

Hardware Locality (hwloc) Reference Manual

1.1rc6r2967

Generated by Doxygen 1.3.9.1

Tue Dec 21 21:04:53 2010

Contents

1	Hardware Locality	1
1.1	Introduction	1
1.2	Installation	2
1.3	CLI Examples	3
1.4	Programming Interface	10
1.5	Questions and Bugs	18
1.6	History / Credits	18
1.7	Further Reading	18
1.8	lstopo	22
1.9	hwloc-bind	22
1.10	hwloc-calc	22
1.11	hwloc-distrib	23
1.12	hwloc-ps	23
1.13	Using hwloc's M4 Embedding Capabilities	27
1.14	Example Embedding hwloc	29
1.15	Topology Context vs. Caching	30
1.16	Hierarchy vs. Core@Socket	30
1.17	Logical vs. Physical/OS Indexes	30
1.18	Counting Specification	31
1.19	I do not want hwloc to rediscover my enormous machine topology ev- erytime I rerun a process	31
2	Hardware Locality (hwloc) Module Index	33

2.1	Hardware Locality (hwloc) Modules	33
3	Hardware Locality (hwloc) Data Structure Index	35
3.1	Hardware Locality (hwloc) Data Structures	35
4	Hardware Locality (hwloc) File Index	37
4.1	Hardware Locality (hwloc) File List	37
5	Hardware Locality (hwloc) Module Documentation	39
5.1	API version	39
5.2	Topology context	40
5.3	Object sets	41
5.4	Topology Object Types	43
5.5	Topology Objects	46
5.6	Create and Destroy Topologies	47
5.7	Configure Topology Detection	49
5.8	Tinker with topologies.	55
5.9	Get some Topology Information	57
5.10	Retrieve Objects	60
5.11	Object/String Conversion	61
5.12	CPU binding	64
5.13	Memory binding	68
5.14	Object Type Helpers	76
5.15	Basic Traversal Helpers	77
5.16	Finding Objects Inside a CPU set	80
5.17	Finding a single Object covering at least CPU set	83
5.18	Finding a set of similar Objects covering at least a CPU set	84
5.19	Cache-specific Finding Helpers	85
5.20	Advanced Traversal Helpers	86
5.21	Binding Helpers	88
5.22	Cpuset Helpers	90
5.23	Nodeset Helpers	92

5.24	Conversion between cpuset and nodeset	93
5.25	The bitmap API	95
5.26	Helpers for manipulating glibc sched affinity	106
5.27	Linux-only helpers	107
5.28	Helpers for manipulating Linux libnuma unsigned long masks	109
5.29	Helpers for manipulating Linux libnuma bitmask	111
5.30	Helpers for manipulating Linux libnuma nodemask_t	113
5.31	CUDA Driver API Specific Functions	115
5.32	CUDA Runtime API Specific Functions	116
5.33	OpenFabrics-Specific Functions	117
5.34	Myrinet Express-Specific Functions	118
6	Hardware Locality (hwloc) Data Structure Documentation	119
6.1	hwloc_obj Struct Reference	119
6.2	hwloc_obj_attr_u Union Reference	127
6.3	hwloc_obj_attr_u::hwloc_cache_attr_s Struct Reference	128
6.4	hwloc_obj_attr_u::hwloc_group_attr_s Struct Reference	129
6.5	hwloc_obj_info_s Struct Reference	130
6.6	hwloc_obj_memory_s Struct Reference	131
6.7	hwloc_obj_memory_s::hwloc_obj_memory_page_type_s Struct Reference	133
6.8	hwloc_topology_cpubind_support Struct Reference	134
6.9	hwloc_topology_discovery_support Struct Reference	136
6.10	hwloc_topology_membind_support Struct Reference	137
6.11	hwloc_topology_support Struct Reference	140
7	Hardware Locality (hwloc) File Documentation	141
7.1	bitmap.h File Reference	141
7.2	cuda.h File Reference	146
7.3	cuda.h File Reference	147
7.4	glibc-sched.h File Reference	148
7.5	helper.h File Reference	149

7.6	hwloc.doxy File Reference	154
7.7	hwloc.h File Reference	155
7.8	linux-libnuma.h File Reference	164
7.9	linux.h File Reference	166
7.10	myriexpress.h File Reference	167
7.11	openfabrics-verbs.h File Reference	168

Chapter 1

Hardware Locality

Portable abstraction of hierarchical architectures for high-performance computing

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 sockets, processor cores, and processing units (logical processors or "threads"). hwloc also gathers various attributes such as cache and memory information, and is portable across a variety of different operating systems and platforms.

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.

Note that the hwloc project represents the merger of the libtopology project from INRIA and the Portable Linux Processor Affinity (PLPA) sub-project from Open MPI. *Both of these prior projects are now deprecated.* The first hwloc release is essentially a "re-branding" of the libtopology code base, but with both a few genuinely new features and a few PLPA-like features added in. More new features and more PLPA-like features will be added to hwloc over time. See [switchfromplpa](#) for more details about converting your application from PLPA to hwloc.

hwloc supports the following operating systems:

- Linux (including old kernels not having sysfs topology information, with knowledge of cpusets, offline cpus, ScaleMP vSMP, and Kerrighed support)
- Solaris

- AIX
- Darwin / OS X
- FreeBSD and its variants, such as kFreeBSD/GNU
- OSF/1 (a.k.a., Tru64)
- HP-UX
- Microsoft Windows

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
- Remote machine simulation through the gathering of Linux sysfs topology 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 [CLI 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.

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.

Nightly development snapshots are available on the web site. Additionally, the code can be directly checked out of Subversion:

```
shell$ svn checkout http://svn.open-mpi.org/svn/hwloc/trunk hwloc-trunk
shell$ cd hwloc-trunk
shell$ ./autogen.sh
```

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


Installation by itself 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 can be found when hwloc is configured and build. Similarly, lstopo's XML support requires the libxml2 development package.

1.3 CLI Examples

On a 4-socket 2-core machine with hyperthreading, the lstopo tool may show the following graphical output:



Here's the equivalent output in textual form:

```
Machine (16GB)
Socket 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)
Socket 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)
Socket 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)
Socket 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)

```

Finally, here's the equivalent output in XML. Long lines were artificially broken for document clarity (in the real output, each XML tag is on a single line), and only socket #0 is shown for brevity:

```

<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE topology SYSTEM "hwloc.dtd">
<topology>
  <object type="Machine" os_level="-1" os_index="0" cpuset="0x0000ffff"
    complete_cpuset="0x0000ffff" online_cpuset="0x0000ffff"
    allowed_cpuset="0x0000ffff"
    dmi_board_vendor="Dell Computer Corporation" dmi_board_name="ORD318"
    local_memory="16648183808">
    <page_type size="4096" count="4064498"/>
    <page_type size="2097152" count="0"/>
    <object type="Socket" os_level="-1" os_index="0" cpuset="0x00001111"
      complete_cpuset="0x00001111" online_cpuset="0x00001111"
      allowed_cpuset="0x00001111">
      <object type="Cache" os_level="-1" cpuset="0x00001111"
        complete_cpuset="0x00001111" online_cpuset="0x00001111"
        allowed_cpuset="0x00001111" cache_size="4194304" depth="3"
        cache_linesize="64">
        <object type="Cache" os_level="-1" cpuset="0x00000101"
          complete_cpuset="0x00000101" online_cpuset="0x00000101"
          allowed_cpuset="0x00000101" cache_size="1048576" depth="2"
          cache_linesize="64">
          <object type="Cache" os_level="-1" cpuset="0x00000101"
            complete_cpuset="0x00000101" online_cpuset="0x00000101"
            allowed_cpuset="0x00000101" cache_size="16384" depth="1"
            cache_linesize="64">
            <object type="Core" os_level="-1" os_index="0" cpuset="0x00000101"
              complete_cpuset="0x00000101" online_cpuset="0x00000101"
              allowed_cpuset="0x00000101">
              <object type="PU" os_level="-1" os_index="0" cpuset="0x00000001"
                complete_cpuset="0x00000001" online_cpuset="0x00000001"
                allowed_cpuset="0x00000001"/>
              <object type="PU" os_level="-1" os_index="8" cpuset="0x00000100"
                complete_cpuset="0x00000100" online_cpuset="0x00000100"
                allowed_cpuset="0x00000100"/>
            </object>
          </object>
        </object>
      </object>
    </object>
  </object>
  <object type="Cache" os_level="-1" cpuset="0x00001010"
    complete_cpuset="0x00001010" online_cpuset="0x00001010"
    allowed_cpuset="0x00001010" cache_size="1048576" depth="2"
    cache_linesize="64">
    <object type="Cache" os_level="-1" cpuset="0x00001010"

```

```
complete_cpuset="0x00001010" online_cpuset="0x00001010"
allowed_cpuset="0x00001010" cache_size="16384" depth="1"
cache_linesize="64">
<object type="Core" os_level="-1" os_index="1" cpuset="0x00001010"
  complete_cpuset="0x00001010" online_cpuset="0x00001010"
  allowed_cpuset="0x00001010">
  <object type="PU" os_level="-1" os_index="4" cpuset="0x00000010"
    complete_cpuset="0x00000010" online_cpuset="0x00000010"
    allowed_cpuset="0x00000010"/>
  <object type="PU" os_level="-1" os_index="12" cpuset="0x00001000"
    complete_cpuset="0x00001000" online_cpuset="0x00001000"
    allowed_cpuset="0x00001000"/>
  </object>
</object>
</object>
</object>
</object>
<!-- ...other sockets listed here ... -->
</object>
</topology>
```

On a 4-socket 2-core Opteron NUMA machine, the `lstopo` tool may show the following graphical output:




Figure 1.1: width=

Here's the equivalent output in textual form:

```
Machine (32GB)
  NUMANode L#0 (P#0 8190MB) + Socket 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) + Socket 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) + Socket 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) + Socket 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)
```

And here's the equivalent output in XML. Similar to above, line breaks were added and only PU #0 is shown for brevity:


```
<?xml version="1.0" encoding="UTF-8"?>
```

```

<!DOCTYPE topology SYSTEM "hwloc.dtd">
<topology>
  <object type="Machine" os_level="-1" os_index="0" cpuset="0x000000ff"
    complete_cpuset="0x000000ff" online_cpuset="0x000000ff"
    allowed_cpuset="0x000000ff" nodeset="0x000000ff"
    complete_nodeset="0x000000ff" allowed_nodeset="0x000000ff"
    dmi_board_vendor="TYAN Computer Corp" dmi_board_name="S4881 ">
  <page_type size="4096" count="0"/>
  <page_type size="2097152" count="0"/>
  <object type="NUMANode" os_level="-1" os_index="0" cpuset="0x00000003"
    complete_cpuset="0x00000003" online_cpuset="0x00000003"
    allowed_cpuset="0x00000003" nodeset="0x00000001"
    complete_nodeset="0x00000001" allowed_nodeset="0x00000001"
    local_memory="7514177536">
  <page_type size="4096" count="1834516"/>
  <page_type size="2097152" count="0"/>
  <object type="Socket" os_level="-1" os_index="0" cpuset="0x00000003"
    complete_cpuset="0x00000003" online_cpuset="0x00000003"
    allowed_cpuset="0x00000003" nodeset="0x00000001"
    complete_nodeset="0x00000001" allowed_nodeset="0x00000001">
  <object type="Cache" os_level="-1" cpuset="0x00000001"
    complete_cpuset="0x00000001" online_cpuset="0x00000001"
    allowed_cpuset="0x00000001" nodeset="0x00000001"
    complete_nodeset="0x00000001" allowed_nodeset="0x00000001"
    cache_size="1048576" depth="2" cache_linesize="64">
  <object type="Cache" os_level="-1" cpuset="0x00000001"
    complete_cpuset="0x00000001" online_cpuset="0x00000001"
    allowed_cpuset="0x00000001" nodeset="0x00000001"
    complete_nodeset="0x00000001" allowed_nodeset="0x00000001"
    cache_size="65536" depth="1" cache_linesize="64">
  <object type="Core" os_level="-1" os_index="0"
    cpuset="0x00000001" complete_cpuset="0x00000001"
    online_cpuset="0x00000001" allowed_cpuset="0x00000001"
    nodeset="0x00000001" complete_nodeset="0x00000001"
    allowed_nodeset="0x00000001">
  <object type="PU" os_level="-1" os_index="0" cpuset="0x00000001"
    complete_cpuset="0x00000001" online_cpuset="0x00000001"
    allowed_cpuset="0x00000001" nodeset="0x00000001"
    complete_nodeset="0x00000001" allowed_nodeset="0x00000001"/>
  </object>
</object>
</object>
</object>
  <!-- ...more objects listed here ... -->
</topology>

```

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



Here's the same output in textual form:

```
Machine (16GB)
Socket 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)
Socket 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)
```

And the same output in XML (line breaks added, only PU #0 shown):

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE topology SYSTEM "hwloc.dtd">
<topology>
  <object type="Machine" os_level="-1" os_index="0" cpuset="0x000000ff"
    complete_cpuset="0x000000ff" online_cpuset="0x000000ff"
    allowed_cpuset="0x000000ff" dmi_board_vendor="Dell Inc."
    dmi_board_name="0NR282" local_memory="16865292288">
    <page_type size="4096" count="4117503"/>
    <page_type size="2097152" count="0"/>
    <object type="Socket" os_level="-1" os_index="0" cpuset="0x00000055"
      complete_cpuset="0x00000055" online_cpuset="0x00000055"
      allowed_cpuset="0x00000055">
      <object type="Cache" os_level="-1" cpuset="0x00000011"
        complete_cpuset="0x00000011" online_cpuset="0x00000011"
```

```

    allowed_cpuset="0x00000011" cache_size="4194304" depth="2"
    cache_linesize="64">
<object type="Cache" os_level="-1" cpuset="0x00000001"
  complete_cpuset="0x00000001" online_cpuset="0x00000001"
  allowed_cpuset="0x00000001" cache_size="32768" depth="1"
  cache_linesize="64">
  <object type="Core" os_level="-1" os_index="0" cpuset="0x00000001"
    complete_cpuset="0x00000001" online_cpuset="0x00000001"
    allowed_cpuset="0x00000001">
    <object type="PU" os_level="-1" os_index="0" cpuset="0x00000001"
      complete_cpuset="0x00000001" online_cpuset="0x00000001"
      allowed_cpuset="0x00000001"/>
    </object>
  </object>
</object>
<object type="Cache" os_level="-1" cpuset="0x00000010"
  complete_cpuset="0x00000010" online_cpuset="0x00000010"
  allowed_cpuset="0x00000010" cache_size="32768" depth="1"
  cache_linesize="64">
  <object type="Core" os_level="-1" os_index="1" cpuset="0x00000010"
    complete_cpuset="0x00000010" online_cpuset="0x00000010"
    allowed_cpuset="0x00000010">
    <object type="PU" os_level="-1" os_index="4" cpuset="0x00000010"
      complete_cpuset="0x00000010" online_cpuset="0x00000010"
      allowed_cpuset="0x00000010"/>
    </object>
  </object>
</object>
</object>
<!-- ...more objects listed here ... -->
</topology>

```

1.4 Programming Interface

The basic interface is available in [hwloc.h](#). It essentially offers low-level routines for advanced programmers that want to manually manipulate objects and follow links between them. Documentation for everything in [hwloc.h](#) are provided later in this document. Developers should also look at [hwloc/helper.h](#) (and also in this document, which provides good higher-level topology traversal examples).

To precisely define the vocabulary used by hwloc, a `termsanddefs` 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 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 Subversion checkout, 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


As shown in [CLI 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 32 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, sockets, or cores is available.

Similarly, Operating Systems 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.

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




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



Notice that hwloc only sees half the PUs when SMT is disabled. PU #15, 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. Hence, PU #15 in the SMT-disabled picture probably corresponds to PU #30 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, there may someday be an element "lower" than a PU, or perhaps a new element may exist between a core and a PU.

1.4.2 API Example

The following small C example (named "hwloc-hello.c") prints the topology of the machine and bring the process to the first logical processor of the second core of the machine.

```

/* Example hwloc API program.
 *
 * Copyright Â© 2009-2010 INRIA
 * Copyright Â© 2009-2010 UniversitÃ© Bordeaux 1
 * Copyright Â© 2009-2010 Cisco Systems, Inc. All rights reserved.
 *
 * 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 string[128];
    unsigned i;

    hwloc_obj_snprintf(string, sizeof(string), topology, obj, "#", 0);
    printf("%s%s\n", 2*depth, "", string);
    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_snprintf(string, sizeof(string), topology,
                           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 sockets.
 *****/
depth = hwloc_get_type_depth(topology, HWLOC_OBJ_SOCKET);
if (depth == HWLOC_TYPE_DEPTH_UNKNOWN) {
    printf("*** The number of sockets is unknown\n");
} else {
    printf("*** %u socket(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_nbobjs_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/hyperthreaded). */
    hwloc_bitmap_singlify(cpuset);

    /* And try to bind ourself there. */
    if (hwloc_set_cpupbind(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_NODE);
if (n) {
    void *m;
    size_t size = 1024*1024;

```

```
    obj = hwloc_get_obj_by_type(topology, HWLOC_OBJ_NODE, n - 1);
    m = hwloc_alloc_membind_nodset(topology, size, obj->nodset,
        HWLOC_MEMBIND_DEFAULT, 0);
    hwloc_free(topology, m, size);

    m = malloc(size);
    hwloc_set_area_membind_nodset(topology, m, size, obj->nodset,
        HWLOC_MEMBIND_DEFAULT, 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 += $(pkg-config --cflags hwloc)
LDLIBS += $(pkg-config --libs hwloc)
cc hwloc-hello.c $(CFLAGS) -o hwloc-hello $(LDLIBS)
```

On a machine with 4GB of RAM and 2 processor sockets – each socket 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: Socket#0
Index 1: Socket#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)
  Socket#0
    Core#0
      PU#0
    Core#1
      PU#1
  Socket#1
    Core#3
```

```
    PU#2
    Core#2
    PU#3
*** 2 socket(s)
shell$
```

1.5 Questions and Bugs

Questions should be sent to the devel mailing list (<http://www.open-mpi.org/community/lists/hwloc.php>).

Bug reports should be reported in the tracker (<https://svn.open-mpi.org/trac/hwloc/>).

If hwloc discovers an incorrect topology for your machine, the very first thing you should check is to ensure that you have the most recent updates installed for your operating system. Indeed, most of hwloc topology discovery relies on hardware information retrieved through the operation system (e.g., via the /sys virtual filesystem of the Linux kernel). If upgrading your OS or Linux kernel does not solve your problem, you may also want to ensure that you are running the most recent version of the BIOS for your machine.

If those things fail, contact us on the mailing list for additional help. Please attach the output of `lstopo` after having given the `-enable-debug` option to `./configure` and rebuilt completely, to get debugging output.

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

- [termsanddefs](#)
- [tools](#)
- [envvar](#)
- [cpu_mem_bind](#)
- [interoperability](#)
- [threadsafety](#)
- [embed](#)
- [switchfromplpa](#)
- [faq](#)

Make sure to have had a look at those too!

[termsanddefs](#) Terms and Definitions

Object Interesting kind of part of the system, such as a Core, a Cache, a 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.

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.

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 of the same type which have the same parent.

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

Cousin objects Objects of the same type as the current object.

Level Set of objects of the same type.

OS or physical index The index that the operating system (OS) uses to identify the object. This may be completely arbitrary, or it may depend on the BIOS configuration.


Depth Nesting level in the object tree, starting from the 0th object.

Logical index Index to uniquely identify objects of the same type. It expresses proximity in a generic way. This index is always linear and in the range [0, num_objs_same_type_same_level). 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 processor

Processing unit 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 SMT processor.

The following diagram can help to understand the vocabulary of the relationships by showing the example of a machine with two dual core sockets (with no hardware threads); thus, a topology with 4 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.)



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

tools 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.

1.8 `lstopo`

`lstopo` (also known as `hwloc-info` and `hwloc-ls`) displays the hierarchical topology map of the current system. The output may be graphical or textual, and can also be exported to numerous file formats such as PDF, PNG, XML, and others.

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

1.9 `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 sockets 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.

1.10 `hwloc-calc`

`hwloc-calc` is generally used to create bitmap strings to pass to `hwloc-bind`. Although `hwloc-bind` accepts many forms of object specification (i.e., bitmap strings are one of many forms that `hwloc-bind` understands), they can be useful, compact representations in shell scripts, for example.

`hwloc-calc` generates bitmap strings from given hardware objects with the ability to aggregate them, intersect them, and more. `hwloc-calc` generally uses the same syntax than `hwloc-bind`, but multiple instances may be composed to generate complex combinations.

Note that `hwloc-calc` can also generate lists of logical processors or NUMA nodes that are convenient to pass to some external tools such as `taskset` or `numactl`.

1.11 hwloc-distrib

hwloc-distrib generates a set of bitmap strings 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.

1.12 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.

envvar 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.

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\(\)](#). 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.

cpu_mem_bind CPU Binding and Memory Binding

Some OSes 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, and changing the memory binding policy may change the CPU binding of the current thread. This is often not a problem for the application, so by default hwloc will make use of these functions when they provide better binding support.

If the application does not want any CPU binding 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. Conversely, `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, this will thus reduce hwloc's support for binding, so their use is discouraged.

One can however avoid using these flags but still closely control both memory and CPU binding, by allocating memory and touching it, and then changing the CPU binding. The already-really-allocated memory will not be migrated, thus even if the memory binding policy gets changed by the CPU binding order, the effect will have been achieved. On 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 OSes support the notion of a current memory binding policy for the current process but those often still provide a way to allocate data on a given node set. Conversely, some OSes support the notion of a current memory binding policy, and do not permit to allocate data on a given node set without just changing the current policy and allocate the data. Hwloc provides functions that set the current memory binding policies (if supported) as well as functions which allocate memory bound to given node set. By default, it does not use the former to achieve the latter, so that users can use both on OSes where they are both supported, and get both effects at the same time. For convenience, hwloc however also provides the `hwloc_alloc_membind_policy` and `hwloc_alloc_membind_policy_noderset` helpers which are allowed to change the current memory binding policy of the process, in order to achieve memory binding even if that means having to change the current memory binding policy.

interoperability 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.

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 nodemasks and 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_setaffinity().

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.

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.

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.

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.

threadsafety 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 [Create and Destroy Topologies](#)) 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 [Tinker with topologies](#).) may modify the topology significantly by adding objects inside the tree, changing the topology depth, etc. Although references to former objects *may* still be valid after insertion, 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 [Configure Topology Detection](#)) do not modify the topology directly, but they do modify internal structures describing the behavior of the next invocation of `hwloc_topology_load()`. Hence, all of these functions should not be used concurrently.

