a href="#">complete_nodeset</a>	224
20.8.2.8	<a href="#">cpuset</a>	224
20.8.2.9	<a href="#">depth</a>	225
20.8.2.10	<a href="#">distances</a>	225
20.8.2.11	<a href="#">distances_count</a>	225
20.8.2.12	<a href="#">first_child</a>	225
20.8.2.13	<a href="#">infos</a>	225
20.8.2.14	<a href="#">infos_count</a>	225
20.8.2.15	<a href="#">last_child</a>	225
20.8.2.16	<a href="#">logical_index</a>	225
20.8.2.17	<a href="#">memory</a>	226
20.8.2.18	<a href="#">name</a>	226
20.8.2.19	<a href="#">next_cousin</a>	226
20.8.2.20	<a href="#">next_sibling</a>	226
20.8.2.21	<a href="#">nodeset</a>	226
20.8.2.22	<a href="#">online_cpuset</a>	226
20.8.2.23	<a href="#">os_index</a>	227
20.8.2.24	<a href="#">os_level</a>	227
20.8.2.25	<a href="#">parent</a>	227
20.8.2.26	<a href="#">prev_cousin</a>	227
20.8.2.27	<a href="#">prev_sibling</a>	227
20.8.2.28	<a href="#">sibling_rank</a>	227
20.8.2.29	<a href="#">symmetric_subtree</a>	227
20.8.2.30	<a href="#">type</a>	227
20.8.2.31	<a href="#">userdata</a>	227
20.9	<a href="#">hwloc_obj_attr_u Union Reference</a>	228
20.9.1	<a href="#">Detailed Description</a>	228
20.9.2	<a href="#">Field Documentation</a>	228
20.9.2.1	<a href="#">bridge</a>	228
20.9.2.2	<a href="#">cache</a>	228
20.9.2.3	<a href="#">group</a>	228
20.9.2.4	<a href="#">osdev</a>	228
20.9.2.5	<a href="#">pcidev</a>	228
20.10	<a href="#">hwloc_obj_info_s Struct Reference</a>	229

---

20.10.1 Detailed Description	229
20.10.2 Field Documentation	229
20.10.2.1 name	229
20.10.2.2 value	229
20.11hwloc_obj_memory_s::hwloc_obj_memory_page_type_s Struct Reference	229
20.11.1 Detailed Description	230
20.11.2 Field Documentation	230
20.11.2.1 count	230
20.11.2.2 size	230
20.12hwloc_obj_memory_s Struct Reference	230
20.12.1 Detailed Description	230
20.12.2 Field Documentation	231
20.12.2.1 local_memory	231
20.12.2.2 page_types	231
20.12.2.3 page_types_len	231
20.12.2.4 total_memory	231
20.13hwloc_obj_attr_u::hwloc_osdev_attr_s Struct Reference	231
20.13.1 Detailed Description	231
20.13.2 Field Documentation	231
20.13.2.1 type	231
20.14hwloc_obj_attr_u::hwloc_pcidev_attr_s Struct Reference	232
20.14.1 Detailed Description	232
20.14.2 Field Documentation	232
20.14.2.1 bus	232
20.14.2.2 class_id	232
20.14.2.3 dev	232
20.14.2.4 device_id	232
20.14.2.5 domain	232
20.14.2.6 func	232
20.14.2.7 linkspeed	232
20.14.2.8 revision	232
20.14.2.9 subdevice_id	232
20.14.2.10subvendor_id	232

---

20.14.2.1	vendor_id	233
20.15	hwloc_topology_cpubind_support Struct Reference	233
20.15.1	Detailed Description	233
20.15.2	Field Documentation	233
20.15.2.1	get_proc_cpubind	233
20.15.2.2	get_proc_last_cpu_location	233
20.15.2.3	get_thisproc_cpubind	234
20.15.2.4	get_thisproc_last_cpu_location	234
20.15.2.5	get_thisthread_cpubind	234
20.15.2.6	get_thisthread_last_cpu_location	234
20.15.2.7	get_thread_cpubind	234
20.15.2.8	set_proc_cpubind	234
20.15.2.9	set_thisproc_cpubind	234
20.15.2.10	set_thisthread_cpubind	234
20.15.2.11	set_thread_cpubind	234
20.16	hwloc_topology_diff_u::hwloc_topology_diff_generic_s Struct Reference	235
20.16.1	Field Documentation	235
20.16.1.1	next	235
20.16.1.2	type	235
20.17	hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_generic_s Struct Reference	235
20.17.1	Field Documentation	235
20.17.1.1	type	235
20.18	hwloc_topology_diff_u::hwloc_topology_diff_obj_attr_s Struct Reference	236
20.18.1	Field Documentation	236
20.18.1.1	diff	236
20.18.1.2	next	236
20.18.1.3	obj_depth	236
20.18.1.4	obj_index	236
20.18.1.5	type	236
20.19	hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_string_s Struct Reference	236
20.19.1	Detailed Description	237
20.19.2	Field Documentation	237

20.19.2.1 name	237
20.19.2.2 newvalue	237
20.19.2.3 oldvalue	237
20.19.2.4 type	237
20.20hwloc_topology_diff_obj_attr_u Union Reference	237
20.20.1 Detailed Description	238
20.20.2 Field Documentation	238
20.20.2.1 generic	238
20.20.2.2 string	238
20.20.2.3 uint64	238
20.21hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_uint64_s Struct Reference	238
20.21.1 Detailed Description	238
20.21.2 Field Documentation	238
20.21.2.1 index	238
20.21.2.2 newvalue	239
20.21.2.3 oldvalue	239
20.21.2.4 type	239
20.22hwloc_topology_diff_u::hwloc_topology_diff_too_complex_s Struct - Reference	239
20.22.1 Field Documentation	239
20.22.1.1 next	239
20.22.1.2 obj_depth	239
20.22.1.3 obj_index	239
20.22.1.4 type	239
20.23hwloc_topology_diff_u Union Reference	240
20.23.1 Detailed Description	240
20.23.2 Field Documentation	240
20.23.2.1 generic	240
20.23.2.2 obj_attr	240
20.23.2.3 too_complex	240
20.24hwloc_topology_discovery_support Struct Reference	240
20.24.1 Detailed Description	241
20.24.2 Field Documentation	241



20.24.2.1 pu . . . . .	241
20.25hwloc_topology_membind_support Struct Reference . . . . .	241
20.25.1 Detailed Description . . . . .	241
20.25.2 Field Documentation . . . . .	242
20.25.2.1 alloc_membind . . . . .	242
20.25.2.2 bind_membind . . . . .	242
20.25.2.3 firsttouch_membind . . . . .	242
20.25.2.4 get_area_membind . . . . .	242
20.25.2.5 get_area_memlocation . . . . .	242
20.25.2.6 get_proc_membind . . . . .	242
20.25.2.7 get_thisproc_membind . . . . .	242
20.25.2.8 get_thisthread_membind . . . . .	242
20.25.2.9 interleave_membind . . . . .	242
20.25.2.10migrate_membind . . . . .	243
20.25.2.11nexttouch_membind . . . . .	243
20.25.2.12replicate_membind . . . . .	243
20.25.2.13set_area_membind . . . . .	243
20.25.2.14set_proc_membind . . . . .	243
20.25.2.15set_thisproc_membind . . . . .	243
20.25.2.16set_thisthread_membind . . . . .	243
20.26hwloc_topology_support Struct Reference . . . . .	243
20.26.1 Detailed Description . . . . .	244
20.26.2 Field Documentation . . . . .	244
20.26.2.1 cpubind . . . . .	244
20.26.2.2 discovery . . . . .	244
20.26.2.3 membind . . . . .	244



# Chapter 1

## Hardware Locality

### Portable abstraction of hierarchical architectures for high-performance computing

See also [Further Reading](#) for links to more sections about hwloc concepts.

#### 1.1 Introduction

hwloc provides command line tools and a C API to obtain the hierarchical map of key computing elements, such as: NUMA memory nodes, shared caches, processor packages, processor cores, processing units (logical processors or "threads") and even I/O devices. hwloc also gathers various attributes such as cache and memory information, and is portable across a variety of different operating systems and platforms. - Additionally it may assemble the topologies of multiple machines into a single one so as to let applications consult the topology of an entire fabric or cluster at once.

hwloc primarily aims at helping high-performance computing (HPC) applications, but is also applicable to any project seeking to exploit code and/or data locality on modern computing platforms.

hwloc supports the following operating systems:

- Linux (including old kernels not having sysfs topology information, with knowledge of cpusets, ScaleMP vSMP and Kerrighed support, etc.) on all supported hardware, including Intel Xeon Phi and NumaScale NumaConnect.
- Solaris (with support for processor sets and logical domains)
- AIX
- Darwin / OS X
- FreeBSD and its variants (such as kFreeBSD/GNU)

- NetBSD
- OSF/1 (a.k.a., Tru64)
- HP-UX
- Microsoft Windows
- IBM BlueGene/Q Compute Node Kernel (CNK)

Since it uses standard Operating System information, hwloc's support is mostly independent from the processor type (x86, powerpc, ...) and just relies on the Operating System support. The main exception is BSD operating systems (NetBSD, FreeBSD, etc.) because they do not provide support topology information, hence hwloc uses an x86-only CPUID-based backend (which can be used for other OSes too, see the [Components and plugins](#) section).

To check whether hwloc works on a particular machine, just try to build it and run `lstopo` or `lstopo-no-graphics`. If some things do not look right (e.g. bogus or missing cache information), see [Questions and Bugs](#) below.

hwloc only reports the number of processors on unsupported operating systems; no topology information is available.

For development and debugging purposes, hwloc also offers the ability to work on "fake" topologies:

- Symmetrical tree of resources generated from a list of level arities, see [Synthetic topologies](#).
- Remote machine simulation through the gathering of topology as XML files, see [Importing and exporting topologies from/to XML files](#).

hwloc can display the topology in a human-readable format, either in graphical mode (`-X11`), or by exporting in one of several different formats, including: plain text, PDF, PNG, and FIG (see [Command-line Examples](#) below). Note that some of the export formats require additional support libraries.

hwloc offers a programming interface for manipulating topologies and objects. It also brings a powerful CPU bitmap API that is used to describe topology objects location on physical/logical processors. See the [Programming Interface](#) below. It may also be used to binding applications onto certain cores or memory nodes. Several utility programs are also provided to ease command-line manipulation of topology objects, binding of processes, and so on.

Perl bindings are available from Bernd Kallies on [CPAN](#).

Python bindings are available from Guy Streeter:

- [Fedora RPM and tarball](#).
- [git tree \(html\)](#).

## 1.2 Installation

hwloc (<http://www.open-mpi.org/projects/hwloc/>) is available under the BSD license. It is hosted as a sub-project of the overall Open MPI project (<http://www.open-mpi.org/>). Note that hwloc does not require any functionality from Open MPI -- it is a wholly separate (and much smaller!) project and code base. It just happens to be hosted as part of the overall Open MPI project.

### 1.2.1 Basic Installation

Installation is the fairly common GNU-based process:

```
shell$ ./configure --prefix=...
shell$ make
shell$ make install
```

The hwloc command-line tool "lstopo" produces human-readable topology maps, as mentioned above. It can also export maps to the "fig" file format. Support for PDF, Postscript, and PNG exporting is provided if the "Cairo" development package (usually `cairo-devel` or `libcairo2-dev`) can be found in "lstopo" when hwloc is configured and build.

### 1.2.2 Installing from a Git clone

Additionally, the code can be directly cloned from Git:

```
shell$ git clone https://github.com/open-mpi/hwloc.git
shell$ cd hwloc
shell$ ./autogen.sh
```

Note that GNU Autoconf  $\geq 2.63$ , Automake  $\geq 1.10$  and Libtool  $\geq 2.2.6$  are required when building from a Git clone.

Nightly development snapshots are available on the web site, they can be configured and built without any need for Git or GNU Autotools.

### 1.2.3 Optional Dependencies

The hwloc core may also benefit from the following development packages:

- `libnuma` for memory binding and migration support on Linux (`numactl-devel` or `libnuma-dev` package).
- `libpciaccess` for full I/O device discovery (`libpciaccess-devel` or `libpciaccess-dev` package). On Linux, PCI discovery may still be performed (without vendor/device names) even if `libpciaccess` cannot be used.
- the AMD OpenCL implementation for OpenCL device discovery.

- the NVIDIA CUDA Toolkit for CUDA device discovery.
- the NVIDIA Management Library (NVML) for NVML device discovery. It is included in CUDA since version 8.0. Older NVML releases were available within the NVIDIA GPU Deployment Kit from <https://developer.nvidia.com/gpu-deployment-kit>.
- the NV-CONTROL X extension library (NVCtrl) for NVIDIA display discovery. The relevant development package is usually `libXNVCtrl-devel` or `libxnvctrl-dev`. It is also available within `nvidia-settings` from <ftp://download.nvidia.com/XFree86/nvidia-settings/> and <https://github.com/NVIDIA/nvidia-settings/>.
- `libxml2` for full XML import/export support (otherwise, the internal minimalistic parser will only be able to import XML files that were exported by the same `hwloc` release). See [Importing and exporting topologies from/to XML files](#) for details. The relevant development package is usually `libxml2-devel` or `libxml2-dev`.
- `libudev` on Linux for easier discovery of OS device information (otherwise `hwloc` will try to manually parse `udev` raw files). The relevant development package is usually `libudev-devel` or `libudev-dev`.
- `libtool`'s `ltdl` library for dynamic plugin loading. The relevant development package is usually `libtool-ltdl-devel` or `libltdl-dev`.

PCI and XML support may be statically built inside the main `hwloc` library, or as separate dynamically-loaded plugins (see the [Components and plugins](#) section).

Note that because of the possibility of GPL taint, the `pciutils` library `libpci` will not be used (remember that `hwloc` is BSD-licensed).

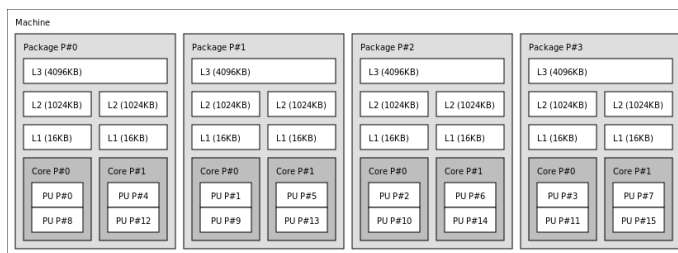
Also note that if you install supplemental libraries in non-standard locations, `hwloc`'s configure script may not be able to find them without some help. You may need to specify additional `CPPFLAGS`, `LDFLAGS`, or `PKG_CONFIG_PATH` values on the configure command line.

For example, if `libpciaccess` was installed into `/opt/pciaccess`, `hwloc`'s configure script may not find it by default. Try adding `PKG_CONFIG_PATH` to the `./configure` command line, like this:

```
./configure PKG_CONFIG_PATH=/opt/pciaccess/lib/pkgconfig ...
```

### 1.3 Command-line Examples

On a 4-package 2-core machine with hyper-threading, the `lstopo` tool may show the following graphical output:

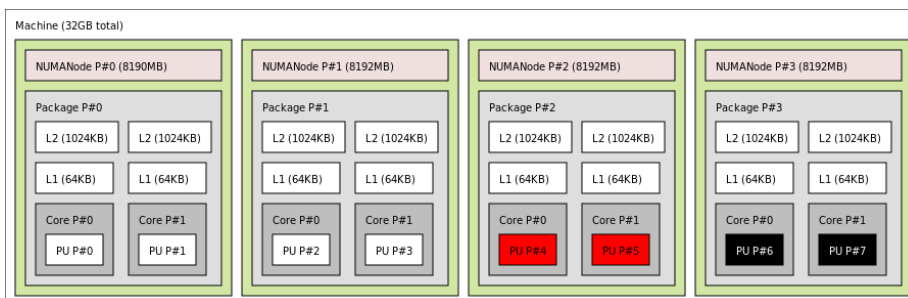


Here's the equivalent output in textual form:

```
Machine
Package L#0 + L3 L#0 (4096KB)
  L2 L#0 (1024KB) + L1 L#0 (16KB) + Core L#0
    PU L#0 (P#0)
    PU L#1 (P#8)
  L2 L#1 (1024KB) + L1 L#1 (16KB) + Core L#1
    PU L#2 (P#4)
    PU L#3 (P#12)
Package L#1 + L3 L#1 (4096KB)
  L2 L#2 (1024KB) + L1 L#2 (16KB) + Core L#2
    PU L#4 (P#1)
    PU L#5 (P#9)
  L2 L#3 (1024KB) + L1 L#3 (16KB) + Core L#3
    PU L#6 (P#5)
    PU L#7 (P#13)
Package L#2 + L3 L#2 (4096KB)
  L2 L#4 (1024KB) + L1 L#4 (16KB) + Core L#4
    PU L#8 (P#2)
    PU L#9 (P#10)
  L2 L#5 (1024KB) + L1 L#5 (16KB) + Core L#5
    PU L#10 (P#6)
    PU L#11 (P#14)
Package L#3 + L3 L#3 (4096KB)
  L2 L#6 (1024KB) + L1 L#6 (16KB) + Core L#6
    PU L#12 (P#3)
    PU L#13 (P#11)
  L2 L#7 (1024KB) + L1 L#7 (16KB) + Core L#7
    PU L#14 (P#7)
    PU L#15 (P#15)
```

Note that there is also an equivalent output in XML that is meant for exporting/importing topologies but it is hardly readable to human-beings (see [Importing and exporting topologies from/to XML files](#) for details).

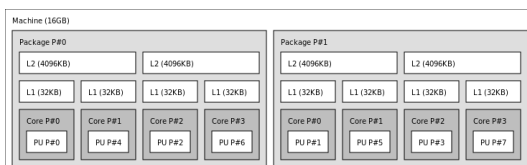
On a 4-package 2-core Opteron NUMA machine (with two core cores disallowed by the administrator, and two offline cores), the `lstopo` tool may show the following graphical output (with `\--whole-system` for displaying disallowed/offline objects):



Here's the equivalent output in textual form:

```
Machine (32GB total)
  NUMANode L#0 (P#0 8190MB) + Package L#0
    L2 L#0 (1024KB) + L1 L#0 (64KB) + Core L#0 + PU L#0 (P#0)
    L2 L#1 (1024KB) + L1 L#1 (64KB) + Core L#1 + PU L#1 (P#1)
  NUMANode L#1 (P#1 8192MB) + Package L#1
    L2 L#2 (1024KB) + L1 L#2 (64KB) + Core L#2 + PU L#2 (P#2)
    L2 L#3 (1024KB) + L1 L#3 (64KB) + Core L#3 + PU L#3 (P#3)
  NUMANode L#2 (P#2 8192MB) + Package L#2
    L2 L#4 (1024KB) + L1 L#4 (64KB) + Core L#4 + PU L#4 (P#4)
    L2 L#5 (1024KB) + L1 L#5 (64KB) + Core L#5 + PU L#5 (P#5)
  NUMANode L#3 (P#3 8192MB) + Package L#3
    L2 L#6 (1024KB) + L1 L#6 (64KB) + Core L#6 + PU L#6 (P#6)
    L2 L#7 (1024KB) + L1 L#7 (64KB) + Core L#7 + PU L#7 (P#7)
```

On a 2-package quad-core Xeon (pre-Nehalem, with 2 dual-core dies into each package):



Here's the same output in textual form:

```
Machine (16GB)
  Package L#0
    L2 L#0 (4096KB)
    L1 L#0 (32KB) + Core L#0 + PU L#0 (P#0)
    L1 L#1 (32KB) + Core L#1 + PU L#1 (P#4)
    L2 L#1 (4096KB)
    L1 L#2 (32KB) + Core L#2 + PU L#2 (P#2)
    L1 L#3 (32KB) + Core L#3 + PU L#3 (P#6)
  Package L#1
    L2 L#2 (4096KB)
    L1 L#4 (32KB) + Core L#4 + PU L#4 (P#1)
    L1 L#5 (32KB) + Core L#5 + PU L#5 (P#5)
    L2 L#3 (4096KB)
    L1 L#6 (32KB) + Core L#6 + PU L#6 (P#3)
    L1 L#7 (32KB) + Core L#7 + PU L#7 (P#7)
```



## 1.4 Programming Interface

The basic interface is available in [hwloc.h](#). Some higher-level functions are available in [hwloc/helper.h](#) to reduce the need to manually manipulate objects and follow links between them. Documentation for all these is provided later in this document. - Developers may also want to look at [hwloc/inlines.h](#) which contains the actual inline code of some [hwloc.h](#) routines, and at this document, which provides good higher-level topology traversal examples.

To precisely define the vocabulary used by hwloc, a [Terms and Definitions](#) section is available and should probably be read first.

Each hwloc object contains a cpuset describing the list of processing units that it contains. These bitmaps may be used for [CPU binding](#) and [Memory binding](#). hwloc offers an extensive bitmap manipulation interface in [hwloc/bitmap.h](#).

Moreover, hwloc also comes with additional helpers for interoperability with several commonly used environments. See the [Interoperability With Other Software](#) section for details.

The complete API documentation is available in a full set of HTML pages, man pages, and self-contained PDF files (formatted for both both US letter and A4 formats) in the source tarball in [doc/doxygen-doc/](#).

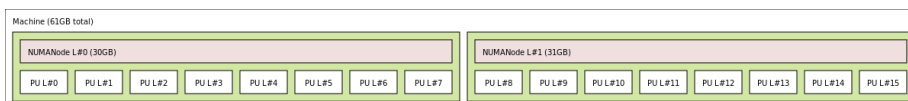
**NOTE:** If you are building the documentation from a Git clone, you will need to have Doxygen and pdflatex installed -- the documentation will be built during the normal "make" process. The documentation is installed during "make install" to `$prefix/share/doc/hwloc/` and your systems default man page tree (under `$prefix`, of course).

### 1.4.1 Portability

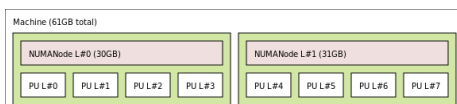
Operating System have varying support for CPU and memory binding, e.g. while some Operating Systems provide interfaces for all kinds of CPU and memory bindings, some others provide only interfaces for a limited number of kinds of CPU and memory binding, and some do not provide any binding interface at all. Hwloc's binding functions would then simply return the ENOSYS error (Function not implemented), meaning that the underlying Operating System does not provide any interface for them. [CPU binding](#) and [Memory binding](#) provide more information on which hwloc binding functions should be preferred because interfaces for them are usually available on the supported Operating Systems.

Similarly, the ability of reporting topology information varies from one platform to another. As shown in [Command-line Examples](#), hwloc can obtain information on a wide variety of hardware topologies. However, some platforms and/or operating system versions will only report a subset of this information. For example, on an PPC64-based system with 8 cores (each with 2 hardware threads) running a default 2.6.18-based kernel from RHEL 5.4, hwloc is only able to glean information about NUMA nodes and processor units (PUs). No information about caches, packages, or cores is available.

Here's the graphical output from `lstopo` on this platform when Simultaneous Multi-Threading (SMT) is enabled:



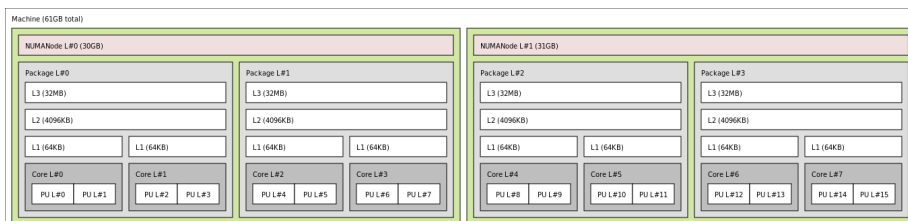
And here's the graphical output from Istopo on this platform when SMT is disabled:



Notice that hwloc only sees half the PUs when SMT is disabled. PU L#6, for example, seems to change location from NUMA node #0 to #1. In reality, no PUs "moved" -- they were simply re-numbered when hwloc only saw half as many (see also Logical index in [Indexes and Sets](#)). Hence, PU L#6 in the SMT-disabled picture probably corresponds to PU L#12 in the SMT-enabled picture.

This same "PUs have disappeared" effect can be seen on other platforms -- even platforms / OSs that provide much more information than the above PPC64 system. This is an unfortunate side-effect of how operating systems report information to hwloc.

Note that upgrading the Linux kernel on the same PPC64 system mentioned above to 2.6.34, hwloc is able to discover all the topology information. The following picture shows the entire topology layout when SMT is enabled:



Developers using the hwloc API or XML output for portable applications should therefore be extremely careful to not make any assumptions about the structure of data that is returned. For example, per the above reported PPC topology, it is not safe to assume that PUs will always be descendants of cores.

Additionally, future hardware may insert new topology elements that are not available in this version of hwloc. Long-lived applications that are meant to span multiple different hardware platforms should also be careful about making structure assumptions. For example, a new element may someday exist between a core and a PU.

## 1.4.2 API Example

The following small C example (available in the source tree as "doc/examples/hwloc-hello.c") prints the topology of the machine and performs some thread and memory binding. More examples are available in the doc/examples/ directory of the source tree.

```
/* Example hwloc API program.
 *
 * See other examples under doc/examples/ in the source tree
```

```

* for more details.
*
* Copyright © 2009–2015 Inria. All rights reserved.
* Copyright © 2009–2011 Université Bordeaux
* Copyright © 2009–2010 Cisco Systems, Inc. All rights reserved.
* See COPYING in top-level directory.
*
* hwloc-hello.c
*/

#include <hwloc.h>
#include <errno.h>
#include <stdio.h>
#include <string.h>

static void print_children(hwloc_topology_t topology, hwloc_obj_t obj,
                          int depth)
{
    char type[32], attr[1024];
    unsigned i;

    hwloc_obj_type_snprintf(type, sizeof(type), obj, 0);
    printf("%*s%s", 2*depth, "", type);
    if (obj->os_index != (unsigned) -1)
        printf("#%u", obj->os_index);
    hwloc_obj_attr_snprintf(attr, sizeof(attr), obj, "", 0);
    if (*attr)
        printf("(%s)", attr);
    printf("\n");
    for (i = 0; i < obj->arity; i++) {
        print_children(topology, obj->children[i], depth + 1);
    }
}

int main(void)
{
    int depth;
    unsigned i, n;
    unsigned long size;
    int levels;
    char string[128];
    int topodepth;
    hwloc_topology_t topology;
    hwloc_cpuset_t cpuset;
    hwloc_obj_t obj;

    /* Allocate and initialize topology object. */
    hwloc_topology_init(&topology);

    /* ... Optionally, put detection configuration here to ignore
       some objects types, define a synthetic topology, etc....

       The default is to detect all the objects of the machine that
       the caller is allowed to access. See Configure Topology
       Detection. */

    /* Perform the topology detection. */
    hwloc_topology_load(topology);

    /* Optionally, get some additional topology information
       in case we need the topology depth later. */
    topodepth = hwloc_topology_get_depth(topology);

    /******
     * First example:
     * Walk the topology with an array style, from level 0 (always
     * the system level) to the lowest level (always the proc level).
     *****/
    for (depth = 0; depth < topodepth; depth++) {
        printf("*** Objects at level %d\n", depth);
        for (i = 0; i < hwloc_get_nobjs_by_depth(topology, depth);
             i++) {
            hwloc_obj_type_snprintf(string, sizeof(string),
                                     hwloc_get_obj_by_depth(topology, depth, i),

```

```

    0);
    printf("Index %u: %s\n", i, string);
}

/*****
 * Second example:
 * Walk the topology with a tree style.
 *****/
printf("*** Printing overall tree\n");
print_children(topology, hwloc_get_root_obj(topology), 0);

/*****
 * Third example:
 * Print the number of packages.
 *****/
depth = hwloc_get_type_depth(topology, HWLOC_OBJ_PACKAGE);
if (depth == HWLOC_TYPE_DEPTH_UNKNOWN) {
    printf("*** The number of packages is unknown\n");
} else {
    printf("*** %u package(s)\n",
           hwloc_get_nobjs_by_depth(topology, depth));
}

/*****
 * Fourth example:
 * Compute the amount of cache that the first logical processor
 * has above it.
 *****/
levels = 0;
size = 0;
for (obj = hwloc_get_obj_by_type(topology, HWLOC_OBJ_PU, 0);
     obj;
     obj = obj->parent)
    if (obj->type == HWLOC_OBJ_CACHE) {
        levels++;
        size += obj->attr->cache.size;
    }
printf("*** Logical processor 0 has %d caches totaling %luKB\n",
       levels, size / 1024);

/*****
 * Fifth example:
 * Bind to only one thread of the last core of the machine.
 *
 * First find out where cores are, or else smaller sets of CPUs if
 * the OS doesn't have the notion of a "core".
 *****/
depth = hwloc_get_type_or_below_depth(topology, HWLOC_OBJ_CORE);

/* Get last core. */
obj = hwloc_get_obj_by_depth(topology, depth,
                             hwloc_get_nobjs_by_depth(topology, depth) - 1);
if (obj) {
    /* Get a copy of its cpuset that we may modify. */
    cpuset = hwloc_bitmap_dup(obj->cpuset);

    /* Get only one logical processor (in case the core is
       SMT/hyper-threaded). */
    hwloc_bitmap_singlify(cpuset);

    /* And try to bind ourself there. */
    if (hwloc_set_cpubind(topology, cpuset, 0)) {
        char *str;
        int error = errno;
        hwloc_bitmap_asprintf(&str, obj->cpuset);
        printf("Couldn't bind to cpuset %s: %s\n", str, strerror(error));
        free(str);
    }

    /* Free our cpuset copy */
    hwloc_bitmap_free(cpuset);
}

```

```

/*****
 * Sixth example:
 * Allocate some memory on the last NUMA node, bind some existing
 * memory to the last NUMA node.
 *****/
/* Get last node. */
n = hwloc_get_nbobjs_by_type(topology, HWLOC_OBJ_NUMANODE);
if (n) {
    void *m;
    size = 1024*1024;

    obj = hwloc_get_obj_by_type(topology, HWLOC_OBJ_NUMANODE, n - 1);
    m = hwloc_alloc_membind_nodeset(topology, size, obj->nodeset,
        HWLOC_MEMBIND_BIND, 0);
    hwloc_free(topology, m, size);

    m = malloc(size);
    hwloc_set_area_membind_nodeset(topology, m, size, obj->nodeset,
        HWLOC_MEMBIND_BIND, 0);
    free(m);
}

/* Destroy topology object. */
hwloc_topology_destroy(topology);

return 0;
}

```

hwloc provides a `pkg-config` executable to obtain relevant compiler and linker flags. For example, it can be used thusly to compile applications that utilize the hwloc library (assuming GNU Make):

```

CFLAGS += $(shell pkg-config --cflags hwloc)
LDLIBS += $(shell pkg-config --libs hwloc)

hwloc-hello: hwloc-hello.c
    $(CC) hwloc-hello.c $(CFLAGS) -o hwloc-hello $(LDLIBS)

```

On a machine with 4GB of RAM and 2 processor packages -- each package of which has two processing cores -- the output from running `hwloc-hello` could be something like the following:

```

shell$ ./hwloc-hello
*** Objects at level 0
Index 0: Machine(3938MB)
*** Objects at level 1
Index 0: Package#0
Index 1: Package#1
*** Objects at level 2
Index 0: Core#0
Index 1: Core#1
Index 2: Core#3
Index 3: Core#2
*** Objects at level 3
Index 0: PU#0
Index 1: PU#1
Index 2: PU#2
Index 3: PU#3
*** Printing overall tree
Machine(3938MB)
  Package#0
    Core#0
      PU#0

```

```
Core#1
  PU#1
Package#1
  Core#3
    PU#2
  Core#2
    PU#3
*** 2 package(s)
shell$
```

## 1.5 Questions and Bugs

Bugs should be reported in the tracker (<https://github.com/open-mpi/hwloc/issues>). Opening a new issue automatically displays lots of hints about how to debug and report issues.

Questions may be sent to the users or developers mailing lists (<http://www.open-mpi.org/community/lists/hwloc.php>).

There is also a #hwloc IRC channel on Freenode (<irc.freenode.net>).

## 1.6 History / Credits

hwloc is the evolution and merger of the libtopology (<http://runtime.bordeaux.inria.fr/libtopology/>) project and the Portable Linux Processor Affinity (PLPA) (<http://www.open-mpi.org/projects/plpa/>) project. - Because of functional and ideological overlap, these two code bases and ideas were merged and released under the name "hwloc" as an Open MPI sub-project.

libtopology was initially developed by the inria Runtime Team-Project (<http://runtime.bordeaux.inria.fr/>) (headed by Raymond Namyst (<http://dept-info.labri.fr/~namyst/>)). PLPA was initially developed by the Open MPI development team as a sub-project. Both are now deprecated in favor of hwloc, which is distributed as an Open MPI sub-project.

## 1.7 Further Reading

The documentation chapters include

- [Terms and Definitions](#)
- [Command-Line Tools](#)
- [Environment Variables](#)
- [CPU and Memory Binding Overview](#)
- [I/O Devices](#)

- [Miscellaneous objects](#)
- [Multi-node Topologies](#)
- [Object attributes](#)
- [Importing and exporting topologies from/to XML files](#)
- [Synthetic topologies](#)
- [Interoperability With Other Software](#)
- [Thread Safety](#)
- [Components and plugins](#)
- [Embedding hwloc in Other Software](#)
- [Frequently Asked Questions](#)

Make sure to have had a look at those too!





## Chapter 2

# Terms and Definitions

### 2.1 Objects

**Object** Interesting kind of part of the system, such as a Core, a Cache, a NUMA memory node, etc. The different types detected by hwloc are detailed in the [hwloc\\_obj\\_type\\_t](#) enumeration.

They are topologically sorted by CPU set into a tree (see [Hierarchy, Tree and Levels](#)).

**Processing Unit (or Logical Processor)** The smallest processing element that can be represented by a hwloc object. It may be a single-core processor, a core of a multicore processor, or a single thread in a SMT processor. hwloc's PU acronym stands for Processing Unit.

"Logical processor" should not be confused with "Logical index of a processor".

**Package** A processor Package is the physical package that usually gets inserted into a socket on the motherboard. It is also often called a physical processor or a CPU even if these names bring confusion with respect to cores and processing units. A processor package usually contains multiple cores (and may also be composed of multiple dies). hwloc Package objects were called Sockets up to hwloc 1.10.

**NUMA Node** An object that contains memory that is directly and byte-accessible to the host processors. It is usually close to some cores (the corresponding objects are descendants of the NUMA node object in the hwloc tree, see [Hierarchy, Tree and Levels](#)) as specified by its CPU set.

### 2.2 Indexes and Sets

**OS or physical index** The index that the operating system (OS) uses to identify the object. This may be completely arbitrary, non-unique, non-contiguous, not representative of logical proximity, and may depend on the BIOS configuration. That is why hwloc almost never uses them, only in the default `lstopo` output (`P#x`) and `cpuset` masks. See also [Should I use logical or physical/OS indexes? and how?](#).

**Logical index** Index to uniquely identify objects of the same type and depth, automatically computed by hwloc according to the topology. It expresses logical proximity in a generic way, i.e. objects which have adjacent logical indexes are adjacent in the topology. That is why hwloc almost always uses it in its API, since it expresses logical proximity. They can be shown (as `L#x`) by `lstopo` thanks to the `-l` option. This index is always linear and in the range `[0, num_objs_same_type_same_level-1]`. Think of it as "cousin rank." The ordering is based on topology first, and then on OS CPU numbers, so it is stable across everything except firmware CPU renumbering. "Logical index" should not be confused with "Logical processor". A "Logical processor" (which in hwloc we rather call "processing unit" to avoid the confusion) has both a physical index (as chosen arbitrarily by BIOS/OS) and a logical index (as computed according to logical proximity by hwloc). See also [Should I use logical or physical/OS indexes? and how?](#).

**CPU set** The set of logical processors (or processing units) logically included in an object (if it makes sense). They are always expressed using physical logical processor numbers (as announced by the OS). They are implemented as the `hwloc_bitmap_t` opaque structure. hwloc CPU sets are just masks, they do *not* have any relation with an operating system actual binding notion like Linux' cpusets.

**Node set** The set of NUMA memory nodes logically included in an object (if it makes sense). They are always expressed using physical node numbers (as announced by the OS). They are implemented with the `hwloc_bitmap_t` opaque structure. as bitmaps.

**Bitmap** A possibly-infinite set of bits used for describing sets of objects such as CPUs (CPU sets) or memory nodes (Node sets). They are implemented with the `hwloc_bitmap_t` opaque structure.

## 2.3 Hierarchy, Tree and Levels

**Parent object** The object logically containing the current object, for example because its CPU set includes the CPU set of the current object.

**Ancestor object** The parent object, or its own parent object, and so on.

**Children object(s)** The object (or objects) contained in the current object because their CPU set is included in the CPU set of the current object.

**Arity** The number of children of an object.

**Sibling objects** Objects which have the same parent. They usually have the same type (and hence are cousins, as well), but they may not if the topology is asymmetric.

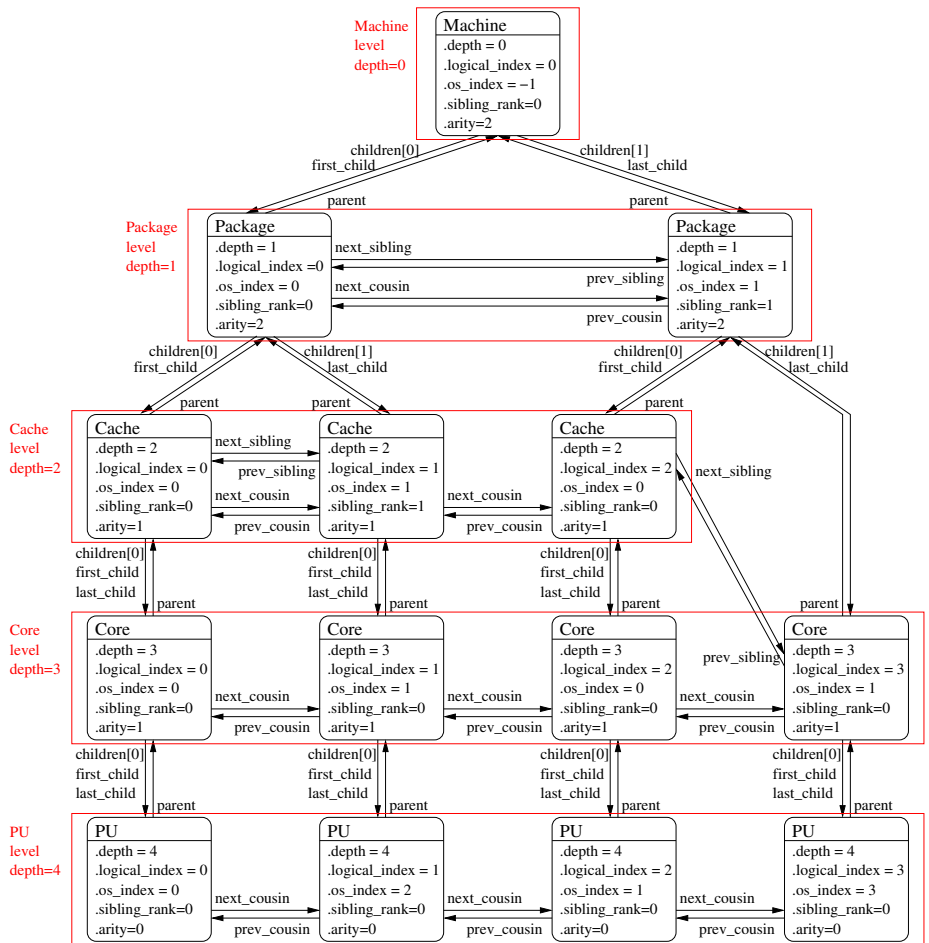
**Sibling rank** Index to uniquely identify objects which have the same parent, and is always in the range `[0, parent_arity)`.

**Cousin objects** Objects of the same type (and depth) as the current object, even if they do not have the same parent.

**Level** Set of objects of the same type and depth. All these objects are cousins.

**Depth** Nesting level in the object tree, starting from the root object. If the topology is symmetric, the depth of a child is equal to the parent depth plus one, and an object depth is also equal to the number of parent/child links between the root object and the given object. If the topology is asymmetric, the difference between some parent and child depths may be larger than one when some intermediate levels (for instance groups) are missing in only some parts of the machine.

The following diagram can help to understand the vocabulary of the relationships by showing the example of a machine with two dual core packages (with no hardware threads); thus, a topology with 5 levels. Each box with rounded corner corresponds to one `hwloc_obj_t`, containing the values of the different integer fields (depth, logical\_index, etc.), and arrows show to which other `hwloc_obj_t` pointers point to (first\_child, parent, etc.). The L2 cache of the last core is intentionally missing to show how asymmetric topologies are handled. See [What happens if my topology is asymmetric?](#) for more information about such strange topologies.



It should be noted that for PU objects, the logical index -- as computed linearly by hwloc -- is not the same as the OS index.



## Chapter 3

# Command-Line Tools

hwloc comes with an extensive C programming interface and several command line utilities. Each of them is fully documented in its own manual page; the following is a summary of the available command line tools.

### 3.1 Istopo and Istopo-no-graphics

Istopo (also known as hwloc-ls) displays the hierarchical topology map of the current system. The output may be graphical, ascii-art or textual, and can also be exported to numerous file formats such as PDF, PNG, XML, and others. Advanced graphical outputs require the "Cairo" development package (usually `cairo-devel` or `libcairo2-dev`).

Istopo and Istopo-no-graphics accept the same command-line options. However graphical outputs are only available in Istopo. Textual outputs (those that do not depend on heavy external libraries such as Cairo) are supported in both Istopo and Istopo-no-graphics.

This command can also display the processes currently bound to a part of the machine (via the `\--ps` option).

Note that Istopo can read XML files and/or alternate chroot filesystems and display topological maps representing those systems (e.g., use Istopo to output an XML file on one system, and then use Istopo to read in that XML file and display it on a different system).

### 3.2 hwloc-bind

hwloc-bind binds processes to specific hardware objects through a flexible syntax. A simple example is binding an executable to specific cores (or packages or bitmaps or ...). The `hwloc-bind(1)` man page provides much more detail on what is possible.

hwloc-bind can also be used to retrieve the current process' binding, or retrieve the last

CPU(s) where a process ran, or operate on memory binding.

Just like `hwloc-calc`, the input locations given to `hwloc-bind` may be either objects or cpusets (bitmaps as reported by `hwloc-calc` or `hwloc-distrib`).

### 3.3 `hwloc-calc`

`hwloc-calc` is `hwloc`'s Swiss Army Knife command-line tool for converting things. The input may be either objects or cpusets (bitmaps as reported by another `hwloc-calc` instance or by `hwloc-distrib`), that may be combined by addition, intersection or subtraction. The output kinds include:

- a cpuset bitmap: This compact opaque representation of objects is useful for shell scripts etc. It may be passed to `hwloc` command-line tools such as `hwloc-calc` or `hwloc-bind`, or to `hwloc` command-line options such as `lstopo \--restrict`.
- the amount of the equivalent `hwloc` objects from a specific type, or the list of their indexes. This is useful for iterating over all similar objects (for instance all cores) within a given part of a platform.
- a hierarchical description of objects, for instance a thread index within a core within a package. This gives a better view of the actual location of an object.

Moreover, input and/or output may be used either physical/OS object indexes or as `hwloc`'s logical object indexes. It eases cooperation with external tools such as `taskset` or `numactl` by exporting `hwloc` specifications into list of processor or NUMA node physical indexes. See also [Should I use logical or physical/OS indexes? and how?](#).

### 3.4 `hwloc-info`

`hwloc-info` dumps information about the given objects, as well as all its specific attributes. It is intended to be used with tools such as `grep` for filtering certain attribute lines. When no object is specified, or when `\--topology` is passed, `hwloc-info` prints a summary of the topology. When `\--support` is passed, `hwloc-info` lists the supported features for the topology.

### 3.5 `hwloc-distrib`

`hwloc-distrib` generates a set of cpuset bitmaps that are uniformly distributed across the machine for the given number of processes. These strings may be used with `hwloc-bind` to run processes to maximize their memory bandwidth by properly distributing them across the machine.

## 3.6 hwloc-ps

hwloc-ps is a tool to display the bindings of processes that are currently running on the local machine. By default, hwloc-ps only lists processes that are bound; unbound process (and Linux kernel threads) are not displayed.

## 3.7 hwloc-distances

hwloc-distances displays all distance matrices attached to the topology. Note that lstopo may also display distance matrices in its verbose textual output. However lstopo only prints matrices that cover the entire topology while hwloc-distances also displays matrices that ignore part of the topology.

## 3.8 hwloc-annotate

hwloc-annotate may modify object attributes such as string information (see [Custom string infos](#) for details) or Misc children objects. It reads an input topology from a XML file and outputs the annotated topology as another XML file.

## 3.9 hwloc-diff, hwloc-patch and hwloc-compress-dir

hwloc-diff computes the difference between two topologies and outputs it to another XML file.

hwloc-patch reads such a difference file and applies to another topology.

hwloc-compress-dir compresses an entire directory of XML files by using hwloc-diff to save the differences between topologies instead of entire topologies.

## 3.10 hwloc-assembler

hwloc-assembler combines several XML topology files into a single multi-node XML topology. It may then be used later as input with [hwloc\\_topology\\_set\\_xml\(\)](#) or with the HWLOC\_XMLFILE environment variable. See [Multi-node Topologies](#) for details.

## 3.11 hwloc-assembler-remote

hwloc-assembler-remote is a frontend to hwloc-assembler. It takes care of contacting the given list of remote hosts (through ssh) and retrieving their topologies as XML before assembling them with hwloc-assembler.

### 3.12 hwloc-dump-hwdata

hwloc-dump-hwdata is a Linux and x86-specific tool that dumps (during boot, privileged) some topology and locality information from raw hardware files (SMBIOS and ACP-I tables) to human-readable and world-accessible files that the hwloc library will later reuse.

Currently only used on Intel Xeon Phi processor platforms. See [Why do I need hwloc-dump-hwdata for memory on Intel Xeon Phi processor?](#).

See `HWLOC_DUMPED_HWDATA_DIR` in [Environment Variables](#) for details about the location of dumped files.

### 3.13 hwloc-gather-topology

hwloc-gather-topology is a Linux-specific tool that saves the relevant topology files of the current machine into a tarball (and the corresponding lstopo outputs).

These files may be used later (possibly offline) for simulating or debugging a machine without actually running on it.



## Chapter 4

# Environment Variables

The behavior of the hwloc library and tools may be tuned thanks to the following environment variables.

**HWLOC\_XMLFILE=/path/to/file.xml** enforces the discovery from the given XML file as if `hwloc_topology_set_xml()` had been called. This file may have been generated earlier with `lstopo file.xml`. For convenience, this backend provides empty binding hooks which just return success. To have hwloc still actually call OS-specific hooks, `HWLOC_THISSYSTEM` should be set 1 in the environment too, to assert that the loaded file is really the underlying system. See also [Importing and exporting topologies from/to XML files](#).

**HWLOC\_XML\_VERBOSE=1**

**HWLOC\_SYNTHETIC\_VERBOSE=1** enables verbose messages in the XML or synthetic topology backends. hwloc XML backends (see [Importing and exporting topologies from/to XML files](#)) can emit some error messages to the error output stream. Enabling these verbose messages within hwloc can be useful for understanding failures to parse input XML topologies. Similarly, enabling verbose messages in the synthetic topology backend can help understand why the description string is invalid. See also [Synthetic topologies](#).

**HWLOC\_FSROOT=/path/to/linux/filesystem-root/** switches to reading the topology from the specified Linux filesystem root instead of the main file-system root, as if `hwloc_topology_set_fsroot()` had been called. Not using the main file-system root causes `hwloc_topology_is_thissystem()` to return 0. For convenience, this backend provides empty binding hooks which just return success. To have hwloc still actually call OS-specific hooks, `HWLOC_THISSYSTEM` should be set 1 in the environment too, to assert that the loaded file is really the underlying system.

**HWLOC\_THISSYSTEM=1** enforces the return value of `hwloc_topology_is_thissystem()`, as if `HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM` was set with `hwloc_topology_set_flags()`. It means that it makes hwloc assume that the selected backend provides the topology for the system on which we are running, even if it is not the OS-specific backend but the XML backend for instance. This means making the binding functions actually call the OS-specific system calls and

really do binding, while the XML backend would otherwise provide empty hooks just returning success. This can be used for efficiency reasons to first detect the topology once, save it to an XML file, and quickly reload it later through the XML backend, but still having binding functions actually do bind. This also enables support for the variable `HWLOC_THISSYSTEM_ALLOWED_RESOURCES`.

**HWLOC\_THISSYSTEM\_ALLOWED\_RESOURCES=1** Get the set of allowed resources from the native operating system even if the topology was loaded from XML or synthetic description, as if `HWLOC_TOPOLOGY_FLAG_THISSYSTEM_ALLOWED_RESOURCES` was set with `hwloc_topology_set_flags()`. This variable requires the topology to match the current system (see the variable `HWLOC_THISSYSTEM`). This is useful when the topology is not loaded directly from the local machine (e.g. for performance reason) and it comes with all resources, but the running process is restricted to only a part of the machine (for instance because of Linux Cgroup/Cpuset).

**HWLOC\_HIDE\_ERRORS=0** enables or disables verbose reporting of errors. The `hwloc` library may issue warnings to the standard error stream when it detects a problem during topology discovery, for instance if the operating system (or user) gives contradictory topology information. Setting this environment variable to 1 removes the actual displaying of these error messages.

**HWLOC\_DONT\_MERGE\_DIE\_GROUPS=0** always keep Group objects that describe CPU Dies. CPU Dies are represented as generic Group objects. By default, Groups are merged with identical children or parent unless they bring some additional structure. Setting this environment variable to 1 ensures these Die Groups will never be merged.

**HWLOC\_GROUPING=1** enables or disables objects grouping based on distances. By default, `hwloc` uses distance matrices between objects (either read from the OS or given by the user) to find groups of close objects. These groups are described by adding intermediate Group objects in the topology. Setting this environment variable to 0 will disable this grouping. This variable supersedes the obsolete `HWLOC_IGNORE_DISTANCES` variable.

**HWLOC\_GROUPING\_ACCURACY=0.05** relaxes distance comparison during grouping. By default, objects may be grouped if their distances form a minimal distance graph. When setting this variable to 0.02, these distances do not have to be strictly equal anymore, they may just be equal with a 2% error. If set to `try` instead of a numerical value, `hwloc` will try to group with perfect accuracy (0, the default), then with 0.01, 0.02, 0.05 and finally 0.1. Numbers given in this environment variable should always use a dot as a decimal mark (for instance 0.01 instead of 0,01).

**HWLOC\_GROUPING\_VERBOSE=0** enables or disables some verbose messages during grouping. If this variable is set to 1, some debug messages will be displayed during distance-based grouping of objects even if debug was not specific at configure time. This is useful when trying to find an interesting distance grouping accuracy.

**HWLOC\_<type>\_DISTANCES=index,...:X\*Y**

**HWLOC\_<type>\_DISTANCES=begin-end:X\*Y\*Z**

**HWLOC\_<type>\_DISTANCES=index,...:distance,...** sets a distance matrix for objects of the given type and physical indexes. The type should be given as its case-sensitive stringified value (e.g. `NUMANode`, `Package`, `Cache`, `Core`, `P-U`). If another distance matrix already exists for the given type, either because the user specified it or because the OS offers it, it will be replaced by the given one.

If the variable value is `none`, the existing distance matrix for the given type is removed. Otherwise, the variable value first consists in a list of physical indexes that may be specified as a comma-separated list (e.g. `0, 2, 4, 1, 3, 5`) or as a range of consecutive indexes (`0-5`). It is followed by a colon and the corresponding distances:

- If `X*Y` is given, X groups of Y close objects are specified.
- If `X*Y*Z` is given, X groups of Y groups of Z close objects are specified.
- Otherwise, the comma-separated list of distances should be given. If N objects are considered, the `i*N+j`-th value gives the distance from the `i`-th object to the `j`-th object. These distance values must use a dot as a decimal separator.

Note that distances are ignored in multi-node topologies.

**HWLOC\_PCI\_<domain>\_<bus>\_LOCALCPUS=<cpuset>** changes the locality of I/O devices behind the specified PCI hostbridge. If no I/O locality information is available or if the BIOS reports incorrect information, it is possible to move a I/O device tree (the entire set of objects behind a host bridge) near a custom set of processors. `domain` and `bus` are the PCI domain and primary bus of the corresponding host bridge.

**HWLOC\_PLUGINS\_PATH=/path/to/hwloc/plugins/...** changes the default search directory for plugins. By default, `$libdir/hwloc` is used. The variable may contain several colon-separated directories.

**HWLOC\_PLUGINS\_VERBOSE=1** displays verbose information about plugins. List which directories are scanned, which files are loaded, and which components are successfully loaded.

**HWLOC\_PLUGINS\_BLACKLIST=filename1,filename2,...** prevents plugins from being loaded if their filename (without path) is listed. Plugin filenames may be found in verbose messages outputted when `HWLOC_PLUGINS_VERBOSE=1`.

**HWLOC\_DUMPED\_HWDATA\_DIR=/path/to/dumped/files/** loads files dumped by `hwloc-dump-hwdata` (on Linux) from the given directory. The default dump/load directory is configured during build based on `--runstatedir`, `--localstatedir`, and `--prefix` options. It usually points to `/var/run/hwloc/` in Linux distribution packages, but it may also point to `$prefix/var/run/hwloc/` when manually installing and only specifying `--prefix`.

**HWLOC\_COMPONENTS=list,of,components** forces a list of components to enable or disable. Enable or disable the given comma-separated list of components (if they do not conflict with each other). Component names prefixed with `-` are disabled. Once the end of the list is reached, `hwloc` falls back to enabling the remaining components (sorted by priority) that do not conflict with the already enabled ones, and unless explicitly disabled in the list. If `stop` is met, the enabling

loop immediately stops, no more component is enabled. If the variable is set to an empty string, no specific component is loaded first, all components are loaded in priority order, this is strictly identical to not specifying any variable. The `xml` component name may be followed by a XML file to load (`xml=file.xml`). The synthetic component may be followed by a basic synthetic topology description (`synthetic=node:2 pu:3`, see [Synthetic topologies](#)). This variable does not take precedence over the application selecting components with functions such as `hwloc_topology_set_xml()`. See [Components and plugins](#) for details.

**HWLOC\_COMPONENTS\_VERBOSE=1** displays verbose information about components. Display messages when components are registered or enabled. This is the recommended way to list the available components with their priority (all of them are *registered* at startup).

**HWLOC\_DEBUG\_VERBOSE=0** disables all verbose messages that are enabled by default when `--enable-debug` is passed to configure.

## Chapter 5

# CPU and Memory Binding Overview

Some operating systems do not systematically provide separate functions for CPU and memory binding. This means that CPU binding functions may have effects on the memory binding policy. Likewise, changing the memory binding policy may change the CPU binding of the current thread. This is often not a problem for applications, so by default hwloc will make use of these functions when they provide better binding support.

If the application does not want the CPU binding to change when changing the memory policy, it needs to use the [HWLOC\\_MEMBIND\\_NOCPUBIND](#) flag to prevent hwloc from using OS functions which would change the CPU binding. Additionally, [HWLOC\\_CPUBIND\\_NOMEMBIND](#) can be passed to CPU binding function to prevent hwloc from using OS functions would change the memory binding policy. Of course, using these flags will reduce hwloc's overall support for binding, so their use is discouraged.

One can avoid using these flags but still closely control both memory and CPU binding by allocating memory, touching each page in the allocated memory, and then changing the CPU binding. The already-really-allocated memory will then be "locked" to physical memory and will not be migrated. Thus, even if the memory binding policy gets changed by the CPU binding order, the already-allocated memory will not change with it. When binding and allocating further memory, the CPU binding should be performed again in case the memory binding altered the previously-selected CPU binding.

Not all operating systems support the notion of a "current" memory binding policy for the current process, but such operating systems often still provide a way to allocate data on a given node set. Conversely, some operating systems support the notion of a "current" memory binding policy and do not permit allocating data on a specific node set without changing the current policy and allocate the data. To provide the most powerful coverage of these facilities, hwloc provides:

- functions that set/get the current memory binding policies (if supported): [hwloc\\_set/get\\_membind\\_\\*](#)() and [hwloc\\_set/get\\_proc\\_membind\(\)](#)
- functions that allocate memory bound to specific node set without changing the current memory binding policy (if supported): [hwloc\\_alloc\\_membind\(\)](#) and [hwloc\\_alloc\\_membind\\_nodeseq\(\)](#).
- helpers which, if needed, change the current memory binding policy of the pro-

cess in order to obtain memory binding: [hwloc\\_alloc\\_membind\\_policy\(\)](#) and [hwloc\\_alloc\\_membind\\_policy\\_nodeset\(\)](#)

An application can thus use the two first sets of functions if it wants to manage separately the global process binding policy and directed allocation, or use the third set of functions if it does not care about the process memory binding policy.

See [CPU binding](#) and [Memory binding](#) for hwloc's API functions regarding CPU and memory binding, respectively. There are some examples under `doc/examples/` in the source tree.

## Chapter 6

# I/O Devices

hwloc usually manipulates processing units and memory but it can also discover I/O devices and report their locality as well. This is useful for placing I/O intensive applications on cores near the I/O devices they use, or for gathering information about all platform components.

