

HPC2013

HPC2013 is a machine is a 781 node machine from which 768 nodes are serving as compute node. This machine had a rank of 130 in the top 500 list published www.top500.org in November 2013. In the initial ratings it had a Rpeak of 307.2 Terra-Flops and Rmax of 282.6 Terra-Flops. Extensive testing of this machine was carried out and we were able to achieve an efficiency of 96% on the Linpack benchmark. The new rated Rmax would be around 292.9 Terra-Flops. It is based on Intel Xeon E5-2670V 2.5 GHz 2 CPU-20-core-IvyBridge on HP-Proliant-SL-230s-Gen8 servers with 128GB of RAM per node E5-2670v2x10 core2.5 GHz. The nodes are connected by Mellanox FDR Infiniband chassis based switches that can provide 56Gbps of throughput. It also has 500 Terra-Bytes of storage with an aggregate performance of around 23Gbps on write and 15 GBps on read. It is divided into a home(13/7 w/r GBps) and a scratch(22/12 w/r GBps) file system. The home file system is around 169 Terra-Bytes and the scratch file system is around 332 Terra-Bytes. It has PBS Pro Scheduler from Altair and is divided into queues as follows:

| queue | nodes | Min-Max nodes | Min-Max cores | Wall time |
|-------------|-------|---------------|---------------|--------------------|
| large | 362 | 6-32 | 120-640 | 96 hours |
| medium | 256 | 2-6 | 40-120 | 96 hours |
| small | 96 | 1-2 | 20-40 | 120 hours |
| mini | 32 | 1-2 | 20-40 | 2 hours |
| hyperthread | 16 | 1-2 | 40-80 | 120 hours |
| workq | 4 | 1 | 1-20 | 24 hours 2 hrs cpu |
| highmem | 5 | 1-1 | 2-20 | 120 hours |
| test | 2 | NA | NA | NA |

HPC2010

HPC2010 is a machine is a 376 node machine from which 368 nodes are serving as compute node. This machine had a rank of 369 in the top 500 list published www.top500.org in June 2010. In the initial ratings it had a Rpeak of 34.05 Terra-Flops and Rmax of 29.01 Terra-Flops. It is based on Intel Xeon X5570 2.93 GHz 2 CPU-8-core-Nehalem on HP-Proliant-BL-280c-Gen6 servers with 48 GB of RAM per node. The nodes are connected by Qlogic QDR Infiniband federated switches that can provide 40Gbps of throughput. It also has 100 Terra-Bytes of storage with an aggregate performance of around 5GBps on write performance. It is divided into a home(1.7/1.3 w/r GBps) and a scratch(3.4/2.4 w/r GBps) file system. The home file system is around 60 Terra-Bytes and the scratch file system is around 40 Terra-Bytes. It has PBS Pro Scheduler from Altair and is divided into queues as follows:

| queue | nodes | Min-Max nodes | Min-Max cores | Wall time |
|--------|-------|---------------|---------------|--------------------|
| large | 184 | 16-32 | 128-256 | 72 hours |
| medium | 100 | 4-12 | 32-96 | 96 hours |
| small | 59 | 1-4 | 2-32 | 120 hours |
| workq | 6 | 1 | 1-6 | 24 hours 2 hrs cpu |
| test | 3 | NA | NA | NA |

This cluster was later augmented with 96 nodes of Intel Xeon E-52670 2.6 GHz 2 CPU-16-core-Sandy-Bridge on HP-Proliant-SL-230s-Gen8 servers with 64 GB of RAM per node that add an additional theoretical 31 Terra-Flops to the above 2010 cluster. PBS Pro is the scheduler of choice. Though it has FDR Infiniband cards it is connected to the QDR Infiniband fabric seamlessly.

| queue | nodes | Min-Max nodes | Min-Max cores | Wall time |
|----------|-------|---------------|---------------|----------------------|
| mediumsb | 47 | 2-6 | 32-96 | 96 hours |
| smallsb | 47 | 1-2 | 2-32 | 120 hours |
| workqsb | 1 | 1 | 1-1 | 24 hours 2 hours cpu |
| testsb | 1 | NA | NA | NA |