Note that these functions do not modify the current topology until it is actually reloaded; it is possible to use them while other threads are only read the current topology.

embed 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).

1.13 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 `"/path/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_SUBSTs` 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_SYSTEM` becomes `FOO_HWLOC_OBJ_SYSTEM`. 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`s to build itself, and `AM_CONDITIONALS`s 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.

1.14 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`.

switchfromplpa Switching from PLPA to `hwloc`

Although PLPA and `hwloc` share some of the same ideas, their programming interfaces are quite different. After much debate, it was decided *not* to emulate the PLPA API with `hwloc`'s API because `hwloc`'s API is already far more rich than PLPA's.

More specifically, exploiting modern computing architecture *requires* the flexible functionality provided by the `hwloc` API – the PLPA API is too rigid in its definitions and practices to handle the evolving server hardware landscape (e.g., PLPA only understands cores and sockets; `hwloc` understands a much larger set of hardware objects).

As such, even though it is fully possible to emulate the PLPA API with `hwloc` (e.g., only deal with sockets and cores), and while the documentation below describes how to do this, we encourage any existing PLPA application authors to actually re-think their application in terms of more than just sockets and cores. In short, we encourage you to use the full `hwloc` API to exploit *all* the hardware.

1.15 Topology Context vs. Caching

First, all hwloc functions take a `topology` parameter. This parameter serves as an internal storage for the result of the topology discovery. It replaces PLPA's caching abilities and even lets you manipulate multiple topologies as the same time, if needed.

Thus, all programs should first run `hwloc_topology_init()` and `hwloc_topology_destroy()` as they did `plpa_init()` and `plpa_finalize()` in the past.

1.16 Hierarchy vs. Core@Socket

PLPA was designed to understand only cores and sockets. hwloc offers many more different types of objects (e.g., cores, sockets, hardware threads, NUMA nodes, and others) and stores them within a tree of resources.

To emulate the PLPA model, it is possible to find sockets using functions such as `hwloc_get_obj_by_type()`. Iterating over sockets is also possible using `hwloc_get_next_obj_by_type()`. Then, finding a core within a socket may be done using `hwloc_get_obj_inside_cpuset_by_type()` or `hwloc_get_next_obj_inside_cpuset_by_type()`.

It is also possible to directly find an object "below" another object using `hwloc_get_obj_below_by_type()` (or `hwloc_get_obj_below_array_by_type()`).

1.17 Logical vs. Physical/OS Indexes

hwloc manipulates logical indexes, meaning indexes specified with regard to the ordering of objects in the hwloc-provided hierarchical tree. Physical or OS indexes may be entirely hidden if not strictly required. The reason for this is that physical/OS indexes may change with the OS or with the BIOS version. They may be non-consecutive, multiple objects may have the same physical/OS indexes, making their manipulation tricky and highly non-portable.

Note that hwloc tries very hard to always present a hierarchical tree with the same logical ordering, regardless of physical or OS index ordering.

It is still possible to retrieve physical/OS indexes through the `os_index` field of objects, but such practice should be avoided as much as possible for the reasons described above (except perhaps for prettyprinting / debugging purposes).

HWLOC_OBJ_PU objects are supposed to have different physical/OS indexes since the OS uses them for binding. The `os_index` field of these objects provides the identifier that may be used for such binding, and `hwloc_get_proc_obj_by_os_index()` finds the object associated with a specific OS index.

But as mentioned above, we discourage the use of these conversion methods for actual binding. Instead, hwloc offers its own binding model using the `cpuset` field of ob-

jects. These cpusets may be duplicated, modified, combined, etc. (see `hwloc/bitmap.h` for details) and then passed to `hwloc_set_cpubind()` for binding.

1.18 Counting Specification

PLPA offers a `countspec` parameter to specify whether counting all CPUs, only the online ones or only the offline ones. However, some operating systems do not expose the topology of offline CPUs (i.e., offline CPUs are not reported at all by the OS). Also, some processors may not be visible to the current application due to administrative restrictions. Finally, some processors let you shutdown a single hardware thread in a core, making some of the PLPA features irrelevant.

`hwloc` stores in the hierarchical tree of objects all CPUs that have known topology information. It then provides the applications with several cpusets that contain the list of CPUs that are actually known, that have topology information, that are online, or that are available to the application. These cpusets may be retrieved with `hwloc_topology_get_online_cpuset()` and other similar functions to filter the object that are relevant or not.

faq Frequently Asked Questions

1.19 I do not want `hwloc` to rediscover my enormous machine topology everytime 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. 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 socket).

Chapter 2

Hardware Locality (hwloc) Module Index

2.1 Hardware Locality (hwloc) Modules

Here is a list of all modules:

API version	39
Topology context	40
Object sets	41
Topology Object Types	43
Topology Objects	46
Create and Destroy Topologies	47
Configure Topology Detection	49
Tinker with topologies.	55
Get some Topology Information	57
Retrieve Objects	60
Object/String Conversion	61
CPU binding	64
Memory binding	68
Object Type Helpers	76
Basic Traversal Helpers	77
Finding Objects Inside a CPU set	80
Finding a single Object covering at least CPU set	83
Finding a set of similar Objects covering at least a CPU set	84
Cache-specific Finding Helpers	85
Advanced Traversal Helpers	86
Binding Helpers	88
Cpuset Helpers	90
Nodeset Helpers	92

Conversion between cpuset and nodeset	93
The bitmap API	95
Helpers for manipulating glibc sched affinity	106
Linux-only helpers	107
Helpers for manipulating Linux libnuma unsigned long masks	109
Helpers for manipulating Linux libnuma bitmask	111
Helpers for manipulating Linux libnuma nodemask_t	113
CUDA Driver API Specific Functions	115
CUDA Runtime API Specific Functions	116
OpenFabrics-Specific Functions	117
Myrinet Express-Specific Functions	118

Chapter 3

Hardware Locality (hwloc) Data Structure Index

3.1 Hardware Locality (hwloc) Data Structures

Here are the data structures with brief descriptions:

hwloc_obj (Structure of a topology object)	119
hwloc_obj_attr_u (Object type-specific Attributes)	127
hwloc_obj_attr_u::hwloc_cache_attr_s (Cache-specific Object Attributes) . .	128
hwloc_obj_attr_u::hwloc_group_attr_s (Group-specific Object Attributes) . .	129
hwloc_obj_info_s (Object info)	130
hwloc_obj_memory_s (Object memory)	131
hwloc_obj_memory_s::hwloc_obj_memory_page_type_s (Array of local memory page types, NULL if no local memory and <code>page_types</code> is 0)	133
hwloc_topology_cpubind_support (Flags describing actual PU binding sup- port for this topology)	134
hwloc_topology_discovery_support (Flags describing actual discovery sup- port for this topology)	136
hwloc_topology_membind_support (Flags describing actual memory bind- ing support for this topology)	137
hwloc_topology_support (Set of flags describing actual support for this topology)	140

Chapter 4

Hardware Locality (hwloc) File Index

4.1 Hardware Locality (hwloc) File List

Here is a list of all files with brief descriptions:

bitmap.h (The bitmap API, for use in hwloc itself)	141
cuda.h (Macros to help interaction between hwloc and the CUDA Driver API)	146
cudart.h (Macros to help interaction between hwloc and the CUDA Runtime API)	147
glibc-sched.h (Macros to help interaction between hwloc and glibc scheduling routines)	148
helper.h (High-level hwloc traversal helpers)	149
hwloc.doxy	154
hwloc.h (The hwloc API)	155
linux-libnuma.h (Macros to help interaction between hwloc and Linux libnuma)	164
linux.h (Macros to help interaction between hwloc and Linux)	166
myriexpress.h (Macros to help interaction between hwloc and Myrinet Express)	167
openfabrics-verbs.h (Macros to help interaction between hwloc and OpenFabrics verbs)	168

Chapter 5

Hardware Locality (hwloc) Module Documentation

5.1 API version

Defines

- `#define HWLOC_API_VERSION 0x00010100`
Indicate at build time which hwloc API version is being used.

5.1.1 Define Documentation

5.1.1.1 `#define HWLOC_API_VERSION 0x00010100`

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

5.2 Topology context

Typedefs

- typedef hwloc_topology * [hwloc_topology_t](#)
Topology context.

5.2.1 Typedef Documentation

5.2.1.1 typedef struct hwloc_topology* [hwloc_topology_t](#)

Topology context.

To be initialized with [hwloc_topology_init\(\)](#) and built with [hwloc_topology_load\(\)](#).

5.3 Object sets

Typedefs

- typedef [hwloc_bitmap_t](#) [hwloc_cpuset_t](#)
A CPU set is a bitmap whose bits are set according to CPU physical OS indexes.
- typedef [hwloc_const_bitmap_t](#) [hwloc_const_cpuset_t](#)
A non-modifiable [hwloc_cpuset_t](#).
- 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.
- typedef [hwloc_const_bitmap_t](#) [hwloc_const_nodeset_t](#)
A non-modifiable [hwloc_nodeset_t](#).

5.3.1 Typedef Documentation

5.3.1.1 typedef [hwloc_const_bitmap_t](#) [hwloc_const_cpuset_t](#)

A non-modifiable [hwloc_cpuset_t](#).

5.3.1.2 typedef [hwloc_const_bitmap_t](#) [hwloc_const_nodeset_t](#)

A non-modifiable [hwloc_nodeset_t](#).

5.3.1.3 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](#)).

5.3.1.4 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](#)).

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 [Conversion between cpuset and nodeset](#).

5.4 Topology Object Types

Enumerations

- enum `hwloc_obj_type_t` {
 `HWLOC_OBJ_SYSTEM`, `HWLOC_OBJ_MACHINE`, `HWLOC_OBJ_NODE`,
 `HWLOC_OBJ_SOCKET`,
 `HWLOC_OBJ_CACHE`, `HWLOC_OBJ_CORE`, `HWLOC_OBJ_PU`,
 `HWLOC_OBJ_GROUP`,
 `HWLOC_OBJ_MISC` }
 Type of topology object.
- enum `hwloc_compare_types_e` { `HWLOC_TYPE_UNORDERED` }

Functions

- int `hwloc_compare_types` (`hwloc_obj_type_t` type1, `hwloc_obj_type_t` type2)
 Compare the depth of two object types.

5.4.1 Enumeration Type Documentation

5.4.1.1 enum `hwloc_compare_types_e`

Enumeration values:

`HWLOC_TYPE_UNORDERED` Value returned by `hwloc_compare_types` when types can not be compared.

5.4.1.2 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.

Enumeration values:

`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_NODE NUMA node.

A set of processors around memory which the processors can directly access.

HWLOC_OBJ_SOCKET Socket, physical package, or chip.

In the physical meaning, i.e. that you can add or remove physically.

HWLOC_OBJ_CACHE Data cache.

Can be L1, 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 OSes expose their arbitrary processors aggregation this way. And hwloc may insert such objects to group NUMA nodes according to their distances.

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.

5.4.2 Function Documentation

5.4.2.1 `int hwloc_compare_types (hwloc_obj_type_t type1, hwloc_obj_type_t type2) const`

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 sockets 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 sockets may also contain nodes. This is thus just to be seen as a fallback comparison method.

5.5 Topology Objects

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_obj_info_s](#)
Object info.

Typedefs

- typedef [hwloc_obj](#) * [hwloc_obj_t](#)
Convenience typedef; a pointer to a struct [hwloc_obj](#).

5.5.1 Typedef Documentation

5.5.1.1 typedef struct [hwloc_obj](#)* [hwloc_obj_t](#)

Convenience typedef; a pointer to a struct [hwloc_obj](#).

5.6 Create and Destroy Topologies

Functions

- int `hwloc_topology_init` (`hwloc_topology_t *topologyp`)
Allocate a topology context.
- int `hwloc_topology_load` (`hwloc_topology_t topology`)
Build the actual topology.
- void `hwloc_topology_destroy` (`hwloc_topology_t topology`)
Terminate and free a topology context.
- void `hwloc_topology_check` (`hwloc_topology_t topology`)
Run internal checks on a topology structure.

5.6.1 Function Documentation

5.6.1.1 void `hwloc_topology_check` (`hwloc_topology_t topology`)

Run internal checks on a topology structure.

Parameters:

topology is the topology to be checked

5.6.1.2 void `hwloc_topology_destroy` (`hwloc_topology_t topology`)

Terminate and free a topology context.

Parameters:

topology is the topology to be freed

5.6.1.3 int `hwloc_topology_init` (`hwloc_topology_t *topologyp`)

Allocate a topology context.

Parameters:

→ *topologyp* is assigned a pointer to the new allocated context.

Returns:

0 on success, -1 on error.

5.6.1.4 `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 [Configure Topology Detection](#) routines. No other routine may be called earlier using this topology context.

Parameters:

topology is the topology to be loaded with objects.

Returns:

0 on success, -1 on error.

See also:

[hwlocality_configuration](#)

5.7 Configure Topology Detection

Data Structures

- struct `hwloc_topology_discovery_support`
Flags describing actual discovery support for this topology.
- struct `hwloc_topology_cpubind_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` }
- Flags to be set onto a topology context before load.*

Functions

- int `hwloc_topology_ignore_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Ignore an object type.
- 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.
- int `hwloc_topology_ignore_all_keep_structure` (`hwloc_topology_t` topology)
Ignore all objects that do not bring any structure.
- int `hwloc_topology_set_flags` (`hwloc_topology_t` topology, unsigned long flags)
Set OR'ed flags to non-yet-loaded topology.

- 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.

- int [hwloc_topology_set_pid](#) ([hwloc_topology_t](#) restrict topology, [hwloc_pid_t](#) pid)

Change which pid the topology is viewed from.

- int [hwloc_topology_set_synthetic](#) ([hwloc_topology_t](#) restrict topology, const char *restrict description)

Enable synthetic topology.

- int [hwloc_topology_set_xml](#) ([hwloc_topology_t](#) restrict topology, const char *restrict xmlpath)

Enable XML-file based topology.

- 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.

- const struct [hwloc_topology_support](#) * [hwloc_topology_get_support](#) ([hwloc_topology_t](#) restrict topology)

Retrieve the topology support.

5.7.1 Detailed Description

These 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\(\)](#).

5.7.2 Enumeration Type Documentation

5.7.2.1 enum `hwloc_topology_flags_e`

Flags to be set onto a topology context before load.

Flags should be given to `hwloc_topology_set_flags()`.

Enumeration values:

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 Cpusets and gather all processors and memory nodes, and ignore the fact that some resources may be offline.

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.

5.7.3 Function Documentation

5.7.3.1 `const struct hwloc_topology_support* hwloc_topology_get_support (hwloc_topology_t restrict topology)`

Retrieve the topology support.

5.7.3.2 `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: Each ignored object should have a single children or be the only child of its parent.

5.7.3.3 `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.

5.7.3.4 `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.

5.7.3.5 `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.

5.7.3.6 `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:

For conveniency, 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.

5.7.3.7 `int hwloc_topology_set_pid (hwloc_topology_t restrict topology, hwloc_pid_t pid)`

Change which pid 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.

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

Enable synthetic topology.

Gather topology information from the given `description` which should be a comma 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.

Note:

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

5.7.3.9 `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 `lstopo file.xml`.

Note:

For conveniency, 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.

5.7.3.10 `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.

Gather topology information from the XML memory buffer given at `buffer` and of length `length`.

5.8 Tinker with topologies.

Functions

- void `hwloc_topology_export_xml` (`hwloc_topology_t` topology, const char *xmlpath)

Export the topology into an XML file.

- void `hwloc_topology_export_xmlbuffer` (`hwloc_topology_t` topology, char **xmlbuffer, int *buflen)

Export the topology into a newly-allocated XML memory buffer.

- `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.

- `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 to the topology.

5.8.1 Function Documentation

5.8.1.1 void `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()`.

5.8.1.2 void `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 `xmlFree` later in the caller.

This memory buffer may be loaded later through `hwloc_topology_set_xmlbuffer()`.

5.8.1.3 `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`.

`cpuset` and `name` will be copied.

Returns:

the newly-created object

5.8.1.4 `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 to the topology.

A new MISC object will be created and inserted into the topology at the position given by `parent`.

`name` will be copied.

Returns:

the newly-created object

5.9 Get some Topology Information

Enumerations

- enum `hwloc_get_type_depth_e` { `HWLOC_TYPE_DEPTH_UNKNOWN`, `HWLOC_TYPE_DEPTH_MULTIPLE` }

Functions

- unsigned `hwloc_topology_get_depth` (`hwloc_topology_t` restrict topology)
Get the depth of the hierachical tree of objects.
- int `hwloc_get_type_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type type.
- `hwloc_obj_type_t` `hwloc_get_depth_type` (`hwloc_topology_t` topology, unsigned depth)
Returns the type of objects at depth depth.
- unsigned `hwloc_get_nobjs_by_depth` (`hwloc_topology_t` topology, unsigned depth)
Returns the width of level at depth depth.
- inline int `hwloc_get_nobjs_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the width of level type type.
- int `hwloc_topology_is_thissystem` (`hwloc_topology_t` restrict topology)
Does the topology context come from this system?

5.9.1 Enumeration Type Documentation

5.9.1.1 enum `hwloc_get_type_depth_e`

Enumeration values:

`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.

5.9.2 Function Documentation

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

Returns the type of objects at depth `depth`.

Returns:

-1 if depth `depth` does not exist.

5.9.2.2 `unsigned` `hwloc_get_nbojbs_by_depth` (`hwloc_topology_t topology`, `unsigned depth`)

Returns the width of level at depth `depth`.

5.9.2.3 `inline int` `hwloc_get_nbojbs_by_type` (`hwloc_topology_t topology`, `hwloc_obj_type_t type`) [`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.

5.9.2.4 `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()`.

5.9.2.5 `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.

5.9.2.6 `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).

5.10 Retrieve Objects

Functions

- [hwloc_obj_t hwloc_get_obj_by_depth](#) ([hwloc_topology_t](#) topology, unsigned depth, unsigned idx)

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

- inline [hwloc_obj_t hwloc_get_obj_by_type](#) ([hwloc_topology_t](#) topology, [hwloc_obj_type_t](#) type, unsigned idx)

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

5.10.1 Function Documentation

5.10.1.1 [hwloc_obj_t hwloc_get_obj_by_depth](#) ([hwloc_topology_t](#) topology, unsigned depth, unsigned idx)

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

5.10.1.2 inline [hwloc_obj_t hwloc_get_obj_by_type](#) ([hwloc_topology_t](#) topology, [hwloc_obj_type_t](#) type, unsigned idx) [static]

Returns the topology object at index `index` 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 their caller may fallback to [hwloc_get_obj_by_depth\(\)](#).

5.11 Object/String Conversion

Functions

- `const char * hwloc_obj_type_string (hwloc_obj_type_t type)`
Return a stringified topology object type.
- `hwloc_obj_type_t hwloc_obj_type_of_string (const char *string)`
Return an object type from the string.
- `int hwloc_obj_type_snprintf (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.
- `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.
- `int hwloc_obj_snprintf (char *restrict string, size_t size, hwloc_topology_t topology, hwloc_obj_t obj, const char *restrict indexprefix, int verbose)`
Stringify a given topology object into a human-readable form.
- `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.
- `inline char * hwloc_obj_get_info_by_name (hwloc_obj_t obj, const char *name)`
Search the given key name in object infos and return the corresponding value.

5.11.1 Function Documentation

5.11.1.1 `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.

Returns:

how many characters were actually written (not including the ending `\0`), or -1 on error.

5.11.1.2 `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.

Returns:

how many characters were actually written (not including the ending `\0`).

5.11.1.3 `inline char* hwloc_obj_get_info_by_name (hwloc_obj_t obj, const char *name) [static]`

Search the given key name in object infos and return the corresponding value.

Returns:

NULL if no such key exists.

5.11.1.4 `int hwloc_obj_snprintf (char *restrict string, size_t size, hwloc_topology_t topology, hwloc_obj_t obj, const char *restrict indexprefix, int verbose)`

Stringify a given topology object into a human-readable form.

Note:

This function is deprecated in favor of `hwloc_obj_type_snprintf()` and `hwloc_obj_attr_snprintf()` since it is not very flexible and only prints physical/OS indexes.

Fill string `string` up to `size` characters with the description of topology object `obj` in topology `topology`.

If `verbose` is set, a longer description is used. Otherwise a short description is used.

`indexprefix` is used to prefix the `os_index` attribute number of the object in the description. If NULL, the `#` character is used.

Returns:

how many characters were actually written (not including the ending `\0`), or -1 on error.

5.11.1.5 `hwloc_obj_type_t hwloc_obj_type_of_string (const char *string)`

Return an object type from the string.

Returns:

-1 if unrecognized.

5.11.1.6 `int hwloc_obj_type_snprintf (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.

It differs from `hwloc_obj_type_string()` because it prints type attributes such as cache depth.

Returns:

how many characters were actually written (not including the ending `\0`), or -1 on error.

5.11.1.7 `const char* hwloc_obj_type_string (hwloc_obj_type_t type) const`

Return a stringified topology object type.

5.12 CPU binding

Enumerations

- enum `hwloc_cpubind_flags_t` { `HWLOC_CPUBIND_PROCESS`, `HWLOC_CPUBIND_THREAD`, `HWLOC_CPUBIND_STRICT`, `HWLOC_CPUBIND_NOMEMBIND` }

Process/Thread binding flags.

Functions

- int `hwloc_set_cpubind` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, int flags)
Bind current process or thread on cpus given in bitmap set.
- int `hwloc_get_cpubind` (`hwloc_topology_t` topology, `hwloc_cpuset_t` set, int flags)
Get current process or thread binding.
- 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 bitmap set.
- int `hwloc_get_proc_cpubind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_cpuset_t` set, int flags)
Get the current binding of process pid.
- int `hwloc_set_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` tid, `hwloc_const_cpuset_t` set, int flags)
Bind a thread tid on cpus given in bitmap set.
- int `hwloc_get_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` tid, `hwloc_cpuset_t` set, int flags)
Get the current binding of thread tid.

5.12.1 Detailed Description

It is often useful to call `hwloc_bitmap_singlify()` first so that a single CPU remains in the set. This way, the process will not even migrate between different CPUs. Some OSes also only support that kind of binding.

Note:

Some OSes do not provide all ways to bind processes, threads, etc and the corresponding binding functions may fail. -1 is returned and `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).

The most portable version that should be preferred over the others, whenever possible, is

```
hwloc_set_cpubind(topology, set, 0),
```

as it just binds the current program, assuming it is monothread, or

```
hwloc_set_cpubind(topology, set, HWLOC_CPUBIND_THREAD),
```

which binds the current thread of the current program (which may be multithreaded).

Note:

To unbind, just call the binding function with either a full `cpuset` or a `cpuset` equal to the system `cpuset`.

