# Intel® Xeon® CPU Max Series Configuration and Tuning Guide

# REVISION HISTORY

| Date | Revision | Description |
|:---:|:---:|:---|
| February 2023 | 001 | Initial release of document |

# Contents

## TABLES

# FIGURES

# CHAPTER 1
# INTRODUCTION

## 1.1    AUDIENCE

This document is for system administrators and application engineers running and optimizing applications on the Intel® Xeon® CPU Max Series.

## 1.2    GLOSSARY

*Table 1. Acronym Definition*

| Acronym | Term | Definition |
|---|---|---|
| BIOS | Basic Input Output Service | |
| HBM | High Bandwidth Memory | |
| 1LM | 1-Level Memory Mode or FLAT Mode | Mode where HBM and DDR are exposed to the software as separate address spaces. |
| 2LM | 2-Level Memory (2LM) mode or Cache Mode | Mode where HBM is used as a memory side cache for DDR. In this mode, only DDR address space is visible to software and HBM functions as a transparent memory-side cache for DDR. |
| | "fake" NUMA Node | This feature is enabled using a Linux kernel boot option (numa=fake). It allows the physical memory of a system to be divided into "fake" NUMA nodes. In other words, with fake-NUMA, a physical NUMA node, which is a uniform physical memory region, can be exposed as multiple NUMA nodes to applications. |

## 1.3    REFERENCES

*Table 2. References*

| Description | URL |
|---|---|
| Intel® Architecture Instruction Set Extensions Programming Reference | https://software.intel.com/content/www/us/en/develop/download/intel-architecture-instruction-set extensions-programming-reference.html |
| memkind library | http://memkind.github.io/memkind/ |
| libnuma API | https://man7.org/linux/man-pages/man3/numa.3.html |
| hbwmalloc API | http://memkind.github.io/memkind/man_pages/hbwmalloc.html |
| Intel® MLC | https://www.intel.com/content/www/us/en/developer/articles/tool/intelr-memory-latency-checker.html |

| STREAM benchmark | https://www.cs.virginia.edu/stream/ |
|---|---|
| Intel® oneAPI Math Kernel Library (oneMKL) | https://www.intel.com/content/www/us/en/developer/tools/oneapi/onemkl-download.html |
| Developer Guide for Intel® oneAPI Math Kernel Library for Linux* | https://www.intel.com/content/www/us/en/develop/documentation/onemkl-linux-developer-guide/top/intel-oneapi-math-kernel-library-benchmarks/intel-distribution-for-linpack-benchmark-1/overview-intel-distribution-for-linpack-benchmark.html |

## 2.1 Processor block diagram
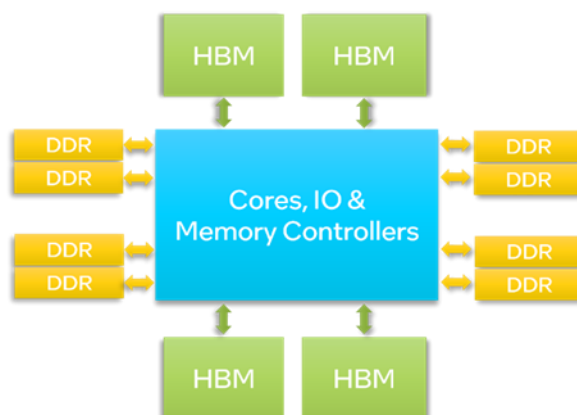


*Figure 1. Processor Block Diagram*

The processors contain four HBM2e stacks totaling 64 GB of High Bandwidth Memory (HBM) capacity per processor in addition to eight channels of DDR memory

Two processors are connected by up to four Intel® Ultra Path Interconnect (Intel® UPI) links in a two-socket system. A two-socket system has a total of 128 GB of HBM capacity.

### 2.1.1 HBM stacks and their specifications
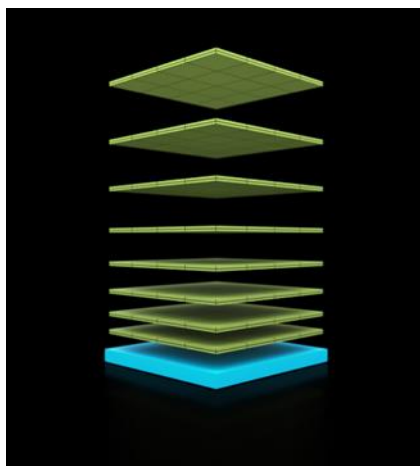


*Figure 2. Memory Stacks Within HBM*

HBM memory is composed of multiple DRAM memory stacks with a wide bus. Each stack contains eight DRAMs stacked on a logic die at the bottom. An Intel® Xeon® CPU Max Series processor has four stacks totaling 64 GB of HBM capacity.

# CHAPTER 3
# HARDWARE CONFIGURATION

## 3.1     CPU CONFIGURATION

The HBM and DDR memory in an Intel® Xeon® CPU Max Series (Package or Socket) can be configured in three memory modes and two clustering modes.

This section describes each of these modes from a hardware point of view. How these modes are configured with the OS is described in Section 5 below and how applications can use them is described in Section 6 below.

### 3.1.1     Memory Modes
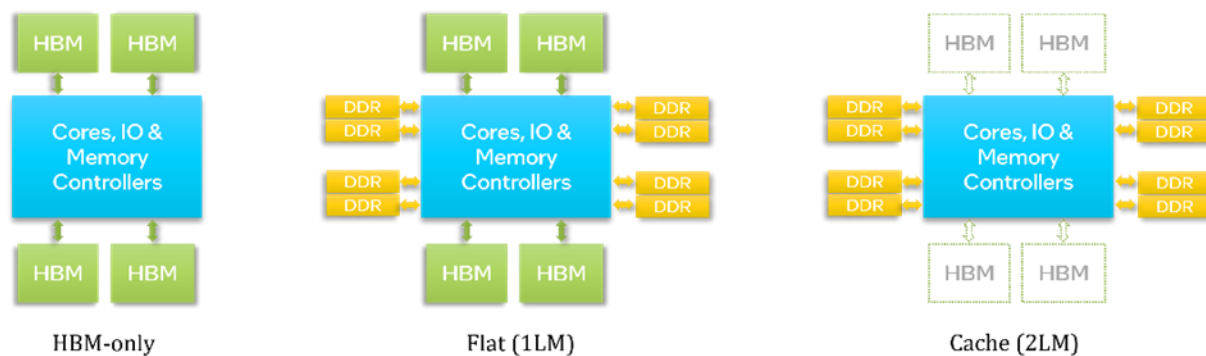


Figure 3. HBM Memory Modes

The processor exposes HBM to software (OS and applications) using three memory modes.

#### 3.1.1.1     HBM-only

When no DDR is installed, HBM-only mode is selected. The only memory available to the OS and applications in this mode is HBM. The OS may see all the installed HBM in this mode, while applications will see what the OS exposes.  Hence the OS and the applications can readily utilize HBM. However, the OS, background services, and applications must share the available HBM capacity (64GB per processor).

#### 3.1.1.2     Flat or 1-Level Memory (1LM) mode

When DDR memory is installed, it is possible to expose both HBM and DDR to software by selecting flat (also known as 1LM) mode from the BIOS menu at boot. HBM and DDR are exposed to software as separate address spaces in this mode. DDR is exposed as a separate address space (NUMA node) and HBM as another address space (NUMA node). Users need to use NUMA-aware tools (e.g., numactl) or libraries to utilize HBM in this mode, as described in Section 6.2 below. Additional OS configuration is necessary before HBM can be accessed as part of the regular memory pool (See section 5.2).

### 3.1.1.3    Cache or 2-Level Memory (2LM) mode

When DDR is installed, it is possible to use HBM as a memory side cache for DDR by selecting Cache (also known as 2LM) mode from the BIOS menu at boot. In this mode, only DDR address space is visible to software and HBM functions as a transparent memory-side cache for DDR. Therefore, applications and command lines do not need modifications to use the cache mode. The HBM is a direct-mapped cache and may require additional configuration steps to minimize conflict misses (see Section 5.2.1).

### 3.1.2    Cluster (partitioning) modes
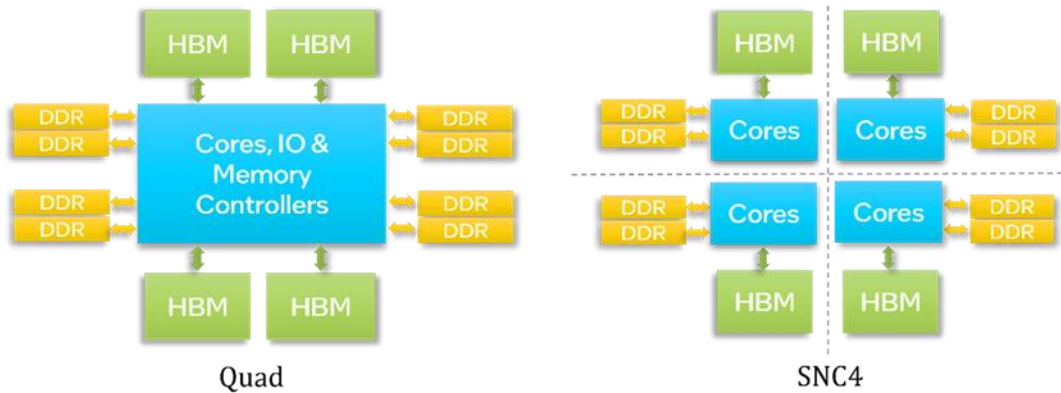


*Figure 4. Cluster Modes*

Cluster modes determine how the processor is partitioned into different address spaces (NUMA nodes).

