WINTECHCON

Storage Offload on SmartNICs

September 27, 2019

Nayana Mariyappa
Network Software Engineer
Intel Data Center Group

Problem Statement

- Increasing demand on network and storage traffic
- Demands manifolds with the virtualization technologies
- More CPU and memory on host

  VMs -> vHost -> n/w and storage controller

  Burst of requests exhaust the resources
- Overall network throughput gets impacted, latency increases and performance degrades

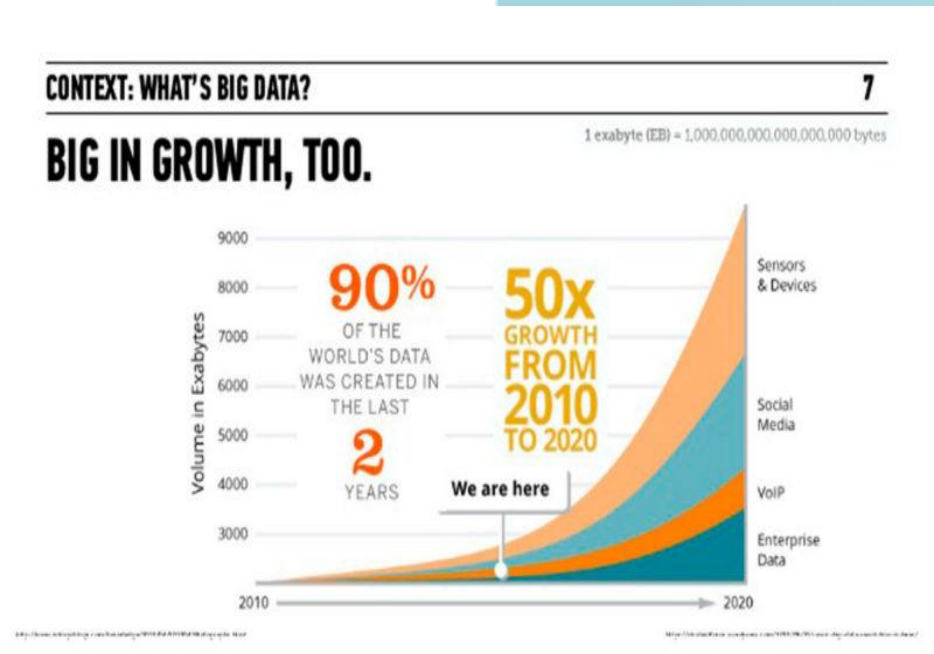

Credit from: https://avora.com/wp-content/uploads/2017/04/big-data-growth-1200x675.jpg

Problem Statement

- CPUs for network and storage are utilized by crypto/ live migration
- Traditionally increase in CPU, memory and storage increases cost and space
- Storage performance can be increased by NVMe
- Latency bottleneck in storage area network
- NVMe-oF solves this by acting as messaging layer b/w the host and target SSDs or shared system network storage over RDMA/Fibre channels

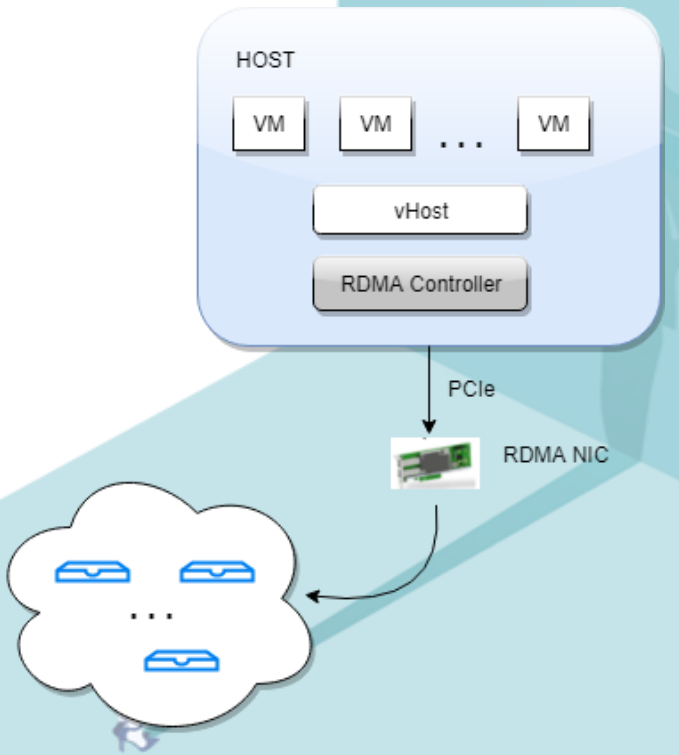

Storage Offload to SmartNICs

- Frees up CPUs on host which can be used for additional VMs
- Network, storage, crypto and live-migration can also be offloaded to SmartNIC
- SmartNIC has it own CPU cores and custom logic to handle control and data path
- Multiple storage protocols possible to be configured in a SmartNIC
- Maintain local cache for performance

