

The relocation of live, running virtual machines from one physical host to another is a new and very desirable function because it provides a variety of features, including resilience to failures and flexibility of location. The data transfers required for live migration are supported in several commercial products by a protocol called iSCSI (Internet SCSI), which runs on top of TCP. We thoroughly tested the performance of a common open source component, the Open-iSCSI initiator, and found a drastic throughput degradation on 100 Mbps networks where the round trip time was more than about 40 ms. We localized the problem to the TCP send buffer size and tested two methods of setting the TCP send buffer size appropriately. Based on our results, we propose a performance tuning scheme that enables users of Open-iSCSI to achieve significant throughput gains. Our scheme results in a dramatic throughput jump from 14 Mbps to 70 Mbps on a 100 Mbps link with an RTT of 100 ms. We also modified one of the data structures internal to Open-iSCSI to handle multiple memory pages in a single scatter-gather list entry. This modification resulted in an additional 20% throughput increase on a 100 Mbps link with an RTT of 200 ms.