Clustering (partitioning) allows cores to have higher bandwidth and lower latency to memory (both HBM and DDR) in the same partition.  Cluster modes are orthogonal to memory modes.  Intel® Xeon® CPU Max Series processors have two clustering modes.

### 3.1.2.1    Quadrant

This mode presents a single address space (NUMA node) to software. Therefore, applications do not have to take additional steps to be NUMA aware in this mode. This mode is preferable for applications that share large data structures among all cores of a processor (e.g., an OpenMP application running on all cores and sharing a large data structure).

### 3.1.2.2    SNC4 (Sub-NUMA Clustering-4) the default clustering mode

This mode partitions each CPU into four sub-NUMA cluster partitions. Each partition is exposed to software as one or more NUMA nodes. Therefore, there are at least four NUMA nodes for each processor. Applications should be NUMA-aware to use this mode, but this mode provides higher bandwidth and lower latencies compared to Quadrant mode. This mode is preferable for NUMA-aware applications (e.g., MPI or MPI+OpenMP applications).
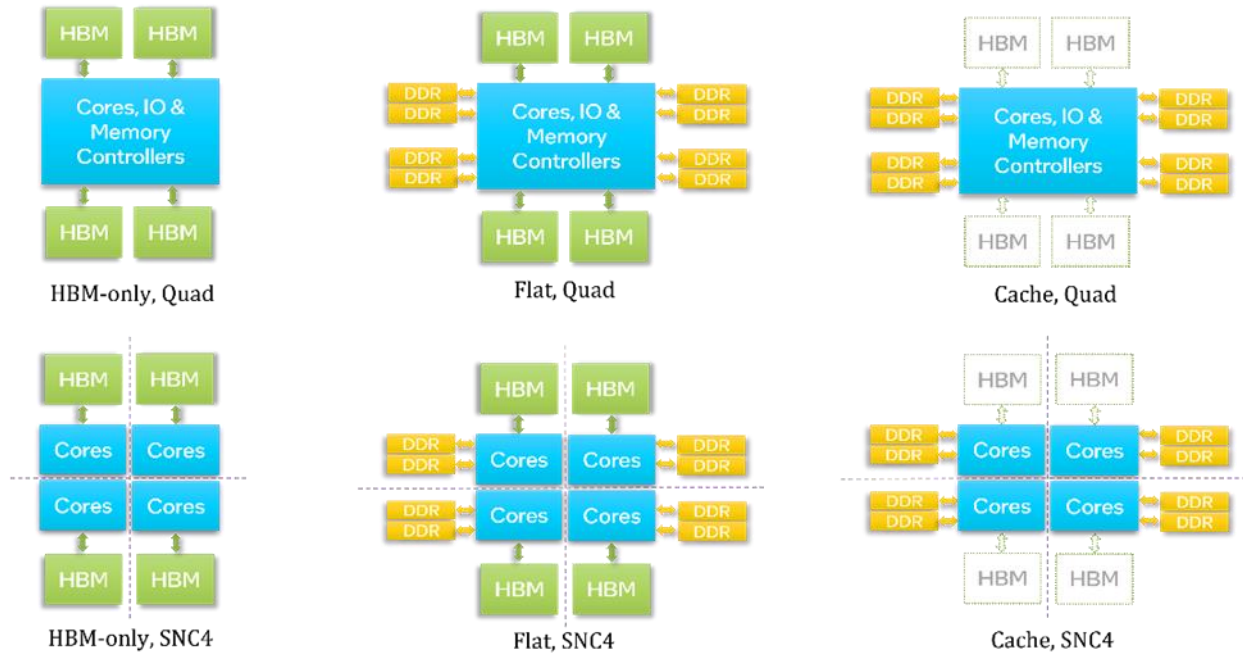
*Figure 5. Memory Mode Configurations*

Combining six memory modes with two clustering modes, Intel® Xeon® CPU Max Series processors have six configuration options, as summarized in Figure 5.

## 3.2　Multi-Socket configuration

Intel® Xeon® CPU Max Series is available in two-socket configurations connected by up to four Intel® UPI links. Each socket (processor) is a separate address space (NUMA node). Therefore, a two-socket system in Quadrant mode has at least two NUMA nodes while a two-socket system in SNC4 has at least eight NUMA nodes.

## 3.3　DIMM configuration

Each processor has four DDR memory controllers, and each memory controller supports two channels for a total of 8-channels per processor.

### 3.3.1　HBM-only mode

To obtain the HBM-only mode, no DIMM must be installed. In some debug BIOS versions, it may be possible to disable the DIMMs via the BIOS options instead of physically removing them.

### 3.3.2　Flat mode

Figure 6 summarizes all DIMM configuration options for flat mode for a single CPU and whether each DIMM configuration supports SNC4 (for HBM, DDR, and HBM+DDR).

| DDR5 +CPS | iMC3 Chan1(7/H) Slot0 | Slot1 | iMC3 Chan0(6/G) Slot0 | Slot1 | iMC2 Chan1(5/F) Slot0 | Slot1 | iMC2 Chan0(4/E) Slot0 | Slot1 | iMC0 Chan0(0/A) Slot1 | Slot0 | iMC0 Chan1(1/B) Slot1 | Slot0 | iMC1 Chan0(2/C) Slot1 | Slot0 | iMC1 Chan1(3/D) Slot1 | Slot0 | SNC2 | SNC4 (DDR & HBM) | Hemi (DDR only) | Quad (HBM) *Excl SNC4 | Quad (DDR) *Excl SNC4 | All2All (DDR only) |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| 0+0 | | | | | | | | | | | | | | | | | | Y | | Y | | |
| 1+0 | | | | | | | | | | DDR5 | | | | | | | | | | Y | | Y |
| 1+0 | | | | | | | | DDR5 | | | | | | | | | | | | Y | | Y |
| 2+0 | | | DDR5 | | | | | | | DDR5 | | | | | | | | | | Y | | Y |
| 2+0 | | | | | | | | DDR5 | | | | | DDR5 | | | | | | | Y | | Y |
| 4+0 | | | DDR5 | | | | DDR5 | | | DDR5 | | | | DDR5 | | | Y | | | Y | Y | |
| 8+0 | DDR5 | | DDR5 | | DDR5 | | DDR5 | | | DDR5 | | DDR5 | | DDR5 | | DDR5 | Y | | | Y | Y | |
| 16+0 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | DDR5 | Y | | | Y | Y | |

*Figure 6. DIMM Configuration Options*

All modes support Quadrant or SNC4 for HBM. However, only the last three modes, which have symmetric DIMM configurations, support SNC4 for both HBM and DDR. In modes with asymmetric DIMM configurations, the DDR space is configured in a special cluster mode called "All-to-All," which produces lower, asymmetric bandwidths and higher, asymmetric latencies.

Populating both DDR slots of a channel (last row) results in lower bandwidth and higher latency than populating only one slot per channel (second to the last row).

### 3.3.3    Cache mode

Cache mode requires a symmetric DIMM configuration across all four memory controllers. Therefore, only the last three rows of the above figure support cache mode.

For best performance dual-rank DIMMs must be used because dual-rank DDR DIMMs provide more bandwidth than single-rank DIMMs.

## 3.4    BIOS settings

This section describes BIOS options for selecting the memory mode and the cluster mode. The menu options shown in this section are for Intel Software Development Platform and the specific menu options **could be different** on your system, depending on the BIOS provider.

Please upgrade to the latest BIOS version if your system has an older BIOS version as the latest BIOS may contain features and performance enhancements.

