

ons

NORTH AMERICA

OPEN NETWORKING //

Enabling Collaborative
Development & Innovation

Are You Insured Against Your Noisy Neighbor?

Sunku Ranganath, Intel Corporation

Sridhar K.N. Rao, Spirent Communications

Hosted By

 THE **LINUX** FOUNDATION |  **LF** NETWORKING



ons
NORTH AMERICA
OPEN NETWORKING //
Enabling Collaborative
Development & Innovation

Legal Disclaimer

© 2018 Intel Corporation. Intel, the Intel logo, Intel Inside, the Intel Inside logo, Intel Experience What's Inside, The Intel Experience What's Inside logo, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

For more complete information about performance and benchmark results, visit www.intel.com/benchmarks.

The cost reduction scenarios described are intended to enable you to get a better understanding of how the purchase of a given Intel based product, combined with a number of situation-specific variables, might affect future costs and savings. Circumstances will vary and there may be unaccounted-for costs related to the use and deployment of a given product. Nothing in this document should be interpreted as either a promise of or contract for a given level of costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804.

No computer system can be absolutely secure.

Intel® Advanced Vector Extensions (Intel® AVX)* provides higher throughput to certain processor operations. Due to varying processor power characteristics, utilizing AVX instructions may cause a) some parts to operate at less than the rated frequency and b) some parts with Intel® Turbo Boost Technology 2.0 to not achieve any or maximum turbo frequencies. Performance varies depending on hardware, software, and system configuration and you can learn more at <http://www.intel.com/go/turbo>.

Available on select Intel® processors. Requires an Intel® HT Technology-enabled system. Your performance varies depending on the specific hardware and software you use. Learn more by visiting <http://www.intel.com/info/hyperthreading>.

§ Configurations: The testing was done on Based on fourth-generation Intel Xeon E5-2699 v4 @2.20 GHz processor with 22 cores, 55 MB LLC and 62 GB memory 16 1G hugepages. The testing was conducted in OPNFV Pharos testbed on Pod 12 by VSPERF community engineers

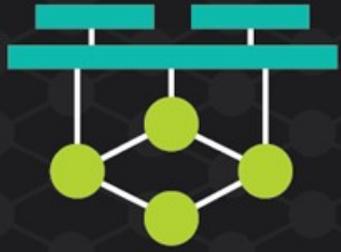
Intel, the Intel logo, [List the Intel trademarks in your document] are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

*Other names and brands may be claimed as the property of others.

© Intel Corporation

Hosted By

 THE LINUX FOUNDATION |  LFNETWORKING



ons

NORTH AMERICA

OPEN NETWORKING //

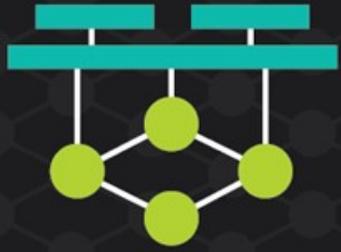
Enabling Collaborative
Development & Innovation

Acknowledgements

- Joseph Gasparakis
- Dakshina Illangovan
- Lin Yang
- Edwin Verplanke
- Priya Autee