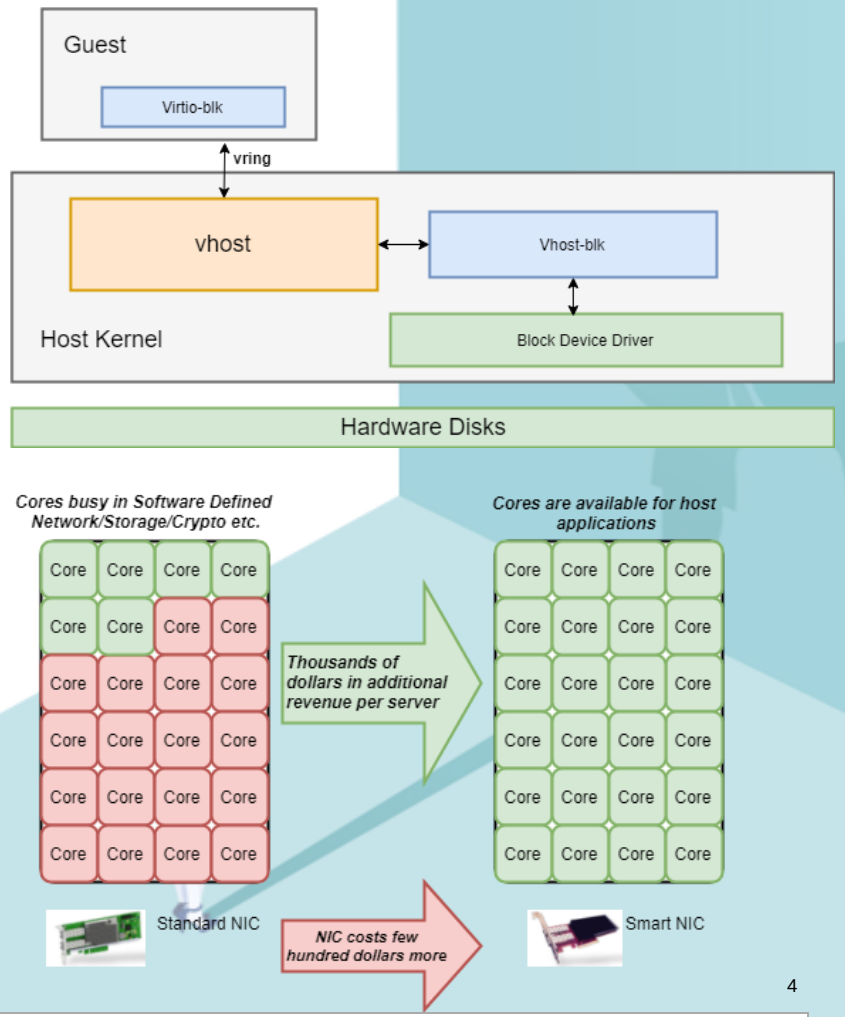

Storage Offload to SmartNICs

- Storage controllers provide high throughput with low latency than a networked or local storage
- These access storage data through NVMeoF and send storage request from Host/VM to cloud
- Reduces software complexity and supports direct data transfer between VMs and SmartNIC over PCIe
- CSPs can configure SPDK application on SoC to transfer the host requests to local NVMe storage array and NVMe-oF