### 3.4.1    Selecting memory mode

Memory mode is selected in BIOS using the following menu selection:

EDKII Menu -> Socket Configuration -> Memory Configurations -> Memory Map -> Volatile Memory Mode -> 1LM/2LM
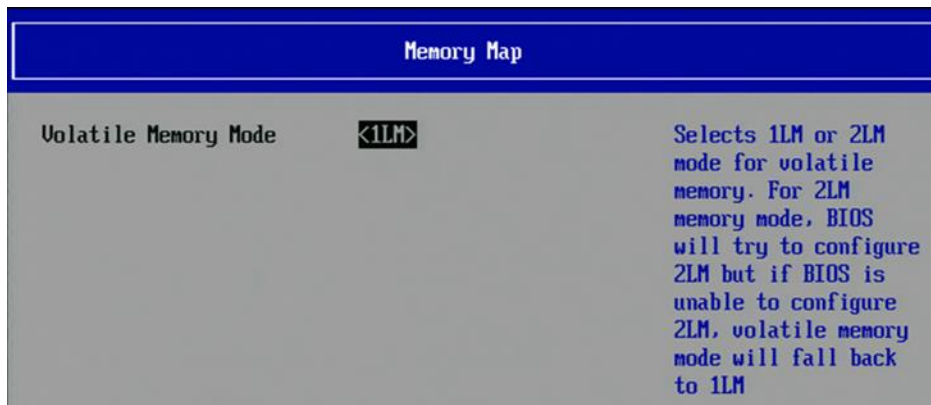
*Figure 7. Selecting Memory Mode*

### 3.4.2　Selecting cluster mode

Cluster mode is selected in BIOS using the following menu selection:

EDKII → Socket configuration → Uncore configuration → Uncore General Configuration → SNC (Sub Numa)
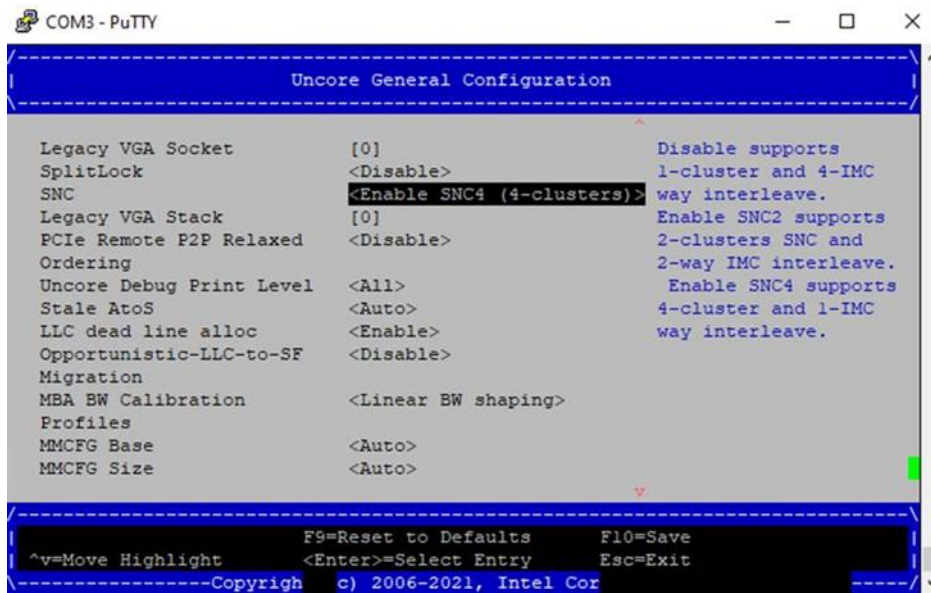


*Figure 8. Selecting Cluster Mode*

Note: Intel SYSCFG utility can be used to save/restore and examine BIOS configuration on Intel server platforms.

# CHAPTER 4
# LINUX System Configuration

## 4.1      Common Configuration Options

Consider the following configuration options for all memory modes:

- Disable swapping. This is especially useful in HBM-only mode, where the capacity is limited. Swapping can severely degrade performance. If an application leads to swapping, consider freeing up memory (e.g., by clearing file system caches) or scaling to more nodes.

- Enable zone-reclaim mode to reduce NUMA misses. This mode is beneficial in situations where the size of a NUMA node is small (e.g., SNC4 clustering mode). This will make the Linux page allocator reclaim easily reusable pages on the requested NUMA node before allocating them on a different NUMA node. This reduces unnecessary NUMA crossings that will degrade performance. However, reclaim activity could introduce small performance variability. Zone-reclaim option can be enabled using the following command. Since it must be done after each reboot, it is recommended that this be automated using initialization scripts.

```
echo 2 > /proc/sys/vm/zone_reclaim_mode
```

- Before each run, consider flushing out file system caches (if cached content from previous runs is not useful) and compacting memory using the following commands. Since these commands need root privileges, system administrators should consider making them part of job prologues of batch systems or providing them as setuid binaries.

```
sync; echo 3 > /proc/sys/vm/drop_caches;
echo 1 > /proc/sys/vm/compact_memory
```

- Consider enabling Transparent Huge Pages (THP). Most HPC applications benefit from THP. Although THP may cause overhead if memory compaction is needed to create large pages, administrators can reduce the overhead by compacting memory before each run as described above.

- Avoid using `/dev/shm` (tmpfs) to store files since it reduces available memory. System administrators should consider clearing `/dev/shm` as part of job prologue to reduce interference between jobs.

- Consider using the latest stable Linux kernel (currently 5.15)

## 4.2      Useful LINUX Tools

### 4.2.1      numactl

```
$ numactl -H
available: 8 nodes (0-7)
node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 112 113 114 115 116 117 118 119 120 121 122 123 124 125
node 0 size: 128734 MB
node 0 free: 128333 MB
node 1 cpus: 14 15 16 17 18 19 20 21 22 23 24 25 26 27 126 127 128 129 130 131 132 133 134 135 136 137 138 139
node 1 size: 129017 MB
node 1 free: 128847 MB
node 2 cpus: 28 29 30 31 32 33 34 35 36 37 38 39 40 41 140 141 142 143 144 145 146 147 148 149 150 151 152 153
node 2 size: 129017 MB
node 2 free: 128834 MB
node 3 cpus: 42 43 44 45 46 47 48 49 50 51 52 53 54 55 154 155 156 157 158 159 160 161 162 163 164 165 166 167
node 3 size: 129017 MB
node 3 free: 128848 MB
node 4 cpus: 56 57 58 59 60 61 62 63 64 65 66 67 68 69 168 169 170 171 172 173 174 175 176 177 178 179 180 181
node 4 size: 128976 MB
node 4 free: 128807 MB
node 5 cpus: 70 71 72 73 74 75 76 77 78 79 80 81 82 83 182 183 184 185 186 187 188 189 190 191 192 193 194 195
node 5 size: 129017 MB
node 5 free: 127974 MB
node 6 cpus: 84 85 86 87 88 89 90 91 92 93 94 95 96 97 196 197 198 199 200 201 202 203 204 205 206 207 208 209
node 6 size: 129017 MB
node 6 free: 128801 MB
node 7 cpus: 98 99 100 101 102 103 104 105 106 107 108 109 110 111 210 211 212 213 214 215 216 217 218 219 220 221 222 223
node 7 size: 129005 MB
node 7 free: 128834 MB
node distances:
node   0    1    2    3    4    5    6    7
  0:  10   12   12   12   21   21   21   21
  1:  12   10   12   12   21   21   21   21
  2:  12   12   10   12   21   21   21   21
  3:  12   12   12   10   21   21   21   21
  4:  21   21   21   21   10   12   12   12
  5:  21   21   21   21   12   10   12   12
  6:  21   21   21   21   12   12   10   12
  7:  21   21   21   21   12   12   12   10
```

*Figure 9. numactl -H Example*

The Linux utility numactl is often used for both observing the NUMA configuration of the system and executing applications on specific NUMA nodes. For observing the NUMA configuration of a system, use numactl -H.  The following figure shows the output of numactl -H showing the number of NUMA nodes, CPU cores, and memory capacity on each NUMA node, followed by a matrix describing the distance of each node from any other node. See man numactl for more information.

## 4.2.2       numastat

The Linux utility numastat provides various statistics about NUMA memory usage. In particular, the following commands are useful (see man numastat for more details).

```
$ numastat -p python

Per-node process memory usage (in MBs)
PID                        Node 0           Node 1           Node 2
----------------   ---------------  ---------------  ---------------
5132 (tuned)                10.54             1.21             1.21
68100 (python3)             11.88             0.00             0.00
----------------   ---------------  ---------------  ---------------
Total                       22.42             1.21             1.21
```

*Figure 10. numastat -p Example*

- numastat -p <process_name gives memory usage of a given process as shown in Figure 11:

- numastat -m shows the memory usage information of an entire system

```
$ numastat
                                node0                node1
numa_hit              7038233469           7552949520
numa_miss               31491495                    0
numa_foreign                   0             31491495
interleave_hit         254896206            254893381
local_node            6905622525           7430407252
other_node             164102439            122542268
```

*Figure 11. numastat -m Example*

- `numastat` (without arguments) shows NUMA hits and misses (cumulative from boot). This is useful in identifying NUMA node crossings (numa_miss) that can lead to unexpected performance degradations. Since these statistics are cumulative from boot, it is necessary to run numastat before and after each run to see whether a given run encountered NUMA misses.

