Linux Plumbers Conference 2019



Contribution ID: 252

Type: not specified

Improving RDMA performance through the use of contiguous memory and larger pages for files.

Wednesday, 11 September 2019 13:00 (30 minutes)

As memory sizes grow so do the sizes of the data transferred between RDMA devices. Generally, the Operating system needs to keep track of the state of each of its pieces of memory and that is on Intel x86 a page of 4 KB. This is also connected to hardware providing memory management features such as the processor page tables as well as the MMU features of the RDMA NIC.

The overhead of the operating system increases as the number of these pages reaches ever higher orders of magnitude. I.e. for 4GB of data one needs 1 million of these page descriptors. Each page descriptor is a 64-byte cache line and thus a 4GB operation requires 64MB of cache lines to be managed.

A lot of efforts on optimization of I/O focuses on avoiding touching these page descriptors through the use of larger contiguous memory or larger page sizes. This talk gives an overview of the current methods in use to avoid these lowdowns and the work in progress to improve the situation and make it less of an effort to avoid these issues.

I agree to abide by the anti-harassment policy

Presenter: LAMETER, Christopher (Jump Trading LLC) Session Classification: RDMA MC