

Database Hardware Model

The goal is to maximize the sustained disk bandwidth of the entire system for the minimum cost. The disk bandwidth can be limited by many things:

- The sustained bandwidth from memory to core.
- The sustained bandwidth of the disk controller.
- The sustained bandwidth available to memory (half used to transfer data from disk and half used to transfer data to the cores). Note that this may be much less than the number of cores times the maximum bandwidth per core. For this reason, decreasing the number of cores per box may be advantageous.
- The maximum bandwidth available across all attached spindles, which is in turn limited by the number of disks that can be attached to a box, the maximum sustained disk bandwidth from a single spindle, and the maximum sustained bandwidth of the disk interconnect.

A secondary goal is to maximize the number of random I/O operations per second, again for minimum cost.

These numbers are needed at (at least) three timeframes: DC4 (around 2011), when large-scale database serving will be required; DR1, the first data release of the survey; and DR11, the last data release of the survey, ten years later.

The "Today (GW)" numbers in the table below are based on the GrayWulf configuration. The "Today (eee)" numbers are based on an Asus eee PC with SSD, which may be able to provide similar performance at lower cost and lower power consumption. Both use numbers from Alex Szalay.

Parameter	Unit	Today (GW)	Today (eee)	DC4 (2011)	DR1 (2016)	DR11 (2026)
Sustained bandwidth from disk to 1 core	MB/s	180	100	?	500	1800
Sustained disk bandwidth from 1 spindle	MB/s	50	200	?	500	1800
Sustained bandwidth from disks to 1 controller	MB/s	720	200	?	1200	2600
Sustained bandwidth to/from memory on 1 box	GB/s	2.88	> 0.4	?	5	24
Disks that can be attached to 1 box	disks	30	1	?	120	1000
Disk size for 1 (cheap) spindle	GB	750	256	?	2000	64000
Maximum random I/O rate for 1 spindle	IO/s	30	180	?	256	2048
Memory per box	GB	> 8	> 1	?	64	2048
Number of cores per box	cores	8	2	?	16	512