## 4.2.3    turbostat

Turbostat can be used to examine the power, frequency, and temperature of x86 architecture processors when executed as root (or as a setuid binary). Turbostat can be useful in identifying system cooling issues. For instance, the following shows the system's power, frequency, and temperature:

- `turbostat -qS`                # average for both CPUs
- `turbostat -qS --cpu package`  # separately for each CPU

```
$ turbostat -qS
Avg_MHz Busy% Bzy_MHz TSC_MHz IPC    IRQ     SMI     POLL    C1ACPI  C2ACPI  POLL%   C1ACPI% C2ACPI% CPU%c1  CPU%c6  CoreTmp PkgTmp  PkgWatt RAMWatt PKG_%   RAM_%
1337    49.92   2679    1800    1.68  599693  0       5       669     8110    0.00    0.05    49.92   50.08   0.00    75      75      697.92  172.62  190.03  0.00
1341    49.93   2676    1807    1.70  601152  0       7       1038    8791    0.00    0.07    50.07   50.07   0.00    73      73      702.72  174.93  191.08  0.00
```

*Figure 12. turbostat Example*

### 4.2.4    lscpu

This standard Linux utility shows high-level system configuration details, including the NUMA nodes, core counts, base frequency, cache sizes, and CPU flags (features).

### 4.2.5    dmidecode and lshw

These Linux utilities can be used (with root privileges) to inspect hardware components installed, including HBM and DDR modules.

### 4.2.6    htop

The `htop utility` is a tool like standard Linux `top utility`; however, it shows individual CPU core/thread usage using a visual format. This is usually helpful in identifying NUMA usage and MPI rank or OpenMP thread placement. In addition, it shows the memory usage of the system.

### 4.2.7    lstopo

The `lstopo utility` is part of the `hwloc` library, which can be installed as a package using a standard package manager (e.g., `dnf install hwloc`). The `lstopo utility` (or `lstopo-no-graphics`) shows the hardware topology of a system.

# CHAPTER 5
# MEMORY MODE SPECIFIC CONFIGURATIONS

This section describes OS configuration options for each memory mode.

## 5.1    HBM-Only Memory Mode

No additional configuration steps are necessary to use HBM in HBM-only mode. However, since HBM capacity is limited, administrators can take additional steps to reduce memory capacity overheads. Consider performing the following steps:

- Consider reducing unnecessary services (daemons), and drivers started at boot (e.g., VNC servers, print/mail/etc. daemons, performance profiling drivers)

- Consider minimizing the size of the OS file caches and MPI buffers

- Before each run, consider clearing out the file system cache and compacting memory (see Section 4.1 above)

### 5.1.1    NUMA Node Enumeration

For a two-Socket system in HBM-only mode, Figure 13 summarizes the NUMA node configuration in Quadrant and SNC4. The Quadrant mode results in two NUMA nodes (one node for each socket), while the SNC4 mode results in eight NUMA nodes (four per socket). Each NUMA node contains both cores and HBM memory.
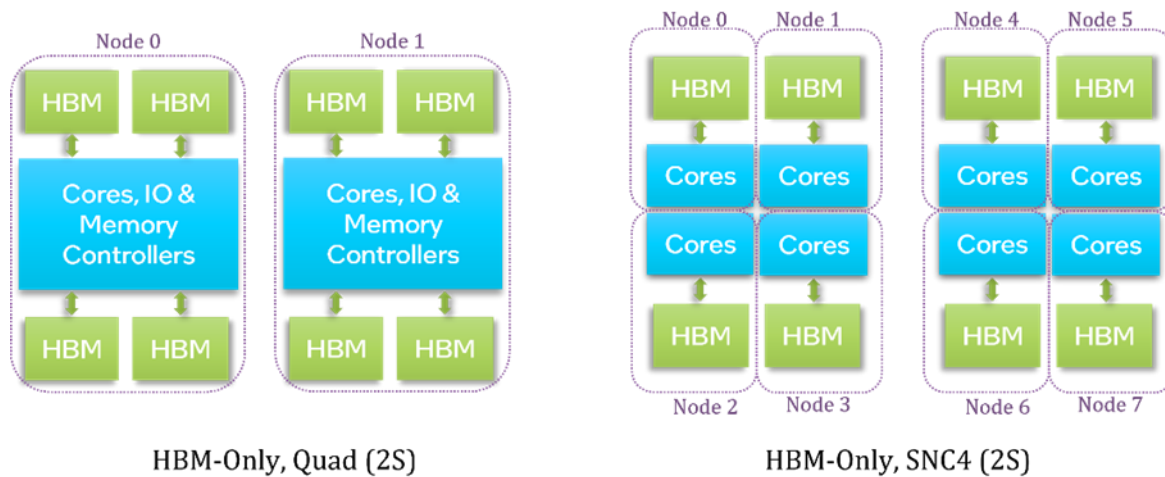


*Figure 13. NUMA Node Configuration for HBM-Only Mode*

Use `numactl –H` to verify the NUMA node configuration and amount of total/free memory on each NUMA node.

## 5.2    Flat Memory Mode

Flat mode is enabled by selecting 1LM in BIOS when DDR is installed as described in Section 3.4.1. However, this alone does not expose HBM to OS and applications via the default memory pool. Additional steps are necessary to enable HBM into the default memory pool.

After selecting 1LM in BIOS, the system boots up with only DDR exposed to the OS and applications. The HBM is still not visible in the default memory pool since HBM is marked as special-purpose memory. This design choice was made to prevent the OS from allocating and reserving valuable HBM memory during the boot process. Since the HBM is "hidden" from the OS during the boot process, the OS cannot allocate or reserve HBM memory.

Here are the steps for booting up and exposing HBM in flat mode:

1.  Select 1LM in the BIOS menu (see Section 3.4.1) and let the OS boot. After the system boots, only DDR is visible in the default memory pool. You can observe that fact using `numactl -H`

2.  Install the following Linux packages:

```
dnf install daxctl ndctl
```

3.  Execute the following daxctl commands for two-socket systems (only the top two commands are required in Quadrant mode, but all are required in SNC4). These commands need root privileges.

```
## Base commands for both Quadrant andSNC4 cluster modes

##

daxctl reconfigure-device -m system-ram dax0.0

daxctl reconfigure-device -m system-ram dax1.0


## For SNC4 cluster mode, use the following additional commands:

##

daxctl reconfigure-device -m system-ram dax2.0

daxctl reconfigure-device -m system-ram dax3.0

daxctl reconfigure-device -m system-ram dax4.0

daxctl reconfigure-device -m system-ram dax5.0

daxctl reconfigure-device -m system-ram dax6.0

daxctl reconfigure-device -m system-ram dax7.0
```

Step 3 needs to be carried out **each time** the system boots. Therefore, it would be convenient to put the above commands in a script that executes when the OS boots.

Use `numactl -H` to verify that the HBM nodes are visible and the entire HBM capacity is free.

## 5.2.1 Flat Mode NUMA Node Enumeration

For a two-Socket system in Flat mode, Figure 14 summarizes the NUMA node configuration in Quadrant and SNC4 modes.

- Quadrant mode results in four NUMA nodes (two DDR nodes with cores and two HBM nodes without cores attached)
- SNC4 mode results in 16 NUMA nodes (8 DDR nodes with cores and 8 HBM nodes without cores attached).
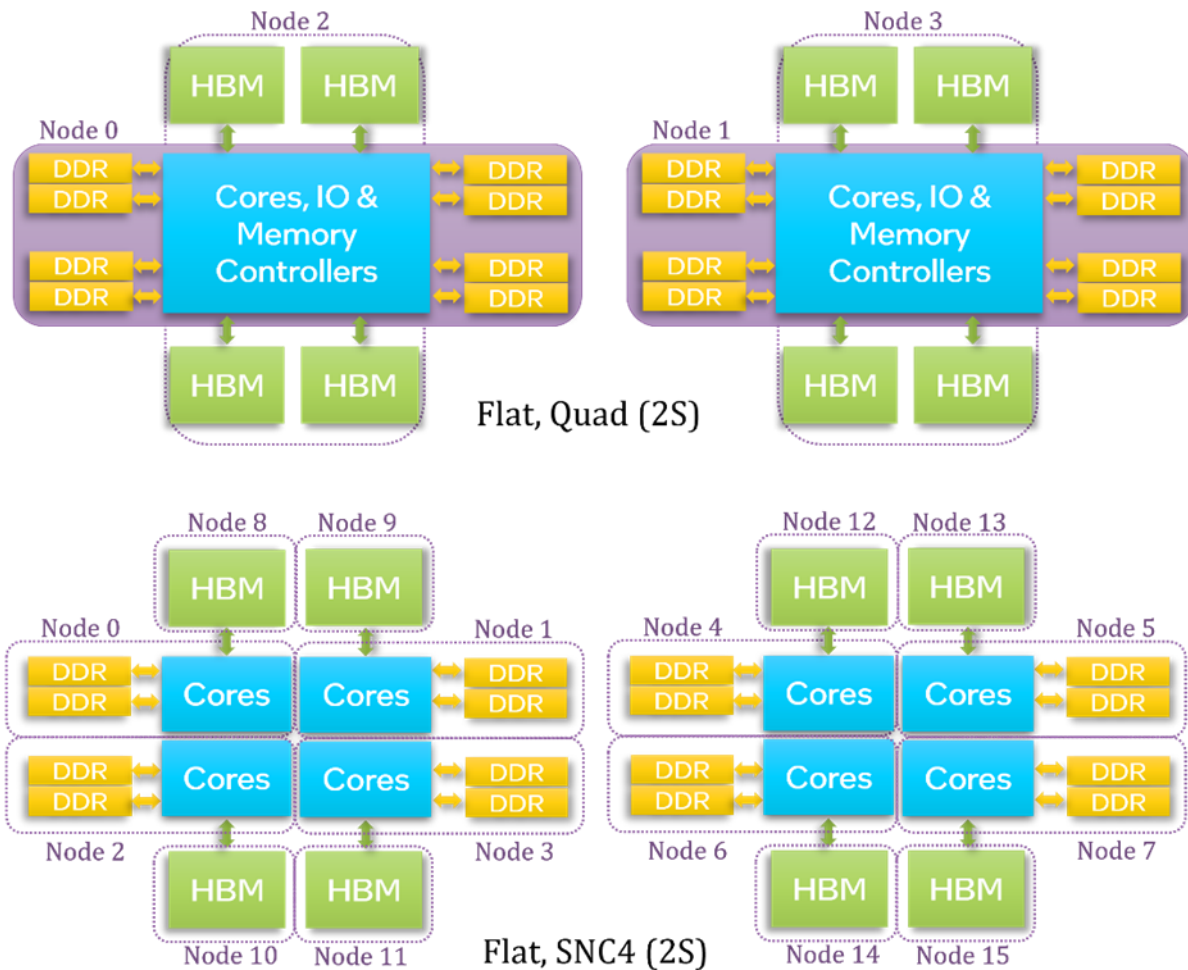