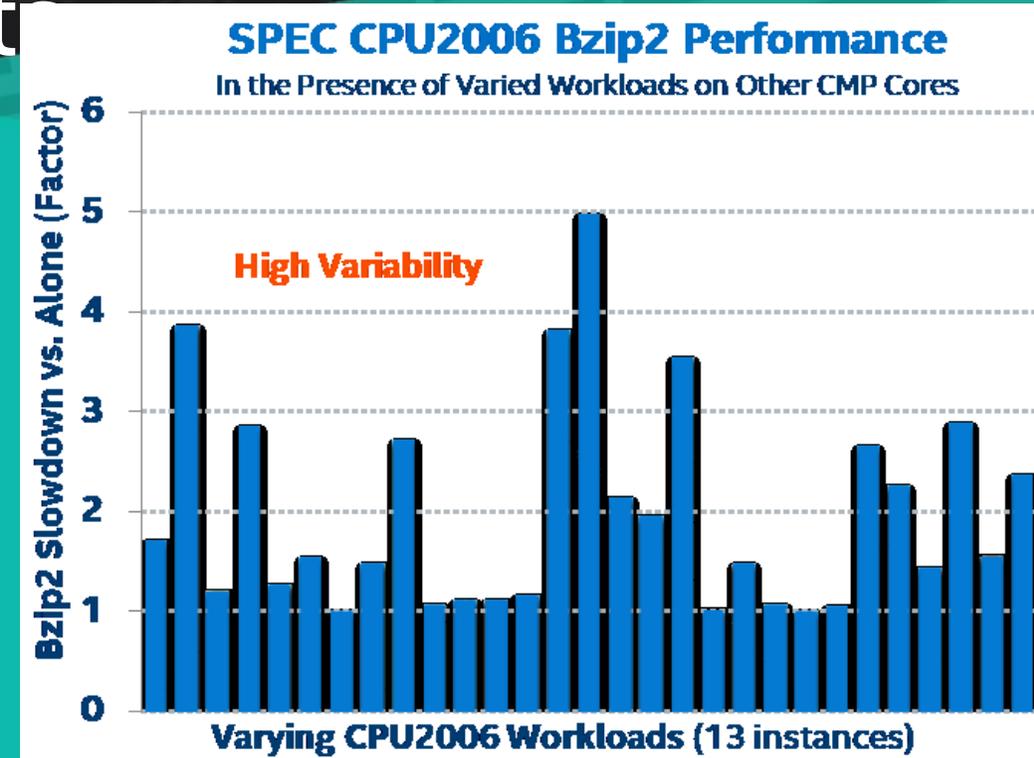
Hosted By

 THE **LINUX** FOUNDATION |  **LF** NETWORKING



Common Contentions in Cloud Deployments

- Minimizing Total Cost of Ownership (TCO) often leads to oversubscription
- Quality of Service (QoS) requirements
- Lack of control of Last Level Cache (LLC) by orchestration layer
- Last-Level Cache Contention Can Lead to 51% Throughput Degradation¹ in Comms Workloads
- Further: Last-Level Cache Contention Can Lead to Almost 5x Performance Variation¹



NFV & RT workloads are Time Sensitive

Hosted By



ons

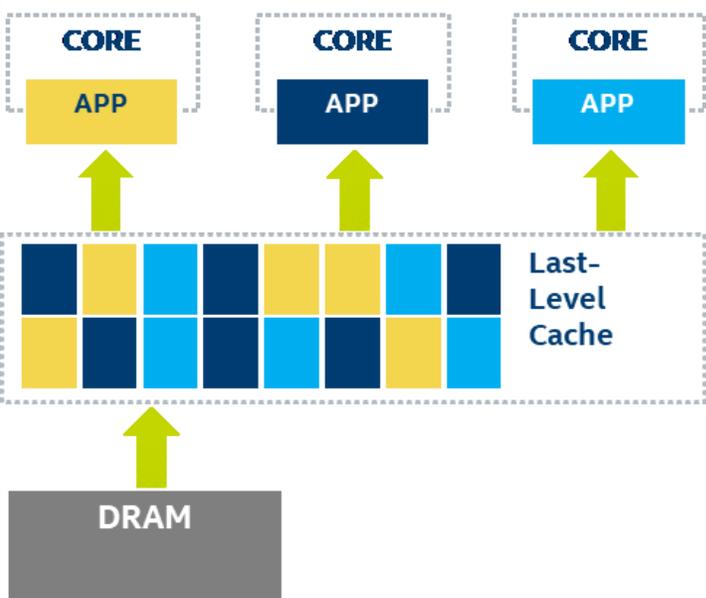
NORTH AMERICA

OPEN NETWORKING //

Enabling Collaborative
Development & Innovation

Intel Resource Director Technology

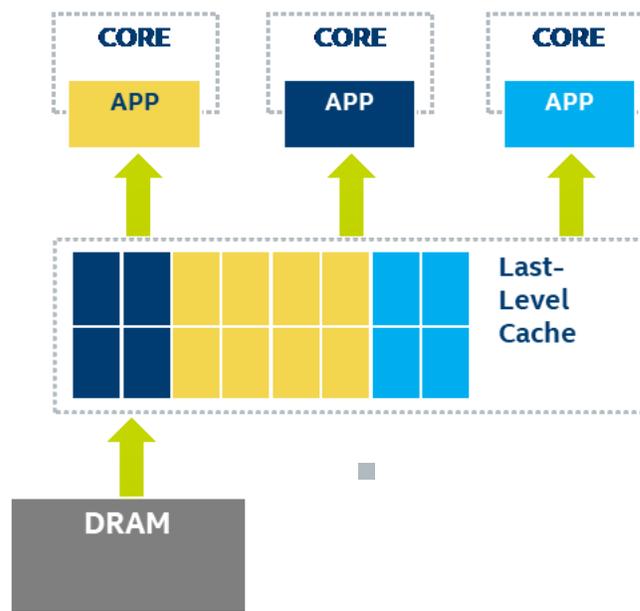
Cache Monitoring Technology (CMT)



