

NUMA Support in ATTO ExpressSAS Host Bus Adapters

What is NUMA?

Non-uniform memory access (NUMA) is a computer memory design that enables multiple processors in the same server to access shared memory over a common, high-speed bus interconnect. Most servers today comprise multiple sockets, each housing a CPU (central processing unit) with multiple processing cores and either local or shared memory. In a NUMA architecture, all memory sharing within the server is managed in a hierarchical manner, a process that improves overall performance by minimizing bus traffic.

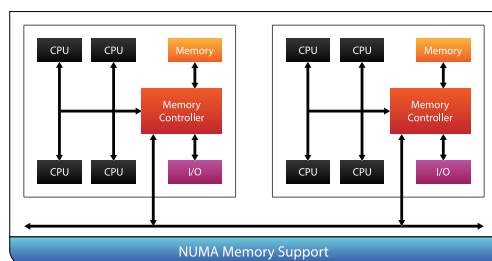
NUMA systems additionally manage data placement so that requested data is located as close to the appropriate CPU as possible. In cases where multiple CPUs are accessing the same data, NUMA maintains coherency by creating efficient memory "clusters" that keep data close to all CPUs that require it.

Support at the hardware and software level is required for NUMA to work properly. Hardware specifications for servers and workstations typically include information on whether NUMA architectures are supported. As for operating systems, Microsoft introduced support in Windows 7 and Windows Server 2008 R2; Linux provided limited NUMA support going back to the 2.5 kernel (significant improvements were made to the 3.8 and 3.13 kernel releases, however); and Java introduced support in its version 7 release. NUMA support for applications differs, so check in with the manufacturer.

Why is NUMA Important?

Modern CPUs operate considerably faster than the main memory they use, which leaves the CPU increasingly starved for data. In addition, many systems rely on outdated methods that limit computer memory access to a single processor at any one time. NUMA provides separate memory for each processor, as well as additional hardware to move data between memory banks.

NUMA additionally provides scalability. With an earlier memory management architecture called Symmetric Multiprocessing (SMP), all memory access was posted to the same shared memory bus. This worked fine for a relatively small number of CPUs and processing cores, but not when you had dozens or even hundreds of processing cores competing for access to the shared memory bus. With NUMA, potential bottlenecks are alleviated by limiting the number of CPUs on any one memory bus, and by providing high speed interconnection between the various nodes.



With NUMA, high speed interconnection between the various server nodes alleviates potential data bottlenecks.

NUMA architectures are often employed in applications such as data mining, business analytics, and other processor-intensive ecosystems where multiple processors collectively work on common databases where multiple operations are parceled out for simultaneous evaluation.

How Does ATTO Enhance Performance in NUMA Ecosystems?

ATTO 12Gb ExpressSAS host bus adapter (HBA) drivers are optimized to ensure that driver resources are located close to the system I/O. In addition, they support NUMA-aware applications and operating systems. The effect of this is to minimize access to non-local NUMA memory by the driver, which reduces latency as I/O is processed in the most efficient way possible.

What is NUMA's Impact on Performance?

While overall NUMA performance is highly dependent on many factors (including hardware architecture, operating system, application support, and configuration parameters), the HBA driver is an especially important consideration since it sits directly in the data path. A hardware device driver like ATTO's ExpressSAS HBA driver must be NUMA-aware to deliver the highest, most consistent performance when used in NUMA-enabled systems.

In less efficient, non-NUMA systems, data must be located and moved across multiple data buses to deliver it to the processing core that requires it to complete an operation. ATTO's NUMA-aware ExpressSAS driver, in contrast, works with the operating system to ensure that essential data is located close to the processor that needs it. In this way, it can be easily ingested into the processing core, resulting in the fastest possible I/O.

The best way to measure the impact of a properly-designed driver is to compare I/O results. Well-designed drivers show little I/O variability across cores, meaning that data is distributed efficiently so that no one single core gets overloaded.

The graphs below demonstrate the impact of NUMA-aware drivers on I/O in a Linux multi-node system. The NUMA-aware driver indicates almost 100% efficiency (green lines), meaning that data is located close to the processor that needs it, regardless of the transfer size. The non-NUMA-aware driver (red lines) shows that data must be moved across primary and secondary memory buses for a significant number of operations (17%-22% of the time) depending on the transfer size.

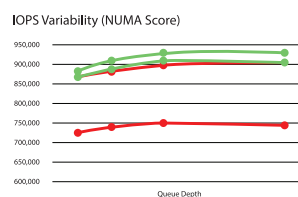


Figure 1: Linux I/O Variability

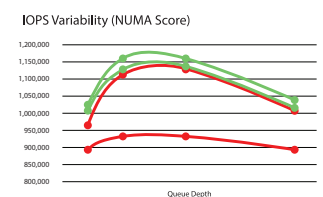


Figure 2: Windows I/O Variability