Figure 14. NUMA Node Configurations in Flat Mode

Use `numactl -H` to verify the NUMA node configuration and amount of total/free memory on each NUMA node.

## 5.3    Cache Memory Mode

No additional configuration is necessary to use Cache mode. However, since the HBM cache is a direct-mapped memory-side cache, additional OS configuration with fake-NUMA is strongly recommended to mitigate the effects of conflict misses on applications.

### 5.3.1    Using Fake-NUMA with cache memory mode

This feature is enabled using a Linux kernel boot option (numa=fake). It allows the physical memory of a system to be divided into "fake" NUMA nodes. In other words, with fake-NUMA, a physical NUMA node, which is a uniform physical memory region, can be exposed as multiple NUMA nodes to applications.

For example, consider 64 GB of HBM configured in cache mode with 128GB of DDR memory. The following figure (on the left) shows how two lines in the 128GB DDR address space map to the same location of HBM, creating a conflict in the 64GB HBM cache. In other words, because the HBM cache is direct-mapped, only one of the two lines can exist in the cache.

Figure 15 on the right shows the effect of creating two fake-NUMA nodes. If an application can fit within a fake-NUMA node (say node zero), it is guaranteed not to encounter any conflict misses in the HBM cache.
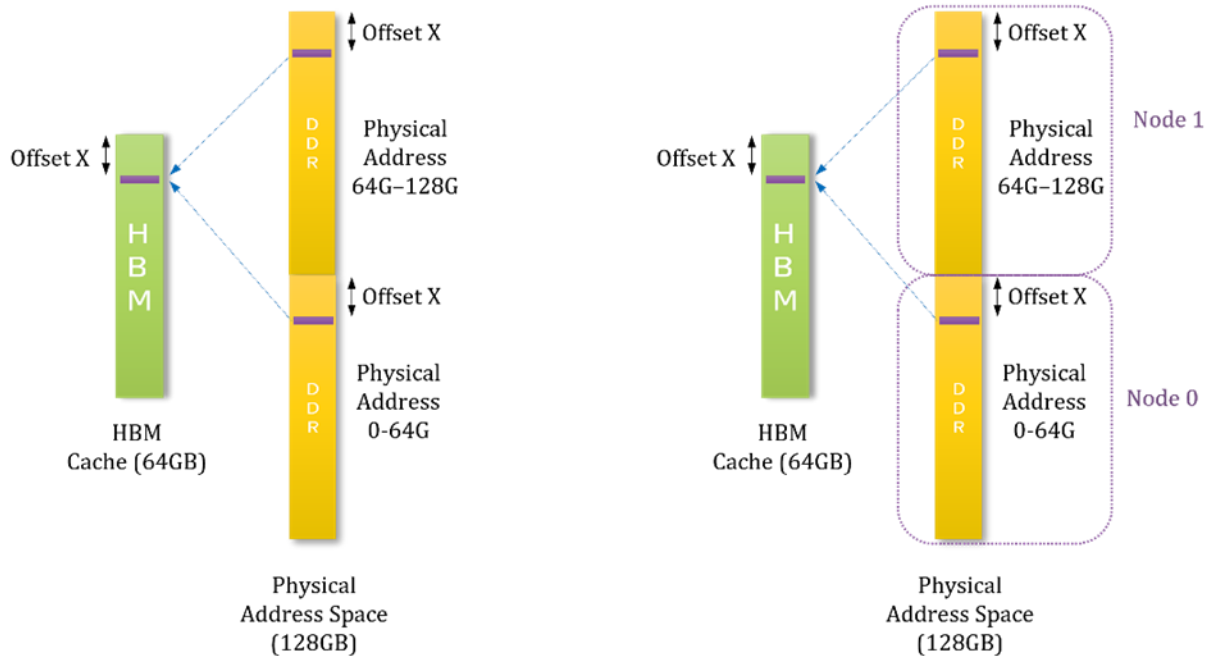


*Figure 15. Fake-NUMA Node Example*

To create two fake-NUMA nodes, each with 64GB capacity, on a system with 128GB of DDR, use

kernel boot option `numa=fake=2U`. This creates two fake-NUMA nodes for each physical NUMA node.

Without fake-NUMA, even applications with footprints smaller than HBM capacity can cause conflict misses in HBM cache due to physical memory fragmentation. When fake-NUMA is enabled, the Linux kernel fills fake-NUMA nodes sequentially. That is, memory is first allocated on NUMA node zero, then on fake-NUMA node one, and so on. This insures conflict-free allocations in the HBM cache for applications that can fit within a fake-NUMA node. Therefore, such applications will see the best possible performance and lower variability.

After the kernel is booted up with the fake-NUMA boot option, verify proper node division using `'numactl -H'`. If fake-NUMA nodes are not visible, ensure that the kernel is built with the kernel config option `'CONFIG_NUMA_EMU=y'`.

In Quadrant cluster mode, the size of a fake-NUMA node should be approximately 64 GB. In SNC4 mode, the size of a fake-NUMA node should be about 16GB.

All fake-NUMA nodes that belong to a physical NUMA node share the same CPU cores. As such, fake-NUMA does not affect application launch commands, although the number of NUMA nodes increases by the ratio between total DDR capacity and total HBM capacity.

It is strongly recommended that **swapping be disabled** when fake-NUMA is used since fake-NUMA could lead to swapping when a fake-NUMA node fills up.

Since fake-NUMA introduces smaller capacity nodes, enabling zone_reclaim with fake-NUMA could cause more frequent reclaim activity when a fake-NUMA node fills up, leading to a small performance variability.

All standard NUMA tools can be used on fake-NUMA nodes. For instance, `'numactl -m 2 ./a.out'` launches an application using the memory of fake-NUMA node two. Similarly, `numastat` will show the properties of fake-NUMA nodes.

## 5.3.2 Page Shuffling (Page Randomization)

Linux provides a feature to randomize page allocations. When fake-NUMA is not used, page shuffling could be useful for achieving more consistent performance results (e.g., between a freshly booted system and a system that has been running for a long time). When page allocations are randomized, pages are allocated at random page addresses in physical memory, changing which pages conflict with each other in the HBM cache, each time an application is launched.

This feature can be enabled in Linux kernel v5.4 or later using kernel boot option `page_alloc.shuffle=y`. Its presence can be checked with file `/sys/module/page_alloc/parameters/shuffle`.