On some OSes, CPU binding may have effects on memory binding, see [HWLOC_CPUBIND_NOMEMBIND](#)

5.12.2 Enumeration Type Documentation

5.12.2.1 enum [hwloc_cpubind_flags_t](#)

Process/Thread binding flags.

These flags can be used to refine the binding policy.

The default (0) is to bind the current process, assumed to be mono-thread, in a non-strict way. This is the most portable way to bind as all OSes usually provide it.

Note:

Not all systems support all kinds of binding.

Enumeration values:

HWLOC_CPUBIND_PROCESS Bind all threads of the current (possibly) multithreaded 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, OSes 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 OSes and implementations, strict binding may not be possible (implementation reason) or not allowed (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 OSes, 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.

5.12.3 Function Documentation

5.12.3.1 `int hwloc_get_cpupbind (hwloc_topology_t topology, hwloc_cpuset_t set, int flags)`

Get current process or thread binding.

5.12.3.2 `int hwloc_get_proc_cpupbind (hwloc_topology_t topology, hwloc_pid_t pid, hwloc_cpuset_t set, int flags)`

Get the current binding of process `pid`.

Note:

`hwloc_pid_t` is `pid_t` on unix platforms, and `HANDLE` on native Windows platforms

`HWLOC_CPUBIND_THREAD` can not be used in `flags`.

5.12.3.3 `int hwloc_get_thread_cpupbind (hwloc_topology_t topology, hwloc_thread_t tid, hwloc_cpuset_t set, int flags)`

Get the current 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.

5.12.3.4 int hwloc_set_cpubind (hwloc_topology_t topology, hwloc_const_cpuset_t set, int flags)

Bind current process or thread on cpus given in 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

5.12.3.5 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 bitmap set.

Note:

hwloc_pid_t is pid_t on unix platforms, and HANDLE on native Windows platforms
HWLOC_CPUBIND_THREAD can not be used in flags.

5.12.3.6 int hwloc_set_thread_cpubind (hwloc_topology_t topology, hwloc_thread_t tid, hwloc_const_cpuset_t set, int flags)

Bind a thread tid on cpus given in 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.

5.13 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`
}

Memory binding policy.
- enum `hwloc_membind_flags_t` {
`HWLOC_MEMBIND_PROCESS`, `HWLOC_MEMBIND_THREAD`,
`HWLOC_MEMBIND_STRICT`, `HWLOC_MEMBIND_MIGRATE`,
`HWLOC_MEMBIND_NOCPUBIND` }

Memory binding flags.

Functions

- int `hwloc_set_membind_nodset` (`hwloc_topology_t` topology, `hwloc_const_nodset_t` nodset, `hwloc_membind_policy_t` policy, int flags)

Bind current process memory on the given nodset nodset.
- int `hwloc_set_membind` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `hwloc_membind_policy_t` policy, int flags)

Bind current process memory on memory nodes near the given cpuset cpuset.
- int `hwloc_get_membind_nodset` (`hwloc_topology_t` topology, `hwloc_nodset_t` nodset, `hwloc_membind_policy_t` *policy, int flags)

Get current process memory binding in nodset nodset.
- int `hwloc_get_membind` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `hwloc_membind_policy_t` *policy, int flags)

Get current process memory binding in cpuset cpuset.
- int `hwloc_set_proc_membind_nodset` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_nodset_t` nodset, `hwloc_membind_policy_t` policy, int flags)

Bind given process memory on the given nodset nodset.

- int `hwloc_set_proc_mbind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_cpuset_t` cpuset, `hwloc_mbind_policy_t` policy, int flags)
Bind given process memory on memory nodes near the given cpuset cpuset.
- int `hwloc_get_proc_mbind_nodeset` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_nodeset_t` nodeset, `hwloc_mbind_policy_t` *policy, int flags)
Get current process memory binding in nodeset nodeset.
- int `hwloc_get_proc_mbind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_cpuset_t` cpuset, `hwloc_mbind_policy_t` *policy, int flags)
Get current process memory binding in cpuset cpuset.
- 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 some memory range on the given nodeset nodeset.
- int `hwloc_set_area_mbind` (`hwloc_topology_t` topology, const void *addr, size_t len, `hwloc_const_cpuset_t` cpuset, `hwloc_mbind_policy_t` policy, int flags)
Bind some memory range on memory nodes near the given cpuset cpuset.
- int `hwloc_get_area_mbind_nodeset` (`hwloc_topology_t` topology, const void *addr, size_t len, `hwloc_nodeset_t` nodeset, `hwloc_mbind_policy_t` *policy, int flags)
Get some memory range memory binding in nodeset nodeset.
- int `hwloc_get_area_mbind` (`hwloc_topology_t` topology, const void *addr, size_t len, `hwloc_cpuset_t` cpuset, `hwloc_mbind_policy_t` *policy, int flags)
Get some memory range memory binding in cpuset cpuset.
- void * `hwloc_alloc` (`hwloc_topology_t` topology, size_t len)
Allocate some memory.
- 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 the given nodeset nodeset.
- void * `hwloc_alloc_mbind` (`hwloc_topology_t` topology, size_t len, `hwloc_const_cpuset_t` cpuset, `hwloc_mbind_policy_t` policy, int flags)
Allocate some memory on memory nodes near the given cpuset cpuset.

- int [hwloc_free](#) ([hwloc_topology_t](#) topology, void *addr, size_t len)
Free some memory allocated by [hwloc_alloc\(\)](#) or [hwloc_alloc_membind\(\)](#).

5.13.1 Detailed Description

Note:

Not all OSes support all ways to bind existing allocated memory (migration), future memory allocation, explicit memory allocation, etc. and the corresponding binding functions may fail. -1 is returned and 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 NUMA node).

The most portable version that should be preferred over the others, whenever possible, is

```
hwloc_alloc_membind_policy(topology, size, set, HWLOC_MEMBIND_DEFAULT, 0),
```

which will try to allocate new data bound to the given set, possibly by changing the current memory binding policy, or at worse allocate memory without binding it at all. Since HWLOC_MEMBIND_STRICT is not given, this will even not fail unless a mere malloc() itself would fail, i.e. ENOMEM.

Each binding is available with a CPU set argument or a NUMA memory node set argument. The name of the latter ends with _nodeset. It is also possible to convert between CPU set and node set using [hwloc_cpuset_to_nodeset](#) or [hwloc_cpuset_from_nodeset](#).

Note:

On some OSes, memory binding may have effects on CPU binding, see [HWLOC_MEMBIND_NOCPUBIND](#)

5.13.2 Enumeration Type Documentation

5.13.2.1 enum [hwloc_membind_flags_t](#)

Memory binding flags.

These flags can be used to refine the binding policy.

Note:

Not all systems support all kinds of binding.

Enumeration values:

HWLOC_MEMBIND_PROCESS Set policy for all threads of the current (possibly multithreaded) process.

HWLOC_MEMBIND_THREAD Set policy for the current thread of the current process.

HWLOC_MEMBIND_STRICT Request strict binding from the OS. The function will fail if the binding can not be completely enforced.

HWLOC_MEMBIND_MIGRATE Migrate existing allocated memory.
If memory can not be migrated and the STRICT flag is passed, an error will be returned.

HWLOC_MEMBIND_NOCPUBIND Avoid any effect on CPU binding.
On some OSes, some memory binding function would also bind the application on the corresponding CPUs. 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 on CPUs. This will however reduce the support of memory bindings, i.e. potentially return ENOSYS in some cases.

5.13.2.2 enum [hwloc_membind_policy_t](#)

Memory binding policy.

These can be used to choose the binding policy.

Note that not all systems support all kinds of binding.

Enumeration values:

HWLOC_MEMBIND_DEFAULT Reset the memory allocation policy to the system default.

HWLOC_MEMBIND_FIRSTTOUCH Allocate memory on the given nodes, but preferably on the node where the first accessor is running.

HWLOC_MEMBIND_BIND Allocate memory on the given nodes.

HWLOC_MEMBIND_INTERLEAVE Allocate memory on the given nodes in a round-robin manner.

HWLOC_MEMBIND_REPLICATE Replicate memory on the given nodes.

HWLOC_MEMBIND_NEXTTOUCH On next touch of existing allocated memory, migrate it to the node where the memory reference happened.

5.13.3 Function Documentation

5.13.3.1 void* [hwloc_alloc](#) ([hwloc_topology_t topology](#), [size_t len](#))

Allocate some memory.

This is equivalent to malloc(), except it tries to allocated page-aligned memory from the OS.

Note:

The allocated memory should be freed with `hwloc_free()`.

5.13.3.2 `void* hwloc_alloc_membind (hwloc_topology_t topology, size_t len, hwloc_const_cpuset_t cpuset, hwloc_membind_policy_t policy, int flags)`

Allocate some memory on memory nodes near the given cpuset `cpuset`.

Returns:

-1 with `errno` set to `ENOSYS` if the action is not supported and `HWLOC_-MEMBIND_STRICT` is given
 -1 with `errno` set to `EXDEV` if the binding cannot be enforced and `HWLOC_-MEMBIND_STRICT` is given

Note:

The allocated memory should be freed with `hwloc_free()`.

5.13.3.3 `void* hwloc_alloc_membind_nodeset (hwloc_topology_t topology, size_t len, hwloc_const_nodeset_t nodeset, hwloc_membind_policy_t policy, int flags)`

Allocate some memory on the given nodeset `nodeset`.

Returns:

-1 with `errno` set to `ENOSYS` if the action is not supported and `HWLOC_-MEMBIND_STRICT` is given
 -1 with `errno` set to `EXDEV` if the binding cannot be enforced and `HWLOC_-MEMBIND_STRICT` is given

Note:

The allocated memory should be freed with `hwloc_free()`.

5.13.3.4 `int hwloc_free (hwloc_topology_t topology, void * addr, size_t len)`

Free some memory allocated by `hwloc_alloc()` or `hwloc_alloc_membind()`.

5.13.3.5 `int hwloc_get_area_membind (hwloc_topology_t topology, const void * addr, size_t len, hwloc_cpuset_t cpuset, hwloc_membind_policy_t * policy, int flags)`

Get some memory range memory binding in cpuset `cpuset`.

5.13.3.6 `int hwloc_get_area_mbind_nodeset (hwloc_topology_t topology, const void * addr, size_t len, hwloc_nodeset_t nodeset, hwloc_mbind_policy_t * policy, int flags)`

Get some memory range memory binding in nodeset nodeset.

5.13.3.7 `int hwloc_get_mbind (hwloc_topology_t topology, hwloc_cpuset_t cpuset, hwloc_mbind_policy_t * policy, int flags)`

Get current process memory binding in cpuset cpuset.

5.13.3.8 `int hwloc_get_mbind_nodeset (hwloc_topology_t topology, hwloc_nodeset_t nodeset, hwloc_mbind_policy_t * policy, int flags)`

Get current process memory binding in nodeset nodeset.

5.13.3.9 `int hwloc_get_proc_mbind (hwloc_topology_t topology, hwloc_pid_t pid, hwloc_cpuset_t cpuset, hwloc_mbind_policy_t * policy, int flags)`

Get current process memory binding in cpuset cpuset.

5.13.3.10 `int hwloc_get_proc_mbind_nodeset (hwloc_topology_t topology, hwloc_pid_t pid, hwloc_nodeset_t nodeset, hwloc_mbind_policy_t * policy, int flags)`

Get current process memory binding in nodeset nodeset.

5.13.3.11 `int hwloc_set_area_mbind (hwloc_topology_t topology, const void * addr, size_t len, hwloc_const_cpuset_t cpuset, hwloc_mbind_policy_t policy, int flags)`

Bind some memory range on memory nodes near the given cpuset 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

5.13.3.12 `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 some memory range on the given nodeset `nodeset`.

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

5.13.3.13 `int hwloc_set_mbind(hwloc_topology_t topology, hwloc_const_cpuset_t cpuset, hwloc_mbind_policy_t policy, int flags)`

Bind current process memory on memory nodes near the given cpuset `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

5.13.3.14 `int hwloc_set_mbind_nodeset(hwloc_topology_t topology, hwloc_const_nodeset_t nodeset, hwloc_mbind_policy_t policy, int flags)`

Bind current process memory on the given nodeset `nodeset`.

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

5.13.3.15 `int hwloc_set_proc_mbind(hwloc_topology_t topology, hwloc_pid_t pid, hwloc_const_cpuset_t cpuset, hwloc_mbind_policy_t policy, int flags)`

Bind given process memory on memory nodes near the given cpuset `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

5.13.3.16 `int hwloc_set_proc_mbind_nodeset (hwloc_topology_t topology, hwloc_pid_t pid, hwloc_const_nodeset_t nodeset, hwloc_mbind_policy_t policy, int flags)`

Bind given process memory on the given nodeset nodeset.

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

5.14 Object Type Helpers

Functions

- inline int `hwloc_get_type_or_below_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type `type` or below.
- inline int `hwloc_get_type_or_above_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type `type` or above.

5.14.1 Function Documentation

5.14.1.1 inline int `hwloc_get_type_or_above_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type) [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`.

5.14.1.2 inline int `hwloc_get_type_or_below_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type) [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`.

5.15 Basic Traversal Helpers

Functions

- inline `hwloc_obj_t hwloc_get_root_obj (hwloc_topology_t topology)`
Returns the top-object of the topology-tree.
- inline `hwloc_obj_t hwloc_get_ancestor_obj_by_depth (hwloc_topology_t topology, unsigned depth, hwloc_obj_t obj)`
Returns the ancestor object of obj at depth depth.
- inline `hwloc_obj_t hwloc_get_ancestor_obj_by_type (hwloc_topology_t topology, hwloc_obj_type_t type, hwloc_obj_t obj)`
Returns the ancestor object of obj with type type.
- inline `hwloc_obj_t hwloc_get_next_obj_by_depth (hwloc_topology_t topology, unsigned depth, hwloc_obj_t prev)`
Returns the next object at depth depth.
- inline `hwloc_obj_t hwloc_get_next_obj_by_type (hwloc_topology_t topology, hwloc_obj_type_t type, hwloc_obj_t prev)`
Returns the next object of type type.
- inline `hwloc_obj_t hwloc_get_pu_obj_by_os_index (hwloc_topology_t topology, unsigned os_index)`
Returns the object of type `HWLOC_OBJ_PU` with `os_index`.
- inline `hwloc_obj_t hwloc_get_next_child (hwloc_topology_t topology, hwloc_obj_t parent, hwloc_obj_t prev)`
Return the next child.
- inline `hwloc_obj_t hwloc_get_common_ancestor_obj (hwloc_topology_t topology, hwloc_obj_t obj1, hwloc_obj_t obj2)`
Returns the common parent object to objects `obj1` and `obj2`.
- inline `int hwloc_obj_is_in_subtree (hwloc_topology_t topology, hwloc_obj_t obj, hwloc_obj_t subtree_root)`
Returns true if `_obj_` is inside the subtree beginning with `subtree_root`.

5.15.1 Function Documentation

5.15.1.1 inline `hwloc_obj_t` `hwloc_get_ancestor_obj_by_depth`
(`hwloc_topology_t` topology , unsigned *depth*, `hwloc_obj_t` obj)
[static]

Returns the ancestor object of obj at depth depth.

5.15.1.2 inline `hwloc_obj_t` `hwloc_get_ancestor_obj_by_type`
(`hwloc_topology_t` topology , `hwloc_obj_type_t` type, `hwloc_obj_t` obj)
[static]

Returns the ancestor object of obj with type type.

5.15.1.3 inline `hwloc_obj_t` `hwloc_get_common_ancestor_obj`
(`hwloc_topology_t` topology , `hwloc_obj_t` obj1, `hwloc_obj_t` obj2)
[static]

Returns the common parent object to objects lvl1 and lvl2.

5.15.1.4 inline `hwloc_obj_t` `hwloc_get_next_child` (`hwloc_topology_t` topology ,
`hwloc_obj_t` parent, `hwloc_obj_t` prev) [static]

Return the next child.

If prev is NULL, return the first child.

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

Returns the next object at depth depth.

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

5.15.1.6 inline `hwloc_obj_t` `hwloc_get_next_obj_by_type` (`hwloc_topology_t`
topology, `hwloc_obj_type_t` type, `hwloc_obj_t` prev) [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()`.

5.15.1.7 inline `hwloc_obj_t` `hwloc_get_pu_obj_by_os_index` (`hwloc_topology_t topology`, unsigned `os_index`) [static]

Returns the object of type `HWLOC_OBJ_PU` with `os_index`.

Note:

The `os_index` field of object should most of the times only be used for pretty-printing purpose. Type `HWLOC_OBJ_PU` is the only case where `os_index` could actually be useful, when manually binding to processors. However, using CPU sets to hide this complexity should often be preferred.

5.15.1.8 inline `hwloc_obj_t` `hwloc_get_root_obj` (`hwloc_topology_t topology`) [static]

Returns the top-object of the topology-tree.

Its type is typically `HWLOC_OBJ_MACHINE` but it could be different for complex topologies. This function replaces the old deprecated `hwloc_get_system_obj()`.

5.15.1.9 inline `int` `hwloc_obj_is_in_subtree` (`hwloc_topology_t topology`, `hwloc_obj_t obj`, `hwloc_obj_t subtree_root`) [static]

Returns true if `_obj_` is inside the subtree beginning with `subtree_root`.

5.16 Finding Objects Inside a CPU set

Functions

- inline `hwloc_obj_t hwloc_get_first_largest_obj_inside_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set)

Get the first largest object included in the given cpuset set.

- 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.

- inline `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)

Return the next object at depth depth included in CPU set set.

- inline `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)

Return the next object of type type included in CPU set set.

- inline `hwloc_obj_t hwloc_get_obj_inside_cpuset_by_depth` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, unsigned depth, unsigned idx)

Return the index -th object at depth depth included in CPU set set.

- inline `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)

Return the idx -th object of type type included in CPU set set.

- inline unsigned `hwloc_get_nbobjs_inside_cpuset_by_depth` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, unsigned depth)

Return the number of objects at depth depth included in CPU set set.

- inline int `hwloc_get_nbobjs_inside_cpuset_by_type` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, `hwloc_obj_type_t` type)

Return the number of objects of type type included in CPU set set.

5.16.1 Function Documentation

5.16.1.1 `inline hwloc_obj_t hwloc_get_first_largest_obj_inside_cpuset`
(`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`) [`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.

5.16.1.2 `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`.

5.16.1.3 `inline unsigned hwloc_get_nbobjs_inside_cpuset_by_depth`
(`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `unsigned depth`)
[`static`]

Return the number of objects at depth `depth` included in CPU set `set`.

5.16.1.4 `inline int hwloc_get_nbobjs_inside_cpuset_by_type` (`hwloc_topology_t`
`topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`)
[`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.

5.16.1.5 `inline 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`]

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`.

5.16.1.6 `inline 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`]

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()`.

5.16.1.7 `inline hwloc_obj_t hwloc_get_obj_inside_cpuset_by_depth`
(`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `unsigned depth`,
`unsigned idx`) [`static`]

Return the `idx`-th object at depth `depth` included in CPU set `set`.

5.16.1.8 `inline 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`]

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()`.

5.17 Finding a single Object covering at least CPU set

Functions

- inline `hwloc_obj_t hwloc_get_child_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_t parent`)
Get the child covering at least CPU set set.
- inline `hwloc_obj_t hwloc_get_obj_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`)
Get the lowest object covering at least CPU set set.

5.17.1 Function Documentation

5.17.1.1 inline `hwloc_obj_t hwloc_get_child_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_t parent`) [`static`]

Get the child covering at least CPU set set.

Returns:

NULL if no child matches or if `set` is empty.

5.17.1.2 inline `hwloc_obj_t hwloc_get_obj_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`) [`static`]

Get the lowest object covering at least CPU set set.

Returns:

NULL if no object matches or if `set` is empty.

5.18 Finding a set of similar Objects covering at least a CPU set

Functions

- inline `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`)

Iterate through same-depth objects covering at least CPU set `set`.

- inline `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`)

Iterate through same-type objects covering at least CPU set `set`.

5.18.1 Function Documentation

5.18.1.1 inline `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]

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`.

5.18.1.2 inline `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`) [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.

5.19 Cache-specific Finding Helpers

Functions

- inline `hwloc_obj_t hwloc_get_cache_covering_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`)
Get the first cache covering a cpuset set.
- inline `hwloc_obj_t hwloc_get_shared_cache_covering_obj` (`hwloc_topology_t topology`, `hwloc_obj_t obj`)
Get the first cache shared between an object and somebody else.

5.19.1 Function Documentation

5.19.1.1 inline `hwloc_obj_t hwloc_get_cache_covering_cpuset`
(`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`) [static]

Get the first cache covering a cpuset set.

Returns:

NULL if no cache matches

5.19.1.2 inline `hwloc_obj_t hwloc_get_shared_cache_covering_obj`
(`hwloc_topology_t topology`, `hwloc_obj_t obj`) [static]

Get the first cache shared between an object and somebody else.

Returns:

NULL if no cache matches

5.20 Advanced Traversal Helpers

Functions

- 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.

- inline `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)

Find an object below another object, both specified by types and indexes.

- inline `hwloc_obj_t hwloc_get_obj_below_array_by_type` (`hwloc_topology_t` topology, int nr, `hwloc_obj_type_t *typev`, unsigned *idxv)

Find an object below a chain of objects specified by types and indexes.

5.20.1 Function Documentation

5.20.1.1 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.

5.20.1.2 inline `hwloc_obj_t hwloc_get_obj_below_array_by_type` (`hwloc_topology_t` topology, int nr, `hwloc_obj_type_t *typev`, unsigned *idxv) [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 index couple in the arrays, look under the previously found object to find the index-th object of the given type. Indexes are specified within the parent, not withing the entire system.

For instance, if `nr` is 3, `typev` contains `NODE`, `SOCKET` and `CORE`, and `idxv` contains 0, 1 and 2, return the third core object below the second socket below the first NUMA node.

5.20.1.3 inline `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]

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 index `idx1`. Then look below this object and find another object of type `type2` and index `idx2`. Indexes are specified within the parent, not within the entire system.

For instance, if `type1` is `SOCKET`, `idx1` is 2, `type2` is `CORE` and `idx2` is 3, return the fourth core object below the third socket.

5.21 Binding Helpers

Functions

- inline void `hwloc_distributev` (`hwloc_topology_t` topology, `hwloc_obj_t` *root, unsigned n_roots, `hwloc_cpuset_t` *cpuset, unsigned n, unsigned until)

Distribute n items over the topology under root.

- inline void `hwloc_distribute` (`hwloc_topology_t` topology, `hwloc_obj_t` root, `hwloc_cpuset_t` *cpuset, unsigned n, unsigned until)
- inline 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)

Allocate some memory on the given nodeset nodeset.