Enabling applications
to be monitored according

- Cache Occupancy reported on a per Resource Monitoring ID (RMID) basis—Advanced Telemetry

Cache Allocation Technology (CAT)

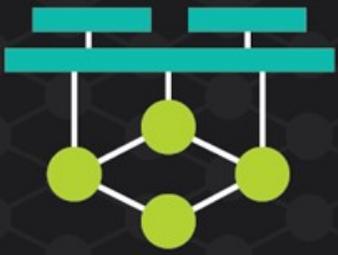


Enabling cache partitioning
and enabling separation
and prioritization of apps or VMs

- Misbehaving threads can be isolated to increase determinism

Hosted By

THE LINUX FOUNDATION | LFN NETWORKING



ons

NORTH AMERICA

OPEN NETWORKING //
Enabling Collaborative
Development & Innovation

Determinism with LLC Management

- Cache Quality of Service (QoS) adjustments with multiple workloads
- Consistency in Throughput & Latency w/ Noisy Neighbor avoidance
- Test setup with OPNFV VSPERF
- Leverage Spirent CloudStress as Noisy Neighbor
- Utilize Collectd for Metrics

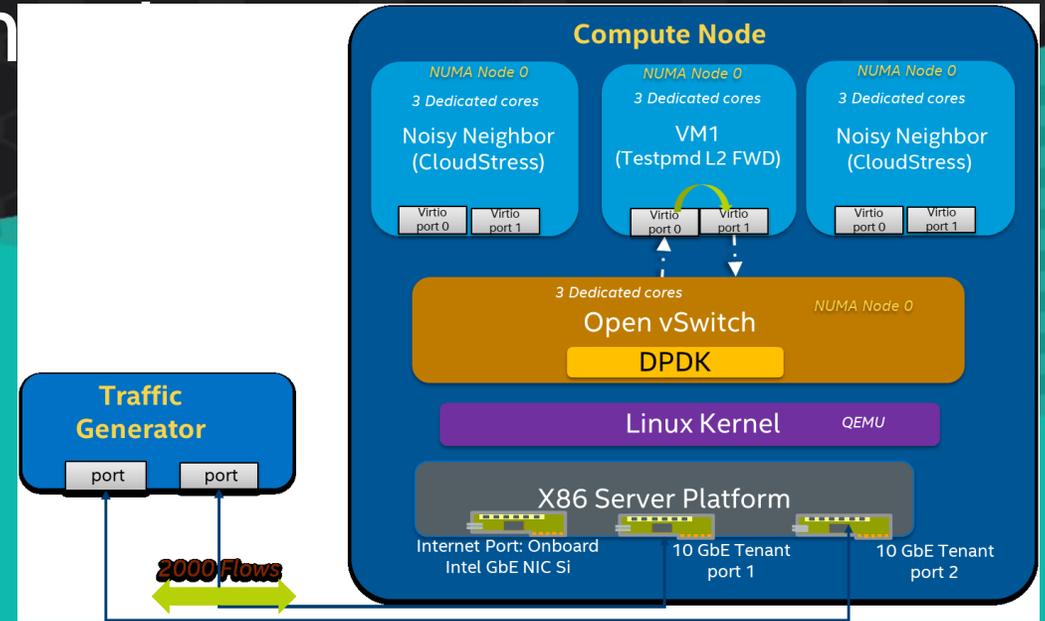


Fig: VSPERF Test Setup

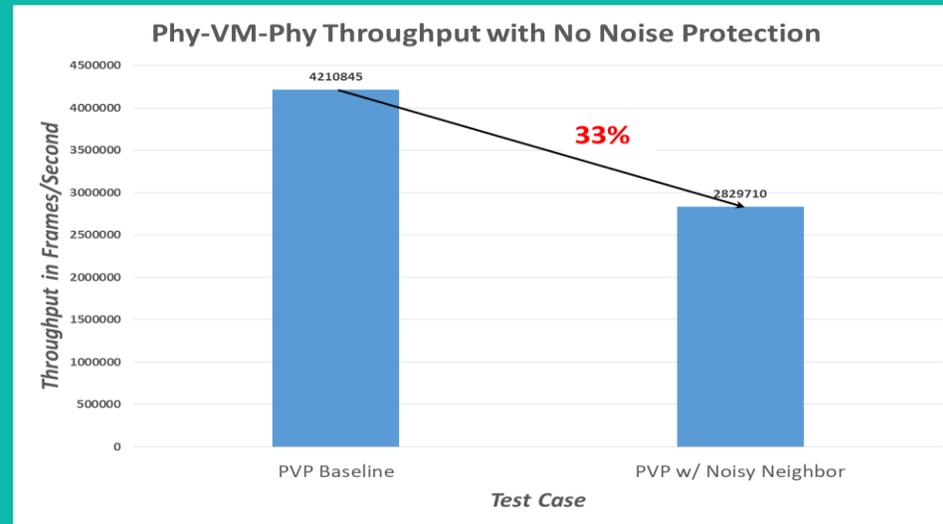


Fig: Throughput of 64B Packets



Policy Driven LLC Allocation - Resource Management Daemon

- Associate workload with dedicated LLC
- Construct LLC policy to protect workload from contention

Resource Management Daemon (RMD) - A Linux daemon that:

- Runs on individual hosts
- REST API, accessible to orchestrator
- Accepts & enforces policy
- Platform Aware

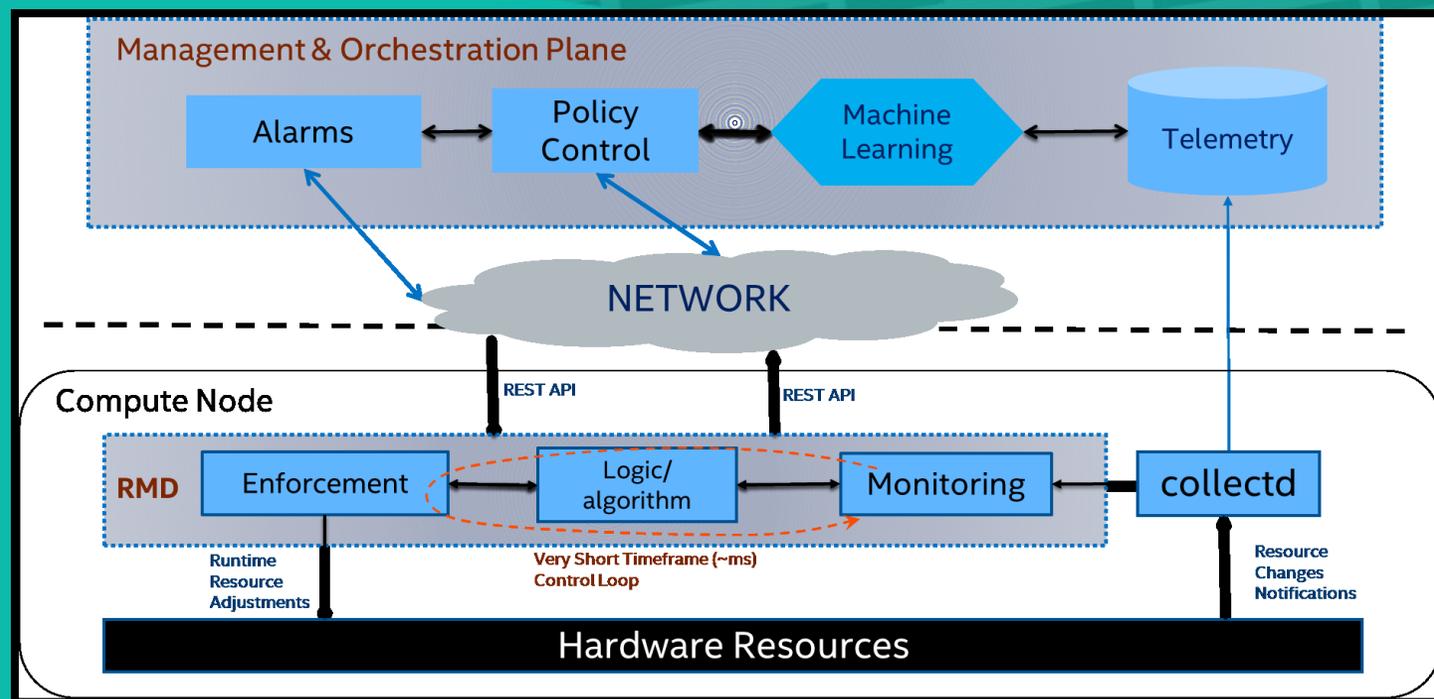
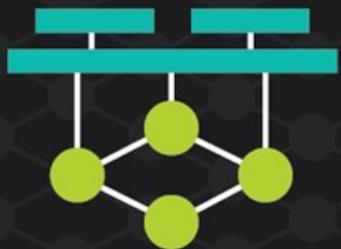