### 6.1 Enabling and requirements

I/O discovery is disabled by default (except in `lstopo`) so as not to break legacy application by adding unexpected I/O objects to the topology. It can be enabled by passing flags such as `HWLOC_TOPOLOGY_FLAG_IO_DEVICES` to `hwloc_topology_set_flags()` before loading the topology.

Note that I/O discovery requires significant help from the operating system. The `pciaccess` library (the development package is usually `libpciaccess-devel` or `libpciaccess-dev`) is needed to fully detect PCI devices and bridges. On Linux, PCI discovery may still be performed even if `libpciaccess` cannot be used. But it misses PCI device names. Moreover, some operating systems require privileges for probing PCI devices, see [Does hwloc require privileged access?](#) for details.

The actual locality of I/O devices is only currently detected on Linux. Other operating system will just report I/O devices as being attached to the topology root object.

### 6.2 I/O objects

When I/O discovery is enabled and supported, some additional objects are added to the topology. The corresponding I/O object types are:

- `HWLOC_OBJ_OS_DEVICE` describes an operating-system-specific handle such as the `sda` drive or the `eth0` network interface. See [OS devices](#).
- `HWLOC_OBJ_PCI_DEVICE` and `HWLOC_OBJ_BRIDGE` build up a PCI hierarchy made of devices and bridges. See [PCI devices and bridges](#).

- [HWLOC\\_OBJ\\_MISC](#) describes miscellaneous devices such as memory modules (DIMMs). These are neither PCI nor OS handles, therefore they appear as Misc objects. See [Misc objects added by hwloc](#).

hwloc tries to attach these new objects to normal objects (usually NUMA nodes) to match their actual physical location. For instance, if a I/O Hub is physically connected to a package, the corresponding hwloc bridge object (and its PCI bridges and devices children) is inserted as a child of the corresponding hwloc Package object.

I/O objects also have neither CPU sets nor node sets (NULL pointers) because they are not directly usable by the user applications for binding. Moreover I/O hierarchies may be highly complex (asymmetric trees of bridges). So I/O objects are placed in specific levels with custom depths. Their lists may still be traversed with regular helpers such as [hwloc\\_get\\_next\\_obj\\_by\\_type\(\)](#). However, hwloc offers some dedicated helpers such as [hwloc\\_get\\_next\\_pcidev\(\)](#) and [hwloc\\_get\\_next\\_osdev\(\)](#) for convenience (see [Finding I/O objects](#)).

### 6.3 OS devices

Although each PCI device is uniquely identified by its bus ID (e.g. 0000:01:02.3), a user-space application can hardly find out which PCI device it is actually using. Applications rather use software handles (such as the *eth0* network interface, the *sda* hard drive, or the *mlx4\_0* OpenFabrics HCA). Therefore hwloc tries to add software devices ([HWLOC\\_OBJ\\_OS\\_DEVICE](#), also known as OS devices) below their PCI objects.

hwloc first tries to discover OS devices from the operating system, e.g. *eth0*, *sda* or *mlx4\_0*. However, this ability is currently only available on Linux for some classes of devices.

hwloc then tries to discover software devices through additional I/O components using external libraries. For instance proprietary graphics drivers do not expose any named OS device, but hwloc may still create one OS object per software handle when supported. For instance the `openc1` and `cuda` components may add some *openc10d0* and *cuda0* OS device objects.

Here is a list of OS device objects commonly created by hwloc components when I/O discovery is enabled and supported.

- Hard disks ([HWLOC\\_OBJ\\_OSDEV\\_BLOCK](#))
  - *sda* (Linux component)
- Network interfaces ([HWLOC\\_OBJ\\_OSDEV\\_NETWORK](#))
  - *eth0*, *wlan0*, *ib0* (Linux component)
- OpenFabrics (InfiniBand, Omni-Path, usNIC, etc) HCAs ([HWLOC\\_OBJ\\_OSDEV\\_OPENFABRICS](#))
  - *mlx5\_0*, *hfi1\_0*, *qib0*, *usnic\_0* (Linux component)
- GPUs ([HWLOC\\_OBJ\\_OSDEV\\_GPU](#))



- *nvm10* for the first NVML device (NVML component, using the NVIDIA - Management Library)
- *.0.0* for the first display (GL component, using the NV-CONTROL X extension library, NVCtrl)
- Co-Processors ([HWLOC\\_OBJ\\_OSDEV\\_COPROC](#))
  - *openc10d0* for the first device of the first OpenCL platform, *openc11d3* for the fourth device of the second OpenCL platform (OpenCL component)
  - *cuda0* for the first NVIDIA CUDA device (CUDA component, using the NVIDIA CUDA Library)
  - *mic0* for the first Intel Xeon Phi (MIC) coprocessor (Linux component)
- DMA engine channel ([HWLOC\\_OBJ\\_OSDEV\\_DMA](#))
  - *dma0chan0* (Linux component) when full I/O discovery is enabled ([HWLOC\\_TOPOLOGY\\_FLAG\\_WHOLE\\_IO](#))

When none of the above strategies is supported and enabled, hwloc cannot place any OS object inside PCI objects. Note that some PCI devices may contain multiple software devices (see the example below).

See also [Interoperability With Other Software](#) for managing these devices without considering them as hwloc objects.

## 6.4 PCI devices and bridges

A PCI hierarchy is usually organized as follows: A hostbridge object ([HWLOC\\_OBJ\\_BRIDGE](#) object with upstream type *Host* and downstream type *PCI*) is attached below a normal object (usually the entire machine or a NUMA node). There may be multiple hostbridges in the machine, attached to different places, but all PCI devices are below one of them.

Each hostbridge contains one or several children, either other bridges (usually PCI to PCI) or PCI devices ([HWLOC\\_OBJ\\_PCI\\_DEVICE](#)). The number of bridges between the hostbridge and a PCI device depends on the machine and on the topology flags.

## 6.5 Consulting I/O devices and binding

I/O devices may be consulted by traversing the topology manually (with usual routines such as [hwloc\\_get\\_obj\\_by\\_type\(\)](#)) or by using dedicated helpers (such as [hwloc\\_get\\_pcidev\\_by\\_busid\(\)](#), see [Finding I/O objects](#)).

I/O objects do not actually contain any locality information because their CPU sets and node sets are NULL. Their locality must be retrieved by walking up the object tree (through the `parent` link) until a non-I/O object is found (see [hwloc\\_get\\_non\\_io\\_ancestor\\_obj\(\)](#)). This regular object should have non-NULL CPU sets and node sets

which describe the processing units and memory that are immediately close to the I/O device. For instance the path from a OS device to its locality may go across a PCI device parent, one or several bridges, up to a NUMA node with the same locality.

Command-line tools are also aware of I/O devices. Istopo displays the interesting ones by default (passing `\--no-io` disables it).

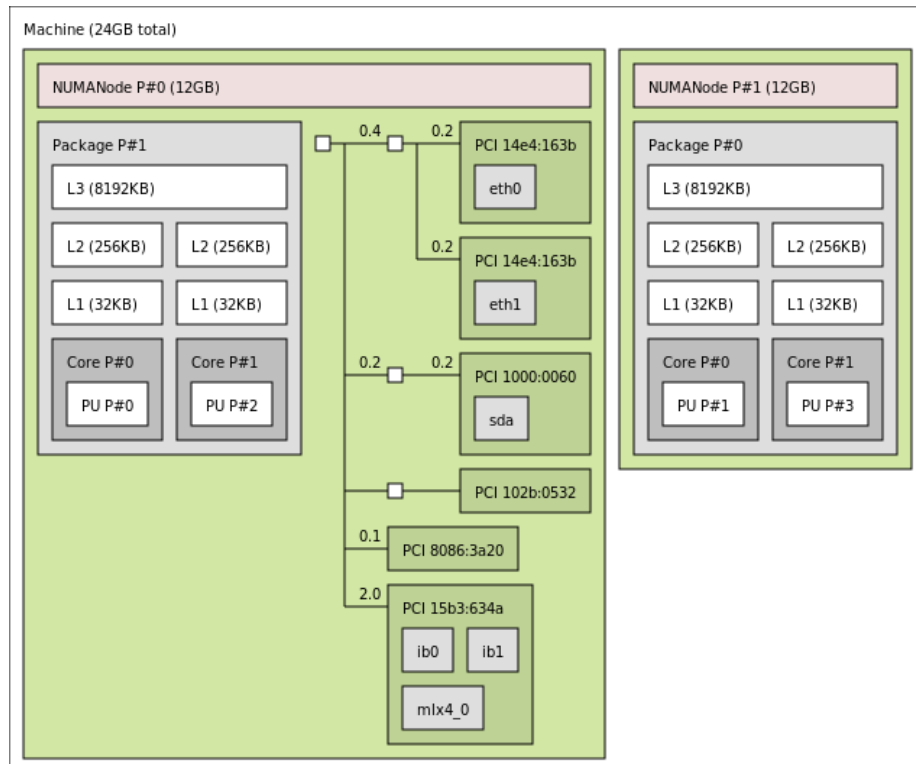
`hwloc-calc` and `hwloc-bind` may manipulate I/O devices specified by PCI bus ID or by OS device name.

- `pci=0000:02:03.0` is replaced by the set of CPUs that are close to the PCI device whose bus ID is given.
- `os=eth0` is replaced by CPUs that are close to the I/O device whose software handle is called `eth0`.

This enables easy binding of I/O-intensive applications near the device they use.

## 6.6 Examples

The following picture shows a dual-package dual-core host whose PCI bus is connected to the first package and NUMA node.



Six interesting PCI devices were discovered. However `hwloc` found some corresponding software devices (`eth0`, `eth1`, `sda`, `mlx4_0`, `ib0`, and `ib1`) for only four of these physical

devices. The other ones (*PCI 102b:0532* and *PCI 8086:3a20*) are an unused IDE controller (no disk attached) and a graphic card (no corresponding software device reported to the user by the operating system).

On the contrary, it should be noted that three different software devices were found for the last PCI device (*PCI 15b3:634a*). Indeed this OpenFabrics HCA PCI device object contains one one OpenFabrics software device (*mlx4\_0*) and two virtual network interface software devices (*ib0* and *ib1*).

Here is the corresponding textual output:

```
Machine (24GB total)
  NUMANode L#0 (P#0 12GB)
    Package L#0 + L3 L#0 (8192KB)
      L2 L#0 (256KB) + L1 L#0 (32KB) + Core L#0 + PU L#0 (P#0)
      L2 L#1 (256KB) + L1 L#1 (32KB) + Core L#1 + PU L#1 (P#2)
    HostBridge L#0
      PCIBridge
        PCI 14e4:163b
          Net L#0 "eth0"
        PCI 14e4:163b
          Net L#1 "eth1"
      PCIBridge
        PCI 1000:0060
          Block L#2 "sda"
      PCIBridge
        PCI 102b:0532
      PCI 8086:3a20
      PCI 15b3:634a
        Net L#3 "ib0"
        Net L#4 "ib1"
        Net L#5 "mlx4_0"
    NUMANode L#1 (P#1 12GB) + Package L#1 + L3 L#1 (8192KB)
      L2 L#2 (256KB) + L1 L#2 (32KB) + Core L#2 + PU L#2 (P#1)
      L2 L#3 (256KB) + L1 L#3 (32KB) + Core L#3 + PU L#3 (P#3)
```



## Chapter 7

# Miscellaneous objects

hwloc topologies may be annotated with Misc objects (of type `HWLOC_OBJ_MISC`) either automatically or by the user. This is an flexible way to annotate topologies with large sets of information since Misc objects may be inserted anywhere in the topology (to annotate specific objects or parts of the topology), even below other Misc objects, and each of them may contain multiple attributes (see also [How do I annotate the topology with private notes?](#)).

These Misc objects may have a `Type` info attribute to replace `Misc` with something else in the `lstopo` output.

### 7.1 Misc objects added by hwloc

hwloc only uses Misc objects when other object types are not sufficient. This currently includes:

- Memory modules (DIMMs), on Linux when privileged and when `dmi-sysfs` is supported by the kernel, and when I/O discovery is enabled. These objects have a `Type` info attribute of value `MemoryModule`. They are currently always attached to the root object. Their attributes describe the DIMM vendor, model, etc. `lstopo -v` displays them as:

```
Misc(MemoryModule) (P#1 Type=MemoryModule DeviceLocation="Bottom-Slot 2(right)"
  BankLocation="BANK 2" Vendor=Elpida SerialNumber=21733667 AssetTag=9876543210
  PartNumber="EBJ81UG8EFU0-GN-F ")
```

- Displaying process binding in `lstopo --top`. These objects have a `Type` info attribute of value `Process` and a name attribute made of their PID and program name. They are attached below the object they are bound to. The textual `lstopo` displays them as:

```
PU L#0 (P#0)
  Misc(Process) 4445 myprogram
```

## 7.2 Annotating topologies with Misc objects

The user may annotate hwloc topologies with its own Misc objects. A Misc object may be inserted anywhere in the topology by specifying its CPU set (using `hwloc_topology_insert_misc_object_by_cpuset()`). Or it may be inserted as a leaf of the topology by specifying its parent (with `hwloc_topology_insert_misc_object_by_parent()`).

## Chapter 8

# Multi-node Topologies

hwloc is usually used for consulting and manipulating single machine topologies. This includes large systems as long as a single instance of the operating system manages the entire system. However it is sometimes desirable to have multiple independent hosts inside the same topology, for instance when applying algorithms to an entire cluster topology. hwloc therefore offers the ability to aggregate multiple host topologies into a single global one.

### Note

**The multinode interface has been removed from hwloc starting with release 2.0.0.**

## 8.1 Multi-node Objects Specificities

A multi-node topology contains several single-node topologies. Those are assembled by making their own root objects (usually Machine object) children of higher objects. These higher objects include at least the root of the global topology (usually a System object). Some intermediate objects may also exist, for instance to represent switches in a large fabric.

There are actually three possible types of objects that have different properties with respect to cpusets, nodesets and binding. Indeed those cpusets and nodesets were designed for execution and memory binding within a single operating system. Binding on another system or across several different systems would be meaningless.

**Local objects** Any object that corresponds to the local machine may be manipulated as usual. Obviously, if the multi-node topology does not contain the local machine topology, no such local object exists.

**Objects from other nodes** Any object that comes from inside another node is represented as usual but its cpusets and nodesets should not be used for binding since binding on another system makes no sense.

**Objects above single nodes** Any object above single-node topologies does not have any `cpuset` or `nodeset` pointer because binding across multiple systems makes no sense. This includes the global root object of a multi-node topology and possibly some intermediate objects between this global root and the local root of single-node topologies.

It is important to keep this in mind before binding using multi-node topologies. To make sure binding on an object is possible, one should first check that its `cpuset` or `nodeset` pointer is not `NULL`. Then, one should check whether the object is indeed local.

To find out which machine a given object corresponds to, one may look at the info attributes of the parent Machine object. The `HostName` info is usually available in Machine objects, it may be retrieved with the following code:

```
hwloc_obj_t machine_obj;
obj = hwloc_get_ancestor_obj_by_type(topology, HWLOC_OBJ_MACHINE, obj);
if (machine_obj)
    return hwloc_obj_get_info_by_name(machine_obj, "HostName");
else
    return NULL;
```

The `hwloc` assembler scripts (see below) also add `AssemblerName` and `AssemblerIndex` info attributes to the Machine objects to identify the corresponding host name and index during assembly.

## 8.2 Assembling topologies with command-line tools

One way to manipulate multinode topologies is to retrieve other nodes' topologies as XML files and combine them as a global XML topology. It may then be loaded with `hwloc_topology_set_xml()` or with the `HWLOC_XMLFILE` environment variable.

The `hwloc-assembler` and `hwloc-assembler-remote` utilities offer the ability to combine XML topologies or remote nodes' topologies (see [Command-Line Tools](#)).

## 8.3 Assembling topologies with the programming interface

The `hwloc` programming interface offers the ability to build multinode topologies using the *custom* interface. A new multinode topology has to be initialized with `hwloc_topology_init()` and then set to custom with `hwloc_topology_set_custom()`. Topologies and objects may then be assembled. Later, the custom topology is finalized as usual with `hwloc_topology_load()`.

A custom topology starts with a single root object of type `System`. It may be modified by inserting a new child object with `hwloc_custom_insert_group_object_by_parent()` or by duplicating another topology with `hwloc_custom_insert_topology()`. Both of these operations require to specify the parent object in the custom topology where the insertion will take place. This parent may be either the root (returned by `hwloc_get_root_obj()`) or an already-inserted object (returned by `hwloc_custom_insert_group_object_by_parent()`).

Ideally, any existing object in the custom topology could be the parent. However, special care should be taken when traversing the topology to find such an object because most



links between objects (children, siblings, cousins) are not setup until `hwloc_topology_load()` is invoked.

## 8.4 Example of assembly with the programming interface

If the topologies of two hosts have been previously gathered in XML files `host1.xml` and `host2.xml`, the global topology may be assembled with the following code.

```
hwloc_topology_t host1, host2, global;

/* initialize global topology */
hwloc_topology_init(&global);
hwloc_topology_set_custom(global);

/* insert host1 entire topology below the global topology root */
hwloc_topology_init(&host1);
hwloc_topology_load(host1);
hwloc_custom_insert_topology(global, hwloc_get_root_obj(global),
                             host1, NULL);
hwloc_topology_destroy(host1);

/* insert host2 entire topology below the global topology root */
hwloc_topology_init(&host2);
hwloc_topology_load(host2);
hwloc_custom_insert_topology(global, hwloc_get_root_obj(global),
                             host2, NULL);
hwloc_topology_destroy(host2);

/* load and play with the global topology */
hwloc_topology_load(global);
...
```

If an intermediate object such as a switch should be inserted above one of the host topologies:

```
...
/* insert a switch object below the global topology root */
hwloc_obj_t sw =
    hwloc_custom_insert_group_object_by_parent(global,
                                               hwloc_get_root_obj(global), 0);

/* insert host2 entire topology below the switch */
hwloc_topology_init(&host2);
hwloc_topology_load(host2);
hwloc_custom_insert_topology(global, sw, host2, NULL);
hwloc_topology_destroy(host2);

/* load and play with the global topology */
hwloc_topology_load(global);
...
```



## Chapter 9

# Object attributes

### 9.1 Normal attributes

hwloc objects have many attributes. The `hwloc_obj` structure contains a common set of attributes that are available for object types, for instance their `type` or `logical_index`.

Each object also contains an `attr` field that, if non NULL, points to a union `hwloc_obj_attr_u` of type-specific attribute structures. For instance, a Cache object `obj` contains cache-specific information in `obj->attr->cache`, such as its size and associativity. See `hwloc_obj_attr_u` for details.

### 9.2 Custom string infos

Aside from the `name` field of each object, hwloc annotates many objects with string attributes that are made of a key and a value. Each object contains a list of such pairs that may be consulted manually (looking at the object `infos` array field) or using the `hwloc_obj_get_info_by_name()`. The user may additionally add new key-value pairs to any object using `hwloc_obj_add_info()` or the `hwloc-annotate` program.

Here is a non-exhaustive list of attributes that may be automatically added by hwloc (with the usual corresponding object in parentheses). Note that these attributes heavily depend on the ability of the operating system to report them. Many of them will therefore be missing on some OS.

**OSName, OSRelease, OSVersion, HostName, Architecture (Machine object)** The operating system name, release, version, the hostname and the architecture name, as reported by the Unix `uname` command.

**Backend (Machine object, topology root object, or specific object added by that backend)** The name of the hwloc backend/component that filled the topology. If several components were combined, multiple Backend keys may exist, with different values, for instance `x86` and `Linux` in the root object and `CUDA` in CUDA OS device objects.

**LinuxCgroup (Machine object)** The name the Linux control group where the calling process is placed.

**SyntheticDescription (topology root object)** The description string that was given to hwloc to build this synthetic topology.

**CPUModel (Package or Machine)** The processor model name. Usually added to Package objects, but can be in Machine instead if hwloc failed to discover any package.

**CPUType (Package)** A Solaris-specific general processor type name, such as "i86pc".

**CPUVendor, CPUModelNumber, CPUFamilyNumber, CPUStepping (Package or Machine)**

The processor vendor name, model number, family number, and stepping number. Currently available for x86 and Xeon Phi processors on most systems, and for ia64 processors on Linux (except CPUStepping). Usually added to Package objects, but can be in Machine instead if hwloc failed to discover any package.

**CPURevision (Package)** A POWER/PowerPC-specific general processor revision number, currently only available on Linux.

**PlatformName, PlatformModel, PlatformVendor, PlatformBoardID, PlatformRevision,**

**SystemVersionRegister, ProcessorVersionRegister (Machine)** Some POWER/PowerPC-specific attributes describing the platform and processor. Currently only available on Linux. Usually added to Package objects, but can be in Machine instead if hwloc failed to discover any package.

**MemoryMode, ClusterMode (topology root object)** Intel Xeon Phi processor configuration modes. Available if hwloc-dump-hwdata was used (see [Why do I need hwloc-dump-hwdata for memory on Intel Xeon Phi processor?](#)) or if hwloc managed to guess them from the NUMA configuration.

The memory mode may be *Cache*, *Flat*, *Hybrid50* (half the MCDRAM is used as a cache) or *Hybrid25* (25% of MCDRAM as cache). The cluster mode may be *Quadrant*, *Hemisphere*, *All2All*, *SNC2* or *SNC4*. See doc/examples/get-knl-modes.c in the source directory for an example of retrieving these attributes.

**Inclusive (Caches)** The inclusiveness of a cache (1 if inclusive, 0 otherwise). - Currently only available on x86 processors.

**SolarisProcessorGroup (Group)** The Solaris kstat processor group name that was used to build this Group object.

**PCIVendor, PCIDevice (PCI devices and bridges)** The vendor and device names of the PCI device.

**PCISlot (PCI devices or Bridges)** The name/number of the physical slot where the device is plugged. If the physical device contains PCI bridges above the actual PCI device, the attribute may be attached to the highest bridge (i.e. the first object that actually appears below the physical slot).

**Vendor, Model, Revision, SerialNumber (Block OS devices)** The vendor and model names, revision, and serial number of a Block OS device.

**LinuxDeviceID (Block OS devices)** The major/minor device number such as 8:0 of Linux device.

**CoProcType (Co-Processor OS devices)** The type of co-processor, for instance "MIC", "CUDA" or "OpenCL".

**GPUNVendor, GPUModel (GPU or Co-Processor OS devices)** The vendor and model names of the GPU device.

**OpenCLDeviceType, OpenCLPlatformIndex,**

**OpenCLPlatformName, OpenCLPlatformDeviceIndex (OpenCL OS devices)** The type of OpenCL device, the OpenCL platform index and name, and the index of the device within the platform.

**OpenCLComputeUnits, OpenCLGlobalMemorySize (OpenCL OS devices)** The number of compute units and global memory size (in kB) of an OpenCL device.

**NVIDIAUUID, NVDIASerial (NVML GPU OS devices)** The UUID and Serial of NVIDIA GPUs.

**CUDAMultiProcessors, CUDACoresPerMP,**

**CUDAGlobalMemorySize, CUDAL2CacheSize, CUDASharedMemorySizePerMP (CUDA OS devices)**

The number of shared multiprocessors, the number of cores per multiprocessor, the global memory size, the (global) L2 cache size, and size of the shared memory in each multiprocessor of a CUDA device. Sizes are in kB.

**MICSerialNumber (MIC coprocessor OS device)** The serial number of an Intel Xeon Phi (MIC) coprocessor. `hwloc` may run either inside the coprocessor itself, or on the host processor. That attribute is set in both cases, so that the exact same coprocessor may be identified from both point of views, even if there are multiple nodes with multiple MICs. When running `hwloc` on the host, each `hwloc` OS device object that corresponds to a Xeon Phi gets such an attribute. When running `hwloc` inside a Xeon Phi coprocessor, the root object of the topology gets this attribute.

**MICFamily, MICSKU, MICActiveCores, MICMemorySize (MIC coprocessor OS device)**

The family, SKU (model), number of active cores, and memory size (in kB) of an Intel Xeon Phi (MIC) coprocessor.

**DMIBoardVendor, DMIBoardName, etc. (Machine object)** DMI hardware information such as the motherboard and chassis models and vendors, the BIOS revision, etc., as reported by Linux under `/sys/class/dmi/id/`.

**Address, Port (Network interface OS devices)** The MAC address and the port number of a software network interface, such as `eth4` on Linux.

**NodeGUID, SysImageGUID, Port1State, Port2LID, Port2LMC, Port3GID1 (OpenFabrics OS devices)**

The node GUID and GUID mask, the state of a port #1 (value is 4 when active), the LID and LID mask count of port #2, and GID #1 of port #3.

**Type** A better type name than the usual one. This may be used to specify where - Groups come from. For instance Linux S/390 *books* appear as Groups of type *Book* (see also [What are these Group objects in my topology?](#)). Block OS devices may have a Type of "Disk", "Tape", "Removable Media Device" or "Other". The Type attribute value is displayed instead of the default object type name in Istopo.

**Vendor, AssetTag, PartNumber, DeviceLocation, BankLocation (MemoryModule Misc objects)** Information about memory modules (DIMMs) extracted from SMBIOS.

**hwlocVersion (topology root, Machine object)** The version number of the hwloc library that was used to generate the topology. If the topology was loaded from XML, this is not the hwloc version that loaded it, but rather the first hwloc instance that exported the topology to XML earlier.

**ProcessName (topology root, Machine object)** The name of the process that contains the hwloc library that was used to generate the topology. If the topology was from XML, this is not the hwloc process that loaded it, but rather the first process that exported the topology to XML earlier.

Here is a non-exhaustive list of user-provided info attributes that have a special meaning:

**IstopoStyle** Enforces the style of an object (background and text colors) in the graphical output of Istopo. See CUSTOM COLORS in the Istopo(1) manpage for details.

## Chapter 10

# Importing and exporting topologies from/to XML files

hwloc offers the ability to export topologies to XML files and reload them later. This is for instance useful for loading topologies faster (see [I do not want hwloc to rediscover my enormous machine topology every time I rerun a process](#)), manipulating other nodes' topology, or avoiding the need for privileged processes (see [Does hwloc require privileged access?](#)).

Topologies may be exported to XML files thanks to `hwloc_topology_export_xml()`, or to a XML memory buffer with `hwloc_topology_export_xmlbuffer()`. The `lstopo` program can also serve as a XML topology export tool.

XML topologies may then be reloaded later with `hwloc_topology_set_xml()` and `hwloc_topology_set_xmlbuffer()`. The `HWLOC_XMLFILE` environment variable also tells hwloc to load the topology from the given XML file.

### Note

Loading XML topologies disables binding because the loaded topology may not correspond to the physical machine that loads it. This behavior may be reverted by asserting that loaded file really matches the underlying system with the `HWLOC_THISSYSTEM` environment variable or the `HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM` topology flag.

The topology flag `HWLOC_TOPOLOGY_FLAG_THISSYSTEM_ALLOWED_RESOURCES` may be used to load a XML topology that contains the entire machine and restrict it to the part that is actually available to the current process (e.g. when Linux Cgroup/Cpuset are used to restrict the set of resources).

hwloc also offers the ability to export/import [Topology differences](#).

XML topology files are not localized. They use a dot as a decimal separator. - Therefore any exported topology can be reloaded on any other machine without requiring to change the locale.

XML exports contain all details about the platform. It means that two very similar nodes still have different XML exports (e.g. some serial numbers or MAC addresses are different). If a less precise exporting/importing is required, one may want to look at [Synthetic topologies](#) instead.

## 10.1 libxml2 and minimalistic XML backends

hwloc offers two backends for importing/exporting XML.

First, it can use the libxml2 library for importing/exporting XML files. It features full XML support, for instance when those files have to be manipulated by non-hwloc software (e.g. a XSLT parser). The libxml2 backend is enabled by default if libxml2 development headers are available (the relevant development package is usually `libxml2-devel` or `libxml2-dev`).

If libxml2 is not available at configure time, or if `--disable-libxml2` is passed, hwloc falls back to a custom backend. Contrary to the aforementioned full XML backend with libxml2, this minimalistic XML backend cannot be guaranteed to work with external programs. It should only be assumed to be compatible with the same hwloc release (even if using the libxml2 backend). Its advantage is however to always be available without requiring any external dependency.

If libxml2 is available but the core hwloc library should not directly depend on it, the libxml2 support may be built as a dynamically-loaded plugin. One should pass `--enable-plugins` to enable plugin support (when supported) and build as plugins all component that support it. Or pass `--enable-plugins=xml_libxml` to only build this libxml2 support as a plugin.

## 10.2 XML import error management

Importing XML files can fail at least because of file access errors, invalid XML syntax or non-hwloc-valid XML contents.

Both backend cannot detect all these errors when the input XML file or buffer is selected (when `hwloc_topology_set_xml()` or `hwloc_topology_set_xmlbuffer()` is called). Some errors such non-hwloc-valid contents can only be detected later when loading the topology with `hwloc_topology_load()`.

It is therefore strongly recommended to check the return value of both `hwloc_topology_set_xml()` (or `hwloc_topology_set_xmlbuffer()`) and `hwloc_topology_load()` to handle all these errors.



# Chapter 11

## Synthetic topologies

hwloc may load fake or remote topologies so as to consult them without having the underlying hardware available. Aside from loading XML topologies, hwloc also enables the building of *synthetic* topologies that are described by a single string listing the arity of each levels.

For instance, lstopo may create a topology made of 2 NUMA nodes, containing a single package each, with one cache above two single-threaded cores:

```
$ lstopo -i "node:2 pack:1 cache:1 core:2 pu:1" -
Machine (2048MB)
  NUMANode L#0 (P#0 1024MB) + Package L#0 + L2 L#0 (4096KB)
    Core L#0 + PU L#0 (P#0)
    Core L#1 + PU L#1 (P#1)
  NUMANode L#1 (P#1 1024MB) + Package L#1 + L2 L#1 (4096KB)
    Core L#2 + PU L#2 (P#2)
    Core L#3 + PU L#3 (P#3)
```

Replacing `-` with `file.xml` in this command line will export this topology to XML as usual.

### Note

Synthetic topologies offer a very basic way to export a topology and reimport it on another machine. It is a lot less precise than XML but may still be enough when only the hierarchy of resources matters.

### 11.1 Synthetic description string

Each item in the description string gives the type of the level and the number of such children under each object of the previous level. That is why the above topology contains 4 cores (2 cores times 2 nodes).

These type names must be written as `machine`, `numanode`, `package`, `core`, `cache`, `pu`, `group`. They do not need to be written case-sensitively, nor entirely (as

long as there is no ambiguity, 2 characters such as `ma` select a Machine level). Type-specific attributes may also be given such as `L2iCache` (`hwloc_obj_type_sscanf()` is used for parsing the type names). Note that I/O and Misc objects are not available.

The root object does not appear in the string. A Machine object is used by default, and a System object replaces it if a Machine level is specified in the string.

Cache level depths are automatically chosen by `hwloc` (only a L2 first, then a L1 under it, then L3 above, then L4 etc.) unless they are specified. Memory and cache sizes are also automatically chosen if needed.

Each item may be followed parentheses containing a list of space-separated attributes. For instance:

- `L2iCache:2(size=32kB)` specifies 2 children of 32kB level-2 instruction caches. The size may be specified in bytes (without any unit suffix) or as TB, GB, MB or kB.
- `NUMANode:3(memory=16MB)` specifies 3 NUMA nodes with 16MB each. - The size may be specified in bytes (without any unit suffix) or as TB, GB, MB or kB.
- `PU:2(indexes=0,2,1,3)` specifies 2 PU children and the full list of OS indexes among the entire set of 4 PU objects.
- `PU:2(indexes=numa:core)` specifies 2 PU children whose OS indexes are interleaved by NUMA node first and then by package.
- Attributes in parentheses at the very beginning of the description apply to the root object.

## 11.2 Loading a synthetic topology

Aside from `lstopo`, the `hwloc` programming interface offers the same ability by passing the synthetic description string to `hwloc_topology_set_synthetic()` before `hwloc_topology_load()`.

Synthetic topologies are created by the `synthetic` component. This component may be enabled by force by setting the `HWLOC_COMPONENTS` environment variable to something such as `synthetic="node:2 core:3 pu:4"`.

Loading a synthetic topology disables binding support since the topology usually does not match the underlying hardware. Binding may be reenabled as usual by setting `HWLOC_THISSYSTEM=1` in the environment or by setting the `HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM` topology flag.

## 11.3 Exporting a topology as a synthetic string

The function `hwloc_topology_export_synthetic()` may export a topology as a synthetic string. It offers a convenient way to quickly describe the contents of a machine. The `lstopo` tool may also perform such an export by forcing the output format.

```
$ lstopo --of synthetic --no-io  
Package:1 Cache:1 Cache:2 Cache:1 Cache:1 Core:1 PU:2
```

The exported string may be passed back to hwloc for recreating another similar topology. The entire tree will be similar, but some attributes such as the processor model will be missing.

Such an export is only possible if the topology is totally symmetric, which means the `symmetric_subtree` field of the root object is set. This usually implies that I/O objects are disabled since attaching I/O busses often cause the topology to become asymmetric. Passing `--no-io` to `lstopo` is therefore often useful to make synthetic export work (as well as not passing any I/O topology flag before exporting with `hwloc_topology_export_synthetic()`).



## Chapter 12

# Interoperability With Other Software

Although hwloc offers its own portable interface, it still may have to interoperate with specific or non-portable libraries that manipulate similar kinds of objects. hwloc therefore offers several specific "helpers" to assist converting between those specific interfaces and hwloc.

Some external libraries may be specific to a particular OS; others may not always be available. The hwloc core therefore generally does not explicitly depend on these types of libraries. However, when a custom application uses or otherwise depends on such a library, it may optionally include the corresponding hwloc helper to extend the hwloc interface with dedicated helpers.

Most of these helpers use structures that are specific to these external libraries and only meaningful on the local machine. If so, the helper requires the input topology to match the current machine. Some helpers also require I/O device discovery to be supported and enabled for the current topology.

**Linux specific features** [hwloc/linux.h](#) offers Linux-specific helpers that utilize some non-portable features of the Linux system, such as binding threads through their thread ID ("tid") or parsing kernel CPU mask files.

**Linux libnuma** [hwloc/linux-libnuma.h](#) provides conversion helpers between hwloc CPU sets and libnuma-specific types, such as bitmasks. It helps you use libnuma memory-binding functions with hwloc CPU sets.

**Glibc** [hwloc/glibc-sched.h](#) offers conversion routines between Glibc and hwloc CPU sets in order to use hwloc with functions such as `sched_getaffinity()` or `pthread_attr_setaffinity_np()`.

**OpenFabrics Verbs** [hwloc/openfabrics-verbs.h](#) helps interoperability with the OpenFabrics Verbs interface. For example, it can return a list of processors near an OpenFabrics device. It may also return the corresponding OS device hwloc object for further information (if I/O device discovery is enabled).

**Myrinet Express** [hwloc/myriexpress.h](#) offers interoperability with the Myrinet Express interface. It can return the list of processors near a Myrinet board managed by

the MX driver. Note that if I/O device discovery is enabled, such boards may also appear as PCI objects in the topology.

**Intel Xeon Phi (MIC)** [hwloc/intel-mic.h](#) helps interoperability with Intel Xeon Phi (MIC) coprocessors by returning the list of processors near these devices. It may also return the corresponding OS device hwloc object for further information (if I/O device discovery is enabled).

**AMD OpenCL** [hwloc/openssl.h](#) enables interoperability with the OpenCL interface. - Only the AMD implementation currently offers locality information. It may return the list of processors near an AMD/ATI GPU given as a `cl_device_id`. It may also return the corresponding OS device hwloc object for further information (if I/O device discovery is enabled).

**NVIDIA CUDA** [hwloc/cuda.h](#) and [hwloc/cudart.h](#) enable interoperability with NVIDIA CUDA Driver and Runtime interfaces. For instance, it may return the list of processors near NVIDIA GPUs. It may also return the corresponding OS device hwloc object for further information (if I/O device discovery is enabled).

**NVIDIA Management Library (NVML)** [hwloc/nvml.h](#) enables interoperability with the NVIDIA NVML interface. It may return the list of processors near a NVIDIA GPU given as a `nvmlDevice_t`. It may also return the corresponding OS device hwloc object for further information (if I/O device discovery is enabled).

**NVIDIA displays** [hwloc/gl.h](#) enables interoperability with NVIDIA displays using the - NV-CONTROL X extension (NVCtrl library). If I/O device discovery is enabled, it may return the OS device hwloc object that corresponds to a display given as a name such as `:0.0` or given as a port/device pair (server/screen).

**Taskset command-line tool** The taskset command-line tool is widely used for binding processes. It manipulates CPU set strings in a format that is slightly different from hwloc's one (it does not divide the string in fixed-size subsets and separates them with commas). To ease interoperability, hwloc offers routines to convert hwloc - CPU sets from/to taskset-specific string format. Most hwloc command-line tools also support the `--taskset` option to manipulate taskset-specific strings.

## Chapter 13

# Thread Safety

Like most libraries that mainly fill data structures, hwloc is not thread safe but rather reentrant: all state is held in a `hwloc_topology_t` instance without mutex protection. That means, for example, that two threads can safely operate on and modify two different `hwloc_topology_t` instances, but they should not simultaneously invoke functions that modify the *same* instance. Similarly, one thread should not modify a `hwloc_topology_t` instance while another thread is reading or traversing it. However, two threads can safely read or traverse the same `hwloc_topology_t` instance concurrently.

When running in multiprocessor environments, be aware that proper thread synchronization and/or memory coherency protection is needed to pass hwloc data (such as `hwloc_topology_t` pointers) from one processor to another (e.g., a mutex, semaphore, or a memory barrier). Note that this is not a hwloc-specific requirement, but it is worth mentioning.

For reference, `hwloc_topology_t` modification operations include (but may not be limited to):

**Creation and destruction** `hwloc_topology_init()`, `hwloc_topology_load()`, `hwloc_topology_destroy()` (see [Topology Creation and Destruction](#)) imply major modifications of the structure, including freeing some objects. No other thread cannot access the topology or any of its objects at the same time.

Also references to objects inside the topology are not valid anymore after these functions return.

**Runtime topology modifications** `hwloc_topology_insert_misc_object_by_*` (see [Modifying a loaded Topology](#)) may modify the topology significantly by adding objects inside the tree, changing the topology depth, etc. `hwloc_topology_restrict()` modifies the topology even more dramatically by removing some objects.

Although references to former objects *may* still be valid after insertion or restriction, it is strongly advised to not rely on any such guarantee and always re-consult the topology to reacquire new instances of objects.

**Locating topologies** `hwloc_topology_ignore*`, `hwloc_topology_set*`

---

(see [Topology Detection Configuration and Query](#)) do not modify the topology directly, but they do modify internal structures describing the behavior of the upcoming invocation of `hwloc_topology_load()`. Hence, all of these functions should not be used concurrently.



## Chapter 14

# Components and plugins

hwloc is organized in components that are responsible for discovering objects. - Depending on the topology configuration, some components will be used, some will be ignored. The usual default is to enable the native operating system component, (e.g. `linux` or `solaris`) and the `pci` miscellaneous component. If available, an architecture-specific component (such as `x86`) may also improve the topology detection.

If a XML topology is loaded, the `xml` discovery component will be used instead of all other components. It internally uses a specific class of components for the actual XML import/export routines (`xml_libxml` and `xml_nolibxml`) but these will not be discussed here (see [libxml2](#) and [minimalistic XML backends](#)).

### 14.1 Components enabled by default

The hwloc core contains a list of components sorted by priority. Each one is enabled as long as it does not conflict with the previously enabled ones. This includes native operating system components, architecture-specific ones, and if available, I/O components such as `pci`.

Usually the native operating system component (when it exists, e.g. `linux` or `aix`) is enabled first. Then hwloc looks for an architecture specific component (e.g. `x86`). Finally there also exist a basic component (`no_os`) that just tries to discover the number of PUs in the system.

Each component discovers as much topology information as possible. Most of them, including most native OS components, do nothing unless the topology is still empty. - Some others, such as `x86` and `pci`, can complete and annotate what other backends found earlier.

Default priorities ensure that clever components are invoked first. Native operating system components have higher priorities, and are therefore invoked first, because they likely offer very detailed topology information. If needed, it will be later extended by architecture-specific information (e.g. from the `x86` component).

If any configuration function such as `hwloc_topology_set_xml()` is used before loading the topology, the corresponding component is enabled first. Then, as usual, hwloc enables any other component (based on priorities) that does not conflict.

Certain components that manage a virtual topology, for instance XML topology import, synthetic topology description, or custom building, conflict with all other components. Therefore, one of them may only be loaded (e.g. with `hwloc_topology_set_xml()`) if no other component is enabled.

The environment variable `HWLOC_COMPONENTS_VERBOSE` may be set to get verbose messages about component registration (including their priority) and enabling.

## 14.2 Selecting which components to use

Once topology configuration functions such as `hwloc_topology_set_custom()` have been taken care of, the priority order of the remaining components may be changed through the `HWLOC_COMPONENTS` environment variable (component names must be separated by commas).

Specifying `x86` in this variable will cause the `x86` component to take precedence over any other component, including the native operating system component. It is therefore loaded first, before hwloc tries to load all remaining non-conflicting components. In this case, `x86` would take care of discovering everything it supports, instead of only completing what the native OS information. This may be useful if the native component is buggy on some platforms.

It is possible to prevent some components from being loaded by prefixing their name with `-` in the list. For instance `x86,-pci` will load the `x86` component, then let hwloc load all the usual components except `pci`.

It is possible to prevent all remaining components from being loaded by placing `stop` in the environment variable. Only the components listed before this keyword will be enabled.

Certain component names (`xml` and `synthetic`) accept an argument (e.g. `xml=file.xml`). These arguments behave exactly as if the corresponding string had been passed to `hwloc_topology_set_xml()` or `hwloc_topology_set_synthetic()`.

## 14.3 Loading components from plugins

Components may optionally be built as plugins so that the hwloc core library does not directly depend on their dependencies (for instance the `libpciaccess` library). Plugin support may be enabled with the `\--enable-plugins` configure option. All components buildable as plugins will then be built as plugins. The configure option may be given a comma-separated list of component names to specify the exact list of components to build as plugins.

Plugins are built as independent dynamic libraries that are installed in `$libdir/hwloc`. All plugins found in this directory are loaded during `topology_init()` (unless

blacklisted in `HWLOC_PLUGINS_BLACKLIST`, see [Environment Variables](#)). A specific list of directories (colon-separated) to scan may be specified in the `HWLOC_PLUGINS_PATH` environment variable.

Note that loading a plugin just means that the corresponding component is registered to the hwloc core. Components are then only enabled if the topology configuration requests it, as explained in the previous sections.

Also note that plugins should carefully be enabled and used when embedding hwloc in another project, see [Embedding hwloc in Other Software](#) for details.

## 14.4 Adding new discovery components and plugins

The types and functions cited below are declared in the `hwloc/plugins.h` header. Components are supposed to only use hwloc public headers (`hwloc.h` and anything under the `include/hwloc` subdirectory) and nothing from the `include/private` subdirectory in the source tree.

### 14.4.1 Basics of discovery components

Each discovery component is defined by a `hwloc_disc_component` structure which contains an `instantiate()` callback. This function is invoked when this component is actually used by a topology. It fills a new `hwloc_backend` structure that usually contains `discover()` and/or `notify_new_object()` callbacks taking care of the actual topology discovery.

#### Note

If two discovery components have the same name, only the highest priority one is actually made available. This offers a way for third-party plugins to override existing components.

### 14.4.2 Registering a new discovery component

Registering components to the hwloc core relies on a `hwloc_component` structure. Its `data` field points to the previously defined `hwloc_disc_component` structure while its `type` should be `HWLOC_COMPONENT_TYPE_DISC`. This structure should be named `hwloc_<name>_component`.

The configure script should be modified to add `<name>` to its `hwloc_components` shell variable so that the component is actually available.

**Note**

The symbol name of the `hwloc_component` structure is independent of the name of the discovery component mentioned in the previous section.

When the component is statically built inside the hwloc library, the symbol `hwloc_<name>_component` is added by configure to the `src/static-components.h`. The core then registers all components listed in this file.

If the new component may be built as a plugin, the configure script should also define the shell variable `hwloc_<name>_component_maybeplugin=1`. When the configure script actually enables the component as a plugin, it will set the variable `hwloc_<name>_component` to `plugin`. The build system may then use this variable to change the way the component is built. It should create a `hwloc_<name>.so` shared object. All these files are loaded in alphabetic order, and the components they contain are registered to the hwloc core.

## 14.5 Existing components and plugins

All components distributed within hwloc are listed below. The list of actually available components may be listed at running with the `HWLOC_COMPONENTS_VERBOSE` environment variable (see [Environment Variables](#)).

**aix, darwin, freebsd, hpux, linux, netbsd, osf, solaris, windows** Each officially supported operating system has its own native component, which is statically built when supported, and which is used by default.

**x86** The x86 architecture (either 32 or 64 bits) has its own component that may complete or replace the previously-found CPU information. It is statically built when supported.

**bgq** This component is specific to IBM BlueGene/Q compute node (running CNK). It is built and enabled by default when `--host=powerpc64-bgq-linux` is passed to configure (see [How do I build hwloc for BlueGene/Q?](#)).

**no\_os** A basic component that just tries to detect the number of processing units in the system. It mostly serves on operating systems that are not natively supported. It is always statically built.

**pci** PCI object discovery uses the external `pciaccess` library (aka `libpciaccess`); see [I/O Devices](#). **It may be built as a plugin.**

**linuxpci** This component can probe PCI devices on Linux without the help of external libraries such as `libpciaccess`. Its priority is lower than the `pci` component because it misses device names.

**openc1** The OpenCL component creates co-processor OS device objects such as `openc10d0` (first device of the first OpenCL platform) or `openc11d3` (fourth device of the second platform). Only the AMD OpenCL implementation currently offers locality information. **It may be built as a plugin.**

- cuda** This component creates co-processor OS device objects such as *cuda0* that correspond to NVIDIA GPUs used with CUDA library. **It may be built as a plugin.**
- nvml** Probing the NVIDIA Management Library creates OS device objects such as *nvml0* that are useful for batch schedulers. It also detects the actual PCIe link bandwidth without depending on power management state and without requiring administrator privileges. **It may be built as a plugin.**
- gl** Probing the NV-CONTROL X extension (NVCtrl library) creates OS device objects such as *:0.0* corresponding to NVIDIA displays. They are useful for graphical applications that need to place computation and/or data near a rendering GPU. **It may be built as a plugin.**
- synthetic** Synthetic topology support (see [Synthetic topologies](#)) is always built statically.
- custom** Custom topology support (see [Multi-node Topologies](#)) is always built statically.
- xml** XML topology import (see [Importing and exporting topologies from/to XML files](#)) is always built statically. It internally uses one of the XML backends (see [libxml2 and minimalistic XML backends](#)).
- **xml\_nolibxml** is a basic and hwloc-specific XML import/export. It is always statically built.
  - **xml\_libxml** relies on the external libxml2 library for providing a feature-complete XML import/export. **It may be built as a plugin.**
- fake** A dummy plugin that does nothing but is used for debugging plugin support.



## Chapter 15

# Embedding hwloc in Other Software

It can be desirable to include hwloc in a larger software package (be sure to check out the LICENSE file) so that users don't have to separately download and install it before installing your software. This can be advantageous to ensure that your software uses a known-tested/good version of hwloc, or for use on systems that do not have hwloc pre-installed.

When used in "embedded" mode, hwloc will:

- not install any header files
- not build any documentation files
- not build or install any executables or tests
- not build `libhwloc.*` -- instead, it will build `libhwloc_embedded.*`

There are two ways to put hwloc into "embedded" mode. The first is directly from the configure command line:

```
shell$ ./configure --enable-embedded-mode ...
```

The second requires that your software project uses the GNU Autoconf / Automake / Libtool tool chain to build your software. If you do this, you can directly integrate hwloc's m4 configure macro into your configure script. You can then invoke hwloc's configuration tests and build setup by calling an m4 macro (see below).

Although hwloc dynamic shared object plugins may be used in embedded mode, the embedder project will have to manually setup `libltdl` in its build system so that hwloc can load its plugins at run time. Also, embedders should be aware of complications that can arise due to public and private linker namespaces (e.g., if the embedder project is loaded into a private namespace and then hwloc tries to dynamically load its plugins, such loading may fail since the hwloc plugins can't find the hwloc symbols they need). The embedder project is **strongly** advised not to use hwloc's dynamically loading plugins / `libltdl` capability.

## 15.1 Using hwloc's M4 Embedding Capabilities

Every project is different, and there are many different ways of integrating hwloc into yours. What follows is *one* example of how to do it.

If your project uses recent versions Autoconf, Automake, and Libtool to build, you can use hwloc's embedded m4 capabilities. We have tested the embedded m4 with projects that use Autoconf 2.65, Automake 1.11.1, and Libtool 2.2.6b. Slightly earlier versions of may also work but are untested. Autoconf versions prior to 2.65 are almost certain to not work.

You can either copy all the config/hwloc\*m4 files from the hwloc source tree to the directory where your project's m4 files reside, or you can tell aclocal to find more m4 files in the embedded hwloc's "config" subdirectory (e.g., add "--lpath/to/embedded/hwloc/config" to your Makefile.am's ACLOCAL\_AMFLAGS).

The following macros can then be used from your configure script (only HWLOC\_SETUP\_CORE *must* be invoked if using the m4 macros):

- `HWLOC_SETUP_CORE(config-dir-prefix, action-upon-success, action-upon-failure, print_banner_or_not)`: Invoke the hwloc configuration tests and setup the hwloc tree to build. The first argument is the prefix to use for AC\_OUTPUT files -- it's where the hwloc tree is located relative to `$top_srcdir`. Hence, if your embedded hwloc is located in the source tree at `contrib/hwloc`, you should pass `[contrib/hwloc]` as the first argument. If `HWLOC_SETUP_CORE` and the rest of `configure` completes successfully, then "make" traversals of the hwloc tree with standard Automake targets (`all`, `clean`, `install`, etc.) should behave as expected. For example, it is safe to list the hwloc directory in the SUBDIRS of a higher-level Makefile.am. The last argument, if not empty, will cause the macro to display an announcement banner that it is starting the hwloc core configuration tests.

`HWLOC_SETUP_CORE` will set the following environment variables and `AC_SUBST` them: `HWLOC_EMBEDDED_CFLAGS`, `HWLOC_EMBEDDED_CPPFLAGS`, and `HWLOC_EMBEDDED_LIBS`. These flags are filled with the values discovered in the hwloc-specific m4 tests, and can be used in your build process as relevant. The `_CFLAGS`, `_CPPFLAGS`, and `_LIBS` variables are necessary to build `libhwloc` (or `libhwloc_embedded`) itself.

`HWLOC_SETUP_CORE` also sets `HWLOC_EMBEDDED_LDADD` environment variable (and `AC_SUBST`s it) to contain the location of the `libhwloc_embedded.la` convenience Libtool archive. It can be used in your build process to link an application or other library against the embedded hwloc library.

**NOTE: If the `HWLOC_SET_SYMBOL_PREFIX` macro is used, it must be invoked *before* `HWLOC_SETUP_CORE`.**

- `HWLOC_BUILD_STANDALONE`: `HWLOC_SETUP_CORE` defaults to building hwloc in an "embedded" mode (described above). If `HWLOC_BUILD_STANDALONE` is invoked *\*before\** `HWLOC_SETUP_CORE`, the embedded definitions will not apply (e.g., `libhwloc.la` will be built, not `libhwloc_embedded.la`).



- `HWLOC_SET_SYMBOL_PREFIX(foo_)`: Tells the hwloc to prefix all of hwloc's types and public symbols with "foo\_"; meaning that function `hwloc_init()` becomes `foo_hwloc_init()`. Enum values are prefixed with an upper-case translation if the prefix supplied; `HWLOC_OBJ_CORE` becomes `FOO_hwloc_OBJ_CORE`. This is recommended behavior if you are including hwloc in middleware -- it is possible that your software will be combined with other software that links to another copy of hwloc. If both uses of hwloc utilize different symbol prefixes, there will be no type/symbol clashes, and everything will compile, link, and run successfully. If you both embed hwloc without changing the symbol prefix and also link against an external hwloc, you may get multiple symbol definitions when linking your final library or application.
- `HWLOC_SETUP_DOCS`, `HWLOC_SETUP_UTILS`, `HWLOC_SETUP_TESTS`: - These three macros only apply when hwloc is built in "standalone" mode (i.e., they should NOT be invoked unless `HWLOC_BUILD_STANDALONE` has already been invoked).
- `HWLOC_DO_AM_CONDITIONALS`: If you embed hwloc in a larger project and build it conditionally with Automake (e.g., if `HWLOC_SETUP_CORE` is invoked conditionally), you must unconditionally invoke `HWLOC_DO_AM_CONDITIONALS` to avoid warnings from Automake (for the cases where hwloc is not selected to be built). This macro is necessary because hwloc uses some `AM_CONDITIONALS` to build itself, and `AM_CONDITIONALS` cannot be defined conditionally. Note that it is safe (but unnecessary) to call `HWLOC_DO_AM_CONDITIONALS` even if `HWLOC_SETUP_CORE` is invoked unconditionally. If you are not using Automake to build hwloc, this macro is unnecessary (and will actually cause errors because it invoked `AM_*` macros that will be undefined).

**NOTE:** When using the `HWLOC_SETUP_CORE` m4 macro, it may be necessary to explicitly invoke `AC_CANONICAL_TARGET` (which requires `config.sub` and `config.guess`) and/or `AC_USE_SYSTEM_EXTENSIONS` macros early in the configure script (e.g., after `AC_INIT` but before `AM_INIT_AUTOMAKE`). See the Autoconf documentation for further information.

Also note that hwloc's top-level `configure.ac` script uses exactly the macros described above to build hwloc in a standalone mode (by default). You may want to examine it for one example of how these macros are used.

## 15.2 Example Embedding hwloc

Here's an example of integrating with a larger project named `sandbox` that already uses Autoconf, Automake, and Libtool to build itself:

```
# First, cd into the sandbox project source tree
shell$ cd sandbox
shell$ cp -r /somewhere/else/hwloc-<version> my-embedded-hwloc
shell$ edit Makefile.am
  1. Add "-Imy-embedded-hwloc/config" to ACLOCAL_AMFLAGS
  2. Add "my-embedded-hwloc" to SUBDIRS
```

```
3. Add "$(HWLOC_EMBEDDED_LDADD)" and "$(HWLOC_EMBEDDED_LIBS)" to
   sandbox's executable's LDADD line. The former is the name of the
   Libtool convenience library that hwloc will generate. The latter
   is any dependent support libraries that may be needed by
   $(HWLOC_EMBEDDED_LDADD).
4. Add "$(HWLOC_EMBEDDED_CFLAGS)" to AM_CFLAGS
5. Add "$(HWLOC_EMBEDDED_CPPFLAGS)" to AM_CPPFLAGS
shell$ edit configure.ac
1. Add "HWLOC_SET_SYMBOL_PREFIX(sandbox_hwloc_)" line
2. Add "HWLOC_SETUP_CORE([my-embedded-hwloc], [happy=yes], [happy=no])" line
3. Add error checking for happy=no case
shell$ edit sandbox.c
1. Add #include <hwloc.h>
2. Add calls to sandbox_hwloc_init() and other hwloc API functions
```

Now you can bootstrap, configure, build, and run the sandbox as normal -- all calls to "sandbox\_hwloc\_\*" will use the embedded hwloc rather than any system-provided copy of hwloc.

# Chapter 16

## Frequently Asked Questions

### 16.1 Concepts

#### 16.1.1 I only need binding, why should I use hwloc ?

hwloc is its portable API that works on a variety of operating systems. It supports binding of threads, processes and memory buffers (see [CPU binding](#) and [Memory binding](#)). - Even if some features are not supported on some systems, using hwloc is much easier than reimplementing your own portability layer.

Moreover, hwloc provides knowledge of cores and hardware threads. It offers easy ways to bind tasks to individual hardware threads, or to entire multithreaded cores, etc. See [How may I ignore symmetric multithreading, hyper-threading, etc. in hwloc?](#). Most alternative software for binding do not even know whether each core is single-threaded, multithreaded or hyper-threaded. They would bind to individual threads without any way to know whether multiple tasks are in the same physical core.

However, using hwloc comes with an overhead since a topology must be loaded before gathering information and binding tasks or memory. This overhead may be reduced by filtering useless information out of the topology. For instance the following code builds a topology that may only contain Cores, hardware threads (PUs), and NUMA nodes (Packages, Caches and Group objects are ignored).

```
hwloc_topology_t topology;
hwloc_topology_init(&topology);
hwloc_topology_ignore_type(topology, HWLOC_OBJ_PACKAGE);
hwloc_topology_ignore_type(topology, HWLOC_OBJ_CACHE);
hwloc_topology_ignore_type(topology, HWLOC_OBJ_GROUP);
hwloc_topology_load(topology);
```

#### 16.1.2 Should I use logical or physical/OS indexes? and how?

One of the original reasons why hwloc was created is that **physical/OS indexes** (`obj->os_index`) are often crazy and unpredictable: logical processors numbers are usually non-contiguous (processors 0 and 1 are not physically close), they vary

from one machine to another, and may even change after a BIOS or system update. This numbers make task placement hardly portable. Moreover some objects have no physical/OS numbers (caches), and some objects have non-unique numbers (core numbers are only unique within a socket). Physical/OS indexes are only guaranteed to exist and be unique for PU and NUMA nodes.

hwloc therefore introduces **logical indexes** (`obj->logical_index`) which are portable, contiguous and logically ordered (based on the resource organization in the locality tree). In general, one should only use logical indexes and just let hwloc do the internal conversion when really needed (when talking to the OS and hardware).

hwloc developers recommends that users do not use physical/OS indexes unless they really know what they are doing. The main reason for still using physical/OS indexes is when interacting with non-hwloc tools such as numactl or taskset, or when reading hardware information from raw sources such as `/proc/cpuinfo`.

lstopo options `-l` and `-p` may be used to switch between logical indexes (prefixed with `L#`) and physical/OS indexes (`P#`). Converting one into the other may also be achieved with `hwloc-calc` which may manipulate either logical or physical indexes as input or output. See also [hwloc-calc](#).

```
# Convert PU with physical number 3 into logical number
$ hwloc-calc -I pu --physical-input --logical-output pu:3
5

# Convert a set of NUMA nodes from logical to physical
# (beware that the output order may not match the input order)
$ hwloc-calc -I numa --logical-input --physical-output numa:2-3 numa:7
0,2,5
```

### 16.1.3 hwloc is only a structural model, it ignores performance models, memory bandwidth, etc.?

hwloc is indeed designed to provide applications with a structural model of the platform. This is an orthogonal approach to describing the machine with performance models, for instance using memory bandwidth or latencies measured by benchmarks. We believe that both approaches are important for helping application make the most of the hardware.