- inline void * `hwloc_alloc_membind_policy` (`hwloc_topology_t` topology, `size_t` len, `hwloc_const_cpuset_t` cpuset, `hwloc_membind_policy_t` policy, int flags)

Allocate some memory on the memory nodes near given cpuset cpuset.

5.21.1 Function Documentation

- 5.21.1.1** inline void* `hwloc_alloc_membind_policy` (`hwloc_topology_t` topology, `size_t` len, `hwloc_const_cpuset_t` cpuset, `hwloc_membind_policy_t` policy, int flags) [static]

Allocate some memory on the memory nodes near given cpuset cpuset.

This is similar to `hwloc_alloc_membind_policy_nodeset`, but for a given cpuset.

- 5.21.1.2** inline 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]

Allocate some memory on the given nodeset nodeset.

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.

5.21.1.3 `inline void hwloc_distribute (hwloc_topology_t topology, hwloc_obj_t root, hwloc_cpuset_t * cpuset, unsigned n, unsigned until) [static]`

5.21.1.4 `inline void hwloc_distributev (hwloc_topology_t topology, hwloc_obj_t * roots, unsigned n_roots, hwloc_cpuset_t * cpuset, unsigned n, unsigned until) [static]`

Distribute `n` items over the topology under `root`.

This is the same as `hwloc_distribute`, but takes an array of roots instead of just one root.

5.22 Cpuset Helpers

Functions

- inline `hwloc_const_cpuset_t hwloc_topology_get_complete_cpuset (hwloc_topology_t topology)`
- inline `hwloc_const_cpuset_t hwloc_topology_get_topology_cpuset (hwloc_topology_t topology)`
- inline `hwloc_const_cpuset_t hwloc_topology_get_online_cpuset (hwloc_topology_t topology)`
Get online CPU set.
- inline `hwloc_const_cpuset_t hwloc_topology_get_allowed_cpuset (hwloc_topology_t topology)`
Get allowed CPU set.

5.22.1 Function Documentation

5.22.1.1 inline `hwloc_const_cpuset_t hwloc_topology_get_allowed_cpuset (hwloc_topology_t topology)` [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_cpuset_dup` must be used to obtain a local copy.

5.22.1.2 inline `hwloc_const_cpuset_t hwloc_topology_get_complete_cpuset (hwloc_topology_t topology)` [static]

5.22.1.3 inline `hwloc_const_cpuset_t hwloc_topology_get_online_cpuset (hwloc_topology_t topology)` [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_cpuset_dup` must be used to obtain a local copy.

5.22.1.4 `inline hwloc_const_cpuset_t hwloc_topology_get_topology_cpuset`
`(hwloc_topology_t topology) [static]`

5.23 Nodeset Helpers

Functions

- inline `hwloc_const_nodeset_t hwloc_topology_get_complete_nodeset (hwloc_topology_t topology)`
- inline `hwloc_const_nodeset_t hwloc_topology_get_topology_nodeset (hwloc_topology_t topology)`
- inline `hwloc_const_nodeset_t hwloc_topology_get_allowed_nodeset (hwloc_topology_t topology)`

Get allowed node set.

5.23.1 Function Documentation

5.23.1.1 inline `hwloc_const_nodeset_t hwloc_topology_get_allowed_nodeset (hwloc_topology_t topology)` [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_nodeset_dup` must be used to obtain a local copy.

5.23.1.2 inline `hwloc_const_nodeset_t hwloc_topology_get_complete_nodeset (hwloc_topology_t topology)` [static]

5.23.1.3 inline `hwloc_const_nodeset_t hwloc_topology_get_topology_nodeset (hwloc_topology_t topology)` [static]

5.24 Conversion between cpuset and nodeset

Functions

- inline void `hwloc_cpuset_to_nodeset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `hwloc_nodeset_t` nodeset)
Convert a CPU set into a NUMA node set and handle non-NUMA cases.
- inline void `hwloc_cpuset_to_nodeset_strict` (`struct hwloc_topology *`topology, `hwloc_const_cpuset_t` cpuset, `hwloc_nodeset_t` nodeset)
Convert a CPU set into a NUMA node set without handling non-NUMA cases.
- inline void `hwloc_cpuset_from_nodeset` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `hwloc_const_nodeset_t` nodeset)
Convert a NUMA node set into a CPU set and handle non-NUMA cases.
- inline void `hwloc_cpuset_from_nodeset_strict` (`struct hwloc_topology *`topology, `hwloc_cpuset_t` cpuset, `hwloc_const_nodeset_t` nodeset)
Convert a NUMA node set into a CPU set without handling non-NUMA cases.

5.24.1 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.

5.24.2 Function Documentation

5.24.2.1 inline void `hwloc_cpuset_from_nodeset` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `hwloc_const_nodeset_t` nodeset) [static]

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.

5.24.2.2 `inline void hwloc_cpuset_from_nodeset_strict (struct hwloc_topology * topology, hwloc_cpuset_t cpuset, hwloc_const_nodeset_t nodeset) [static]`

Convert a NUMA node set into a CPU set without handling non-NUMA cases.

This is the strict variant of `hwloc_cpuset_from_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 `cpuset`.

5.24.2.3 `inline void hwloc_cpuset_to_nodeset (hwloc_topology_t topology, hwloc_const_cpuset_t cpuset, hwloc_nodeset_t nodeset) [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.

5.24.2.4 `inline void hwloc_cpuset_to_nodeset_strict (struct hwloc_topology * topology, hwloc_const_cpuset_t cpuset, hwloc_nodeset_t nodeset) [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`.

5.25 The bitmap API

Defines

- #define `hwloc_bitmap_foreach_begin`(id, bitmap)
Loop macro iterating on bitmap `bitmap`.
- #define `hwloc_bitmap_foreach_end`()
End of loop. Needs a terminating `';`.

Typedefs

- typedef `hwloc_bitmap_s` * `hwloc_bitmap_t`
Set of bits represented as an opaque pointer to an internal bitmap.
- typedef const struct `hwloc_bitmap_s` * `hwloc_const_bitmap_t`

Functions

- `hwloc_bitmap_t` `hwloc_bitmap_alloc` (void)
Allocate a new empty bitmap.
- `hwloc_bitmap_t` `hwloc_bitmap_alloc_full` (void)
Allocate a new full bitmap.
- void `hwloc_bitmap_free` (`hwloc_bitmap_t` bitmap)
Free bitmap `bitmap`.
- `hwloc_bitmap_t` `hwloc_bitmap_dup` (`hwloc_const_bitmap_t` bitmap)
Duplicate bitmap `bitmap` by allocating a new bitmap and copying `bitmap` contents.
- 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`.
- int `hwloc_bitmap_sprintf` (char *restrict buf, size_t buflen, `hwloc_const_bitmap_t` bitmap)
Stringify a bitmap.
- int `hwloc_bitmap_asprintf` (char **strp, `hwloc_const_bitmap_t` bitmap)
Stringify a bitmap into a newly allocated string.

- int `hwloc_bitmap_sscanf` (`hwloc_bitmap_t` bitmap, const char *restrict string)
Parse a bitmap string and stores it in bitmap `bitmap`.
- int `hwloc_bitmap_taskset_sprintf` (char *restrict buf, size_t buflen, `hwloc_const_bitmap_t` bitmap)
Stringify a bitmap in the taskset-specific format.
- int `hwloc_bitmap_taskset_asprintf` (char **strp, `hwloc_const_bitmap_t` bitmap)
Stringify a bitmap into a newly allocated taskset-specific string.
- 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`.
- void `hwloc_bitmap_zero` (`hwloc_bitmap_t` bitmap)
Empty the bitmap `bitmap`.
- 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).
- void `hwloc_bitmap_only` (`hwloc_bitmap_t` bitmap, unsigned id)
Empty the bitmap `bitmap` and add bit `id`.
- void `hwloc_bitmap_allbut` (`hwloc_bitmap_t` bitmap, unsigned id)
Fill the bitmap and clear the index `id`.
- void `hwloc_bitmap_from_ulong` (`hwloc_bitmap_t` bitmap, unsigned long mask)
Setup bitmap `bitmap` from unsigned long `mask`.
- 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.
- void `hwloc_bitmap_set` (`hwloc_bitmap_t` bitmap, unsigned id)
Add index `id` in bitmap `bitmap`.
- void `hwloc_bitmap_set_range` (`hwloc_bitmap_t` bitmap, unsigned begin, unsigned end)
Add indexes from `begin` to `end` in bitmap `bitmap`.

- 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`.
- void `hwloc_bitmap_clr` (`hwloc_bitmap_t` bitmap, unsigned `id`)
Remove index `id` from bitmap `bitmap`.
- void `hwloc_bitmap_clr_range` (`hwloc_bitmap_t` bitmap, unsigned `begin`, unsigned `end`)
Remove index from `begin` to `end` in bitmap `bitmap`.
- void `hwloc_bitmap_singlify` (`hwloc_bitmap_t` bitmap)
Keep a single index among those set in bitmap `bitmap`.
- unsigned long `hwloc_bitmap_to_ulong` (`hwloc_const_bitmap_t` bitmap)
Convert the beginning part of bitmap `bitmap` into unsigned long `mask`.
- 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`.
- int `hwloc_bitmap_isset` (`hwloc_const_bitmap_t` bitmap, unsigned `id`)
Test whether index `id` is part of bitmap `bitmap`.
- int `hwloc_bitmap_iszero` (`hwloc_const_bitmap_t` bitmap)
Test whether bitmap `bitmap` is empty.
- int `hwloc_bitmap_isfull` (`hwloc_const_bitmap_t` bitmap)
Test whether bitmap `bitmap` is completely full.
- int `hwloc_bitmap_first` (`hwloc_const_bitmap_t` bitmap)
Compute the first index (least significant bit) in bitmap `bitmap`.
- int `hwloc_bitmap_next` (`hwloc_const_bitmap_t` bitmap, unsigned `prev`)
Compute the next index in bitmap `bitmap` which is after index `prev`.
- int `hwloc_bitmap_last` (`hwloc_const_bitmap_t` bitmap)
Compute the last index (most significant bit) in bitmap `bitmap`.
- 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).

- 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.
- void `hwloc_bitmap_and` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
And bitmaps bitmap1 and bitmap2 and store the result in bitmap res.
- void `hwloc_bitmap_andnot` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
And bitmap bitmap1 and the negation of bitmap2 and store the result in bitmap res.
- 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.
- void `hwloc_bitmap_not` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap)
Negate bitmap bitmap and store the result in bitmap res.
- int `hwloc_bitmap_intersects` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Test whether bitmaps bitmap1 and bitmap2 intersect.
- 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.
- int `hwloc_bitmap_isequal` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Test whether bitmap bitmap1 is equal to bitmap bitmap2.
- int `hwloc_bitmap_compare_first` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Compare bitmaps bitmap1 and bitmap2 using their lowest index.
- int `hwloc_bitmap_compare` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Compare bitmaps bitmap1 and bitmap2 using their highest index.

5.25.1 Detailed Description

For use in hwloc itself, a `hwloc_bitmap_t` usually represents a set of objects, typically logical processors or memory nodes, indexed by OS physical number.

A bitmap may be infinite.

5.25.2 Define Documentation

5.25.2.1 `#define hwloc_bitmap_foreach_begin(id, bitmap)`

Loop macro iterating on bitmap `bitmap`.

`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 that 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 infinite.

5.25.2.2 `#define hwloc_bitmap_foreach_end()`

End of loop. Needs a terminating `';`.

See also:

[hwloc_bitmap_foreach_begin](#)

5.25.3 Typedef Documentation

5.25.3.1 `typedef struct hwloc_bitmap_s* hwloc_bitmap_t`

Set of bits represented as an opaque pointer to an internal bitmap.

5.25.3.2 `typedef const struct hwloc_bitmap_s* hwloc_const_bitmap_t`

5.25.4 Function Documentation

5.25.4.1 `void hwloc_bitmap_allbut(hwloc_bitmap_t bitmap, unsigned id)`

Fill the bitmap and clear the index `id`.

5.25.4.2 `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()`.

5.25.4.3 `hwloc_bitmap_t hwloc_bitmap_alloc_full (void)`

Allocate a new full bitmap.

5.25.4.4 `void hwloc_bitmap_and (hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2)`

And bitmaps `bitmap1` and `bitmap2` and store the result in bitmap `res`.

5.25.4.5 `void hwloc_bitmap_andnot (hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2)`

And bitmap `bitmap1` and the negation of `bitmap2` and store the result in bitmap `res`.

5.25.4.6 `int hwloc_bitmap_asprintf (char ** strp, hwloc_const_bitmap_t bitmap)`

Stringify a bitmap into a newly allocated string.

Returns:

the number of character that were actually written (not including the ending `\0`).

5.25.4.7 `void hwloc_bitmap_clr (hwloc_bitmap_t bitmap, unsigned id)`

Remove index `id` from bitmap `bitmap`.

5.25.4.8 `void hwloc_bitmap_clr_range (hwloc_bitmap_t bitmap, unsigned begin, unsigned end)`

Remove index from `begin` to `end` in bitmap `bitmap`.

5.25.4.9 `int hwloc_bitmap_compare (hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2)`

Compare bitmaps `bitmap1` and `bitmap2` using their highest index.

Higher most significant bit is higher. The empty bitmap is considered lower than anything.

5.25.4.10 `int hwloc_bitmap_compare_first (hwloc_const_bitmap_t bitmap1, hwloc_const_bitmap_t bitmap2)`

Compare bitmaps `bitmap1` and `bitmap2` using their lowest index.

Smaller least significant bit is smaller. The empty bitmap is considered higher than anything.

5.25.4.11 `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`.

5.25.4.12 `hwloc_bitmap_t hwloc_bitmap_dup (hwloc_const_bitmap_t bitmap)`

Duplicate bitmap `bitmap` by allocating a new bitmap and copying `bitmap` contents.

5.25.4.13 `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).

5.25.4.14 `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.

5.25.4.15 `void hwloc_bitmap_free (hwloc_bitmap_t bitmap)`

Free bitmap `bitmap`.

5.25.4.16 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.

5.25.4.17 void `hwloc_bitmap_from_ulong` ([hwloc_bitmap_t](#) *bitmap*, unsigned long *mask*)

Setup bitmap `bitmap` from unsigned long `mask`.

5.25.4.18 int `hwloc_bitmap_intersects` ([hwloc_const_bitmap_t](#) *bitmap1*, [hwloc_const_bitmap_t](#) *bitmap2*)

Test whether bitmaps `bitmap1` and `bitmap2` intersect.

5.25.4.19 int `hwloc_bitmap_isequal` ([hwloc_const_bitmap_t](#) *bitmap1*, [hwloc_const_bitmap_t](#) *bitmap2*)

Test whether bitmap `bitmap1` is equal to bitmap `bitmap2`.

5.25.4.20 int `hwloc_bitmap_isfull` ([hwloc_const_bitmap_t](#) *bitmap*)

Test whether bitmap `bitmap` is completely full.

5.25.4.21 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`.

5.25.4.22 int `hwloc_bitmap_isset` ([hwloc_const_bitmap_t](#) *bitmap*, unsigned *id*)

Test whether index `id` is part of bitmap `bitmap`.

5.25.4.23 int `hwloc_bitmap_iszero` ([hwloc_const_bitmap_t](#) *bitmap*)

Test whether bitmap `bitmap` is empty.

5.25.4.24 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 bitmap, or if the index bitmap is infinite.

5.25.4.25 `int hwloc_bitmap_next (hwloc_const_bitmap_t bitmap, unsigned prev)`

Compute the next index in bitmap `bitmap` which is after index `prev`.

Returns:

-1 if no index with higher index is bitmap.

5.25.4.26 `void hwloc_bitmap_not (hwloc_bitmap_t res, hwloc_const_bitmap_t bitmap)`

Negate bitmap `bitmap` and store the result in bitmap `res`.

5.25.4.27 `void hwloc_bitmap_only (hwloc_bitmap_t bitmap, unsigned id)`

Empty the bitmap `bitmap` and add bit `id`.

5.25.4.28 `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`.

5.25.4.29 `void hwloc_bitmap_set (hwloc_bitmap_t bitmap, unsigned id)`

Add index `id` in bitmap `bitmap`.

5.25.4.30 `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`.

5.25.4.31 `void hwloc_bitmap_set_range (hwloc_bitmap_t bitmap, unsigned begin, unsigned end)`

Add indexes from `begin` to `end` in bitmap `bitmap`.

5.25.4.32 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.

5.25.4.33 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`.

Returns:

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

5.25.4.34 int hwloc_bitmap_sscanf (hwloc_bitmap_t bitmap, const char *restrict string)

Parse a bitmap string and stores it in bitmap `bitmap`.

Must start and end with a digit.

5.25.4.35 int hwloc_bitmap_taskset_asprintf (char ** strp, hwloc_const_bitmap_t bitmap)

Stringify a bitmap into a newly allocated taskset-specific string.

5.25.4.36 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`.

5.25.4.37 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`.

5.25.4.38 `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.

5.25.4.39 `unsigned long hwloc_bitmap_to_ulong (hwloc_const_bitmap_t bitmap)`

Convert the beginning part of bitmap *bitmap* into unsigned long mask.

5.25.4.40 `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.

5.25.4.41 `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*.

5.25.4.42 `void hwloc_bitmap_zero (hwloc_bitmap_t bitmap)`

Empty the bitmap *bitmap*.

5.26 Helpers for manipulating glibc sched affinity

Functions

- inline int `hwloc_cpuset_to_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_const_cpuset_t` hwlocset, `cpu_set_t` *schedset, `size_t` schedsetsize)
Convert hwloc CPU set toposet into glibc sched affinity CPU set schedset.
- inline int `hwloc_cpuset_from_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_cpuset_t` hwlocset, `const cpu_set_t` *schedset, `size_t` schedsetsize)
Convert glibc sched affinity CPU set schedset into hwloc CPU set.

5.26.1 Function Documentation

5.26.1.1 inline int `hwloc_cpuset_from_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_cpuset_t` hwlocset, `const cpu_set_t` * schedset, `size_t` schedsetsize) [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`

5.26.1.2 inline int `hwloc_cpuset_to_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_const_cpuset_t` hwlocset, `cpu_set_t` * schedset, `size_t` schedsetsize) [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`

5.27 Linux-only helpers

Functions

- int `hwloc_linux_parse_cpumap_file` (FILE *file, `hwloc_cpuset_t` set)
Convert a linux kernel cpumap file `file` into hwloc CPU set.
- 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`.
- int `hwloc_linux_get_tid_cpupbind` (`hwloc_topology_t` topology, pid_t tid, `hwloc_cpuset_t` set)
Get the current binding of thread `tid`.

5.27.1 Detailed Description

This includes helpers for manipulating linux kernel cpumap files, and hwloc equivalents of the Linux `sched_setaffinity` and `sched_getaffinity` system calls.

5.27.2 Function Documentation

5.27.2.1 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_setaffinity` system call, but uses a hwloc cpuset.

5.27.2.2 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.

5.27.2.3 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.

5.28 Helpers for manipulating Linux libnuma unsigned long masks

Functions

- inline int `hwloc_cpuset_to_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, unsigned long *mask, unsigned long *maxnode)

Convert hwloc CPU set cpuset into the array of unsigned long mask.

- inline int `hwloc_node`set_to_linux_libnuma_ulongs (`hwloc_topology_t` topology, `hwloc_const_node`set_t nodeset, unsigned long *mask, unsigned long *maxnode)

Convert hwloc NUMA node set nodeset into the array of unsigned long mask.

- inline int `hwloc_cpuset_from_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const unsigned long *mask, unsigned long maxnode)

Convert the array of unsigned long mask into hwloc CPU set.

- inline int `hwloc_node`set_from_linux_libnuma_ulongs (`hwloc_topology_t` topology, `hwloc_node`set_t nodeset, const unsigned long *mask, unsigned long maxnode)

Convert the array of unsigned long mask into hwloc NUMA node set.

5.28.1 Function Documentation

5.28.1.1 inline int `hwloc_cpuset_from_linux_libnuma_ulong`s (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const unsigned long * mask, unsigned long maxnode) [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).

5.28.1.2 `inline int hwloc_cpuset_to_linux_libnuma_ulongs (hwloc_topology_t topology, hwloc_const_cpuset_t cpuset, unsigned long * mask, unsigned long * maxnode) [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.

5.28.1.3 `inline int hwloc_nodeset_from_linux_libnuma_ulongs (hwloc_topology_t topology, hwloc_nodeset_t nodeset, const unsigned long * mask, unsigned long maxnode) [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).

5.28.1.4 `inline int hwloc_nodeset_to_linux_libnuma_ulongs (hwloc_topology_t topology, hwloc_const_nodeset_t nodeset, unsigned long * mask, unsigned long * maxnode) [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.

5.29 Helpers for manipulating Linux libnuma bitmask

Functions

- inline struct bitmask * `hwloc_cpuset_to_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset)
Convert hwloc CPU set cpuset into the returned libnuma bitmask.
- inline struct bitmask * `hwloc_nodeset_to_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_const_nodeset_t` nodeset)
Convert hwloc NUMA node set nodeset into the returned libnuma bitmask.
- inline int `hwloc_cpuset_from_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const struct bitmask *bitmask)
Convert libnuma bitmask bitmask into hwloc CPU set cpuset.
- inline int `hwloc_nodeset_from_linux_libnuma_bitmask` (`hwloc_topology_t` topology, `hwloc_nodeset_t` nodeset, const struct bitmask *bitmask)
Convert libnuma bitmask bitmask into hwloc NUMA node set nodeset.

5.29.1 Function Documentation

5.29.1.1 inline int hwloc_cpuset_from_linux_libnuma_bitmask
(`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, const struct bitmask * *bitmask*) [`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.

5.29.1.2 inline struct bitmask* hwloc_cpuset_to_linux_libnuma_bitmask
(`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset) [`static`]

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.

5.29.1.3 `inline int hwloc_nodeset_from_linux_libnuma_bitmask`
(`hwloc_topology_t topology`, `hwloc_nodeset_t nodeset`, `const struct`
`bitmask * bitmask`) [`static`]

Convert libnuma bitmask `bitmask` into hwloc NUMA node set `nodeset`.

This function may be used after calling many `numa_` functions that use a struct bitmask as an output parameter.

5.29.1.4 `inline struct bitmask* hwloc_nodeset_to_linux_libnuma_bitmask`
(`hwloc_topology_t topology`, `hwloc_const_nodeset_t nodeset`)
[`static`]

Convert hwloc NUMA node set `nodeset` 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.

