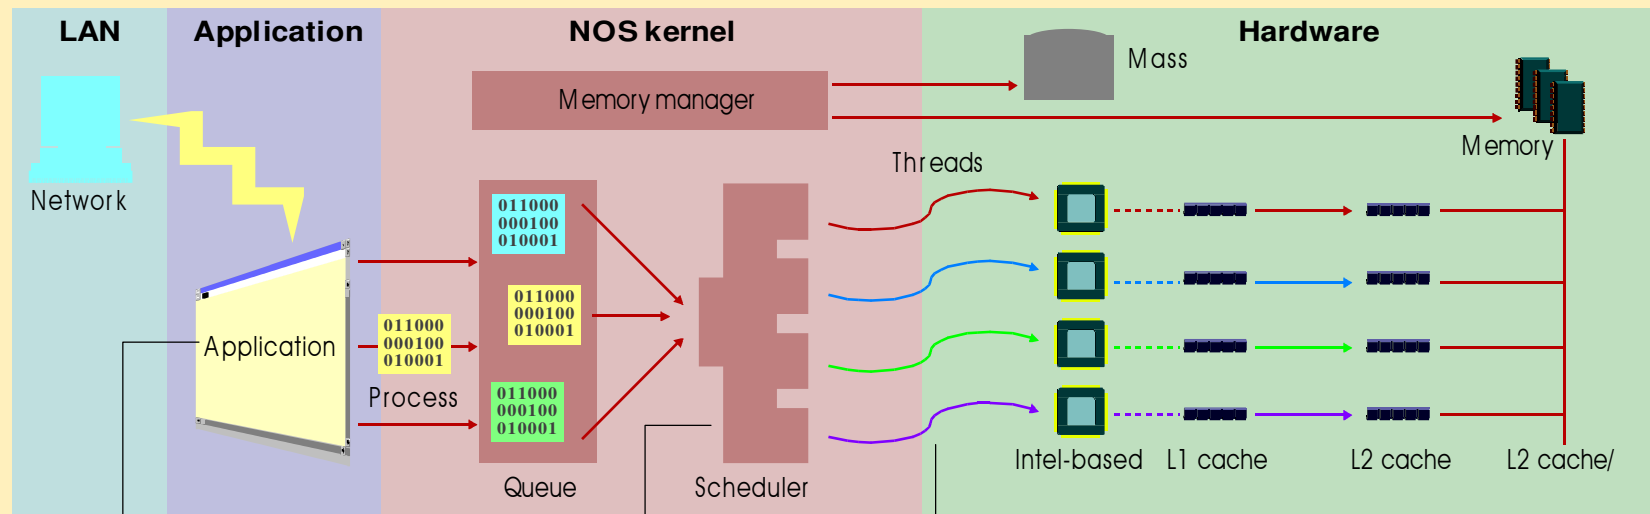


The Complex Art of Multiprocessing

To kick performance-intensive application servers into overdrive, a multiprocessing NOS is essential. Here's how the software does the work.



When an application receives a client request, it issues a process to the kernel's queue.

The scheduler distributes the process, which is made up of multiple instructions called *threads*. The scheduler also looks for high-priority threads, so it can preempt those with low priority.

The scheduler assigns the threads among one or more CPUs, so that no one processor is overburdened.

In NOSs such as Windows NT Server, the memory manager can allocate a protected memory space to each thread, and virtual-memory hard disk space to each process. NT Server also uses asynchronous I/O, which allows further requests to be submitted before the first is complete.