For instance, on a dual-processor host with four cores each, hwloc clearly shows which four cores are together. Latencies between all pairs of cores of the same processor are likely identical, and also likely lower than the latency between cores of different processors. However the structural model cannot guarantee such implementation details. On the other side, performance models would reveal such details without always clearly identifying which cores are in the same processor.

The focus of hwloc is mainly of the structural modeling side. However, hwloc lets user adds performance information to the topology through distances (see [hwloc\\_distances\\_s](#) and [hwloc\\_topology\\_set\\_distance\\_matrix\(\)](#)) or even custom annotations (see [How do I annotate the topology with private notes?](#)). hwloc may also use such distance information for grouping objects together (see [hwloc only has a one-dimensional view of the architecture, it ignores distances?](#) and [What are these Group objects in my topology?](#)).

#### 16.1.4 hwloc only has a one-dimensional view of the architecture, it ignores distances?

hwloc places all objects in a tree. Each level is a one-dimensional view of a set of similar objects. All children of the same object (siblings) are assumed to be equally interconnected (same distance between any of them), while the distance between children of different objects (cousins) is supposed to be larger.

Modern machines exhibit complex hardware interconnects, so this tree may miss some information about the actual physical distances between objects. The hwloc topology may therefore be annotated with distance information that may be used to build a more realistic representation (multi-dimensional) of each level. For instance, the root object may contain a distance matrix that represents the latencies between any pairs of NUMA nodes if the BIOS and/or operating system reports them.

For more information about the distance API, see [hwloc\\_distances\\_s](#) and [hwloc\\_topology\\_set\\_distance\\_matrix\(\)](#).

#### 16.1.5 What are these Group objects in my topology?

hwloc comes with a set of predefined object types (Core, Package, NUMA node, - Caches) that match the vast majority of hardware platforms. The [HWLOC\\_OBJ\\_GROUP](#) type was designed for cases where this set is not sufficient. Groups may be used anywhere to add more structure information to the topology, for instance to show that 2 out of 4 NUMA nodes are actually closer than the others. When applicable, the `Type` info attribute describes why a Group was actually added (see also [Custom string infos](#)).

hwloc currently uses Groups for the following reasons:

- AMD Bulldozer dual-core compute units (`Type=ComputeUnit`, in the x86 backend), but these objects are usually merged with the L2 caches.
- Intel Dies on multi-die CPUs (Linux and x86 backends) and other Extended - Topology Enumeration levels (in the x86 backend).
- Windows processor groups (unless they contain a single NUMA node, or a single Package, etc.).
- IBM S/390 "Books" on Linux (`Type=Book`).
- AIX unknown hierarchy levels.
- Distance-based groups made of close objects.
- I/O parents when I/O locality does not match any existing object.

hwloc Groups are only kept if no other object has the same locality information. It means that a Group containing a single child is merged into that child. And a Group is merged into its parent if it is its only child. For instance a Windows processor group containing a single NUMA node would be merged with that NUMA node since it already contains the relevant hierarchy information.

### 16.1.6 What happens if my topology is asymmetric?

hwloc supports asymmetric topologies even if most platforms are usually symmetric. - For example, there could be different types of processors in a single machine, each with different numbers of cores, symmetric multithreading, or levels of caches.

In practice, asymmetric topologies mostly appear when intermediate groups are added for I/O affinity: on a 4-package machine, an I/O bus may be connected to 2 packages. These packages are below an additional Group object, while the other packages are not (see also [What are these Group objects in my topology?](#)).

Before hwloc v2.0, [hwloc\\_topology\\_ignore\\_type\\_keep\\_structure\(\)](#) and [hwloc\\_topology\\_ignore\\_all\\_keep\\_structure\(\)](#) may also make topologies asymmetric by removing parts of levels, especially when part of the machine is disallowed by administrator restrictions (e.g. Linux cgroups).

To understand how hwloc manages such cases, one should first remember the meaning of levels and cousin objects. All objects of the same type are gathered as horizontal levels with a given depth. They are also connected through the cousin pointers of the [hwloc\\_obj](#) structure. Some types, such as Caches or Groups, are annotated with a depth or level attribute (for instance L2 cache or Group1). Moreover caches have a type attribute (for instance L1i or L1d). Such attributes are also taken in account when gathering objects as horizontal levels. To be clear: there will be one level for L1i caches, another level for L1d caches, another one for L2, etc.

If the topology is asymmetric (e.g., if a group is missing above some processors), a given horizontal level will still exist if there exist any objects of that type. However, some branches of the overall tree may not have an object located in that horizontal level. Note that this specific hole within one horizontal level does not imply anything for other levels. All objects of the same type are gathered in horizontal levels even if their parents or children have different depths and types.

See the diagram in [Terms and Definitions](#) for a graphical representation of such topologies.

Moreover, it is important to understand that a same parent object may have children of different types (and therefore, different depths). **These children are therefore siblings (because they have the same parent), but they are *not* cousins (because they do not belong to the same horizontal level).**

### 16.1.7 What happens to my topology if I disable symmetric multithreading, hyper-threading, etc. in the system?

hwloc creates one PU (processing unit) object per hardware thread. If your machine supports symmetric multithreading, for instance Hyper-Threading, each Core object may contain multiple PU objects:

```
$ lstopo -
...
Core L#0
  PU L#0 (P#0)
  PU L#1 (P#2)
Core L#1
  PU L#2 (P#1)
```

```
PU L#3 (P#3)
```

x86 machines usually offer the ability to disable hyper-threading in the BIOS. Or it can be disabled on the Linux kernel command-line at boot time, or later by writing in sysfs virtual files.

If you do so, the hwloc topology structure does not significantly change, but some PU objects will not appear anymore. No level will disappear, you will see the same number of Core objects, but each of them will contain a single PU now. The PU level does not disappear either (remember that hwloc topologies always contain a PU level at the bottom of the topology) even if there is a single PU object per Core parent.

```
$ lstopo -
...
Core L#0
  PU L#0 (P#0)
Core L#1
  PU L#1 (P#1)
```

### 16.1.8 How may I ignore symmetric multithreading, hyper-threading, etc. in hwloc?

First, see [What happens to my topology if I disable symmetric multithreading, hyper-threading, etc. in the system?](#) for more information about multithreading.

If you need to ignore symmetric multithreading in software, you should likely manipulate hwloc Core objects directly:

```
/* get the number of cores */
unsigned ncores = hwloc_get_nbobjs_by_type(topology, HWLOC_OBJ_CORE);
...
/* get the third core below the first package */
hwloc_obj_t package, core;
package = hwloc_get_obj_by_type(topology, HWLOC_OBJ_PACKAGE, 0);
core = hwloc_get_obj_inside_cpuset_by_type(topology, package->cpuset,
                                           HWLOC_OBJ_CORE, 2);
```

Whenever you want to bind a process or thread to a core, make sure you simplify its cpuset first, so that the task is actually bound to a single thread within this core (to avoid useless migrations).

```
/* bind on the second core */
hwloc_obj_t core = hwloc_get_obj_by_type(topology, HWLOC_OBJ_CORE, 1);
hwloc_cpuset_t set = hwloc_bitmap_dup(core->cpuset);
hwloc_bitmap_simplify(set);
hwloc_set_cpupubind(topology, set, 0);
hwloc_bitmap_free(set);
```

With hwloc-calc or hwloc-bind command-line tools, you may specify that you only want a single-thread within each core by asking for their first PU object:

```
$ hwloc-calc core:4-7
0x0000ff00
$ hwloc-calc core:4-7.pu:0
0x00005500
```

When binding a process on the command-line, you may either specify the exact thread that you want to use, or ask hwloc-bind to simplify the cpuset before binding

```
$ hwloc-bind core:3.pu:0 -- echo "hello from first thread on core #3"
hello from first thread on core #3
...
$ hwloc-bind core:3 --single -- echo "hello from a single thread on core #3"
hello from a single thread on core #3
```

## 16.2 Advanced

### 16.2.1 I do not want hwloc to rediscover my enormous machine topology every time I rerun a process

Although the topology discovery is not expensive on common machines, its overhead may become significant when multiple processes repeat the discovery on large machines (for instance when starting one process per core in a parallel application). The machine topology usually does not vary much, except if some cores are stopped/restarted or if the administrator restrictions are modified. Thus rediscovering the whole topology again and again may look useless.

For this purpose, hwloc offers XML import/export features. It lets you save the discovered topology to a file (for instance with the `lstopo` program) and reload it later by setting the `HWLOC_XMLFILE` environment variable. The `HWLOC_THISSYSTEM` environment variable should also be set to 1 to assert that loaded file is really the underlying system.

Loading a XML topology is usually much faster than querying multiple files or calling multiple functions of the operating system. It is also possible to manipulate such XML files with the C programming interface, and the import/export may also be directed to memory buffer (that may for instance be transmitted between applications through a package). See also [Importing and exporting topologies from/to XML files](#).

#### Note

The environment variable `HWLOC_THISSYSTEM_ALLOWED_RESOURCES` may be used to load a XML topology that contains the entire machine and restrict it to the part that is actually available to the current process (e.g. when Linux `Cgroup`/`Cpuset` are used to restrict the set of resources). See [Environment Variables](#).

### 16.2.2 How many topologies may I use in my program?

hwloc lets you manipulate multiple topologies at the same time. However these topologies consume memory and system resources (for instance file descriptors) until they are destroyed. It is therefore discouraged to open the same topology multiple times.

Sharing a single topology between threads is easy (see [Thread Safety](#)) since the vast majority of accesses are read-only.

If multiple topologies of different (but similar) nodes are needed in your program, have a look at [How to avoid memory waste when manipulating multiple similar topologies?](#).



### 16.2.3 How to avoid memory waste when manipulating multiple similar topologies?

hwloc does not share information between topologies. If multiple similar topologies are loaded in memory, for instance the topologies of different identical nodes of a cluster, lots of information will be duplicated.

[hwloc/diff.h](#) (see also [Topology differences](#)) offers the ability to compute topology differences, apply or unapply them, or export/import to/from XML. However this feature is limited to basic differences such as attribute changes. It does not support complex modifications such as adding or removing some objects.

### 16.2.4 How do I annotate the topology with private notes?

Each hwloc object contains a `userdata` field that may be used by applications to store private pointers. This field is only valid during the lifetime of these container object and topology. It becomes invalid as soon the topology is destroyed, or as soon as the object disappears, for instance when restricting the topology. The `userdata` field is not exported/imported to/from XML by default since hwloc does not know what it contains. This behavior may be changed by specifying application-specific callbacks with `hwloc_topology_set_userdata_export_callback()` and `hwloc_topology_set_userdata_import_callback()`.

Each object may also contain some *info* attributes (key name and value) that are setup by hwloc during discovery and that may be extended by the user with `hwloc_obj_add_info()` (see also [Object attributes](#)). Contrary to the `userdata` field which is unique, multiple `info` attributes may exist for each object, even with the same name. These attributes are always exported to XML. However only character strings may be used as key names and values.

It is also possible to insert Misc objects with a custom name anywhere as a leaf of the topology (see [Miscellaneous objects](#)). And Misc objects may have their own `userdata` and `info` attributes just like any other object.

The `hwloc-annotate` command-line tool may be used for adding Misc objects and `info` attributes.

There is also a topology-specific `userdata` pointer that can be used to recognize different topologies by storing a custom pointer. It may be manipulated with `hwloc_topology_set_userdata()` and `hwloc_topology_get_userdata()`.

## 16.3 Caveats

### 16.3.1 Why is hwloc slow?

Building a hwloc topology on a large machine may be slow because the discovery of hundreds of hardware cores or threads takes time (especially when reading thousands of `sysfs` files on Linux). One should consider using XML (see [I do not want hwloc to rediscover my enormous machine topology every time I rerun a process](#)) to work

around such issues.

Additionally, `lstopo` enables most `hwloc` objects and discovery flags by default so that the output topology is as precise as possible (while `hwloc` disables many of them by default). This includes I/O device discovery through PCI libraries as well as external libraries such as NVML. To speed up `lstopo`, you may disable such features with command-line options such as `\--no-io`.

When NVIDIA GPU probing is enabled with CUDA or NVML, one should make sure that the *Persistent* mode is enabled (with `nvidia-smi -pm 1`) to avoid significant GPU initialization overhead.

When AMD GPU discovery is enabled with OpenCL and `hwloc` is used remotely over `ssh`, some spurious round-trips on the network may significantly increase the discovery time. Forcing the `DISPLAY` environment variable to the remote X server display (usually `:0`) instead of only setting the `COMPUTE` variable may avoid this.

Also remember that these components may be disabled at build-time with configure flags such as `\--disable-openssl`, `\--disable-cuda` or `\--disable-nvml`, and at runtime with the environment variable `HWLOC_COMPONENTS=-openssl,-cuda,-nvml`.

### 16.3.2 Does `hwloc` require privileged access?

`hwloc` discovers the topology by querying the operating system. Some minor features may require privileged access to the operation system. For instance memory module discovery on Linux is reserved to root, and the entire PCI discovery on Solaris and BSDs requires access to some special files that are usually restricted to root (`/dev/pci*` or `/devices/pci*`).

To workaround this limitation, it is recommended to export the topology as a XML file generated by the administrator (with the `lstopo` program) and make it available to all users (see [Importing and exporting topologies from/to XML files](#)). It will offer all discovery information to any application without requiring any privileged access anymore. Only the necessary hardware characteristics will be exported, no sensitive information will be disclosed through this XML export.

This XML-based model also has the advantage of speeding up the discovery because reading a XML topology is usually much faster than querying the operating system again.

The utility `hwloc-dump-hwdata` is also involved in gathering privileged information at boot time and making it available to non-privileged users (note that this may require a specific SELinux MLS policy module). However, it only applies to Intel Xeon Phi processors for now (see [Why do I need `hwloc-dump-hwdata` for memory on Intel Xeon Phi processor?](#)). See also `HWLOC_DUMPED_HWDATA_DIR` in [Environment Variables](#) for details about the location of dumped files.

### 16.3.3 What should I do when hwloc reports "operating system" warnings?

When the operating system reports invalid locality information (because of either software or hardware bugs), hwloc may fail to insert some objects in the topology because they cannot fit in the already built tree of resources. If so, hwloc will report a warning like the following. The object causing this error is ignored, the discovery continues but the resulting topology will miss some objects and may be asymmetric (see also [What happens if my topology is asymmetric?](#)).

```
*****
* hwloc received invalid information from the operating system.
*
* L3 (cpuset 0x000003f0) intersects with NUMANode (P#0 cpuset 0x000003f) without inclusion!
* Error occurred in topology.c line 940
*
* Please report this error message to the hwloc user's mailing list,
* along with the files generated by the hwloc-gather-topology script.
*
* hwloc will now ignore this invalid topology information and continue.
*****
```

These errors are common on large AMD platforms because of BIOS and/or Linux kernel bugs causing invalid L3 cache information. In the above example, the hardware reports a L3 cache that is shared by 2 cores in the first NUMA node and 4 cores in the second NUMA node. That's wrong, it should actually be shared by all 6 cores in a single NUMA node. The resulting topology will miss some L3 caches.

If your application does not care about cache sharing, or if you do not plan to request cache-aware binding in your process launcher, you may likely ignore this error (and hide it by setting `HWLOC_HIDE_ERRORS=1` in your environment).

Some platforms report similar warnings about conflicting Packages and NUMANodes.

On x86 hosts, passing `HWLOC_COMPONENTS=x86` in the environment may workaround some of these issues by switching to a different way to discover the topology.

Upgrading the BIOS and/or the operating system may help. Otherwise, as explained in the message, reporting this issue to the hwloc developers (by sending the tarball that is generated by the `hwloc-gather-topology` script on this platform) is a good way to make sure that this is a software (operating system) or hardware bug (BIOS, etc).

See also [Questions and Bugs](#). Opening an issue on GitHub automatically displays hints on what information you should provide when reporting such bugs.

### 16.3.4 Why does Valgrind complain about hwloc memory leaks?

If you are debugging your application with Valgrind, you want to avoid memory leak reports that are caused by hwloc and not by your program.

hwloc itself is often checked with Valgrind to make sure it does not leak memory. - However some global variables in hwloc dependencies are never freed. For instance `libz` allocates its global state once at startup and never frees it so that it may be reused later. Some `libxml2` global state is also never freed because hwloc does not know

whether it can safely ask libxml2 to free it (the application may also be using libxml2 outside of hwloc).

These unfreed variables cause leak reports in Valgrind. hwloc installs a Valgrind *suppressions* file to hide them. You should pass the following command-line option to Valgrind to use it:

```
--suppressions=/path/to/hwloc-valgrind.supp
```

### 16.3.5 How do I handle ABI breaks and API upgrades?

The hwloc interface is extended with every new major release. Any application using the hwloc API should be prepared to check at compile-time whether some features are available in the currently installed hwloc distribution.

For instance, to check whether the hwloc version is at least 1.10, you should use:

```
#include <hwloc.h>
#if HWLOC_API_VERSION >= 0x00010a00
...
#endif
```

To check for the API of release X.Y.Z at build time, you may compare [HWLOC\\_API\\_VERSION](#) with  $(X \ll 16) + (Y \ll 8) + Z$ .

For supporting older releases that do not have `HWLOC_OBJ_NUMANODE` and `HWLOC_OBJ_PACKAGE` yet, you may use:

```
#include <hwloc.h>
#if HWLOC_API_VERSION < 0x00010b00
#define HWLOC_OBJ_NUMANODE HWLOC_OBJ_NODE
#define HWLOC_OBJ_PACKAGE HWLOC_OBJ_SOCKET
#endif
```

The hwloc interface will be deeply modified in release 2.0 to fix several issues of the 1.x interface. The ABI will be broken, which means **applications must be recompiled against the new 2.0 interface**.

To check that you are not mixing old/recent headers with a recent/old runtime library, check the major revision number in the API version:

```
#include <hwloc.h>
unsigned version = hwloc_get_api_version();
if ((version >> 16) != (HWLOC_API_VERSION >> 16)) {
    fprintf(stderr,
        "%s compiled for hwloc API 0x%x but running on library API 0x%x.\n"
        "You may need to point LD_LIBRARY_PATH to the right hwloc library.\n"
        "Aborting since the new ABI is not backward compatible.\n",
        callname, HWLOC_API_VERSION, version);
    exit(EXIT_FAILURE);
}
```

To specifically detect v2.0 issues:

```
#include <hwloc.h>
#if HWLOC_API_VERSION >= 0x00020000
```

```
/* headers are recent */
if (hwloc_get_api_version() < 0x20000)
    ... error out, the hwloc runtime library is older than 2.0 ...
#else
/* headers are pre-2.0 */
if (hwloc_get_api_version() >= 0x20000)
    ... error out, the hwloc runtime library is more recent than 2.0 ...
#endif
```

You should not try to remain compatible with very old releases such as 1.1.x or earlier because `HWLOC_API_VERSION` was added in 1.0.0 and `hwloc_get_api_version()` came only in 1.1.1. Also do not use the old `cpuset` API since it was deprecated and superseded by the `bitmap` API in 1.1, and later removed in 1.5.

## 16.4 Platform-specific

### 16.4.1 How do I find the local MCDRAM NUMA node on Intel Xeon Phi processor?

Intel Xeon Phi processors introduced a new memory architecture by possibly having two distinct local memories: some normal memory (DDR) and some high-bandwidth on-package memory (MCDRAM). Processors can be configured in various clustering modes to have up to 4 *Clusters*. Moreover, each *Cluster* (quarter, half or whole processor) of the processor may have its own local parts of the DDR and of the MCDRAM. This memory and clustering configuration may be probed by looking at `MemoryMode` and `ClusterMode` attributes, see [Custom string infos](#) and `doc/examples/get-knl-modes.c` in the source directory.

The upcoming hwloc 2.0 will address this new architecture by presenting memory in an improved way. For now, starting with 1.11.2, hwloc releases use the following approximate representation:

If a cluster only contains DDR or MCDRAM but not both, that memory is available as a local NUMA node above cores as usual.

If a cluster contains both, two distinct NUMA nodes appear. They are sibling children of a `Group` object of type `Cluster` (or sibling children of the `Package` object for non-clustered processors).

The DDR memory is the local NUMA node above cores as usual. Allocating memory from one core to its local NUMA node will therefore actually allocate it on the normal memory by default.

The local high-bandwidth MCDRAM is the second NUMA node (without any `Core` or `PU` below it). It is the next sibling of the local DDR NUMA node below the same parent object. To allocate on the faster MCDRAM, one should first find the local NUMA node (the DDR memory, by looking up parent objects), and then take the next sibling to reach the local MCDRAM NUMA node (if any).

The MCDRAM NUMA nodes may also be identified thanks to the `Type` info attribute which is set to `MCDRAM`.

Command-line tools such as `hwloc-bind` may bind memory on the MCDRAM by using the `hbm` keyword. For instance, to bind on the first MCDRAM NUMA node:

```
$ hwloc-bind --membind --hbm numa:0 -- myprogram
$ hwloc-bind --membind numa:0 -- myprogram
```

### 16.4.2 Why do I need hwloc-dump-hwdata for memory on Intel Xeon Phi processor?

Intel Xeon Phi processors may use the on-package memory (MCDRAM) as either memory or a memory-side cache (currently reported as a L3 cache by hwloc). There are also several clustering modes that significantly affect the memory organization (see [How do I find the local MCDRAM NUMA node on Intel Xeon Phi processor?](#) for more information about these modes). Details about these are currently only available to privileged users. Without them, hwloc relies on a heuristic for guessing the modes.

The hwloc-dump-hwdata utility may be used to dump this privileged binary information into human-readable and world-accessible files that the hwloc library will later load. - The utility should usually run as root once during boot, in order to update dumped information (stored under /var/run/hwloc by default) in case the MCDRAM or clustering configuration changed between reboots.

When SELinux MLS policy is enabled, a specific hwloc policy module may be required so that all users get access to the dumped files (in /var/run/hwloc by default). One may use hwloc policy files from the SELinux Reference Policy at <https://github.com/TresysTechnology/refpolicy-contrib> (see also the documentation at <https://github.com/TresysTechnology/refpolicy/wiki/GettingStarted>).

hwloc-dump-hwdata requires `dmi-sysfs` kernel module loaded.

The utility is currently unneeded on platforms without Intel Xeon Phi processors.

See `HWLOC_DUMPED_HWDATA_DIR` in [Environment Variables](#) for details about the location of dumped files.

### 16.4.3 How do I build for Intel Xeon Phi coprocessor?

#### Note

This section does not apply to standalone Intel Xeon Phi processors (Knights - Landing and Knights Mill).

Intel Xeon Phi coprocessors (Knights Corner) usually runs a Linux environment but cross-compiling from the host is required. hwloc uses standard autotools options for cross-compiling.

If building with `icc`:

```
./configure CC="icc -mmic" --host=x86_64-k1om-linux --build=x86_64-unknown-linux-gnu
```

If building with the Xeon Phi-specific GCC that comes with the MPSS environment, for instance `/usr/linux-k1om-4.7/bin/x86_64-k1om-linux-gcc`:

```
export PATH=$PATH:/usr/linux-k1om-4.7/bin/
./configure --host=x86_64-k1om-linux --build=x86_64-unknown-linux-gnu
```

#### 16.4.4 How do I build hwloc for BlueGene/Q?

IBM BlueGene/Q machines run a standard Linux on the login/frontend nodes and a custom CNK (*Compute Node Kernel*) on the compute nodes.

To discover the topology of a login/frontend node, hwloc should be configured as usual, without any BlueGene/Q-specific option.

However, one would likely rather discover the topology of the compute nodes where parallel jobs are actually running. If so, hwloc must be cross-compiled with the following configuration line:

```
./configure --host=powerpc64-bgq-linux --disable-shared --enable-static \  
  CPPFLAGS='-I/bgsys/drivers/ppcfloor -I/bgsys/drivers/ppcfloor/spi/include/kernel/cnk/'
```

CPPFLAGS may have to be updated if your platform headers are installed in a different directory.

#### 16.4.5 How do I build hwloc for Windows?

hwloc releases are available as pre-built ZIPs for Windows on both 32bits and 64bits x86 platforms. They are built using MSYS2 and MinGW on a Windows host. Such an environment allows using the Unix-like `configure`, `make` and `make install` steps without having to tweak too many variables or options. One may look at `contrib/ci.inria.fr/job-3-mingw.sh` in the hwloc repository for an example used for nightly testing.

hwloc releases also contain a basic Microsoft Visual Studio solution under `contrib/windows/`.

#### 16.4.6 How to get useful topology information on NetBSD?

The NetBSD (and FreeBSD) backend uses x86-specific topology discovery (through the x86 component). This implementation requires CPU binding so as to query topology information from each individual logical processor. This means that hwloc cannot find any useful topology information unless user-level process binding is allowed by the NetBSD kernel. The `security.models.extensions.user_set_cpu_affinity` sysctl variable must be set to 1 to do so. Otherwise, only the number of logical processors will be detected.

#### 16.4.7 Why does binding fail on AIX?

The AIX operating system requires specific user capabilities for attaching processes to resource sets (CAP\_NUMA\_ATTACH). Otherwise functions such as `hwloc_set_cpubind()` fail (return -1 with `errno` set to `EPERM`).

This capability must also be inherited (through the additional `CAP_PROPAGATE` capability) if you plan to bind a process before forking another process, for instance with `hwloc-bind`.

These capabilities may be given by the administrator with:

```
chuser "capabilities=CAP_PROPAGATE,CAP_NUMA_ATTACH" <username>
```



# Chapter 17

## Module Index

### 17.1 Modules

Here is a list of all modules:

API version . . . . .	83
Object Sets (hwloc_cpuset_t and hwloc_nodeset_t) . . . . .	85
Object Types . . . . .	87
Object Structure and Attributes . . . . .	91
Topology Creation and Destruction . . . . .	92
Topology Detection Configuration and Query . . . . .	95
Object levels, depths and types . . . . .	104
Converting between Object Types, Sets and Attributes, and Strings . . . . .	108
Consulting and Adding Key-Value Info Attributes . . . . .	111
CPU binding . . . . .	112
Memory binding . . . . .	117
Modifying a loaded Topology . . . . .	130
Building Custom Topologies . . . . .	133
Exporting Topologies to XML . . . . .	135
Exporting Topologies to Synthetic . . . . .	139
Finding Objects inside a CPU set . . . . .	141
Finding Objects covering at least CPU set . . . . .	145
Looking at Ancestor and Child Objects . . . . .	147
Looking at Cache Objects . . . . .	149
Finding objects, miscellaneous helpers . . . . .	151
Distributing items over a topology . . . . .	153
CPU and node sets of entire topologies . . . . .	155
Converting between CPU sets and node sets . . . . .	158
Manipulating Distances . . . . .	160
Finding I/O objects . . . . .	163
The bitmap API . . . . .	165
Linux-specific helpers . . . . .	177
Interoperability with Linux libnuma unsigned long masks . . . . .	179
Interoperability with Linux libnuma bitmask . . . . .	181

---

Interoperability with glibc sched affinity . . . . .	183
Interoperability with OpenCL . . . . .	184
Interoperability with the CUDA Driver API . . . . .	186
Interoperability with the CUDA Runtime API . . . . .	188
Interoperability with the NVIDIA Management Library . . . . .	190
Interoperability with OpenGL displays . . . . .	192
Interoperability with Intel Xeon Phi (MIC) . . . . .	194
Interoperability with OpenFabrics . . . . .	195
Interoperability with Myrinet Express . . . . .	197
Topology differences . . . . .	198
Components and Plugins: Discovery components . . . . .	204
Components and Plugins: Discovery backends . . . . .	205
Components and Plugins: Generic components . . . . .	207
Components and Plugins: Core functions to be used by components . . . . .	208
Components and Plugins: PCI functions to be used by components . . . . .	211

# Chapter 18

## Data Structure Index

### 18.1 Data Structures

Here are the data structures with brief descriptions:

<a href="#">hwloc_backend</a>	Discovery backend structure . . . . .	213
<a href="#">hwloc_obj_attr_u::hwloc_bridge_attr_s</a>	Bridge specific Object Attributes . . . . .	215
<a href="#">hwloc_obj_attr_u::hwloc_cache_attr_s</a>	Cache-specific Object Attributes . . . . .	216
<a href="#">hwloc_component</a>	Generic component structure . . . . .	217
<a href="#">hwloc_disc_component</a>	Discovery component structure . . . . .	218
<a href="#">hwloc_distances_s</a>	Distances between objects . . . . .	220
<a href="#">hwloc_obj_attr_u::hwloc_group_attr_s</a>	Group-specific Object Attributes . . . . .	221
<a href="#">hwloc_obj</a>	Structure of a topology object . . . . .	222
<a href="#">hwloc_obj_attr_u</a>	Object type-specific Attributes . . . . .	228
<a href="#">hwloc_obj_info_s</a>	Object info . . . . .	229
<a href="#">hwloc_obj_memory_s::hwloc_obj_memory_page_type_s</a>	Array of local memory page types, NULL if no local memory and page_types is 0 . . . . .	229
<a href="#">hwloc_obj_memory_s</a>	Object memory . . . . .	230
<a href="#">hwloc_obj_attr_u::hwloc_osdev_attr_s</a>	OS Device specific Object Attributes . . . . .	231
<a href="#">hwloc_obj_attr_u::hwloc_pcidev_attr_s</a>	PCI Device specific Object Attributes . . . . .	232

<a href="#">hwloc_topology_cpupbind_support</a>	
Flags describing actual PU binding support for this topology	233
<a href="#">hwloc_topology_diff_u::hwloc_topology_diff_generic_s</a>	235
<a href="#">hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_generic_s</a>	235
<a href="#">hwloc_topology_diff_u::hwloc_topology_diff_obj_attr_s</a>	236
<a href="#">hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_string_s</a>	
String attribute modification with an optional name	236
<a href="#">hwloc_topology_diff_obj_attr_u</a>	
One object attribute difference	237
<a href="#">hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_uint64_s</a>	
Integer attribute modification with an optional index	238
<a href="#">hwloc_topology_diff_u::hwloc_topology_diff_too_complex_s</a>	239
<a href="#">hwloc_topology_diff_u</a>	
One element of a difference list between two topologies	240
<a href="#">hwloc_topology_discovery_support</a>	
Flags describing actual discovery support for this topology	240
<a href="#">hwloc_topology_membind_support</a>	
Flags describing actual memory binding support for this topology	241
<a href="#">hwloc_topology_support</a>	
Set of flags describing actual support for this topology	243

# Chapter 19

## Module Documentation

### 19.1 API version

#### Defines

- `#define HWLOC_API_VERSION 0x00010b06`
- `#define HWLOC_COMPONENT_ABI 4`

#### Functions

- unsigned `hwloc_get_api_version` (void)

#### 19.1.1 Define Documentation

##### 19.1.1.1 `#define HWLOC_API_VERSION 0x00010b06`

Indicate at build time which hwloc API version is being used.

This number is updated to  $(X \ll 16) + (Y \ll 8) + Z$  when a new release X.Y.Z actually modifies the API.

Users may check for available features at build time using this number (see [How do I handle ABI breaks and API upgrades?](#)).

##### 19.1.1.2 `#define HWLOC_COMPONENT_ABI 4`

Current component and plugin ABI version (see [hwloc/plugins.h](#))

#### 19.1.2 Function Documentation

#### 19.1.2.1 unsigned hwloc\_get\_api\_version ( void )

Indicate at runtime which hwloc API version was used at build time.

Should be [HWLOC\\_API\\_VERSION](#) if running on the same version.

## 19.2 Object Sets ([hwloc\\_cpuset\\_t](#) and [hwloc\\_nodeset\\_t](#))

### Typedefs

- typedef [hwloc\\_bitmap\\_t](#) [hwloc\\_cpuset\\_t](#)
- typedef [hwloc\\_const\\_bitmap\\_t](#) [hwloc\\_const\\_cpuset\\_t](#)
- typedef [hwloc\\_bitmap\\_t](#) [hwloc\\_nodeset\\_t](#)
- typedef [hwloc\\_const\\_bitmap\\_t](#) [hwloc\\_const\\_nodeset\\_t](#)

### 19.2.1 Detailed Description

Hwloc uses bitmaps to represent two distinct kinds of object sets: CPU sets ([hwloc\\_cpuset\\_t](#)) and NUMA node sets ([hwloc\\_nodeset\\_t](#)). These types are both typedefs to a common back end type ([hwloc\\_bitmap\\_t](#)), and therefore all the hwloc bitmap functions are applicable to both [hwloc\\_cpuset\\_t](#) and [hwloc\\_nodeset\\_t](#) (see [The bitmap API](#)).

The rationale for having two different types is that even though the actions one wants to perform on these types are the same (e.g., enable and disable individual items in the set/mask), they're used in very different contexts: one for specifying which processors to use and one for specifying which NUMA nodes to use. Hence, the name difference is really just to reflect the intent of where the type is used.

### 19.2.2 Typedef Documentation

#### 19.2.2.1 typedef [hwloc\\_const\\_bitmap\\_t](#) [hwloc\\_const\\_cpuset\\_t](#)

A non-modifiable [hwloc\\_cpuset\\_t](#).

#### 19.2.2.2 typedef [hwloc\\_const\\_bitmap\\_t](#) [hwloc\\_const\\_nodeset\\_t](#)

A non-modifiable [hwloc\\_nodeset\\_t](#).

#### 19.2.2.3 typedef [hwloc\\_bitmap\\_t](#) [hwloc\\_cpuset\\_t](#)

A CPU set is a bitmap whose bits are set according to CPU physical OS indexes.

It may be consulted and modified with the bitmap API as any [hwloc\\_bitmap\\_t](#) (see [hwloc/bitmap.h](#)).

Each bit may be converted into a PU object using [hwloc\\_get\\_pu\\_obj\\_by\\_os\\_index\(\)](#).

#### 19.2.2.4 typedef [hwloc\\_bitmap\\_t](#) [hwloc\\_nodeset\\_t](#)

A node set is a bitmap whose bits are set according to NUMA memory node physical OS indexes.

It may be consulted and modified with the bitmap API as any [hwloc\\_bitmap\\_t](#) (see [hwloc/bitmap.h](#)). Each bit may be converted into a NUMA node object using [hwloc\\_get\\_numanode\\_obj\\_by\\_os\\_index\(\)](#).

When binding memory on a system without any NUMA node (when the whole memory is considered as a single memory bank), the nodeset may be either empty (no memory selected) or full (whole system memory selected).

See also [Converting between CPU sets and node sets](#).



## 19.3 Object Types

### Typedefs

- typedef enum `hwloc_obj_cache_type_e` `hwloc_obj_cache_type_t`
- typedef enum `hwloc_obj_bridge_type_e` `hwloc_obj_bridge_type_t`
- typedef enum `hwloc_obj_osdev_type_e` `hwloc_obj_osdev_type_t`

### Enumerations

- enum `hwloc_obj_type_t` { `HWLOC_OBJ_SYSTEM`, `HWLOC_OBJ_MACHINE`, `HWLOC_OBJ_NUMANODE`, `HWLOC_OBJ_PACKAGE`, `HWLOC_OBJ_CACHE`, `HWLOC_OBJ_CORE`, `HWLOC_OBJ_PU`, `HWLOC_OBJ_GROUP`, `HWLOC_OBJ_MISC`, `HWLOC_OBJ_BRIDGE`, `HWLOC_OBJ_PCI_DEVICE`, `HWLOC_OBJ_OS_DEVICE`, `HWLOC_OBJ_TYPE_MAX` }
- enum `hwloc_obj_cache_type_e` { `HWLOC_OBJ_CACHE_UNIFIED`, `HWLOC_OBJ_CACHE_DATA`, `HWLOC_OBJ_CACHE_INSTRUCTION` }
- enum `hwloc_obj_bridge_type_e` { `HWLOC_OBJ_BRIDGE_HOST`, `HWLOC_OBJ_BRIDGE_PCI` }
- enum `hwloc_obj_osdev_type_e` { `HWLOC_OBJ_OSDEV_BLOCK`, `HWLOC_OBJ_OSDEV_GPU`, `HWLOC_OBJ_OSDEV_NETWORK`, `HWLOC_OBJ_OSDEV_OPENFABRICS`, `HWLOC_OBJ_OSDEV_DMA`, `HWLOC_OBJ_OSDEV_COPROC` }
- enum `hwloc_compare_types_e` { `HWLOC_TYPE_UNORDERED` }

### Functions

- int `hwloc_compare_types` (`hwloc_obj_type_t` type1, `hwloc_obj_type_t` type2)

#### 19.3.1 Typedef Documentation

##### 19.3.1.1 typedef enum `hwloc_obj_bridge_type_e` `hwloc_obj_bridge_type_t`

Type of one side (upstream or downstream) of an I/O bridge.

##### 19.3.1.2 typedef enum `hwloc_obj_cache_type_e` `hwloc_obj_cache_type_t`

Cache type.

##### 19.3.1.3 typedef enum `hwloc_obj_osdev_type_e` `hwloc_obj_osdev_type_t`

Type of a OS device.

## 19.3.2 Enumeration Type Documentation

### 19.3.2.1 enum hwloc\_compare\_types\_e

Enumerator:

**HWLOC\_TYPE\_UNORDERED** Value returned by `hwloc_compare_types()` when types can not be compared.

### 19.3.2.2 enum hwloc\_obj\_bridge\_type\_e

Type of one side (upstream or downstream) of an I/O bridge.

Enumerator:

**HWLOC\_OBJ\_BRIDGE\_HOST** Host-side of a bridge, only possible upstream.

**HWLOC\_OBJ\_BRIDGE\_PCI** PCI-side of a bridge.

### 19.3.2.3 enum hwloc\_obj\_cache\_type\_e

Cache type.

Enumerator:

**HWLOC\_OBJ\_CACHE\_UNIFIED** Unified cache.

**HWLOC\_OBJ\_CACHE\_DATA** Data cache.

**HWLOC\_OBJ\_CACHE\_INSTRUCTION** Instruction cache. Only used when the `HWLOC_TOPOLOGY_FLAG_ICACHES` topology flag is set.

### 19.3.2.4 enum hwloc\_obj\_osdev\_type\_e

Type of a OS device.

Enumerator:

**HWLOC\_OBJ\_OSDEV\_BLOCK** Operating system block device. For instance "sda" on Linux.

**HWLOC\_OBJ\_OSDEV\_GPU** Operating system GPU device. For instance ":0.0" for a GL display, "card0" for a Linux DRM device.

**HWLOC\_OBJ\_OSDEV\_NETWORK** Operating system network device. For instance the "eth0" interface on Linux.

**HWLOC\_OBJ\_OSDEV\_OPENFABRICS** Operating system openfabrics device. - For instance the "mlx4\_0" InfiniBand HCA, or "hfi1\_0" Omni-Path interface on Linux.

**HWLOC\_OBJ\_OSDEV\_DMA** Operating system dma engine device. For instance the "dma0chan0" DMA channel on Linux.

**HWLOC\_OBJ\_OSDEV\_COPROC** Operating system co-processor device. For instance "mic0" for a Xeon Phi (MIC) on Linux, "opencl0d0" for a OpenCL device, "cuda0" for a CUDA device.

### 19.3.2.5 enum hwloc\_obj\_type\_t

Type of topology object.

#### Note

Do not rely on the ordering or completeness of the values as new ones may be defined in the future! If you need to compare types, use [hwloc\\_compare\\_types\(\)](#) instead.

#### Enumerator:

**HWLOC\_OBJ\_SYSTEM** Whole system (may be a cluster of machines). The whole system that is accessible to hwloc. That may comprise several machines in SSI systems like Kerrighed.

**HWLOC\_OBJ\_MACHINE** Machine. The typical root object type. A set of processors and memory with cache coherency.

**HWLOC\_OBJ\_NUMANODE** NUMA node. An object that contains memory that is directly and byte-accessible to the host processors. It is usually close to some cores (the corresponding objects are descendants of the NUMA node object in the hwloc tree). There is always at one such object in the topology even if the machine is not NUMA.

**HWLOC\_OBJ\_PACKAGE** Physical package. The physical package that usually gets inserted into a socket on the motherboard. A processor package usually contains multiple cores.

**HWLOC\_OBJ\_CACHE** Cache. Can be L1i, L1d, L2, L3, ...

**HWLOC\_OBJ\_CORE** Core. A computation unit (may be shared by several logical processors).

**HWLOC\_OBJ\_PU** Processing Unit, or (Logical) Processor. An execution unit (may share a core with some other logical processors, e.g. in the case of an SMT core). Objects of this kind are always reported and can thus be used as fallback when others are not.

**HWLOC\_OBJ\_GROUP** Group objects. Objects which do not fit in the above but are detected by hwloc and are useful to take into account for affinity. For instance, some operating systems expose their arbitrary processors aggregation this way. And hwloc may insert such objects to group NUMA nodes according to their distances. See also [What are these Group objects in my topology?](#). These objects are ignored when they do not bring any structure.

**HWLOC\_OBJ\_MISC** Miscellaneous objects. Objects without particular meaning, that can e.g. be added by the application for its own use, or by hwloc for miscellaneous objects such as MemoryModule (DIMMs).

**HWLOC\_OBJ\_BRIDGE** Bridge. Any bridge that connects the host or an I/O bus, to another I/O bus. Bridge objects have neither CPU sets nor node sets. They are not added to the topology unless I/O discovery is enabled with [hwloc\\_topology\\_set\\_flags\(\)](#).

**HWLOC\_OBJ\_PCI\_DEVICE** PCI device. These objects have neither CPU sets nor node sets. They are not added to the topology unless I/O discovery is enabled with [hwloc\\_topology\\_set\\_flags\(\)](#).

**HWLOC\_OBJ\_OS\_DEVICE** Operating system device. These objects have neither CPU sets nor node sets. They are not added to the topology unless I/O discovery is enabled with [hwloc\\_topology\\_set\\_flags\(\)](#).

**HWLOC\_OBJ\_TYPE\_MAX** Sentinel value

### 19.3.3 Function Documentation

#### 19.3.3.1 `int hwloc_compare_types ( hwloc_obj_type_t type1, hwloc_obj_type_t type2 )`

Compare the depth of two object types.

Types shouldn't be compared as they are, since newer ones may be added in the future. This function returns less than, equal to, or greater than zero respectively if `type1` objects usually include `type2` objects, are the same as `type2` objects, or are included in `type2` objects. If the types can not be compared (because neither is usually contained in the other), [HWLOC\\_TYPE\\_UNORDERED](#) is returned. Object types containing CPUs can always be compared (usually, a system contains machines which contain nodes which contain packages which contain caches, which contain cores, which contain processors).

#### Note

[HWLOC\\_OBJ\\_PU](#) will always be the deepest.

This does not mean that the actual topology will respect that order: e.g. as of today cores may also contain caches, and packages may also contain nodes. This is thus just to be seen as a fallback comparison method.

## 19.4 Object Structure and Attributes

### Data Structures

- struct [hwloc\\_obj\\_memory\\_s](#)  
*Object memory.*
- struct [hwloc\\_obj](#)  
*Structure of a topology object.*
- union [hwloc\\_obj\\_attr\\_u](#)  
*Object type-specific Attributes.*
- struct [hwloc\\_distances\\_s](#)  
*Distances between objects.*
- struct [hwloc\\_obj\\_info\\_s](#)  
*Object info.*

### Typedefs

- typedef struct [hwloc\\_obj](#) \* [hwloc\\_obj\\_t](#)

#### 19.4.1 Typedef Documentation

##### 19.4.1.1 typedef struct [hwloc\\_obj](#)\* [hwloc\\_obj\\_t](#)

Convenience typedef; a pointer to a struct [hwloc\\_obj](#).

## 19.5 Topology Creation and Destruction

### Typedefs

- typedef struct hwloc\_topology \* [hwloc\\_topology\\_t](#)

### Functions

- int [hwloc\\_topology\\_init](#) ([hwloc\\_topology\\_t](#) \*topologyp)
- int [hwloc\\_topology\\_load](#) ([hwloc\\_topology\\_t](#) topology)
- void [hwloc\\_topology\\_destroy](#) ([hwloc\\_topology\\_t](#) topology)
- int [hwloc\\_topology\\_dup](#) ([hwloc\\_topology\\_t](#) \*newtopology, [hwloc\\_topology\\_t](#) oldtopology)
- void [hwloc\\_topology\\_check](#) ([hwloc\\_topology\\_t](#) topology)

### 19.5.1 Typedef Documentation

#### 19.5.1.1 typedef struct hwloc\_topology\* [hwloc\\_topology\\_t](#)

Topology context.

To be initialized with [hwloc\\_topology\\_init\(\)](#) and built with [hwloc\\_topology\\_load\(\)](#).

### 19.5.2 Function Documentation

#### 19.5.2.1 void [hwloc\\_topology\\_check](#) ( [hwloc\\_topology\\_t](#) *topology* )

Run internal checks on a topology structure.

The program aborts if an inconsistency is detected in the given topology.

#### Parameters

<i>topology</i>	is the topology to be checked
-----------------	-------------------------------

#### Note

This routine is only useful to developers.

The input topology should have been previously loaded with [hwloc\\_topology\\_load\(\)](#).

#### 19.5.2.2 void [hwloc\\_topology\\_destroy](#) ( [hwloc\\_topology\\_t](#) *topology* )

Terminate and free a topology context.

## Parameters

<i>topology</i>	is the topology to be freed
-----------------	-----------------------------

19.5.2.3 `int hwloc_topology_dup ( hwloc_topology_t * newtopology,  
hwloc_topology_t oldtopology )`

Duplicate a topology.

The entire topology structure as well as its objects are duplicated into a new one.

This is useful for keeping a backup while modifying a topology.

## Note

Object userdata is not duplicated since hwloc does not know what it point to. The objects of both old and new topologies will point to the same userdata.

19.5.2.4 `int hwloc_topology_init ( hwloc_topology_t * topologyp )`

Allocate a topology context.

## Parameters

<code>out</code>	<i>topologyp</i>	is assigned a pointer to the new allocated context.
------------------	------------------	-----------------------------------------------------

## Returns

0 on success, -1 on error.

19.5.2.5 `int hwloc_topology_load ( hwloc_topology_t topology )`

Build the actual topology.

Build the actual topology once initialized with [hwloc\\_topology\\_init\(\)](#) and tuned with [Topology Detection Configuration and Query](#) routines. No other routine may be called earlier using this topology context.

## Parameters

<i>topology</i>	is the topology to be loaded with objects.
-----------------	--------------------------------------------

## Returns

0 on success, -1 on error.

**Note**

On failure, the topology is reinitialized. It should be either destroyed with [hwloc\\_topology\\_destroy\(\)](#) or configured and loaded again. This function may be called only once per topology. The binding of the current thread or process may temporarily change during this call but it will be restored before it returns.

**See also**

[Topology Detection Configuration and Query](#)



## 19.6 Topology Detection Configuration and Query

### Data Structures

- struct `hwloc_topology_discovery_support`  
*Flags describing actual discovery support for this topology.*
- struct `hwloc_topology_cpupbind_support`  
*Flags describing actual PU binding support for this topology.*
- struct `hwloc_topology_membind_support`  
*Flags describing actual memory binding support for this topology.*
- struct `hwloc_topology_support`  
*Set of flags describing actual support for this topology.*

### Enumerations

- enum `hwloc_topology_flags_e` { `HWLOC_TOPOLOGY_FLAG_WHOLE_SYSTEM`, `HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM`, `HWLOC_TOPOLOGY_FLAG_IO_DEVICES`, `HWLOC_TOPOLOGY_FLAG_IO_BRIDGES`, `HWLOC_TOPOLOGY_FLAG_WHOLE_IO`, `HWLOC_TOPOLOGY_FLAG_ICACHES`, `HWLOC_TOPOLOGY_FLAG_THISSYSTEM_ALLOWED_RESOURCES` }

### Functions

- int `hwloc_topology_ignore_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- int `hwloc_topology_ignore_type_keep_structure` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- int `hwloc_topology_ignore_all_keep_structure` (`hwloc_topology_t` topology)
- int `hwloc_topology_set_flags` (`hwloc_topology_t` topology, unsigned long flags)
- unsigned long `hwloc_topology_get_flags` (`hwloc_topology_t` topology)
- int `hwloc_topology_set_pid` (`hwloc_topology_t` restrict topology, `hwloc_pid_t` pid)
- int `hwloc_topology_set_fsroot` (`hwloc_topology_t` restrict topology, const char \*restrict fsroot\_path)
- int `hwloc_topology_set_synthetic` (`hwloc_topology_t` restrict topology, const char \*restrict description)
- int `hwloc_topology_set_xml` (`hwloc_topology_t` restrict topology, const char \*restrict xmlpath)
- int `hwloc_topology_set_xmlbuffer` (`hwloc_topology_t` restrict topology, const char \*restrict buffer, int size)
- int `hwloc_topology_set_custom` (`hwloc_topology_t` topology)
- int `hwloc_topology_set_distance_matrix` (`hwloc_topology_t` restrict topology, `hwloc_obj_type_t` type, unsigned nbobjs, unsigned \*os\_index, float \*distances)
- int `hwloc_topology_is_thissystem` (`hwloc_topology_t` restrict topology)
- struct `hwloc_topology_support` \* `hwloc_topology_get_support` (`hwloc_topology_t` restrict topology)

- void [hwloc\\_topology\\_set\\_userdata](#) ([hwloc\\_topology\\_t](#) topology, const void \*userdata)
- void \* [hwloc\\_topology\\_get\\_userdata](#) ([hwloc\\_topology\\_t](#) topology)

### 19.6.1 Detailed Description

Several functions can optionally be called between [hwloc\\_topology\\_init\(\)](#) and [hwloc\\_topology\\_load\(\)](#) to configure how the detection should be performed, e.g. to ignore some objects types, define a synthetic topology, etc.

If none of them is called, the default is to detect all the objects of the machine that the caller is allowed to access.

This default behavior may also be modified through environment variables if the application did not modify it already. Setting `HWLOC_XMLFILE` in the environment enforces the discovery from a XML file as if [hwloc\\_topology\\_set\\_xml\(\)](#) had been called. `HWLOC_FSROOT` switches to reading the topology from the specified Linux filesystem root as if [hwloc\\_topology\\_set\\_fsroot\(\)](#) had been called. Finally, `HWLOC_THISSYSTEM` enforces the return value of [hwloc\\_topology\\_is\\_thissystem\(\)](#).

### 19.6.2 Enumeration Type Documentation

#### 19.6.2.1 enum [hwloc\\_topology\\_flags\\_e](#)

Flags to be set onto a topology context before load.

Flags should be given to [hwloc\\_topology\\_set\\_flags\(\)](#). They may also be returned by [hwloc\\_topology\\_get\\_flags\(\)](#).

Enumerator:

***HWLOC\_TOPOLOGY\_FLAG\_WHOLE\_SYSTEM*** Detect the whole system, ignore reservations and offline settings. Gather all resources, even if some were disabled by the administrator. For instance, ignore Linux Cgroup/Cpusets and gather all processors and memory nodes, and ignore the fact that some resources may be offline.

When this flag is not set, PUs that are disallowed are not added to the topology. Parent objects (package, core, cache, etc.) are added only if some of their children are allowed. NUMA nodes are always added but their available memory is set to 0 when disallowed.

If the current topology is exported to XML and reimported later, this flag should be set again in the reimported topology so that disallowed resources are reimported as well.

***HWLOC\_TOPOLOGY\_FLAG\_IS\_THISSYSTEM*** Assume that the selected backend provides the topology for the system on which we are running. This forces [hwloc\\_topology\\_is\\_thissystem\(\)](#) to return 1, i.e. makes hwloc assume that the selected backend provides the topology for the system on which we are running, even if it is not the OS-specific backend but the XML backend for instance. This means making the binding functions actually call the OS-specific

system calls and really do binding, while the XML backend would otherwise provide empty hooks just returning success.

Setting the environment variable `HWLOC_THISSYSTEM` may also result in the same behavior.

This can be used for efficiency reasons to first detect the topology once, save it to an XML file, and quickly reload it later through the XML backend, but still having binding functions actually do bind.

**`HWLOC_TOPOLOGY_FLAG_IO_DEVICES`** Detect PCI devices. By default, I/O devices are ignored. This flag enables I/O device detection using the pci backend. Only the common PCI devices (GPUs, NICs, block devices, ...) and host bridges (objects that connect the host objects to an I/O subsystem) will be added to the topology. Additionally it also enables MemoryModule misc objects. Uncommon devices and other bridges (such as PCI-to-PCI bridges) will be ignored.

**`HWLOC_TOPOLOGY_FLAG_IO_BRIDGES`** Detect PCI bridges. This flag should be combined with [HWLOC\\_TOPOLOGY\\_FLAG\\_IO\\_DEVICES](#) to enable the detection of both common devices and of all useful bridges (bridges that have at least one device behind them).

**`HWLOC_TOPOLOGY_FLAG_WHOLE_IO`** Detect the whole PCI hierarchy. This flag enables detection of all I/O devices (even the uncommon ones such as DMA channels) and bridges (even those that have no device behind them) using the pci backend. This implies [HWLOC\\_TOPOLOGY\\_FLAG\\_IO\\_DEVICES](#).

**`HWLOC_TOPOLOGY_FLAG_ICACHES`** Detect instruction caches. This flag enables detection of Instruction caches, instead of only Data and Unified caches.

**`HWLOC_TOPOLOGY_FLAG_THISSYSTEM_ALLOWED_RESOURCES`** Get the set of allowed resources from the local operating system even if the topology was loaded from XML or synthetic description. If the topology was loaded from XML or from a synthetic string, restrict it by applying the current process restrictions such as Linux Cgroup/Cpuset.

This is useful when the topology is not loaded directly from the local machine (e.g. for performance reason) and it comes with all resources, while the running process is restricted to only parts of the machine.

This flag is ignored unless [HWLOC\\_TOPOLOGY\\_FLAG\\_IS\\_THISSYSTEM](#) is also set since the loaded topology must match the underlying machine where restrictions will be gathered from.

Setting the environment variable `HWLOC_THISSYSTEM_ALLOWED_RESOURCES` would result in the same behavior.

### 19.6.3 Function Documentation

#### 19.6.3.1 unsigned long `hwloc_topology_get_flags` ( `hwloc_topology_t topology` )

Get OR'ed flags of a topology.

Get the OR'ed set of [hwloc\\_topology\\_flags\\_e](#) of a topology.

**Returns**

the flags previously set with [hwloc\\_topology\\_set\\_flags\(\)](#).

**19.6.3.2** `struct hwloc_topology_support* hwloc_topology_get_support ( hwloc_topology_t restrict topology )` [read]

Retrieve the topology support.

Each flag indicates whether a feature is supported. If set to 0, the feature is not supported. If set to 1, the feature is supported, but the corresponding call may still fail in some corner cases.

These features are also listed by `hwloc-info --support`

**19.6.3.3** `void* hwloc_topology_get_userdata ( hwloc_topology_t topology )`

Retrieve the topology-specific userdata pointer.

Retrieve the application-given private data pointer that was previously set with [hwloc\\_topology\\_set\\_userdata\(\)](#).

**19.6.3.4** `int hwloc_topology_ignore_all_keep_structure ( hwloc_topology_t topology )`

Ignore all objects that do not bring any structure.

Ignore all objects that do not bring any structure: This is equivalent to calling [hwloc\\_topology\\_ignore\\_type\\_keep\\_structure\(\)](#) for all object types.

**19.6.3.5** `int hwloc_topology_ignore_type ( hwloc_topology_t topology, hwloc_obj_type_t type )`

Ignore an object type.

Ignore all objects from the given type. The bottom-level type `HWLOC_OBJ_PU` may not be ignored. The top-level object of the hierarchy will never be ignored, even if this function succeeds. Group objects are always ignored if they do not bring any structure since they are designed to add structure to the topology. I/O objects may not be ignored, topology flags should be used to configure their discovery instead.

**19.6.3.6** `int hwloc_topology_ignore_type_keep_structure ( hwloc_topology_t topology, hwloc_obj_type_t type )`

Ignore an object type if it does not bring any structure.

Ignore all objects from the given type as long as they do not bring any structure: - Each ignored object should have a single children or be the only child of its parent.

The bottom-level type `HWLOC_OBJ_PU` may not be ignored. I/O objects may not be ignored, topology flags should be used to configure their discovery instead.

#### 19.6.3.7 `int hwloc_topology_is_thissystem ( hwloc_topology_t restrict topology )`

Does the topology context come from this system?

##### Returns

1 if this topology context was built using the system running this program.  
0 instead (for instance if using another file-system root, a XML topology file, or a synthetic topology).

