

Hardware Overview

The Anselm cluster consists of 209 computational nodes named cn[1-209] of which 180 are regular compute nodes, 23 GPU Kepler K20 accelerated nodes, 4 MIC Xeon Phi 5110 accelerated nodes and 2 fat nodes. Each node is a powerful x86-64 computer, equipped with 16 cores (two eight-core Intel Sandy Bridge processors), at least 64GB RAM, and local hard drive. The user access to the Anselm cluster is provided by two login nodes login[1,2]. The nodes are interlinked by high speed InfiniBand and Ethernet networks. All nodes share 320TB /home disk storage to store the user files. The 146TB shared /scratch storage is available for the scratch data.

The Fat nodes are equipped with large amount (512GB) of memory. Virtualization infrastructure provides resources to run long term servers and services in virtual mode. Fat nodes and virtual servers may access 45 TB of dedicated block storage. Accelerated nodes, fat nodes, and virtualization infrastructure are available upon request made by a PI.

Schematic representation of the Anselm cluster. Each box represents a node (computer) or storage capacity:

User-oriented infrastructure Storage Management infrastructure — login1
login2 dm1 —

Rack 01, Switch isw5

cn186	cn187		cn188	cn189
cn181	cn182	cn183	cn184	cn185

Rack 01, Switch isw4

cn29 cn30 cn31 cn32 cn33 cn34 cn35 cn36 cn19 cn20 cn21 cn22 cn23 cn24 cn25
cn26 cn27 cn28

|

Lustre FS

/home320TB

| |Lustre FS

/scratch146TB |

Management nodes Block storage 45 TB Virtualization infrastructure servers ...
Srv node Srv node Srv node ... Rack 01, Switch isw0

cn11 cn12 cn13 cn14 cn15 cn16 cn17 cn18 cn1 cn2 cn3 cn4 cn5 cn6 cn7 cn8 cn9
cn10 Rack 02, Switch isw10

cn73 cn74 cn75 cn76 cn77 cn78 cn79 cn80 cn190 cn191 cn192 cn205 cn206 Rack
02, Switch isw9

cn65 cn66 cn67 cn68 cn69 cn70 cn71 cn72 cn55 cn56 cn57 cn58 cn59 cn60 cn61
cn62 cn63 cn64 Rack 02, Switch isw6

cn47 cn48 cn49 cn50 cn51 cn52 cn53 cn54 cn37 cn38 cn39 cn40 cn41 cn42 cn43
cn44 cn45 cn46 Rack 03, Switch isw15

cn193 cn194 cn195 cn207 cn117 cn118 cn119 cn120 cn121 cn122 cn123 cn124
cn125 cn126 Rack 03, Switch isw14

cn109 cn110 cn111 cn112 cn113 cn114 cn115 cn116 cn99 cn100 cn101 cn102
cn103 cn104 cn105 cn106 cn107 cn108 Rack 03, Switch isw11

cn91 cn92 cn93 cn94 cn95 cn96 cn97 cn98 cn81 cn82 cn83 cn84 cn85 cn86 cn87
cn88 cn89 cn90 Rack 04, Switch isw20

cn173 cn174 cn175 cn176 cn177 cn178 cn179 cn180 cn163 cn164 cn165 cn166
cn167 cn168 cn169 cn170 cn171 cn172 Rack 04, **Switch** isw19

cn155 cn156 cn157 cn158 cn159 cn160 cn161 cn162 cn145 cn146 cn147 cn148
cn149 cn150 cn151 cn152 cn153 cn154 Rack 04, Switch isw16

cn137 cn138 cn139 cn140 cn141 cn142 cn143 cn144 cn127 cn128 cn129 cn130
cn131 cn132 cn133 cn134 cn135 cn136 Rack 05, Switch isw21

cn201	cn202		cn203	cn204
cn196	cn197	cn198	cn199	cn200

Fat node cn208
Fat node cn209
...

The cluster compute nodes cn[1-207] are organized within 13 chassis.

There are four types of compute nodes:

- 180 compute nodes without the accelerator
- 23 compute nodes with GPU accelerator - equipped with NVIDIA Tesla Kepler K20
- 4 compute nodes with MIC accelerator - equipped with Intel Xeon Phi 5110P
- 2 fat nodes - equipped with 512GB RAM and two 100GB SSD drives

More about Compute nodes.

GPU and accelerated nodes are available upon request, see the Resources Allocation Policy.

All these nodes are interconnected by fast InfiniBand class="WYSIWYG_LINK">QDR network and Ethernet network. More about the Network. Every chassis provides Infiniband switch, marked **isw**, connecting all nodes in the chassis, as well as connecting the chassis to the upper level switches.

All nodes share 360TB /home disk storage to store user files. The 146TB shared /scratch storage is available for the scratch data. These file systems are provided by Lustre parallel file system. There is also local disk storage available on all compute nodes /lscratch. [More about

Storage](storage.html).

The user access to the Anselm cluster is provided by two login nodes login1, login2, and data mover node dm1. More about accessing cluster.

The parameters are summarized in the following tables:

In general **Primary purpose High Performance Computing Architecture of compute nodes x86-64 Operating system Linux Compute nodes Totally 209 Processor cores 16 (2x8 cores) RAM min. 64 GB, min. 4 GB per core Local disk drive yes - usually 500 GB Compute network InfiniBand QDR, fully non-blocking, fat-tree w/o accelerator 180, cn[1-180] GPU accelerated 23, cn[181-203] MIC accelerated 4, cn[204-207] Fat compute nodes 2, cn[208-209]**
In total Total theoretical peak performance (Rpeak) 94 Tflop/s Total max. LINPACK performance (Rmax) 73 Tflop/s Total amount of RAM 15.136 TB |Node|Processor|Memory|Accelerator| |—|—|—|—| |w/o accelerator|2x Intel Sandy Bridge E5-2665, 2.4GHz|64GB|-| |GPU accelerated|2x Intel Sandy Bridge E5-2470, 2.3GHz|96GB|NVIDIA Kepler K20| |MIC accelerated|2x Intel Sandy Bridge E5-2470, 2.3GHz|96GB|Intel Xeon Phi P5110| |Fat compute node|2x Intel Sandy Bridge E5-2665, 2.4GHz|512GB|-|

For more details please refer to the Compute nodes, Storage, and Network.