## 5.3.3    Cache memory mode NUMA node enumeration
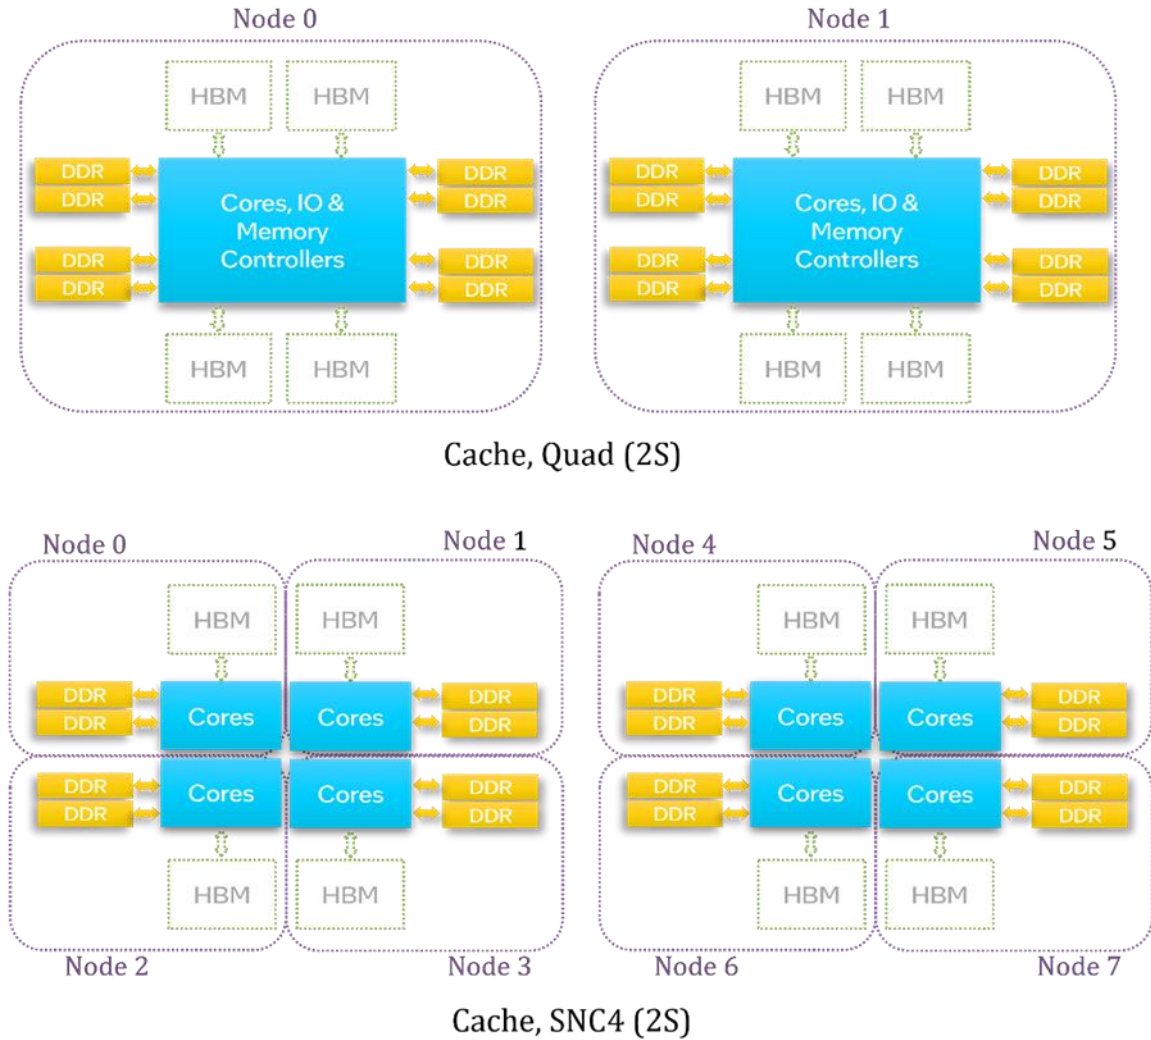


*Figure 16. NUMA Node Configuration in Cache Mode*

For a two-Socket system in Cache mode, the Figure 16 summarizes NUMA node configuration in Quadrant and SNC4. The Quadrant mode results in at least two NUMA nodes (one node for each socket), while the SNC4 mode results in at least eight NUMA nodes (four per socket). Each NUMA node contains both cores and memory.

Use `numactl –H` to verify the NUMA node configuration and amount of total/free memory on each NUMA node.

# CHAPTER 6
## USING MEMORY MODES AND CLUSTER MODES

This section describes how end users can use memory modes and cluster modes.

## 6.1      HBM-Only Memory Mode

No changes to source code or command-line syntax are necessary to use the HBM-only mode. Both the OS and applications use HBM memory, the only available option. However, to fit applications to available memory, applications may have to take some extra steps given below. These are in addition to OS configuration steps given in Section 5.1.

- Balance the number of OpenMP threads with MPI ranks. Since OpenMP threads share memory, utilizing more OpenMP threads can lower the overall memory footprint

- Properly size the OpenMP stack size and MPI communication buffer sizes if an application cannot fit within the HBM capacity

- Free file system cache and compact memory before each run as described in Section 4.1.

- Avoid using `/dev/shm` (tmpfs) to store files since it reduces available memory. Clear files in `/dev/shm` if there are already files from previous jobs.

- Make sure there are no NUMA misses (by running `numastat` before and after a run)

- If an application still fails to fit within HBM capacity after taking the above steps, consider scaling out to more nodes

## 6.2      FLAT Memory Mode

Flat mode exposes DDR and HBM to the user as separate address spaces (NUMA nodes) as described above. The Figure 17 shows the output of `numactl –H` of a flat memory mode system in SNC4 clustering mode.

```
[root@JF5300-B11A345T ~]# numactl -H
available: 16 nodes (0-15)
node 0 cpus: 0 1 2 3 4 5 6 7 8 9 10 11 12 13 112 113 114 115 116 117 118 119 120 121 122 123 124 125
node 0 size: 128725 MB
node 0 free: 127357 MB
node 1 cpus: 14 15 16 17 18 19 20 21 22 23 24 25 26 27 126 127 128 129 130 131 132 133 134 135 136 137 138 139
node 1 size: 128970 MB
node 1 free: 127420 MB
node 2 cpus: 28 29 30 31 32 33 34 35 36 37 38 39 40 41 140 141 142 143 144 145 146 147 148 149 150 151 152 153
node 2 size: 129018 MB
node 2 free: 126361 MB
node 3 cpus: 42 43 44 45 46 47 48 49 50 51 52 53 54 55 154 155 156 157 158 159 160 161 162 163 164 165 166 167
node 3 size: 129018 MB
node 3 free: 127844 MB
node 4 cpus: 56 57 58 59 60 61 62 63 64 65 66 67 68 69 168 169 170 171 172 173 174 175 176 177 178 179 180 181
node 4 size: 129018 MB
node 4 free: 127802 MB
node 5 cpus: 70 71 72 73 74 75 76 77 78 79 80 81 82 83 182 183 184 185 186 187 188 189 190 191 192 193 194 195
node 5 size: 129018 MB
node 5 free: 127788 MB
node 6 cpus: 84 85 86 87 88 89 90 91 92 93 94 95 96 97 196 197 198 199 200 201 202 203 204 205 206 207 208 209
node 6 size: 129018 MB
node 6 free: 126525 MB
node 7 cpus: 98 99 100 101 102 103 104 105 106 107 108 109 110 111 210 211 212 213 214 215 216 217 218 219 220 221 222 223
node 7 size: 128997 MB
node 7 free: 127683 MB
node 8 cpus:
node 8 size: 16384 MB
node 8 free: 16384 MB
node 9 cpus:
node 9 size: 16384 MB
node 9 free: 16384 MB
node 10 cpus:
node 10 size: 16384 MB
node 10 free: 16384 MB
node 11 cpus:
node 11 size: 16384 MB
node 11 free: 16384 MB
node 12 cpus:
node 12 size: 16384 MB
node 12 free: 16384 MB
node 13 cpus:
node 13 size: 16384 MB
node 13 free: 16384 MB
node 14 cpus:
node 14 size: 16384 MB
node 14 free: 16384 MB
node 15 cpus:
node 15 size: 16384 MB
node 15 free: 16384 MB
node distances:
node    0    1    2    3    4    5    6    7    8    9   10   11   12   13   14   15
  0:   10   12   12   12   21   21   21   21   13   14   14   14   23   23   23   23
  1:   12   10   12   12   21   21   21   21   14   13   14   14   23   23   23   23
  2:   12   12   10   12   21   21   21   21   14   14   13   14   23   23   23   23
  3:   12   12   12   10   21   21   21   21   14   14   14   13   23   23   23   23
  4:   21   21   21   21   10   12   12   12   23   23   23   23   13   14   14   14
  5:   21   21   21   21   12   10   12   12   23   23   23   23   14   13   14   14
  6:   21   21   21   21   12   12   10   12   23   23   23   23   14   14   13   14
  7:   21   21   21   21   12   12   12   10   23   23   23   23   14   14   14   13
  8:   13   14   14   14   23   23   23   23   10   14   14   14   23   23   23   23
  9:   14   13   14   14   23   23   23   23   14   10   14   14   23   23   23   23
 10:   14   14   13   14   23   23   23   23   14   14   10   14   23   23   23   23
 11:   14   14   14   13   23   23   23   23   14   14   14   10   23   23   23   23
 12:   23   23   23   23   13   14   14   14   23   23   23   23   10   14   14   14
 13:   23   23   23   23   14   13   14   14   23   23   23   23   14   10   14   14
 14:   23   23   23   23   14   14   13   14   23   23   23   23   14   14   10   14
 15:   23   23   23   23   14   14   14   13   23   23   23   23   14   14   14   10
[root@JF5300-B11A345T ~]# lscpu
```

*Figure 17. numactl -H Example*

As shown in the `numactl -H` output, here are two types of nodes for each socket:

- DDR nodes with CPUs attached (nodes 0-7)
- HBM nodes without any CPUs attached (nodes 8-15)

When an application is launched on CPU cores, memory allocations go to the NUMA nodes that are closest to them, as determined by the 'node distances'. For instance, in the above figure, for CPUs in node zero, the closest memory (distance ten) is in node zero, which is DDR attached to node zero. To use a different kind of memory, users need to use the features of libnuma, made available in the following four ways:

- numactl utility
- Intel® MPI Library
- OpenMP library

- [libnuma API](#)
- [memkind library](#)

The first two methods are used to place the entire application (program code, static data, heap, stack) in HBM, whereas the last two methods can be used to place dynamically allocated (heap-allocated) individual data structures in HBM.