#### 19.6.3.8 `int hwloc_topology_set_custom ( hwloc_topology_t topology )`

Prepare the topology for custom assembly.

The topology then contains a single root object. It must then be built by inserting other topologies with `hwloc_custom_insert_topology()` or single objects with `hwloc_custom_insert_group_object_by_parent()`. `hwloc_topology_load()` must be called to finalize the new topology as usual.

##### Note

If nothing is inserted in the topology, `hwloc_topology_load()` will fail with `errno` set to `EINVAL`.

The `cpuset` and `nodeset` of the root object are `NULL` because these sets are meaningless when assembling multiple topologies.

On success, the custom component replaces the previously enabled component (if any), but the topology is not actually modified until `hwloc_topology_load()`.

#### 19.6.3.9 `int hwloc_topology_set_distance_matrix ( hwloc_topology_t restrict topology, hwloc_obj_type_t type, unsigned nbobjs, unsigned * os_index, float * distances )`

Provide a distance matrix.

Provide the matrix of distances between a set of objects of the given `type`. `nbobjs` must be at least 2. The set may or may not contain all the existing objects of this type. The objects are specified by their OS/physical index in the `os_index` array. The `distances` matrix follows the same order. The distance from object `i` to object `j` in the `i*nbobjs+j`.

A single latency matrix may be defined for each type. If another distance matrix already exists for the given type, either because the user specified it or because the OS offers it, it will be replaced by the given one. If `nbobjs` is 0, `os_index` is `NULL` and `distances` is `NULL`, the existing distance matrix for the given type is removed.

**Note**

Distance matrices are ignored in multi-node topologies.

**19.6.3.10** `int hwloc_topology_set_flags ( hwloc_topology_t topology, unsigned long flags )`

Set OR'ed flags to non-yet-loaded topology.

Set a OR'ed set of [hwloc\\_topology\\_flags\\_e](#) onto a topology that was not yet loaded.

If this function is called multiple times, the last invocation will erase and replace the set of flags that was previously set.

The flags set in a topology may be retrieved with [hwloc\\_topology\\_get\\_flags\(\)](#)

**19.6.3.11** `int hwloc_topology_set_fsroot ( hwloc_topology_t restrict topology, const char *restrict fsroot_path )`

Change the file-system root path when building the topology from sysfs/procfs.

On Linux system, use sysfs and procfs files as if they were mounted on the given `fsroot_path` instead of the main file-system root. Setting the environment variable `HWLOC_FSROOT` may also result in this behavior. Not using the main file-system root causes [hwloc\\_topology\\_is\\_thissystem\(\)](#) to return 0.

Note that this function does not actually load topology information; it just tells hwloc where to load it from. You'll still need to invoke [hwloc\\_topology\\_load\(\)](#) to actually load the topology information.

**Returns**

-1 with `errno` set to `ENOSYS` on non-Linux and on Linux systems that do not support it.

-1 with the appropriate `errno` if `fsroot_path` cannot be used.

**Note**

For convenience, this backend provides empty binding hooks which just return success. To have hwloc still actually call OS-specific hooks, the [HWLOC\\_TOPOLOGY\\_FLAG\\_IS\\_THISSYSTEM](#) has to be set to assert that the loaded file is really the underlying system.

On success, the Linux component replaces the previously enabled component (if any), but the topology is not actually modified until [hwloc\\_topology\\_load\(\)](#).

**19.6.3.12** `int hwloc_topology_set_pid ( hwloc_topology_t restrict topology, hwloc_pid_t pid )`

Change which process the topology is viewed from.

On some systems, processes may have different views of the machine, for instance the set of allowed CPUs. By default, hwloc exposes the view from the current process. Calling [hwloc\\_topology\\_set\\_pid\(\)](#) permits to make it expose the topology of the machine from the point of view of another process.

**Note**

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

-1 is returned and `errno` is set to `ENOSYS` on platforms that do not support this feature.

**19.6.3.13** `int hwloc_topology_set_synthetic ( hwloc_topology_t restrict topology, const char *restrict description )`

Enable synthetic topology.

Gather topology information from the given `description`, a space-separated string of numbers describing the arity of each level. Each number may be prefixed with a type and a colon to enforce the type of a level. If only some level types are enforced, hwloc will try to choose the other types according to usual topologies, but it may fail and you may have to specify more level types manually. See also the [Synthetic topologies](#).

If `description` was properly parsed and describes a valid topology configuration, this function returns 0. Otherwise -1 is returned and `errno` is set to `EINVAL`.

Note that this function does not actually load topology information; it just tells hwloc where to load it from. You'll still need to invoke [hwloc\\_topology\\_load\(\)](#) to actually load the topology information.

**Note**

For convenience, this backend provides empty binding hooks which just return success.

On success, the synthetic component replaces the previously enabled component (if any), but the topology is not actually modified until [hwloc\\_topology\\_load\(\)](#).

**19.6.3.14** `void hwloc_topology_set_userdata ( hwloc_topology_t topology, const void * userdata )`

Set the topology-specific userdata pointer.

Each topology may store one application-given private data pointer. It is initialized to `NULL`. hwloc will never modify it.

Use it as you wish, after [hwloc\\_topology\\_init\(\)](#) and until [hwloc\\_topolog\\_destroy\(\)](#).

This pointer is not exported to XML.

**19.6.3.15** `int hwloc_topology_set_xml ( hwloc_topology_t restrict topology, const char *restrict xmlpath )`

Enable XML-file based topology.

Gather topology information from the XML file given at `xmlpath`. Setting the environment variable `HWLOC_XMLFILE` may also result in this behavior. This file may have been generated earlier with `hwloc_topology_export_xml()` or `lstopo file.xml`.

Note that this function does not actually load topology information; it just tells `hwloc` where to load it from. You'll still need to invoke `hwloc_topology_load()` to actually load the topology information.

#### Returns

-1 with `errno` set to `EINVAL` on failure to read the XML file.

#### Note

See also `hwloc_topology_set_userdata_import_callback()` for importing application-specific object userdata.

For convenience, this backend provides empty binding hooks which just return success. To have `hwloc` still actually call OS-specific hooks, the `HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM` has to be set to assert that the loaded file is really the underlying system.

On success, the XML component replaces the previously enabled component (if any), but the topology is not actually modified until `hwloc_topology_load()`.

**19.6.3.16** `int hwloc_topology_set_xmlbuffer ( hwloc_topology_t restrict topology, const char *restrict buffer, int size )`

Enable XML based topology using a memory buffer (instead of a file, as with `hwloc_topology_set_xml()`).

Gather topology information from the XML memory buffer given at `buffer` and of length `size`. This buffer may have been filled earlier with `hwloc_topology_export_xmlbuffer()`.

Note that this function does not actually load topology information; it just tells `hwloc` where to load it from. You'll still need to invoke `hwloc_topology_load()` to actually load the topology information.

#### Returns

-1 with `errno` set to `EINVAL` on failure to read the XML buffer.



**Note**

See also [hwloc\\_topology\\_set\\_userdata\\_import\\_callback\(\)](#) for importing application-specific object userdata.

For convenience, this backend provides empty binding hooks which just return success. To have hwloc still actually call OS-specific hooks, the [HWLOC\\_TOPOLOGY\\_FLAG\\_IS\\_THISSYSTEM](#) has to be set to assert that the loaded file is really the underlying system.

On success, the XML component replaces the previously enabled component (if any), but the topology is not actually modified until [hwloc\\_topology\\_load\(\)](#).

## 19.7 Object levels, depths and types

### Enumerations

- enum `hwloc_get_type_depth_e` { `HWLOC_TYPE_DEPTH_UNKNOWN`, `HWLOC_TYPE_DEPTH_MULTIPLE`, `HWLOC_TYPE_DEPTH_BRIDGE`, `HWLOC_TYPE_DEPTH_PCI_DEVICE`, `HWLOC_TYPE_DEPTH_OS_DEVICE` }

### Functions

- unsigned `hwloc_topology_get_depth` (`hwloc_topology_t` restrict topology)
- int `hwloc_get_type_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- static int `hwloc_get_type_or_below_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- static int `hwloc_get_type_or_above_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- `hwloc_obj_type_t` `hwloc_get_depth_type` (`hwloc_topology_t` topology, unsigned depth)
- unsigned `hwloc_get_nbobjs_by_depth` (`hwloc_topology_t` topology, unsigned depth)
- static int `hwloc_get_nbobjs_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- static `hwloc_obj_t` `hwloc_get_root_obj` (`hwloc_topology_t` topology)
- `hwloc_obj_t` `hwloc_get_obj_by_depth` (`hwloc_topology_t` topology, unsigned depth, unsigned idx)
- static `hwloc_obj_t` `hwloc_get_obj_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type, unsigned idx)
- static `hwloc_obj_t` `hwloc_get_next_obj_by_depth` (`hwloc_topology_t` topology, unsigned depth, `hwloc_obj_t` prev)
- static `hwloc_obj_t` `hwloc_get_next_obj_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type, `hwloc_obj_t` prev)

### 19.7.1 Detailed Description

Be sure to see the figure in [Terms and Definitions](#) that shows a complete topology tree, including depths, child/sibling/cousin relationships, and an example of an asymmetric topology where one package has fewer caches than its peers.

### 19.7.2 Enumeration Type Documentation

#### 19.7.2.1 enum `hwloc_get_type_depth_e`

Enumerator:

**`HWLOC_TYPE_DEPTH_UNKNOWN`** No object of given type exists in the topology.

**HWLOC\_TYPE\_DEPTH\_MULTIPLE** Objects of given type exist at different depth in the topology.

**HWLOC\_TYPE\_DEPTH\_BRIDGE** Virtual depth for bridge object level.

**HWLOC\_TYPE\_DEPTH\_PCI\_DEVICE** Virtual depth for PCI device object level.

**HWLOC\_TYPE\_DEPTH\_OS\_DEVICE** Virtual depth for software device object level.

### 19.7.3 Function Documentation

19.7.3.1 `hwloc_obj_type_t hwloc_get_depth_type ( hwloc_topology_t topology, unsigned depth )`

Returns the type of objects at depth `depth`.

`depth` should be between 0 and `hwloc_topology_get_depth()-1`.

#### Returns

-1 if depth `depth` does not exist.

19.7.3.2 `unsigned hwloc_get_nobjs_by_depth ( hwloc_topology_t topology, unsigned depth )`

Returns the width of level at depth `depth`.

19.7.3.3 `static int hwloc_get_nobjs_by_type ( hwloc_topology_t topology, hwloc_obj_type_t type ) [inline, static]`

Returns the width of level `type`.

If no object for that type exists, 0 is returned. If there are several levels with objects of that type, -1 is returned.

19.7.3.4 `static hwloc_obj_t hwloc_get_next_obj_by_depth ( hwloc_topology_t topology, unsigned depth, hwloc_obj_t prev ) [inline, static]`

Returns the next object at depth `depth`.

If `prev` is NULL, return the first object at depth `depth`.

19.7.3.5 `static hwloc_obj_t hwloc_get_next_obj_by_type ( hwloc_topology_t topology, hwloc_obj_type_t type, hwloc_obj_t prev ) [inline, static]`

Returns the next object of type `type`.

If `prev` is NULL, return the first object at type `type`. If there are multiple or no depth for given type, return NULL and let the caller fallback to `hwloc_get_next_obj_by_depth()`.

19.7.3.6 `hwloc_obj_t hwloc_get_obj_by_depth ( hwloc_topology_t topology, unsigned depth, unsigned idx )`

Returns the topology object at logical index `idx` from depth `depth`.

19.7.3.7 `static hwloc_obj_t hwloc_get_obj_by_type ( hwloc_topology_t topology, hwloc_obj_type_t type, unsigned idx ) [inline, static]`

Returns the topology object at logical index `idx` with type `type`.

If no object for that type exists, `NULL` is returned. If there are several levels with objects of that type, `NULL` is returned and the caller may fallback to [hwloc\\_get\\_obj\\_by\\_depth\(\)](#).

19.7.3.8 `static hwloc_obj_t hwloc_get_root_obj ( hwloc_topology_t topology ) [inline, static]`

Returns the top-object of the topology-tree.

Its type is typically `HWLOC_OBJ_MACHINE` but it could be different for complex topologies.

19.7.3.9 `int hwloc_get_type_depth ( hwloc_topology_t topology, hwloc_obj_type_t type )`

Returns the depth of objects of type `type`.

If no object of this type is present on the underlying architecture, or if the OS doesn't provide this kind of information, the function returns `HWLOC_TYPE_DEPTH_UNKNOWN`.

If type is absent but a similar type is acceptable, see also [hwloc\\_get\\_type\\_or\\_below\\_depth\(\)](#) and [hwloc\\_get\\_type\\_or\\_above\\_depth\(\)](#).

If some objects of the given type exist in different levels, for instance L1 and L2 caches, or L1i and L1d caches, the function returns `HWLOC_TYPE_DEPTH_MULTIPLE`. See [hwloc\\_get\\_cache\\_type\\_depth\(\)](#) in `hwloc/helper.h` to better handle this case.

If an I/O object type is given, the function returns a virtual value because I/O objects are stored in special levels that are not CPU-related. This virtual depth may be passed to other hwloc functions such as [hwloc\\_get\\_obj\\_by\\_depth\(\)](#) but it should not be considered as an actual depth by the application. In particular, it should not be compared with any other object depth or with the entire topology depth.

If `HWLOC_OBJ_MISC` is given, the function returns `HWLOC_TYPE_DEPTH_UNKNOWN`.

19.7.3.10 `static int hwloc_get_type_or_above_depth ( hwloc_topology_t topology,  
hwloc_obj_type_t type ) [inline, static]`

Returns the depth of objects of type `type` or above.

If no object of this type is present on the underlying architecture, the function returns the depth of the first "present" object typically containing `type`.

This function is only meaningful for normal object types. If an I/O object type is given, the corresponding virtual depth is always returned (see [hwloc\\_get\\_type\\_depth\(\)](#)). If [HWLOC\\_OBJ\\_MISC](#) is given, the function returns [HWLOC\\_TYPE\\_DEPTH\\_UNKNOWN](#).

If some objects of the given type exist in different levels, for instance L1 and L2 caches, the function returns [HWLOC\\_TYPE\\_DEPTH\\_MULTIPLE](#).

19.7.3.11 `static int hwloc_get_type_or_below_depth ( hwloc_topology_t topology,  
hwloc_obj_type_t type ) [inline, static]`

Returns the depth of objects of type `type` or below.

If no object of this type is present on the underlying architecture, the function returns the depth of the first "present" object typically found inside `type`.

This function is only meaningful for normal object types. If an I/O object type is given, the corresponding virtual depth is always returned (see [hwloc\\_get\\_type\\_depth\(\)](#)). If [HWLOC\\_OBJ\\_MISC](#) is given, the function returns [HWLOC\\_TYPE\\_DEPTH\\_UNKNOWN](#).

If some objects of the given type exist in different levels, for instance L1 and L2 caches, the function returns [HWLOC\\_TYPE\\_DEPTH\\_MULTIPLE](#).

19.7.3.12 `unsigned hwloc_topology_get_depth ( hwloc_topology_t restrict topology )`

Get the depth of the hierarchical tree of objects.

This is the depth of [HWLOC\\_OBJ\\_PU](#) objects plus one.

#### Note

I/O and Misc objects are ignored when computing the depth of the tree (they are placed on special levels, or none).

## 19.8 Converting between Object Types, Sets and Attributes, and - Strings

### Functions

- const char \* [hwloc\\_obj\\_type\\_string](#) ([hwloc\\_obj\\_type\\_t](#) type)
- int [hwloc\\_obj\\_type\\_snprintf](#) (char \*restrict string, size\_t size, [hwloc\\_obj\\_t](#) obj, int verbose)
- int [hwloc\\_obj\\_attr\\_snprintf](#) (char \*restrict string, size\_t size, [hwloc\\_obj\\_t](#) obj, const char \*restrict separator, int verbose)
- int [hwloc\\_obj\\_cpuset\\_snprintf](#) (char \*restrict str, size\_t size, size\_t nobj, const [hwloc\\_obj\\_t](#) \*restrict objs)
- int [hwloc\\_obj\\_type\\_sscanf](#) (const char \*string, [hwloc\\_obj\\_type\\_t](#) \*typep, int \*depthattrp, void \*typeattrp, size\_t typeattrsize)

### 19.8.1 Function Documentation

**19.8.1.1** int [hwloc\\_obj\\_attr\\_snprintf](#) ( char \*restrict *string*, size\_t *size*, [hwloc\\_obj\\_t](#) *obj*, const char \*restrict *separator*, int *verbose* )

Stringify the attributes of a given topology object into a human-readable form.

Attribute values are separated by *separator*.

Only the major attributes are printed in non-verbose mode.

If *size* is 0, *string* may safely be NULL.

#### Returns

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

**19.8.1.2** int [hwloc\\_obj\\_cpuset\\_snprintf](#) ( char \*restrict *str*, size\_t *size*, size\_t *nobj*, const [hwloc\\_obj\\_t](#) \*restrict *objs* )

Stringify the cpuset containing a set of objects.

If *size* is 0, *string* may safely be NULL.

#### Returns

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

**19.8.1.3** `int hwloc_obj_type_sprintf ( char *restrict string, size_t size, hwloc_obj_t obj, int verbose )`

Stringify the type of a given topology object into a human-readable form.

Contrary to `hwloc_obj_type_string()`, this function includes object-specific attributes (such as the Group depth, the Bridge type, or OS device type) in the output, and it requires the caller to provide the output buffer.

The output is guaranteed to be the same for all objects of a same topology level.

If `size` is 0, `string` may safely be `NULL`.

#### Returns

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

**19.8.1.4** `int hwloc_obj_type_sscanf ( const char * string, hwloc_obj_type_t * typep, int * depthattrp, void * typeattrp, size_t typeattrsize )`

Return an object type and attributes from a type string.

Convert strings such as "Package" or "Cache" into the corresponding types. Matching is case-insensitive, and only the first letters are actually required to match.

This function is guaranteed to match any string returned by `hwloc_obj_type_string()` or `hwloc_obj_type_sprintf()`.

Types that have specific attributes, for instance caches and groups, may be returned in `depthattrp` and `typeattrp`. They are ignored when these pointers are `NULL`.

For instance "L2i" or "L2iCache" would return type `HWLOC_OBJ_CACHE` in `typep`, 2 in `depthattrp`, and `HWLOC_OBJ_CACHE_TYPE_INSTRUCTION` in `typeattrp` (this last pointer should point to a `hwloc_obj_cache_type_t`). "Group3" would return type `HWLOC_OBJ_GROUP` type and 3 in `depthattrp`. Attributes that are not specified in the string (for instance "Group" without a depth, or "L2Cache" without a cache type) are set to -1.

`typeattrp` is only filled if the size specified in `typeattrsize` is large enough. It is currently only used for caches, and the required size is at least the size of `hwloc_obj_cache_type_t`.

#### Returns

0 if a type was correctly identified, otherwise -1.

#### Note

This is an extended version of the now deprecated `hwloc_obj_type_of_string()`

**19.8.1.5** `const char* hwloc_obj_type_string ( hwloc_obj_type_t type )`

Return a constant stringified object type.

This function is the basic way to convert a generic type into a string.

[hwloc\\_obj\\_type\\_snprintf\(\)](#) may return a more precise output for a specific object, but it requires the caller to provide the output buffer.



## 19.9 Consulting and Adding Key-Value Info Attributes

### Functions

- static const char \* [hwloc\\_obj\\_get\\_info\\_by\\_name](#) ([hwloc\\_obj\\_t](#) obj, const char \*name)
- void [hwloc\\_obj\\_add\\_info](#) ([hwloc\\_obj\\_t](#) obj, const char \*name, const char \*value)

### 19.9.1 Function Documentation

19.9.1.1 void [hwloc\\_obj\\_add\\_info](#) ( [hwloc\\_obj\\_t](#) obj, const char \* name, const char \* value )

Add the given info name and value pair to the given object.

The info is appended to the existing info array even if another key with the same name already exists.

The input strings are copied before being added in the object infos.

#### Note

This function may be used to enforce object colors in the lstopo graphical output by using "lstopoStyle" as a name and "Background=#rrggbb" as a value. See CUSTOM COLORS in the lstopo(1) manpage for details.

If `value` contains some non-printable characters, they will be dropped when exporting to XML, see [hwloc\\_topology\\_export\\_xml\(\)](#).

19.9.1.2 static const char\* [hwloc\\_obj\\_get\\_info\\_by\\_name](#) ( [hwloc\\_obj\\_t](#) obj, const char \* name ) [inline, static]

Search the given key name in object infos and return the corresponding value.

If multiple keys match the given name, only the first one is returned.

#### Returns

NULL if no such key exists.

## 19.10 CPU binding

### Enumerations

- enum `hwloc_cpubind_flags_t` { `HWLOC_CPUBIND_PROCESS`, `HWLOC_CPUBIND_THREAD`, `HWLOC_CPUBIND_STRICT`, `HWLOC_CPUBIND_NOMEMBIND` }

### Functions

- int `hwloc_set_cpubind` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, int flags)
- int `hwloc_get_cpubind` (`hwloc_topology_t` topology, `hwloc_cpuset_t` set, int flags)
- int `hwloc_set_proc_cpubind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_cpuset_t` set, int flags)
- int `hwloc_get_proc_cpubind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_cpuset_t` set, int flags)
- int `hwloc_set_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` thread, `hwloc_const_cpuset_t` set, int flags)
- int `hwloc_get_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` thread, `hwloc_cpuset_t` set, int flags)
- int `hwloc_get_last_cpu_location` (`hwloc_topology_t` topology, `hwloc_cpuset_t` set, int flags)
- int `hwloc_get_proc_last_cpu_location` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_cpuset_t` set, int flags)

#### 19.10.1 Detailed Description

Some operating systems only support binding threads or processes to a single PU. - Others allow binding to larger sets such as entire Cores or Packages or even random sets of individual PUs. In such operating system, the scheduler is free to run the task on one of these PU, then migrate it to another PU, etc. It is often useful to call `hwloc_bitmap_singlify()` on the target CPU set before passing it to the binding function to avoid these expensive migrations. See the documentation of `hwloc_bitmap_singlify()` for details.

Some operating systems do not provide all hwloc-supported mechanisms to bind processes, threads, etc. `hwloc_topology_get_support()` may be used to query about the actual CPU binding support in the currently used operating system.

When the requested binding operation is not available and the `HWLOC_CPUBIND_STRICT` flag was passed, the function returns -1. `errno` is set to `ENOSYS` when it is not possible to bind the requested kind of object processes/threads. `errno` is set to `EXDEV` when the requested cpuset can not be enforced (e.g. some systems only allow one CPU, and some other systems only allow one NUMA node).

If `HWLOC_CPUBIND_STRICT` was not passed, the function may fail as well, or the operating system may use a slightly different operation (with side-effects, smaller binding set, etc.) when the requested operation is not exactly supported.

The most portable version that should be preferred over the others, whenever possible, is the following one which just binds the current program, assuming it is single-threaded:

```
hwloc_set_cpubind(topology, set, 0),
```

If the program may be multithreaded, the following one should be preferred to only bind the current thread:

```
hwloc_set_cpubind(topology, set, HWLOC_CPUBIND_THREAD),
```

#### See also

Some example codes are available under `doc/examples/` in the source tree.

#### Note

To unbind, just call the binding function with either a full cpuset or a cpuset equal to the system cpuset.

On some operating systems, CPU binding may have effects on memory binding, see [HWLOC\\_CPUBIND\\_NOMEMBIND](#)

Running `lstopo --top` or `hwloc-ps` can be a very convenient tool to check how binding actually happened.

## 19.10.2 Enumeration Type Documentation

### 19.10.2.1 enum `hwloc_cpubind_flags_t`

Process/Thread binding flags.

These bit flags can be used to refine the binding policy.

The default (0) is to bind the current process, assumed to be single-threaded, in a non-strict way. This is the most portable way to bind as all operating systems usually provide it.

#### Note

Not all systems support all kinds of binding. See the "Detailed Description" section of [CPU binding](#) for a description of errors that can occur.

#### Enumerator:

***HWLOC\_CPUBIND\_PROCESS*** Bind all threads of the current (possibly) multi-threaded process.

***HWLOC\_CPUBIND\_THREAD*** Bind current thread of current process.

***HWLOC\_CPUBIND\_STRICT*** Request for strict binding from the OS. By default, when the designated CPUs are all busy while other CPUs are idle, operating systems may execute the thread/process on those other CPUs instead of the designated CPUs, to let them progress anyway. Strict binding means that the thread/process will *never* execute on other cpus than the designated CPUs, even when those are busy with other tasks and other CPUs are idle.

**Note**

Depending on the operating system, strict binding may not be possible (e.g., the OS does not implement it) or not allowed (e.g., for administrative reasons), and the function will fail in that case.

When retrieving the binding of a process, this flag checks whether all its threads actually have the same binding. If the flag is not given, the binding of each thread will be accumulated.

**Note**

This flag is meaningless when retrieving the binding of a thread.

**HWLOC\_CPUBIND\_NOMEMBIND** Avoid any effect on memory binding. On some operating systems, some CPU binding function would also bind the memory on the corresponding NUMA node. It is often not a problem for the application, but if it is, setting this flag will make hwloc avoid using OS functions that would also bind memory. This will however reduce the support of CPU bindings, i.e. potentially return -1 with errno set to ENOSYS in some cases.

This flag is only meaningful when used with functions that set the CPU binding. It is ignored when used with functions that get CPU binding information.

**19.10.3 Function Documentation**

**19.10.3.1** `int hwloc_get_cpupind ( hwloc_topology_t topology, hwloc_cpuset_t set, int flags )`

Get current process or thread binding.

Writes into `set` the physical cpuset which the process or thread (according to `flags`) was last bound to.

**19.10.3.2** `int hwloc_get_last_cpu_location ( hwloc_topology_t topology, hwloc_cpuset_t set, int flags )`

Get the last physical CPU where the current process or thread ran.

The operating system may move some tasks from one processor to another at any time according to their binding, so this function may return something that is already outdated.

`flags` can include either `HWLOC_CPUBIND_PROCESS` or `HWLOC_CPUBIND_THREAD` to specify whether the query should be for the whole process (union of all CPUs on which all threads are running), or only the current thread. If the process is single-threaded, flags can be set to zero to let hwloc use whichever method is available on the underlying OS.

**19.10.3.3** `int hwloc_get_proc_cpupind ( hwloc_topology_t topology, hwloc_pid_t pid, hwloc_cpuset_t set, int flags )`

Get the current physical binding of process `pid`.

**Note**

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

As a special case on Linux, if a `tid` (thread ID) is supplied instead of a `pid` (process ID) and `HWLOC_CPUBIND_THREAD` is passed in `flags`, the binding for that specific thread is returned.

On non-Linux systems, `HWLOC_CPUBIND_THREAD` can not be used in `flags`.

**19.10.3.4** `int hwloc_get_proc_last_cpu_location ( hwloc_topology_t topology,  
hwloc_pid_t pid, hwloc_cpuset_t set, int flags )`

Get the last physical CPU where a process ran.

The operating system may move some tasks from one processor to another at any time according to their binding, so this function may return something that is already outdated.

**Note**

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

As a special case on Linux, if a `tid` (thread ID) is supplied instead of a `pid` (process ID) and `HWLOC_CPUBIND_THREAD` is passed in `flags`, the last CPU location of that specific thread is returned.

On non-Linux systems, `HWLOC_CPUBIND_THREAD` can not be used in `flags`.

**19.10.3.5** `int hwloc_get_thread_cpupbind ( hwloc_topology_t topology, hwloc_thread_t  
thread, hwloc_cpuset_t set, int flags )`

Get the current physical binding of thread `tid`.

**Note**

`hwloc_thread_t` is `pthread_t` on Unix platforms, and `HANDLE` on native Windows platforms.

`HWLOC_CPUBIND_PROCESS` can not be used in `flags`.

**19.10.3.6** `int hwloc_set_cpupbind ( hwloc_topology_t topology,  
hwloc_const_cpuset_t set, int flags )`

Bind current process or thread on cpus given in physical bitmap `set`.

**Returns**

- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

19.10.3.7 `int hwloc_set_proc_cpubind ( hwloc_topology_t topology, hwloc_pid_t pid, hwloc_const_cpuset_t set, int flags )`

Bind a process `pid` on cpus given in physical bitmap `set`.

**Note**

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

As a special case on Linux, if a `tid` (thread ID) is supplied instead of a `pid` (process ID) and `HWLOC_CPUBIND_THREAD` is passed in `flags`, the binding is applied to that specific thread.

On non-Linux systems, `HWLOC_CPUBIND_THREAD` can not be used in `flags`.

19.10.3.8 `int hwloc_set_thread_cpubind ( hwloc_topology_t topology, hwloc_thread_t thread, hwloc_const_cpuset_t set, int flags )`

Bind a thread `thread` on cpus given in physical bitmap `set`.

**Note**

`hwloc_thread_t` is `pthread_t` on Unix platforms, and `HANDLE` on native Windows platforms.

`HWLOC_CPUBIND_PROCESS` can not be used in `flags`.

## 19.11 Memory binding

### Enumerations

- enum `hwloc_membind_policy_t` { `HWLOC_MEMBIND_DEFAULT`, `HWLOC_MEMBIND_FIRSTTOUCH`, `HWLOC_MEMBIND_BIND`, `HWLOC_MEMBIND_INTERLEAVE`, `HWLOC_MEMBIND_REPLICATE`, `HWLOC_MEMBIND_NEXTTOUCH`, `HWLOC_MEMBIND_MIXED` }
- enum `hwloc_membind_flags_t` { `HWLOC_MEMBIND_PROCESS`, `HWLOC_MEMBIND_THREAD`, `HWLOC_MEMBIND_STRICT`, `HWLOC_MEMBIND_MIGRATE`, `HWLOC_MEMBIND_NOCPUBIND`, `HWLOC_MEMBIND_BYNODESET` }

### Functions

- int `hwloc_set_membind_nodese`t (`hwloc_topology_t` topology, `hwloc_const_nodese`t\_t nodeset, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_set_membind` (`hwloc_topology_t` topology, `hwloc_const_bitmap_t` set, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_get_membind_nodese`t (`hwloc_topology_t` topology, `hwloc_nodese`t\_t nodeset, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_get_membind` (`hwloc_topology_t` topology, `hwloc_bitmap_t` set, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_set_proc_membind_nodese`t (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_nodese`t\_t nodeset, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_set_proc_membind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_bitmap_t` set, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_get_proc_membind_nodese`t (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_nodese`t\_t nodeset, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_get_proc_membind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_bitmap_t` set, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_set_area_membind_nodese`t (`hwloc_topology_t` topology, const void \*addr, size\_t len, `hwloc_const_nodese`t\_t nodeset, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_set_area_membind` (`hwloc_topology_t` topology, const void \*addr, size\_t len, `hwloc_const_bitmap_t` set, `hwloc_membind_policy_t` policy, int flags)
- int `hwloc_get_area_membind_nodese`t (`hwloc_topology_t` topology, const void \*addr, size\_t len, `hwloc_nodese`t\_t nodeset, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_get_area_membind` (`hwloc_topology_t` topology, const void \*addr, size\_t len, `hwloc_bitmap_t` set, `hwloc_membind_policy_t` \*policy, int flags)
- int `hwloc_get_area_memlocation` (`hwloc_topology_t` topology, const void \*addr, size\_t len, `hwloc_bitmap_t` set, int flags)
- void \* `hwloc_alloc` (`hwloc_topology_t` topology, size\_t len)
- void \* `hwloc_alloc_membind_nodese`t (`hwloc_topology_t` topology, size\_t len, `hwloc_const_nodese`t\_t nodeset, `hwloc_membind_policy_t` policy, int flags)
- void \* `hwloc_alloc_membind` (`hwloc_topology_t` topology, size\_t len, `hwloc_const_bitmap_t` set, `hwloc_membind_policy_t` policy, int flags)

- static void \* [hwloc\\_alloc\\_membind\\_policy\\_nodeset](#) ([hwloc\\_topology\\_t](#) topology, size\_t len, [hwloc\\_const\\_nodeset\\_t](#) nodeset, [hwloc\\_membind\\_policy\\_t](#) policy, int flags)
- static void \* [hwloc\\_alloc\\_membind\\_policy](#) ([hwloc\\_topology\\_t](#) topology, size\_t len, [hwloc\\_const\\_bitmap\\_t](#) set, [hwloc\\_membind\\_policy\\_t](#) policy, int flags)
- int [hwloc\\_free](#) ([hwloc\\_topology\\_t](#) topology, void \*addr, size\_t len)

### 19.11.1 Detailed Description

Memory binding can be done three ways:

- explicit memory allocation thanks to [hwloc\\_alloc\\_membind\(\)](#) and friends: the binding will have effect on the memory allocated by these functions.
- implicit memory binding through binding policy: [hwloc\\_set\\_membind\(\)](#) and friends only define the current policy of the process, which will be applied to the subsequent calls to `malloc()` and friends.
- migration of existing memory ranges, thanks to [hwloc\\_set\\_area\\_membind\(\)](#) and friends, which move already-allocated data.

Not all operating systems support all three ways. [hwloc\\_topology\\_get\\_support\(\)](#) may be used to query about the actual memory binding support in the currently used operating system.

When the requested binding operation is not available and the [HWLOC\\_MEMBIND\\_STRICT](#) flag was passed, the function returns -1. `errno` will be set to `ENOSYS` when the system does support the specified action or policy (e.g., some systems only allow binding memory on a per-thread basis, whereas other systems only allow binding memory for all threads in a process). `errno` will be set to `EXDEV` when the requested set can not be enforced (e.g., some systems only allow binding memory to a single NUMA node).

If [HWLOC\\_MEMBIND\\_STRICT](#) was not passed, the function may fail as well, or the operating system may use a slightly different operation (with side-effects, smaller binding set, etc.) when the requested operation is not exactly supported.

The most portable form that should be preferred over the others whenever possible is as follows. It allocates some memory hopefully bound to the specified set. To do so, `hwloc` will possibly have to change the current memory binding policy in order to actually get the memory bound, if the OS does not provide any other way to simply allocate bound memory without changing the policy for all allocations. That is the difference with [hwloc\\_alloc\\_membind\(\)](#), which will never change the current memory binding policy.

```
hwloc_alloc_membind_policy(topology, size, set,
                           HWLOC_MEMBIND_BIND, 0);
```

Each `hwloc` memory binding function is available in two forms: one that takes a bitmap argument (a CPU set by default, or a NUMA memory node set if the flag [HWLOC\\_MEMBIND\\_BYNODESET](#) is specified), and another one (whose name ends with `_nodeset`) that always takes a NUMA memory node set. See [Object Sets \(`hwloc\_cpuset\_t` and](#)



[hwloc\\_nodese\\_t](#)) and [The bitmap API](#) for a discussion of CPU sets and NUMA memory node sets. It is also possible to convert between CPU set and node set using [hwloc\\_cpuset\\_to\\_nodese\\_t\(\)](#) or [hwloc\\_cpuset\\_from\\_nodese\\_t\(\)](#).

Memory binding by CPU set cannot work for CPU-less NUMA memory nodes. Binding by node set should therefore be preferred whenever possible.

#### See also

Some example codes are available under `doc/examples/` in the source tree.

#### Note

On some operating systems, memory binding affects the CPU binding; see [HWLOC\\_MEMBIND\\_NOCPUBIND](#)

## 19.11.2 Enumeration Type Documentation

### 19.11.2.1 enum `hwloc_membind_flags_t`

Memory binding flags.

These flags can be used to refine the binding policy. All flags can be logically OR'ed together with the exception of [HWLOC\\_MEMBIND\\_PROCESS](#) and [HWLOC\\_MEMBIND\\_THREAD](#); these two flags are mutually exclusive.

Not all systems support all kinds of binding. [hwloc\\_topology\\_get\\_support\(\)](#) may be used to query about the actual memory binding support in the currently used operating system. See the "Detailed Description" section of [Memory binding](#) for a description of errors that can occur.

#### Enumerator:

**`HWLOC_MEMBIND_PROCESS`** Set policy for all threads of the specified (possibly multithreaded) process. This flag is mutually exclusive with [HWLOC\\_MEMBIND\\_THREAD](#).

**`HWLOC_MEMBIND_THREAD`** Set policy for a specific thread of the current process. This flag is mutually exclusive with [HWLOC\\_MEMBIND\\_PROCESS](#).

**`HWLOC_MEMBIND_STRICT`** Request strict binding from the OS. The function will fail if the binding can not be guaranteed / completely enforced.

This flag has slightly different meanings depending on which function it is used with.

**`HWLOC_MEMBIND_MIGRATE`** Migrate existing allocated memory. If the memory cannot be migrated and the [HWLOC\\_MEMBIND\\_STRICT](#) flag is passed, an error will be returned.

**`HWLOC_MEMBIND_NOCPUBIND`** Avoid any effect on CPU binding. On some operating systems, some underlying memory binding functions also bind the application to the corresponding CPU(s). Using this flag will cause `hwloc` to avoid using OS functions that could potentially affect CPU bindings. Note,

however, that using NOCPUBIND may reduce hwloc's overall memory binding support. Specifically: some of hwloc's memory binding functions may fail with errno set to ENOSYS when used with NOCPUBIND.

**HWLOC\_MEMBIND\_BYNODESET** Consider the bitmap argument as a nodeset. Functions whose name ends with `_nodeset()` take a nodeset argument. Other functions take a bitmap argument that is considered a nodeset if this flag is given, or a cpuset otherwise.

Memory binding by CPU set cannot work for CPU-less NUMA memory nodes. Binding by nodeset should therefore be preferred whenever possible.

### 19.11.2.2 enum hwloc\_membind\_policy\_t

Memory binding policy.

These constants can be used to choose the binding policy. Only one policy can be used at a time (i.e., the values cannot be OR'ed together).

Not all systems support all kinds of binding. `hwloc_topology_get_support()` may be used to query about the actual memory binding policy support in the currently used operating system. See the "Detailed Description" section of [Memory binding](#) for a description of errors that can occur.

Enumerator:

**HWLOC\_MEMBIND\_DEFAULT** Reset the memory allocation policy to the system default. Depending on the operating system, this may correspond to [HWLOC\\_MEMBIND\\_FIRSTTOUCH](#) (Linux), or [HWLOC\\_MEMBIND\\_BIND](#) (-AIX, HP-UX, OSF, Solaris, Windows). This policy is never returned by get membind functions when running on normal machines. It is only returned when binding hooks are empty because the topology was loaded from XML, or `HWLOC_THISSYSTEM=0`, etc.

**HWLOC\_MEMBIND\_FIRSTTOUCH** Allocate memory but do not immediately bind it to a specific locality. Instead, each page in the allocation is bound only when it is first touched. Pages are individually bound to the local NUMA node of the first thread that touches it. If there is not enough memory on the node, allocation may be done in the specified nodes before allocating on other nodes.

**HWLOC\_MEMBIND\_BIND** Allocate memory on the specified nodes.

**HWLOC\_MEMBIND\_INTERLEAVE** Allocate memory on the given nodes in an interleaved / round-robin manner. The precise layout of the memory across multiple NUMA nodes is OS/system specific. Interleaving can be useful when threads distributed across the specified NUMA nodes will all be accessing the whole memory range concurrently, since the interleave will then balance the memory references.

**HWLOC\_MEMBIND\_REPLICATE** Replicate memory on the given nodes; reads from this memory will attempt to be serviced from the NUMA node local to the reading thread. Replicating can be useful when multiple threads from the specified NUMA nodes will be sharing the same read-only data. This

policy can only be used with existing memory allocations (i.e., the `hwloc_set_*mbind*()` functions); it cannot be used with functions that allocate new memory (i.e., the `hwloc_alloc*()` functions).

**HWLOC\_MEBIND\_NEXTTOUCH** For each page bound with this policy, by next time it is touched (and next time only), it is moved from its current location to the local NUMA node of the thread where the memory reference occurred (if it needs to be moved at all).

**HWLOC\_MEBIND\_MIXED** Returned by `get_mbind()` functions when multiple threads or parts of a memory area have differing memory binding policies.

### 19.11.3 Function Documentation

#### 19.11.3.1 `void* hwloc_alloc ( hwloc_topology_t topology, size_t len )`

Allocate some memory.

This is equivalent to `malloc()`, except that it tries to allocate page-aligned memory from the OS.

#### Note

The allocated memory should be freed with `hwloc_free()`.

#### 19.11.3.2 `void* hwloc_alloc_mbind ( hwloc_topology_t topology, size_t len, hwloc_const_bitmap_t set, hwloc_mbind_policy_t policy, int flags )`

Allocate some memory on NUMA memory nodes specified by `set`.

#### Returns

NULL with `errno` set to `ENOSYS` if the action is not supported and `HWLOC_MEBIND_STRICT` is given

NULL with `errno` set to `EXDEV` if the binding cannot be enforced and `HWLOC_MEBIND_STRICT` is given

NULL with `errno` set to `ENOMEM` if the memory allocation failed even before trying to bind.

If `HWLOC_MEBIND_BYNODESET` is specified, `set` is considered a `nodeset`. - Otherwise it's a `cpuset`.

#### Note

The allocated memory should be freed with `hwloc_free()`.

#### 19.11.3.3 `void* hwloc_alloc_mbind_nodeset ( hwloc_topology_t topology, size_t len, hwloc_const_nodeset_t nodeset, hwloc_mbind_policy_t policy, int flags )`

Allocate some memory on NUMA memory nodes specified by `nodeset`.

**Returns**

NULL with `errno` set to `ENOSYS` if the action is not supported and `HWLOC_MEMBIND_STRICT` is given  
 NULL with `errno` set to `EXDEV` if the binding cannot be enforced and `HWLOC_MEMBIND_STRICT` is given  
 NULL with `errno` set to `ENOMEM` if the memory allocation failed even before trying to bind.

**Note**

The allocated memory should be freed with `hwloc_free()`.

```
19.11.3.4 static void* hwloc_alloc_membind_policy ( hwloc_topology_t topology,
size_t len, hwloc_const_bitmap_t set, hwloc_membind_policy_t policy, int
flags ) [inline, static]
```

Allocate some memory on NUMA memory nodes specified by `set`.

This is similar to `hwloc_alloc_membind_nodeseq()` except that it is allowed to change the current memory binding policy, thus providing more binding support, at the expense of changing the current state.

If `HWLOC_MEMBIND_BYNODESET` is specified, `set` is considered a `nodeseq`. - Otherwise it's a `cpuset`.

```
19.11.3.5 static void* hwloc_alloc_membind_policy_nodeseq (
hwloc_topology_t topology, size_t len, hwloc_const_nodeseq_t nodeseq,
hwloc_membind_policy_t policy, int flags ) [inline, static]
```

Allocate some memory on NUMA memory nodes specified by `nodeseq`.

This is similar to `hwloc_alloc_membind()` except that it is allowed to change the current memory binding policy, thus providing more binding support, at the expense of changing the current state.

```
19.11.3.6 int hwloc_free ( hwloc_topology_t topology, void * addr, size_t len )
```

Free memory that was previously allocated by `hwloc_alloc()` or `hwloc_alloc_membind()`.

```
19.11.3.7 int hwloc_get_area_membind ( hwloc_topology_t topology, const void *
addr, size_t len, hwloc_bitmap_t set, hwloc_membind_policy_t * policy, int
flags )
```

Query the CPUs near the physical NUMA node(s) and binding policy of the memory identified by `(addr, len)`.

This function has two output parameters: `set` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the memory binding policies and nodesets of the pages in the address range.

If `HWLOC_MEBIND_STRICT` is specified, the target pages are first checked to see if they all have the same memory binding policy and nodeset. If they do not, -1 is returned and `errno` is set to `EXDEV`. If they are identical across all pages, the `set` and `policy` are returned in `set` and `policy`, respectively.

If `HWLOC_MEBIND_STRICT` is not specified, the union of all NUMA node(s) containing pages in the address range is calculated. If all pages in the target have the same policy, it is returned in `policy`. Otherwise, `policy` is set to `HWLOC_MEBIND_MIXED`.

If `HWLOC_MEBIND_BYNODESET` is specified, `set` is considered a nodeset. - Otherwise it's a cpuset.

If `len` is 0, -1 is returned and `errno` is set to `EINVAL`.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

```
19.11.3.8 int hwloc_get_area_membind_nodeset( hwloc_topology_t topology, const
void * addr, size_t len, hwloc_nodeset_t nodeset, hwloc_membind_policy_t
* policy, int flags )
```

Query the physical NUMA node(s) and binding policy of the memory identified by (`addr`, `len`).

This function has two output parameters: `nodeset` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the memory binding policies and nodesets of the pages in the address range.

If `HWLOC_MEBIND_STRICT` is specified, the target pages are first checked to see if they all have the same memory binding policy and nodeset. If they do not, -1 is returned and `errno` is set to `EXDEV`. If they are identical across all pages, the `nodeset` and `policy` are returned in `nodeset` and `policy`, respectively.

If `HWLOC_MEBIND_STRICT` is not specified, `nodeset` is set to the union of all NUMA node(s) containing pages in the address range. If all pages in the target have the same policy, it is returned in `policy`. Otherwise, `policy` is set to `HWLOC_MEBIND_MIXED`.

If `len` is 0, -1 is returned and `errno` is set to `EINVAL`.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

```
19.11.3.9 int hwloc_get_area_memlocation( hwloc_topology_t topology, const void
* addr, size_t len, hwloc_bitmap_t set, int flags )
```

Get the NUMA nodes where memory identified by (`addr`, `len`) is physically allocated.

Fills `set` according to the NUMA nodes where the memory area pages are physically allocated. If no page is actually allocated yet, `set` may be empty.

If pages spread to multiple nodes, it is not specified whether they spread equitably, or whether most of them are on a single node, etc.

The operating system may move memory pages from one processor to another at any time according to their binding, so this function may return something that is already outdated.

If `HWLOC_MEMBIND_BYNODESET` is specified in `flags`, `set` is considered a node-set. Otherwise it's a cpuset.

If `len` is 0, `set` is emptied.

19.11.3.10 `int hwloc_get_membind ( hwloc_topology_t topology, hwloc_bitmap_t set, hwloc_membind_policy_t * policy, int flags )`

Query the default memory binding policy and physical locality of the current process or thread.

This function has two output parameters: `set` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the current memory binding policies and nodesets in the queried target.

Passing the `HWLOC_MEMBIND_PROCESS` flag specifies that the query target is the current policies and nodesets for all the threads in the current process. Passing `HWLOC_MEMBIND_THREAD` specifies that the query target is the current policy and nodeset for only the thread invoking this function.

If neither of these flags are passed (which is the most portable method), the process is assumed to be single threaded. This allows `hwloc` to use either process-based OS functions or thread-based OS functions, depending on which are available.

`HWLOC_MEMBIND_STRICT` is only meaningful when `HWLOC_MEMBIND_PROCESS` is also specified. In this case, `hwloc` will check the default memory policies and nodesets for all threads in the process. If they are not identical, -1 is returned and `errno` is set to `EXDEV`. If they are identical, the values are returned in `set` and `policy`.

Otherwise, if `HWLOC_MEMBIND_PROCESS` is specified (and `HWLOC_MEMBIND_STRICT` is *not* specified), the default set from each thread is logically OR'ed together. If all threads' default policies are the same, `policy` is set to that policy. If they are different, `policy` is set to `HWLOC_MEMBIND_MIXED`.

In the `HWLOC_MEMBIND_THREAD` case (or when neither `HWLOC_MEMBIND_PROCESS` or `HWLOC_MEMBIND_THREAD` is specified), there is only one set and policy; they are returned in `set` and `policy`, respectively.

If `HWLOC_MEMBIND_BYNODESET` is specified, `set` is considered a nodeset. - Otherwise it's a cpuset.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

```
19.11.3.11 int hwloc_get_membind_nodeset ( hwloc_topology_t topology,  
      hwloc_nodeset_t nodeset, hwloc_membind_policy_t* policy, int flags )
```

Query the default memory binding policy and physical locality of the current process or thread.

This function has two output parameters: `nodeset` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the current memory binding policies and nodesets in the queried target.

Passing the `HWLOC_MEMBIND_PROCESS` flag specifies that the query target is the current policies and nodesets for all the threads in the current process. Passing `HWLOC_MEMBIND_THREAD` specifies that the query target is the current policy and nodeset for only the thread invoking this function.

If neither of these flags are passed (which is the most portable method), the process is assumed to be single threaded. This allows `hwloc` to use either process-based OS functions or thread-based OS functions, depending on which are available.

`HWLOC_MEMBIND_STRICT` is only meaningful when `HWLOC_MEMBIND_PROCESS` is also specified. In this case, `hwloc` will check the default memory policies and nodesets for all threads in the process. If they are not identical, -1 is returned and `errno` is set to `EXDEV`. If they are identical, the values are returned in `nodeset` and `policy`.

Otherwise, if `HWLOC_MEMBIND_PROCESS` is specified (and `HWLOC_MEMBIND_STRICT` is *not* specified), `nodeset` is set to the logical OR of all threads' default nodeset. If all threads' default policies are the same, `policy` is set to that policy. If they are different, `policy` is set to `HWLOC_MEMBIND_MIXED`.

In the `HWLOC_MEMBIND_THREAD` case (or when neither `HWLOC_MEMBIND_PROCESS` or `HWLOC_MEMBIND_THREAD` is specified), there is only one nodeset and policy; they are returned in `nodeset` and `policy`, respectively.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

```
19.11.3.12 int hwloc_get_proc_membind ( hwloc_topology_t topology, hwloc_pid_t  
      pid, hwloc_bitmap_t set, hwloc_membind_policy_t* policy, int flags )
```

Query the default memory binding policy and physical locality of the specified process.

This function has two output parameters: `set` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the current memory binding policies and nodesets in the queried target.

Passing the `HWLOC_MEMBIND_PROCESS` flag specifies that the query target is the current policies and nodesets for all the threads in the specified process. If `HWLOC_MEMBIND_PROCESS` is not specified (which is the most portable method), the process is assumed to be single threaded. This allows `hwloc` to use either process-based OS functions or thread-based OS functions, depending on which are available.

Note that it does not make sense to pass `HWLOC_MEMBIND_THREAD` to this function.

If `HWLOC_MEMBIND_STRICT` is specified, `hwloc` will check the default memory poli-

cies and nodesets for all threads in the specified process. If they are not identical, -1 is returned and `errno` is set to `EXDEV`. If they are identical, the values are returned in `set` and `policy`.

Otherwise, `set` is set to the logical OR of all threads' default set. If all threads' default policies are the same, `policy` is set to that policy. If they are different, `policy` is set to `HWLOC_MEBIND_MIXED`.

If `HWLOC_MEBIND_BYNODESET` is specified, `set` is considered a nodeset. - Otherwise it's a cpuset.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

#### Note

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

19.11.3.13 `int hwloc_get_proc_mebind_nodeset ( hwloc_topology_t topology, hwloc_pid_t pid, hwloc_nodeset_t nodeset, hwloc_mebind_policy_t * policy, int flags )`

Query the default memory binding policy and physical locality of the specified process.

This function has two output parameters: `nodeset` and `policy`. The values returned in these parameters depend on both the `flags` passed in and the current memory binding policies and nodesets in the queried target.

Passing the `HWLOC_MEBIND_PROCESS` flag specifies that the query target is the current policies and nodesets for all the threads in the specified process. If `HWLOC_MEBIND_PROCESS` is not specified (which is the most portable method), the process is assumed to be single threaded. This allows `hwloc` to use either process-based OS functions or thread-based OS functions, depending on which are available.

Note that it does not make sense to pass `HWLOC_MEBIND_THREAD` to this function.

If `HWLOC_MEBIND_STRICT` is specified, `hwloc` will check the default memory policies and nodesets for all threads in the specified process. If they are not identical, -1 is returned and `errno` is set to `EXDEV`. If they are identical, the values are returned in `nodeset` and `policy`.

Otherwise, `nodeset` is set to the logical OR of all threads' default nodeset. If all threads' default policies are the same, `policy` is set to that policy. If they are different, `policy` is set to `HWLOC_MEBIND_MIXED`.

If any other flags are specified, -1 is returned and `errno` is set to `EINVAL`.

#### Note

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.



19.11.3.14 `int hwloc_set_area_mbind ( hwloc_topology_t topology, const void *  
addr, size_t len, hwloc_const_bitmap_t set, hwloc_mbind_policy_t  
policy, int flags )`

Bind the already-allocated memory identified by (addr, len) to the NUMA node(s) specified by `set`.

If `HWLOC_MEBIND_BYNODESET` is specified, `set` is considered a nodeset. -  
Otherwise it's a cpuset.

#### Returns

- 0 if `len` is 0.
- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

19.11.3.15 `int hwloc_set_area_mbind_nodeset ( hwloc_topology_t  
topology, const void * addr, size_t len, hwloc_const_nodeset_t nodeset,  
hwloc_mbind_policy_t policy, int flags )`

Bind the already-allocated memory identified by (addr, len) to the NUMA node(s) specified by `nodeset`.

#### Returns

- 0 if `len` is 0.
- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

19.11.3.16 `int hwloc_set_mbind ( hwloc_topology_t topology,  
hwloc_const_bitmap_t set, hwloc_mbind_policy_t policy, int flags )`

Set the default memory binding policy of the current process or thread to prefer the NUMA node(s) specified by `set`.

If neither `HWLOC_MEBIND_PROCESS` nor `HWLOC_MEBIND_THREAD` is specified, the current process is assumed to be single-threaded. This is the most portable form as it permits `hwloc` to use either process-based OS functions or thread-based OS functions, depending on which are available.

If `HWLOC_MEBIND_BYNODESET` is specified, `set` is considered a nodeset. -  
Otherwise it's a cpuset.

#### Returns

- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

```
19.11.3.17 int hwloc_set_membind_nodese
t ( hwloc_topology_t topology,
    hwloc_const_nodese
t_t nodese
t, hwloc_membind_policy_t policy, int
    flags )
```

Set the default memory binding policy of the current process or thread to prefer the NUMA node(s) specified by `nodese
t`.

If neither `HWLOC_MEMBIND_PROCESS` nor `HWLOC_MEMBIND_THREAD` is specified, the current process is assumed to be single-threaded. This is the most portable form as it permits hwloc to use either process-based OS functions or thread-based OS functions, depending on which are available.

#### Returns

- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

```
19.11.3.18 int hwloc_set_proc_membind
( hwloc_topology_t topology, hwloc_pid_t
    pid, hwloc_const_bitmap_t se
t, hwloc_membind_policy_t policy, int flags
    )
```

Set the default memory binding policy of the specified process to prefer the NUMA node(s) specified by `se
t`.

If `HWLOC_MEMBIND_BYNODESET` is specified, `se
t` is considered a `nodese
t`. - Otherwise it's a `cpuset`.

#### Returns

- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

#### Note

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

```
19.11.3.19 int hwloc_set_proc_membind_nodese
t ( hwloc_topology_t
    topology, hwloc_pid_t pid, hwloc_const_nodese
t_t nodese
t, hwloc_membind_policy_t policy, int flags )
```

Set the default memory binding policy of the specified process to prefer the NUMA node(s) specified by `nodese
t`.

#### Returns

- 1 with `errno` set to `ENOSYS` if the action is not supported
- 1 with `errno` set to `EXDEV` if the binding cannot be enforced

**Note**

`hwloc_pid_t` is `pid_t` on Unix platforms, and `HANDLE` on native Windows platforms.

## 19.12 Modifying a loaded Topology

### Enumerations

- enum `hwloc_restrict_flags_e` { `HWLOC_RESTRICT_FLAG_ADAPT_DISTANCES`, `HWLOC_RESTRICT_FLAG_ADAPT_MISC`, `HWLOC_RESTRICT_FLAG_ADAPT_IO` }

### Functions

- `hwloc_obj_t hwloc_topology_insert_misc_object_by_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t cpuset`, `const char *name`)
- `hwloc_obj_t hwloc_topology_insert_misc_object_by_parent` (`hwloc_topology_t topology`, `hwloc_obj_t parent`, `const char *name`)
- int `hwloc_topology_restrict` (`hwloc_topology_t restrict topology`, `hwloc_const_cpuset_t cpuset`, unsigned long flags)

### 19.12.1 Enumeration Type Documentation

#### 19.12.1.1 enum `hwloc_restrict_flags_e`

Flags to be given to `hwloc_topology_restrict()`.

Enumerator:

**`HWLOC_RESTRICT_FLAG_ADAPT_DISTANCES`** Adapt distance matrices according to objects being removed during restriction. If this flag is not set, distance matrices are removed.

**`HWLOC_RESTRICT_FLAG_ADAPT_MISC`** Move Misc objects to ancestors if their parents are removed during restriction. If this flag is not set, Misc objects are removed when their parents are removed.

**`HWLOC_RESTRICT_FLAG_ADAPT_IO`** Move I/O objects to ancestors if their parents are removed during restriction. If this flag is not set, I/O devices and bridges are removed when their parents are removed.

### 19.12.2 Function Documentation

#### 19.12.2.1 `hwloc_obj_t hwloc_topology_insert_misc_object_by_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t cpuset`, `const char *name`)

Add a MISC object to the topology.

A new MISC object will be created and inserted into the topology at the position given by bitmap `cpuset`. This offers a way to add new intermediate levels to the topology hierarchy.

`cpuset` and `name` will be copied to setup the new object attributes.

**Returns**

the newly-created object.  
NULL if the insertion conflicts with the existing topology tree.

**Note**

If `name` contains some non-printable characters, they will be dropped when exporting to XML, see [hwloc\\_topology\\_export\\_xml\(\)](#).