Fig: RMD interactions with Platform & MANO Layer

Open Source: <https://github.com/intel/rmd>

Hosted By



ONS

NORTH AMERICA

OPEN NETWORKING //

Enabling Collaborative
Development & Innovation

Workload Sensitivity & Policy Mapping

- Apply LLC policy at run time using RMD
 - LLC for VM under test – “Guaranteed” -2.5 MB
 - LLC for CloudStress VMs – “Best-effort” - 2.5 MB/VM
- Re-run the performance tests

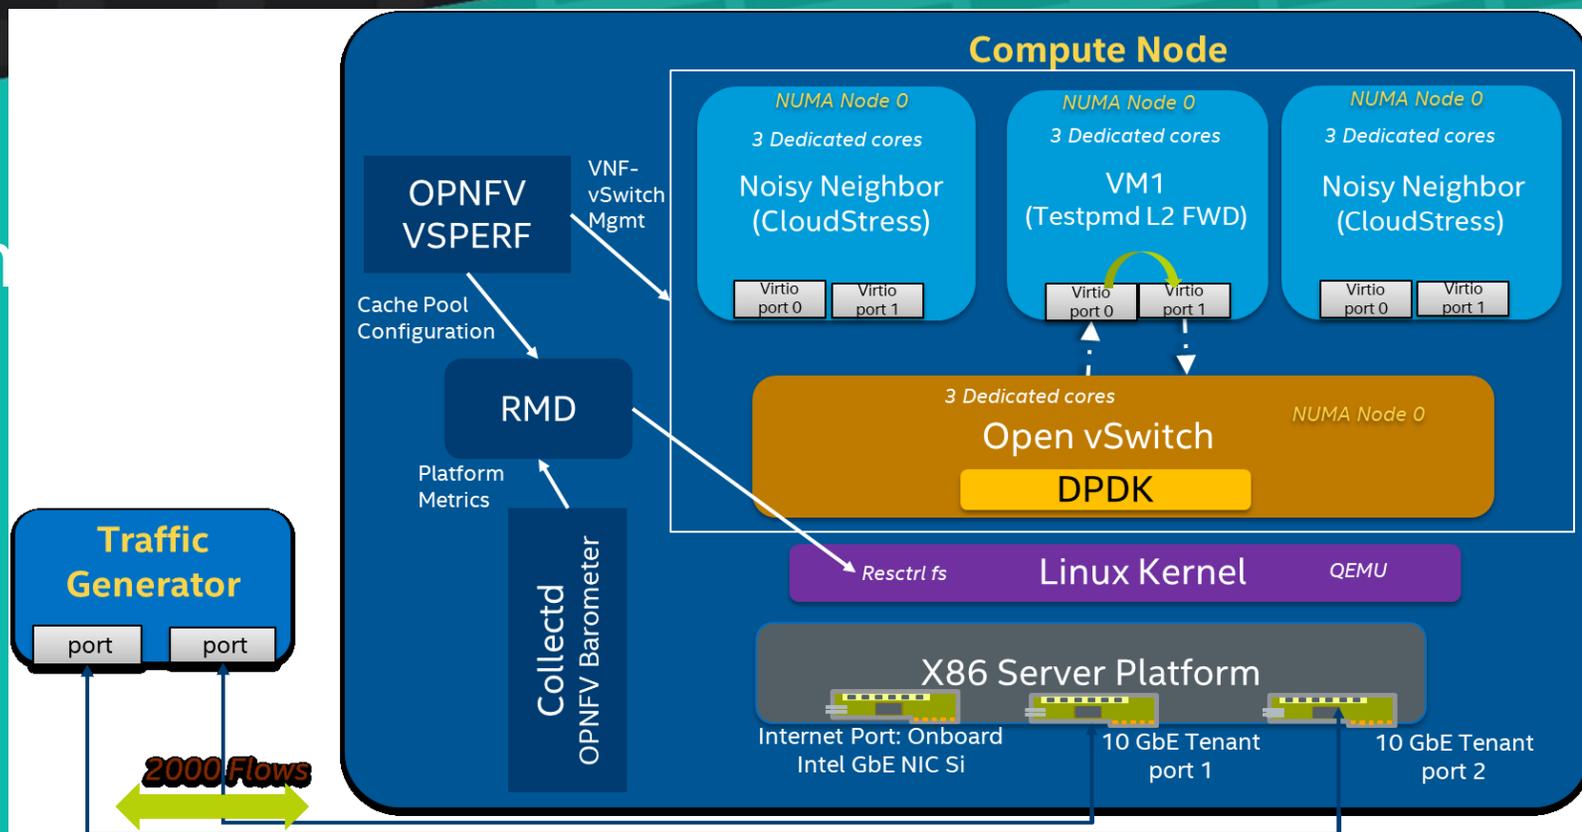
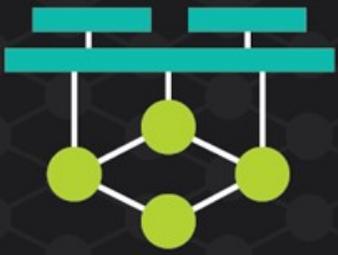