5.30 Helpers for manipulating Linux libnuma nodemask_t

Functions

- inline int `hwloc_cpuset_to_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `nodemask_t *nodemask`)
Convert hwloc CPU set cpuset into libnuma nodemask nodemask.
- inline int `hwloc_nodeseq_to_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_const_nodeseq_t` nodeseq, `nodemask_t *nodemask`)
Convert hwloc NUMA node set nodeseq into libnuma nodemask nodemask.
- inline int `hwloc_cpuset_from_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `const nodemask_t *nodemask`)
Convert libnuma nodemask nodemask into hwloc CPU set cpuset.
- inline int `hwloc_nodeseq_from_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_nodeseq_t` nodeseq, `const nodemask_t *nodemask`)
Convert libnuma nodemask nodemask into hwloc NUMA node set nodeseq.

5.30.1 Function Documentation

5.30.1.1 inline int `hwloc_cpuset_from_linux_libnuma_nodemask`
(`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `const nodemask_t *nodemask`) [static]

Convert libnuma nodemask nodemask into hwloc CPU set cpuset.

This function may be used before calling some old libnuma functions that use a `nodemask_t` as an output parameter.

5.30.1.2 inline int `hwloc_cpuset_to_linux_libnuma_nodemask`
(`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `nodemask_t *nodemask`) [static]

Convert hwloc CPU set cpuset into libnuma nodemask nodemask.

This function may be used before calling some old libnuma functions that use a `nodemask_t` as an input parameter.

5.30.1.3 `inline int hwloc_nodese`*t_from_linux_libnuma_nodemask*
(`hwloc_topology_t topology`, `hwloc_nodese`*t nodeset*, `const`
`nodemask_t * nodemask`) [`static`]

Convert libnuma nodemask `nodemask` into hwloc NUMA node set `nodeset`.

This function may be used before calling some old libnuma functions that use a `nodemask_t` as an output parameter.

5.30.1.4 `inline int hwloc_nodese`*t_to_linux_libnuma_nodemask*
(`hwloc_topology_t topology`, `hwloc_const_nodese`*t nodeset*,
`nodemask_t * nodemask`) [`static`]

Convert hwloc NUMA node set `nodeset` into libnuma nodemask `nodemask`.

This function may be used before calling some old libnuma functions that use a `nodemask_t` as an input parameter.

5.31 CUDA Driver API Specific Functions

Functions

- inline int `hwloc_cuda_get_device_cpuset` (`hwloc_topology_t` topology , CUdevice `cudevice`, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device `cudevice`.

5.31.1 Function Documentation

5.31.1.1 inline int `hwloc_cuda_get_device_cpuset` (`hwloc_topology_t` topology , CUdevice `cudevice`, `hwloc_cpuset_t` set) [static]

Get the CPU set of logical processors that are physically close to device `cudevice`.

For the given CUDA Driver API device `cudevice`, read the corresponding kernel-provided `cpumap` file and return the corresponding CPU set. This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full `cpuset`.

5.32 CUDA Runtime API Specific Functions

Functions

- inline int `hwloc_cudart_get_device_cpuset` (`hwloc_topology_t` topology , int device, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device `cudevice`.

5.32.1 Function Documentation

5.32.1.1 inline int `hwloc_cudart_get_device_cpuset` (`hwloc_topology_t` topology , int device, `hwloc_cpuset_t` set) [static]

Get the CPU set of logical processors that are physically close to device `cudevice`.

For the given CUDA Runtime API device `cudevice`, read the corresponding kernel-provided `cpumap` file and return the corresponding CPU set. This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full `cpuset`.

5.33 OpenFabrics-Specific Functions

Functions

- inline int `hwloc_ibv_get_device_cpuset` (`hwloc_topology_t` topology , struct `ibv_device *ibdev`, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device `ibdev`.

5.33.1 Function Documentation

- 5.33.1.1** inline int `hwloc_ibv_get_device_cpuset` (`hwloc_topology_t` topology , struct `ibv_device *ibdev`, `hwloc_cpuset_t` set) [static]

Get the CPU set of logical processors that are physically close to device `ibdev`.

For the given OpenFabrics device `ibdev`, read the corresponding kernel-provided `cpumap` file and return the corresponding CPU set. This function is currently only implemented in a meaningful way for Linux; other systems will simply get a full `cpuset`.

5.34 Myrinet Express-Specific Functions

Functions

- inline int `hwloc_mx_board_get_device_cpuset` (`hwloc_topology_t` topology, unsigned id, `hwloc_cpuset_t` set)
Get the CPU set of logical processors that are physically close the MX board id.
- inline int `hwloc_mx_endpoint_get_device_cpuset` (`hwloc_topology_t` topology, `mx_endpoint_t` endpoint, `hwloc_cpuset_t` set)
Get the CPU set of logical processors that are physically close to endpoint endpoint.

5.34.1 Function Documentation

5.34.1.1 inline int `hwloc_mx_board_get_device_cpuset` (`hwloc_topology_t topology`, unsigned `id`, `hwloc_cpuset_t set`) [static]

Get the CPU set of logical processors that are physically close the MX board id.

For the given Myrinet Express board index `id`, read the OS-provided NUMA node and return the corresponding CPU set.

5.34.1.2 inline int `hwloc_mx_endpoint_get_device_cpuset` (`hwloc_topology_t topology`, `mx_endpoint_t endpoint`, `hwloc_cpuset_t set`) [static]

Get the CPU set of logical processors that are physically close to endpoint `endpoint`.

For the given Myrinet Express endpoint `endpoint`, read the OS-provided NUMA node and return the corresponding CPU set.

Chapter 6

Hardware Locality (hwloc) Data Structure Documentation

6.1 hwloc_obj Struct Reference

Structure of a topology object.

```
#include <hwloc.h>
```

Data Fields

- [hwloc_obj_type_t](#) type
Type of object.
- unsigned [os_index](#)
OS-provided physical index number.
- char * [name](#)
Object description if any.
- [hwloc_obj_memory_s](#) memory
Memory attributes.
- [hwloc_obj_attr_u](#) * attr
Object type-specific Attributes, may be NULL if no attribute value was found.
- unsigned [depth](#)
Vertical index in the hierarchy.

- unsigned `logical_index`
Horizontal index in the whole list of similar objects, could be a "cousin_rank" since it's the rank within the "cousin" list below.
- signed `os_level`
OS-provided physical level, -1 if unknown or meaningless.
- `hwloc_obj * next_cousin`
Next object of same type.
- `hwloc_obj * prev_cousin`
Previous object of same type.
- `hwloc_obj * parent`
Parent, NULL if root (system object).
- unsigned `sibling_rank`
Index in parent's `children[]` array.
- `hwloc_obj * next_sibling`
Next object below the same parent.
- `hwloc_obj * prev_sibling`
Previous object below the same parent.
- unsigned `arity`
Number of children.
- `hwloc_obj ** children`
Children, `children[0]`.
- `hwloc_obj * first_child`
First child.
- `hwloc_obj * last_child`
Last child.
- void * `userdata`
Application-given private data pointer, initialized to NULL, use it as you wish.
- `hwloc_cpuset_t cpuset`

CPUs covered by this object.

- [hwloc_cpuset_t complete_cpuset](#)
The complete CPU set of logical processors of this object,.
- [hwloc_cpuset_t online_cpuset](#)
The CPU set of online logical processors.
- [hwloc_cpuset_t allowed_cpuset](#)
The CPU set of allowed logical processors.
- [hwloc_nodeset_t nodeset](#)
NUMA nodes covered by this object or containing this object.
- [hwloc_nodeset_t complete_nodeset](#)
The complete NUMA node set of this object,.
- [hwloc_nodeset_t allowed_nodeset](#)
The set of allowed NUMA memory nodes.
- [hwloc_obj_info_s * infos](#)
Array of stringified info type=name.
- unsigned [infos_count](#)
Size of infos array.

6.1.1 Detailed Description

Structure of a topology object.

Applications mustn't modify any field except `::userdata` .

6.1.2 Field Documentation

6.1.2.1 [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.

6.1.2.2 `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.

6.1.2.3 `unsigned hwloc_obj::arity`

Number of children.

6.1.2.4 `union hwloc_obj_attr_u* hwloc_obj::attr`

Object type-specific Attributes, may be `NULL` if no attribute value was found.

6.1.2.5 `struct hwloc_obj hwloc_obj::children`**

Children, `children[0`.

`arity - 1]`

6.1.2.6 `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 the 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.

6.1.2.7 `hwloc_noderset_t hwloc_obj::complete_noderset`

The complete NUMA node set of this object,.

This includes not only the same as the `noderset` field, but also the 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 `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_noderset` is full.

Note:

Its value must not be changed, `hwloc_bitmap_dup` must be used instead.

6.1.2.8 `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.

6.1.2.9 `unsigned hwloc_obj::depth`

Vertical index in the hierarchy.

6.1.2.10 `struct hwloc_obj* hwloc_obj::first_child`

First child.

6.1.2.11 `struct hwloc_obj_info_s* hwloc_obj::infos`

Array of stringified info `type=name`.

6.1.2.12 `unsigned hwloc_obj::infos_count`

Size of infos array.

6.1.2.13 `struct hwloc_obj* hwloc_obj::last_child`

Last child.

6.1.2.14 `unsigned hwloc_obj::logical_index`

Horizontal index in the whole list of similar objects, could be a "cousin_rank" since it's the rank within the "cousin" list below.

6.1.2.15 `struct hwloc_obj_memory_s hwloc_obj::memory`

Memory attributes.

6.1.2.16 `char* hwloc_obj::name`

Object description if any.

6.1.2.17 `struct hwloc_obj* hwloc_obj::next_cousin`

Next object of same type.

6.1.2.18 `struct hwloc_obj* hwloc_obj::next_sibling`

Next object below the same parent.

6.1.2.19 `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 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 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.

6.1.2.20 `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.

6.1.2.21 `unsigned hwloc_obj::os_index`

OS-provided physical index number.

6.1.2.22 `signed hwloc_obj::os_level`

OS-provided physical level, -1 if unknown or meaningless.

6.1.2.23 `struct hwloc_obj* hwloc_obj::parent`

Parent, NULL if root (system object).

6.1.2.24 `struct hwloc_obj* hwloc_obj::prev_cousin`

Previous object of same type.

6.1.2.25 `struct hwloc_obj* hwloc_obj::prev_sibling`

Previous object below the same parent.

6.1.2.26 `unsigned hwloc_obj::sibling_rank`

Index in parent's `children[]` array.

6.1.2.27 [hwloc_obj_type_t hwloc_obj::type](#)

Type of object.

6.1.2.28 **void*** [hwloc_obj::userdata](#)

Application-given private data pointer, initialized to `NULL`, use it as you wish.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.2 hwloc_obj_attr_u Union Reference

Object type-specific Attributes.

```
#include <hwloc.h>
```

Data Fields

- [hwloc_obj_attr_u::hwloc_cache_attr_s cache](#)
Cache-specific Object Attributes.
- [hwloc_obj_attr_u::hwloc_group_attr_s group](#)
Group-specific Object Attributes.

6.2.1 Detailed Description

Object type-specific Attributes.

6.2.2 Field Documentation

6.2.2.1 `struct hwloc_obj_attr_u::hwloc_cache_attr_s hwloc_obj_attr_u::cache`

Cache-specific Object Attributes.

6.2.2.2 `struct hwloc_obj_attr_u::hwloc_group_attr_s hwloc_obj_attr_u::group`

Group-specific Object Attributes.

The documentation for this union was generated from the following file:

- [hwloc.h](#)

6.3 hwloc_obj_attr_u::hwloc_cache_attr_s Struct Reference

Cache-specific Object Attributes.

```
#include <hwloc.h>
```

Data Fields

- `uint64_t size`
Size of cache in bytes.
- `unsigned depth`
Depth of cache.
- `unsigned linesize`
Cache-line size in bytes.

6.3.1 Detailed Description

Cache-specific Object Attributes.

6.3.2 Field Documentation

6.3.2.1 `unsigned hwloc_obj_attr_u::hwloc_cache_attr_s::depth`

Depth of cache.

6.3.2.2 `unsigned hwloc_obj_attr_u::hwloc_cache_attr_s::linesize`

Cache-line size in bytes.

6.3.2.3 `uint64_t hwloc_obj_attr_u::hwloc_cache_attr_s::size`

Size of cache in bytes.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.4 hwloc_obj_attr_u::hwloc_group_attr_s Struct Reference

Group-specific Object Attributes.

```
#include <hwloc.h>
```

Data Fields

- unsigned [depth](#)
Depth of group object.

6.4.1 Detailed Description

Group-specific Object Attributes.

6.4.2 Field Documentation

6.4.2.1 unsigned [hwloc_obj_attr_u::hwloc_group_attr_s::depth](#)

Depth of group object.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.5 hwloc_obj_info_s Struct Reference

Object info.

```
#include <hwloc.h>
```

Data Fields

- `char * name`
Info name.
- `char * value`
Info value.

6.5.1 Detailed Description

Object info.

6.5.2 Field Documentation

6.5.2.1 `char* hwloc_obj_info_s::name`

Info name.

6.5.2.2 `char* hwloc_obj_info_s::value`

Info value.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.6 hwloc_obj_memory_s Struct Reference

Object memory.

```
#include <hwloc.h>
```

Data Fields

- `uint64_t total_memory`
Total memory (in bytes) in this object and its children.
- `uint64_t local_memory`
Local memory (in bytes).
- `unsigned page_types_len`
Size of array `page_types`.
- `hwloc_obj_memory_s::hwloc_obj_memory_page_type_s * page_types`
Array of local memory page types, NULL if no local memory and `page_types` is 0.

6.6.1 Detailed Description

Object memory.

6.6.2 Field Documentation

6.6.2.1 `uint64_t hwloc_obj_memory_s::local_memory`

Local memory (in bytes).

6.6.2.2 `struct hwloc_obj_memory_s::hwloc_obj_memory_page_type_s * hwloc_obj_memory_s::page_types`

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.

6.6.2.3 `unsigned hwloc_obj_memory_s::page_types_len`

Size of array `page_types`.

6.6.2.4 `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](#)

6.7 hwloc_obj_memory_s::hwloc_obj_memory_page_type_s Struct Reference

6.7 hwloc_obj_memory_s::hwloc_obj_memory_page_type_s Struct Reference

Array of local memory page types, NULL if no local memory and `page_types` is 0.

```
#include <hwloc.h>
```

Data Fields

- [uint64_t size](#)
Size of pages.
- [uint64_t count](#)
Number of pages of this size.

6.7.1 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.

6.7.2 Field Documentation

6.7.2.1 [uint64_t hwloc_obj_memory_s::hwloc_obj_memory_page_type_s::count](#)

Number of pages of this size.

6.7.2.2 [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](#)

6.8 hwloc_topology_cpupbind_support Struct Reference

Flags describing actual PU binding support for this topology.

```
#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](#)

6.8.1 Detailed Description

Flags describing actual PU binding support for this topology.

6.8.2 Field Documentation

6.8.2.1 unsigned char [hwloc_topology_cpupbind_support::get_proc_cpupbind](#)

Getting the binding of a whole given process is supported.

6.8.2.2 unsigned char [hwloc_topology_cpupbind_support::get_thisproc_cpupbind](#)

Getting the binding of the whole current process is supported.

6.8.2.3 unsigned char [hwloc_topology_cpupbind_support::get_thisthread_cpupbind](#)

Getting the binding of the current thread only is supported.

6.8.2.4 unsigned char [hwloc_topology_cpupbind_support::get_thread_cpupbind](#)

Getting the binding of a given thread only is supported.

6.8.2.5 unsigned char [hwloc_topology_cpuid_support::set_proc_cpuid](#)

Binding a whole given process is supported.

6.8.2.6 unsigned char [hwloc_topology_cpuid_support::set_thisproc_cpuid](#)

Binding the whole current process is supported.

6.8.2.7 unsigned char [hwloc_topology_cpuid_support::set_thisthread_cpuid](#)

Binding the current thread only is supported.

6.8.2.8 unsigned char [hwloc_topology_cpuid_support::set_thread_cpuid](#)

Binding a given thread only is supported.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.9 hwloc_topology_discovery_support Struct Reference

Flags describing actual discovery support for this topology.

```
#include <hwloc.h>
```

Data Fields

- unsigned char [pu](#)
Detecting the number of PU objects is supported.

6.9.1 Detailed Description

Flags describing actual discovery support for this topology.

6.9.2 Field Documentation

6.9.2.1 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](#)

6.10 hwloc_topology_membind_support Struct Reference

Flags describing actual memory binding support for this topology.

```
#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](#)

6.10.1 Detailed Description

Flags describing actual memory binding support for this topology.

6.10.2 Field Documentation

6.10.2.1 unsigned char [hwloc_topology_membind_support::alloc_membind](#)

Allocating a bound memory area is supported.

6.10.2.2 unsigned char [hwloc_topology_membind_support::bind_membind](#)

Bind policy is supported.

6.10.2.3 unsigned char `hwloc_topology_membind_support::firsttouch_membind`

First-touch policy is supported.

6.10.2.4 unsigned char `hwloc_topology_membind_support::get_area_membind`

Getting the binding of a given memory area is supported.

6.10.2.5 unsigned char `hwloc_topology_membind_support::get_proc_membind`

Getting the binding of a whole given process is supported.

6.10.2.6 unsigned char `hwloc_topology_membind_support::get_thisproc_membind`

Getting the binding of the whole current process is supported.

6.10.2.7 unsigned char `hwloc_topology_membind_support::get_thisthread_membind`

Getting the binding of the current thread only is supported.

6.10.2.8 unsigned char `hwloc_topology_membind_support::interleave_membind`

Interleave policy is supported.

6.10.2.9 unsigned char `hwloc_topology_membind_support::migrate_membind`

Migration flags is supported.

6.10.2.10 unsigned char `hwloc_topology_membind_support::nexttouch_membind`

Next-touch migration policy is supported.

6.10.2.11 unsigned char [hwloc_topology_membind_support::replicate_membind](#)

Replication policy is supported.

6.10.2.12 unsigned char [hwloc_topology_membind_support::set_area_membind](#)

Binding a given memory area is supported.

6.10.2.13 unsigned char [hwloc_topology_membind_support::set_proc_membind](#)

Binding a whole given process is supported.

6.10.2.14 unsigned char [hwloc_topology_membind_support::set_thisproc_membind](#)

Binding the whole current process is supported.

6.10.2.15 unsigned char [hwloc_topology_membind_support::set_thisthread_membind](#)

Binding the current thread only is supported.

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

6.11 hwloc_topology_support Struct Reference

Set of flags describing actual support for this topology.

```
#include <hwloc.h>
```

Data Fields

- [hwloc_topology_discovery_support](#) * [discovery](#)
- [hwloc_topology_cpubind_support](#) * [cpubind](#)
- [hwloc_topology_membind_support](#) * [membind](#)

6.11.1 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.

6.11.2 Field Documentation

6.11.2.1 [struct hwloc_topology_cpubind_support*](#)
[hwloc_topology_support::cpubind](#)

6.11.2.2 [struct hwloc_topology_discovery_support*](#)
[hwloc_topology_support::discovery](#)

6.11.2.3 [struct hwloc_topology_membind_support*](#)
[hwloc_topology_support::membind](#)

The documentation for this struct was generated from the following file:

- [hwloc.h](#)

Chapter 7

Hardware Locality (hwloc) File Documentation

7.1 `bitmap.h` File Reference

The bitmap API, for use in hwloc itself.

```
#include <hwloc/config.h>
#include <assert.h>
```

Defines

- #define `hwloc_bitmap_foreach_begin(id, bitmap)`
Loop macro iterating on bitmap `bitmap`.
- #define `hwloc_bitmap_foreach_end()`
End of loop. Needs a terminating `';`.

Typedefs

- typedef `hwloc_bitmap_s * hwloc_bitmap_t`
Set of bits represented as an opaque pointer to an internal bitmap.
- typedef `const struct hwloc_bitmap_s * hwloc_const_bitmap_t`

Functions

- [hwloc_bitmap_t hwloc_bitmap_alloc](#) (void)
Allocate a new empty bitmap.
- [hwloc_bitmap_t hwloc_bitmap_alloc_full](#) (void)
Allocate a new full bitmap.
- void [hwloc_bitmap_free](#) ([hwloc_bitmap_t](#) bitmap)
Free bitmap `bitmap`.
- [hwloc_bitmap_t hwloc_bitmap_dup](#) ([hwloc_const_bitmap_t](#) bitmap)
Duplicate bitmap `bitmap` by allocating a new bitmap and copying `bitmap` contents.
- 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`.
- int [hwloc_bitmap_snprintf](#) (char *restrict buf, size_t buflen, [hwloc_const_bitmap_t](#) bitmap)
Stringify a bitmap.
- int [hwloc_bitmap_asprintf](#) (char **strp, [hwloc_const_bitmap_t](#) bitmap)
Stringify a bitmap into a newly allocated string.
- int [hwloc_bitmap_sscanf](#) ([hwloc_bitmap_t](#) bitmap, const char *restrict string)
Parse a bitmap string and stores it in bitmap `bitmap`.
- int [hwloc_bitmap_taskset_snprintf](#) (char *restrict buf, size_t buflen, [hwloc_const_bitmap_t](#) bitmap)
Stringify a bitmap in the taskset-specific format.
- int [hwloc_bitmap_taskset_asprintf](#) (char **strp, [hwloc_const_bitmap_t](#) bitmap)
Stringify a bitmap into a newly allocated taskset-specific string.
- 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`.
- void [hwloc_bitmap_zero](#) ([hwloc_bitmap_t](#) bitmap)
Empty the bitmap `bitmap`.

- 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).
- void `hwloc_bitmap_only` (`hwloc_bitmap_t` bitmap, unsigned id)
Empty the bitmap `bitmap` and add bit `id`.
- void `hwloc_bitmap_allbut` (`hwloc_bitmap_t` bitmap, unsigned id)
Fill the bitmap and clear the index `id`.
- void `hwloc_bitmap_from_ulong` (`hwloc_bitmap_t` bitmap, unsigned long mask)
Setup bitmap `bitmap` from unsigned long `mask`.
- 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.
- void `hwloc_bitmap_set` (`hwloc_bitmap_t` bitmap, unsigned id)
Add index `id` in bitmap `bitmap`.
- void `hwloc_bitmap_set_range` (`hwloc_bitmap_t` bitmap, unsigned begin, unsigned end)
Add indexes from `begin` to `end` in bitmap `bitmap`.
- 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`.
- void `hwloc_bitmap_clr` (`hwloc_bitmap_t` bitmap, unsigned id)
Remove index `id` from bitmap `bitmap`.
- void `hwloc_bitmap_clr_range` (`hwloc_bitmap_t` bitmap, unsigned begin, unsigned end)
Remove index from `begin` to `end` in bitmap `bitmap`.
- void `hwloc_bitmap_singlify` (`hwloc_bitmap_t` bitmap)
Keep a single index among those set in bitmap `bitmap`.
- unsigned long `hwloc_bitmap_to_ulong` (`hwloc_const_bitmap_t` bitmap)
Convert the beginning part of bitmap `bitmap` into unsigned long `mask`.