**19.12.2.2 hwloc\_obj\_t hwloc\_topology\_insert\_misc\_object\_by\_parent ( hwloc\_topology\_t topology, hwloc\_obj\_t parent, const char \* name )**

Add a MISC object as a leaf of the topology.

A new MISC object will be created and inserted into the topology at the position given by `parent`. It is appended to the list of existing children, without ever adding any intermediate hierarchy level. This is useful for annotating the topology without actually changing the hierarchy.

`name` will be copied to the setup the new object attributes. However, the new leaf object will not have any `cpuset`.

**Returns**

the newly-created object

**Note**

If `name` contains some non-printable characters, they will be dropped when exporting to XML, see [hwloc\\_topology\\_export\\_xml\(\)](#).

**19.12.2.3 int hwloc\_topology\_restrict ( hwloc\_topology\_t restrict\_topology, hwloc\_const\_cpuset\_t cpuset, unsigned long flags )**

Restrict the topology to the given CPU set.

Topology `restrict_topology` is modified so as to remove all objects that are not included (or partially included) in the CPU set `cpuset`. All objects CPU and node sets are restricted accordingly.

`flags` is a OR'ed set of [hwloc\\_restrict\\_flags\\_e](#).

**Note**

This call may not be reverted by restricting back to a larger `cpuset`. Once dropped during restriction, objects may not be brought back, except by loading another topology with [hwloc\\_topology\\_load\(\)](#).

**Returns**

0 on success.

-1 with `errno` set to `EINVAL` if the input `cpuset` is invalid. The topology is not modified in this case.

-1 with `errno` set to `ENOMEM` on failure to allocate internal data. The topology is reinitialized in this case. It should be either destroyed with [hwloc\\_topology\\_destroy\(\)](#) or configured and loaded again.

## 19.13 Building Custom Topologies

### Functions

- `int hwloc_custom_insert_topology (hwloc_topology_t newtopology, hwloc_obj_t newparent, hwloc_topology_t oldtopology, hwloc_obj_t oldroot)`
- `hwloc_obj_t hwloc_custom_insert_group_object_by_parent (hwloc_topology_t topology, hwloc_obj_t parent, int groupdepth)`

#### 19.13.1 Detailed Description

A custom topology may be initialized by calling `hwloc_topology_set_custom()` after `hwloc_topology_init()`. It may then be modified by inserting objects or entire topologies. Once done assembling, `hwloc_topology_load()` should be invoked as usual to finalize the topology.

#### 19.13.2 Function Documentation

**19.13.2.1** `hwloc_obj_t hwloc_custom_insert_group_object_by_parent (hwloc_topology_t topology, hwloc_obj_t parent, int groupdepth )`

Insert a new group object inside a custom topology.

An object with type `HWLOC_OBJ_GROUP` is inserted as a new child of object `parent`.

`groupdepth` is the depth attribute to be given to the new object. It may for instance be 0 for top-level groups, 1 for their children, and so on.

The custom topology `newtopology` must have been prepared with `hwloc_topology_set_custom()` and not loaded with `hwloc_topology_load()` yet.

`parent` may be either the root of `topology` or an object that was added earlier through `hwloc_custom_insert_group_object_by_parent()`.

#### Note

The `cpuset` and `nodeset` of the new group object are `NULL` because these sets are meaningless when assembling multiple topologies.

The `cpuset` and `nodeset` of the `parent` object are not modified.

**19.13.2.2** `int hwloc_custom_insert_topology ( hwloc_topology_t newtopology, hwloc_obj_t newparent, hwloc_topology_t oldtopology, hwloc_obj_t oldroot )`

Insert an existing topology inside a custom topology.

Duplicate the existing topology `oldtopology` inside a new custom topology `newtopology` as a leaf of object `newparent`.

If `oldroot` is not `NULL`, duplicate `oldroot` and all its children instead of the entire `oldtopology`. Passing the root object of `oldtopology` in `oldroot` is equivalent to passing `NULL`.

The custom topology `newtopology` must have been prepared with [hwloc\\_topology\\_set\\_custom\(\)](#) and not loaded with [hwloc\\_topology\\_load\(\)](#) yet.

`newparent` may be either the root of `newtopology` or an object that was added through [hwloc\\_custom\\_insert\\_group\\_object\\_by\\_parent\(\)](#).

**Note**

The `cpuset` and `nodeset` of the `newparent` object are not modified based on the contents of `oldtopology`.



## 19.14 Exporting Topologies to XML

### Functions

- int `hwloc_topology_export_xml` (`hwloc_topology_t` topology, const char \*xmlpath)
- int `hwloc_topology_export_xmlbuffer` (`hwloc_topology_t` topology, char \*\*xmlbuffer, int \*buflen)
- void `hwloc_free_xmlbuffer` (`hwloc_topology_t` topology, char \*xmlbuffer)
- void `hwloc_topology_set_userdata_export_callback` (`hwloc_topology_t` topology, void(\*export\_cb)(void \*reserved, `hwloc_topology_t` topology, `hwloc_obj_t` obj))
- int `hwloc_export_obj_userdata` (void \*reserved, `hwloc_topology_t` topology, `hwloc_obj_t` obj, const char \*name, const void \*buffer, size\_t length)
- int `hwloc_export_obj_userdata_base64` (void \*reserved, `hwloc_topology_t` topology, `hwloc_obj_t` obj, const char \*name, const void \*buffer, size\_t length)
- void `hwloc_topology_set_userdata_import_callback` (`hwloc_topology_t` topology, void(\*import\_cb)(`hwloc_topology_t` topology, `hwloc_obj_t` obj, const char \*name, const void \*buffer, size\_t length))

### 19.14.1 Function Documentation

19.14.1.1 int `hwloc_export_obj_userdata` ( void \* *reserved*, `hwloc_topology_t` *topology*, `hwloc_obj_t` *obj*, const char \* *name*, const void \* *buffer*, size\_t *length* )

Export some object userdata to XML.

This function may only be called from within the `export()` callback passed to `hwloc_topology_set_userdata_export_callback()`. It may be invoked one of multiple times to export some userdata to XML. The `buffer` content of length `length` is stored with optional name `name`.

When importing this XML file, the `import()` callback (if set) will be called exactly as many times as `hwloc_export_obj_userdata()` was called during `export()`. It will receive the corresponding `name`, `buffer` and `length` arguments.

`reserved`, `topology` and `obj` must be the first three parameters that were given to the export callback.

Only printable characters may be exported to XML string attributes. If a non-printable character is passed in `name` or `buffer`, the function returns -1 with `errno` set to `EINVAL`.

If exporting binary data, the application should first encode into printable characters only (or use `hwloc_export_obj_userdata_base64()`). It should also take care of portability issues if the export may be reimported on a different architecture.

19.14.1.2 int `hwloc_export_obj_userdata_base64` ( void \* *reserved*, `hwloc_topology_t` *topology*, `hwloc_obj_t` *obj*, const char \* *name*, const void \* *buffer*, size\_t *length* )

Encode and export some object userdata to XML.

This function is similar to [hwloc\\_export\\_obj\\_userdata\(\)](#) but it encodes the input buffer into printable characters before exporting. On import, decoding is automatically performed before the data is given to the `import()` callback if any.

This function may only be called from within the `export()` callback passed to [hwloc\\_topology\\_set\\_userdata\\_export\\_callback\(\)](#).

The function does not take care of portability issues if the export may be reimported on a different architecture.

19.14.1.3 `void hwloc_free_xmlbuffer ( hwloc_topology_t topology, char * xmlbuffer )`

Free a buffer allocated by [hwloc\\_topology\\_export\\_xmlbuffer\(\)](#)

19.14.1.4 `int hwloc_topology_export_xml ( hwloc_topology_t topology, const char * xmlpath )`

Export the topology into an XML file.

This file may be loaded later through [hwloc\\_topology\\_set\\_xml\(\)](#).

#### Returns

-1 if a failure occurred.

#### Note

See also [hwloc\\_topology\\_set\\_userdata\\_export\\_callback\(\)](#) for exporting application-specific object userdata.

The topology-specific userdata pointer is ignored when exporting to XML.

Only printable characters may be exported to XML string attributes. Any other character, especially any non-ASCII character, will be silently dropped.

If `name` is "-", the XML output is sent to the standard output.

19.14.1.5 `int hwloc_topology_export_xmlbuffer ( hwloc_topology_t topology, char ** xmlbuffer, int * buflen )`

Export the topology into a newly-allocated XML memory buffer.

`xmlbuffer` is allocated by the callee and should be freed with [hwloc\\_free\\_xmlbuffer\(\)](#) later in the caller.

This memory buffer may be loaded later through [hwloc\\_topology\\_set\\_xmlbuffer\(\)](#).

The returned buffer ends with a `\0` that is included in the returned length.

#### Returns

-1 if a failure occurred.

**Note**

See also [hwloc\\_topology\\_set\\_userdata\\_export\\_callback\(\)](#) for exporting application-specific object userdata.

The topology-specific userdata pointer is ignored when exporting to XML.

Only printable characters may be exported to XML string attributes. Any other character, especially any non-ASCII character, will be silently dropped.

**19.14.1.6** `void hwloc_topology_set_userdata_export_callback ( hwloc_topology_t topology, void (*)(void *reserved, hwloc_topology_t topology, hwloc_obj_t obj) export_cb )`

Set the application-specific callback for exporting object userdata.

The object userdata pointer is not exported to XML by default because hwloc does not know what it contains.

This function lets applications set `export_cb` to a callback function that converts this opaque userdata into an exportable string.

`export_cb` is invoked during XML export for each object whose `userdata` pointer is not `NULL`. The callback should use [hwloc\\_export\\_obj\\_userdata\(\)](#) or [hwloc\\_export\\_obj\\_userdata\\_base64\(\)](#) to actually export something to XML (possibly multiple times per object).

`export_cb` may be set to `NULL` if userdata should not be exported to XML.

**Note**

The topology-specific userdata pointer is ignored when exporting to XML.

**19.14.1.7** `void hwloc_topology_set_userdata_import_callback ( hwloc_topology_t topology, void (*)(hwloc_topology_t topology, hwloc_obj_t obj, const char *name, const void *buffer, size_t length) import_cb )`

Set the application-specific callback for importing userdata.

On XML import, userdata is ignored by default because hwloc does not know how to store it in memory.

This function lets applications set `import_cb` to a callback function that will get the XML-stored userdata and store it in the object as expected by the application.

`import_cb` is called during [hwloc\\_topology\\_load\(\)](#) as many times as [hwloc\\_export\\_obj\\_userdata\(\)](#) was called during export. The topology is not entirely setup yet. Object attributes are ready to consult, but links between objects are not.

`import_cb` may be `NULL` if userdata should be ignored during import.

**Note**

`buffer` contains `length` characters followed by a null byte (`'\0'`).  
This function should be called before [hwloc\\_topology\\_load\(\)](#).  
The topology-specific userdata pointer is ignored when importing from XML.

## 19.15 Exporting Topologies to Synthetic

### Enumerations

- enum `hwloc_topology_export_synthetic_flags_e` { `HWLOC_TOPOLOGY_EXPORT_SYNTHETIC_FLAG_NO_EXTENDED_TYPES`, `HWLOC_TOPOLOGY_EXPORT_SYNTHETIC_FLAG_NO_ATTRS` }

### Functions

- int `hwloc_topology_export_synthetic` (`hwloc_topology_t` topology, char \*buffer, size\_t buflen, unsigned long flags)

#### 19.15.1 Enumeration Type Documentation

##### 19.15.1.1 enum `hwloc_topology_export_synthetic_flags_e`

Flags for exporting synthetic topologies.

Flags to be given as a OR'ed set to `hwloc_topology_export_synthetic()`.

Enumerator:

***HWLOC\_TOPOLOGY\_EXPORT\_SYNTHETIC\_FLAG\_NO\_EXTENDED\_TYPES***

Export extended types such as L2dcache as basic types such as Cache. This is required if loading the synthetic description with `hwloc < 1.9`.

***HWLOC\_TOPOLOGY\_EXPORT\_SYNTHETIC\_FLAG\_NO\_ATTRS***

Do not export level attributes. Ignore level attributes such as memory/cache sizes or PU indexes. This is required if loading the synthetic description with `hwloc < 1.10`.

#### 19.15.2 Function Documentation

##### 19.15.2.1 int `hwloc_topology_export_synthetic` ( `hwloc_topology_t` topology, char \*buffer, size\_t buflen, unsigned long flags )

Export the topology as a synthetic string.

At most `buflen` characters will be written in `buffer`, including the terminating `\0`.

This exported string may be given back to `hwloc_topology_set_synthetic()`.

`flags` is a OR'ed set of `hwloc_topology_export_synthetic_flags_e`.

Returns

The number of characters that were written, not including the terminating `\0`.

-1 if the topology could not be exported, for instance if it is not symmetric.

**Note**

A 1024-byte buffer should be large enough for exporting topologies in the vast majority of cases.

## 19.16 Finding Objects inside a CPU set

### Functions

- static `hwloc_obj_t hwloc_get_first_largest_obj_inside_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`)
- int `hwloc_get_largest_objs_inside_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_t *restrict objs`, int `max`)
- static `hwloc_obj_t hwloc_get_next_obj_inside_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, unsigned `depth`, `hwloc_obj_t prev`)
- static `hwloc_obj_t hwloc_get_next_obj_inside_cpuset_by_type` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`, `hwloc_obj_t prev`)
- static `hwloc_obj_t hwloc_get_obj_inside_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, unsigned `depth`, unsigned `idx`)
- static `hwloc_obj_t hwloc_get_obj_inside_cpuset_by_type` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`, unsigned `idx`)
- static unsigned `hwloc_get_nbojs_inside_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, unsigned `depth`)
- static int `hwloc_get_nbojs_inside_cpuset_by_type` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`)
- static int `hwloc_get_obj_index_inside_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_t obj`)

### 19.16.1 Function Documentation

19.16.1.1 `static hwloc_obj_t hwloc_get_first_largest_obj_inside_cpuset ( hwloc_topology_t topology, hwloc_const_cpuset_t set ) [inline, static]`

Get the first largest object included in the given cpuset `set`.

#### Returns

the first object that is included in `set` and whose parent is not.

This is convenient for iterating over all largest objects within a CPU set by doing a loop getting the first largest object and clearing its CPU set from the remaining CPU set.

#### Note

This function cannot work if the root object does not have a CPU set, e.g. if the topology is made of different machines.

19.16.1.2 `int hwloc_get_largest_objs_inside_cpuset ( hwloc_topology_t topology, hwloc_const_cpuset_t set, hwloc_obj_t *restrict objs, int max )`

Get the set of largest objects covering exactly a given cpuset `set`.

**Returns**

the number of objects returned in `objs`.

**Note**

This function cannot work if the root object does not have a CPU set, e.g. if the topology is made of different machines.

```
19.16.1.3 static unsigned hwloc_get_nbobjs_inside_cpuset_by_depth (
    hwloc_topology_t topology, hwloc_const_cpuset_t set, unsigned depth )
    [inline, static]
```

Return the number of objects at depth `depth` included in CPU set `set`.

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects at the given depth do not have CPU sets or if the topology is made of different machines.

```
19.16.1.4 static int hwloc_get_nbobjs_inside_cpuset_by_type ( hwloc_topology_t
    topology, hwloc_const_cpuset_t set, hwloc_obj_type_t type )
    [inline, static]
```

Return the number of objects of type `type` included in CPU set `set`.

If no object for that type exists inside CPU set `set`, 0 is returned. If there are several levels with objects of that type inside CPU set `set`, -1 is returned.

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects of the given type do not have CPU sets or if the topology is made of different machines.

```
19.16.1.5 static hwloc_obj_t hwloc_get_next_obj_inside_cpuset_by_depth (
    hwloc_topology_t topology, hwloc_const_cpuset_t set, unsigned depth,
    hwloc_obj_t prev ) [inline, static]
```

Return the next object at depth `depth` included in CPU set `set`.

If `prev` is `NULL`, return the first object at depth `depth` included in `set`. The next invocation should pass the previous return value in `prev` so as to obtain the next object in `set`.



**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects at the given depth do not have CPU sets or if the topology is made of different machines.

```
19.16.1.6 static hwloc_obj_t hwloc_get_next_obj_inside_cpuset_by_type (
    hwloc_topology_t topology, hwloc_const_cpuset_t set, hwloc_obj_type_t
    type, hwloc_obj_t prev ) [inline, static]
```

Return the next object of type `type` included in CPU set `set`.

If there are multiple or no depth for given type, return `NULL` and let the caller fallback to [hwloc\\_get\\_next\\_obj\\_inside\\_cpuset\\_by\\_depth\(\)](#).

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects of the given type do not have CPU sets or if the topology is made of different machines.

```
19.16.1.7 static int hwloc_get_obj_index_inside_cpuset ( hwloc_topology_t
    topology, hwloc_const_cpuset_t set, hwloc_obj_t obj ) [inline,
    static]
```

Return the logical index among the objects included in CPU set `set`.

Consult all objects in the same level as `obj` and inside CPU set `set` in the logical order, and return the index of `obj` within them. If `set` covers the entire topology, this is the logical index of `obj`. Otherwise, this is similar to a logical index within the part of the topology defined by CPU set `set`.

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

```
19.16.1.8 static hwloc_obj_t hwloc_get_obj_inside_cpuset_by_depth (
    hwloc_topology_t topology, hwloc_const_cpuset_t set, unsigned depth,
    unsigned idx ) [inline, static]
```

Return the (logically) `idx`-th object at depth `depth` included in CPU set `set`.

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects at the given depth do not have CPU sets or if the topology is made of different machines.

```
19.16.1.9 static hwloc_obj_t hwloc_get_obj_inside_cpuset_by_type (
    hwloc_topology_t topology, hwloc_const_cpuset_t set, hwloc_obj_type_t
    type, unsigned idx ) [inline, static]
```

Return the `idx`-th object of type `type` included in CPU set `set`.

If there are multiple or no depth for given type, return `NULL` and let the caller fallback to [hwloc\\_get\\_obj\\_inside\\_cpuset\\_by\\_depth\(\)](#).

**Note**

Objects with empty CPU sets are ignored (otherwise they would be considered included in any given set).

This function cannot work if objects of the given type do not have CPU sets or if the topology is made of different machines.

## 19.17 Finding Objects covering at least CPU set

### Functions

- static `hwloc_obj_t hwloc_get_child_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_t parent`)
- static `hwloc_obj_t hwloc_get_obj_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`)
- static `hwloc_obj_t hwloc_get_next_obj_covering_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, unsigned `depth`, `hwloc_obj_t prev`)
- static `hwloc_obj_t hwloc_get_next_obj_covering_cpuset_by_type` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`, `hwloc_obj_t prev`)

### 19.17.1 Function Documentation

19.17.1.1 `static hwloc_obj_t hwloc_get_child_covering_cpuset ( hwloc_topology_t topology, hwloc_const_cpuset_t set, hwloc_obj_t parent ) [inline, static]`

Get the child covering at least CPU set `set`.

#### Returns

NULL if no child matches or if `set` is empty.

#### Note

This function cannot work if `parent` does not have a CPU set.

19.17.1.2 `static hwloc_obj_t hwloc_get_next_obj_covering_cpuset_by_depth ( hwloc_topology_t topology, hwloc_const_cpuset_t set, unsigned depth, hwloc_obj_t prev ) [inline, static]`

Iterate through same-depth objects covering at least CPU set `set`.

If object `prev` is NULL, return the first object at depth `depth` covering at least part of CPU set `set`. The next invocation should pass the previous return value in `prev` so as to obtain the next object covering at least another part of `set`.

#### Note

This function cannot work if objects at the given depth do not have CPU sets or if the topology is made of different machines.

19.17.1.3 `static hwloc_obj_t hwloc_get_next_obj_covering_cpuset_by_type ( hwloc_topology_t topology, hwloc_const_cpuset_t set, hwloc_obj_type_t type, hwloc_obj_t prev ) [inline, static]`

Iterate through same-type objects covering at least CPU set `set`.

If object `prev` is `NULL`, return the first object of type `type` covering at least part of CPU set `set`. The next invocation should pass the previous return value in `prev` so as to obtain the next object of type `type` covering at least another part of `set`.

If there are no or multiple depths for type `type`, `NULL` is returned. The caller may fallback to [hwloc\\_get\\_next\\_obj\\_covering\\_cpuset\\_by\\_depth\(\)](#) for each depth.

#### Note

This function cannot work if objects of the given type do not have CPU sets or if the topology is made of different machines.

19.17.1.4 `static hwloc_obj_t hwloc_get_obj_covering_cpuset ( hwloc_topology_t topology, hwloc_const_cpuset_t set ) [inline, static]`

Get the lowest object covering at least CPU set `set`.

#### Returns

`NULL` if no object matches or if `set` is empty.

#### Note

This function cannot work if the root object does not have a CPU set, e.g. if the topology is made of different machines.

## 19.18 Looking at Ancestor and Child Objects

### Functions

- static `hwloc_obj_t hwloc_get_ancestor_obj_by_depth` (`hwloc_topology_t topology`, unsigned `depth`, `hwloc_obj_t obj`)
- static `hwloc_obj_t hwloc_get_ancestor_obj_by_type` (`hwloc_topology_t topology`, `hwloc_obj_type_t type`, `hwloc_obj_t obj`)
- static `hwloc_obj_t hwloc_get_common_ancestor_obj` (`hwloc_topology_t topology`, `hwloc_obj_t obj1`, `hwloc_obj_t obj2`)
- static int `hwloc_obj_is_in_subtree` (`hwloc_topology_t topology`, `hwloc_obj_t obj`, `hwloc_obj_t subtree_root`)
- static `hwloc_obj_t hwloc_get_next_child` (`hwloc_topology_t topology`, `hwloc_obj_t parent`, `hwloc_obj_t prev`)

#### 19.18.1 Detailed Description

Be sure to see the figure in [Terms and Definitions](#) that shows a complete topology tree, including depths, child/sibling/cousin relationships, and an example of an asymmetric topology where one package has fewer caches than its peers.

#### 19.18.2 Function Documentation

**19.18.2.1** `static hwloc_obj_t hwloc_get_ancestor_obj_by_depth ( hwloc_topology_t topology, unsigned depth, hwloc_obj_t obj ) [inline, static]`

Returns the ancestor object of `obj` at depth `depth`.

**19.18.2.2** `static hwloc_obj_t hwloc_get_ancestor_obj_by_type ( hwloc_topology_t topology, hwloc_obj_type_t type, hwloc_obj_t obj ) [inline, static]`

Returns the ancestor object of `obj` with type `type`.

**19.18.2.3** `static hwloc_obj_t hwloc_get_common_ancestor_obj ( hwloc_topology_t topology, hwloc_obj_t obj1, hwloc_obj_t obj2 ) [inline, static]`

Returns the common parent object to objects `obj1` and `obj2`.

**19.18.2.4** `static hwloc_obj_t hwloc_get_next_child ( hwloc_topology_t topology, hwloc_obj_t parent, hwloc_obj_t prev ) [inline, static]`

Return the next child.

If `prev` is `NULL`, return the first child.

19.18.2.5 `static int hwloc_obj_is_in_subtree ( hwloc_topology_t topology,  
hwloc_obj_t obj, hwloc_obj_t subtree_root ) [inline, static]`

Returns true if `obj` is inside the subtree beginning with ancestor object `subtree_root`.

**Note**

This function assumes that both `obj` and `subtree_root` have a `cpuset`.

## 19.19 Looking at Cache Objects

### Functions

- static int `hwloc_get_cache_type_depth` (`hwloc_topology_t` topology, unsigned cachelevel, `hwloc_obj_cache_type_t` cachetype)
- static `hwloc_obj_t` `hwloc_get_cache_covering_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set)
- static `hwloc_obj_t` `hwloc_get_shared_cache_covering_obj` (`hwloc_topology_t` topology, `hwloc_obj_t` obj)

### 19.19.1 Function Documentation

19.19.1.1 static `hwloc_obj_t` `hwloc_get_cache_covering_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set) [`inline`, `static`]

Get the first cache covering a cpuset `set`.

#### Returns

NULL if no cache matches.

#### Note

This function cannot work if the root object does not have a CPU set, e.g. if the topology is made of different machines.

19.19.1.2 static int `hwloc_get_cache_type_depth` (`hwloc_topology_t` topology, unsigned cachelevel, `hwloc_obj_cache_type_t` cachetype) [`inline`, `static`]

Find the depth of cache objects matching cache depth and type.

Return the depth of the topology level that contains cache objects whose attributes match `cachedepth` and `cachetype`. This function intends to disambiguate the case where `hwloc_get_type_depth()` returns `HWLOC_TYPE_DEPTH_MULTIPLE`.

If no cache level matches, `HWLOC_TYPE_DEPTH_UNKNOWN` is returned.

If `cachetype` is `HWLOC_OBJ_CACHE_UNIFIED`, the depth of the unique matching unified cache level is returned.

If `cachetype` is `HWLOC_OBJ_CACHE_DATA` or `HWLOC_OBJ_CACHE_INSTRUCTION`, either a matching cache, or a unified cache is returned.

If `cachetype` is `-1`, it is ignored and multiple levels may match. The function returns either the depth of a uniquely matching level or `HWLOC_TYPE_DEPTH_MULTIPLE`.

19.19.1.3 `static hwloc_obj_t hwloc_get_shared_cache_covering_obj (`  
`hwloc_topology_t topology, hwloc_obj_t obj ) [inline, static]`

Get the first cache shared between an object and somebody else.

**Returns**

NULL if no cache matches or if an invalid object is given.



## 19.20 Finding objects, miscellaneous helpers

### Functions

- static `hwloc_obj_t hwloc_get_pu_obj_by_os_index` (`hwloc_topology_t topology`, unsigned `os_index`)
- static `hwloc_obj_t hwloc_get_numanode_obj_by_os_index` (`hwloc_topology_t topology`, unsigned `os_index`)
- unsigned `hwloc_get_closest_objs` (`hwloc_topology_t topology`, `hwloc_obj_t src`, `hwloc_obj_t *restrict objs`, unsigned `max`)
- static `hwloc_obj_t hwloc_get_obj_below_by_type` (`hwloc_topology_t topology`, `hwloc_obj_type_t type1`, unsigned `idx1`, `hwloc_obj_type_t type2`, unsigned `idx2`)
- static `hwloc_obj_t hwloc_get_obj_below_array_by_type` (`hwloc_topology_t topology`, int `nr`, `hwloc_obj_type_t *typev`, unsigned `*idxv`)

### 19.20.1 Detailed Description

Be sure to see the figure in [Terms and Definitions](#) that shows a complete topology tree, including depths, child/sibling/cousin relationships, and an example of an asymmetric topology where one package has fewer caches than its peers.

### 19.20.2 Function Documentation

19.20.2.1 `unsigned hwloc_get_closest_objs ( hwloc_topology_t topology, hwloc_obj_t src, hwloc_obj_t *restrict objs, unsigned max )`

Do a depth-first traversal of the topology to find and sort.

all objects that are at the same depth than `src`. Report in `objs` up to `max` physically closest ones to `src`.

#### Returns

the number of objects returned in `objs`.  
0 if `src` is an I/O object.

#### Note

This function requires the `src` object to have a CPU set.

19.20.2.2 `static hwloc_obj_t hwloc_get_numanode_obj_by_os_index ( hwloc_topology_t topology, unsigned os_index ) [inline, static]`

Returns the object of type `HWLOC_OBJ_NUMANODE` with `os_index`.

This function is useful for converting a nodeset into the NUMA node objects it contains. When retrieving the current binding (e.g. with `hwloc_get_membind_nodeset()`), one may iterate over the bits of the resulting nodeset with `hwloc_bitmap_foreach_begin()`, and find the corresponding NUMA nodes with this function.

```
19.20.2.3 static hwloc_obj_t hwloc_get_obj_below_array_by_type (
    hwloc_topology_t topology, int nr, hwloc_obj_type_t * typev, unsigned * idxv
) [inline, static]
```

Find an object below a chain of objects specified by types and indexes.

This is a generalized version of [hwloc\\_get\\_obj\\_below\\_by\\_type\(\)](#).

Arrays `typev` and `idxv` must contain `nr` types and indexes.

Start from the top system object and walk the arrays `typev` and `idxv`. For each type and logical index couple in the arrays, look under the previously found object to find the `idx`-th object of the given type. Indexes are specified within the parent, not within the entire system.

For instance, if `nr` is 3, `typev` contains `NODE`, `PACKAGE` and `CORE`, and `idxv` contains 0, 1 and 2, return the third core object below the second package below the first NUMA node.

#### Note

This function requires all these objects and the root object to have a CPU set.

```
19.20.2.4 static hwloc_obj_t hwloc_get_obj_below_by_type ( hwloc_topology_t
    topology, hwloc_obj_type_t type1, unsigned idx1, hwloc_obj_type_t type2,
    unsigned idx2 ) [inline, static]
```

Find an object below another object, both specified by types and indexes.

Start from the top system object and find object of type `type1` and logical index `idx1`. Then look below this object and find another object of type `type2` and logical index `idx2`. Indexes are specified within the parent, not within the entire system.

For instance, if `type1` is `PACKAGE`, `idx1` is 2, `type2` is `CORE` and `idx2` is 3, return the fourth core object below the third package.

#### Note

This function requires these objects to have a CPU set.

```
19.20.2.5 static hwloc_obj_t hwloc_get_pu_obj_by_os_index ( hwloc_topology_t
    topology, unsigned os_index ) [inline, static]
```

Returns the object of type `HWLOC_OBJ_PU` with `os_index`.

This function is useful for converting a CPU set into the PU objects it contains. When retrieving the current binding (e.g. with [hwloc\\_get\\_cpupbind\(\)](#)), one may iterate over the bits of the resulting CPU set with [hwloc\\_bitmap\\_foreach\\_begin\(\)](#), and find the corresponding PUs with this function.

## 19.21 Distributing items over a topology

### Enumerations

- enum `hwloc_distrib_flags_e` { `HWLOC_DISTRIB_FLAG_REVERSE` }

### Functions

- static int `hwloc_distrib` (`hwloc_topology_t` topology, `hwloc_obj_t` \*roots, unsigned `n_roots`, `hwloc_cpuset_t` \*set, unsigned `n`, unsigned `until`, unsigned long `flags`)

#### 19.21.1 Enumeration Type Documentation

##### 19.21.1.1 enum `hwloc_distrib_flags_e`

Flags to be given to `hwloc_distrib()`.

Enumerator:

**`HWLOC_DISTRIB_FLAG_REVERSE`** Distrib in reverse order, starting from the last objects.

#### 19.21.2 Function Documentation

19.21.2.1 static int `hwloc_distrib` ( `hwloc_topology_t` topology, `hwloc_obj_t` \* roots, unsigned `n_roots`, `hwloc_cpuset_t` \* set, unsigned `n`, unsigned `until`, unsigned long `flags` ) [inline, static]

Distribute `n` items over the topology under `roots`.

Array `set` will be filled with `n` cpusets recursively distributed linearly over the topology under objects `roots`, down to depth `until` (which can be `INT_MAX` to distribute down to the finest level).

`n_roots` is usually 1 and `roots` only contains the topology root object so as to distribute over the entire topology.

This is typically useful when an application wants to distribute `n` threads over a machine, giving each of them as much private cache as possible and keeping them locally in number order.

The caller may typically want to also call `hwloc_bitmap_singlify()` before binding a thread so that it does not move at all.

`flags` should be 0 or a OR'ed set of `hwloc_distrib_flags_e`.

**Note**

This function requires the `roots` objects to have a CPU set.  
This function replaces the now deprecated `hwloc_distribute()` and `hwloc_distributev()` functions.

## 19.22 CPU and node sets of entire topologies

### Functions

- static `hwloc_const_cpuset_t hwloc_topology_get_complete_cpuset (hwloc_topology_t topology)`
- static `hwloc_const_cpuset_t hwloc_topology_get_topology_cpuset (hwloc_topology_t topology)`
- static `hwloc_const_cpuset_t hwloc_topology_get_online_cpuset (hwloc_topology_t topology)`
- static `hwloc_const_cpuset_t hwloc_topology_get_allowed_cpuset (hwloc_topology_t topology)`
- static `hwloc_const_nodeset_t hwloc_topology_get_complete_nodeset (hwloc_topology_t topology)`
- static `hwloc_const_nodeset_t hwloc_topology_get_topology_nodeset (hwloc_topology_t topology)`
- static `hwloc_const_nodeset_t hwloc_topology_get_allowed_nodeset (hwloc_topology_t topology)`

### 19.22.1 Function Documentation

19.22.1.1 `static hwloc_const_cpuset_t hwloc_topology_get_allowed_cpuset (hwloc_topology_t topology) [inline, static]`

Get allowed CPU set.

#### Returns

the CPU set of allowed logical processors of the system. If the topology is the result of a combination of several systems, NULL is returned.

#### Note

The returned cpuset is not newly allocated and should thus not be changed or freed, `hwloc_bitmap_dup()` must be used to obtain a local copy.

19.22.1.2 `static hwloc_const_nodeset_t hwloc_topology_get_allowed_nodeset (hwloc_topology_t topology) [inline, static]`

Get allowed node set.

#### Returns

the node set of allowed memory of the system. If the topology is the result of a combination of several systems, NULL is returned.

**Note**

The returned nodeset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

19.22.1.3 `static hwloc_const_cpuset_t hwloc_topology_get_complete_cpuset ( hwloc_topology_t topology ) [inline, static]`

Get complete CPU set.

**Returns**

the complete CPU set of logical processors of the system. If the topology is the result of a combination of several systems, NULL is returned.

**Note**

The returned cpuset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

19.22.1.4 `static hwloc_const_nodeset_t hwloc_topology_get_complete_nodeset ( hwloc_topology_t topology ) [inline, static]`

Get complete node set.

**Returns**

the complete node set of memory of the system. If the topology is the result of a combination of several systems, NULL is returned.

**Note**

The returned nodeset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

19.22.1.5 `static hwloc_const_cpuset_t hwloc_topology_get_online_cpuset ( hwloc_topology_t topology ) [inline, static]`

Get online CPU set.

**Returns**

the CPU set of online logical processors of the system. If the topology is the result of a combination of several systems, NULL is returned.

**Note**

The returned cpuset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

19.22.1.6 `static hwloc_const_cpuset_t hwloc_topology_get_topology_cpuset ( hwloc_topology_t topology ) [inline, static]`

Get topology CPU set.

#### Returns

the CPU set of logical processors of the system for which hwloc provides topology information. This is equivalent to the cpuset of the system object. If the topology is the result of a combination of several systems, NULL is returned.

#### Note

The returned cpuset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

19.22.1.7 `static hwloc_const_nodeset_t hwloc_topology_get_topology_nodeset ( hwloc_topology_t topology ) [inline, static]`

Get topology node set.

#### Returns

the node set of memory of the system for which hwloc provides topology information. This is equivalent to the nodeset of the system object. If the topology is the result of a combination of several systems, NULL is returned.

#### Note

The returned nodeset is not newly allocated and should thus not be changed or freed; [hwloc\\_bitmap\\_dup\(\)](#) must be used to obtain a local copy.

## 19.23 Converting between CPU sets and node sets

### Functions

- static void `hwloc_cpuset_to_nodese`t (hwloc\_topology\_t topology, hwloc\_const\_cpuset\_t cpuset, hwloc\_nodese
- static void `hwloc_cpuset_to_nodese`t\_strict (struct hwloc\_topology \*topology, hwloc\_const\_cpuset\_t cpuset, hwloc\_nodese
- static void `hwloc_cpuset_from_nodese`t (hwloc\_topology\_t topology, hwloc\_cpuset\_t cpuset, hwloc\_const\_nodese
- static void `hwloc_cpuset_from_nodese`t\_strict (struct hwloc\_topology \*topology, hwloc\_cpuset\_t cpuset, hwloc\_const\_nodese

### 19.23.1 Detailed Description

There are two semantics for converting cpusets to nodesets depending on how non-NUMA machines are handled.

When manipulating nodesets for memory binding, non-NUMA machines should be considered as having a single NUMA node. The standard conversion routines below should be used so that marking the first bit of the nodeset means that memory should be bound to a non-NUMA whole machine.

When manipulating nodesets as an actual list of NUMA nodes without any need to handle memory binding on non-NUMA machines, the strict conversion routines may be used instead.

### 19.23.2 Function Documentation

19.23.2.1 `static void hwloc_cpuset_from_nodese`t ( hwloc\_topology\_t topology, hwloc\_cpuset\_t cpuset, hwloc\_const\_nodese

Convert a NUMA node set into a CPU set and handle non-NUMA cases.

If the topology contains no NUMA nodes, the machine is considered as a single memory node, and the following behavior is used: If `nodeset` is empty, `cpuset` will be emptied as well. Otherwise `cpuset` will be entirely filled. This is useful for manipulating memory binding sets.

19.23.2.2 `static void hwloc_cpuset_from_nodese`t\_strict ( struct hwloc\_topology \* topology, hwloc\_cpuset\_t cpuset, hwloc\_const\_nodese

Convert a NUMA node set into a CPU set without handling non-NUMA cases.

This is the strict variant of `hwloc_cpuset_from_nodese`t(). It does not fix non-NUMA cases. If the topology contains some NUMA nodes, behave exactly the same. However, if the topology contains no NUMA nodes, return an empty cpuset.



```
19.23.2.3 static void hwloc_cpuset_to_nodeset ( hwloc_topology_t topology,  
hwloc_const_cpuset_t cpuset, hwloc_nodeset_t nodeset ) [inline,  
static]
```

Convert a CPU set into a NUMA node set and handle non-NUMA cases.

If some NUMA nodes have no CPUs at all, this function never sets their indexes in the output node set, even if a full CPU set is given in input.

If the topology contains no NUMA nodes, the machine is considered as a single memory node, and the following behavior is used: If `cpuset` is empty, `nodeset` will be emptied as well. Otherwise `nodeset` will be entirely filled.

```
19.23.2.4 static void hwloc_cpuset_to_nodeset_strict ( struct hwloc_topology * topology,  
hwloc_const_cpuset_t cpuset, hwloc_nodeset_t nodeset ) [inline,  
static]
```

Convert a CPU set into a NUMA node set without handling non-NUMA cases.

This is the strict variant of [hwloc\\_cpuset\\_to\\_nodeset\(\)](#). It does not fix non-NUMA cases. If the topology contains some NUMA nodes, behave exactly the same. However, if the topology contains no NUMA nodes, return an empty `nodeset`.

## 19.24 Manipulating Distances

### Functions

- static struct `hwloc_distances_s` \* `hwloc_get_whole_distance_matrix_by_depth` (`hwloc_topology_t` topology, unsigned depth)
- static struct `hwloc_distances_s` \* `hwloc_get_whole_distance_matrix_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
- static struct `hwloc_distances_s` \* `hwloc_get_distance_matrix_covering_obj_by_depth` (`hwloc_topology_t` topology, `hwloc_obj_t` obj, unsigned depth, unsigned \*firstp)
- static int `hwloc_get_latency` (`hwloc_topology_t` topology, `hwloc_obj_t` obj1, `hwloc_obj_t` obj2, float \*latency, float \*reverse\_latency)

### 19.24.1 Function Documentation

19.24.1.1 static struct `hwloc_distances_s`\* `hwloc_get_distance_matrix_covering_obj_by_depth` ( `hwloc_topology_t` topology, `hwloc_obj_t` obj, unsigned depth, unsigned \* firstp ) [static, read]

Get distances for the given depth and covering some objects.

Return a distance matrix that describes depth `depth` and covers at least object `obj` and all its children.

When looking for the distance between some objects, a common ancestor should be passed in `obj`.

`firstp` is set to logical index of the first object described by the matrix.

The returned structure belongs to the hwloc library. The caller should not modify or free it.

19.24.1.2 static int `hwloc_get_latency` ( `hwloc_topology_t` topology, `hwloc_obj_t` obj1, `hwloc_obj_t` obj2, float \* latency, float \* reverse\_latency ) [inline, static]

Get the latency in both directions between two objects.

Look at ancestor objects from the bottom to the top until one of them contains a distance matrix that matches the objects exactly.

`latency` gets the value from object `obj1` to `obj2`, while `reverse_latency` gets the reverse-direction value, which may be different on some architectures.

### Returns

- 1 if no ancestor contains a matching latency matrix.

19.24.1.3 `static struct hwloc_distances_s* hwloc_get_whole_distance_matrix_by_depth ( hwloc_topology_t topology, unsigned depth )` [*static*, *read*]

Get the distances between all objects at the given depth.

#### Returns

a distances structure containing a matrix with all distances between all objects at the given depth.

Slot `i+nbobjs*j` contains the distance from the object of logical index `i` the object of logical index `j`.

#### Note

This function only returns matrices covering the whole topology, without any unknown distance value. Those matrices are available in top-level object of the hierarchy. Matrices of lower objects are not reported here since they cover only part of the machine.

The returned structure belongs to the hwloc library. The caller should not modify or free it.

#### Returns

NULL if no such distance matrix exists.

19.24.1.4 `static struct hwloc_distances_s* hwloc_get_whole_distance_matrix_by_type ( hwloc_topology_t topology, hwloc_obj_type_t type )` [*static*, *read*]

Get the distances between all objects of a given type.

#### Returns

a distances structure containing a matrix with all distances between all objects of the given type.

Slot `i+nbobjs*j` contains the distance from the object of logical index `i` the object of logical index `j`.

#### Note

This function only returns matrices covering the whole topology, without any unknown distance value. Those matrices are available in top-level object of the hierarchy. Matrices of lower objects are not reported here since they cover only part of the machine.

The returned structure belongs to the hwloc library. The caller should not modify or free it.

**Returns**

NULL if no such distance matrix exists.

## 19.25 Finding I/O objects

### Functions

- static `hwloc_obj_t hwloc_get_non_io_ancestor_obj (hwloc_topology_t topology, hwloc_obj_t iobj)`
- static `hwloc_obj_t hwloc_get_next_pcidev (hwloc_topology_t topology, hwloc_obj_t prev)`
- static `hwloc_obj_t hwloc_get_pcidev_by_busid (hwloc_topology_t topology, unsigned domain, unsigned bus, unsigned dev, unsigned func)`
- static `hwloc_obj_t hwloc_get_pcidev_by_busidstring (hwloc_topology_t topology, const char *busid)`
- static `hwloc_obj_t hwloc_get_next_osdev (hwloc_topology_t topology, hwloc_obj_t prev)`
- static `hwloc_obj_t hwloc_get_next_bridge (hwloc_topology_t topology, hwloc_obj_t prev)`
- static `int hwloc_bridge_covers_pcibus (hwloc_obj_t bridge, unsigned domain, unsigned bus)`
- static `hwloc_obj_t hwloc_get_hostbridge_by_pcibus (hwloc_topology_t topology, unsigned domain, unsigned bus)`

### 19.25.1 Function Documentation

19.25.1.1 `static int hwloc_bridge_covers_pcibus ( hwloc_obj_t bridge, unsigned domain, unsigned bus ) [inline, static]`

19.25.1.2 `static hwloc_obj_t hwloc_get_hostbridge_by_pcibus ( hwloc_topology_t topology, unsigned domain, unsigned bus ) [inline, static]`

Find the hostbridge that covers the given PCI bus.

This is useful for finding the locality of a bus because it is the hostbridge parent cpuset.

19.25.1.3 `static hwloc_obj_t hwloc_get_next_bridge ( hwloc_topology_t topology, hwloc_obj_t prev ) [inline, static]`

Get the next bridge in the system.

#### Returns

the first bridge if `prev` is `NULL`.

19.25.1.4 `static hwloc_obj_t hwloc_get_next_osdev ( hwloc_topology_t topology, hwloc_obj_t prev ) [inline, static]`

Get the next OS device in the system.

**Returns**

the first OS device if `prev` is `NULL`.

```
19.25.1.5 static hwloc_obj_t hwloc_get_next_pcidev ( hwloc_topology_t topology,
hwloc_obj_t prev ) [inline, static]
```

Get the next PCI device in the system.

**Returns**

the first PCI device if `prev` is `NULL`.

```
19.25.1.6 static hwloc_obj_t hwloc_get_non_io_ancestor_obj ( hwloc_topology_t
topology, hwloc_obj_t iobj ) [inline, static]
```

Get the first non-I/O ancestor object.

Given the I/O object `iobj`, find the smallest non-I/O ancestor object. This regular object may then be used for binding because its locality is the same as `iobj`.

```
19.25.1.7 static hwloc_obj_t hwloc_get_pcidev_by_busid ( hwloc_topology_t
topology, unsigned domain, unsigned bus, unsigned dev, unsigned func )
[inline, static]
```

Find the PCI device object matching the PCI bus id given domain, bus device and function PCI bus id.

```
19.25.1.8 static hwloc_obj_t hwloc_get_pcidev_by_busidstring (
hwloc_topology_t topology, const char * busid ) [inline, static]
```

Find the PCI device object matching the PCI bus id given as a string `xxxx:yy:zz.t` or `yy:zz.t`.

## 19.26 The bitmap API

### Defines

- #define `hwloc_bitmap_foreach_begin`(id, bitmap)
- #define `hwloc_bitmap_foreach_end`()

### Typedefs

- typedef struct `hwloc_bitmap_s` \* `hwloc_bitmap_t`
- typedef struct `hwloc_bitmap_s` \* `hwloc_const_bitmap_t`

### Functions

- `hwloc_bitmap_t` `hwloc_bitmap_alloc` (void)
- `hwloc_bitmap_t` `hwloc_bitmap_alloc_full` (void)
- void `hwloc_bitmap_free` (`hwloc_bitmap_t` bitmap)
- `hwloc_bitmap_t` `hwloc_bitmap_dup` (`hwloc_const_bitmap_t` bitmap)
- void `hwloc_bitmap_copy` (`hwloc_bitmap_t` dst, `hwloc_const_bitmap_t` src)
- int `hwloc_bitmap_snprintf` (char \*restrict buf, size\_t buflen, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_asprintf` (char \*\*strp, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_sscanf` (`hwloc_bitmap_t` bitmap, const char \*restrict string)
- int `hwloc_bitmap_list_snprintf` (char \*restrict buf, size\_t buflen, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_list_asprintf` (char \*\*strp, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_list_sscanf` (`hwloc_bitmap_t` bitmap, const char \*restrict string)
- int `hwloc_bitmap_taskset_snprintf` (char \*restrict buf, size\_t buflen, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_taskset_asprintf` (char \*\*strp, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_taskset_sscanf` (`hwloc_bitmap_t` bitmap, const char \*restrict string)
- void `hwloc_bitmap_zero` (`hwloc_bitmap_t` bitmap)
- void `hwloc_bitmap_fill` (`hwloc_bitmap_t` bitmap)
- void `hwloc_bitmap_only` (`hwloc_bitmap_t` bitmap, unsigned id)
- void `hwloc_bitmap_allbut` (`hwloc_bitmap_t` bitmap, unsigned id)
- void `hwloc_bitmap_from_ulong` (`hwloc_bitmap_t` bitmap, unsigned long mask)
- void `hwloc_bitmap_from_ith_ulong` (`hwloc_bitmap_t` bitmap, unsigned i, unsigned long mask)
- void `hwloc_bitmap_set` (`hwloc_bitmap_t` bitmap, unsigned id)
- void `hwloc_bitmap_set_range` (`hwloc_bitmap_t` bitmap, unsigned begin, int end)
- void `hwloc_bitmap_set_ith_ulong` (`hwloc_bitmap_t` bitmap, unsigned i, unsigned long mask)
- void `hwloc_bitmap_clr` (`hwloc_bitmap_t` bitmap, unsigned id)
- void `hwloc_bitmap_clr_range` (`hwloc_bitmap_t` bitmap, unsigned begin, int end)
- void `hwloc_bitmap_singlify` (`hwloc_bitmap_t` bitmap)

- unsigned long `hwloc_bitmap_to_ulong` (`hwloc_const_bitmap_t` bitmap)
- unsigned long `hwloc_bitmap_to_ith_ulong` (`hwloc_const_bitmap_t` bitmap, unsigned i)
- int `hwloc_bitmap_isset` (`hwloc_const_bitmap_t` bitmap, unsigned id)
- int `hwloc_bitmap_iszero` (`hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_isfull` (`hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_first` (`hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_next` (`hwloc_const_bitmap_t` bitmap, int prev)
- int `hwloc_bitmap_last` (`hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_weight` (`hwloc_const_bitmap_t` bitmap)
- void `hwloc_bitmap_or` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- void `hwloc_bitmap_and` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- void `hwloc_bitmap_andnot` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- void `hwloc_bitmap_xor` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- void `hwloc_bitmap_not` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap)
- int `hwloc_bitmap_intersects` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- int `hwloc_bitmap_isincluded` (`hwloc_const_bitmap_t` sub\_bitmap, `hwloc_const_bitmap_t` super\_bitmap)
- int `hwloc_bitmap_isequal` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- int `hwloc_bitmap_compare_first` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
- int `hwloc_bitmap_compare` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)

### 19.26.1 Detailed Description

The `hwloc_bitmap_t` type represents a set of integers (positive or null). A bitmap may be of infinite size (all bits are set after some point). A bitmap may even be full if all bits are set.

Bitmaps are used by hwloc for sets of OS processors (which may actually be hardware threads) as by `hwloc_cpuset_t` (a typedef for `hwloc_bitmap_t`), or sets of NUMA memory nodes as `hwloc_nodeset_t` (also a typedef for `hwloc_bitmap_t`). Those are used for cpuset and nodeset fields in the `hwloc_obj` structure, see [Object Sets \(`hwloc\_cpuset\_t` and `hwloc\_nodeset\_t`\)](#).

*Both CPU and node sets are always indexed by OS physical number.* However users should usually not build CPU and node sets manually (e.g. with `hwloc_bitmap_set()`). One should rather use existing object sets and combine them with `hwloc_bitmap_or()`, etc. For instance, binding the current thread on a pair of cores may be performed with:



```
hwloc_obj_t core1 = ... , core2 = ... ;
hwloc_bitmap_t set = hwloc_bitmap_alloc();
hwloc_bitmap_or(set, core1->cpuset, core2->cpuset);
hwloc_set_cpupbind(topology, set, HWLOC_CPUBIND_THREAD);
hwloc_bitmap_free(set);
```

#### Note

Several examples of using the bitmap API are available under the `doc/examples/` directory in the source tree. Regression tests such as `tests/hwloc_bitmap*.c` also make intensive use of this API.

### 19.26.2 Define Documentation

#### 19.26.2.1 `#define hwloc_bitmap_foreach_begin( id, bitmap )`

Loop macro iterating on bitmap `bitmap`.

The loop must start with `hwloc_bitmap_foreach_begin()` and end with `hwloc_bitmap_foreach_end()` followed by a terminating `;`.

`index` is the loop variable; it should be an unsigned int. The first iteration will set `index` to the lowest index in the bitmap. Successive iterations will iterate through, in order, all remaining indexes set in the bitmap. To be specific: each iteration will return a value for `index` such that `hwloc_bitmap_isset(bitmap, index)` is true.

The assert prevents the loop from being infinite if the bitmap is infinitely set.

#### 19.26.2.2 `#define hwloc_bitmap_foreach_end( )`

End of loop macro iterating on a bitmap.

Needs a terminating `;`.

#### See also

[hwloc\\_bitmap\\_foreach\\_begin\(\)](#)

### 19.26.3 Typedef Documentation

#### 19.26.3.1 `typedef struct hwloc_bitmap_s* hwloc_bitmap_t`

Set of bits represented as an opaque pointer to an internal bitmap.

#### 19.26.3.2 `typedef struct hwloc_bitmap_s* hwloc_const_bitmap_t`

a non-modifiable [hwloc\\_bitmap\\_t](#)

## 19.26.4 Function Documentation

### 19.26.4.1 void hwloc\_bitmap\_allbut ( hwloc\_bitmap\_t *bitmap*, unsigned *id* )

Fill the *bitmap* and clear the index *id*.

### 19.26.4.2 hwloc\_bitmap\_t hwloc\_bitmap\_alloc ( void )

Allocate a new empty *bitmap*.

#### Returns

A valid *bitmap* or NULL.

The *bitmap* should be freed by a corresponding call to [hwloc\\_bitmap\\_free\(\)](#).

### 19.26.4.3 hwloc\_bitmap\_t hwloc\_bitmap\_alloc\_full ( void )

Allocate a new full *bitmap*.

### 19.26.4.4 void hwloc\_bitmap\_and ( hwloc\_bitmap\_t *res*, hwloc\_const\_bitmap\_t *bitmap1*, hwloc\_const\_bitmap\_t *bitmap2* )

And *bitmap1* and *bitmap2* and store the result in *res*.

*res* can be the same as *bitmap1* or *bitmap2*

### 19.26.4.5 void hwloc\_bitmap\_andnot ( hwloc\_bitmap\_t *res*, hwloc\_const\_bitmap\_t *bitmap1*, hwloc\_const\_bitmap\_t *bitmap2* )

And *bitmap1* and the negation of *bitmap2* and store the result in *res*.

*res* can be the same as *bitmap1* or *bitmap2*

### 19.26.4.6 int hwloc\_bitmap\_asprintf ( char \*\* *strp*, hwloc\_const\_bitmap\_t *bitmap* )

Stringify a *bitmap* into a newly allocated string.

#### Returns

-1 on error.

### 19.26.4.7 void hwloc\_bitmap\_clr ( hwloc\_bitmap\_t *bitmap*, unsigned *id* )

Remove index *id* from *bitmap*.

19.26.4.8 `void hwloc_bitmap_clr_range ( hwloc_bitmap_t bitmap, unsigned begin, int end )`

Remove indexes from `begin` to `end` in bitmap `bitmap`.

If `end` is `-1`, the range is infinite.

19.26.4.9 `int hwloc_bitmap_compare ( hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2 )`

Compare bitmaps `bitmap1` and `bitmap2` in lexicographic order.

Lexicographic comparison of bitmaps, starting for their highest indexes. Compare last indexes first, then second, etc. The empty bitmap is considered lower than anything.

#### Returns

- 1 if `bitmap1` is considered smaller than `bitmap2`.
- 1 if `bitmap1` is considered larger than `bitmap2`.
- 0 if bitmaps are equal (contrary to [hwloc\\_bitmap\\_compare\\_first\(\)](#)).

For instance comparing binary bitmaps 0011 and 0110 returns -1 (hence 0011 is considered smaller than 0110). Comparing 00101 and 01010 returns -1 too.

#### Note

This is different from the non-existing `hwloc_bitmap_compare_last()` which would only compare the highest index of each bitmap.

19.26.4.10 `int hwloc_bitmap_compare_first ( hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2 )`

Compare bitmaps `bitmap1` and `bitmap2` using their lowest index.

A bitmap is considered smaller if its least significant bit is smaller. The empty bitmap is considered higher than anything (because its least significant bit does not exist).

#### Returns

- 1 if `bitmap1` is considered smaller than `bitmap2`.
- 1 if `bitmap1` is considered larger than `bitmap2`.

For instance comparing binary bitmaps 0011 and 0110 returns -1 (hence 0011 is considered smaller than 0110) because least significant bit of 0011 (0001) is smaller than least significant bit of 0110 (0010). Comparing 01001 and 00110 would also return -1 for the same reason.

#### Returns

- 0 if bitmaps are considered equal, even if they are not strictly equal. They just need to have the same least significant bit. For instance, comparing binary bitmaps 0010 and 0110 returns 0 because they have the same least significant bit.

19.26.4.11 `void hwloc_bitmap_copy ( hwloc_bitmap_t dst, hwloc_const_bitmap_t src )`

Copy the contents of bitmap `src` into the already allocated bitmap `dst`.

19.26.4.12 `hwloc_bitmap_t hwloc_bitmap_dup ( hwloc_const_bitmap_t bitmap )`

Duplicate bitmap `bitmap` by allocating a new bitmap and copying `bitmap` contents. If `bitmap` is `NULL`, `NULL` is returned.

19.26.4.13 `void hwloc_bitmap_fill ( hwloc_bitmap_t bitmap )`

Fill bitmap `bitmap` with all possible indexes (even if those objects don't exist or are otherwise unavailable)

19.26.4.14 `int hwloc_bitmap_first ( hwloc_const_bitmap_t bitmap )`

Compute the first index (least significant bit) in bitmap `bitmap`.

#### Returns

-1 if no index is set in `bitmap`.

19.26.4.15 `void hwloc_bitmap_free ( hwloc_bitmap_t bitmap )`

Free bitmap `bitmap`.

If `bitmap` is `NULL`, no operation is performed.

19.26.4.16 `void hwloc_bitmap_from_ith_ulong ( hwloc_bitmap_t bitmap, unsigned i, unsigned long mask )`

Setup bitmap `bitmap` from unsigned long `mask` used as `i`-th subset.

19.26.4.17 `void hwloc_bitmap_from_ulong ( hwloc_bitmap_t bitmap, unsigned long mask )`

Setup bitmap `bitmap` from unsigned long `mask`.

19.26.4.18 `int hwloc_bitmap_intersects ( hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2 )`

Test whether bitmaps `bitmap1` and `bitmap2` intersects.

**Returns**

1 if bitmaps intersect, 0 otherwise.

19.26.4.19 `int hwloc_bitmap_isequal ( hwloc_const_bitmap_t bitmap1,  
hwloc_const_bitmap_t bitmap2 )`

Test whether bitmap `bitmap1` is equal to bitmap `bitmap2`.

**Returns**

1 if bitmaps are equal, 0 otherwise.

19.26.4.20 `int hwloc_bitmap_isfull ( hwloc_const_bitmap_t bitmap )`

Test whether bitmap `bitmap` is completely full.

**Returns**

1 if bitmap is full, 0 otherwise.

**Note**

A full bitmap is always infinitely set.