Fig: VSPERF Test setup with RMD & Collectd

Run Time LLC Control via RMD Policy Mapping

Hosted By

THE LINUX FOUNDATION | LFNWORKING



Noisy Neighbor Protection & Takeaways

- Guaranteed LLC policy helped preserve VM performance
- Throughput improvement of ~40% with noise protection
- Integrating into OpenStack & Kubernetes

Takeaways

- Noisy Neighbor affects are real and here to persist
- Intel RDT & RMD enables hardware infrastructure for LLC QoS control
- OPNFV VSPERF with RMD enables LLC QoS analysis for NFVi

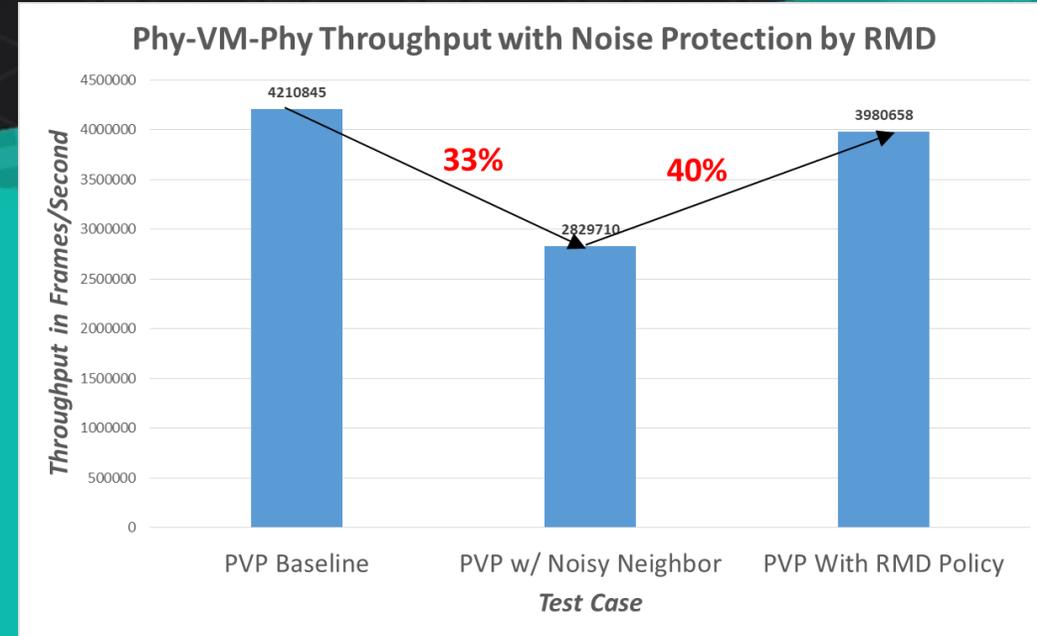


Fig: Throughput of 64B Packets

Optimal Cache Policy Ensures Deterministic Performance

Hosted By