

Title: NetVM: High Performance and Flexible Networking using Virtualization on Commodity Platforms

Abstract:

NetVM brings virtualization to the Network, by enabling high bandwidth network functions to operate at near line speed, while taking advantage of the flexibility and customization of low cost commodity servers. NetVM exploits Intel's DPDK libraries to minimize the overhead of packet processing, and extends this framework and the KVM hypervisor to provide high throughput, low latency networking in virtualized environments. These techniques allow customizable processing capabilities such as firewalls, proxies, and routers to be embedded within the network, thus enabling the vision of software defined networking. NetVM runs network functions on commodity hardware using virtualization software, making it easy to dynamically scale, deploy, and reprogram services. This provides far greater flexibility than existing purpose-built, sometimes proprietary hardware. Complementing the control plane capabilities of SDNs, NetVM allows the system to perform the full range of packet processing required in network elements, including forwarding, implementing policies such as firewalls, and full packet inspection to determine subsequent processing. It does so with dramatically higher throughput than existing software router platforms. We have implemented NetVM on the KVM virtualization platform and detail many of the challenges we have solved such as adding support for high-speed inter-VM communication through shared huge pages and enhancing the CPU scheduler to prevent overheads caused by inter-core communication and context switching. NetVM allows true zero-copy delivery of data to VMs both for packet processing and messaging among VMs within a trust boundary. Our evaluation shows how NetVM can compose complex network functionality from multiple pipelined VMs and still obtain throughputs up to 10 Gbps, an improvement of more than 250% compared to existing techniques that use SR-IOV for virtualized networking.