19.26.4.21 `int hwloc_bitmap_isincluded ( hwloc_const_bitmap_t sub_bitmap,  
hwloc_const_bitmap_t super_bitmap )`

Test whether bitmap `sub_bitmap` is part of bitmap `super_bitmap`.

**Returns**

1 if `sub_bitmap` is included in `super_bitmap`, 0 otherwise.

**Note**

The empty bitmap is considered included in any other bitmap.

19.26.4.22 `int hwloc_bitmap_isset ( hwloc_const_bitmap_t bitmap, unsigned id )`

Test whether index `id` is part of bitmap `bitmap`.

**Returns**

1 if the bit at index `id` is set in bitmap `bitmap`, 0 otherwise.

19.26.4.23 `int hwloc_bitmap_iszero ( hwloc_const_bitmap_t bitmap )`

Test whether bitmap `bitmap` is empty.

**Returns**

1 if bitmap is empty, 0 otherwise.

19.26.4.24 `int hwloc_bitmap_last ( hwloc_const_bitmap_t bitmap )`

Compute the last index (most significant bit) in bitmap `bitmap`.

**Returns**

-1 if no index is set in `bitmap`, or if `bitmap` is infinitely set.

19.26.4.25 `int hwloc_bitmap_list_asprintf ( char ** strp, hwloc_const_bitmap_t bitmap )`

Stringify a bitmap into a newly allocated list string.

**Returns**

-1 on error.

19.26.4.26 `int hwloc_bitmap_list_sprintf ( char *restrict buf, size_t buflen, hwloc_const_bitmap_t bitmap )`

Stringify a bitmap in the list format.

Lists are comma-separated indexes or ranges. Ranges are dash separated indexes. The last range may not have an ending indexes if the bitmap is infinitely set.

Up to `buflen` characters may be written in buffer `buf`.

If `buflen` is 0, `buf` may safely be NULL.

**Returns**

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

19.26.4.27 `int hwloc_bitmap_list_sscanf ( hwloc_bitmap_t bitmap, const char *restrict string )`

Parse a list string and stores it in bitmap `bitmap`.

**19.26.4.28** `int hwloc_bitmap_next ( hwloc_const_bitmap_t bitmap, int prev )`

Compute the next index in bitmap `bitmap` which is after index `prev`.

If `prev` is -1, the first index is returned.

**Returns**

-1 if no index with higher index is set in `bitmap`.

**19.26.4.29** `void hwloc_bitmap_not ( hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap )`

Negate bitmap `bitmap` and store the result in bitmap `res`.

`res` can be the same as `bitmap`

**19.26.4.30** `void hwloc_bitmap_only ( hwloc_bitmap_t bitmap, unsigned id )`

Empty the bitmap `bitmap` and add bit `id`.

**19.26.4.31** `void hwloc_bitmap_or ( hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2 )`

Or bitmaps `bitmap1` and `bitmap2` and store the result in bitmap `res`.

`res` can be the same as `bitmap1` or `bitmap2`

**19.26.4.32** `void hwloc_bitmap_set ( hwloc_bitmap_t bitmap, unsigned id )`

Add index `id` in bitmap `bitmap`.

**19.26.4.33** `void hwloc_bitmap_set_ith_ulong ( hwloc_bitmap_t bitmap, unsigned i, unsigned long mask )`

Replace `i`-th subset of bitmap `bitmap` with unsigned long `mask`.

**19.26.4.34** `void hwloc_bitmap_set_range ( hwloc_bitmap_t bitmap, unsigned begin, int end )`

Add indexes from `begin` to `end` in bitmap `bitmap`.

If `end` is -1, the range is infinite.

**19.26.4.35 void hwloc\_bitmap\_singlify ( hwloc\_bitmap\_t *bitmap* )**

Keep a single index among those set in bitmap *bitmap*.

May be useful before binding so that the process does not have a chance of migrating between multiple logical CPUs in the original mask. Instead of running the task on any PU inside the given CPU set, the operating system scheduler will be forced to run it on a single of these PUs. It avoids a migration overhead and cache-line ping-pongs between PUs.

**Note**

This function is NOT meant to distribute multiple processes within a single CPU set. It always return the same single bit when called multiple times on the same input set. [hwloc\\_distrib\(\)](#) may be used for generating CPU sets to distribute multiple tasks below a single multi-PU object.

This function cannot be applied to an object set directly. It should be applied to a copy (which may be obtained with [hwloc\\_bitmap\\_dup\(\)](#)).

**19.26.4.36 int hwloc\_bitmap\_sprintf ( char \*restrict *buf*, size\_t *buflen*, hwloc\_const\_bitmap\_t *bitmap* )**

Stringify a bitmap.

Up to *buflen* characters may be written in buffer *buf*.

If *buflen* is 0, *buf* may safely be NULL.

**Returns**

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

**19.26.4.37 int hwloc\_bitmap\_sscanf ( hwloc\_bitmap\_t *bitmap*, const char \*restrict *string* )**

Parse a bitmap string and stores it in bitmap *bitmap*.

**19.26.4.38 int hwloc\_bitmap\_taskset\_asprintf ( char \*\* *strp*, hwloc\_const\_bitmap\_t *bitmap* )**

Stringify a bitmap into a newly allocated taskset-specific string.

**Returns**

-1 on error.



19.26.4.39 `int hwloc_bitmap_taskset_sprintf ( char *restrict buf, size_t buflen, hwloc_const_bitmap_t bitmap )`

Stringify a bitmap in the taskset-specific format.

The taskset command manipulates bitmap strings that contain a single (possible very long) hexadecimal number starting with 0x.

Up to `buflen` characters may be written in buffer `buf`.

If `buflen` is 0, `buf` may safely be `NULL`.

#### Returns

the number of character that were actually written if not truncating, or that would have been written (not including the ending `\0`).

19.26.4.40 `int hwloc_bitmap_taskset_sscanf ( hwloc_bitmap_t bitmap, const char *restrict string )`

Parse a taskset-specific bitmap string and stores it in bitmap `bitmap`.

19.26.4.41 `unsigned long hwloc_bitmap_to_ith_ulong ( hwloc_const_bitmap_t bitmap, unsigned i )`

Convert the `i`-th subset of bitmap `bitmap` into unsigned long mask.

19.26.4.42 `unsigned long hwloc_bitmap_to_ulong ( hwloc_const_bitmap_t bitmap )`

Convert the beginning part of bitmap `bitmap` into unsigned long mask.

19.26.4.43 `int hwloc_bitmap_weight ( hwloc_const_bitmap_t bitmap )`

Compute the "weight" of bitmap `bitmap` (i.e., number of indexes that are in the bitmap).

#### Returns

the number of indexes that are in the bitmap.

-1 if `bitmap` is infinitely set.

19.26.4.44 `void hwloc_bitmap_xor ( hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2 )`

Xor bitmaps `bitmap1` and `bitmap2` and store the result in bitmap `res`.

`res` can be the same as `bitmap1` or `bitmap2`

19.26.4.45 void `hwloc_bitmap_zero` ( `hwloc_bitmap_t` *bitmap* )

Empty the bitmap `bitmap`.

## 19.27 Linux-specific helpers

### Functions

- int `hwloc_linux_parse_cpumap_file` (FILE \*file, `hwloc_cpuset_t` set)
- int `hwloc_linux_set_tid_cpupbind` (`hwloc_topology_t` topology, pid\_t tid, `hwloc_const_cpuset_t` set)
- int `hwloc_linux_get_tid_cpupbind` (`hwloc_topology_t` topology, pid\_t tid, `hwloc_cpuset_t` set)
- int `hwloc_linux_get_tid_last_cpu_location` (`hwloc_topology_t` topology, pid\_t tid, `hwloc_bitmap_t` set)

### 19.27.1 Detailed Description

This includes helpers for manipulating Linux kernel cpumap files, and hwloc equivalents of the Linux `sched_setaffinity` and `sched_getaffinity` system calls.

### 19.27.2 Function Documentation

19.27.2.1 int `hwloc_linux_get_tid_cpupbind` ( `hwloc_topology_t` topology, pid\_t tid, `hwloc_cpuset_t` set )

Get the current binding of thread `tid`.

The behavior is exactly the same as the Linux `sched_getaffinity` system call, but uses a `hwloc_cpuset`.

#### Note

This is equivalent to calling `hwloc_get_proc_cpupbind()` with `HWLOC_CPUBIND_THREAD` as flags.

19.27.2.2 int `hwloc_linux_get_tid_last_cpu_location` ( `hwloc_topology_t` topology, pid\_t tid, `hwloc_bitmap_t` set )

Get the last physical CPU where thread `tid` ran.

#### Note

This is equivalent to calling `hwloc_get_proc_last_cpu_location()` with `HWLOC_CPUBIND_THREAD` as flags.

19.27.2.3 int `hwloc_linux_parse_cpumap_file` ( FILE \* file, `hwloc_cpuset_t` set )

Convert a linux kernel cpumap file `file` into hwloc CPU set.

Might be used when reading CPU set from sysfs attributes such as topology and caches for processors, or `local_cpus` for devices.

19.27.2.4 `int hwloc_linux_set_tid_cpupbind ( hwloc_topology_t topology, pid_t tid, hwloc_const_cpuset_t set )`

Bind a thread `tid` on cpus given in cpuset `set`.

The behavior is exactly the same as the Linux `sched_setaffinity` system call, but uses a `hwloc` cpuset.

**Note**

This is equivalent to calling [hwloc\\_set\\_proc\\_cpupbind\(\)](#) with `HWLOC_CPUBIND_THREAD` as flags.

## 19.28 Interoperability with Linux libnuma unsigned long masks

### Functions

- static int `hwloc_cpuset_to_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, unsigned long \*mask, unsigned long \*maxnode)
- static int `hwloc_nodeset_to_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_const_nodeset_t` nodeset, unsigned long \*mask, unsigned long \*maxnode)
- static int `hwloc_cpuset_from_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const unsigned long \*mask, unsigned long maxnode)
- static int `hwloc_nodeset_from_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_nodeset_t` nodeset, const unsigned long \*mask, unsigned long maxnode)

### 19.28.1 Detailed Description

This interface helps converting between Linux libnuma unsigned long masks and hwloc cpusets and nodesets.

It also offers a consistent behavior on non-NUMA machines or non-NUMA-aware kernels by assuming that the machines have a single NUMA node.

#### Note

Topology `topology` must match the current machine.  
The behavior of libnuma is undefined if the kernel is not NUMA-aware. (when `CONFIG_NUMA` is not set in the kernel configuration). This helper and libnuma may thus not be strictly compatible in this case, which may be detected by checking whether `numa_available()` returns -1.

### 19.28.2 Function Documentation

**19.28.2.1** `static int hwloc_cpuset_from_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const unsigned long \* mask, unsigned long maxnode ) [`inline, static`]

Convert the array of unsigned long `mask` into hwloc CPU set.

`mask` is a array of unsigned long that will be read. `maxnode` contains the maximal node number that may be read in `mask`.

This function may be used after calling `get_mempolicy` or any other function that takes an array of unsigned long as output parameter (and possibly a maximal node number as input parameter).

```
19.28.2.2 static int hwloc_cpuset_to_linux_libnuma_ulong ( hwloc_topology_t
              topology, hwloc_const_cpuset_t cpuset, unsigned long * mask, unsigned long
              * maxnode ) [inline, static]
```

Convert hwloc CPU set `cpuset` into the array of unsigned long `mask`.

`mask` is the array of unsigned long that will be filled. `maxnode` contains the maximal node number that may be stored in `mask`. `maxnode` will be set to the maximal node number that was found, plus one.

This function may be used before calling `set_mempolicy`, `mbind`, `migrate_pages` or any other function that takes an array of unsigned long and a maximal node number as input parameter.

```
19.28.2.3 static int hwloc_nodeset_from_linux_libnuma_ulong (
              hwloc_topology_t topology, hwloc_nodeset_t nodeset, const unsigned long *
              mask, unsigned long maxnode ) [inline, static]
```

Convert the array of unsigned long `mask` into hwloc NUMA node set.

`mask` is a array of unsigned long that will be read. `maxnode` contains the maximal node number that may be read in `mask`.

This function may be used after calling `get_mempolicy` or any other function that takes an array of unsigned long as output parameter (and possibly a maximal node number as input parameter).

```
19.28.2.4 static int hwloc_nodeset_to_linux_libnuma_ulong ( hwloc_topology_t
              topology, hwloc_const_nodeset_t nodeset, unsigned long * mask, unsigned
              long * maxnode ) [inline, static]
```

Convert hwloc NUMA node set `nodeset` into the array of unsigned long `mask`.

`mask` is the array of unsigned long that will be filled. `maxnode` contains the maximal node number that may be stored in `mask`. `maxnode` will be set to the maximal node number that was found, plus one.

This function may be used before calling `set_mempolicy`, `mbind`, `migrate_pages` or any other function that takes an array of unsigned long and a maximal node number as input parameter.

## 19.29 Interoperability with Linux libnuma bitmask

### Functions

- static struct bitmask \* `hwloc_cpuset_to_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset)
- static struct bitmask \* `hwloc_nodeset_to_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_const_nodeset_t` nodeset)
- static int `hwloc_cpuset_from_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const struct bitmask \*bitmask)
- static int `hwloc_nodeset_from_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_nodeset_t` nodeset, const struct bitmask \*bitmask)

### 19.29.1 Detailed Description

This interface helps converting between Linux libnuma bitmasks and hwloc cpusets and nodesets.

It also offers a consistent behavior on non-NUMA machines or non-NUMA-aware kernels by assuming that the machines have a single NUMA node.

#### Note

Topology `topology` must match the current machine.

The behavior of libnuma is undefined if the kernel is not NUMA-aware. (when `CONFIG_NUMA` is not set in the kernel configuration). This helper and libnuma may thus not be strictly compatible in this case, which may be detected by checking whether `numa_available()` returns -1.

### 19.29.2 Function Documentation

**19.29.2.1** `static int hwloc_cpuset_from_linux_libnuma_bitmask ( hwloc_topology_t topology, hwloc_cpuset_t cpuset, const struct bitmask * bitmask )`  
`[inline, static]`

Convert libnuma bitmask `bitmask` into hwloc CPU set `cpuset`.

This function may be used after calling many `numa_*` functions that use a struct bitmask as an output parameter.

**19.29.2.2** `static struct bitmask * hwloc_cpuset_to_linux_libnuma_bitmask ( hwloc_topology_t topology, hwloc_const_cpuset_t cpuset )` `[static, read]`

Convert hwloc CPU set `cpuset` into the returned libnuma bitmask.

The returned bitmask should later be freed with `numa_bitmask_free`.

This function may be used before calling many `numa_` functions that use a struct bitmask as an input parameter.

**Returns**

newly allocated struct bitmask.

```
19.29.2.3 static int hwloc_nodese_t_from_linux_libnuma_bitmask (
    hwloc_topology_t topology, hwloc_nodese_t nodese_t, const struct bitmask *
    bitmask ) [inline, static]
```

Convert libnuma bitmask `bitmask` into hwloc NUMA node set `nodese_t`.

This function may be used after calling many `numa_` functions that use a struct bitmask as an output parameter.

```
19.29.2.4 static struct bitmask * hwloc_nodese_t_to_linux_libnuma_bitmask
    ( hwloc_topology_t topology, hwloc_const_nodese_t nodese_t )
    [static, read]
```

Convert hwloc NUMA node set `nodese_t` into the returned libnuma bitmask.

The returned bitmask should later be freed with `numa_bitmask_free`.

This function may be used before calling many `numa_` functions that use a struct bitmask as an input parameter.

**Returns**

newly allocated struct bitmask.



## 19.30 Interoperability with glibc sched affinity

### Functions

- static int `hwloc_cpuset_to_glibc_sched_affinity` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` hwlocset, `cpu_set_t` \*schedset, `size_t` schedsetsize)
- static int `hwloc_cpuset_from_glibc_sched_affinity` (`hwloc_topology_t` topology, `hwloc_cpuset_t` hwlocset, `const cpu_set_t` \*schedset, `size_t` schedsetsize)

### 19.30.1 Detailed Description

This interface offers ways to convert between hwloc cpusets and glibc cpusets such as those manipulated by `sched_getaffinity()` or `pthread_attr_setaffinity_np()`.

#### Note

Topology `topology` must match the current machine.

### 19.30.2 Function Documentation

**19.30.2.1** static int `hwloc_cpuset_from_glibc_sched_affinity` (`hwloc_topology_t` topology, `hwloc_cpuset_t` hwlocset, `const cpu_set_t` \* schedset, `size_t` schedsetsize ) [`inline`, `static`]

Convert glibc sched affinity CPU set `schedset` into hwloc CPU set.

This function may be used before calling `sched_setaffinity` or any other function that takes a `cpu_set_t` as input parameter.

`schedsetsize` should be `sizeof(cpu_set_t)` unless `schedset` was dynamically allocated with `CPU_ALLOC`

**19.30.2.2** static int `hwloc_cpuset_to_glibc_sched_affinity` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` hwlocset, `cpu_set_t` \* schedset, `size_t` schedsetsize ) [`inline`, `static`]

Convert hwloc CPU set `toposet` into glibc sched affinity CPU set `schedset`.

This function may be used before calling `sched_setaffinity` or any other function that takes a `cpu_set_t` as input parameter.

`schedsetsize` should be `sizeof(cpu_set_t)` unless `schedset` was dynamically allocated with `CPU_ALLOC`

## 19.31 Interoperability with OpenCL

### Functions

- static int [hwloc\\_opencl\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, [cl\\_device\\_id](#) device, [hwloc\\_cpuset\\_t](#) set)
- static [hwloc\\_obj\\_t](#) [hwloc\\_opencl\\_get\\_device\\_osdev\\_by\\_index](#) ([hwloc\\_topology\\_t](#) topology, unsigned platform\_index, unsigned device\_index)
- static [hwloc\\_obj\\_t](#) [hwloc\\_opencl\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, [cl\\_device\\_id](#) device)

### 19.31.1 Detailed Description

This interface offers ways to retrieve topology information about OpenCL devices.

Only the AMD OpenCL interface currently offers useful locality information about its devices.

### 19.31.2 Function Documentation

**19.31.2.1** static int [hwloc\\_opencl\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, [cl\\_device\\_id](#) device, [hwloc\\_cpuset\\_t](#) set) [*inline, static*]

Get the CPU set of logical processors that are physically close to OpenCL device *device*.

Return the CPU set describing the locality of the OpenCL device *device*.

Topology *topology* and device *device* must match the local machine. I/O devices detection and the OpenCL component are not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see [hwloc\\_opencl\\_get\\_device\\_osdev\(\)](#) and [hwloc\\_opencl\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

This function is currently only implemented in a meaningful way for Linux with the AMD OpenCL implementation; other systems will simply get a full cpuset.

**19.31.2.2** static [hwloc\\_obj\\_t](#) [hwloc\\_opencl\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, [cl\\_device\\_id](#) device) [*inline, static*]

Get the hwloc OS device object corresponding to OpenCL device *device*.

Return the hwloc OS device object that describes the given OpenCL device *device*. Return NULL if there is none.

Topology *topology* and device *device* must match the local machine. I/O devices detection and the OpenCL component must be enabled in the topology. If not, the locality of the object may still be found using [hwloc\\_opencl\\_get\\_device\\_cpuset\(\)](#).

**Note**

The corresponding hwloc PCI device may be found by looking at the result parent pointer.

```
19.31.2.3 static hwloc_obj_t hwloc_opencl_get_device_osdev_by_index (
    hwloc_topology_t topology, unsigned platform_index, unsigned device_index )
    [inline, static]
```

Get the hwloc OS device object corresponding to the OpenCL device for the given indexes.

Return the OS device object describing the OpenCL device whose platform index is `platform_index`, and whose device index within this platform is `device_index`. Return NULL if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the OpenCL component must be enabled in the topology.

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

## 19.32 Interoperability with the CUDA Driver API

### Functions

- static int [hwloc\\_cuda\\_get\\_device\\_pci\\_ids](#) ([hwloc\\_topology\\_t](#) topology, CUdevice cudevice, int \*domain, int \*bus, int \*dev)
- static int [hwloc\\_cuda\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, CUdevice cudevice, [hwloc\\_cpuset\\_t](#) set)
- static [hwloc\\_obj\\_t](#) [hwloc\\_cuda\\_get\\_device\\_pcidev](#) ([hwloc\\_topology\\_t](#) topology, - CUdevice cudevice)
- static [hwloc\\_obj\\_t](#) [hwloc\\_cuda\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, - CUdevice cudevice)
- static [hwloc\\_obj\\_t](#) [hwloc\\_cuda\\_get\\_device\\_osdev\\_by\\_index](#) ([hwloc\\_topology\\_t](#) topology, unsigned idx)

### 19.32.1 Detailed Description

This interface offers ways to retrieve topology information about CUDA devices when using the CUDA Driver API.

### 19.32.2 Function Documentation

**19.32.2.1** static int [hwloc\\_cuda\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, CUdevice cudevice, [hwloc\\_cpuset\\_t](#) set ) [inline, static]

Get the CPU set of logical processors that are physically close to device cudevice.

Return the CPU set describing the locality of the CUDA device cudevice.

Topology topology and device cudevice must match the local machine. I/O devices detection and the CUDA component are not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see [hwloc\\_cuda\\_get\\_device\\_osdev\(\)](#) and [hwloc\\_cuda\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full cpuset.

**19.32.2.2** static [hwloc\\_obj\\_t](#) [hwloc\\_cuda\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, CUdevice cudevice ) [inline, static]

Get the hwloc OS device object corresponding to CUDA device cudevice.

Return the hwloc OS device object that describes the given CUDA device cudevice. Return NULL if there is none.

Topology topology and device cudevice must match the local machine. I/O devices detection and the CUDA component must be enabled in the topology. If not, the locality of the object may still be found using [hwloc\\_cuda\\_get\\_device\\_cpuset\(\)](#).

**Note**

The corresponding hwloc PCI device may be found by looking at the result parent pointer.

**19.32.2.3** `static hwloc_obj_t hwloc_cuda_get_device_osdev_by_index ( hwloc_topology_t topology, unsigned idx ) [inline, static]`

Get the hwloc OS device object corresponding to the CUDA device whose index is `idx`.

Return the OS device object describing the CUDA device whose index is `idx`. Return NULL if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the CUDA component must be enabled in the topology.

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

This function is identical to [hwloc\\_cudart\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

**19.32.2.4** `static int hwloc_cuda_get_device_pci_ids ( hwloc_topology_t topology, CUdevice cudevice, int * domain, int * bus, int * dev ) [inline, static]`

Return the domain, bus and device IDs of the CUDA device `cudevice`.

Device `cudevice` must match the local machine.

**19.32.2.5** `static hwloc_obj_t hwloc_cuda_get_device_pcidev ( hwloc_topology_t topology, CUdevice cudevice ) [inline, static]`

Get the hwloc PCI device object corresponding to the CUDA device `cudevice`.

Return the PCI device object describing the CUDA device `cudevice`. Return NULL if there is none.

Topology `topology` and device `cudevice` must match the local machine. I/O devices detection must be enabled in topology `topology`. The CUDA component is not needed in the topology.

## 19.33 Interoperability with the CUDA Runtime API

### Functions

- static int [hwloc\\_cudart\\_get\\_device\\_pci\\_ids](#) ([hwloc\\_topology\\_t](#) topology, int idx, int \*domain, int \*bus, int \*dev)
- static int [hwloc\\_cudart\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, int idx, [hwloc\\_cpuset\\_t](#) set)
- static [hwloc\\_obj\\_t](#) [hwloc\\_cudart\\_get\\_device\\_pcidev](#) ([hwloc\\_topology\\_t](#) topology, int idx)
- static [hwloc\\_obj\\_t](#) [hwloc\\_cudart\\_get\\_device\\_osdev\\_by\\_index](#) ([hwloc\\_topology\\_t](#) topology, unsigned idx)

#### 19.33.1 Detailed Description

This interface offers ways to retrieve topology information about CUDA devices when using the CUDA Runtime API.

#### 19.33.2 Function Documentation

**19.33.2.1** static int [hwloc\\_cudart\\_get\\_device\\_cpuset](#) ( [hwloc\\_topology\\_t](#) topology, int idx, [hwloc\\_cpuset\\_t](#) set ) [inline, static]

Get the CPU set of logical processors that are physically close to device *idx*.

Return the CPU set describing the locality of the CUDA device whose index is *idx*.

Topology *topology* and device *idx* must match the local machine. I/O devices detection and the CUDA component are not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see [hwloc\\_cudart\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full cpuset.

**19.33.2.2** static [hwloc\\_obj\\_t](#) [hwloc\\_cudart\\_get\\_device\\_osdev\\_by\\_index](#) ( [hwloc\\_topology\\_t](#) topology, unsigned idx ) [inline, static]

Get the hwloc OS device object corresponding to the CUDA device whose index is *idx*.

Return the OS device object describing the CUDA device whose index is *idx*. Return NULL if there is none.

The topology *topology* does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the CUDA component must be enabled in the topology. If not, the locality of the object may still be found using [hwloc\\_cudart\\_get\\_device\\_cpuset\(\)](#).

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

This function is identical to [hwloc\\_cuda\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

**19.33.2.3** `static int hwloc_cudart_get_device_pci_ids ( hwloc_topology_t topology,  
int idx, int * domain, int * bus, int * dev ) [inline, static]`

Return the domain, bus and device IDs of the CUDA device whose index is `idx`.

Device index `idx` must match the local machine.

**19.33.2.4** `static hwloc_obj_t hwloc_cudart_get_device_pcidev ( hwloc_topology_t  
topology, int idx ) [inline, static]`

Get the hwloc PCI device object corresponding to the CUDA device whose index is `idx`.

Return the PCI device object describing the CUDA device whose index is `idx`. Return NULL if there is none.

Topology `topology` and device `idx` must match the local machine. I/O devices detection must be enabled in topology `topology`. The CUDA component is not needed in the topology.

## 19.34 Interoperability with the NVIDIA Management Library

### Functions

- static int `hwloc_nvml_get_device_cpuset` (`hwloc_topology_t` topology, `nvmlDevice_t` device, `hwloc_cpuset_t` set)
- static `hwloc_obj_t` `hwloc_nvml_get_device_osdev_by_index` (`hwloc_topology_t` topology, unsigned idx)
- static `hwloc_obj_t` `hwloc_nvml_get_device_osdev` (`hwloc_topology_t` topology, `nvmlDevice_t` device)

#### 19.34.1 Detailed Description

This interface offers ways to retrieve topology information about devices managed by the NVIDIA Management Library (NVML).

#### 19.34.2 Function Documentation

**19.34.2.1** static int `hwloc_nvml_get_device_cpuset` (`hwloc_topology_t` topology, `nvmlDevice_t` device, `hwloc_cpuset_t` set) [`inline`, `static`]

Get the CPU set of logical processors that are physically close to NVML device `device`.

Return the CPU set describing the locality of the NVML device `device`.

Topology `topology` and device `device` must match the local machine. I/O devices detection and the NVML component are not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see `hwloc_nvml_get_device_osdev()` and `hwloc_nvml_get_device_osdev_by_index()`.

This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full cpuset.

**19.34.2.2** static `hwloc_obj_t` `hwloc_nvml_get_device_osdev` (`hwloc_topology_t` topology, `nvmlDevice_t` device) [`inline`, `static`]

Get the hwloc OS device object corresponding to NVML device `device`.

Return the hwloc OS device object that describes the given NVML device `device`. Return NULL if there is none.

Topology `topology` and device `device` must match the local machine. I/O devices detection and the NVML component must be enabled in the topology. If not, the locality of the object may still be found using `hwloc_nvml_get_device_cpuset()`.



**Note**

The corresponding hwloc PCI device may be found by looking at the result parent pointer.

**19.34.2.3** `static hwloc_obj_t hwloc_nvml_get_device_osdev_by_index (`  
`hwloc_topology_t topology, unsigned idx ) [inline, static]`

Get the hwloc OS device object corresponding to the NVML device whose index is `idx`.

Return the OS device object describing the NVML device whose index is `idx`. Returns NULL if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the NVML component must be enabled in the topology.

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

## 19.35 Interoperability with OpenGL displays

### Functions

- static `hwloc_obj_t hwloc_gl_get_display_osdev_by_port_device` (`hwloc_topology_t topology`, unsigned `port`, unsigned `device`)
- static `hwloc_obj_t hwloc_gl_get_display_osdev_by_name` (`hwloc_topology_t topology`, const char `*name`)
- static int `hwloc_gl_get_display_by_osdev` (`hwloc_topology_t topology`, `hwloc_obj_t osdev`, unsigned `*port`, unsigned `*device`)

### 19.35.1 Detailed Description

This interface offers ways to retrieve topology information about OpenGL displays.

Only the NVIDIA display locality information is currently available, using the NV-CONTROL X11 extension and the NVCtrl library.

### 19.35.2 Function Documentation

**19.35.2.1** static int `hwloc_gl_get_display_by_osdev` (`hwloc_topology_t topology`, `hwloc_obj_t osdev`, unsigned `*port`, unsigned `*device`) [`inline`, `static`]

Get the OpenGL display port and device corresponding to the given hwloc OS object.

Return the OpenGL display port (server) in `port` and device (screen) in `screen` that correspond to the given hwloc OS device object. Return `-1` if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the GL component must be enabled in the topology.

**19.35.2.2** static `hwloc_obj_t hwloc_gl_get_display_osdev_by_name` (`hwloc_topology_t topology`, const char `*name`) [`inline`, `static`]

Get the hwloc OS device object corresponding to the OpenGL display given by name.

Return the OS device object describing the OpenGL display whose name is `name`, built as `":port.device"` such as `":0.0"`. Return `NULL` if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the GL component must be enabled in the topology.

#### Note

The corresponding PCI device object can be obtained by looking at the OS device parent object.

```
19.35.2.3 static hwloc_obj_t hwloc_gl_get_display_osdev_by_port_device (  
    hwloc_topology_t topology, unsigned port, unsigned device ) [inline,  
    static]
```

Get the hwloc OS device object corresponding to the OpenGL display given by port and device index.

Return the OS device object describing the OpenGL display whose port (server) is `port` and device (screen) is `device`. Return NULL if there is none.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection and the GL component must be enabled in the topology.

#### Note

The corresponding PCI device object can be obtained by looking at the OS device parent object.

## 19.36 Interoperability with Intel Xeon Phi (MIC)

### Functions

- static int [hwloc\\_intel\\_mic\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, int idx, [hwloc\\_cpuset\\_t](#) set)
- static [hwloc\\_obj\\_t](#) [hwloc\\_intel\\_mic\\_get\\_device\\_osdev\\_by\\_index](#) ([hwloc\\_topology\\_t](#) topology, unsigned idx)

### 19.36.1 Detailed Description

This interface offers ways to retrieve topology information about Intel Xeon Phi (MIC) devices.

### 19.36.2 Function Documentation

**19.36.2.1** static int [hwloc\\_intel\\_mic\\_get\\_device\\_cpuset](#) ( [hwloc\\_topology\\_t](#) topology, int idx, [hwloc\\_cpuset\\_t](#) set ) [inline, static]

Get the CPU set of logical processors that are physically close to MIC device whose index is *idx*.

Return the CPU set describing the locality of the MIC device whose index is *idx*.

Topology *topology* and device index *idx* must match the local machine. I/O devices detection is not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see [hwloc\\_intel\\_mic\\_get\\_device\\_osdev\\_by\\_index\(\)](#).

This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full cpuset.

**19.36.2.2** static [hwloc\\_obj\\_t](#) [hwloc\\_intel\\_mic\\_get\\_device\\_osdev\\_by\\_index](#) ( [hwloc\\_topology\\_t](#) topology, unsigned idx ) [inline, static]

Get the hwloc OS device object corresponding to the MIC device for the given index.

Return the OS device object describing the MIC device whose index is *idx*. Return NULL if there is none.

The topology *topology* does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection must be enabled in the topology.

#### Note

The corresponding PCI device object can be obtained by looking at the OS device parent object.

## 19.37 Interoperability with OpenFabrics

### Functions

- static int [hwloc\\_ibv\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, struct [ibv\\_device](#) \*ibdev, [hwloc\\_cpuset\\_t](#) set)
- static [hwloc\\_obj\\_t](#) [hwloc\\_ibv\\_get\\_device\\_osdev\\_by\\_name](#) ([hwloc\\_topology\\_t](#) topology, const char \*ibname)
- static [hwloc\\_obj\\_t](#) [hwloc\\_ibv\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, struct [ibv\\_device](#) \*ibdev)

### 19.37.1 Detailed Description

This interface offers ways to retrieve topology information about OpenFabrics devices (InfiniBand, Omni-Path, usNIC, etc).

### 19.37.2 Function Documentation

**19.37.2.1** static int [hwloc\\_ibv\\_get\\_device\\_cpuset](#) ([hwloc\\_topology\\_t](#) topology, struct [ibv\\_device](#) \* *ibdev*, [hwloc\\_cpuset\\_t](#) set ) [inline, static]

Get the CPU set of logical processors that are physically close to device *ibdev*.

Return the CPU set describing the locality of the OpenFabrics device *ibdev* (InfiniBand, etc).

Topology *topology* and device *ibdev* must match the local machine. I/O devices detection is not needed in the topology.

The function only returns the locality of the device. If more information about the device is needed, OS objects should be used instead, see [hwloc\\_ibv\\_get\\_device\\_osdev\(\)](#) and [hwloc\\_ibv\\_get\\_device\\_osdev\\_by\\_name\(\)](#).

This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full cpuset.

**19.37.2.2** static [hwloc\\_obj\\_t](#) [hwloc\\_ibv\\_get\\_device\\_osdev](#) ([hwloc\\_topology\\_t](#) topology, struct [ibv\\_device](#) \* *ibdev* ) [inline, static]

Get the hwloc OS device object corresponding to the OpenFabrics device *ibdev*.

Return the OS device object describing the OpenFabrics device *ibdev* (InfiniBand, etc). Returns NULL if there is none.

Topology *topology* and device *ibdev* must match the local machine. I/O devices detection must be enabled in the topology. If not, the locality of the object may still be found using [hwloc\\_ibv\\_get\\_device\\_cpuset\(\)](#).

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

**19.37.2.3** `static hwloc_obj_t hwloc_ibv_get_device_osdev_by_name (`  
`hwloc_topology_t topology, const char * ibname ) [inline, static]`

Get the hwloc OS device object corresponding to the OpenFabrics device named `ibname`.

Return the OS device object describing the OpenFabrics device (InfiniBand, Omni-Path, usNIC, etc) whose name is `ibname` (`mlx5_0`, `hfi1_0`, `usnic_0`, `qib0`, etc). Returns NULL if there is none. The name `ibname` is usually obtained from `ibv_get_device_name()`.

The topology `topology` does not necessarily have to match the current machine. For instance the topology may be an XML import of a remote host. I/O devices detection must be enabled in the topology.

**Note**

The corresponding PCI device object can be obtained by looking at the OS device parent object.

## 19.38 Interoperability with Myrinet Express

### Functions

- static int `hwloc_mx_board_get_device_cpuset` (`hwloc_topology_t` topology, unsigned id, `hwloc_cpuset_t` set)
- static int `hwloc_mx_endpoint_get_device_cpuset` (`hwloc_topology_t` topology, `mx_endpoint_t` endpoint, `hwloc_cpuset_t` set)

### 19.38.1 Detailed Description

This interface offers ways to retrieve topology information about Myrinet Express hardware.

### 19.38.2 Function Documentation

**19.38.2.1** `static int hwloc_mx_board_get_device_cpuset ( hwloc_topology_t topology, unsigned id, hwloc_cpuset_t set ) [inline, static]`

Get the CPU set of logical processors that are physically close the MX board `id`.

Return the CPU set describing the locality of the Myrinet Express board whose index is `id`.

Topology `topology` and device `id` must match the local machine. I/O devices detection is not needed in the topology.

The function only returns the locality of the device. No additional information about the device is available.

**19.38.2.2** `static int hwloc_mx_endpoint_get_device_cpuset ( hwloc_topology_t topology, mx_endpoint_t endpoint, hwloc_cpuset_t set ) [inline, static]`

Get the CPU set of logical processors that are physically close the MX endpoint `endpoint`.

Return the CPU set describing the locality of the Myrinet Express board that runs the MX endpoint `endpoint`.

Topology `topology` and device `id` must match the local machine. I/O devices detection is not needed in the topology.

The function only returns the locality of the endpoint. No additional information about the endpoint or device is available.

## 19.39 Topology differences

### Data Structures

- union [hwloc\\_topology\\_diff\\_obj\\_attr\\_u](#)  
*One object attribute difference.*
- union [hwloc\\_topology\\_diff\\_u](#)  
*One element of a difference list between two topologies.*

### Typedefs

- typedef enum [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_e](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#)
- typedef enum [hwloc\\_topology\\_diff\\_type\\_e](#) [hwloc\\_topology\\_diff\\_type\\_t](#)
- typedef union [hwloc\\_topology\\_diff\\_u](#) \* [hwloc\\_topology\\_diff\\_t](#)

### Enumerations

- enum [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_e](#) { [HWLOC\\_TOPOLOGY\\_DIFF\\_OBJ\\_ATTR\\_SIZE](#), [HWLOC\\_TOPOLOGY\\_DIFF\\_OBJ\\_ATTR\\_NAME](#), [HWLOC\\_TOPOLOGY\\_DIFF\\_OBJ\\_ATTR\\_INFO](#) }
- enum [hwloc\\_topology\\_diff\\_type\\_e](#) { [HWLOC\\_TOPOLOGY\\_DIFF\\_OBJ\\_ATTR](#), [HWLOC\\_TOPOLOGY\\_DIFF\\_TOO\\_COMPLEX](#) }
- enum [hwloc\\_topology\\_diff\\_apply\\_flags\\_e](#) { [HWLOC\\_TOPOLOGY\\_DIFF\\_APPLY\\_REVERSE](#) }

### Functions

- int [hwloc\\_topology\\_diff\\_build](#) ([hwloc\\_topology\\_t](#) topology, [hwloc\\_topology\\_t](#) new\_topology, unsigned long flags, [hwloc\\_topology\\_diff\\_t](#) \*diff)
- int [hwloc\\_topology\\_diff\\_apply](#) ([hwloc\\_topology\\_t](#) topology, [hwloc\\_topology\\_diff\\_t](#) diff, unsigned long flags)
- int [hwloc\\_topology\\_diff\\_destroy](#) ([hwloc\\_topology\\_t](#) topology, [hwloc\\_topology\\_diff\\_t](#) diff)
- int [hwloc\\_topology\\_diff\\_load\\_xml](#) ([hwloc\\_topology\\_t](#) topology, const char \*xmlpath, [hwloc\\_topology\\_diff\\_t](#) \*diff, char \*\*refname)
- int [hwloc\\_topology\\_diff\\_export\\_xml](#) ([hwloc\\_topology\\_t](#) topology, [hwloc\\_topology\\_diff\\_t](#) diff, const char \*refname, const char \*xmlpath)
- int [hwloc\\_topology\\_diff\\_load\\_xmlbuffer](#) ([hwloc\\_topology\\_t](#) topology, const char \*xmlbuffer, int buflen, [hwloc\\_topology\\_diff\\_t](#) \*diff, char \*\*refname)
- int [hwloc\\_topology\\_diff\\_export\\_xmlbuffer](#) ([hwloc\\_topology\\_t](#) topology, [hwloc\\_topology\\_diff\\_t](#) diff, const char \*refname, char \*\*xmlbuffer, int \*buflen)



### 19.39.1 Detailed Description

Applications that manipulate many similar topologies, for instance one for each node of a homogeneous cluster, may want to compress topologies to reduce the memory footprint.

This file offers a way to manipulate the difference between topologies and export/import it to/from XML. Compression may therefore be achieved by storing one topology entirely while the others are only described by their differences with the former. The actual topology can be reconstructed when actually needed by applying the precomputed difference to the reference topology.

This interface targets very similar nodes. Only very simple differences between topologies are actually supported, for instance a change in the memory size, the name of the object, or some info attribute. More complex differences such as adding or removing objects cannot be represented in the difference structures and therefore return errors. Differences between object sets cannot be represented either.

It means that there is no need to apply the difference when looking at the tree organization (how many levels, how many objects per level, what kind of objects, CPU and node sets, etc) and when binding to objects. However the difference must be applied when looking at object attributes such as the name, the memory size or info attributes.

### 19.39.2 Typedef Documentation

19.39.2.1 `typedef enum hwloc_topology_diff_obj_attr_type_e  
hwloc_topology_diff_obj_attr_type_t`

Type of one object attribute difference.

19.39.2.2 `typedef union hwloc_topology_diff_u * hwloc_topology_diff_t`

One element of a difference list between two topologies.

19.39.2.3 `typedef enum hwloc_topology_diff_type_e hwloc_topology_diff_type_t`

Type of one element of a difference list.

### 19.39.3 Enumeration Type Documentation

19.39.3.1 `enum hwloc_topology_diff_apply_flags_e`

Flags to be given to `hwloc_topology_diff_apply()`.

Enumerator:

**`HWLOC_TOPOLOGY_DIFF_APPLY_REVERSE`** Apply topology diff in reverse direction.

### 19.39.3.2 enum `hwloc_topology_diff_obj_attr_type_e`

Type of one object attribute difference.

Enumerator:

**`HWLOC_TOPOLOGY_DIFF_OBJ_ATTR_SIZE`** The object local memory is modified. The union is a `hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_uint64_s` (and the index field is ignored).

**`HWLOC_TOPOLOGY_DIFF_OBJ_ATTR_NAME`** The object name is modified. The union is a `hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_string_s` (and the name field is ignored).

**`HWLOC_TOPOLOGY_DIFF_OBJ_ATTR_INFO`** the value of an info attribute is modified. The union is a `hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_string_s`.

### 19.39.3.3 enum `hwloc_topology_diff_type_e`

Type of one element of a difference list.

Enumerator:

**`HWLOC_TOPOLOGY_DIFF_OBJ_ATTR`** An object attribute was changed. The union is a `hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_obj_attr_s`.

**`HWLOC_TOPOLOGY_DIFF_TOO_COMPLEX`** The difference is too complex, it cannot be represented. The difference below this object has not been checked. `hwloc_topology_diff_build()` will return 1. The union is a `hwloc_topology_diff_obj_attr_u::hwloc_topology_diff_too_complex_s`.

## 19.39.4 Function Documentation

### 19.39.4.1 `int hwloc_topology_diff_apply ( hwloc_topology_t topology, hwloc_topology_diff_t diff, unsigned long flags )`

Apply a topology diff to an existing topology.

`flags` is an OR'ed set of `hwloc_topology_diff_apply_flags_e`.

The new topology is modified in place. `hwloc_topology_dup()` may be used to duplicate it before patching.

If the difference cannot be applied entirely, all previous applied elements are unapplied before returning.

Returns

0 on success.

-N if applying the difference failed while trying to apply the N-th part of the difference. For instance -1 is returned if the very first difference element could not be applied.

```
19.39.4.2 int hwloc_topology_diff_build ( hwloc_topology_t topology,
    hwloc_topology_t newtopology, unsigned long flags, hwloc_topology_diff_t
    * diff )
```

Compute the difference between 2 topologies.

The difference is stored as a list of [hwloc\\_topology\\_diff\\_t](#) entries starting at `diff`. It is computed by doing a depth-first traversal of both topology trees simultaneously.

If the difference between 2 objects is too complex to be represented (for instance if some objects have different types, or different numbers of children), a special diff entry of type [HWLOC\\_TOPOLOGY\\_DIFF\\_TOO\\_COMPLEX](#) is queued. The computation of the diff does not continue below these objects. So each such diff entry means that the difference between two subtrees could not be computed.

#### Returns

0 if the difference can be represented properly.  
0 with `diff` pointing to NULL if there is no difference between the topologies.  
1 if the difference is too complex (see above). Some entries in the list will be of type [HWLOC\\_TOPOLOGY\\_DIFF\\_TOO\\_COMPLEX](#).  
-1 on any other error.

#### Note

`flags` is currently not used. It should be 0.  
The output diff has to be freed with [hwloc\\_topology\\_diff\\_destroy\(\)](#).  
The output diff can only be exported to XML or passed to [hwloc\\_topology\\_diff\\_apply\(\)](#) if 0 was returned, i.e. if no entry of type [HWLOC\\_TOPOLOGY\\_DIFF\\_TOO\\_COMPLEX](#) is listed.  
The output diff may be modified by removing some entries from the list. The removed entries should be freed by passing them to [hwloc\\_topology\\_diff\\_destroy\(\)](#) (possible as another list).

```
19.39.4.3 int hwloc_topology_diff_destroy ( hwloc_topology_t topology,
    hwloc_topology_diff_t diff )
```

Destroy a list of topology differences.

#### Note

The `topology` parameter must be a valid topology but it is not required that it is related to `diff`.

```
19.39.4.4 int hwloc_topology_diff_export_xml ( hwloc_topology_t topology,
    hwloc_topology_diff_t diff, const char * refname, const char * xmlpath )
```

Export a list of topology differences to a XML file.

If not `NULL`, `refname` defines an identifier string for the reference topology which was used as a base when computing this difference. This identifier is usually the name of the other XML file that contains the reference topology. This attribute is given back when reading the diff from XML.

**Note**

The `topology` parameter must be a valid topology but it is not required that it is related to `diff`.

```
19.39.4.5 int hwloc_topology_diff_export_xmlbuffer ( hwloc_topology_t topology,
          hwloc_topology_diff_t diff, const char * refname, char ** xmlbuffer, int * buflen
          )
```

Export a list of topology differences to a XML buffer.

If not `NULL`, `refname` defines an identifier string for the reference topology which was used as a base when computing this difference. This identifier is usually the name of the other XML file that contains the reference topology. This attribute is given back when reading the diff from XML.

The returned buffer ends with a `\0` that is included in the returned length.

**Note**

The XML buffer should later be freed with [hwloc\\_free\\_xmlbuffer\(\)](#).  
The `topology` parameter must be a valid topology but it is not required that it is related to `diff`.

```
19.39.4.6 int hwloc_topology_diff_load_xml ( hwloc_topology_t topology, const char
          * xmlpath, hwloc_topology_diff_t * diff, char ** refname )
```

Load a list of topology differences from a XML file.

If not `NULL`, `refname` will be filled with the identifier string of the reference topology for the difference file, if any was specified in the XML file. This identifier is usually the name of the other XML file that contains the reference topology.

**Note**

The `topology` parameter must be a valid topology but it is not required that it is related to `diff`.  
the pointer returned in `refname` should later be freed by the caller.

```
19.39.4.7 int hwloc_topology_diff_load_xmlbuffer ( hwloc_topology_t topology,
          const char * xmlbuffer, int buflen, hwloc_topology_diff_t * diff, char ** refname
          )
```

Load a list of topology differences from a XML buffer.

If not `NULL`, `refname` will be filled with the identifier string of the reference topology for the difference file, if any was specified in the XML file. This identifier is usually the name of the other XML file that contains the reference topology.

**Note**

The `topology` parameter must be a valid topology but it is not required that it is related to `diff`.

the pointer returned in `refname` should later be freed by the caller.

## 19.40 Components and Plugins: Discovery components

### Data Structures

- struct `hwloc_disc_component`  
*Discovery component structure.*

### Typedefs

- typedef enum `hwloc_disc_component_type_e` `hwloc_disc_component_type_t`

### Enumerations

- enum `hwloc_disc_component_type_e` { `HWLOC_DISC_COMPONENT_TYPE_CPU`, `HWLOC_DISC_COMPONENT_TYPE_GLOBAL`, `HWLOC_DISC_COMPONENT_TYPE_MISC` }

#### 19.40.1 Typedef Documentation

##### 19.40.1.1 typedef enum `hwloc_disc_component_type_e` `hwloc_disc_component_type_t`

Discovery component type.

#### 19.40.2 Enumeration Type Documentation

##### 19.40.2.1 enum `hwloc_disc_component_type_e`

Discovery component type.

Enumerator:

**`HWLOC_DISC_COMPONENT_TYPE_CPU`** CPU-only discovery through the OS, or generic no-OS support.

**`HWLOC_DISC_COMPONENT_TYPE_GLOBAL`** xml, synthetic or custom, platform-specific components such as bgq. Anything the discovers CPU and everything else. No misc backend is expected to complement a global component.

**`HWLOC_DISC_COMPONENT_TYPE_MISC`** OpenCL, Cuda, etc.

## 19.41 Components and Plugins: Discovery backends

### Data Structures

- struct `hwloc_backend`  
*Discovery backend structure.*

### Enumerations

- enum `hwloc_backend_flag_e` { `HWLOC_BACKEND_FLAG_NEED_LEVELS` }

### Functions

- struct `hwloc_backend` \* `hwloc_backend_alloc` (struct `hwloc_disc_component` \*`component`)
- int `hwloc_backend_enable` (struct `hwloc_topology` \*`topology`, struct `hwloc_backend` \*`backend`)
- int `hwloc_backends_get_obj_cpuset` (struct `hwloc_backend` \*`caller`, struct `hwloc_obj` \*`obj`, `hwloc_bitmap_t` `cpuset`)
- int `hwloc_backends_notify_new_object` (struct `hwloc_backend` \*`caller`, struct `hwloc_obj` \*`obj`)

#### 19.41.1 Enumeration Type Documentation

##### 19.41.1.1 enum `hwloc_backend_flag_e`

Backend flags.

Enumerator:

**`HWLOC_BACKEND_FLAG_NEED_LEVELS`** Levels should be reconnected before this backend `discover()` is used.

#### 19.41.2 Function Documentation

##### 19.41.2.1 struct `hwloc_backend`\* `hwloc_backend_alloc` ( struct `hwloc_disc_component` \* `component` ) [read]

Allocate a backend structure, set good default values, initialize `backend->component` and `topology`, etc. The caller will then modify whatever needed, and call `hwloc_backend_enable()`.

##### 19.41.2.2 int `hwloc_backend_enable` ( struct `hwloc_topology` \* `topology`, struct `hwloc_backend` \* `backend` )

Enable a previously allocated and setup backend.

19.41.2.3 `int hwloc_backends_get_obj_cpuset ( struct hwloc_backend * caller,  
struct hwloc_obj * obj, hwloc_bitmap_t cpuset )`

Used by backends discovery callbacks to request locality information from others.

Traverse the list of enabled backends until one has a `get_obj_cpuset()` method, and call it.

19.41.2.4 `int hwloc_backends_notify_new_object ( struct hwloc_backend * caller,  
struct hwloc_obj * obj )`

Used by backends discovery callbacks to notify other backends of new objects.

Traverse the list of enabled backends (all but caller) and invoke their `notify_new_object()` method to notify them that a new object just got added to the topology.

Currently only used for notifying of new PCI device objects.



## 19.42 Components and Plugins: Generic components

### Data Structures

- struct [hwloc\\_component](#)  
*Generic component structure.*

### Typedefs

- typedef enum [hwloc\\_component\\_type\\_e](#) [hwloc\\_component\\_type\\_t](#)

### Enumerations

- enum [hwloc\\_component\\_type\\_e](#) { [HWLOC\\_COMPONENT\\_TYPE\\_DISC](#), [HWLOC\\_COMPONENT\\_TYPE\\_XML](#) }

#### 19.42.1 Typedef Documentation

##### 19.42.1.1 typedef enum [hwloc\\_component\\_type\\_e](#) [hwloc\\_component\\_type\\_t](#)

Generic component type.

#### 19.42.2 Enumeration Type Documentation

##### 19.42.2.1 enum [hwloc\\_component\\_type\\_e](#)

Generic component type.

Enumerator:

**[HWLOC\\_COMPONENT\\_TYPE\\_DISC](#)** The data field must point to a struct [hwloc\\_disc\\_component](#).

**[HWLOC\\_COMPONENT\\_TYPE\\_XML](#)** The data field must point to a struct [hwloc\\_xml\\_component](#).

## 19.43 Components and Plugins: Core functions to be used by components

### Typedefs

- typedef void(\* [hwloc\\_report\\_error\\_t](#))(const char \*msg, int line)

### Functions

- struct [hwloc\\_obj](#) \* [hwloc\\_insert\\_object\\_by\\_cpuset](#) (struct [hwloc\\_topology](#) \*topology, [hwloc\\_obj\\_t](#) obj)
- void [hwloc\\_report\\_os\\_error](#) (const char \*msg, int line)
- int [hwloc\\_hide\\_errors](#) (void)
- struct [hwloc\\_obj](#) \* [hwloc\\_\\_insert\\_object\\_by\\_cpuset](#) (struct [hwloc\\_topology](#) \*topology, [hwloc\\_obj\\_t](#) obj, [hwloc\\_report\\_error\\_t](#) report\_error)
- void [hwloc\\_insert\\_object\\_by\\_parent](#) (struct [hwloc\\_topology](#) \*topology, [hwloc\\_obj\\_t](#) parent, [hwloc\\_obj\\_t](#) obj)
- static struct [hwloc\\_obj](#) \* [hwloc\\_alloc\\_setup\\_object](#) ([hwloc\\_obj\\_type\\_t](#) type, signed *os\_index*)
- int [hwloc\\_fill\\_object\\_sets](#) ([hwloc\\_obj\\_t](#) obj)
- static int [hwloc\\_plugin\\_check\\_namespace](#) (const char \*pluginname, const char \*symbol)

### 19.43.1 Typedef Documentation

#### 19.43.1.1 typedef void(\* [hwloc\\_report\\_error\\_t](#))(const char \*msg, int line)

Type of error callbacks during object insertion.

### 19.43.2 Function Documentation

#### 19.43.2.1 struct [hwloc\\_obj](#)\* [hwloc\\_\\_insert\\_object\\_by\\_cpuset](#) ( struct [hwloc\\_topology](#) \* *topology*, [hwloc\\_obj\\_t](#) *obj*, [hwloc\\_report\\_error\\_t](#) *report\_error* ) [read]

Add an object to the topology and specify which error callback to use.

Aside from the error callback selection, this function is identical to [hwloc\\_insert\\_object\\_by\\_cpuset\(\)](#)

#### 19.43.2.2 static struct [hwloc\\_obj](#)\* [hwloc\\_alloc\\_setup\\_object](#) ( [hwloc\\_obj\\_type\\_t](#) *type*, signed *os\_index* ) [static, read]

Allocate and initialize an object of the given type and physical index.

## **19.43 Components and Plugins: Core functions to be used by components 209**

### **19.43.2.3 int hwloc\_fill\_object\_sets ( hwloc\_obj\_t obj )**

Setup object cpusets/nodesets by OR'ing its children.

Used when adding an object late in the topology, after propagating sets up and down. The caller should use this after inserting by cpuset (which means the cpusets is already OK). Typical case: PCI backend adding a hostbridge parent.

### **19.43.2.4 int hwloc\_hide\_errors ( void )**

Check whether insertion errors are hidden.

### **19.43.2.5 struct hwloc\_obj\* hwloc\_insert\_object\_by\_cpuset ( struct hwloc\_topology \* topology, hwloc\_obj\_t obj ) [read]**

Add an object to the topology.

It is sorted along the tree of other objects according to the inclusion of cpusets, to eventually be added as a child of the smallest object including this object.

If the cpuset is empty, the type of the object (and maybe some attributes) must be enough to find where to insert the object. This is especially true for NUMA nodes with memory and no CPUs.

The given object should not have children.

This shall only be called before levels are built.

In case of error, [hwloc\\_report\\_os\\_error\(\)](#) is called.

Returns the object on success. Returns NULL and frees obj on error. Returns another object and frees obj if it was merged with an identical pre-existing object.

### **19.43.2.6 void hwloc\_insert\_object\_by\_parent ( struct hwloc\_topology \* topology, hwloc\_obj\_t parent, hwloc\_obj\_t obj )**

Insert an object somewhere in the topology.

It is added as the last child of the given parent. The cpuset is completely ignored, so strange objects such as I/O devices should preferably be inserted with this.

When used for "normal" children with cpusets (when importing from XML when duplicating a topology), the caller should make sure children are inserted in order.

The given object may have children.

Remember to call [topology\\_connect\(\)](#) afterwards to fix handy pointers.

### **19.43.2.7 static int hwloc\_plugin\_check\_namespace ( const char \* pluginname, const char \* symbol ) [inline, static]**

Make sure that plugins can lookup core symbols.

This is a sanity check to avoid lazy-lookup failures when libhwloc is loaded within a plugin, and later tries to load its own plugins. This may fail (and abort the program) if libhwloc symbols are in a private namespace.

#### Returns

- 0 on success.
- 1 if the plugin cannot be successfully loaded. The caller plugin init() callback should return a negative error code as well.

Plugins should call this function in their init() callback to avoid later crashes if lazy symbol resolution is used by the upper layer that loaded hwloc (e.g. OpenCL implementations using dlopen with RTLD\_LAZY).

#### Note

- The build system must define HWLOC\_INSIDE\_PLUGIN if and only if building the caller as a plugin.
- This function should remain inline so plugins can call it even when they cannot find libhwloc symbols.

#### 19.43.2.8 void hwloc\_report\_os\_error ( const char \* *msg*, int *line* )

Report an insertion error from a backend.

## 19.44 Components and Plugins: PCI functions to be used by components

### Functions

- int `hwloc_insert_pci_device_list` (struct `hwloc_backend` \*backend, struct `hwloc_obj` \*first\_obj)
- unsigned `hwloc_pci_find_cap` (const unsigned char \*config, unsigned cap)
- int `hwloc_pci_find_linkspeed` (const unsigned char \*config, unsigned offset, float \*linkspeed)
- int `hwloc_pci_prepare_bridge` (`hwloc_obj_t` obj, const unsigned char \*config)

#### 19.44.1 Function Documentation

19.44.1.1 int `hwloc_insert_pci_device_list` ( struct `hwloc_backend` \* *backend*, struct `hwloc_obj` \* *first\_obj* )

Insert a list of PCI devices and bridges in the backend topology.