- 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.
- int `hwloc_bitmap_isset` (`hwloc_const_bitmap_t` bitmap, unsigned `id`)
Test whether index `id` is part of bitmap `bitmap`.
- int `hwloc_bitmap_iszero` (`hwloc_const_bitmap_t` bitmap)
Test whether bitmap `bitmap` is empty.
- int `hwloc_bitmap_isfull` (`hwloc_const_bitmap_t` bitmap)
Test whether bitmap `bitmap` is completely full.
- int `hwloc_bitmap_first` (`hwloc_const_bitmap_t` bitmap)
Compute the first index (least significant bit) in bitmap `bitmap`.
- int `hwloc_bitmap_next` (`hwloc_const_bitmap_t` bitmap, unsigned `prev`)
Compute the next index in bitmap `bitmap` which is after index `prev`.
- int `hwloc_bitmap_last` (`hwloc_const_bitmap_t` bitmap)
Compute the last index (most significant bit) in bitmap `bitmap`.
- 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).
- 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`.
- void `hwloc_bitmap_and` (`hwloc_bitmap_t` `res`, `hwloc_const_bitmap_t` `bitmap1`, `hwloc_const_bitmap_t` `bitmap2`)
And bitmaps `bitmap1` and `bitmap2` and store the result in bitmap `res`.
- void `hwloc_bitmap_andnot` (`hwloc_bitmap_t` `res`, `hwloc_const_bitmap_t` `bitmap1`, `hwloc_const_bitmap_t` `bitmap2`)
And bitmap `bitmap1` and the negation of `bitmap2` and store the result in bitmap `res`.
- 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`.

- void `hwloc_bitmap_not` (`hwloc_bitmap_t` res, `hwloc_const_bitmap_t` bitmap)
Negate bitmap `bitmap` and store the result in bitmap `res`.
- int `hwloc_bitmap_intersects` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Test whether bitmaps `bitmap1` and `bitmap2` intersect.
- 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`.
- int `hwloc_bitmap_isequal` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Test whether bitmap `bitmap1` is equal to bitmap `bitmap2`.
- int `hwloc_bitmap_compare_first` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Compare bitmaps `bitmap1` and `bitmap2` using their lowest index.
- int `hwloc_bitmap_compare` (`hwloc_const_bitmap_t` bitmap1, `hwloc_const_bitmap_t` bitmap2)
Compare bitmaps `bitmap1` and `bitmap2` using their highest index.

7.1.1 Detailed Description

The bitmap API, for use in hwloc itself.

7.2 cuda.h File Reference

Macros to help interaction between hwloc and the CUDA Driver API.

```
#include <hwloc.h>
#include <hwloc/config.h>
#include <hwloc/linux.h>
#include <cuda.h>
```

Functions

- inline int `hwloc_cuda_get_device_cpuset` (`hwloc_topology_t` topology , CUdevice cudevice, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device cudevice.

7.2.1 Detailed Description

Macros to help interaction between hwloc and the CUDA Driver API.

Applications that use both hwloc and the CUDA Driver API may want to include this file so as to get topology information for CUDA devices.

7.3 `cuda.h` File Reference

Macros to help interaction between `hwloc` and the CUDA Runtime API.

```
#include <hwloc.h>
#include <hwloc/config.h>
#include <hwloc/linux.h>
#include <cuda_runtime_api.h>
```

Functions

- inline int `hwloc_cuda_get_device_cpuset` (`hwloc_topology_t` topology , int device, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device `device`.

7.3.1 Detailed Description

Macros to help interaction between `hwloc` and the CUDA Runtime API.

Applications that use both `hwloc` and the CUDA Runtime API may want to include this file so as to get topology information for CUDA devices.

7.4 glibc-sched.h File Reference

Macros to help interaction between hwloc and glibc scheduling routines.

```
#include <hwloc.h>
#include <hwloc/helper.h>
#include <assert.h>
```

Functions

- inline int `hwloc_cpuset_to_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_const_cpuset_t` hwlocset, `cpu_set_t` *schedset, `size_t` schedsetsize)
Convert hwloc CPU set toposet into glibc sched affinity CPU set schedset.
- inline int `hwloc_cpuset_from_glibc_sched_affinity` (`hwloc_topology_t` topology , `hwloc_cpuset_t` hwlocset, `const cpu_set_t` *schedset, `size_t` schedsetsize)
Convert glibc sched affinity CPU set schedset into hwloc CPU set.

7.4.1 Detailed Description

Macros to help interaction between hwloc and glibc scheduling routines.

Applications that use both hwloc and glibc scheduling routines such as `sched_getaffinity` may want to include this file so as to ease conversion between their respective types.

7.5 helper.h File Reference

High-level hwloc traversal helpers.

```
#include <stdlib.h>
#include <errno.h>
```

Functions

- inline int `hwloc_get_type_or_below_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type `type` or below.
- inline int `hwloc_get_type_or_above_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type `type` or above.
- inline `hwloc_obj_t` `hwloc_get_root_obj` (`hwloc_topology_t` topology)
Returns the top-object of the topology-tree.
- inline `hwloc_obj_t` `hwloc_get_ancestor_obj_by_depth` (`hwloc_topology_t` topology, unsigned depth, `hwloc_obj_t` obj)
Returns the ancestor object of `obj` at depth `depth`.
- inline `hwloc_obj_t` `hwloc_get_ancestor_obj_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type, `hwloc_obj_t` obj)
Returns the ancestor object of `obj` with type `type`.
- inline `hwloc_obj_t` `hwloc_get_next_obj_by_depth` (`hwloc_topology_t` topology, unsigned depth, `hwloc_obj_t` prev)
Returns the next object at depth `depth`.
- inline `hwloc_obj_t` `hwloc_get_next_obj_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type, `hwloc_obj_t` prev)
Returns the next object of type `type`.
- inline `hwloc_obj_t` `hwloc_get_pu_obj_by_os_index` (`hwloc_topology_t` topology, unsigned `os_index`)
Returns the object of type `HWLOC_OBJ_PU` with `os_index`.
- inline `hwloc_obj_t` `hwloc_get_next_child` (`hwloc_topology_t` topology, `hwloc_obj_t` parent, `hwloc_obj_t` prev)
Return the next child.

- inline `hwloc_obj_t hwloc_get_common_ancestor_obj` (`hwloc_topology_t topology`, `hwloc_obj_t obj1`, `hwloc_obj_t obj2`)
Returns the common parent object to objects `obj1` and `obj2`.
- inline `int hwloc_obj_is_in_subtree` (`hwloc_topology_t topology`, `hwloc_obj_t obj`, `hwloc_obj_t subtree_root`)
Returns true if `obj` is inside the subtree beginning with `subtree_root`.
- inline `hwloc_obj_t hwloc_get_first_largest_obj_inside_cpuset` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`)
Get the first largest object included in the given `set`.
- `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 `set`.
- inline `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`)
Return the next object at depth `depth` included in CPU set `set`.
- inline `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`)
Return the next object of type `type` included in CPU set `set`.
- inline `hwloc_obj_t hwloc_get_obj_inside_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `unsigned depth`, `unsigned idx`)
Return the `idx`-th object at depth `depth` included in CPU set `set`.
- inline `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`)
Return the `idx`-th object of type `type` included in CPU set `set`.
- inline `unsigned hwloc_get_nbobjs_inside_cpuset_by_depth` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `unsigned depth`)
Return the number of objects at depth `depth` included in CPU set `set`.
- inline `int hwloc_get_nbobjs_inside_cpuset_by_type` (`hwloc_topology_t topology`, `hwloc_const_cpuset_t set`, `hwloc_obj_type_t type`)
Return the number of objects of type `type` included in CPU set `set`.

- inline `hwloc_obj_t hwloc_get_child_covering_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set, `hwloc_obj_t` parent)
Get the child covering at least CPU set set.
- inline `hwloc_obj_t hwloc_get_obj_covering_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set)
Get the lowest object covering at least CPU set set.
- inline `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)
Iterate through same-depth objects covering at least CPU set set.
- inline `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)
Iterate through same-type objects covering at least CPU set set.
- inline `hwloc_obj_t hwloc_get_cache_covering_cpuset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` set)
Get the first cache covering a cpuset set.
- inline `hwloc_obj_t hwloc_get_shared_cache_covering_obj` (`hwloc_topology_t` topology, `hwloc_obj_t` obj)
Get the first cache shared between an object and somebody else.
- 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.
- inline `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)
Find an object below another object, both specified by types and indexes.
- inline `hwloc_obj_t hwloc_get_obj_below_array_by_type` (`hwloc_topology_t` topology, int nr, `hwloc_obj_type_t *`typev, unsigned `*idxv`)
Find an object below a chain of objects specified by types and indexes.
- inline void `hwloc_distributev` (`hwloc_topology_t` topology, `hwloc_obj_t *`root, unsigned n_roots, `hwloc_cpuset_t *`cpuset, unsigned n, unsigned until)
Distribute n items over the topology under root.

- inline void `hwloc_distribute` (`hwloc_topology_t` topology, `hwloc_obj_t` root, `hwloc_cpuset_t` *cpuset, unsigned n, unsigned until)
- inline void * `hwloc_alloc_mbind_policy_nodeset` (`hwloc_topology_t` topology, `size_t` len, `hwloc_const_nodeset_t` nodeset, `hwloc_mbind_policy_t` policy, int flags)

Allocate some memory on the given nodeset nodeset.

- inline void * `hwloc_alloc_mbind_policy` (`hwloc_topology_t` topology, `size_t` len, `hwloc_const_cpuset_t` cpuset, `hwloc_mbind_policy_t` policy, int flags)

Allocate some memory on the memory nodes near given cpuset cpuset.

- inline `hwloc_const_cpuset_t` `hwloc_topology_get_complete_cpuset` (`hwloc_topology_t` topology)
- inline `hwloc_const_cpuset_t` `hwloc_topology_get_topology_cpuset` (`hwloc_topology_t` topology)
- inline `hwloc_const_cpuset_t` `hwloc_topology_get_online_cpuset` (`hwloc_topology_t` topology)

Get online CPU set.

- inline `hwloc_const_cpuset_t` `hwloc_topology_get_allowed_cpuset` (`hwloc_topology_t` topology)

Get allowed CPU set.

- inline `hwloc_const_nodeset_t` `hwloc_topology_get_complete_nodeset` (`hwloc_topology_t` topology)
- inline `hwloc_const_nodeset_t` `hwloc_topology_get_topology_nodeset` (`hwloc_topology_t` topology)
- inline `hwloc_const_nodeset_t` `hwloc_topology_get_allowed_nodeset` (`hwloc_topology_t` topology)

Get allowed node set.

- inline void `hwloc_cpuset_to_nodeset` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `hwloc_nodeset_t` nodeset)

Convert a CPU set into a NUMA node set and handle non-NUMA cases.

- inline void `hwloc_cpuset_to_nodeset_strict` (`struct hwloc_topology *`topology, `hwloc_const_cpuset_t` cpuset, `hwloc_nodeset_t` nodeset)

Convert a CPU set into a NUMA node set without handling non-NUMA cases.

- inline void `hwloc_cpuset_from_nodeset` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `hwloc_const_nodeset_t` nodeset)

Convert a NUMA node set into a CPU set and handle non-NUMA cases.

- inline void `hwloc_cpuset_from_nodeset_strict` (struct `hwloc_topology` *topology, `hwloc_cpuset_t` cpuset, `hwloc_const_nodeset_t` nodeset)
Convert a NUMA node set into a CPU set without handling non-NUMA cases.

7.5.1 Detailed Description

High-level hwloc traversal helpers.

7.6 hwloc.doxy File Reference

7.7 hwloc.h File Reference

The hwloc API.

```
#include <hwloc/config.h>
#include <sys/types.h>
#include <stdio.h>
#include <string.h>
#include <limits.h>
#include <hwloc/rename.h>
#include <hwloc/bitmap.h>
#include <hwloc/cpuset.h>
#include <hwloc/helper.h>
```

Data Structures

- struct [hwloc_obj_memory_s](#)
Object memory.
- struct [hwloc_obj_memory_s::hwloc_obj_memory_page_type_s](#)
Array of local memory page types, NULL if no local memory and page_types is 0.
- struct [hwloc_obj](#)
Structure of a topology object.
- union [hwloc_obj_attr_u](#)
Object type-specific Attributes.
- struct [hwloc_obj_attr_u::hwloc_cache_attr_s](#)
Cache-specific Object Attributes.
- struct [hwloc_obj_attr_u::hwloc_group_attr_s](#)
Group-specific Object Attributes.
- struct [hwloc_obj_info_s](#)
Object info.
- struct [hwloc_topology_discovery_support](#)
Flags describing actual discovery support for this topology.

- struct [hwloc_topology_cpubind_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.

Defines

- #define [HWLOC_API_VERSION](#) 0x00010100
Indicate at build time which hwloc API version is being used.

Typedefs

- typedef [hwloc_topology](#) * [hwloc_topology_t](#)
Topology context.
- typedef [hwloc_bitmap_t](#) [hwloc_cpuset_t](#)
A CPU set is a bitmap whose bits are set according to CPU physical OS indexes.
- typedef [hwloc_const_bitmap_t](#) [hwloc_const_cpuset_t](#)
A non-modifiable [hwloc_cpuset_t](#).
- typedef [hwloc_bitmap_t](#) [hwloc_nodese_t](#)
A node set is a bitmap whose bits are set according to NUMA memory node physical OS indexes.
- typedef [hwloc_const_bitmap_t](#) [hwloc_const_nodese_t](#)
A non-modifiable [hwloc_nodese_t](#).
- typedef [hwloc_obj](#) * [hwloc_obj_t](#)
Convenience typedef; a pointer to a struct [hwloc_obj](#).

Enumerations

- enum `hwloc_obj_type_t` {
 [HWLOC_OBJ_SYSTEM](#), [HWLOC_OBJ_MACHINE](#), [HWLOC_OBJ_NODE](#),
 [HWLOC_OBJ_SOCKET](#),
 [HWLOC_OBJ_CACHE](#), [HWLOC_OBJ_CORE](#), [HWLOC_OBJ_PU](#),
 [HWLOC_OBJ_GROUP](#),
 [HWLOC_OBJ_MISC](#) }
Type of topology object.
- enum `hwloc_compare_types_e` { [HWLOC_TYPE_UNORDERED](#) }
- enum `hwloc_topology_flags_e` { [HWLOC_TOPOLOGY_FLAG_WHOLE_SYSTEM](#), [HWLOC_TOPOLOGY_FLAG_IS_THISSYSTEM](#) }
Flags to be set onto a topology context before load.
- enum `hwloc_get_type_depth_e` { [HWLOC_TYPE_DEPTH_UNKNOWN](#),
 [HWLOC_TYPE_DEPTH_MULTIPLE](#) }
- enum `hwloc_cpubind_flags_t` { [HWLOC_CPUBIND_PROCESS](#), [HWLOC_CPUBIND_THREAD](#), [HWLOC_CPUBIND_STRICT](#), [HWLOC_CPUBIND_NOMEMBIND](#) }
Process/Thread binding flags.
- enum `hwloc_membind_policy_t` {
 [HWLOC_MEMBIND_DEFAULT](#), [HWLOC_MEMBIND_FIRSTTOUCH](#),
 [HWLOC_MEMBIND_BIND](#), [HWLOC_MEMBIND_INTERLEAVE](#),
 [HWLOC_MEMBIND_REPLICATE](#), [HWLOC_MEMBIND_NEXTTOUCH](#)
 }
Memory binding policy.
- enum `hwloc_membind_flags_t` {
 [HWLOC_MEMBIND_PROCESS](#), [HWLOC_MEMBIND_THREAD](#),
 [HWLOC_MEMBIND_STRICT](#), [HWLOC_MEMBIND_MIGRATE](#),
 [HWLOC_MEMBIND_NOCPUBIND](#) }
Memory binding flags.

Functions

- int `hwloc_compare_types` ([hwloc_obj_type_t](#) type1, [hwloc_obj_type_t](#) type2)
Compare the depth of two object types.

- int `hwloc_topology_init` (`hwloc_topology_t *topology`)
Allocate a topology context.
- int `hwloc_topology_load` (`hwloc_topology_t topology`)
Build the actual topology.
- void `hwloc_topology_destroy` (`hwloc_topology_t topology`)
Terminate and free a topology context.
- void `hwloc_topology_check` (`hwloc_topology_t topology`)
Run internal checks on a topology structure.
- int `hwloc_topology_ignore_type` (`hwloc_topology_t topology`, `hwloc_obj_type_t type`)
Ignore an object type.
- 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.
- int `hwloc_topology_ignore_all_keep_structure` (`hwloc_topology_t topology`)
Ignore all objects that do not bring any structure.
- int `hwloc_topology_set_flags` (`hwloc_topology_t topology`, unsigned long flags)
Set OR'ed flags to non-yet-loaded topology.
- 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.
- int `hwloc_topology_set_pid` (`hwloc_topology_t restrict topology`, `hwloc_pid_t pid`)
Change which pid the topology is viewed from.
- int `hwloc_topology_set_synthetic` (`hwloc_topology_t restrict topology`, const char *restrict description)
Enable synthetic topology.
- int `hwloc_topology_set_xml` (`hwloc_topology_t restrict topology`, const char *restrict xmlpath)
Enable XML-file based topology.

- 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.
- const struct `hwloc_topology_support` * `hwloc_topology_get_support` (`hwloc_topology_t` restrict topology)
Retrieve the topology support.
- void `hwloc_topology_export_xml` (`hwloc_topology_t` topology, const char *xmlpath)
Export the topology into an XML file.
- void `hwloc_topology_export_xmlbuffer` (`hwloc_topology_t` topology, char **xmlbuffer, int *buflen)
Export the topology into a newly-allocated XML memory buffer.
- `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.
- `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 to the topology.
- unsigned `hwloc_topology_get_depth` (`hwloc_topology_t` restrict topology)
Get the depth of the hierachical tree of objects.
- int `hwloc_get_type_depth` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the depth of objects of type type.
- `hwloc_obj_type_t` `hwloc_get_depth_type` (`hwloc_topology_t` topology, unsigned depth)
Returns the type of objects at depth depth.
- unsigned `hwloc_get_nbobjs_by_depth` (`hwloc_topology_t` topology, unsigned depth)
Returns the width of level at depth depth.
- inline int `hwloc_get_nbobjs_by_type` (`hwloc_topology_t` topology, `hwloc_obj_type_t` type)
Returns the width of level type type.

- `int hwloc_topology_is_thissystem (hwloc_topology_t restrict topology)`
Does the topology context come from this system?
- `hwloc_obj_t hwloc_get_obj_by_depth (hwloc_topology_t topology, unsigned depth, unsigned idx)`
Returns the topology object at index `index` from depth `depth`.
- `inline hwloc_obj_t hwloc_get_obj_by_type (hwloc_topology_t topology, hwloc_obj_type_t type, unsigned idx)`
Returns the topology object at index `index` with type `type`.
- `const char * hwloc_obj_type_string (hwloc_obj_type_t type)`
Return a stringified topology object type.
- `hwloc_obj_type_t hwloc_obj_type_of_string (const char *string)`
Return an object type from the string.
- `int hwloc_obj_type_snprintf (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.
- `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.
- `int hwloc_obj_snprintf (char *restrict string, size_t size, hwloc_topology_t topology, hwloc_obj_t obj, const char *restrict indexprefix, int verbose)`
Stringify a given topology object into a human-readable form.
- `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.
- `inline char * hwloc_obj_get_info_by_name (hwloc_obj_t obj, const char *name)`
Search the given key name in object infos and return the corresponding value.
- `int hwloc_set_cpupbind (hwloc_topology_t topology, hwloc_const_cpuset_t set, int flags)`
Bind current process or thread on cpus given in bitmap `set`.
- `int hwloc_get_cpupbind (hwloc_topology_t topology, hwloc_cpuset_t set, int flags)`

Get current process or thread binding.

- 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 bitmap set.

- int `hwloc_get_proc_cpubind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_cpuset_t` set, int flags)

Get the current binding of process pid.

- int `hwloc_set_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` tid, `hwloc_const_cpuset_t` set, int flags)

Bind a thread tid on cpus given in bitmap set.

- int `hwloc_get_thread_cpubind` (`hwloc_topology_t` topology, `hwloc_thread_t` tid, `hwloc_cpuset_t` set, int flags)

Get the current binding of thread tid.

- int `hwloc_set_membind_nodeset` (`hwloc_topology_t` topology, `hwloc_const_nodeset_t` nodeset, `hwloc_membind_policy_t` policy, int flags)

Bind current process memory on the given nodeset nodeset.

- int `hwloc_set_membind` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `hwloc_membind_policy_t` policy, int flags)

Bind current process memory on memory nodes near the given cpuset cpuset.

- int `hwloc_get_membind_nodeset` (`hwloc_topology_t` topology, `hwloc_nodeset_t` nodeset, `hwloc_membind_policy_t` *policy, int flags)

Get current process memory binding in nodeset nodeset.

- int `hwloc_get_membind` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `hwloc_membind_policy_t` *policy, int flags)

Get current process memory binding in cpuset cpuset.

- int `hwloc_set_proc_membind_nodeset` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_nodeset_t` nodeset, `hwloc_membind_policy_t` policy, int flags)

Bind given process memory on the given nodeset nodeset.

- int `hwloc_set_proc_membind` (`hwloc_topology_t` topology, `hwloc_pid_t` pid, `hwloc_const_cpuset_t` cpuset, `hwloc_membind_policy_t` policy, int flags)

Bind given process memory on memory nodes near the given cpuset cpuset.

- 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)
Get current process memory binding in nodeset nodeset.
- int `hwloc_get_proc_membind` (hwloc_topology_t topology, hwloc_pid_t pid, hwloc_cpuse
- t_t cpuset, hwloc_membind_policy_t *policy, int flags)
Get current process memory binding in cpuset cpuset.
- 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)
Bind some memory range on the given nodeset nodeset.
- int `hwloc_set_area_membind` (hwloc_topology_t topology, const void *addr, size_t len, hwloc_const_cpuse
- t_t cpuset, hwloc_membind_policy_t policy, int flags)
Bind some memory range on memory nodes near the given cpuset cpuset.
- 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)
Get some memory range memory binding in nodeset nodeset.
- int `hwloc_get_area_membind` (hwloc_topology_t topology, const void *addr, size_t len, hwloc_cpuse
- t_t cpuset, hwloc_membind_policy_t *policy, int flags)
Get some memory range memory binding in cpuset cpuset.
- void * `hwloc_alloc` (hwloc_topology_t topology, size_t len)
Allocate some memory.
- 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)
Allocate some memory on the given nodeset nodeset.
- void * `hwloc_alloc_membind` (hwloc_topology_t topology, size_t len, hwloc_const_cpuse
- t_t cpuset, hwloc_membind_policy_t policy, int flags)
Allocate some memory on memory nodes near the given cpuset cpuset.
- int `hwloc_free` (hwloc_topology_t topology, void *addr, size_t len)
Free some memory allocated by hwloc_alloc() or hwloc_alloc_membind().