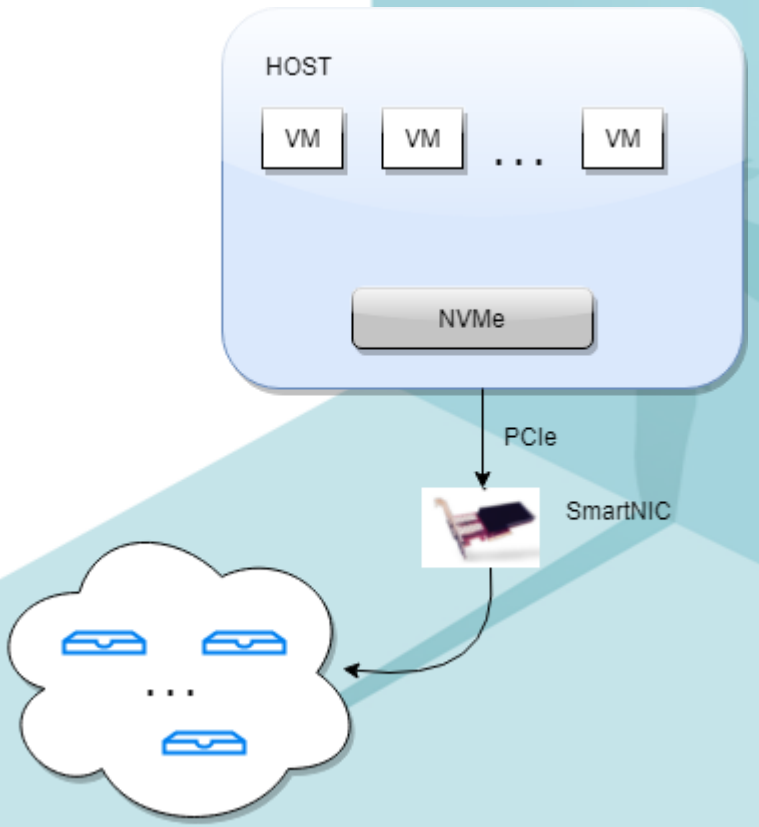

Experimental Results

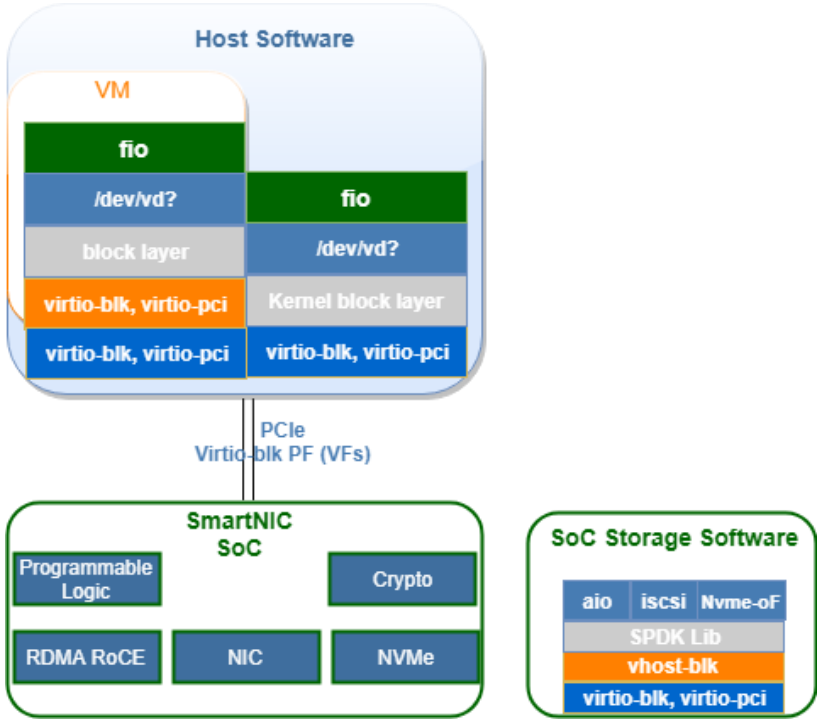

TABLE I: Storage Offload Performance Testing Results with SPDK

Block Size	Queue Size	IOPS
512	32	1267k
512	128	2314k
4096	32	896k
4096	128	1547k

TABLE II: Performance and latency for Host accessing cloud storage over iSCSI

IOPS	Bandwidth(Gbps)	Latency(ms)	
		P50	P99
1318k	40.24	38.14	48.89
1274k	38.8	36.6	52.99
1771k	54.04	26.75	34.04
1711k	52.24	18.816	40.704

Conclusion

- SmartNIC emulates the storage capability of the host and frees up storage and packet processing memory and CPUs on the host
- SmartNIC can do virtualized networked storage in an easier to manage and look-a-like to local physical storage
- Allows CSPs to deploy specific cryptography for storage traffic and maintain local cache for improved performance
- Moreover custom logic in SmartNIC will be able to assist the Live Migration of running VMs to different Host Server
- However, SmartNIC does not give better latency results compared to other methods discussed

References

- Mazhar Ali, Samee U. Khan, and Athanasios V. Vasilakos, ”Security in cloud computing: Opportunities and challenges”, *Information Sciences*, Vol. 305, pp. 357–383, 2015.

- Zhengyu Yang, Morteza Hoseinzadeh, Ping Wong, John Artoux, Clay Mayers, David Thomas Evans, Rory Thomas Bolt, Janki Bhimani, Ningfang Mi, and Steven Swanson, ”H-NVMe: A hybrid framework of NVMe-based storage system in cloud computing environment”, *2017 IEEE 36th International Performance Computing and Communications Conference (IPCCC)*, 2017.

- Daniel Firestone, Andrew Putnam, Sambhrama Mundkur, Derek Chiou, Alireza Dabagh, Mike Andrewartha, Hari Angepat, Vivek Bhanu, Adrian Caulfield, Eric Chung, Harish Kumar Chandrappa, Somesh Chaturmohta, Matt Humphrey, Jack Lavier, Norman Lam, Fengfen Liu, Kalin Ovtcharov, Jitu Padhye, Gautham Popuri, Shachar Raindel, Tejas Sapre, Mark Shaw, Gabriel Silva, Madhan Sivakumar, Nisheeth Srivastava, Anshuman Verma, Qasim Zuhair, Deepak Bansal, Doug Burger, Kushagra Vaid, David A. Maltz, and Albert Greenberg, ”Azure Accelerated Networking: SmartNICs in the Public Cloud”, *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI ’18)*, pp. 51–64, 2018