Insert a list of objects (either PCI device or bridges) starting at `first_obj` (linked by `next_sibling` in the topology, and ending with NULL). Objects are placed under the right bridges, and the remaining upstream bridges are then inserted in the topology by calling the `get_obj_cpuset()` callback to find their locality.

19.44.1.2 unsigned `hwloc_pci_find_cap` ( const unsigned char \* *config*, unsigned *cap* )

Return the offset of the given capability in the PCI config space buffer.

This function requires a 256-bytes config space. Unknown/unavailable bytes should be set to 0xff.

19.44.1.3 int `hwloc_pci_find_linkspeed` ( const unsigned char \* *config*, unsigned *offset*, float \* *linkspeed* )

Fill `linkspeed` by reading the PCI config space where `PCI_CAP_ID_EXP` is at position `offset`.

Needs 20 bytes of EXP capability block starting at `offset` in the config space for registers up to link status.

19.44.1.4 int `hwloc_pci_prepare_bridge` ( `hwloc_obj_t` *obj*, const unsigned char \* *config* )

Modify the PCI device object into a bridge and fill its attribute if a bridge is found in the PCI config space.

This function requires 64 bytes of common configuration header at the beginning of config.

Returns -1 and destroys /p obj if bridge fields are invalid.

## Chapter 20

# Data Structure Documentation

### 20.1 hwloc\_backend Struct Reference

```
#include <plugins.h>
```

#### Data Fields

- unsigned long [flags](#)
- int [is\\_custom](#)
- int [is\\_thissystem](#)
- void \* [private\\_data](#)
- void(\* [disable](#) )(struct [hwloc\\_backend](#) \*backend)
- int(\* [discover](#) )(struct [hwloc\\_backend](#) \*backend)
- int(\* [get\\_obj\\_cpuset](#) )(struct [hwloc\\_backend](#) \*backend, struct [hwloc\\_backend](#) \*caller, struct [hwloc\\_obj](#) \*obj, [hwloc\\_bitmap\\_t](#) cpuset)
- int(\* [notify\\_new\\_object](#) )(struct [hwloc\\_backend](#) \*backend, struct [hwloc\\_backend](#) \*caller, struct [hwloc\\_obj](#) \*obj)

#### 20.1.1 Detailed Description

Discovery backend structure.

A backend is the instantiation of a discovery component. When a component gets enabled for a topology, its `instantiate()` callback creates a backend.

[hwloc\\_backend\\_alloc\(\)](#) initializes all fields to default values that the component may change (except "component" and "next") before enabling the backend with [hwloc\\_backend\\_enable\(\)](#).

#### 20.1.2 Field Documentation

**20.1.2.1** `void(* hwloc_backend::disable)(struct hwloc_backend *backend)`

Callback for freeing the `private_data`. May be NULL.

**20.1.2.2** `int(* hwloc_backend::discover)(struct hwloc_backend *backend)`

Main discovery callback. returns  $> 0$  if it modified the topology tree,  $-1$  on error,  $0$  otherwise. May be NULL if type is `HWLOC_DISC_COMPONENT_TYPE_MISC`.

**20.1.2.3** `unsigned long hwloc_backend::flags`

Backend flags, as an OR'ed set of `hwloc_backend_flag_e`.

**20.1.2.4** `int(* hwloc_backend::get_obj_cpuset)(struct hwloc_backend *backend, struct hwloc_backend *caller, struct hwloc_obj *obj, hwloc_bitmap_t cpuset)`

Callback used by the PCI backend to retrieve the locality of a PCI object from the OS/cpu backend. May be NULL.

**20.1.2.5** `int hwloc_backend::is_custom`

Backend-specific 'is\_custom' property. Shortcut on `!strcmp(..->component->name, "custom")`. Only the custom component should touch this.

**20.1.2.6** `int hwloc_backend::is_thissystem`

Backend-specific 'is\_thissystem' property. Set to  $0$  or  $1$  if the backend should enforce the `thissystem` flag when it gets enabled. Set to  $-1$  if the backend doesn't care (default).

**20.1.2.7** `int(* hwloc_backend::notify_new_object)(struct hwloc_backend *backend, struct hwloc_backend *caller, struct hwloc_obj *obj)`

Callback called by backends to notify this backend that a new object was added. returns  $> 0$  if it modified the topology tree,  $0$  otherwise. May be NULL.

**20.1.2.8** `void* hwloc_backend::private_data`

Backend private data, or NULL if none.

The documentation for this struct was generated from the following file:

- `plugins.h`



## 20.2 hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- union {
  - struct [hwloc\\_pcidev\\_attr\\_s](#) [pci](#)
- [upstream](#)
- [hwloc\\_obj\\_bridge\\_type\\_t](#) [upstream\\_type](#)
- union {
  - struct {
    - unsigned short [domain](#)
    - unsigned char [secondary\\_bus](#)
    - unsigned char [subordinate\\_bus](#)
  - [pci](#)
- [downstream](#)
- [hwloc\\_obj\\_bridge\\_type\\_t](#) [downstream\\_type](#)
- unsigned [depth](#)

### 20.2.1 Detailed Description

Bridge specific Object Attributes.

### 20.2.2 Field Documentation

20.2.2.1 unsigned [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::depth](#)

20.2.2.2 unsigned short [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::domain](#)

20.2.2.3 union { ... } [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::downstream](#)

20.2.2.4 [hwloc\\_obj\\_bridge\\_type\\_t](#) [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::downstream\\_type](#)

20.2.2.5 struct [hwloc\\_pcidev\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::pci](#)

20.2.2.6 struct { ... } [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::pci](#)

20.2.2.7 unsigned char [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::secondary\\_bus](#)

20.2.2.8 unsigned char [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s::subordinate\\_bus](#)

20.2.2.9 union { ... } hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s::upstream

20.2.2.10 hwloc\_obj\_bridge\_type\_t hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s::upstream\_type

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.3 hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- hwloc\_uint64\_t [size](#)
- unsigned [depth](#)
- unsigned [linesize](#)
- int [associativity](#)
- [hwloc\\_obj\\_cache\\_type\\_t](#) type

#### 20.3.1 Detailed Description

Cache-specific Object Attributes.

#### 20.3.2 Field Documentation

20.3.2.1 int hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s::associativity

Ways of associativity.

-1 if fully associative, 0 if unknown

20.3.2.2 unsigned hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s::depth

Depth of cache (e.g., L1, L2, ...etc.)

20.3.2.3 unsigned hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s::linesize

Cache-line size in bytes. 0 if unknown.

20.3.2.4 hwloc\_uint64\_t hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s::size

Size of cache in bytes.

### 20.3.2.5 hwloc\_obj\_cache\_type\_t hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s::type

Cache type.

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.4 hwloc\_component Struct Reference

```
#include <plugins.h>
```

### Data Fields

- unsigned [abi](#)
- int(\* [init](#) )(unsigned long [flags](#))
- void(\* [finalize](#) )(unsigned long [flags](#))
- [hwloc\\_component\\_type\\_t](#) type
- unsigned long [flags](#)
- void \* [data](#)

### 20.4.1 Detailed Description

Generic component structure.

Generic components structure, either statically listed by configure in static-components.h or dynamically loaded as a plugin.

### 20.4.2 Field Documentation

#### 20.4.2.1 unsigned hwloc\_component::abi

Component ABI version, set to [HWLOC\\_COMPONENT\\_ABI](#).

#### 20.4.2.2 void\* hwloc\_component::data

Component data, pointing to a struct [hwloc\\_disc\\_component](#) or struct hwloc\_xml\_component.

#### 20.4.2.3 void(\* hwloc\_component::finalize)(unsigned long flags)

Process-wide component termination callback.

This optional callback is called after unregistering the component from the hwloc core (before unloading the plugin).

`flags` is always 0 for now.

#### Note

If the component uses ltdl for loading its own plugins, it should load/unload them only in `init()` and `finalize()`, to avoid race conditions with hwloc's use of ltdl.

#### 20.4.2.4 unsigned long `hwloc_component::flags`

Component flags, unused for now.

#### 20.4.2.5 `int(* hwloc_component::init)(unsigned long flags)`

Process-wide component initialization callback.

This optional callback is called when the component is registered to the hwloc core (after loading the plugin).

When the component is built as a plugin, this callback should call `hwloc_check_plugin_namespace()` and return a negative error code on error.

`flags` is always 0 for now.

#### Returns

0 on success, or a negative code on error.

#### Note

If the component uses ltdl for loading its own plugins, it should load/unload them only in `init()` and `finalize()`, to avoid race conditions with hwloc's use of ltdl.

#### 20.4.2.6 `hwloc_component_type_t hwloc_component::type`

Component type.

The documentation for this struct was generated from the following file:

- `plugins.h`

## 20.5 `hwloc_disc_component` Struct Reference

```
#include <plugins.h>
```

## Data Fields

- [hwloc\\_disc\\_component\\_type\\_t](#) type
- const char \* [name](#)
- unsigned [excludes](#)
- struct [hwloc\\_backend](#) \*(\* [instantiate](#))(struct [hwloc\\_disc\\_component](#) \*component, const void \*data1, const void \*data2, const void \*data3)
- unsigned [priority](#)

### 20.5.1 Detailed Description

Discovery component structure.

This is the major kind of components, taking care of the discovery. They are registered by generic components, either statically-built or as plugins.

### 20.5.2 Field Documentation

#### 20.5.2.1 unsigned hwloc\_disc\_component::excludes

Component types to exclude, as an OR'ed set of [hwloc\\_disc\\_component\\_type\\_e](#).

For a GLOBAL component, this usually includes all other types (~0).

Other components only exclude types that may bring conflicting topology information. MISC components should likely not be excluded since they usually bring non-primary additional information.

#### 20.5.2.2 struct hwloc\_backend>(\* hwloc\_disc\_component::instantiate)(struct hwloc\_disc\_component \*component, const void \*data1, const void \*data2, const void \*data3) [read]

Instantiate callback to create a backend from the component. Parameters data1, data2, data3 are NULL except for components that have special enabling routines such as [hwloc\\_topology\\_set\\_xml\(\)](#).

#### 20.5.2.3 const char\* hwloc\_disc\_component::name

Name. If this component is built as a plugin, this name does not have to match the plugin filename.

#### 20.5.2.4 unsigned hwloc\_disc\_component::priority

Component priority. Used to sort topology->components, higher priority first. Also used to decide between two components with the same name.

Usual values are 50 for native OS (or platform) components, 45 for x86, 40 for no-OS fallback, 30 for global components (xml/synthetic/custom), 20 for pci, 10 for other misc components (opencl etc.).

#### 20.5.2.5 `hwloc_disc_component_type_t hwloc_disc_component::type`

Discovery component type.

The documentation for this struct was generated from the following file:

- `plugins.h`

## 20.6 `hwloc_distances_s` Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned `relative_depth`
- unsigned `nbobjs`
- float \* `latency`
- float `latency_max`
- float `latency_base`

#### 20.6.1 Detailed Description

Distances between objects.

One object may contain a distance structure describing distances between all its descendants at a given relative depth. If the containing object is the root object of the topology, then the distances are available for all objects in the machine.

If the `latency` pointer is not `NULL`, the pointed array contains memory latencies (non-zero values), see below.

In the future, some other types of distances may be considered. In these cases, `latency` may be `NULL`.

#### 20.6.2 Field Documentation

##### 20.6.2.1 `float* hwloc_distances_s::latency`

Matrix of latencies between objects, stored as a one-dimension array. May be `NULL` if the distances considered here are not latencies.

Unless defined by the user, this currently contains latencies between NUMA nodes (as reported in the System Locality Distance Information Table (SLIT) in the ACPI specification), which may or may not be accurate. It corresponds to the latency for accessing the memory of one node from a core in another node.

Values are normalized to get 1.0 as the minimal value in the matrix. Latency from i-th to j-th object is stored in slot  $i * nbobjs + j$ .

#### 20.6.2.2 float hwloc\_distances\_s::latency\_base

The multiplier that should be applied to latency matrix to retrieve the original OS-provided latencies. Usually 10 on Linux since ACPI SLIT uses 10 for local latency.

#### 20.6.2.3 float hwloc\_distances\_s::latency\_max

The maximal value in the latency matrix.

#### 20.6.2.4 unsigned hwloc\_distances\_s::nbobjs

Number of objects considered in the matrix. It is the number of descendant objects at `relative_depth` below the containing object. It corresponds to the result of [hwloc\\_get\\_nbobjs\\_inside\\_cpuset\\_by\\_depth\(\)](#).

#### 20.6.2.5 unsigned hwloc\_distances\_s::relative\_depth

Relative depth of the considered objects below the object containing this distance information.

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.7 hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned [depth](#)
- unsigned char [dont\\_merge](#)

### 20.7.1 Detailed Description

Group-specific Object Attributes.

## 20.7.2 Field Documentation

### 20.7.2.1 unsigned hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s::depth

Depth of group object.

### 20.7.2.2 unsigned char hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s::dont\_merge

Flag preventing groups from being automatically merged with identical parent or children.

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.8 hwloc\_obj Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- [hwloc\\_obj\\_type\\_t](#) type
- unsigned [os\\_index](#)
- char \* [name](#)
- struct [hwloc\\_obj\\_memory\\_s](#) memory
- union [hwloc\\_obj\\_attr\\_u](#) \* attr
- unsigned [depth](#)
- unsigned [logical\\_index](#)
- signed [os\\_level](#)
- struct [hwloc\\_obj](#) \* [next\\_cousin](#)
- struct [hwloc\\_obj](#) \* [prev\\_cousin](#)
- struct [hwloc\\_obj](#) \* [parent](#)
- unsigned [sibling\\_rank](#)
- struct [hwloc\\_obj](#) \* [next\\_sibling](#)
- struct [hwloc\\_obj](#) \* [prev\\_sibling](#)
- unsigned [arity](#)
- struct [hwloc\\_obj](#) \*\* [children](#)
- struct [hwloc\\_obj](#) \* [first\\_child](#)
- struct [hwloc\\_obj](#) \* [last\\_child](#)
- void \* [userdata](#)
- [hwloc\\_cpuset\\_t](#) [cpuset](#)
- [hwloc\\_cpuset\\_t](#) [complete\\_cpuset](#)
- [hwloc\\_cpuset\\_t](#) [online\\_cpuset](#)
- [hwloc\\_cpuset\\_t](#) [allowed\\_cpuset](#)
- [hwloc\\_nodese\\_t](#) [nodeset](#)



- [hwloc\\_nodese\\_t complete\\_nodese](#)
- [hwloc\\_nodese\\_t allowed\\_nodese](#)
- struct [hwloc\\_distances\\_s](#) \*\* [distances](#)
- unsigned [distances\\_count](#)
- struct [hwloc\\_obj\\_info\\_s](#) \* [infos](#)
- unsigned [infos\\_count](#)
- int [symmetric\\_subtree](#)

### 20.8.1 Detailed Description

Structure of a topology object.

Applications must not modify any field except [hwloc\\_obj.userdata](#).

### 20.8.2 Field Documentation

#### 20.8.2.1 [hwloc\\_cpuset\\_t hwloc\\_obj::allowed\\_cpuset](#)

The CPU set of allowed logical processors.

This includes the CPUs contained in this object which are allowed for binding, i.e. passing them to the hwloc binding functions should not return permission errors. This is usually restricted by administration rules. Some of them may however be offline so binding to them may still not be possible, see [online\\_cpuset](#).

#### Note

Its value must not be changed, [hwloc\\_bitmap\\_dup\(\)](#) must be used instead.

#### 20.8.2.2 [hwloc\\_nodese\\_t hwloc\\_obj::allowed\\_nodese](#)

The set of allowed NUMA memory nodes.

This includes the NUMA memory nodes contained in this object which are allowed for memory allocation, i.e. passing them to NUMA node-directed memory allocation should not return permission errors. This is usually restricted by administration rules.

If there are no NUMA nodes in the machine, all the memory is close to this object, so [allowed\\_nodese](#) is full.

#### Note

Its value must not be changed, [hwloc\\_bitmap\\_dup\(\)](#) must be used instead.

#### 20.8.2.3 [unsigned hwloc\\_obj::arity](#)

Number of children.

#### 20.8.2.4 union `hwloc_obj_attr_u*` `hwloc_obj::attr`

Object type-specific Attributes, may be `NULL` if no attribute value was found.

#### 20.8.2.5 struct `hwloc_obj**` `hwloc_obj::children`

Children, `children[0 .. arity -1]`.

#### 20.8.2.6 `hwloc_cpuset_t` `hwloc_obj::complete_cpuset`

The complete CPU set of logical processors of this object,.

This includes not only the same as the `cpuset` field, but also some CPUs for which topology information is unknown or incomplete, and the CPUs that are ignored when the `HWLOC_TOPOLOGY_FLAG_WHOLE_SYSTEM` flag is not set. Thus no corresponding PU object may be found in the topology, because the precise position is undefined. It is however known that it would be somewhere under this object.

##### Note

Its value must not be changed, `hwloc_bitmap_dup()` must be used instead.

#### 20.8.2.7 `hwloc_nodest_t` `hwloc_obj::complete_nodest`

The complete NUMA node set of this object,.

This includes not only the same as the `nodest` field, but also some NUMA nodes for which topology information is unknown or incomplete, and the nodes that are ignored when the `HWLOC_TOPOLOGY_FLAG_WHOLE_SYSTEM` flag is not set. Thus no corresponding NUMA node object may be found in the topology, because the precise position is undefined. It is however known that it would be somewhere under this object.

If there are no NUMA nodes in the machine, all the memory is close to this object, so `complete_nodest` is full.

##### Note

Its value must not be changed, `hwloc_bitmap_dup()` must be used instead.

#### 20.8.2.8 `hwloc_cpuset_t` `hwloc_obj::cpuset`

CPUs covered by this object.

This is the set of CPUs for which there are PU objects in the topology under this object, i.e. which are known to be physically contained in this object and known how (the children path between this object and the PU objects).

If the `HWLOC_TOPOLOGY_FLAG_WHOLE_SYSTEM` configuration flag is set, some of these CPUs may be offline, or not allowed for binding, see `online_cpuset` and `allowed_cpuset`.

**Note**

Its value must not be changed, [hwloc\\_bitmap\\_dup\(\)](#) must be used instead.

**20.8.2.9 unsigned hwloc\_obj::depth**

Vertical index in the hierarchy.

For normal objects, this is the depth of the horizontal level that contains this object and its cousins of the same type. If the topology is symmetric, this is equal to the parent depth plus one, and also equal to the number of parent/child links from the root object to here.

For special objects (I/O and Misc) that are not in the main tree, this is a special negative value that corresponds to their dedicated level, see [hwloc\\_get\\_type\\_depth\(\)](#) and [hwloc\\_get\\_type\\_depth\\_e](#). Those special values can be passed to hwloc functions such [hwloc\\_get\\_nbobjs\\_by\\_depth\(\)](#) as usual.

**20.8.2.10 struct hwloc\_distances\_s\*\* hwloc\_obj::distances**

Distances between all objects at same depth below this object.

**20.8.2.11 unsigned hwloc\_obj::distances\_count****20.8.2.12 struct hwloc\_obj\* hwloc\_obj::first\_child**

First child.

**20.8.2.13 struct hwloc\_obj\_info\_s\* hwloc\_obj::infos**

Array of stringified info type=name.

**20.8.2.14 unsigned hwloc\_obj::infos\_count**

Size of infos array.

**20.8.2.15 struct hwloc\_obj\* hwloc\_obj::last\_child**

Last child.

**20.8.2.16 unsigned hwloc\_obj::logical\_index**

Horizontal index in the whole list of similar objects, hence guaranteed unique across the entire machine. Could be a "cousin\_rank" since it's the rank within the "cousin" list below.

**20.8.2.17 struct hwloc\_obj\_memory\_s hwloc\_obj::memory**

Memory attributes.

**20.8.2.18 char\* hwloc\_obj::name**

Object-specific name if any. Mostly used for identifying OS devices and Misc objects where a name string is more useful than numerical indexes.

**20.8.2.19 struct hwloc\_obj\* hwloc\_obj::next\_cousin**

Next object of same type and depth.

**20.8.2.20 struct hwloc\_obj\* hwloc\_obj::next\_sibling**

Next object below the same parent.

**20.8.2.21 hwloc\_nodese\_t hwloc\_obj::nodeset**

NUMA nodes covered by this object or containing this object.

This is the set of NUMA nodes for which there are NUMA node objects in the topology under or above this object, i.e. which are known to be physically contained in this object or containing it and known how (the children path between this object and the NUMA node objects).

In the end, these nodes are those that are close to the current object.

If the [HWLOC\\_TOPOLOGY\\_FLAG\\_WHOLE\\_SYSTEM](#) configuration flag is set, some of these nodes may not be allowed for allocation, see `allowed_nodese`.

If there are no NUMA nodes in the machine, all the memory is close to this object, so `nodeset` is full.

**Note**

Its value must not be changed, [hwloc\\_bitmap\\_dup\(\)](#) must be used instead.

**20.8.2.22 hwloc\_cpuset\_t hwloc\_obj::online\_cpuset**

The CPU set of online logical processors.

This includes the CPUs contained in this object that are online, i.e. draw power and can execute threads. It may however not be allowed to bind to them due to administration rules, see `allowed_cpuset`.

**Note**

Its value must not be changed, [hwloc\\_bitmap\\_dup\(\)](#) must be used instead.

**20.8.2.23 unsigned hwloc\_obj::os\_index**

OS-provided physical index number. It is not guaranteed unique across the entire machine, except for PUs and NUMA nodes.

**20.8.2.24 signed hwloc\_obj::os\_level**

OS-provided physical level, -1 if unknown or meaningless.

**20.8.2.25 struct hwloc\_obj\* hwloc\_obj::parent**

Parent, `NULL` if root (system object)

**20.8.2.26 struct hwloc\_obj\* hwloc\_obj::prev\_cousin**

Previous object of same type and depth.

**20.8.2.27 struct hwloc\_obj\* hwloc\_obj::prev\_sibling**

Previous object below the same parent.

**20.8.2.28 unsigned hwloc\_obj::sibling\_rank**

Index in parent's `children[]` array.

**20.8.2.29 int hwloc\_obj::symmetric\_subtree**

Set if the subtree of objects below this object is symmetric, which means all children and their children have identical subtrees. If set in the topology root object, `lstopo` may export the topology as a synthetic string.

**20.8.2.30 hwloc\_obj\_type\_t hwloc\_obj::type**

Type of object.

**20.8.2.31 void\* hwloc\_obj::userdata**

Application-given private data pointer, initialized to `NULL`, use it as you wish. See [hwloc\\_topology\\_set\\_userdata\\_export\\_callback\(\)](#) if you wish to export this field to XML.

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.9 hwloc\_obj\_attr\_u Union Reference

```
#include <hwloc.h>
```

### Data Structures

- struct [hwloc\\_bridge\\_attr\\_s](#)  
*Bridge specific Object Attributes.*
- struct [hwloc\\_cache\\_attr\\_s](#)  
*Cache-specific Object Attributes.*
- struct [hwloc\\_group\\_attr\\_s](#)  
*Group-specific Object Attributes.*
- struct [hwloc\\_osdev\\_attr\\_s](#)  
*OS Device specific Object Attributes.*
- struct [hwloc\\_pcidev\\_attr\\_s](#)  
*PCI Device specific Object Attributes.*

### Data Fields

- struct [hwloc\\_obj\\_attr\\_u::hwloc\\_cache\\_attr\\_s](#) cache
- struct [hwloc\\_obj\\_attr\\_u::hwloc\\_group\\_attr\\_s](#) group
- struct [hwloc\\_obj\\_attr\\_u::hwloc\\_pcidev\\_attr\\_s](#) pcidev
- struct [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s](#) bridge
- struct [hwloc\\_obj\\_attr\\_u::hwloc\\_osdev\\_attr\\_s](#) osdev

#### 20.9.1 Detailed Description

Object type-specific Attributes.

#### 20.9.2 Field Documentation

20.9.2.1 struct [hwloc\\_obj\\_attr\\_u::hwloc\\_bridge\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::bridge](#)

20.9.2.2 struct [hwloc\\_obj\\_attr\\_u::hwloc\\_cache\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::cache](#)

20.9.2.3 struct [hwloc\\_obj\\_attr\\_u::hwloc\\_group\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::group](#)

20.9.2.4 struct [hwloc\\_obj\\_attr\\_u::hwloc\\_osdev\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::osdev](#)

20.9.2.5 struct [hwloc\\_obj\\_attr\\_u::hwloc\\_pcidev\\_attr\\_s](#) [hwloc\\_obj\\_attr\\_u::pcidev](#)

The documentation for this union was generated from the following file:

- [hwloc.h](#)

## 20.10 hwloc\_obj\_info\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- char \* [name](#)
- char \* [value](#)

### 20.10.1 Detailed Description

Object info.

#### See also

[Consulting and Adding Key-Value Info Attributes](#)

### 20.10.2 Field Documentation

#### 20.10.2.1 char\* hwloc\_obj\_info\_s::name

Info name.

#### 20.10.2.2 char\* hwloc\_obj\_info\_s::value

Info value.

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.11 hwloc\_obj\_memory\_s::hwloc\_obj\_memory\_page\_type\_s - Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- hwloc\_uint64\_t [size](#)
- hwloc\_uint64\_t [count](#)

### 20.11.1 Detailed Description

Array of local memory page types, `NULL` if no local memory and `page_types` is 0. The array is sorted by increasing `size` fields. It contains `page_types_len` slots.

### 20.11.2 Field Documentation

#### 20.11.2.1 `hwloc_uint64_t hwloc_obj_memory_s::hwloc_obj_memory_page_type_s::count`

Number of pages of this size.

#### 20.11.2.2 `hwloc_uint64_t hwloc_obj_memory_s::hwloc_obj_memory_page_type_s::size`

Size of pages.

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.12 `hwloc_obj_memory_s` Struct Reference

```
#include <hwloc.h>
```

### Data Structures

- struct [hwloc\\_obj\\_memory\\_page\\_type\\_s](#)

*Array of local memory page types, `NULL` if no local memory and `page_types` is 0.*

### Data Fields

- `hwloc_uint64_t` [total\\_memory](#)
- `hwloc_uint64_t` [local\\_memory](#)
- unsigned [page\\_types\\_len](#)
- struct [hwloc\\_obj\\_memory\\_s::hwloc\\_obj\\_memory\\_page\\_type\\_s](#) \* [page\\_types](#)

### 20.12.1 Detailed Description

Object memory.



## 20.12.2 Field Documentation

### 20.12.2.1 hwloc\_uint64\_t hwloc\_obj\_memory\_s::local\_memory

Local memory (in bytes)

### 20.12.2.2 struct hwloc\_obj\_memory\_s::hwloc\_obj\_memory\_page\_type\_s \* hwloc\_obj\_memory\_s::page\_types

### 20.12.2.3 unsigned hwloc\_obj\_memory\_s::page\_types\_len

Size of array `page_types`.

### 20.12.2.4 hwloc\_uint64\_t hwloc\_obj\_memory\_s::total\_memory

Total memory (in bytes) in this object and its children.

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.13 hwloc\_obj\_attr\_u::hwloc\_osdev\_attr\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- [hwloc\\_obj\\_osdev\\_type\\_t](#) type

### 20.13.1 Detailed Description

OS Device specific Object Attributes.

### 20.13.2 Field Documentation

#### 20.13.2.1 hwloc\_obj\_osdev\_type\_t hwloc\_obj\_attr\_u::hwloc\_osdev\_attr\_s::type

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.14 hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned short [domain](#)
- unsigned char [bus](#)
- unsigned char [dev](#)
- unsigned char [func](#)
- unsigned short [class\\_id](#)
- unsigned short [vendor\\_id](#)
- unsigned short [device\\_id](#)
- unsigned short [subvendor\\_id](#)
- unsigned short [subdevice\\_id](#)
- unsigned char [revision](#)
- float [linkspeed](#)

### 20.14.1 Detailed Description

PCI Device specific Object Attributes.

### 20.14.2 Field Documentation

20.14.2.1 unsigned char `hwloc_obj_attr_u::hwloc_pcidev_attr_s::bus`

20.14.2.2 unsigned short `hwloc_obj_attr_u::hwloc_pcidev_attr_s::class_id`

20.14.2.3 unsigned char `hwloc_obj_attr_u::hwloc_pcidev_attr_s::dev`

20.14.2.4 unsigned short `hwloc_obj_attr_u::hwloc_pcidev_attr_s::device_id`

20.14.2.5 unsigned short `hwloc_obj_attr_u::hwloc_pcidev_attr_s::domain`

20.14.2.6 unsigned char `hwloc_obj_attr_u::hwloc_pcidev_attr_s::func`

20.14.2.7 float `hwloc_obj_attr_u::hwloc_pcidev_attr_s::linkspeed`

20.14.2.8 unsigned char `hwloc_obj_attr_u::hwloc_pcidev_attr_s::revision`

20.14.2.9 unsigned short `hwloc_obj_attr_u::hwloc_pcidev_attr_s::subdevice_id`

20.14.2.10 unsigned short `hwloc_obj_attr_u::hwloc_pcidev_attr_s::subvendor_id`

## 20.14.2.11 unsigned short hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s::vendor\_id

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.15 hwloc\_topology\_cpupbind\_support Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned char [set\\_thisproc\\_cpupbind](#)
- unsigned char [get\\_thisproc\\_cpupbind](#)
- unsigned char [set\\_proc\\_cpupbind](#)
- unsigned char [get\\_proc\\_cpupbind](#)
- unsigned char [set\\_thisthread\\_cpupbind](#)
- unsigned char [get\\_thisthread\\_cpupbind](#)
- unsigned char [set\\_thread\\_cpupbind](#)
- unsigned char [get\\_thread\\_cpupbind](#)
- unsigned char [get\\_thisproc\\_last\\_cpu\\_location](#)
- unsigned char [get\\_proc\\_last\\_cpu\\_location](#)
- unsigned char [get\\_thisthread\\_last\\_cpu\\_location](#)

### 20.15.1 Detailed Description

Flags describing actual PU binding support for this topology.

A flag may be set even if the feature isn't supported in all cases (e.g. binding to random sets of non-contiguous objects).

### 20.15.2 Field Documentation

#### 20.15.2.1 unsigned char hwloc\_topology\_cpupbind\_support::get\_proc\_cpupbind

Getting the binding of a whole given process is supported.

#### 20.15.2.2 unsigned char hwloc\_topology\_cpupbind\_support::get\_proc\_last\_cpu\_location

Getting the last processors where a whole process ran is supported

**20.15.2.3 unsigned char hwloc\_topology\_cpupbind\_support::get\_thisproc\_cpupbind**

Getting the binding of the whole current process is supported.

**20.15.2.4 unsigned char hwloc\_topology\_cpupbind\_support::get\_thisproc\_last\_cpu\_location**

Getting the last processors where the whole current process ran is supported

**20.15.2.5 unsigned char hwloc\_topology\_cpupbind\_support::get\_thisthread\_cpupbind**

Getting the binding of the current thread only is supported.

**20.15.2.6 unsigned char hwloc\_topology\_cpupbind\_support::get\_thisthread\_last\_cpu\_location**

Getting the last processors where the current thread ran is supported

**20.15.2.7 unsigned char hwloc\_topology\_cpupbind\_support::get\_thread\_cpupbind**

Getting the binding of a given thread only is supported.

**20.15.2.8 unsigned char hwloc\_topology\_cpupbind\_support::set\_proc\_cpupbind**

Binding a whole given process is supported.

**20.15.2.9 unsigned char hwloc\_topology\_cpupbind\_support::set\_thisproc\_cpupbind**

Binding the whole current process is supported.

**20.15.2.10 unsigned char hwloc\_topology\_cpupbind\_support::set\_thisthread\_cpupbind**

Binding the current thread only is supported.

**20.15.2.11 unsigned char hwloc\_topology\_cpupbind\_support::set\_thread\_cpupbind**

Binding a given thread only is supported.

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.16 hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_generic\_s Struct Reference 285

### 20.16 hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_generic\_s - Struct Reference

```
#include <diff.h>
```

#### Data Fields

- [hwloc\\_topology\\_diff\\_type\\_t](#) type
- union [hwloc\\_topology\\_diff\\_u](#) \* next

#### 20.16.1 Field Documentation

20.16.1.1 [union hwloc\\_topology\\_diff\\_u](#)\* [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_generic\\_s::next](#)

20.16.1.2 [hwloc\\_topology\\_diff\\_type\\_t](#) [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_generic\\_s::type](#)

The documentation for this struct was generated from the following file:

- [diff.h](#)

### 20.17 hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_generic\_s Struct Reference

```
#include <diff.h>
```

#### Data Fields

- [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#) type

#### 20.17.1 Field Documentation

20.17.1.1 [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_generic\\_s::type](#)

The documentation for this struct was generated from the following file:

- [diff.h](#)

## 20.18 hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s Struct Reference

```
#include <diff.h>
```

### Data Fields

- [hwloc\\_topology\\_diff\\_type\\_t](#) type
- union [hwloc\\_topology\\_diff\\_u](#) \* next
- unsigned [obj\\_depth](#)
- unsigned [obj\\_index](#)
- union [hwloc\\_topology\\_diff\\_obj\\_attr\\_u](#) diff

### 20.18.1 Field Documentation

20.18.1.1 union [hwloc\\_topology\\_diff\\_obj\\_attr\\_u](#) [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s::diff](#)

20.18.1.2 union [hwloc\\_topology\\_diff\\_u](#)\* [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s::next](#)

20.18.1.3 unsigned [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s::obj\\_depth](#)

20.18.1.4 unsigned [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s::obj\\_index](#)

20.18.1.5 [hwloc\\_topology\\_diff\\_type\\_t](#) [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s::type](#)

The documentation for this struct was generated from the following file:

- [diff.h](#)

## 20.19 hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_string\_s Struct Reference

```
#include <diff.h>
```

### Data Fields

- [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#) type
- char \* [name](#)

- char \* [oldvalue](#)
- char \* [newvalue](#)

### 20.19.1 Detailed Description

String attribute modification with an optional name.

### 20.19.2 Field Documentation

20.19.2.1 char\* [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s::name](#)

20.19.2.2 char\* [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s::newvalue](#)

20.19.2.3 char\* [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s::oldvalue](#)

20.19.2.4 [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s::type](#)

The documentation for this struct was generated from the following file:

- [diff.h](#)

## 20.20 hwloc\_topology\_diff\_obj\_attr\_u Union Reference

```
#include <diff.h>
```

### Data Structures

- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_generic\\_s](#)
- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s](#)  
*String attribute modification with an optional name.*
- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_uint64\\_s](#)  
*Integer attribute modification with an optional index.*

### Data Fields

- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_generic\\_s](#) [generic](#)
- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_uint64\\_s](#) [uint64](#)

- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s](#)  
string

### 20.20.1 Detailed Description

One object attribute difference.

### 20.20.2 Field Documentation

20.20.2.1 struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_generic\\_s](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::generic](#)

20.20.2.2 struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_string\\_s](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::string](#)

20.20.2.3 struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_uint64\\_s](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::uint64](#)

The documentation for this union was generated from the following file:

- [diff.h](#)

## 20.21 [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_uint64\\_s](#) Struct Reference

```
#include <diff.h>
```

### Data Fields

- [hwloc\\_topology\\_diff\\_obj\\_attr\\_type\\_t](#) type
- [hwloc\\_uint64\\_t](#) [index](#)
- [hwloc\\_uint64\\_t](#) [oldvalue](#)
- [hwloc\\_uint64\\_t](#) [newvalue](#)

### 20.21.1 Detailed Description

Integer attribute modification with an optional index.

### 20.21.2 Field Documentation

20.21.2.1 [hwloc\\_uint64\\_t](#) [hwloc\\_topology\\_diff\\_obj\\_attr\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_uint64\\_s::index](#)



## 20.22 hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s Struct Reference

239

20.21.2.2 hwloc\_uint64\_t hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s::newvalue

20.21.2.3 hwloc\_uint64\_t hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s::oldvalue

20.21.2.4 hwloc\_topology\_diff\_obj\_attr\_type\_t hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s::type

The documentation for this struct was generated from the following file:

- [diff.h](#)

## 20.22 hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s - Struct Reference

```
#include <diff.h>
```

### Data Fields

- [hwloc\\_topology\\_diff\\_type\\_t](#) type
- union [hwloc\\_topology\\_diff\\_u](#) \* next
- unsigned [obj\\_depth](#)
- unsigned [obj\\_index](#)

### 20.22.1 Field Documentation

20.22.1.1 union [hwloc\\_topology\\_diff\\_u](#)\* [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s::next](#)

20.22.1.2 unsigned [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s::obj\\_depth](#)

20.22.1.3 unsigned [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s::obj\\_index](#)

20.22.1.4 [hwloc\\_topology\\_diff\\_type\\_t](#) [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s::type](#)

The documentation for this struct was generated from the following file:

- [diff.h](#)

## 20.23 hwloc\_topology\_diff\_u Union Reference

```
#include <diff.h>
```

### Data Structures

- struct [hwloc\\_topology\\_diff\\_generic\\_s](#)
- struct [hwloc\\_topology\\_diff\\_obj\\_attr\\_s](#)
- struct [hwloc\\_topology\\_diff\\_too\\_complex\\_s](#)

### Data Fields

- struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_generic\\_s](#) generic
- struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s](#) obj\_attr
- struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s](#) too\_complex

#### 20.23.1 Detailed Description

One element of a difference list between two topologies.

#### 20.23.2 Field Documentation

20.23.2.1 struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_generic\\_s](#)  
[hwloc\\_topology\\_diff\\_u::generic](#)

20.23.2.2 struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_obj\\_attr\\_s](#)  
[hwloc\\_topology\\_diff\\_u::obj\\_attr](#)

20.23.2.3 struct [hwloc\\_topology\\_diff\\_u::hwloc\\_topology\\_diff\\_too\\_complex\\_s](#)  
[hwloc\\_topology\\_diff\\_u::too\\_complex](#)

The documentation for this union was generated from the following file:

- [diff.h](#)

## 20.24 hwloc\_topology\_discovery\_support Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned char [pu](#)

### 20.24.1 Detailed Description

Flags describing actual discovery support for this topology.

### 20.24.2 Field Documentation

#### 20.24.2.1 unsigned char hwloc\_topology\_discovery\_support::pu

Detecting the number of PU objects is supported.

The documentation for this struct was generated from the following file:

- hwloc.h

## 20.25 hwloc\_topology\_membind\_support Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- unsigned char [set\\_thisproc\\_membind](#)
- unsigned char [get\\_thisproc\\_membind](#)
- unsigned char [set\\_proc\\_membind](#)
- unsigned char [get\\_proc\\_membind](#)
- unsigned char [set\\_thisthread\\_membind](#)
- unsigned char [get\\_thisthread\\_membind](#)
- unsigned char [set\\_area\\_membind](#)
- unsigned char [get\\_area\\_membind](#)
- unsigned char [alloc\\_membind](#)
- unsigned char [firsttouch\\_membind](#)
- unsigned char [bind\\_membind](#)
- unsigned char [interleave\\_membind](#)
- unsigned char [replicate\\_membind](#)
- unsigned char [nexttouch\\_membind](#)
- unsigned char [migrate\\_membind](#)
- unsigned char [get\\_area\\_memlocation](#)

### 20.25.1 Detailed Description

Flags describing actual memory binding support for this topology.

A flag may be set even if the feature isn't supported in all cases (e.g. binding to random sets of non-contiguous objects).

## 20.25.2 Field Documentation

### 20.25.2.1 unsigned char hwloc\_topology\_mbind\_support::alloc\_mbind

Allocating a bound memory area is supported.

### 20.25.2.2 unsigned char hwloc\_topology\_mbind\_support::bind\_mbind

Bind policy is supported.

### 20.25.2.3 unsigned char hwloc\_topology\_mbind\_support::firsttouch\_mbind

First-touch policy is supported.

### 20.25.2.4 unsigned char hwloc\_topology\_mbind\_support::get\_area\_mbind

Getting the binding of a given memory area is supported.

### 20.25.2.5 unsigned char hwloc\_topology\_mbind\_support::get\_area\_mlocation

Getting the last NUMA nodes where a memory area was allocated is supported

### 20.25.2.6 unsigned char hwloc\_topology\_mbind\_support::get\_proc\_mbind

Getting the binding of a whole given process is supported.

### 20.25.2.7 unsigned char hwloc\_topology\_mbind\_support::get\_thisproc\_mbind

Getting the binding of the whole current process is supported.

### 20.25.2.8 unsigned char hwloc\_topology\_mbind\_support::get\_thisthread\_mbind

Getting the binding of the current thread only is supported.

### 20.25.2.9 unsigned char hwloc\_topology\_mbind\_support::interleave\_mbind

Interleave policy is supported.

20.25.2.10 unsigned char `hwloc_topology_membed_support::migrate_membed`

Migration flags is supported.

20.25.2.11 unsigned char `hwloc_topology_membed_support::nexttouch_membed`

Next-touch migration policy is supported.

20.25.2.12 unsigned char `hwloc_topology_membed_support::replicate_membed`

Replication policy is supported.

20.25.2.13 unsigned char `hwloc_topology_membed_support::set_area_membed`

Binding a given memory area is supported.

20.25.2.14 unsigned char `hwloc_topology_membed_support::set_proc_membed`

Binding a whole given process is supported.

20.25.2.15 unsigned char `hwloc_topology_membed_support::set_thisproc_membed`

Binding the whole current process is supported.

20.25.2.16 unsigned char `hwloc_topology_membed_support::set_thisthread_membed`

Binding the current thread only is supported.

The documentation for this struct was generated from the following file:

- `hwloc.h`

## 20.26 hwloc\_topology\_support Struct Reference

```
#include <hwloc.h>
```

### Data Fields

- struct [hwloc\\_topology\\_discovery\\_support](#) \* `discovery`
- struct [hwloc\\_topology\\_cpubind\\_support](#) \* `cpubind`
- struct [hwloc\\_topology\\_membed\\_support](#) \* `membed`

### 20.26.1 Detailed Description

Set of flags describing actual support for this topology.

This is retrieved with [hwloc\\_topology\\_get\\_support\(\)](#) and will be valid until the topology object is destroyed. Note: the values are correct only after discovery.

### 20.26.2 Field Documentation

20.26.2.1 `struct hwloc_topology_cpupbind_support*`  
`hwloc_topology_support::cpubind`

20.26.2.2 `struct hwloc_topology_discovery_support*`  
`hwloc_topology_support::discovery`

20.26.2.3 `struct hwloc_topology_membind_support*`  
`hwloc_topology_support::membind`

The documentation for this struct was generated from the following file:

- `hwloc.h`

# Index

- CPU binding
  - HWLOC\_CPUBIND\_NOMEMBIND, [114](#)
  - HWLOC\_CPUBIND\_PROCESS, [113](#)
  - HWLOC\_CPUBIND\_STRICT, [113](#)
  - HWLOC\_CPUBIND\_THREAD, [113](#)
- Components and Plugins: Discovery backends
  - HWLOC\_BACKEND\_FLAG\_NEED\_LEVELS, [205](#)
- Components and Plugins: Discovery components
  - HWLOC\_DISC\_COMPONENT\_TYPE\_CPU, [204](#)
  - HWLOC\_DISC\_COMPONENT\_TYPE\_GLOBAL, [204](#)
  - HWLOC\_DISC\_COMPONENT\_TYPE\_MISC, [204](#)
- Components and Plugins: Generic components
  - HWLOC\_COMPONENT\_TYPE\_DISC, [207](#)
  - HWLOC\_COMPONENT\_TYPE\_XML, [207](#)
- Distributing items over a topology
  - HWLOC\_DISTRIB\_FLAG\_REVERSE, [153](#)
- Exporting Topologies to Synthetic
  - HWLOC\_TOPOLOGY\_EXPORT\_SYNTHETIC\_FLAG\_NO\_ATTRS, [139](#)
  - HWLOC\_TOPOLOGY\_EXPORT\_SYNTHETIC\_FLAG\_NO\_EXTENDED\_TYPES, [139](#)
- HWLOC\_BACKEND\_FLAG\_NEED\_LEVELS
  - Components and Plugins: Discovery backends, [205](#)
- HWLOC\_COMPONENT\_TYPE\_DISC
  - Components and Plugins: Generic components, [207](#)
- HWLOC\_COMPONENT\_TYPE\_XML
  - Components and Plugins: Generic components, [207](#)
- HWLOC\_CPUBIND\_NOMEMBIND
  - CPU binding, [114](#)
- HWLOC\_CPUBIND\_PROCESS
  - CPU binding, [113](#)
- HWLOC\_CPUBIND\_STRICT
  - CPU binding, [113](#)
- HWLOC\_CPUBIND\_THREAD
  - CPU binding, [113](#)
- HWLOC\_DISC\_COMPONENT\_TYPE\_CPU
  - Components and Plugins: Discovery components, [204](#)
- HWLOC\_DISC\_COMPONENT\_TYPE\_GLOBAL
  - Components and Plugins: Discovery components, [204](#)
- HWLOC\_DISC\_COMPONENT\_TYPE\_MISC
  - Components and Plugins: Discovery components, [204](#)
- HWLOC\_DISTRIB\_FLAG\_REVERSE
  - Distributing items over a topology, [153](#)
- HWLOC\_MEMBIND\_BIND
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_BYNODESET
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_DEFAULT
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_FIRSTTOUCH
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_INTERLEAVE
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_MIGRATE
  - Memory binding, [119](#)
- HWLOC\_MEMBIND\_MIXED
  - Memory binding, [121](#)
- HWLOC\_MEMBIND\_NEXTTOUCH

- Memory binding, [121](#)
- HWLOC\_MEMBIND\_NOCPUBIND
  - Memory binding, [119](#)
- HWLOC\_MEMBIND\_PROCESS
  - Memory binding, [119](#)
- HWLOC\_MEMBIND\_REPLICATE
  - Memory binding, [120](#)
- HWLOC\_MEMBIND\_STRICT
  - Memory binding, [119](#)
- HWLOC\_MEMBIND\_THREAD
  - Memory binding, [119](#)
- HWLOC\_OBJ\_BRIDGE
  - Object Types, [89](#)
- HWLOC\_OBJ\_BRIDGE\_HOST
  - Object Types, [88](#)
- HWLOC\_OBJ\_BRIDGE\_PCI
  - Object Types, [88](#)
- HWLOC\_OBJ\_CACHE
  - Object Types, [89](#)
- HWLOC\_OBJ\_CACHE\_DATA
  - Object Types, [88](#)
- HWLOC\_OBJ\_CACHE\_INSTRUCTION
  - Object Types, [88](#)
- HWLOC\_OBJ\_CACHE\_UNIFIED
  - Object Types, [88](#)
- HWLOC\_OBJ\_CORE
  - Object Types, [89](#)
- HWLOC\_OBJ\_GROUP
  - Object Types, [89](#)
- HWLOC\_OBJ\_MACHINE
  - Object Types, [89](#)
- HWLOC\_OBJ\_MISC
  - Object Types, [89](#)
- HWLOC\_OBJ\_NUMANODE
  - Object Types, [89](#)
- HWLOC\_OBJ\_OSDEV\_BLOCK
  - Object Types, [88](#)
- HWLOC\_OBJ\_OSDEV\_COPROC
  - Object Types, [89](#)
- HWLOC\_OBJ\_OSDEV\_DMA
  - Object Types, [88](#)
- HWLOC\_OBJ\_OSDEV\_GPU
  - Object Types, [88](#)
- HWLOC\_OBJ\_OSDEV\_NETWORK
  - Object Types, [88](#)
- HWLOC\_OBJ\_OSDEV\_OPENFABRICS
  - Object Types, [88](#)
- HWLOC\_OBJ\_OS\_DEVICE
  - Object Types, [90](#)
- HWLOC\_OBJ\_PACKAGE
  - Object Types, [89](#)
- HWLOC\_OBJ\_PCI\_DEVICE
  - Object Types, [90](#)
- HWLOC\_OBJ\_PU
  - Object Types, [89](#)
- HWLOC\_OBJ\_SYSTEM
  - Object Types, [89](#)
- HWLOC\_OBJ\_TYPE\_MAX
  - Object Types, [90](#)
- HWLOC\_RESTRICT\_FLAG\_ADAPT\_DISTANCES
  - Modifying a loaded Topology, [130](#)
- HWLOC\_RESTRICT\_FLAG\_ADAPT\_IO
  - Modifying a loaded Topology, [130](#)
- HWLOC\_RESTRICT\_FLAG\_ADAPT\_MISC
  - Modifying a loaded Topology, [130](#)
- HWLOC\_TOPOLOGY\_DIFF\_APPLY\_REVERSE
  - Topology differences, [199](#)
- HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR
  - Topology differences, [200](#)
- HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_INFO
  - Topology differences, [200](#)
- HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_NAME
  - Topology differences, [200](#)
- HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_SIZE
  - Topology differences, [200](#)
- HWLOC\_TOPOLOGY\_DIFF\_TOO\_COMPLEX
  - Topology differences, [200](#)
- HWLOC\_TOPOLOGY\_EXPORT\_SYNETHETIC\_FLAG\_NO\_ATTRS
  - Exporting Topologies to Synthetic, [139](#)
- HWLOC\_TOPOLOGY\_EXPORT\_SYNETHETIC\_FLAG\_NO\_EXTENDED\_TYPES
  - Exporting Topologies to Synthetic, [139](#)
- HWLOC\_TOPOLOGY\_FLAG\_ICACHES
  - Topology Detection Configuration and Query, [97](#)
- HWLOC\_TOPOLOGY\_FLAG\_IO\_BRIDGES
  - Topology Detection Configuration and Query, [97](#)



- HWLOC\_TOPOLOGY\_FLAG\_IO\_DEVICES
  - Topology Detection Configuration and Query, [97](#)
- HWLOC\_TOPOLOGY\_FLAG\_IS\_THISSYSTEM
  - Topology Detection Configuration and Query, [96](#)
- HWLOC\_TOPOLOGY\_FLAG\_THISSYSTEM\_ALLOWED\_RESOURCES
  - Topology Detection Configuration and Query, [97](#)
- HWLOC\_TOPOLOGY\_FLAG\_WHOLE\_IO
  - Topology Detection Configuration and Query, [97](#)
- HWLOC\_TOPOLOGY\_FLAG\_WHOLE\_SYSTEM
  - Topology Detection Configuration and Query, [96](#)
- HWLOC\_TYPE\_DEPTH\_BRIDGE
  - Object levels, depths and types, [105](#)
- HWLOC\_TYPE\_DEPTH\_MULTIPLE
  - Object levels, depths and types, [104](#)
- HWLOC\_TYPE\_DEPTH\_OS\_DEVICE
  - Object levels, depths and types, [105](#)
- HWLOC\_TYPE\_DEPTH\_PCI\_DEVICE
  - Object levels, depths and types, [105](#)
- HWLOC\_TYPE\_DEPTH\_UNKNOWN
  - Object levels, depths and types, [104](#)
- HWLOC\_TYPE\_UNORDERED
  - Object Types, [88](#)
- Memory binding
  - HWLOC\_MEMBIND\_BIND, [120](#)
  - HWLOC\_MEMBIND\_BYNODESET, [120](#)
  - HWLOC\_MEMBIND\_DEFAULT, [120](#)
  - HWLOC\_MEMBIND\_FIRSTTOUCH, [120](#)
  - HWLOC\_MEMBIND\_INTERLEAVE, [120](#)
  - HWLOC\_MEMBIND\_MIGRATE, [119](#)
  - HWLOC\_MEMBIND\_MIXED, [121](#)
  - HWLOC\_MEMBIND\_NEXTTOUCH, [121](#)
  - HWLOC\_MEMBIND\_NOCPUBIND, [119](#)
  - HWLOC\_MEMBIND\_PROCESS, [119](#)
  - HWLOC\_MEMBIND\_REPLICATE, [120](#)
  - HWLOC\_MEMBIND\_STRICT, [119](#)
  - HWLOC\_MEMBIND\_THREAD, [119](#)
- Modifying a loaded Topology
  - HWLOC\_RESTRICT\_FLAG\_ADAPT\_DISTANCES, [130](#)
  - HWLOC\_RESTRICT\_FLAG\_ADAPT\_IO, [130](#)
  - HWLOC\_RESTRICT\_FLAG\_ADAPT\_MISC, [130](#)
- Object Types
  - HWLOC\_OBJ\_BRIDGE, [89](#)
  - HWLOC\_OBJ\_BRIDGE\_HOST, [88](#)
  - HWLOC\_OBJ\_BRIDGE\_PCI, [88](#)
  - HWLOC\_OBJ\_CACHE, [89](#)
  - HWLOC\_OBJ\_CACHE\_DATA, [88](#)
  - HWLOC\_OBJ\_CACHE\_INSTRUCTION, [88](#)
  - HWLOC\_OBJ\_CACHE\_UNIFIED, [88](#)
  - HWLOC\_OBJ\_CORE, [89](#)
  - HWLOC\_OBJ\_GROUP, [89](#)
  - HWLOC\_OBJ\_MACHINE, [89](#)
  - HWLOC\_OBJ\_MISC, [89](#)
  - HWLOC\_OBJ\_NUMANODE, [89](#)
  - HWLOC\_OBJ\_OSDEV\_BLOCK, [88](#)
  - HWLOC\_OBJ\_OSDEV\_COPROC, [89](#)
  - HWLOC\_OBJ\_OSDEV\_DMA, [88](#)
  - HWLOC\_OBJ\_OSDEV\_GPU, [88](#)
  - HWLOC\_OBJ\_OSDEV\_NETWORK, [88](#)
  - HWLOC\_OBJ\_OSDEV\_OPENFABRICS, [88](#)
  - HWLOC\_OBJ\_OS\_DEVICE, [90](#)
  - HWLOC\_OBJ\_PACKAGE, [89](#)
  - HWLOC\_OBJ\_PCI\_DEVICE, [90](#)
  - HWLOC\_OBJ\_PU, [89](#)
  - HWLOC\_OBJ\_SYSTEM, [89](#)
  - HWLOC\_OBJ\_TYPE\_MAX, [90](#)
  - HWLOC\_TYPE\_UNORDERED, [88](#)
- Object levels, depths and types
  - HWLOC\_TYPE\_DEPTH\_BRIDGE, [105](#)
  - HWLOC\_TYPE\_DEPTH\_MULTIPLE, [104](#)
  - HWLOC\_TYPE\_DEPTH\_OS\_DEVICE, [105](#)

- HWLOC\_TYPE\_DEPTH\_PCI\_DEVICE, 105
- HWLOC\_TYPE\_DEPTH\_UNKNOWN, 104
- Topology Detection Configuration and Query
  - HWLOC\_TOPOLOGY\_FLAG\_ICACHES, 97
  - HWLOC\_TOPOLOGY\_FLAG\_IO\_BRIDGES, 97
  - HWLOC\_TOPOLOGY\_FLAG\_IO\_DEVICES, 97
  - HWLOC\_TOPOLOGY\_FLAG\_IS\_THISSYSTEM, 96
  - HWLOC\_TOPOLOGY\_FLAG\_THISSYSTEM\_ALLOWED\_RESOURCES, 97
  - HWLOC\_TOPOLOGY\_FLAG\_WHOLE\_IO, 97
  - HWLOC\_TOPOLOGY\_FLAG\_WHOLE\_SYSTEM, 96
- Topology differences
  - HWLOC\_TOPOLOGY\_DIFF\_APPLY\_REVERSE, 199
  - HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR, 200
  - HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_INFO, 200
  - HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_NAME, 200
  - HWLOC\_TOPOLOGY\_DIFF\_OBJ\_ATTR\_SIZE, 200
  - HWLOC\_TOPOLOGY\_DIFF\_TOO\_COMPLEX, 200
- API version, 83
  - hwloc\_get\_api\_version, 83
- Building Custom Topologies, 133
  - hwloc\_custom\_insert\_group\_object\_by\_parent, 133
  - hwloc\_custom\_insert\_topology, 133
- CPU and node sets of entire topologies, 155
  - hwloc\_topology\_get\_allowed\_cpuset, 155
  - hwloc\_topology\_get\_allowed\_nodeset, 155
  - hwloc\_topology\_get\_complete\_cpuset, 156
  - hwloc\_topology\_get\_complete\_nodeset, 156
  - hwloc\_topology\_get\_online\_cpuset, 156
  - hwloc\_topology\_get\_topology\_cpuset, 156
  - hwloc\_topology\_get\_topology\_nodeset, 157
- CPU binding, 112
  - hwloc\_cpupbind\_flags\_t, 113
  - hwloc\_get\_cpupbind, 114
  - hwloc\_get\_last\_cpu\_location, 114
  - hwloc\_get\_proc\_cpupbind, 114
  - hwloc\_get\_proc\_last\_cpu\_location, 115
  - hwloc\_get\_thread\_cpupbind, 115
  - hwloc\_set\_cpupbind, 115
  - hwloc\_set\_proc\_cpupbind, 115
  - hwloc\_set\_thread\_cpupbind, 116
- Components and Plugins: Core functions to be used by components, 208
  - hwloc\_\_insert\_object\_by\_cpuset, 208
  - hwloc\_alloc\_setup\_object, 208
  - hwloc\_fill\_object\_sets, 208
  - hwloc\_hide\_errors, 209
  - hwloc\_insert\_object\_by\_cpuset, 209
  - hwloc\_insert\_object\_by\_parent, 209
  - hwloc\_plugin\_check\_namespace, 209
  - hwloc\_report\_error\_t, 208
  - hwloc\_report\_os\_error, 210
- Components and Plugins: Discovery backends, 205
  - hwloc\_backend\_alloc, 205
  - hwloc\_backend\_enable, 205
  - hwloc\_backend\_flag\_e, 205
  - hwloc\_backends\_get\_obj\_cpuset, 205
  - hwloc\_backends\_notify\_new\_object, 206
- Components and Plugins: Discovery components, 204
  - hwloc\_disc\_component\_type\_e, 204
  - hwloc\_disc\_component\_type\_t, 204
- Components and Plugins: Generic components, 207
  - hwloc\_component\_type\_e, 207
  - hwloc\_component\_type\_t, 207
- Components and Plugins: PCI functions to be used by components, 211
  - hwloc\_insert\_pci\_device\_list, 211