7.7.1 Detailed Description

The hwloc API.

See hwloc/bitmap.h for bitmap specific macros. See hwloc/helper.h for high-level topology traversal helpers.

7.8 linux-libnuma.h File Reference

Macros to help interaction between hwloc and Linux libnuma.

```
#include <hwloc.h>
```

```
#include <numa.h>
```

Functions

- inline int [hwloc_cpuset_to_linux_libnuma_ulong](#)s (hwloc_topology_t topology, hwloc_const_cpuset_t cpuset, unsigned long *mask, unsigned long *maxnode)
Convert hwloc CPU set cpuset into the array of unsigned long mask.
- inline int [hwloc_node](#)set_to_linux_libnuma_ulong
s (hwloc_topology_t topology, hwloc_const_nodeset_t nodeset, unsigned long *mask, unsigned long *maxnode)
Convert hwloc NUMA node set nodeset into the array of unsigned long mask.- inline int [hwloc_cpuset_from_linux_libnuma_ulong](#)s (hwloc_topology_t topology, hwloc_cpuset_t cpuset, const unsigned long *mask, unsigned long maxnode)
Convert the array of unsigned long mask into hwloc CPU set.
- inline int [hwloc_node](#)set_from_linux_libnuma_ulong
s (hwloc_topology_t topology, hwloc_nodeset_t nodeset, const unsigned long *mask, unsigned long maxnode)
Convert the array of unsigned long mask into hwloc NUMA node set.- inline struct bitmask * [hwloc_cpuset_to_linux_libnuma_bitmask](#) (hwloc_topo
logy_t topology, hwloc_const_cpuset_t cpuset)
Convert hwloc CPU set cpuset into the returned libnuma bitmask.- inline struct bitmask * [hwloc_node](#)set_to_linux_libnuma_bitmask (hwloc_topo
logy_t topology, hwloc_const_nodeset_t nodeset)
Convert hwloc NUMA node set nodeset into the returned libnuma bitmask.- inline int [hwloc_cpuset_from_linux_libnuma_bitmask](#) (hwloc_topo
logy_t topology, hwloc_cpuset_t cpuset, const struct bitmask *bitmask)
Convert libnuma bitmask bitmask into hwloc CPU set cpuset.- inline int [hwloc_node](#)set_from_linux_libnuma_bitmask (hwloc_topo
logy_t topology, hwloc_nodeset_t nodeset, const struct bitmask *bitmask)
Convert libnuma bitmask bitmask into hwloc NUMA node set nodeset.

- inline int `hwloc_cpuset_to_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_const_cpuset_t` cpuset, `nodemask_t *nodemask`)
Convert hwloc CPU set cpuset into libnuma nodemask nodemask.
- inline int `hwloc_nodese_t_to_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_const_nodese_t` nodese_t, `nodemask_t *nodemask`)
Convert hwloc NUMA node set nodese_t into libnuma nodemask nodemask.
- inline int `hwloc_cpuset_from_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_cpuset_t` cpuset, `const nodemask_t *nodemask`)
Convert libnuma nodemask nodemask into hwloc CPU set cpuset.
- inline int `hwloc_nodese_t_from_linux_libnuma_nodemask` (`hwloc_topology_t` topology, `hwloc_nodese_t` nodese_t, `const nodemask_t *nodemask`)
Convert libnuma nodemask nodemask into hwloc NUMA node set nodese_t.

7.8.1 Detailed Description

Macros to help interaction between hwloc and Linux libnuma.

Applications that use both Linux libnuma and hwloc may want to include this file so as to ease conversion between their respective types.

7.9 linux.h File Reference

Macros to help interaction between hwloc and Linux.

```
#include <hwloc.h>
#include <stdio.h>
```

Functions

- int `hwloc_linux_parse_cpumap_file` (FILE *file, hwloc_cpuset_t set)
Convert a linux kernel cpumap file file into hwloc CPU set.
- int `hwloc_linux_set_tid_cpupubind` (hwloc_topology_t topology, pid_t tid, hwloc_const_cpuset_t set)
Bind a thread tid on cpus given in cpuset set.
- int `hwloc_linux_get_tid_cpupubind` (hwloc_topology_t topology, pid_t tid, hwloc_cpuset_t set)
Get the current binding of thread tid.

7.9.1 Detailed Description

Macros to help interaction between hwloc and Linux.

Applications that use hwloc on Linux may want to include this file if using some low-level Linux features.

7.10 myriexpress.h File Reference

Macros to help interaction between hwloc and Myrinet Express.

```
#include <hwloc.h>
#include <hwloc/config.h>
#include <hwloc/linux.h>
#include <myriexpress.h>
```

Functions

- inline int `hwloc_mx_board_get_device_cpuset` (`hwloc_topology_t` topology, unsigned id, `hwloc_cpuset_t` set)
Get the CPU set of logical processors that are physically close the MX board id.
- inline int `hwloc_mx_endpoint_get_device_cpuset` (`hwloc_topology_t` topology, `mx_endpoint_t` endpoint, `hwloc_cpuset_t` set)
Get the CPU set of logical processors that are physically close to endpoint endpoint.

7.10.1 Detailed Description

Macros to help interaction between hwloc and Myrinet Express.

Applications that use both hwloc and Myrinet Express verbs may want to include this file so as to get topology information for Myrinet hardware.

7.11 openfabrics-verbs.h File Reference

Macros to help interaction between hwloc and OpenFabrics verbs.

```
#include <hwloc.h>
#include <hwloc/config.h>
#include <hwloc/linux.h>
#include <infiniband/verbs.h>
```

Functions

- inline int `hwloc_ibv_get_device_cpuset` (`hwloc_topology_t` topology , struct `ibv_device` *ibdev, `hwloc_cpuset_t` set)

Get the CPU set of logical processors that are physically close to device ibdev.

7.11.1 Detailed Description

Macros to help interaction between hwloc and OpenFabrics verbs.

Applications that use both hwloc and OpenFabrics verbs may want to include this file so as to get topology information for OpenFabrics hardware.

Index

- Advanced Traversal Helpers, [86](#)
- alloc_membind
 - hwloc_topology_membind_
support, [137](#)
- allowed_cpuset
 - hwloc_obj, [121](#)
- allowed_nodeset
 - hwloc_obj, [122](#)
- API version, [39](#)
- arity
 - hwloc_obj, [122](#)
- attr
 - hwloc_obj, [122](#)

- Basic Traversal Helpers, [77](#)
- bind_membind
 - hwloc_topology_membind_
support, [137](#)
- Binding Helpers, [88](#)
- bitmap.h, [141](#)

- cache
 - hwloc_obj_attr_u, [127](#)
- Cache-specific Finding Helpers, [85](#)
- children
 - hwloc_obj, [122](#)
- complete_cpuset
 - hwloc_obj, [122](#)
- complete_nodeset
 - hwloc_obj, [122](#)
- Configure Topology Detection, [49](#)
- Conversion between cpuset and nodeset, [93](#)

- count
 - hwloc_obj_memory_s::hwloc_
obj_memory_page_type_s,
[133](#)

- CPU binding, [64](#)
- cpubind
 - hwloc_topology_support, [140](#)
- cpuset
 - hwloc_obj, [123](#)
- Cpuset Helpers, [90](#)
- Create and Destroy Topologies, [47](#)
- CUDA Driver API Specific Functions,
[115](#)
- CUDA Runtime API Specific Functions,
[116](#)
- cuda.h, [146](#)
- cuda.h, [147](#)

- depth
 - hwloc_obj, [123](#)
 - hwloc_obj_attr_u::hwloc_
cache_attr_s, [128](#)
 - hwloc_obj_attr_u::hwloc_
group_attr_s, [129](#)
- discovery
 - hwloc_topology_support, [140](#)

- Finding a set of similar Objects covering at least a CPU set, [84](#)
- Finding a single Object covering at least CPU set, [83](#)
- Finding Objects Inside a CPU set, [80](#)
- first_child
 - hwloc_obj, [123](#)
- firsttouch_membind
 - hwloc_topology_membind_
support, [137](#)

- Get some Topology Information, [57](#)
- get_area_membind
 - hwloc_topology_membind_
support, [138](#)

- get_proc_cpupind
 - hwloc_topology_cpupind_-support, 134
- get_proc_membind
 - hwloc_topology_membind_-support, 138
- get_thisproc_cpupind
 - hwloc_topology_cpupind_-support, 134
- get_thisproc_membind
 - hwloc_topology_membind_-support, 138
- get_thisthread_cpupind
 - hwloc_topology_cpupind_-support, 134
- get_thisthread_membind
 - hwloc_topology_membind_-support, 138
- get_thread_cpupind
 - hwloc_topology_cpupind_-support, 134
- glibc-sched.h, 148
- group
 - hwloc_obj_attr_u, 127

- helper.h, 149
- Helpers for manipulating glibc sched affinity, 106
- Helpers for manipulating Linux libnuma bitmask, 111
- Helpers for manipulating Linux libnuma nodemask_t, 113
- Helpers for manipulating Linux libnuma unsigned long masks, 109
- hwloc.doxy, 154
- hwloc.h, 155
- hwloc_alloc
 - hwlocality_membinding, 71
- hwloc_alloc_membind
 - hwlocality_membinding, 72
- hwloc_alloc_membind_nodese
 - hwlocality_membinding, 72
- hwloc_alloc_membind_policy
 - hwlocality_helper_binding, 88
- hwloc_alloc_membind_policy_-nodeset
 - hwlocality_helper_binding, 88
- HWLOC_API_VERSION
 - hwlocality_api_version, 39
- hwloc_bitmap_allbut
 - hwlocality_bitmap, 99
- hwloc_bitmap_alloc
 - hwlocality_bitmap, 99
- hwloc_bitmap_alloc_full
 - hwlocality_bitmap, 100
- hwloc_bitmap_and
 - hwlocality_bitmap, 100
- hwloc_bitmap_andnot
 - hwlocality_bitmap, 100
- hwloc_bitmap_asprintf
 - hwlocality_bitmap, 100
- hwloc_bitmap_clr
 - hwlocality_bitmap, 100
- hwloc_bitmap_clr_range
 - hwlocality_bitmap, 100
- hwloc_bitmap_compare
 - hwlocality_bitmap, 100
- hwloc_bitmap_compare_first
 - hwlocality_bitmap, 101
- hwloc_bitmap_copy
 - hwlocality_bitmap, 101
- hwloc_bitmap_dup
 - hwlocality_bitmap, 101
- hwloc_bitmap_fill
 - hwlocality_bitmap, 101
- hwloc_bitmap_first
 - hwlocality_bitmap, 101
- hwloc_bitmap_foreach_begin
 - hwlocality_bitmap, 99
- hwloc_bitmap_foreach_end
 - hwlocality_bitmap, 99
- hwloc_bitmap_free
 - hwlocality_bitmap, 101
- hwloc_bitmap_from_ith_ulong
 - hwlocality_bitmap, 101
- hwloc_bitmap_from_ulong
 - hwlocality_bitmap, 102
- hwloc_bitmap_intersects
 - hwlocality_bitmap, 102
- hwloc_bitmap_isequal

- hwlocality_bitmap, 102
- hwloc_bitmap_isfull
 - hwlocality_bitmap, 102
- hwloc_bitmap_isincluded
 - hwlocality_bitmap, 102
- hwloc_bitmap_isset
 - hwlocality_bitmap, 102
- hwloc_bitmap_iszero
 - hwlocality_bitmap, 102
- hwloc_bitmap_last
 - hwlocality_bitmap, 102
- hwloc_bitmap_next
 - hwlocality_bitmap, 103
- hwloc_bitmap_not
 - hwlocality_bitmap, 103
- hwloc_bitmap_only
 - hwlocality_bitmap, 103
- hwloc_bitmap_or
 - hwlocality_bitmap, 103
- hwloc_bitmap_set
 - hwlocality_bitmap, 103
- hwloc_bitmap_set_ith_ulong
 - hwlocality_bitmap, 103
- hwloc_bitmap_set_range
 - hwlocality_bitmap, 103
- hwloc_bitmap_singlify
 - hwlocality_bitmap, 103
- hwloc_bitmap_snprintf
 - hwlocality_bitmap, 104
- hwloc_bitmap_sscanf
 - hwlocality_bitmap, 104
- hwloc_bitmap_t
 - hwlocality_bitmap, 99
- hwloc_bitmap_taskset_asprintf
 - hwlocality_bitmap, 104
- hwloc_bitmap_taskset_snprintf
 - hwlocality_bitmap, 104
- hwloc_bitmap_taskset_sscanf
 - hwlocality_bitmap, 104
- hwloc_bitmap_to_ith_ulong
 - hwlocality_bitmap, 104
- hwloc_bitmap_to_ulong
 - hwlocality_bitmap, 105
- hwloc_bitmap_weight
 - hwlocality_bitmap, 105
- hwloc_bitmap_xor
 - hwlocality_bitmap, 105
- hwloc_bitmap_zero
 - hwlocality_bitmap, 105
- hwloc_compare_types
 - hwlocality_types, 44
- hwloc_compare_types_e
 - hwlocality_types, 43
- hwloc_const_bitmap_t
 - hwlocality_bitmap, 99
- hwloc_const_cpuset_t
 - hwlocality_sets, 41
- hwloc_const_nodeset_t
 - hwlocality_sets, 41
- hwloc_cpupbind_flags_t
 - hwlocality_cpupbinding, 65
- HWLOC_CPUBIND_NOMEMBIND
 - hwlocality_cpupbinding, 66
- HWLOC_CPUBIND_PROCESS
 - hwlocality_cpupbinding, 65
- HWLOC_CPUBIND_STRICT
 - hwlocality_cpupbinding, 65
- HWLOC_CPUBIND_THREAD
 - hwlocality_cpupbinding, 65
- hwloc_cpuset_from_glibc_sched_
 - affinity
 - hwlocality_glibc_sched, 106
- hwloc_cpuset_from_linux_libnuma_
 - bitmask
 - hwlocality_linux_libnuma_
 - bitmask, 111
- hwloc_cpuset_from_linux_libnuma_
 - nodemask
 - hwlocality_linux_libnuma_
 - nodemask, 113
- hwloc_cpuset_from_linux_libnuma_
 - ulongs
 - hwlocality_linux_libnuma_
 - ulongs, 109
- hwloc_cpuset_from_nodeset
 - hwlocality_helper_nodeset_
 - convert, 93
- hwloc_cpuset_from_nodeset_strict
 - hwlocality_helper_nodeset_
 - convert, 94
- hwloc_cpuset_t
 - hwlocality_sets, 41

- hwloc_cpuset_to_glibc_sched_affinity
 - hwlocality_glibc_sched, 106
- hwloc_cpuset_to_linux_libnuma_-
 - bitmask
 - hwlocality_linux_libnuma_-
 - bitmask, 111
- hwloc_cpuset_to_linux_libnuma_-
 - nodemask
 - hwlocality_linux_libnuma_-
 - nodemask, 113
- hwloc_cpuset_to_linux_libnuma_-
 - ulongs
 - hwlocality_linux_libnuma_-
 - ulongs, 109
- hwloc_cpuset_to_nodeset
 - hwlocality_helper_nodeset_-
 - convert, 94
- hwloc_cpuset_to_nodeset_strict
 - hwlocality_helper_nodeset_-
 - convert, 94
- hwloc_cuda_get_device_cpuset
 - hwlocality_cuda, 115
- hwloc_cudart_get_device_cpuset
 - hwlocality_cudart, 116
- hwloc_distribute
 - hwlocality_helper_binding, 88
- hwloc_distributev
 - hwlocality_helper_binding, 89
- hwloc_free
 - hwlocality_membinding, 72
- hwloc_get_ancestor_obj_by_depth
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_ancestor_obj_by_type
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_area_membind
 - hwlocality_membinding, 72
- hwloc_get_area_membind_nodeset
 - hwlocality_membinding, 72
- hwloc_get_cache_covering_cpuset
 - hwlocality_helper_find_cache,
 - 85
- hwloc_get_child_covering_cpuset
 - hwlocality_helper_find_covering,
 - 83
- hwloc_get_closest_objs
 - hwlocality_helper_traversal, 86
- hwloc_get_common_ancestor_obj
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_cpupbind
 - hwlocality_cpupbinding, 66
- hwloc_get_depth_type
 - hwlocality_information, 58
- hwloc_get_first_largest_obj_inside_-
 - cpuset
 - hwlocality_helper_find_inside,
 - 81
- hwloc_get_largest_objs_inside_cpuset
 - hwlocality_helper_find_inside,
 - 81
- hwloc_get_membind
 - hwlocality_membinding, 73
- hwloc_get_membind_nodeset
 - hwlocality_membinding, 73
- hwloc_get_nbobjs_by_depth
 - hwlocality_information, 58
- hwloc_get_nbobjs_by_type
 - hwlocality_information, 58
- hwloc_get_nbobjs_inside_cpuset_-
 - by_depth
 - hwlocality_helper_find_inside,
 - 81
- hwloc_get_nbobjs_inside_cpuset_-
 - by_type
 - hwlocality_helper_find_inside,
 - 81
- hwloc_get_next_child
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_next_obj_by_depth
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_next_obj_by_type
 - hwlocality_helper_traversal_-
 - basic, 78
- hwloc_get_next_obj_covering_-
 - cpuset_by_depth
 - hwlocality_helper_find_-
 - coverings, 84

- hwloc_get_next_obj_covering_-
 cpuset_by_type
 hwlocality_helper_find_-
 coverings, 84
- hwloc_get_next_obj_inside_cpuset_-
 by_depth
 hwlocality_helper_find_inside,
 81
- hwloc_get_next_obj_inside_cpuset_-
 by_type
 hwlocality_helper_find_inside,
 82
- hwloc_get_obj_below_array_by_type
 hwlocality_helper_traversal, 86
- hwloc_get_obj_below_by_type
 hwlocality_helper_traversal, 87
- hwloc_get_obj_by_depth
 hwlocality_traversal, 60
- hwloc_get_obj_by_type
 hwlocality_traversal, 60
- hwloc_get_obj_covering_cpuset
 hwlocality_helper_find_covering,
 83
- hwloc_get_obj_inside_cpuset_by_-
 depth
 hwlocality_helper_find_inside,
 82
- hwloc_get_obj_inside_cpuset_by_-
 type
 hwlocality_helper_find_inside,
 82
- hwloc_get_proc_cpupbind
 hwlocality_cpupbinding, 66
- hwloc_get_proc_membind
 hwlocality_membinding, 73
- hwloc_get_proc_membind_nodeset
 hwlocality_membinding, 73
- hwloc_get_pu_obj_by_os_index
 hwlocality_helper_traversal_-
 basic, 78
- hwloc_get_root_obj
 hwlocality_helper_traversal_-
 basic, 79
- hwloc_get_shared_cache_covering_-
 obj
 hwlocality_helper_find_cache,
 85
- hwloc_get_thread_cpupbind
 hwlocality_cpupbinding, 66
- hwloc_get_type_depth
 hwlocality_information, 58
- hwloc_get_type_depth_e
 hwlocality_information, 57
- hwloc_get_type_or_above_depth
 hwlocality_helper_types, 76
- hwloc_get_type_or_below_depth
 hwlocality_helper_types, 76
- hwloc_ibv_get_device_cpuset
 hwlocality_openfabrics, 117
- hwloc_linux_get_tid_cpupbind
 hwlocality_linux, 107
- hwloc_linux_parse_cpumap_file
 hwlocality_linux, 107
- hwloc_linux_set_tid_cpupbind
 hwlocality_linux, 107
- HWLOC_MEMBIND_BIND
 hwlocality_membinding, 71
- HWLOC_MEMBIND_DEFAULT
 hwlocality_membinding, 71
- HWLOC_MEMBIND_-
 FIRSTTOUCH
 hwlocality_membinding, 71
- hwloc_membind_flags_t
 hwlocality_membinding, 70
- HWLOC_MEMBIND_-
 INTERLEAVE
 hwlocality_membinding, 71
- HWLOC_MEMBIND_MIGRATE
 hwlocality_membinding, 71
- HWLOC_MEMBIND_-
 NEXTTOUCH
 hwlocality_membinding, 71
- HWLOC_MEMBIND_NOCPUBIND
 hwlocality_membinding, 71
- hwloc_membind_policy_t
 hwlocality_membinding, 71
- HWLOC_MEMBIND_PROCESS
 hwlocality_membinding, 70
- HWLOC_MEMBIND_REPLICATE
 hwlocality_membinding, 71
- HWLOC_MEMBIND_STRICT

- hwlocality_membinding, 71
- HWLOC_MEMBIND_THREAD
 - hwlocality_membinding, 70
- hwloc_mx_board_get_device_cpuset
 - hwlocality_myriexpress, 118
- hwloc_mx_endpoint_get_device_-
 - cpuset
 - hwlocality_myriexpress, 118
- hwloc_nodeset_from_linux_-
 - libnuma_bitmask
 - hwlocality_linux_libnuma_-
 - bitmask, 111
- hwloc_nodeset_from_linux_-
 - libnuma_nodemask
 - hwlocality_linux_libnuma_-
 - nodemask, 113
- hwloc_nodeset_from_linux_-
 - libnuma_ulongs
 - hwlocality_linux_libnuma_-
 - ulongs, 110
- hwloc_nodeset_t
 - hwlocality_sets, 41
- hwloc_nodeset_to_linux_libnuma_-
 - bitmask
 - hwlocality_linux_libnuma_-
 - bitmask, 112
- hwloc_nodeset_to_linux_libnuma_-
 - nodemask
 - hwlocality_linux_libnuma_-
 - nodemask, 114
- hwloc_nodeset_to_linux_libnuma_-
 - ulongs
 - hwlocality_linux_libnuma_-
 - ulongs, 110
- hwloc_obj, 119
 - allowed_cpuset, 121
 - allowed_nodeset, 122
 - arity, 122
 - attr, 122
 - children, 122
 - complete_cpuset, 122
 - complete_nodeset, 122
 - cpuset, 123
 - depth, 123
 - first_child, 123
 - infos, 123
 - infos_count, 123
 - last_child, 124
 - logical_index, 124
 - memory, 124
 - name, 124
 - next_cousin, 124
 - next_sibling, 124
 - nodeset, 124
 - online_cpuset, 125
 - os_index, 125
 - os_level, 125
 - parent, 125
 - prev_cousin, 125
 - prev_sibling, 125
 - sibling_rank, 125
 - type, 125
 - userdata, 126
- hwloc_obj_attr_sprintf
 - hwlocality_conversion, 61
- hwloc_obj_attr_u, 127
 - cache, 127
 - group, 127
- hwloc_obj_attr_u::hwloc_cache_attr_-
 - s, 128
 - depth, 128
 - linesize, 128
 - size, 128
- hwloc_obj_attr_u::hwloc_group_-
 - attr_s, 129
 - depth, 129
- HWLOC_OBJ_CACHE
 - hwlocality_types, 44
- HWLOC_OBJ_CORE
 - hwlocality_types, 44
- hwloc_obj_cpuset_sprintf
 - hwlocality_conversion, 61
- hwloc_obj_get_info_by_name
 - hwlocality_conversion, 62
- HWLOC_OBJ_GROUP
 - hwlocality_types, 44
- hwloc_obj_info_s, 130
 - name, 130
 - value, 130
- hwloc_obj_is_in_subtree
 - hwlocality_helper_traversal_-
 - basic, 79