## 6.2.1     Using numactl For HBM Placement of Entire Application

The standard Linux utility numactl can be used to place an application's memory in a NUMA node. There are several policies to consider:

- membind (`numactl --membind hbm_node1, hbm_node2, ... ./a.out`): This forces all memory for the application to be allocated from the specified nodes. If the application exceeds the capacity of the specified nodes, the application will terminate. As such, the user must guarantee that the application will not exceed the *available* HBM capacity, which could be less than the maximum HBM capacity.

- preferred (`numactl --preferred hbm_node ./a.out`): This will make an application allocate memory first from the specified preferred node, until it fills up. After a node fills up, the subsequent allocations will go to the default node, which is always a DDR node in flat memory mode. Notice that only one preferred node can be specified. Since Linux uses a first-touch policy, the application needs to allocate and touch (e.g., initialize) pages for them to be placed in the preferred node. As an example, on a two-socket system in Quadrant clustering mode, the user can specify a different HBM node for different ranks in mpiexec command using MPI colon syntax:

```
mpiexec -n 1 numactl ---preferred hbm_node1 ./a.out : -n 1 numactl --
preferred hbm_node2 ./a.out
```

- preferred-many (`numactl --preferred-many hbm_node1, hbm_node2, ... ./a.out`): Similar to `--preferred` option but allows multiple preferred nodes. Therefore, this option is especially useful for SNC4 and multiple sockets. However, this requires Linux kernel 5.15+ and numactl 2.0.15+.

- interleaved (`numactl --interleave hbm_node,DDR_node`): This allows interleaving memory between any two NUMA nodes, and is especially useful when the memory footprint is roughly twice as large as the size of HBM. If we interleave between DDR and HBM, the maximum bandwidth we can expect is twice the bandwidth of DDR.

The best approach is usually using the membind policy to place the entire application within HBM when it can fit within the HBM capacity. When that is not possible, scaling out to more nodes should be considered before using preferred or interleaved policies.

### 6.2.1.1     Special Considerations with SNC4

Special attention should be paid when using numactl in SNC4 to accommodate multiple NUMA nodes, based on whether membind or preferred is used:

- membind: When running an MPI application in SNC4 mode, the user can specify all HBM nodes as an argument to numactl (e.g., `mpiexec -np 8 numactl -m 4-7 ./a.out`), and HBM memory will be allocated from the closest node to each rank (process).

- Preferred or interleaved: When using SNC4 in flat mode, if we want to place part of an application in HBM either with preferred or interleaved methods discussed above, we have to resort to MPI's colon syntax, because preferred accepts only 1 NUMA node (unless `--preferred-many` is available) and interleaving should be done with corresponding HBM and DDR node. As an example, to use preferred on a single socket with SNC4, when running 56 MPI ranks, we can use:

```
     mpirun -n 14 numactl -N 0 -p 4 ./a.out : -n 14 numactl -N 1 -p 5 ./a.out : -n 14
numactl -N 2 -p 6 ./a.out : -n 14 numactl -N 3 -p 7 ./a.out
```

Similarly, to interleave, on the same system, we can use:

```
     mpirun -n 14 numactl -N 0 -i 0,4 ./a.out : -n 14 numactl -N 1 -i 1,5 ./a.out :
-n 14 numactl -N 2 -i 2,6 ./a.out : -n 14 numactl -N 3 -i 3,7 ./a.out
```

## 6.2.2     Using Intel MPI for HBM Placement of Entire Application

For MPI applications, environment variable `I_MPI_HBW_POLICY` can allocate HBM for MPI ranks (instead of `numactl`). More information about this environment variable can be found in the reference page for `I_MPI_HBW_POLICY`.

```
    mpirun -genv I_MPI_HBW_POLICY hbw_bind -n 2 ./a.out
    mpirun -genv I_MPI_HBW_POLICY hbw_preferred -n 2 ./a.out
    mpirun -genv I_MPI_HBW_POLICY hbw_interleave -n 2 ./a.out
```

`I_MPI_HBW_POLICY` environment variable also accepts an allocation policy for memory allocated by MPI itself (e.g., MPI buffers). For instance, the following uses hbw_bind policy for both user and MPI library allocations.

```
    mpirun -genv I_MPI_HBW_POLICY hbw_bind,hbw_bind -n 2 ./a.out
```

## 6.2.3     Placing of Individual Data Structures in HBM (In FLAT-Mode)

For finer control, it is possible to place dynamically allocated individual data structures in HBM using the following methods. These methods should be used only when the user needs fine-grained control and

placing the entire application memory in HBM is not possible using `numactl` or MPI environment variables (for instance, when the application exceeds the total amount of HBM capacity).

These methods require source code modifications. Only dynamically allocated data structures (i.e., allocated on the heap) can be placed in HBM. Stack, static data, and code cannot be placed in HBM using these methods.

### 6.2.3.1     Using OpenMP for HBM Placement

This is available in Intel classic compilers (any recent version) and Intel® oneAPI compilers starting with version 2021.3. The compiler's OpenMP pragmas and directives depend on the [memkind library](#) as an interface to libnuma.

This OpenMP feature is also available in gcc version 11 or higher.

**C/C++**

```
#include <omp.h>

float *x = (float *)omp_aligned_alloc(64, N*sizeof(float), omp_high_bw_mem_alloc);

omp_free(x, omp_high_bw_mem_alloc);
```

To get "membind" behavior, set fallback to null_fb or abort_fb

```
omp_alloctrait_t traits[2] = { {omp_atk_alignment, 64}, {omp_atk_fallback, omp_atv_null_fb} };

omp_allocator_handle_t my_high_bw_mem_alloc = omp_init_allocator(omp_high_bw_mem_space, 2, traits);

float *x = (float *)omp_alloc(N*sizeof(float), my_high_bw_mem_alloc);

omp_free(x, my_high_bw_mem_alloc);

omp_destroy_allocator(my_high_bw_mem_alloc);
```

**FORTRAN**

```fortran
real, allocatable ::x(:)

!dir$ omp allocate(x) allocator(omp_high_bw_mem_alloc) align(64)

allocate(x(N))
```

## 6.2.3.2    Using hbwmalloc from memkind Library for HBM Placement

You can use hbwmalloc API  provided by the  memkind library for allocating individual data structures in HBM.

Link with –lmemkind

```c
#include <hbwmalloc.h>

float* x = (float *)hbwmalloc(N * sizeof(float));

hbw_free(x);
```

There is also hbw_posix_mem_align,

```c
#include <hbwmalloc.h>

float* x; hbw_posix_memalign((void**) &x, 64, N * sizeof(float));

hbw_free(x);
```

and an allocator,

```cpp
#include <hbw_allocator.h>

std::vector<float, hbw::allocator<float>> x;
```

Two FORTRAN examples,

```
    !dir$ attributes memkind:hbw :: x

    real, allocatable :: x(:)

    allocate(x(N))



    real, allocatable :: x(:)

    !dir$ memkind : hbw, align:64

    allocate(x(N))
```

## 6.3    Cache Memory Mode

No changes to source code or command-line syntax are necessary to use the cache mode. Since HBM cache is transparent to software, applications see only DDR memory space. Therefore, users can run their applications as if they are using a DDR-only machine.

However, to get the best possible performance, the user should be aware that the HBM acts as a cache for DDR and this HBM cache is direct-mapped. As a result, applications may see conflict misses in the HBM cache. A conflict miss occurs when two DDR addresses map to the same location (set) in the HBM cache. Since the HBM is a direct –mapped memory-side cache, only one of those addresses can be cached at a given time. This can lead to frequent cache misses (not finding the required line in the cache). Cache misses increase memory latency and reduce effective bandwidth.

Since HBM capacity is 64 GB per socket, an application that can fit within this capacity ideally should not incur any conflict misses. However, in practice, this is not the case due to physical memory fragmentation. Operating systems can allocate physical memory from anywhere in the physical address range. Physical memory is not always allocated contiguously starting from address zero. Frequent allocations and deallocations lead to allocated memory not being contiguous. This is called memory fragmentation. When physical memory is fragmented, an application can end up with memory addresses that conflict with each other in the HBM direct-mapped cache (i.e., addresses map to the same location in the cache), even though the total memory footprint of the application is less than the size of HBM.

For applications with memory footprints smaller than 64 GB, we can use a Linux kernel feature called fake-NUMA to avoid these unnecessary conflicts occurring due to physical memory fragmentation, as described in Section 5.3.1. Using fake-NUMA, we can divide the physical memory address space into contiguous 64 GB regions (fake-NUMA nodes). Addresses within a given 64 GB NUMA node are guaranteed to be conflict-free (i.e., they cannot map to the same location in the HBM cache). Therefore, if an application can run within a single fake-NUMA node, it can avoid conflict misses.