- hwloc\_pci\_find\_cap, 211
- hwloc\_pci\_find\_linkspeed, 211
- hwloc\_pci\_prepare\_bridge, 211
- Consulting and Adding Key-Value Info Attributes, 111
  - hwloc\_obj\_add\_info, 111
  - hwloc\_obj\_get\_info\_by\_name, 111
- Converting between CPU sets and node sets, 158
  - hwloc\_cpuset\_from\_nodese, 158
  - hwloc\_cpuset\_from\_nodese\_strict, 158
  - hwloc\_cpuset\_to\_nodese, 158
  - hwloc\_cpuset\_to\_nodese\_strict, 159
- Converting between Object Types, Sets and Attributes, and Strings, 108
  - hwloc\_obj\_attr\_snprintf, 108
  - hwloc\_obj\_cpuset\_snprintf, 108
  - hwloc\_obj\_type\_snprintf, 108
  - hwloc\_obj\_type\_sscanf, 109
  - hwloc\_obj\_type\_string, 109
- Distributing items over a topology, 153
  - hwloc\_distrib, 153
  - hwloc\_distrib\_flags\_e, 153
- Exporting Topologies to Synthetic, 139
  - hwloc\_topology\_export\_synthetic, 139
  - hwloc\_topology\_export\_synthetic\_flags\_e, 139
- Exporting Topologies to XML, 135
  - hwloc\_export\_obj\_userdata, 135
  - hwloc\_export\_obj\_userdata\_base64, 135
  - hwloc\_free\_xmlbuffer, 136
  - hwloc\_topology\_export\_xml, 136
  - hwloc\_topology\_export\_xmlbuffer, 136
  - hwloc\_topology\_set\_userdata\_export\_callback, 137
  - hwloc\_topology\_set\_userdata\_import\_callback, 137
- Finding I/O objects, 163
  - hwloc\_bridge\_covers\_pcibus, 163
  - hwloc\_get\_hostbridge\_by\_pcibus, 163
  - hwloc\_get\_next\_bridge, 163
  - hwloc\_get\_next\_osdev, 163
  - hwloc\_get\_next\_pcidev, 164
  - hwloc\_get\_non\_io\_ancestor\_obj, 164
  - hwloc\_get\_pcidev\_by\_busid, 164
  - hwloc\_get\_pcidev\_by\_busidstring, 164
- Finding Objects covering at least CPU set, 145
  - hwloc\_get\_child\_covering\_cpuset, 145
  - hwloc\_get\_next\_obj\_covering\_cpuset\_by\_depth, 145
  - hwloc\_get\_next\_obj\_covering\_cpuset\_by\_type, 145
  - hwloc\_get\_obj\_covering\_cpuset, 146
- Finding Objects inside a CPU set, 141
  - hwloc\_get\_first\_largest\_obj\_inside\_cpuset, 141
  - hwloc\_get\_largest\_objs\_inside\_cpuset, 141
  - hwloc\_get\_nbobjs\_inside\_cpuset\_by\_depth, 142
  - hwloc\_get\_nbobjs\_inside\_cpuset\_by\_type, 142
  - hwloc\_get\_next\_obj\_inside\_cpuset\_by\_depth, 142
  - hwloc\_get\_next\_obj\_inside\_cpuset\_by\_type, 143
  - hwloc\_get\_obj\_index\_inside\_cpuset, 143
  - hwloc\_get\_obj\_inside\_cpuset\_by\_depth, 143
  - hwloc\_get\_obj\_inside\_cpuset\_by\_type, 144
- Finding objects, miscellaneous helpers, 151
  - hwloc\_get\_closest\_objs, 151
  - hwloc\_get\_numanode\_obj\_by\_os\_index, 151
  - hwloc\_get\_obj\_below\_array\_by\_type, 151
  - hwloc\_get\_obj\_below\_by\_type, 152
  - hwloc\_get\_pu\_obj\_by\_os\_index, 152
- HWLOC\_API\_VERSION
  - API version, 83
- Interoperability with Intel Xeon Phi (MIC), 194
  - hwloc\_intel\_mic\_get\_device\_cpuset, 194
  - hwloc\_intel\_mic\_get\_device\_osdev\_by\_index, 194

- Interoperability with Linux libnuma bit-  
mask, [181](#)
  - [hwloc\\_cpuset\\_from\\_linux\\_libnuma\\_-  
bitmask](#), [181](#)
  - [hwloc\\_cpuset\\_to\\_linux\\_libnuma\\_-  
bitmask](#), [181](#)
  - [hwloc\\_nodeset\\_from\\_linux\\_libnuma-  
bitmask](#), [182](#)
  - [hwloc\\_nodeset\\_to\\_linux\\_libnuma\\_-  
bitmask](#), [182](#)
- Interoperability with Linux libnuma un-  
signed long masks, [179](#)
  - [hwloc\\_cpuset\\_from\\_linux\\_libnuma\\_-  
ulongs](#), [179](#)
  - [hwloc\\_cpuset\\_to\\_linux\\_libnuma\\_-  
ulongs](#), [179](#)
  - [hwloc\\_nodeset\\_from\\_linux\\_libnuma-  
ulongs](#), [180](#)
  - [hwloc\\_nodeset\\_to\\_linux\\_libnuma\\_-  
ulongs](#), [180](#)
- Interoperability with Myrinet Express, [197](#)
  - [hwloc\\_mx\\_board\\_get\\_device\\_-  
cpuset](#), [197](#)
  - [hwloc\\_mx\\_endpoint\\_get\\_device\\_-  
cpuset](#), [197](#)
- Interoperability with OpenCL, [184](#)
  - [hwloc\\_opencl\\_get\\_device\\_cpuset](#),  
[184](#)
  - [hwloc\\_opencl\\_get\\_device\\_osdev](#),  
[184](#)
  - [hwloc\\_opencl\\_get\\_device\\_osdev\\_-  
by\\_index](#), [185](#)
- Interoperability with OpenFabrics, [195](#)
  - [hwloc\\_ibv\\_get\\_device\\_cpuset](#), [195](#)
  - [hwloc\\_ibv\\_get\\_device\\_osdev](#), [195](#)
  - [hwloc\\_ibv\\_get\\_device\\_osdev\\_by\\_-  
name](#), [196](#)
- Interoperability with OpenGL displays, [192](#)
  - [hwloc\\_gl\\_get\\_display\\_by\\_osdev](#),  
[192](#)
  - [hwloc\\_gl\\_get\\_display\\_osdev\\_by\\_-  
name](#), [192](#)
  - [hwloc\\_gl\\_get\\_display\\_osdev\\_by\\_-  
port\\_device](#), [192](#)
- Interoperability with glibc sched affinity,  
[183](#)
  - [hwloc\\_cpuset\\_from\\_glibc\\_sched\\_-  
affinity](#), [183](#)
  - [hwloc\\_cpuset\\_to\\_glibc\\_sched\\_-  
affinity](#), [183](#)
- Interoperability with the CUDA Driver API,  
[186](#)
  - [hwloc\\_cuda\\_get\\_device\\_cpuset](#), [186](#)
  - [hwloc\\_cuda\\_get\\_device\\_osdev](#), [186](#)
  - [hwloc\\_cuda\\_get\\_device\\_osdev\\_by\\_-  
index](#), [187](#)
  - [hwloc\\_cuda\\_get\\_device\\_pci\\_ids](#),  
[187](#)
  - [hwloc\\_cuda\\_get\\_device\\_pcidev](#), [187](#)
- Interoperability with the CUDA Runtime A-  
PI, [188](#)
  - [hwloc\\_cudart\\_get\\_device\\_cpuset](#),  
[188](#)
  - [hwloc\\_cudart\\_get\\_device\\_osdev\\_-  
by\\_index](#), [188](#)
  - [hwloc\\_cudart\\_get\\_device\\_pci\\_ids](#),  
[189](#)
  - [hwloc\\_cudart\\_get\\_device\\_pcidev](#),  
[189](#)
- Interoperability with the NVIDIA Manage-  
ment Library, [190](#)
  - [hwloc\\_nvml\\_get\\_device\\_cpuset](#), [190](#)
  - [hwloc\\_nvml\\_get\\_device\\_osdev](#), [190](#)
  - [hwloc\\_nvml\\_get\\_device\\_osdev\\_by\\_-  
index](#), [191](#)
- Linux-specific helpers, [177](#)
  - [hwloc\\_linux\\_get\\_tid\\_cpubind](#), [177](#)
  - [hwloc\\_linux\\_get\\_tid\\_last\\_cpu\\_-  
location](#), [177](#)
  - [hwloc\\_linux\\_parse\\_cpumap\\_file](#), [177](#)
  - [hwloc\\_linux\\_set\\_tid\\_cpubind](#), [177](#)
- Looking at Ancestor and Child Objects,  
[147](#)
  - [hwloc\\_get\\_ancestor\\_obj\\_by\\_depth](#),  
[147](#)
  - [hwloc\\_get\\_ancestor\\_obj\\_by\\_type](#),  
[147](#)
  - [hwloc\\_get\\_common\\_ancestor\\_obj](#),  
[147](#)
  - [hwloc\\_get\\_next\\_child](#), [147](#)
  - [hwloc\\_obj\\_is\\_in\\_subtree](#), [148](#)
- Looking at Cache Objects, [149](#)
  - [hwloc\\_get\\_cache\\_covering\\_cpuset](#),  
[149](#)
  - [hwloc\\_get\\_cache\\_type\\_depth](#), [149](#)
  - [hwloc\\_get\\_shared\\_cache\\_covering-  
\\_obj](#), [149](#)
- Manipulating Distances, [160](#)
  - [hwloc\\_get\\_distance\\_matrix\\_covering-  
\\_obj\\_by\\_depth](#), [160](#)

- hwloc\_get\_latency, 160
- hwloc\_get\_whole\_distance\_matrix\_-  
by\_depth, 160
- hwloc\_get\_whole\_distance\_matrix\_-  
by\_type, 161
- Memory binding, 117
  - hwloc\_alloc, 121
  - hwloc\_alloc\_membind, 121
  - hwloc\_alloc\_membind\_nodeset, 121
  - hwloc\_alloc\_membind\_policy, 122
  - hwloc\_alloc\_membind\_policy\_-  
nodeset, 122
  - hwloc\_free, 122
  - hwloc\_get\_area\_membind, 122
  - hwloc\_get\_area\_membind\_nodeset,  
123
  - hwloc\_get\_area\_memlocation, 123
  - hwloc\_get\_membind, 124
  - hwloc\_get\_membind\_nodeset, 124
  - hwloc\_get\_proc\_membind, 125
  - hwloc\_get\_proc\_membind\_nodeset,  
126
  - hwloc\_membind\_flags\_t, 119
  - hwloc\_membind\_policy\_t, 120
  - hwloc\_set\_area\_membind, 126
  - hwloc\_set\_area\_membind\_nodeset,  
127
  - hwloc\_set\_membind, 127
  - hwloc\_set\_membind\_nodeset, 127
  - hwloc\_set\_proc\_membind, 128
  - hwloc\_set\_proc\_membind\_nodeset,  
128
- Modifying a loaded Topology, 130
  - hwloc\_restrict\_flags\_e, 130
  - hwloc\_topology\_insert\_misc\_object\_-  
by\_cpuset, 130
  - hwloc\_topology\_insert\_misc\_object\_-  
by\_parent, 131
  - hwloc\_topology\_restrict, 131
- Object Sets (hwloc\_cpuset\_t and hwloc\_-  
nodeset\_t), 85
  - hwloc\_const\_cpuset\_t, 85
  - hwloc\_const\_nodeset\_t, 85
  - hwloc\_cpuset\_t, 85
  - hwloc\_nodeset\_t, 85
- Object Structure and Attributes, 91
  - hwloc\_obj\_t, 91
- Object Types, 87
  - hwloc\_compare\_types, 90
  - hwloc\_compare\_types\_e, 88
  - hwloc\_obj\_bridge\_type\_e, 88
  - hwloc\_obj\_bridge\_type\_t, 87
  - hwloc\_obj\_cache\_type\_e, 88
  - hwloc\_obj\_cache\_type\_t, 87
  - hwloc\_obj\_osdev\_type\_e, 88
  - hwloc\_obj\_osdev\_type\_t, 87
  - hwloc\_obj\_type\_t, 89
- Object levels, depths and types, 104
  - hwloc\_get\_depth\_type, 105
  - hwloc\_get\_nbobjs\_by\_depth, 105
  - hwloc\_get\_nbobjs\_by\_type, 105
  - hwloc\_get\_next\_obj\_by\_depth, 105
  - hwloc\_get\_next\_obj\_by\_type, 105
  - hwloc\_get\_obj\_by\_depth, 105
  - hwloc\_get\_obj\_by\_type, 106
  - hwloc\_get\_root\_obj, 106
  - hwloc\_get\_type\_depth, 106
  - hwloc\_get\_type\_depth\_e, 104
  - hwloc\_get\_type\_or\_above\_depth,  
106
  - hwloc\_get\_type\_or\_below\_depth,  
107
  - hwloc\_topology\_get\_depth, 107
- The bitmap API, 165
  - hwloc\_bitmap\_allbut, 168
  - hwloc\_bitmap\_alloc, 168
  - hwloc\_bitmap\_alloc\_full, 168
  - hwloc\_bitmap\_and, 168
  - hwloc\_bitmap\_andnot, 168
  - hwloc\_bitmap\_asprintf, 168
  - hwloc\_bitmap\_clr, 168
  - hwloc\_bitmap\_clr\_range, 168
  - hwloc\_bitmap\_compare, 169
  - hwloc\_bitmap\_compare\_first, 169
  - hwloc\_bitmap\_copy, 169
  - hwloc\_bitmap\_dup, 170
  - hwloc\_bitmap\_fill, 170
  - hwloc\_bitmap\_first, 170
  - hwloc\_bitmap\_foreach\_begin, 167
  - hwloc\_bitmap\_foreach\_end, 167
  - hwloc\_bitmap\_free, 170
  - hwloc\_bitmap\_from\_ith\_ulong, 170
  - hwloc\_bitmap\_from\_ulong, 170
  - hwloc\_bitmap\_intersects, 170
  - hwloc\_bitmap\_isequal, 171
  - hwloc\_bitmap\_isfull, 171
  - hwloc\_bitmap\_isincluded, 171
  - hwloc\_bitmap\_isset, 171
  - hwloc\_bitmap\_iszero, 171
  - hwloc\_bitmap\_last, 172

- hwloc\_bitmap\_list\_asprintf, 172
- hwloc\_bitmap\_list\_snprintf, 172
- hwloc\_bitmap\_list\_sscanf, 172
- hwloc\_bitmap\_next, 172
- hwloc\_bitmap\_not, 173
- hwloc\_bitmap\_only, 173
- hwloc\_bitmap\_or, 173
- hwloc\_bitmap\_set, 173
- hwloc\_bitmap\_set\_ith\_ulong, 173
- hwloc\_bitmap\_set\_range, 173
- hwloc\_bitmap\_singlify, 173
- hwloc\_bitmap\_snprintf, 174
- hwloc\_bitmap\_sscanf, 174
- hwloc\_bitmap\_t, 167
- hwloc\_bitmap\_taskset\_asprintf, 174
- hwloc\_bitmap\_taskset\_snprintf, 174
- hwloc\_bitmap\_taskset\_sscanf, 175
- hwloc\_bitmap\_to\_ith\_ulong, 175
- hwloc\_bitmap\_to\_ulong, 175
- hwloc\_bitmap\_weight, 175
- hwloc\_bitmap\_xor, 175
- hwloc\_bitmap\_zero, 175
- hwloc\_const\_bitmap\_t, 167
- Topology Creation and Destruction, 92
  - hwloc\_topology\_check, 92
  - hwloc\_topology\_destroy, 92
  - hwloc\_topology\_dup, 93
  - hwloc\_topology\_init, 93
  - hwloc\_topology\_load, 93
  - hwloc\_topology\_t, 92
- Topology Detection Configuration and Query, 95
  - hwloc\_topology\_flags\_e, 96
  - hwloc\_topology\_get\_flags, 97
  - hwloc\_topology\_get\_support, 98
  - hwloc\_topology\_get\_userdata, 98
  - hwloc\_topology\_ignore\_all\_keep\_-structure, 98
  - hwloc\_topology\_ignore\_type, 98
  - hwloc\_topology\_ignore\_type\_keep\_-structure, 98
  - hwloc\_topology\_is\_thissystem, 99
  - hwloc\_topology\_set\_custom, 99
  - hwloc\_topology\_set\_distance\_-matrix, 99
  - hwloc\_topology\_set\_flags, 100
  - hwloc\_topology\_set\_fsroot, 100
  - hwloc\_topology\_set\_pid, 100
  - hwloc\_topology\_set\_synthetic, 101
  - hwloc\_topology\_set\_userdata, 101
- hwloc\_topology\_set\_xml, 101
- hwloc\_topology\_set\_xmlbuffer, 102
- Topology differences, 198
  - hwloc\_topology\_diff\_apply, 200
  - hwloc\_topology\_diff\_apply\_flags\_e, 199
  - hwloc\_topology\_diff\_build, 200
  - hwloc\_topology\_diff\_destroy, 201
  - hwloc\_topology\_diff\_export\_xml, 201
  - hwloc\_topology\_diff\_export\_xmlbuffer, 202
  - hwloc\_topology\_diff\_load\_xml, 202
  - hwloc\_topology\_diff\_load\_xmlbuffer, 202
  - hwloc\_topology\_diff\_obj\_attr\_type\_-e, 199
  - hwloc\_topology\_diff\_obj\_attr\_type\_t, 199
  - hwloc\_topology\_diff\_t, 199
  - hwloc\_topology\_diff\_type\_e, 200
  - hwloc\_topology\_diff\_type\_t, 199
- abi
  - hwloc\_component, 217
- alloc\_membind
  - hwloc\_topology\_membind\_support, 242
- allowed\_cpuset
  - hwloc\_obj, 223
- allowed\_nodeset
  - hwloc\_obj, 223
- arity
  - hwloc\_obj, 223
- associativity
  - hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_-s, 216
- attr
  - hwloc\_obj, 223
- bind\_membind
  - hwloc\_topology\_membind\_support, 242
- bridge
  - hwloc\_obj\_attr\_u, 228
- bus
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_-attr\_s, 232
- cache

- hwloc\_obj\_attr\_u, [228](#)
- children
  - hwloc\_obj, [224](#)
- class\_id
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- complete\_cpuset
  - hwloc\_obj, [224](#)
- complete\_nodeset
  - hwloc\_obj, [224](#)
- count
  - hwloc\_obj\_memory\_s::hwloc\_obj\_memory\_page\_type\_s, [230](#)
- cpubind
  - hwloc\_topology\_support, [244](#)
- cpuset
  - hwloc\_obj, [224](#)
- data
  - hwloc\_component, [217](#)
- depth
  - hwloc\_obj, [225](#)
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
  - hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s, [216](#)
  - hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s, [222](#)
- dev
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- device\_id
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- diff
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, [236](#)
- disable
  - hwloc\_backend, [213](#)
- discover
  - hwloc\_backend, [214](#)
- discovery
  - hwloc\_topology\_support, [244](#)
- distances
  - hwloc\_obj, [225](#)
- distances\_count
  - hwloc\_obj, [225](#)
- domain
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- dont\_merge
  - hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s, [222](#)
- downstream
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
- excludes
  - hwloc\_disc\_component, [219](#)
- finalize
  - hwloc\_component, [217](#)
- first\_child
  - hwloc\_obj, [225](#)
- firsttouch\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- flags
  - hwloc\_backend, [214](#)
  - hwloc\_component, [218](#)
- func
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- generic
  - hwloc\_topology\_diff\_obj\_attr\_u, [238](#)
  - hwloc\_topology\_diff\_u, [240](#)
- get\_area\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- get\_area\_memlocation
  - hwloc\_topology\_membind\_support, [242](#)
- get\_obj\_cpuset
  - hwloc\_backend, [214](#)
- get\_proc\_cpuset
  - hwloc\_topology\_cpuset\_support, [233](#)
- get\_proc\_last\_cpu\_location
  - hwloc\_topology\_cpuset\_support, [233](#)
- get\_proc\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- get\_thisproc\_cpuset

- hwloc\_topology\_cpupbind\_support, [233](#)
- get\_thisproc\_last\_cpu\_location
  - hwloc\_topology\_cpupbind\_support, [234](#)
- get\_thisproc\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- get\_thisthread\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)
- get\_thisthread\_last\_cpu\_location
  - hwloc\_topology\_cpupbind\_support, [234](#)
- get\_thisthread\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- get\_thread\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)
- group
  - hwloc\_obj\_attr\_u, [228](#)
- hwloc\_\_insert\_object\_by\_cpuset
  - Components and Plugins: Core functions to be used by components, [208](#)
- hwloc\_alloc
  - Memory binding, [121](#)
- hwloc\_alloc\_membind
  - Memory binding, [121](#)
- hwloc\_alloc\_membind\_nodeset
  - Memory binding, [121](#)
- hwloc\_alloc\_membind\_policy
  - Memory binding, [122](#)
- hwloc\_alloc\_membind\_policy\_nodeset
  - Memory binding, [122](#)
- hwloc\_alloc\_setup\_object
  - Components and Plugins: Core functions to be used by components, [208](#)
- hwloc\_backend, [213](#)
  - disable, [213](#)
  - discover, [214](#)
  - flags, [214](#)
  - get\_obj\_cpuset, [214](#)
  - is\_custom, [214](#)
  - is\_thissystem, [214](#)
  - notify\_new\_object, [214](#)
  - private\_data, [214](#)
- hwloc\_backend\_alloc
  - Components and Plugins: Discovery backends, [205](#)
- hwloc\_backend\_enable
  - Components and Plugins: Discovery backends, [205](#)
- hwloc\_backend\_flag\_e
  - Components and Plugins: Discovery backends, [205](#)
- hwloc\_backends\_get\_obj\_cpuset
  - Components and Plugins: Discovery backends, [205](#)
- hwloc\_backends\_notify\_new\_object
  - Components and Plugins: Discovery backends, [206](#)
- hwloc\_bitmap\_allbut
  - The bitmap API, [168](#)
- hwloc\_bitmap\_alloc
  - The bitmap API, [168](#)
- hwloc\_bitmap\_alloc\_full
  - The bitmap API, [168](#)
- hwloc\_bitmap\_and
  - The bitmap API, [168](#)
- hwloc\_bitmap\_andnot
  - The bitmap API, [168](#)
- hwloc\_bitmap\_asprintf
  - The bitmap API, [168](#)
- hwloc\_bitmap\_clr
  - The bitmap API, [168](#)
- hwloc\_bitmap\_clr\_range
  - The bitmap API, [168](#)
- hwloc\_bitmap\_compare
  - The bitmap API, [169](#)
- hwloc\_bitmap\_compare\_first
  - The bitmap API, [169](#)
- hwloc\_bitmap\_copy
  - The bitmap API, [169](#)
- hwloc\_bitmap\_dup
  - The bitmap API, [170](#)
- hwloc\_bitmap\_fill
  - The bitmap API, [170](#)
- hwloc\_bitmap\_first
  - The bitmap API, [170](#)
- hwloc\_bitmap\_foreach\_begin
  - The bitmap API, [167](#)
- hwloc\_bitmap\_foreach\_end
  - The bitmap API, [167](#)
- hwloc\_bitmap\_free
  - The bitmap API, [170](#)
- hwloc\_bitmap\_from\_ith\_ulong



- The bitmap API, [170](#)
- `hwloc_bitmap_from_ulong`
  - The bitmap API, [170](#)
- `hwloc_bitmap_intersects`
  - The bitmap API, [170](#)
- `hwloc_bitmap_isequal`
  - The bitmap API, [171](#)
- `hwloc_bitmap_isfull`
  - The bitmap API, [171](#)
- `hwloc_bitmap_isincluded`
  - The bitmap API, [171](#)
- `hwloc_bitmap_isset`
  - The bitmap API, [171](#)
- `hwloc_bitmap_iszero`
  - The bitmap API, [171](#)
- `hwloc_bitmap_last`
  - The bitmap API, [172](#)
- `hwloc_bitmap_list_asprintf`
  - The bitmap API, [172](#)
- `hwloc_bitmap_list_sprintf`
  - The bitmap API, [172](#)
- `hwloc_bitmap_list_sscanf`
  - The bitmap API, [172](#)
- `hwloc_bitmap_next`
  - The bitmap API, [172](#)
- `hwloc_bitmap_not`
  - The bitmap API, [173](#)
- `hwloc_bitmap_only`
  - The bitmap API, [173](#)
- `hwloc_bitmap_or`
  - The bitmap API, [173](#)
- `hwloc_bitmap_set`
  - The bitmap API, [173](#)
- `hwloc_bitmap_set_ith_ulong`
  - The bitmap API, [173](#)
- `hwloc_bitmap_set_range`
  - The bitmap API, [173](#)
- `hwloc_bitmap_singlify`
  - The bitmap API, [173](#)
- `hwloc_bitmap_sprintf`
  - The bitmap API, [174](#)
- `hwloc_bitmap_sscanf`
  - The bitmap API, [174](#)
- `hwloc_bitmap_t`
  - The bitmap API, [167](#)
- `hwloc_bitmap_taskset_asprintf`
  - The bitmap API, [174](#)
- `hwloc_bitmap_taskset_sprintf`
  - The bitmap API, [174](#)
- `hwloc_bitmap_taskset_sscanf`
  - The bitmap API, [175](#)
- `hwloc_bitmap_to_ith_ulong`
  - The bitmap API, [175](#)
- `hwloc_bitmap_to_ulong`
  - The bitmap API, [175](#)
- `hwloc_bitmap_weight`
  - The bitmap API, [175](#)
- `hwloc_bitmap_xor`
  - The bitmap API, [175](#)
- `hwloc_bitmap_zero`
  - The bitmap API, [175](#)
- `hwloc_bridge_covers_pcibus`
  - Finding I/O objects, [163](#)
- `hwloc_compare_types`
  - Object Types, [90](#)
- `hwloc_compare_types_e`
  - Object Types, [88](#)
- `hwloc_component`, [217](#)
  - abi, [217](#)
  - data, [217](#)
  - finalize, [217](#)
  - flags, [218](#)
  - init, [218](#)
  - type, [218](#)
- `hwloc_component_type_e`
  - Components and Plugins: Generic components, [207](#)
- `hwloc_component_type_t`
  - Components and Plugins: Generic components, [207](#)
- `hwloc_const_bitmap_t`
  - The bitmap API, [167](#)
- `hwloc_const_cpuset_t`
  - Object Sets (`hwloc_cpuset_t` and `hwloc_nodeuset_t`), [85](#)
- `hwloc_const_nodeuset_t`
  - Object Sets (`hwloc_cpuset_t` and `hwloc_nodeuset_t`), [85](#)
- `hwloc_cpubind_flags_t`
  - CPU binding, [113](#)
- `hwloc_cpuset_from_glibc_sched_affinity`
  - Interoperability with glibc sched affinity, [183](#)
- `hwloc_cpuset_from_linux_libnuma_-bitmask`
  - Interoperability with Linux libnuma bitmask, [181](#)
- `hwloc_cpuset_from_linux_libnuma_-ulongs`
  - ulongs

- Interoperability with Linux libnuma unsigned long masks, [179](#)
- hwloc\_cpuset\_from\_nodese
  - Converting between CPU sets and node sets, [158](#)
- hwloc\_cpuset\_from\_nodese\_strict
  - Converting between CPU sets and node sets, [158](#)
- hwloc\_cpuset\_t
  - Object Sets (hwloc\_cpuset\_t and hwloc\_nodese\_t), [85](#)
- hwloc\_cpuset\_to\_glibc\_sched\_affinity
  - Interoperability with glibc sched affinity, [183](#)
- hwloc\_cpuset\_to\_linux\_libnuma\_bitmask
  - Interoperability with Linux libnuma bitmask, [181](#)
- hwloc\_cpuset\_to\_linux\_libnuma\_ulong
  - Interoperability with Linux libnuma unsigned long masks, [179](#)
- hwloc\_cpuset\_to\_nodese
  - Converting between CPU sets and node sets, [158](#)
- hwloc\_cpuset\_to\_nodese\_strict
  - Converting between CPU sets and node sets, [159](#)
- hwloc\_cuda\_get\_device\_cpuset
  - Interoperability with the CUDA Driver API, [186](#)
- hwloc\_cuda\_get\_device\_osdev
  - Interoperability with the CUDA Driver API, [186](#)
- hwloc\_cuda\_get\_device\_osdev\_by\_index
  - Interoperability with the CUDA Driver API, [187](#)
- hwloc\_cuda\_get\_device\_pci\_ids
  - Interoperability with the CUDA Driver API, [187](#)
- hwloc\_cuda\_get\_device\_pcidev
  - Interoperability with the CUDA Driver API, [187](#)
- hwloc\_cudart\_get\_device\_cpuset
  - Interoperability with the CUDA Runtime API, [188](#)
- hwloc\_cudart\_get\_device\_osdev\_by\_index
  - Interoperability with the CUDA Runtime API, [188](#)
- hwloc\_cudart\_get\_device\_pci\_ids
  - Interoperability with the CUDA Runtime API, [189](#)
- hwloc\_cudart\_get\_device\_pcidev
  - Interoperability with the CUDA Runtime API, [189](#)
- hwloc\_custom\_insert\_group\_object\_by\_parent
  - Building Custom Topologies, [133](#)
- hwloc\_custom\_insert\_topology
  - Building Custom Topologies, [133](#)
- hwloc\_disc\_component, [218](#)
  - excludes, [219](#)
  - instantiate, [219](#)
  - name, [219](#)
  - priority, [219](#)
  - type, [220](#)
- hwloc\_disc\_component\_type\_e
  - Components and Plugins: Discovery components, [204](#)
- hwloc\_disc\_component\_type\_t
  - Components and Plugins: Discovery components, [204](#)
- hwloc\_distances\_s, [220](#)
  - latency, [220](#)
  - latency\_base, [221](#)
  - latency\_max, [221](#)
  - nbobjs, [221](#)
  - relative\_depth, [221](#)
- hwloc\_distrib
  - Distributing items over a topology, [153](#)
- hwloc\_distrib\_flags\_e
  - Distributing items over a topology, [153](#)
- hwloc\_export\_obj\_userdata
  - Exporting Topologies to XML, [135](#)
- hwloc\_export\_obj\_userdata\_base64
  - Exporting Topologies to XML, [135](#)
- hwloc\_fill\_object\_sets
  - Components and Plugins: Core functions to be used by components, [208](#)
- hwloc\_free
  - Memory binding, [122](#)
- hwloc\_free\_xmlbuffer
  - Exporting Topologies to XML, [136](#)
- hwloc\_get\_ancestor\_obj\_by\_depth
  - Looking at Ancestor and Child Objects, [147](#)
- hwloc\_get\_ancestor\_obj\_by\_type

- Looking at Ancestor and Child Objects, [147](#)
- `hwloc_get_api_version`
  - API version, [83](#)
- `hwloc_get_area_membind`
  - Memory binding, [122](#)
- `hwloc_get_area_membind_nodese`
  - Memory binding, [123](#)
- `hwloc_get_area_memlocation`
  - Memory binding, [123](#)
- `hwloc_get_cache_covering_cpuset`
  - Looking at Cache Objects, [149](#)
- `hwloc_get_cache_type_depth`
  - Looking at Cache Objects, [149](#)
- `hwloc_get_child_covering_cpuset`
  - Finding Objects covering at least CPU set, [145](#)
- `hwloc_get_closest_objs`
  - Finding objects, miscellaneous helpers, [151](#)
- `hwloc_get_common_ancestor_obj`
  - Looking at Ancestor and Child Objects, [147](#)
- `hwloc_get_cpubind`
  - CPU binding, [114](#)
- `hwloc_get_depth_type`
  - Object levels, depths and types, [105](#)
- `hwloc_get_distance_matrix_covering_obj_by_depth`
  - Manipulating Distances, [160](#)
- `hwloc_get_first_largest_obj_inside_cpuset`
  - Finding Objects inside a CPU set, [141](#)
- `hwloc_get_hostbridge_by_pcibus`
  - Finding I/O objects, [163](#)
- `hwloc_get_largest_objs_inside_cpuset`
  - Finding Objects inside a CPU set, [141](#)
- `hwloc_get_last_cpu_location`
  - CPU binding, [114](#)
- `hwloc_get_latency`
  - Manipulating Distances, [160](#)
- `hwloc_get_membind`
  - Memory binding, [124](#)
- `hwloc_get_membind_nodese`
  - Memory binding, [124](#)
- `hwloc_get_nbobjs_by_depth`
  - Object levels, depths and types, [105](#)
- `hwloc_get_nbobjs_by_type`
  - Object levels, depths and types, [105](#)
- `hwloc_get_nbobjs_inside_cpuset_by_depth`
  - Finding Objects inside a CPU set, [142](#)
- `hwloc_get_nbobjs_inside_cpuset_by_type`
  - Finding Objects inside a CPU set, [142](#)
- `hwloc_get_next_bridge`
  - Finding I/O objects, [163](#)
- `hwloc_get_next_child`
  - Looking at Ancestor and Child Objects, [147](#)
- `hwloc_get_next_obj_by_depth`
  - Object levels, depths and types, [105](#)
- `hwloc_get_next_obj_by_type`
  - Object levels, depths and types, [105](#)
- `hwloc_get_next_obj_covering_cpuset_by_depth`
  - Finding Objects covering at least CPU set, [145](#)
- `hwloc_get_next_obj_covering_cpuset_by_type`
  - Finding Objects covering at least CPU set, [145](#)
- `hwloc_get_next_obj_inside_cpuset_by_depth`
  - Finding Objects inside a CPU set, [142](#)
- `hwloc_get_next_obj_inside_cpuset_by_type`
  - Finding Objects inside a CPU set, [143](#)
- `hwloc_get_next_osdev`
  - Finding I/O objects, [163](#)
- `hwloc_get_next_pcidev`
  - Finding I/O objects, [164](#)
- `hwloc_get_non_io_ancestor_obj`
  - Finding I/O objects, [164](#)
- `hwloc_get_numanode_obj_by_os_index`
  - Finding objects, miscellaneous helpers, [151](#)
- `hwloc_get_obj_below_array_by_type`
  - Finding objects, miscellaneous helpers, [151](#)
- `hwloc_get_obj_below_by_type`
  - Finding objects, miscellaneous helpers, [152](#)
- `hwloc_get_obj_by_depth`

- Object levels, depths and types, [105](#)
- `hwloc_get_obj_by_type`
  - Object levels, depths and types, [106](#)
- `hwloc_get_obj_covering_cpuset`
  - Finding Objects covering at least CPU set, [146](#)
- `hwloc_get_obj_index_inside_cpuset`
  - Finding Objects inside a CPU set, [143](#)
- `hwloc_get_obj_inside_cpuset_by_depth`
  - Finding Objects inside a CPU set, [143](#)
- `hwloc_get_obj_inside_cpuset_by_type`
  - Finding Objects inside a CPU set, [144](#)
- `hwloc_get_pcidev_by_busid`
  - Finding I/O objects, [164](#)
- `hwloc_get_pcidev_by_busidstring`
  - Finding I/O objects, [164](#)
- `hwloc_get_proc_cpupbind`
  - CPU binding, [114](#)
- `hwloc_get_proc_last_cpu_location`
  - CPU binding, [115](#)
- `hwloc_get_proc_membind`
  - Memory binding, [125](#)
- `hwloc_get_proc_membind_nodeset`
  - Memory binding, [126](#)
- `hwloc_get_pu_obj_by_os_index`
  - Finding objects, miscellaneous helpers, [152](#)
- `hwloc_get_root_obj`
  - Object levels, depths and types, [106](#)
- `hwloc_get_shared_cache_covering_obj`
  - Looking at Cache Objects, [149](#)
- `hwloc_get_thread_cpupbind`
  - CPU binding, [115](#)
- `hwloc_get_type_depth`
  - Object levels, depths and types, [106](#)
- `hwloc_get_type_depth_e`
  - Object levels, depths and types, [104](#)
- `hwloc_get_type_or_above_depth`
  - Object levels, depths and types, [106](#)
- `hwloc_get_type_or_below_depth`
  - Object levels, depths and types, [107](#)
- `hwloc_get_whole_distance_matrix_by_-depth`
  - Manipulating Distances, [160](#)
- `hwloc_get_whole_distance_matrix_by_-type`
  - Manipulating Distances, [161](#)
- `hwloc_gl_get_display_by_osdev`
  - Interoperability with OpenGL displays, [192](#)
- `hwloc_gl_get_display_osdev_by_name`
  - Interoperability with OpenGL displays, [192](#)
- `hwloc_gl_get_display_osdev_by_port_-device`
  - Interoperability with OpenGL displays, [192](#)
- `hwloc_hide_errors`
  - Components and Plugins: Core functions to be used by components, [209](#)
- `hwloc_ibv_get_device_cpuset`
  - Interoperability with OpenFabrics, [195](#)
- `hwloc_ibv_get_device_osdev`
  - Interoperability with OpenFabrics, [195](#)
- `hwloc_ibv_get_device_osdev_by_name`
  - Interoperability with OpenFabrics, [196](#)
- `hwloc_insert_object_by_cpuset`
  - Components and Plugins: Core functions to be used by components, [209](#)
- `hwloc_insert_object_by_parent`
  - Components and Plugins: Core functions to be used by components, [209](#)
- `hwloc_insert_pci_device_list`
  - Components and Plugins: PCI functions to be used by components, [211](#)
- `hwloc_intel_mic_get_device_cpuset`
  - Interoperability with Intel Xeon Phi (-MIC), [194](#)
- `hwloc_intel_mic_get_device_osdev_by_-index`
  - Interoperability with Intel Xeon Phi (-MIC), [194](#)
- `hwloc_linux_get_tid_cpupbind`
  - Linux-specific helpers, [177](#)
- `hwloc_linux_get_tid_last_cpu_location`
  - Linux-specific helpers, [177](#)
- `hwloc_linux_parse_cpumap_file`
  - Linux-specific helpers, [177](#)
- `hwloc_linux_set_tid_cpupbind`
  - Linux-specific helpers, [177](#)

- hwloc\_membind\_flags\_t
  - Memory binding, 119
- hwloc\_membind\_policy\_t
  - Memory binding, 120
- hwloc\_mx\_board\_get\_device\_cpuset
  - Interoperability with Myrinet Express, 197
- hwloc\_mx\_endpoint\_get\_device\_cpuset
  - Interoperability with Myrinet Express, 197
- hwloc\_nodese\_t\_from\_linux\_libnuma\_ - bitmask
  - Interoperability with Linux libnuma bitmask, 182
- hwloc\_nodese\_t\_from\_linux\_libnuma\_ - ulongs
  - Interoperability with Linux libnuma unsigned long masks, 180
- hwloc\_nodese\_t
  - Object Sets (hwloc\_cpuset\_t and hwloc\_nodese\_t), 85
- hwloc\_nodese\_t\_to\_linux\_libnuma\_ - bitmask
  - Interoperability with Linux libnuma bitmask, 182
- hwloc\_nodese\_t\_to\_linux\_libnuma\_ ulongs
  - Interoperability with Linux libnuma unsigned long masks, 180
- hwloc\_nvml\_get\_device\_cpuset
  - Interoperability with the NVIDIA Management Library, 190
- hwloc\_nvml\_get\_device\_osdev
  - Interoperability with the NVIDIA Management Library, 190
- hwloc\_nvml\_get\_device\_osdev\_by\_index
  - Interoperability with the NVIDIA Management Library, 191
- hwloc\_obj, 222
  - allowed\_cpuset, 223
  - allowed\_nodese\_t, 223
  - arity, 223
  - attr, 223
  - children, 224
  - complete\_cpuset, 224
  - complete\_nodese\_t, 224
  - cpuset, 224
  - depth, 225
  - distances, 225
  - distances\_count, 225
  - first\_child, 225
  - infos, 225
  - infos\_count, 225
  - last\_child, 225
  - logical\_index, 225
  - memory, 225
  - name, 226
  - next\_cousin, 226
  - next\_sibling, 226
  - nodese\_t, 226
  - online\_cpuset, 226
  - os\_index, 226
  - os\_level, 227
  - parent, 227
  - prev\_cousin, 227
  - prev\_sibling, 227
  - sibling\_rank, 227
  - symmetric\_subtree, 227
  - type, 227
  - userdata, 227
- hwloc\_obj\_add\_info
  - Consulting and Adding Key-Value Info Attributes, 111
- hwloc\_obj\_attr\_snprintf
  - Converting between Object Types, Sets and Attributes, and Strings, 108
- hwloc\_obj\_attr\_u, 228
  - bridge, 228
  - cache, 228
  - group, 228
  - osdev, 228
  - pcidev, 228
- hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, 215
  - depth, 215
  - domain, 215
  - downstream, 215
  - downstream\_type, 215
  - pci, 215
  - secondary\_bus, 215
  - subordinate\_bus, 215
  - upstream, 215
  - upstream\_type, 216
- hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s, 216
  - associativity, 216
  - depth, 216
  - linesize, 216
  - size, 216
  - type, 216

- hwloc\_obj\_attr\_u::hwloc\_group\_attr\_s,
  - 221
  - depth, 222
  - dont\_merge, 222
- hwloc\_obj\_attr\_u::hwloc\_osdev\_attr\_s,
  - 231
  - type, 231
- hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s,
  - 232
  - bus, 232
  - class\_id, 232
  - dev, 232
  - device\_id, 232
  - domain, 232
  - func, 232
  - linkspeed, 232
  - revision, 232
  - subdevice\_id, 232
  - subvendor\_id, 232
  - vendor\_id, 232
- hwloc\_obj\_bridge\_type\_e
  - Object Types, 88
- hwloc\_obj\_bridge\_type\_t
  - Object Types, 87
- hwloc\_obj\_cache\_type\_e
  - Object Types, 88
- hwloc\_obj\_cache\_type\_t
  - Object Types, 87
- hwloc\_obj\_cpuset\_sprintf
  - Converting between Object Types, Sets and Attributes, and Strings, 108
- hwloc\_obj\_get\_info\_by\_name
  - Consulting and Adding Key-Value Info Attributes, 111
- hwloc\_obj\_info\_s, 229
  - name, 229
  - value, 229
- hwloc\_obj\_is\_in\_subtree
  - Looking at Ancestor and Child Objects, 148
- hwloc\_obj\_memory\_s, 230
  - local\_memory, 231
  - page\_types, 231
  - page\_types\_len, 231
  - total\_memory, 231
- hwloc\_obj\_memory\_s::hwloc\_obj\_memory\_page\_type\_s, 229
  - count, 230
  - size, 230
- hwloc\_obj\_osdev\_type\_e
  - Object Types, 88
- hwloc\_obj\_osdev\_type\_t
  - Object Types, 87
- hwloc\_obj\_t
  - Object Structure and Attributes, 91
- hwloc\_obj\_type\_sprintf
  - Converting between Object Types, Sets and Attributes, and Strings, 108
- hwloc\_obj\_type\_sscanf
  - Converting between Object Types, Sets and Attributes, and Strings, 109
- hwloc\_obj\_type\_string
  - Converting between Object Types, Sets and Attributes, and Strings, 109
- hwloc\_obj\_type\_t
  - Object Types, 89
- hwloc\_opencl\_get\_device\_cpuset
  - Interoperability with OpenCL, 184
- hwloc\_opencl\_get\_device\_osdev
  - Interoperability with OpenCL, 184
- hwloc\_opencl\_get\_device\_osdev\_by\_index
  - Interoperability with OpenCL, 185
- hwloc\_pci\_find\_cap
  - Components and Plugins: PCI functions to be used by components, 211
- hwloc\_pci\_find\_linkspeed
  - Components and Plugins: PCI functions to be used by components, 211
- hwloc\_pci\_prepare\_bridge
  - Components and Plugins: PCI functions to be used by components, 211
- hwloc\_plugin\_check\_namespace
  - Components and Plugins: Core functions to be used by components, 209
- hwloc\_report\_error\_t
  - Components and Plugins: Core functions to be used by components, 208
- hwloc\_report\_os\_error
  - Components and Plugins: Core functions to be used by components,

- 210
- hwloc\_restrict\_flags\_e
  - Modifying a loaded Topology, 130
- hwloc\_set\_area\_membind
  - Memory binding, 126
- hwloc\_set\_area\_membind\_nodese
  - Memory binding, 127
- hwloc\_set\_cpubind
  - CPU binding, 115
- hwloc\_set\_membind
  - Memory binding, 127
- hwloc\_set\_membind\_nodese
  - Memory binding, 127
- hwloc\_set\_proc\_cp
  - CPU binding, 115
- hwloc\_set\_proc\_membind
  - Memory binding, 128
- hwloc\_set\_proc\_membind\_nodese
  - Memory binding, 128
- hwloc\_set\_thread\_cp
  - CPU binding, 116
- hwloc\_topology\_check
  - Topology Creation and Destruction, 92
- hwloc\_topology\_cp
  - get\_proc\_cp
    - 233
  - get\_proc\_last\_cpu\_location, 233
  - get\_thisproc\_cp
    - 233
  - get\_thisproc\_last\_cpu\_location, 234
  - get\_thisthread\_cp
    - 234
  - get\_thisthread\_last\_cpu\_location, 234
  - get\_thread\_cp
    - 234
  - set\_proc\_cp
    - 234
  - set\_thisproc\_cp
    - 234
  - set\_thisthread\_cp
    - 234
  - set\_thread\_cp
    - 234
- hwloc\_topology\_destroy
  - Topology Creation and Destruction, 92
- hwloc\_topology\_diff\_apply
  - Topology differences, 200
- hwloc\_topology\_diff\_apply\_flags\_e
  - Topology differences, 199
- hwloc\_topology\_diff\_build
  - Topology differences, 200
- hwloc\_topology\_diff\_destroy
  - Topology differences, 201
- hwloc\_topology\_diff\_export\_xml
  - Topology differences, 201
- hwloc\_topology\_diff\_export\_xmlbuffer
  - Topology differences, 202
- hwloc\_topology\_diff\_load\_xml
  - Topology differences, 202
- hwloc\_topology\_diff\_load\_xmlbuffer
  - Topology differences, 202
- hwloc\_topology\_diff\_obj\_attr\_type\_e
  - Topology differences, 199
- hwloc\_topology\_diff\_obj\_attr\_type\_t
  - Topology differences, 199
- hwloc\_topology\_diff\_obj\_attr\_u, 237
  - generic, 238
  - string, 238
  - uint64, 238
- hwloc\_topology\_diff\_t
  - Topology differences, 199
- hwloc\_topology\_diff\_type\_e
  - Topology differences, 200
- hwloc\_topology\_diff\_type\_t
  - Topology differences, 199
- hwloc\_topology\_diff\_u, 240
  - generic, 240
  - obj\_attr, 240
  - too\_complex, 240
- hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_generic\_s, 235
  - next, 235
  - type, 235
- hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, 236
  - diff, 236
  - next, 236
  - obj\_depth, 236
  - obj\_index, 236
  - type, 236
- hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s, 239
  - next, 239
  - obj\_depth, 239
  - obj\_index, 239
  - type, 239
- hwloc\_topology\_discovery\_support, 240
  - pu, 241
- hwloc\_topology\_dup
  - Topology Creation and Destruction, 93
- hwloc\_topology\_export\_synthetic
  - Exporting Topologies to Synthetic, 139
- hwloc\_topology\_export\_synthetic\_flags\_e

- Exporting Topologies to Synthetic, [139](#)
- hwloc\_topology\_export\_xml
  - Exporting Topologies to XML, [136](#)
- hwloc\_topology\_export\_xmlbuffer
  - Exporting Topologies to XML, [136](#)
- hwloc\_topology\_flags\_e
  - Topology Detection Configuration and Query, [96](#)
- hwloc\_topology\_get\_allowed\_cpuset
  - CPU and node sets of entire topologies, [155](#)
- hwloc\_topology\_get\_allowed\_nodeset
  - CPU and node sets of entire topologies, [155](#)
- hwloc\_topology\_get\_complete\_cpuset
  - CPU and node sets of entire topologies, [156](#)
- hwloc\_topology\_get\_complete\_nodeset
  - CPU and node sets of entire topologies, [156](#)
- hwloc\_topology\_get\_depth
  - Object levels, depths and types, [107](#)
- hwloc\_topology\_get\_flags
  - Topology Detection Configuration and Query, [97](#)
- hwloc\_topology\_get\_online\_cpuset
  - CPU and node sets of entire topologies, [156](#)
- hwloc\_topology\_get\_support
  - Topology Detection Configuration and Query, [98](#)
- hwloc\_topology\_get\_topology\_cpuset
  - CPU and node sets of entire topologies, [156](#)
- hwloc\_topology\_get\_topology\_nodeset
  - CPU and node sets of entire topologies, [157](#)
- hwloc\_topology\_get\_userdata
  - Topology Detection Configuration and Query, [98](#)
- hwloc\_topology\_ignore\_all\_keep\_-structure
  - Topology Detection Configuration and Query, [98](#)
- hwloc\_topology\_ignore\_type
  - Topology Detection Configuration and Query, [98](#)
- hwloc\_topology\_ignore\_type\_keep\_-structure
  - Topology Detection Configuration and Query, [98](#)
- Topology Detection Configuration and Query, [98](#)
- hwloc\_topology\_init
  - Topology Creation and Destruction, [93](#)
- hwloc\_topology\_insert\_misc\_object\_by\_-cpuset
  - Modifying a loaded Topology, [130](#)
- hwloc\_topology\_insert\_misc\_object\_by\_-parent
  - Modifying a loaded Topology, [131](#)
- hwloc\_topology\_is\_thissystem
  - Topology Detection Configuration and Query, [99](#)
- hwloc\_topology\_load
  - Topology Creation and Destruction, [93](#)
- hwloc\_topology\_membind\_support, [241](#)
  - alloc\_membind, [242](#)
  - bind\_membind, [242](#)
  - firsttouch\_membind, [242](#)
  - get\_area\_membind, [242](#)
  - get\_area\_memlocation, [242](#)
  - get\_proc\_membind, [242](#)
  - get\_thisproc\_membind, [242](#)
  - get\_thisthread\_membind, [242](#)
  - interleave\_membind, [242](#)
  - migrate\_membind, [242](#)
  - nexttouch\_membind, [243](#)
  - replicate\_membind, [243](#)
  - set\_area\_membind, [243](#)
  - set\_proc\_membind, [243](#)
  - set\_thisproc\_membind, [243](#)
  - set\_thisthread\_membind, [243](#)
- hwloc\_topology\_restrict
  - Modifying a loaded Topology, [131](#)
- hwloc\_topology\_set\_custom
  - Topology Detection Configuration and Query, [99](#)
- hwloc\_topology\_set\_distance\_matrix
  - Topology Detection Configuration and Query, [99](#)
- hwloc\_topology\_set\_flags
  - Topology Detection Configuration and Query, [100](#)
- hwloc\_topology\_set\_fsroot
  - Topology Detection Configuration and Query, [100](#)
- hwloc\_topology\_set\_pid



- Topology Detection Configuration and Query, [100](#)
- hwloc\_topology\_set\_synthetic
  - Topology Detection Configuration and Query, [101](#)
- hwloc\_topology\_set\_userdata
  - Topology Detection Configuration and Query, [101](#)
- hwloc\_topology\_set\_userdata\_export\_callback
  - Exporting Topologies to XML, [137](#)
- hwloc\_topology\_set\_userdata\_import\_callback
  - Exporting Topologies to XML, [137](#)
- hwloc\_topology\_set\_xml
  - Topology Detection Configuration and Query, [101](#)
- hwloc\_topology\_set\_xmlbuffer
  - Topology Detection Configuration and Query, [102](#)
- hwloc\_topology\_support, [243](#)
  - cpubind, [244](#)
  - discovery, [244](#)
  - membind, [244](#)
- hwloc\_topology\_t
  - Topology Creation and Destruction, [92](#)
- index
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s, [238](#)
- infos
  - hwloc\_obj, [225](#)
- infos\_count
  - hwloc\_obj, [225](#)
- init
  - hwloc\_component, [218](#)
- instantiate
  - hwloc\_disc\_component, [219](#)
- interleave\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- is\_custom
  - hwloc\_backend, [214](#)
- is\_thissystem
  - hwloc\_backend, [214](#)
- last\_child
  - hwloc\_obj, [225](#)
- latency
  - hwloc\_distances\_s, [220](#)
- latency\_base
  - hwloc\_distances\_s, [221](#)
- latency\_max
  - hwloc\_distances\_s, [221](#)
- linesize
  - hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s, [216](#)
- linkspeed
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- local\_memory
  - hwloc\_obj\_memory\_s, [231](#)
- logical\_index
  - hwloc\_obj, [225](#)
- membind
  - hwloc\_topology\_support, [244](#)
- memory
  - hwloc\_obj, [225](#)
- migrate\_membind
  - hwloc\_topology\_membind\_support, [242](#)
- name
  - hwloc\_disc\_component, [219](#)
  - hwloc\_obj, [226](#)
  - hwloc\_obj\_info\_s, [229](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_string\_s, [237](#)
- nbobjs
  - hwloc\_distances\_s, [221](#)
- newvalue
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_string\_s, [237](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s, [238](#)
- next
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_generic\_s, [235](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, [236](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s, [239](#)
- next\_cousin

- hwloc\_obj, [226](#)
- next\_sibling
  - hwloc\_obj, [226](#)
- nexttouch\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- nodeset
  - hwloc\_obj, [226](#)
- notify\_new\_object
  - hwloc\_backend, [214](#)
- obj\_attr
  - hwloc\_topology\_diff\_u, [240](#)
- obj\_depth
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, [236](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s, [239](#)
- obj\_index
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, [236](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s, [239](#)
- oldvalue
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_string\_s, [237](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s, [239](#)
- online\_cpuset
  - hwloc\_obj, [226](#)
- os\_index
  - hwloc\_obj, [226](#)
- os\_level
  - hwloc\_obj, [227](#)
- osdev
  - hwloc\_obj\_attr\_u, [228](#)
- page\_types
  - hwloc\_obj\_memory\_s, [231](#)
- page\_types\_len
  - hwloc\_obj\_memory\_s, [231](#)
- parent
  - hwloc\_obj, [227](#)
- pci
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
- pcidev
  - hwloc\_obj\_attr\_u, [228](#)
- prev\_cousin
  - hwloc\_obj, [227](#)
- prev\_sibling
  - hwloc\_obj, [227](#)
- priority
  - hwloc\_disc\_component, [219](#)
- private\_data
  - hwloc\_backend, [214](#)
- pu
  - hwloc\_topology\_discovery\_support, [241](#)
- relative\_depth
  - hwloc\_distances\_s, [221](#)
- replicate\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- revision
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- secondary\_bus
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
- set\_area\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- set\_proc\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)
- set\_proc\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- set\_thisproc\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)
- set\_thisproc\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- set\_thisthread\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)
- set\_thisthread\_membind
  - hwloc\_topology\_membind\_support, [243](#)
- set\_thread\_cpupbind
  - hwloc\_topology\_cpupbind\_support, [234](#)

- sibling\_rank
  - hwloc\_obj, [227](#)
- size
  - hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s, [216](#)
  - hwloc\_obj\_memory\_s::hwloc\_obj\_memory\_page\_type\_s, [230](#)
- string
  - hwloc\_topology\_diff\_obj\_attr\_u, [238](#)
- subdevice\_id
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- subordinate\_bus
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
- subvendor\_id
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)
- symmetric\_subtree
  - hwloc\_obj, [227](#)
  
- too\_complex
  - hwloc\_topology\_diff\_u, [240](#)
- total\_memory
  - hwloc\_obj\_memory\_s, [231](#)
- type
  - hwloc\_component, [218](#)
  - hwloc\_disc\_component, [220](#)
  - hwloc\_obj, [227](#)
  - hwloc\_obj\_attr\_u::hwloc\_cache\_attr\_s, [216](#)
  - hwloc\_obj\_attr\_u::hwloc\_osdev\_attr\_s, [231](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_generic\_s, [235](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_string\_s, [237](#)
  - hwloc\_topology\_diff\_obj\_attr\_u::hwloc\_topology\_diff\_obj\_attr\_uint64\_s, [239](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_generic\_s, [235](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_obj\_attr\_s, [236](#)
  - hwloc\_topology\_diff\_u::hwloc\_topology\_diff\_too\_complex\_s, [239](#)
- uint64
  - hwloc\_topology\_diff\_obj\_attr\_u, [238](#)
- upstream
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [215](#)
- upstream\_type
  - hwloc\_obj\_attr\_u::hwloc\_bridge\_attr\_s, [216](#)
- userdata
  - hwloc\_obj, [227](#)
- value
  - hwloc\_obj\_info\_s, [229](#)
- vendor\_id
  - hwloc\_obj\_attr\_u::hwloc\_pcidev\_attr\_s, [232](#)