- HWLOC_OBJ_MACHINE
 - hwlocality_types, 43
- hwloc_obj_memory_s, 131
 - local_memory, 131
 - page_types, 131
 - page_types_len, 131
 - total_memory, 131
- hwloc_obj_memory_s::hwloc_obj_
memory_page_type_s,
133
 - count, 133
 - size, 133
- HWLOC_OBJ_MISC
 - hwlocality_types, 44
- HWLOC_OBJ_NODE
 - hwlocality_types, 44
- HWLOC_OBJ_PU
 - hwlocality_types, 44
- hwloc_obj_snprintf
 - hwlocality_conversion, 62
- HWLOC_OBJ_SOCKET
 - hwlocality_types, 44
- HWLOC_OBJ_SYSTEM
 - hwlocality_types, 43
- hwloc_obj_t
 - hwlocality_objects, 46
- hwloc_obj_type_of_string
 - hwlocality_conversion, 62
- hwloc_obj_type_snprintf
 - hwlocality_conversion, 62
- hwloc_obj_type_string
 - hwlocality_conversion, 63
- hwloc_obj_type_t
 - hwlocality_types, 43
- hwloc_set_area_membind
 - hwlocality_membinding, 73
- hwloc_set_area_membind_nodese
t
 - hwlocality_membinding, 73
- hwloc_set_cpupbind
 - hwlocality_cpupbinding, 67
- hwloc_set_membind
 - hwlocality_membinding, 74
- hwloc_set_membind_nodese
t
 - hwlocality_membinding, 74
- hwloc_set_proc_cpupbind
 - hwlocality_cpupbinding, 67
- hwloc_set_proc_membind
 - hwlocality_membinding, 74
- hwloc_set_proc_membind_nodese
t
 - hwlocality_membinding, 74
- hwloc_set_thread_cpupbind
 - hwlocality_cpupbinding, 67
- hwloc_topology_check
 - hwlocality_creation, 47
- hwloc_topology_cpupbind_s
upport,
134
 - get_proc_cpupbind, 134
 - get_thisproc_cpupbind, 134
 - get_thisthread_cpupbind, 134
 - get_thread_cpupbind, 134
 - set_proc_cpupbind, 134
 - set_thisproc_cpupbind, 135
 - set_thisthread_cpupbind, 135
 - set_thread_cpupbind, 135
- hwloc_topology_destroy
 - hwlocality_creation, 47
- hwloc_topology_discovery_s
upport,
136
 - pu, 136
- hwloc_topology_export_xml
 - hwlocality_tinker, 55
- hwloc_topology_export_xmlbu
ffer
 - hwlocality_tinker, 55
- HWLOC_TOPOLOGY_FLAG_IS_
THISSYSTEM
 - hwlocality_configuration, 51
- HWLOC_TOPOLOGY_FLAG_
WHOLE_SYSTEM
 - hwlocality_configuration, 51
- hwloc_topology_flags_e
 - hwlocality_configuration, 51
- hwloc_topology_get_allowed_cp
uset
 - hwlocality_helper_cpuset, 90
- hwloc_topology_get_allowed_n
odese
t
 - hwlocality_helper_nodese
t, 92
- hwloc_topology_get_complete_cp
uset
 - hwlocality_helper_cpuset, 90
- hwloc_topology_get_complete_
n
odese
t
 - hwlocality_helper_nodese
t, 92
- hwloc_topology_get_depth
 - hwlocality_information, 58

- hwloc_topology_get_online_cpuset
 - hwlocality_helper_cpuset, 90
- hwloc_topology_get_support
 - hwlocality_configuration, 51
- hwloc_topology_get_topology_cpuset
 - hwlocality_helper_cpuset, 91
- hwloc_topology_get_topology_-
 - nodeset
 - hwlocality_helper_nodeset, 92
- hwloc_topology_ignore_all_keep_-
 - structure
 - hwlocality_configuration, 51
- hwloc_topology_ignore_type
 - hwlocality_configuration, 51
- hwloc_topology_ignore_type_keep_-
 - structure
 - hwlocality_configuration, 52
- hwloc_topology_init
 - hwlocality_creation, 47
- hwloc_topology_insert_misc_object_-
 - by_cpuset
 - hwlocality_tinker, 55
- hwloc_topology_insert_misc_object_-
 - by_parent
 - hwlocality_tinker, 56
- hwloc_topology_is_thissystem
 - hwlocality_information, 58
- hwloc_topology_load
 - hwlocality_creation, 47
- hwloc_topology_membind_support,
 - 137
 - alloc_membind, 137
 - bind_membind, 137
 - firsttouch_membind, 137
 - get_area_membind, 138
 - get_proc_membind, 138
 - get_thisproc_membind, 138
 - get_thisthread_membind, 138
 - interleave_membind, 138
 - migrate_membind, 138
 - nexttouch_membind, 138
 - replicate_membind, 138
 - set_area_membind, 139
 - set_proc_membind, 139
 - set_thisproc_membind, 139
 - set_thisthread_membind, 139
- hwloc_topology_set_flags
 - hwlocality_configuration, 52
- hwloc_topology_set_fsroot
 - hwlocality_configuration, 52
- hwloc_topology_set_pid
 - hwlocality_configuration, 52
- hwloc_topology_set_synthetic
 - hwlocality_configuration, 53
- hwloc_topology_set_xml
 - hwlocality_configuration, 53
- hwloc_topology_set_xmlbuffer
 - hwlocality_configuration, 53
- hwloc_topology_support, 140
 - cpubind, 140
 - discovery, 140
 - membind, 140
- hwloc_topology_t
 - hwlocality_topology, 40
- HWLOC_TYPE_DEPTH_-
 - MULTIPLE
 - hwlocality_information, 57
- HWLOC_TYPE_DEPTH_-
 - UNKNOWN
 - hwlocality_information, 57
- HWLOC_TYPE_UNORDERED
 - hwlocality_types, 43
- hwlocality_api_version
 - HWLOC_API_VERSION, 39
- hwlocality_bitmap
 - hwloc_bitmap_allbut, 99
 - hwloc_bitmap_alloc, 99
 - hwloc_bitmap_alloc_full, 100
 - hwloc_bitmap_and, 100
 - hwloc_bitmap_andnot, 100
 - hwloc_bitmap_asprintf, 100
 - hwloc_bitmap_clr, 100
 - hwloc_bitmap_clr_range, 100
 - hwloc_bitmap_compare, 100
 - hwloc_bitmap_compare_first,
 - 101
 - hwloc_bitmap_copy, 101
 - hwloc_bitmap_dup, 101
 - hwloc_bitmap_fill, 101
 - hwloc_bitmap_first, 101
 - hwloc_bitmap_foreach_begin, 99
 - hwloc_bitmap_foreach_end, 99

- hwloc_bitmap_free, 101
- hwloc_bitmap_from_ith_ulong, 101
- hwloc_bitmap_from_ulong, 102
- hwloc_bitmap_intersects, 102
- hwloc_bitmap_isequal, 102
- hwloc_bitmap_isfull, 102
- hwloc_bitmap_isincluded, 102
- hwloc_bitmap_isset, 102
- hwloc_bitmap_iszero, 102
- hwloc_bitmap_last, 102
- hwloc_bitmap_next, 103
- hwloc_bitmap_not, 103
- hwloc_bitmap_only, 103
- hwloc_bitmap_or, 103
- hwloc_bitmap_set, 103
- hwloc_bitmap_set_ith_ulong, 103
- hwloc_bitmap_set_range, 103
- hwloc_bitmap_singlify, 103
- hwloc_bitmap_snprintf, 104
- hwloc_bitmap_sscanf, 104
- hwloc_bitmap_t, 99
- hwloc_bitmap_taskset_asprintf, 104
- hwloc_bitmap_taskset_snprintf, 104
- hwloc_bitmap_taskset_sscanf, 104
- hwloc_bitmap_to_ith_ulong, 104
- hwloc_bitmap_to_ulong, 105
- hwloc_bitmap_weight, 105
- hwloc_bitmap_xor, 105
- hwloc_bitmap_zero, 105
- hwloc_const_bitmap_t, 99
- hwlocality_configuration
 - HWLOC_TOPOLOGY_FLAG_-IS_THISSYSTEM, 51
 - HWLOC_TOPOLOGY_FLAG_-WHOLE_SYSTEM, 51
- hwlocality_configuration
 - hwloc_topology_flags_e, 51
 - hwloc_topology_get_support, 51
 - hwloc_topology_ignore_all_-keep_structure, 51
 - hwloc_topology_ignore_type, 51
 - hwloc_topology_ignore_type_-keep_structure, 52
 - hwloc_topology_set_flags, 52
 - hwloc_topology_set_fsroot, 52
 - hwloc_topology_set_pid, 52
 - hwloc_topology_set_synthetic, 53
 - hwloc_topology_set_xml, 53
 - hwloc_topology_set_xmlbuffer, 53
- hwlocality_conversion
 - hwloc_obj_attr_snprintf, 61
 - hwloc_obj_cpuset_snprintf, 61
 - hwloc_obj_get_info_by_name, 62
 - hwloc_obj_snprintf, 62
 - hwloc_obj_type_of_string, 62
 - hwloc_obj_type_snprintf, 62
 - hwloc_obj_type_string, 63
- hwlocality_cpupinding
 - HWLOC_CPUBIND_-NOMEMBIND, 66
 - HWLOC_CPUBIND_-PROCESS, 65
 - HWLOC_CPUBIND_STRICT, 65
 - HWLOC_CPUBIND_THREAD, 65
- hwlocality_cpupinding
 - hwloc_cpupind_flags_t, 65
 - hwloc_get_cpupind, 66
 - hwloc_get_proc_cpupind, 66
 - hwloc_get_thread_cpupind, 66
 - hwloc_set_cpupind, 67
 - hwloc_set_proc_cpupind, 67
 - hwloc_set_thread_cpupind, 67
- hwlocality_creation
 - hwloc_topology_check, 47
 - hwloc_topology_destroy, 47
 - hwloc_topology_init, 47
 - hwloc_topology_load, 47
- hwlocality_cuda
 - hwloc_cuda_get_device_cpuset, 115
- hwlocality_cudart

- hwloc_cudart_get_device_-
cpuset, 116
- hwlocality_glibc_sched
 - hwloc_cpuset_from_glibc_-
sched_affinity, 106
 - hwloc_cpuset_to_glibc_sched_-
affinity, 106
- hwlocality_helper_binding
 - hwloc_alloc_membind_policy,
88
 - hwloc_alloc_membind_policy_-
nodeset, 88
 - hwloc_distribute, 88
 - hwloc_distributev, 89
- hwlocality_helper_cpuset
 - hwloc_topology_get_allowed_-
cpuset, 90
 - hwloc_topology_get_complete_-
cpuset, 90
 - hwloc_topology_get_online_-
cpuset, 90
 - hwloc_topology_get_topology_-
cpuset, 91
- hwlocality_helper_find_cache
 - hwloc_get_cache_covering_-
cpuset, 85
 - hwloc_get_shared_cache_-
covering_obj, 85
- hwlocality_helper_find_covering
 - hwloc_get_child_covering_-
cpuset, 83
 - hwloc_get_obj_covering_cpuset,
83
- hwlocality_helper_find_coverings
 - hwloc_get_next_obj_covering_-
cpuset_by_depth, 84
 - hwloc_get_next_obj_covering_-
cpuset_by_type, 84
- hwlocality_helper_find_inside
 - hwloc_get_first_largest_obj_-
inside_cpuset, 81
 - hwloc_get_largest_objs_inside_-
cpuset, 81
 - hwloc_get_nbobjs_inside_-
cpuset_by_depth, 81
- hwloc_get_nbobjs_inside_-
cpuset_by_type, 81
- hwloc_get_next_obj_inside_-
cpuset_by_depth, 81
- hwloc_get_next_obj_inside_-
cpuset_by_type, 82
- hwloc_get_obj_inside_cpuset_-
by_depth, 82
- hwloc_get_obj_inside_cpuset_-
by_type, 82
- hwlocality_helper_nodeset
 - hwloc_topology_get_allowed_-
nodeset, 92
 - hwloc_topology_get_complete_-
nodeset, 92
 - hwloc_topology_get_topology_-
nodeset, 92
- hwlocality_helper_nodeset_convert
 - hwloc_cpuset_from_nodeset, 93
 - hwloc_cpuset_from_nodeset_-
strict, 94
 - hwloc_cpuset_to_nodeset, 94
 - hwloc_cpuset_to_nodeset_strict,
94
- hwlocality_helper_traversal
 - hwloc_get_closest_objs, 86
 - hwloc_get_obj_below_array_-
by_type, 86
 - hwloc_get_obj_below_by_type,
87
- hwlocality_helper_traversal_basic
 - hwloc_get_ancestor_obj_by_-
depth, 78
 - hwloc_get_ancestor_obj_by_-
type, 78
 - hwloc_get_common_ancestor_-
obj, 78
 - hwloc_get_next_child, 78
 - hwloc_get_next_obj_by_depth,
78
 - hwloc_get_next_obj_by_type, 78
 - hwloc_get_pu_obj_by_os_index,
78
 - hwloc_get_root_obj, 79
 - hwloc_obj_is_in_subtree, 79
- hwlocality_helper_types

- hwloc_get_type_or_above_-
depth, 76
- hwloc_get_type_or_below_-
depth, 76
- hwlocality_information
 - HWLOC_TYPE_DEPTH_-
MULTIPLE, 57
 - HWLOC_TYPE_DEPTH_-
UNKNOWN, 57
- hwlocality_information
 - hwloc_get_depth_type, 58
 - hwloc_get_nobjs_by_depth, 58
 - hwloc_get_nobjs_by_type, 58
 - hwloc_get_type_depth, 58
 - hwloc_get_type_depth_e, 57
 - hwloc_topology_get_depth, 58
 - hwloc_topology_is_thissystem,
58
- hwlocality_linux
 - hwloc_linux_get_tid_cpupbind,
107
 - hwloc_linux_parse_cpumap_file,
107
 - hwloc_linux_set_tid_cpupbind,
107
- hwlocality_linux_libnuma_bitmask
 - hwloc_cpuset_from_linux_-
libnuma_bitmask, 111
 - hwloc_cpuset_to_linux_-
libnuma_bitmask, 111
 - hwloc_nodeset_from_linux_-
libnuma_bitmask, 111
 - hwloc_nodeset_to_linux_-
libnuma_bitmask, 112
- hwlocality_linux_libnuma_nodemask
 - hwloc_cpuset_from_linux_-
libnuma_nodemask, 113
 - hwloc_cpuset_to_linux_-
libnuma_nodemask, 113
 - hwloc_nodeset_from_linux_-
libnuma_nodemask, 113
 - hwloc_nodeset_to_linux_-
libnuma_nodemask, 114
- hwlocality_linux_libnuma_ulongs
 - hwloc_cpuset_from_linux_-
libnuma_ulongs, 109
 - hwloc_cpuset_to_linux_-
libnuma_ulongs, 109
 - hwloc_nodeset_from_linux_-
libnuma_ulongs, 110
 - hwloc_nodeset_to_linux_-
libnuma_ulongs, 110
- hwlocality_membinding
 - HWLOC_MEMBIND_BIND, 71
 - HWLOC_MEMBIND_-
DEFAULT, 71
 - HWLOC_MEMBIND_-
FIRSTTOUCH, 71
 - HWLOC_MEMBIND_-
INTERLEAVE, 71
 - HWLOC_MEMBIND_-
MIGRATE, 71
 - HWLOC_MEMBIND_-
NEXTTOUCH, 71
 - HWLOC_MEMBIND_-
NOCPUBIND, 71
 - HWLOC_MEMBIND_-
PROCESS, 70
 - HWLOC_MEMBIND_-
REPLICATE, 71
 - HWLOC_MEMBIND_STRICT,
71
 - HWLOC_MEMBIND_-
THREAD, 70
- hwlocality_membinding
 - hwloc_alloc, 71
 - hwloc_alloc_membind, 72
 - hwloc_alloc_membind_nodeset,
72
 - hwloc_free, 72
 - hwloc_get_area_membind, 72
 - hwloc_get_area_membind_-
nodeset, 72
 - hwloc_get_membind, 73
 - hwloc_get_membind_nodeset, 73
 - hwloc_get_proc_membind, 73
 - hwloc_get_proc_membind_-
nodeset, 73
 - hwloc_membind_flags_t, 70
 - hwloc_membind_policy_t, 71
 - hwloc_set_area_membind, 73

- hwloc_set_area_membind_-nodeset, 73
- hwloc_set_membind, 74
- hwloc_set_membind_nodeset, 74
- hwloc_set_proc_membind, 74
- hwloc_set_proc_membind_-nodeset, 74
- hwlocality_myriexpress
 - hwloc_mx_board_get_device_-cpuset, 118
 - hwloc_mx_endpoint_get_-device_cpuset, 118
- hwlocality_objects
 - hwloc_obj_t, 46
- hwlocality_openfabrics
 - hwloc_ibv_get_device_cpuset, 117
- hwlocality_sets
 - hwloc_const_cpuset_t, 41
 - hwloc_const_nodeset_t, 41
 - hwloc_cpuset_t, 41
 - hwloc_nodeset_t, 41
- hwlocality_tinker
 - hwloc_topology_export_xml, 55
 - hwloc_topology_export_-xmlbuffer, 55
 - hwloc_topology_insert_misc_-object_by_cpuset, 55
 - hwloc_topology_insert_misc_-object_by_parent, 56
- hwlocality_topology
 - hwloc_topology_t, 40
- hwlocality_traversal
 - hwloc_get_obj_by_depth, 60
 - hwloc_get_obj_by_type, 60
- hwlocality_types
 - HWLOC_OBJ_CACHE, 44
 - HWLOC_OBJ_CORE, 44
 - HWLOC_OBJ_GROUP, 44
 - HWLOC_OBJ_MACHINE, 43
 - HWLOC_OBJ_MISC, 44
 - HWLOC_OBJ_NODE, 44
 - HWLOC_OBJ_PU, 44
 - HWLOC_OBJ_SOCKET, 44
 - HWLOC_OBJ_SYSTEM, 43
- HWLOC_TYPE_-UNORDERED, 43
- hwlocality_types
 - hwloc_compare_types, 44
 - hwloc_compare_types_e, 43
 - hwloc_obj_type_t, 43
- infos
 - hwloc_obj, 123
- infos_count
 - hwloc_obj, 123
- interleave_membind
 - hwloc_topology_membind_-support, 138
- last_child
 - hwloc_obj, 124
- linesize
 - hwloc_obj_attr_u::hwloc_-cache_attr_s, 128
- linux-libnuma.h, 164
- Linux-only helpers, 107
- linux.h, 166
- local_memory
 - hwloc_obj_memory_s, 131
- logical_index
 - hwloc_obj, 124
- membind
 - hwloc_topology_support, 140
- memory
 - hwloc_obj, 124
- Memory binding, 68
- migrate_membind
 - hwloc_topology_membind_-support, 138
- myriexpress.h, 167
- Myrinet Express-Specific Functions, 118
- name
 - hwloc_obj, 124
 - hwloc_obj_info_s, 130
- next_cousin
 - hwloc_obj, 124
- next_sibling

- hwloc_obj, [124](#)
- nexttouch_membind
 - hwloc_topology_membind_-
support, [138](#)
- nodeset
 - hwloc_obj, [124](#)
- Nodeset Helpers, [92](#)
- Object sets, [41](#)
- Object Type Helpers, [76](#)
- Object/String Conversion, [61](#)
- online_cpuset
 - hwloc_obj, [125](#)
- OpenFabrics-Specific Functions, [117](#)
- openfabrics-verbs.h, [168](#)
- os_index
 - hwloc_obj, [125](#)
- os_level
 - hwloc_obj, [125](#)
- page_types
 - hwloc_obj_memory_s, [131](#)
- page_types_len
 - hwloc_obj_memory_s, [131](#)
- parent
 - hwloc_obj, [125](#)
- prev_cousin
 - hwloc_obj, [125](#)
- prev_sibling
 - hwloc_obj, [125](#)
- pu
 - hwloc_topology_discovery_-
support, [136](#)
- replicate_membind
 - hwloc_topology_membind_-
support, [138](#)
- Retrieve Objects, [60](#)
- set_area_membind
 - hwloc_topology_membind_-
support, [139](#)
- set_proc_cpupbind
 - hwloc_topology_cpupbind_-
support, [134](#)
- set_proc_membind
 - hwloc_topology_membind_-
support, [139](#)
- set_thisproc_cpupbind
 - hwloc_topology_cpupbind_-
support, [135](#)
- set_thisproc_membind
 - hwloc_topology_membind_-
support, [139](#)
- set_thisthread_cpupbind
 - hwloc_topology_cpupbind_-
support, [135](#)
- set_thisthread_membind
 - hwloc_topology_membind_-
support, [139](#)
- set_thread_cpupbind
 - hwloc_topology_cpupbind_-
support, [135](#)
- sibling_rank
 - hwloc_obj, [125](#)
- size
 - hwloc_obj_attr_u::hwloc_-
cache_attr_s, [128](#)
 - hwloc_obj_memory_s::hwloc_-
obj_memory_page_type_s,
[133](#)
- The bitmap API, [95](#)
- Tinker with topologies., [55](#)
- Topology context, [40](#)
- Topology Object Types, [43](#)
- Topology Objects, [46](#)
- total_memory
 - hwloc_obj_memory_s, [131](#)
- type
 - hwloc_obj, [125](#)
- userdata
 - hwloc_obj, [126](#)
- value
 - hwloc_obj_info_s, [130](#)