If the system is configured with fake-NUMA, no additional steps are necessary to avoid conflict misses for applications with footprints that can fit within the size of HBM. When an application is launched, it will automatically start allocating memory from fake-NUMA node zero, before allocations go to other fake-NUMA nodes. However, the user may place an application in any fake-NUMA node using numactl. In some rare cases, where the application uses almost 64 GB of memory, placing it in fake-NUMA node one, for instance, (instead of zero) can marginally improve performance. This is because the fake-NUMA node zero usually has a little bit less free memory than other fake-NUMA nodes because the OS reserves some memory from the very first fake-NUMA node on each socket. You can see all fake-NUMA nodes using numactl -H.

SNC4 increases the number of fake-NUMA nodes. However, the user does not have to take any extra steps to use fake-NUMA as default behavior guarantees conflict-free placement for applications that can fit within HBM capacity.

If the application's memory footprint exceeds that of the HBM capacity, fake-NUMA continues to allocate memory from fake-NUMA nodes sequentially. This leads to more predictable behavior.

# CHAPTER 7
## APPLICATION CONFIGURATION

This section describes how to configure common benchmarks and tools for Intel® Xeon® CPU Max Series.

## 7.1 Software Environment

Intel® oneAPI toolkits such as Intel® oneAPI Base Toolkit (Base Kit) and Intel® oneAPI HPC Toolkit (HPC Kit) provide compilers, profilers (e.g., Intel® VTune™, Intel® Advisor), and libraries (e.g., Intel® Math Kernel Library (Intel® MKL), Intel® MPI Library) supporting the Intel® Xeon® CPU Max Series.

## 7.2 Smoke Tests

The following benchmarks can be used as smoke tests to verify the proper performance of a system. They should be run after a system is booted to verify performance. In addition, they are fast enough to be run before or after each batch job to verify the expected performance of each system.

The first two tests (Intel® Memory Latency Checker (Intel® MLC) and STREAM measure the memory system performance (bandwidth and latency), whereas HPL measures the floating-point compute performance (GFLOPs). HPCG is mostly sensitive to the memory bandwidth.

### 7.2.1 Intel® Memory Latency Checker (Intel® MLC)

Intel® MLC, provides detailed latency and bandwidth measurements of a single system. The following commands are useful in testing a system:

Peak bandwidth: `mlc --peak_injection_bandwidth -Z -X -t60`

Bandwidth matrix and latency: `mlc`

### 7.2.2 STREAM

The STREAM benchmark provides bandwidth measurements for different routines such as Copy, Scale, Add, and Triad on a single node.

For best performance on Intel® Xeon® CPU Max Series processors, enable software prefetching with the following command line:

```
icc -O3 -xCORE-AVX512 -qopt-zmm-usage=high -mcmodel=large -qopenmp -qopt-
streaming-stores=always -fno-builtin -qopt-prefetch-distance=128,16 -
DSTREAM_ARRAY_SIZE=500000000 -DNTIMES=500 stream.c -o stream
```

Execute the resulting binary using the following command:

```
KMP_HW_SUBSET=1t KMP_AFFINITY=balanced,granularity=core,verbose ./stream
```

### 7.2.3 HPL

The Intel® Distribution for LINPACK* Benchmark is based on modifications and additions to High-Performance LINPACK (HPL).  It is available with Intel® oneAPI Math Kernel Library (oneMKL) public library release (or as a part of Intel® oneAPI Base Toolkit (Base Kit)) with instructions. It measures the

amount of time it takes to factor and solve a random dense system of linear equations in double precision, converts that time into a performance rate (GFLOPS), and tests the results for accuracy.

The benchmark can be run on a single node or a cluster of nodes. Instructions for running the benchmark are available at the above link. As an example, to run in HBM-only mode in SNC4 cluster mode on a single node with two sockets, change the following definitions in `runme_intel64_dynamic` file:

```
export MPI_PROC_NUM=2
export MPI_PER_NODE=2
export NUMA_PER_MPI=4
```

and then run:

```
./runme_intel64_dynamic -p 2 -q 1 -b 384 -n 120000
```

Note: In Cache memory mode with fake-NUMA, `NUMA_PER_MPI` should be equal to the number of fake-NUMA nodes on a socket.

## 7.2.4     HPCG

HPCG benchmark optimized for Intel CPUs (source code and prebuilt binaries) can be downloaded with the latest Intel® oneAPI Math Kernel Library (oneMKL) public library [release](#) (or as a part of Intel® oneAPI Base Toolkit (Base Kit), with the [developer guide](#).

To build from source, use:

```
# source C/C++ compiler, MPI compiler, and MKL library
#

export MKLROOT=/path/to/mkl

export LD_LIBRARY_PATH=${MKLROOT}/lib/intel64:${LD_LIBRARY_PATH}

# build binary for Intel AVX-512 -- bin/xhpcg_skx will be created
#

./configure IMPI_IOMP_SKX
make -j4 MKLROOT=${MKLROOT} MKL_INCLUDE=${MKLROOT}/include
```

To run on a 2-socket system with Intel® Xeon® CPU Max Series processors (each with 56 cores) in HBM-only memory mode and SNC4 cluster mode, you can use the command-line:

```
        # Note: Select the best MPI x OMP decomposition for your case
        #       Following assumes SNC4 (8 NUMA nodes on 2S),
        #       and 14 cores (28 threads) on a NUMA node
        #

        export MKL_NUM_THREADS=28
        export OMP_NUM_THREADS=28

        nprocs_per_node=8
        nnodes=1
        nprocs=$((nnodes*nprocs_per_node))

        problem_size=168        # options: 168, 192, 256
        run_time_in_seconds=100  # 100 used as smoke test.
                                # 1800 is min for official HPCG submission

        export I_MPI_SHM=spr-hbm
        export I_MPI_FABRICS=shm:ofi
        export I_MPI_PIN_DOMAIN=numa
        export I_MPI_DEBUG=10    # print out mpi configuration mapping data

        # for 1 hyper-thread, use 'compact,1,0' instead of 'compact'
        #
        export KMP_AFFINITY=granularity=fine,compact

        echo " ===           nnodes: ${nnodes}"
        echo " ===              ppn: ${nprocs_per_node}"
        echo " ===           nprocs: ${nprocs}"
        echo " === n_omp_per_proc: ${MKL_NUM_THREADS}"
        echo " ===        prob_size: ${problem_size}"
        echo " ===         run_time: ${run_time_in_seconds}"

        # run bin/xhpcg_skx binary (either prebuilt or built by user)
        #
        mpiexec.hydra -genvall -n ${nprocs} -ppn ${nprocs_per_node}  bin/xhpcg_skx -
n$problem_size -t$run_time_in_seconds
```

## 7.3    Finding Out Memory Usage Of An Application

For the best possible performance, it is usually necessary to fit applications within the HBM capacity. To do that, it is important to find the memory footprint of the applications and workloads. The following tools are available for this purpose:

- `top` and `htop` (provides total used/free memory of a system at a given time)

- `numastat -m` (provides total memory on each NUMA node at a given time)

- `numastat -p <binary_name>` (provides memory consumption of a given process at a given time)

- `/usr/bin/time -v <app_cmd_line>` (provides various stats about the application including *maximum resident set size*). This is different from bash built-in 'time' command so path must be specified.

If the application footprint exceeds the available HBM capacity, consider scaling out to more nodes or using the Cache mode.

## 7.4     Optimizing Applications For Memory Bandwidth

Because Intel® Xeon® CPU Max Series processors offer a much higher memory bandwidth compared to previous Intel® Xeon® processors, applications (or routines) that were memory bandwidth bound on previous processors may not be memory bandwidth bound on Intel® Xeon® CPU Max Series processors.

The best way to identify the actual bandwidth utilization of an application is to use memory access analysis of Intel® VTune™ Profiler. If that analysis shows underutilization of memory bandwidth for a given phase or routine of an application, following options could be considered to optimize for memory bandwidth:

Optimize compute: the routine may not be doing computations fast enough to generate sufficient memory bandwidth. Consider optimizing computations such as address calculations and consider vectorization to improve compute throughput. Intel® Advisor can be used to identify vectorization opportunity and improve vectorization.

Optimize memory latency: if the routine produces many memory accesses that miss the last-level (L3) cache but still fails to saturate the HBM bandwidth, the routine is likely to be memory latency bound. Irregular access patterns (e.g., indirect accesses, gathers, scatters) that reduce the effectiveness of hardware prefetchers often lead to high memory latency. Consider using software prefetches for such access patterns. Intel® oneAPI compilers provide compiler flags such as -qopt-prefetch and support explicit prefetch directives and intrinsics (_mm_prefetch in C and mm_prefetch in FORTRAN) that can be used within source code.

Overall architecture and optimization details for Intel® Xeon® processors and Intel® Xeon® CPU Max Series processors can be found in the Intel Software Developer's Manuals and Software Optimization and Reference Manuals and from GitHub.